[28] A.S. Rao and M.P. Georgeff. Asymmetry thesis and side-effect problems in linear time and branching time intention logics. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI91)*, pages 498–504, 1991.

[29] A.S. Rao and M.P. Georgeff. Modeling rational agents within a BDI-architecture. In J. Allen, R. Fikes, and E. Sandewall, editors, *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, pages 473–484, San Mateo CA, 1991. Mogan Kaufmann.

[30] A.S. Rao and M.P. Georgeff. A model-theoretic approach to the verification of situated reasoning systems. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI'93)*, pages 318–324, 1993.

[31] Y. Shoham. Agent-oriented programming. *Artificial Intelligence*, 60:51–92, 1993.

[32] Y. Shoham and S.B. Cousins. Logics of mental attitudes in AI. In G. Lakemeyer and B. Nebel, editors, *Foundations of Knowledge Representation and Reasoning*, pages 296–309. Springer-Verlag, 1994.

[33] R. Stalnaker. A theory of conditionals. In N. Rescher, editor, *Studies in Logical Theory*, number 2 in American Philosophical Quarterly Monograph Series, pages 98–122. Blackwell, Oxford, 1968.

[34] E. Thijsse. *Partial Logic and Knowledge Representation.* PhD thesis, Katholieke Universiteit Brabant, 1992.

[35] S.R. Thomas. *PLACA, An Agent Oriented Programming Language.* PhD thesis, Department of Computer Science, Stanford University, Stanford CA, September 1993. Appeared as technical report STAN-CS-93-1487.

[36] M. Wooldridge. *The Logical Modelling of Computational Multi-Agent Systems.* PhD thesis, Department of Computation, UMIST, Manchester, October 1994. Appeared as technical report MMU-DOC-94-01.

[37] M. Wooldridge and M. Fisher. A first-order branching time logic of multi-agent systems. In B. Neumann, editor, *Proceedings of the 10th European Conference on Artificial Intelligence (ECAI'92)*, pages 234–238. John Wiley & Sons, 1992.

[38] M. .Wooldridge and N.R. Jennings. Intelligent agents: Theory and practice. Submitted to Knowledge Engineering Review.

[39] G.H. von Wright. *Norm and Action.* Routledge & Kegan Paul, London, 1963.

Our papers are available at http://www.cs.ruu.nl/~bernd

[14] Z. Huang. *Logics for Agents with Bounded Rationality.* PhD thesis, Universiteit van Amsterdam, 1994.

[15] Z. Huang, M. Masuch, and L. Pólos. ALX, an action logic for agents with bounded rationality. Technical Report 92-70, Center for Computer Science in Organization and Management, University of Amsterdam, October 1992. To appear in Artificial Intelligence.

[16] G.E. Hughes and M.J. Cresswell. *An Introduction to Modal Logic.* Routledge, London, 1968.

[17] A. Kenny. *Will, Freedom and Power.* Basil Blackwell, Oxford, 1975.

[18] S. Kraus and D. Lehmann. Knowledge, belief and time. *Theoretical Computer Science*, 58:155–174, 1988.

[19] Y. Lesperance, H. Levesque, F. Lin, D. Marcu, R. Reiter, and R. Scherl. A logical approach to high-level robot programming – a progress report. To appear in Control of the Physical World by Intelligent Systems, Working Notes of the 1994 AAAI Fall Symposium, New Orleans, LA, November, 1994.

[20] H. Levesque. Knowledge, action and ability in the situation calculus. Overheads from invited talk at TARK 1994.

[21] I. Levi. Direct inference. *The Journal of Philosophy*, 74:5–29, 1977.

[22] B. van Linder, W. van der Hoek, and J.-J. Ch. Meyer. Communicating rational agents. In B. Nebel and L. Dreschler-Fischer, editors, *KI-94: Advances in Artificial Intelligence*, volume 861 of *Lecture Notes in Computer Science (subseries LNAI)*, pages 202–213. Springer-Verlag, 1994.

[23] B. van Linder, W. van der Hoek, and J.-J. Ch. Meyer. The dynamics of default reasoning. Technical Report UU-CS-1994-48, Utrecht University, October 1994.

[24] B. van Linder, W. van der Hoek, and J.-J. Ch. Meyer. Tests as epistemic updates. In A.G. Cohn, editor, *Proceedings of the 11th European Conference on Artificial Intelligence (ECAI'94)*, pages 331–335. John Wiley & Sons, 1994.

[25] J.-J. Ch. Meyer and W. van der Hoek. *Epistemic Logic for AI and Computer Science.* Cambridge University Press, 1994. To appear.

[26] R.C. Moore. Reasoning about knowledge and action. Technical Report 191, SRI International, 1980.

[27] R.C. Moore. A formal theory of knowledge and action. Technical Report 320, SRI International, 1984.

References

[1] B.F. Chellas. *Modal Logic. An Introduction.* Cambridge University Press, Cambridge, 1980.

[2] P.R. Cohen and H.J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.

[3] D. Elgesem. *Action Theory and Modal Logic.* PhD thesis, Institute for Philosophy, University of Oslo, Oslo, Norway, 1993.

[4] R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning about Knowledge.* MIT Press, Cambridge MA, 1994. To appear.

[5] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* The MIT Press, Cambridge, Massachusetts and London, England, 1988.

[6] R. Goldblatt. *Logics of Time and Computation*, volume 7 of *CSLI Lecture Notes.* CSLI, Stanford, 1992. Second edition.

[7] J.Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artifical Intelligence*, 54:319–379, 1992.

[8] D. Harel. Dynamic logic. In D.M. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, volume 2, chapter 10, pages 497–604. D. Reidel, Dordrecht, 1984.

[9] J. Hintikka. *Knowledge and Belief.* Cornell University Press, Ithaca, NY, 1962.

[10] W. van der Hoek. Systems for knowledge and beliefs. *Journal of Logic and Computation*, 3(2):173–195, 1993.

[11] W. van der Hoek, B. van Linder, and J.-J. Ch. Meyer. A logic of capabilities. In A. Nerode and Yu. V. Matiyasevich, editors, *Proceedings of the Third International Symposium on the Logical Foundations of Computer Science (LFCS'94)*, volume 813 of *Lecture Notes in Computer Science*, pages 366–378. Springer-Verlag, 1994.

[12] W. van der Hoek, B. van Linder, and J.-J. Ch. Meyer. Unravelling nondeterminism: On having the ability to choose (extended abstract). In P. Jorrand and V. Sgurev, editors, *Proceedings of the Sixth International Conference on Artificial Intelligence: Methodology, Systems, Applications (AIMSA'94)*, pages 163–172. World Scientific, 1994.

[13] J. Horty and Y. Shoham, Program Chairs. Reasoning about mental states: Formal theories & applications. Technical Report SS-93-05, AAAI Press, 1993. Papers from the 1993 AAAI Spring Symposium, Stanford CA.

$$\begin{aligned}
&\mathcal{M}, s \models \mathbf{A}_i\texttt{expand}\ \varphi \\
\Leftrightarrow\ &\mathcal{M}, s \models \neg\mathbf{B}_i\neg\varphi \\
\Leftrightarrow\ &\exists s' \in \mathrm{B}(i,s)[\mathcal{M}, s' \models \varphi] \\
\Leftrightarrow\ &\mathrm{B}(i,s) \cap [\![\varphi]\!] \neq \emptyset \\
\Leftrightarrow\ &\mathrm{B}'(i,s) \neq \emptyset \\
\Leftrightarrow\ &\mathcal{M}', s \models \neg\mathbf{B}_i\mathbf{ff} \\
\Leftrightarrow\ &\mathcal{M}, s \models \langle \mathrm{do}_i(\texttt{expand}\ \varphi)\rangle\neg\mathbf{B}_i\mathbf{ff}
\end{aligned}$$

Clause 4: Let $\mathcal{M}', s = \mathtt{r}(i, \texttt{contract}\ \varphi)(\mathcal{M}, s)$.

$$\begin{aligned}
&\mathcal{M}, s \models \mathbf{A}_i\texttt{contract}\ \varphi \\
\Leftrightarrow\ &\mathcal{M}, s \models \neg\mathbf{K}_i\varphi \\
\Rightarrow\ &\varphi \notin \mathcal{B}^-_\varphi(i, \mathcal{M}, s) && \text{(by } B^-4) \\
\Leftrightarrow\ &\mathcal{M}', s \not\models \mathbf{B}_i\varphi \\
\Leftrightarrow\ &\mathcal{M}, s \models \langle \mathrm{do}_i(\texttt{contract}\ \varphi)\rangle\neg\mathbf{B}_i\varphi
\end{aligned}$$

Clause 5:

$$\begin{aligned}
&\mathcal{M}, s \models \mathbf{A}_i\texttt{contract}\ \neg\varphi \\
\Leftrightarrow\ &\mathcal{M}, s \models \mathbf{A}_i\texttt{contract}\ \neg\varphi \wedge \langle \mathrm{do}_i(\texttt{contract}\ \neg\varphi)\rangle\neg\mathbf{B}_i\neg\varphi && \text{(Clause 4)} \\
\Leftrightarrow\ &\mathcal{M}, s \models \mathbf{A}_i\texttt{contract}\ \neg\varphi \wedge [\mathrm{do}_i(\texttt{contract}\ \neg\varphi)]\neg\mathbf{B}_i\neg\varphi && \text{(Proposition 3.17)} \\
\Leftrightarrow\ &\mathcal{M}, s \models \mathbf{A}_i\texttt{contract}\ \neg\varphi \wedge [\mathrm{do}_i(\texttt{contract}\ \neg\varphi)]\mathbf{A}_i\texttt{expand}\ \varphi \\
\Leftrightarrow\ &\mathcal{M}, s \models \mathbf{A}_i(\texttt{contract}\ \neg\varphi; \texttt{expand}\ \varphi) \\
\Leftrightarrow\ &\mathcal{M}, s \models \mathbf{A}_i\texttt{revise}\ \varphi
\end{aligned}$$

⊠

3.43. Proposition. *For all agents i, and for all propositional formulae φ we have:*

- $\models \mathbf{A}_i\texttt{expand}\ \varphi \to \mathbf{Can}_i(\texttt{expand}\ \varphi, \mathbf{B}_i\varphi \wedge \neg\mathbf{B}_i\mathbf{ff})$
- $\models \mathbf{A}_i\texttt{contract}\ \varphi \to \mathbf{Can}_i(\texttt{contract}\ \varphi, \neg\mathbf{B}_i\varphi)$
- $\models \mathbf{A}_i\texttt{contract}\ \neg\varphi \to \mathbf{Can}_i(\texttt{revise}\ \varphi, \mathbf{B}_i\varphi \wedge \neg\mathbf{B}_i\mathbf{ff})$

Proof: We show the first clause; the other clauses are analogous.

$$\begin{aligned}
&\models [\mathrm{do}_i(\texttt{expand}\ \varphi)]\mathbf{B}_i\varphi && \text{Proposition 3.6, clause 1} \\
\Leftrightarrow\ &\models \langle \mathrm{do}_i(\texttt{expand}\ \varphi)\rangle\mathbf{B}_i\varphi && \text{Proposition 3.7, clause 2} \\
\Leftrightarrow\ &\models \mathbf{K}_i\langle \mathrm{do}_i(\texttt{expand}\ \varphi)\rangle\mathbf{B}_i\varphi\ (*) && \text{Proposition 2.6, clause 5}
\end{aligned}$$

$$\begin{aligned}
&\models \mathbf{A}_i\texttt{expand}\ \varphi \to \langle \mathrm{do}_i(\texttt{expand}\ \varphi)\rangle\neg\mathbf{B}_i\mathbf{ff} && \text{Proposition 3.41, clause 2} \\
&\models \mathbf{K}_i(\mathbf{A}_i\texttt{expand}\ \varphi \to \langle \mathrm{do}_i(\texttt{expand}\ \varphi)\rangle\neg\mathbf{B}_i\mathbf{ff}) && \text{Proposition 2.6, clause 5} \\
&\models \mathbf{K}_i\mathbf{A}_i\texttt{expand}\ \varphi \to \mathbf{K}_i\langle \mathrm{do}_i(\texttt{expand}\ \varphi)\rangle\neg\mathbf{B}_i\mathbf{ff}\ (**) && \text{Proposition 2.6, clause 1}
\end{aligned}$$

Now assume $\mathcal{M}, s \models \mathbf{A}_i\texttt{expand}\ \varphi$ for some arbitrary Kripke model $\mathcal{M}$ with state s. Then also $\mathcal{M}, s \models \mathbf{K}_i\mathbf{A}_i\texttt{expand}\ \varphi$ by Clause 1 of Proposition 3.41. Hence using (*) and (**) it follows that $\mathcal{M}, s \models \mathbf{K}_i\mathbf{A}_i\texttt{expand}\ \varphi \wedge \mathbf{K}_i\langle \mathrm{do}_i(\texttt{expand}\ \varphi)\rangle\mathbf{B}_i\varphi \wedge \mathbf{K}_i\langle \mathrm{do}_i(\texttt{expand}\ \varphi)\rangle\neg\mathbf{B}_i\mathbf{ff}$. Hence $\mathcal{M}, s \models \mathbf{K}_i(\langle \mathrm{do}_i(\texttt{expand}\ \varphi)\rangle(\mathbf{B}_i\varphi \wedge \neg\mathbf{B}_i\mathbf{ff}) \wedge \mathbf{A}_i\texttt{expand}\ \varphi)$ and thus $\mathcal{M}, s \models \mathbf{Can}_i(\texttt{expand}\ \varphi, \mathbf{B}_i\varphi \wedge \neg\mathbf{B}_i\mathbf{ff})$.

⊠

Hence in both cases $B_2(i,s) \subseteq B_1(i,s)$, hence $\mathcal{B}(i,\mathcal{M}_1,s) \subseteq \mathcal{B}(i,\mathcal{M}_2,s)$, and thus $\mathcal{B}^{\star+}_{\varphi\psi}(i,\mathcal{M},s) \subseteq \mathcal{B}^{\star}_{\varphi\wedge\psi}(i,\mathcal{M},s)$.

⊠

3.39. PROPOSITION. *For all Kripke models $\mathcal{M}$ with state s, for all agents i, and for all propositional formulae φ we have:*

- *if $\neg\varphi \in \mathcal{K}(i,\mathcal{M},s)$ then $\mathcal{B}^{\star}_{\varphi}(i,\mathcal{M},s) = \mathcal{B}_{\perp}$.*
- *if $\neg\varphi \notin \mathcal{B}(i,\mathcal{M},s)$ then $\mathcal{B}^{\star}_{\varphi}(i,\mathcal{M},s) = \mathrm{Th}(\mathcal{B}(i,\mathcal{M},s) \cup \{\varphi\})$.*
- *if $\neg\varphi \in \mathcal{B}(i,\mathcal{M},s) \setminus \mathcal{K}(i,\mathcal{M},s)$ then $\mathcal{B}^{\star}_{\varphi}(i,\mathcal{M},s) = \mathrm{Th}(\mathcal{K}(i,\mathcal{M},s) \cup \{\varphi\})$ if the definition of* `r` *for the* `contract` *action is based on the AiG function for $\mathcal{M}$.*

PROOF: Let $\mathcal{M}$ be some Kripke model with state s, let i be some agent, and let φ be some arbitrary propositional formula. Let $\mathcal{M}', s = \mathtt{r}(i, \mathtt{contract}\ \neg\varphi)(\mathcal{M},s)$, and let $\mathcal{M}'', s = \mathtt{r}(i, \mathtt{expand}\ \varphi)(\mathcal{M}',s) = \mathtt{r}(i, \mathtt{revise}\ \varphi)(\mathcal{M},s)$. We successively prove the three cases.

- Suppose $\neg\varphi \in \mathcal{K}(i,\mathcal{M},s)$. Then by definition it follows that $\mathcal{M}', s = \mathcal{M}, s$, and hence $\mathcal{M}', s \models \mathbf{B}_i\neg\varphi$. Then the expansion with φ of the beliefs of agent i in $\mathcal{M}', s$ leads to a model $\mathcal{M}''$ such that $B''(i,s) = \emptyset$, and hence $\mathcal{B}^{\star}_{\varphi}(i,\mathcal{M},s) = \mathcal{B}(i,\mathcal{M}'',s) = \mathcal{B}_{\perp}$.
- Suppose $\neg\varphi \notin \mathcal{B}(i,\mathcal{M},s)$. In this case it follows from $B^{\star}3$ and $B^{\star}4$ that $\mathcal{B}^{\star}_{\varphi}(i,\mathcal{M},s) = \mathcal{B}^{+}_{\varphi}(i,\mathcal{M},s)$, and using Proposition 3.13 it follows that $\mathcal{B}^{\star}_{\varphi}(i,\mathcal{M},s) = \mathrm{Th}(\mathcal{B}(i,\mathcal{M},s) \cup \{\varphi\})$.
- Suppose $\neg\varphi \in \mathcal{B}(i,\mathcal{M},s) \setminus \mathcal{K}(i,\mathcal{M},s)$. Then by definition of the AiG function it follows that $B'(i,s) = B(i,s) \cup ([s]_{R(i)} \cap [\![\varphi]\!])$. By definition of $\mathtt{r}(i, \mathtt{expand}\ \varphi)$ it follows that $B''(i,s) = B'(i,s) \cap [\![\varphi]\!]$, hence $B''(i,s) = (B(i,s) \cup ([s]_{R(i)} \cap [\![\varphi]\!])) \cap [\![\varphi]\!]$, and since $\neg\varphi \in \mathcal{B}(i,\mathcal{M},s)$ it follows that $B''(i,s) = [s]_{R(i)} \cap [\![\varphi]\!]$. By an argument similar to that given in the proof of Proposition 3.13 it is shown that $\mathcal{B}^{\star}_{\varphi}(i,\mathcal{M},s) = \mathrm{Th}(\mathcal{K}(i,\mathcal{M},s) \cup \{\varphi\})$.

⊠

3.41. PROPOSITION. *For all agents i and for all propositional formulae φ we have:*

- $\models \mathbf{A}_i\mathtt{expand}\ \varphi \leftrightarrow \mathbf{K}_i\mathbf{A}_i\mathtt{expand}\ \varphi$
- $\models \mathbf{A}_i\mathtt{expand}\ \varphi \to \langle \mathrm{do}_i(\mathtt{expand}\ \varphi)\rangle\neg\mathbf{B}_i\mathbf{ff}$
- $\models \mathbf{A}_i\mathtt{contract}\ \varphi \leftrightarrow \mathbf{K}_i\mathbf{A}_i\mathtt{contract}\ \varphi$
- $\models \mathbf{A}_i\mathtt{contract}\ \varphi \to \langle \mathrm{do}_i(\mathtt{contract}\ \varphi)\rangle\neg\mathbf{B}_i\varphi$
- $\models \mathbf{A}_i\mathtt{revise}\ \varphi \leftrightarrow \mathbf{A}_i\mathtt{contract}\ \neg\varphi$
- $\models \mathbf{A}_i\mathtt{revise}\ \varphi \to \langle \mathrm{do}_i(\mathtt{revise}\ \varphi)\rangle(\mathbf{B}_i\varphi \wedge \neg\mathbf{B}_i\mathbf{ff})$

PROOF: The first and the third clause are straightforward consequences of the sixth and the fifth clause of Proposition 2.8 respectively. Clause 6 is a combination of clauses 2, 4, and 5. The other clauses are proved as follows. Let $\mathcal{M}$ be some Kripke model with state s, let i be some agent and let φ be some propositional formula.
Clause 2: Let $\mathcal{M}', s = \mathtt{r}(i, \mathtt{expand}\ \varphi)(\mathcal{M},s)$.

$\sigma(i,s,\neg\varphi) \cap [\![\psi]\!] \subseteq \sigma(i,s,\neg\varphi \vee \neg\psi)$, then $\mathrm{B}_2(i,s) \subseteq \mathrm{B}_1(i,s)$, and therefore $\mathcal{B}(i,\mathcal{M}_1,s) \subseteq \mathcal{B}(i,\mathcal{M}_2,s)$. So to prove that $\sigma(i,s,\neg\varphi) \cap [\![\psi]\!] \subseteq \sigma(i,s,\neg\varphi \vee \neg\psi)$. Since $\models \mathbf{K}_i((\neg\varphi \vee \neg\psi) \wedge (\neg\varphi \vee \psi) \leftrightarrow \neg\varphi)$, we have by $\Sigma 4$ that $\sigma(i,s,\neg\varphi) = \sigma(i,s,(\neg\varphi \vee \neg\psi) \wedge (\neg\varphi \vee \psi))$. From $\Sigma 5$ it follows that $\sigma(i,s,(\neg\varphi \vee \neg\psi) \wedge (\neg\varphi \vee \psi)) \subseteq \sigma(i,s,\neg\varphi \vee \neg\psi) \cup \sigma(i,s,\neg\varphi \vee \psi)$. From $\Sigma 1$ we conclude that $\sigma(i,s,\neg\varphi \vee \neg\psi) \subseteq [\![\varphi \wedge \psi]\!]$ and $\sigma(i,s,\neg\varphi \vee \psi) \subseteq [\![\varphi \wedge \neg\psi]\!]$. Since $\sigma(i,s,\neg\varphi) \subseteq [\![\varphi]\!]$ we have that $\sigma(i,s,\neg\varphi) \cap [\![\psi]\!] \subseteq [\![\varphi \wedge \psi]\!]$. Hence $\sigma(i,s,\neg\varphi) \cap [\![\psi]\!] \subseteq \sigma(i,s,\neg\varphi \vee \neg\psi)$, which was to be proved.

($B^\star 8$) Suppose $\neg\psi \notin \mathcal{B}^\star_\varphi(i,\mathcal{M},s)$. To keep our proof understandable we introduce the following definitions:

- $\mathcal{M}_{11}, s = \mathtt{r}(i, \mathtt{revise}\ \varphi)(\mathcal{M}, s)$
- $\mathcal{M}_1, s = \mathtt{r}(i, \mathtt{expand}\ \psi)(\mathcal{M}_{11}, s)$
- $\mathcal{M}_{21}, s = \mathtt{r}(i, \mathtt{contract}\ \neg\varphi \vee \neg\psi)(\mathcal{M}, s)$
- $\mathcal{M}_2, s = \mathtt{r}(i, \mathtt{expand}\ \varphi \wedge \psi)(\mathcal{M}_{21}, s) = \mathtt{r}(i, \mathtt{revise}\ \varphi \wedge \psi)(\mathcal{M}, s)$

Using similar arguments as in the proof of $B^\star 7$ we find that:

- $\mathrm{B}_{11}(i,s) = (\mathrm{B}(i,s) \cup \sigma(i,s,\neg\varphi)) \cap [\![\varphi]\!] = (\mathrm{B}(i,s) \cap [\![\varphi]\!]) \cup \sigma(i,s,\neg\varphi)$
- $\mathrm{B}_1(i,s) = \mathrm{B}_{11}(i,s) \cap [\![\psi]\!]$
- $\mathrm{B}_{21}(i,s) = \mathrm{B}(i,s) \cup \sigma(i,s,\neg\varphi \vee \neg\psi)$
- $\mathrm{B}_2(i,s) = (\mathrm{B}(i,s) \cup \sigma(i,s,\neg\varphi \vee \neg\psi)) \cap [\![\varphi \wedge \psi]\!] =$
  $(\mathrm{B}(i,s) \cap [\![\varphi \wedge \psi]\!]) \cup \sigma(i,s,\neg\varphi \vee \neg\psi)$

Now to prove that $\mathrm{B}_2(i,s) \subseteq \mathrm{B}_1(i,s)$. For then $\mathcal{B}(i,\mathcal{M}_1,s) \subseteq \mathcal{B}(i,\mathcal{M}_2,s)$ which means that $\mathcal{B}^{\star +}_{\varphi\psi}(i,\mathcal{M},s) \subseteq \mathcal{B}^\star_{\varphi\wedge\psi}(i,\mathcal{M},s)$. We distinguish two cases:

- $\neg\varphi \in \mathcal{B}(i,\mathcal{M},s)$. Then $\mathrm{B}(i,s) \subseteq [\![\neg\varphi]\!]$, and thus $\mathrm{B}(i,s) \cap [\![\varphi]\!] = \emptyset$. Hence $\mathrm{B}_{11}(i,s) = \sigma(i,s,\neg\varphi)$ and $\mathrm{B}_1(i,s) = \sigma(i,s,\neg\varphi) \cap [\![\psi]\!]$. In this case also $\neg\varphi \vee \neg\psi \in \mathcal{B}(i,\mathcal{M},s)$. Hence $\mathrm{B}(i,s) \cap [\![\varphi \wedge \psi]\!] = \emptyset$, and thus $\mathrm{B}_2(i,s) = \sigma(i,s,\neg\varphi \vee \neg\psi)$. Since $\neg\psi \notin \mathcal{B}^\star_\varphi(i,\mathcal{M},s)$, it follows that $\mathrm{B}_{11}(i,s) \cap [\![\psi]\!] \neq \emptyset$, and thus $\sigma(i,s,\neg\varphi) \cap [\![\psi]\!] \neq \emptyset$. From $\Sigma 1$ it follows that $\sigma(i,s,\neg\varphi) \cap [\![\varphi \wedge \psi]\!] = \sigma(i,s,\neg\varphi) \cap [\![\varphi]\!] \cap [\![\psi]\!] = \sigma(i,s,\neg\varphi) \cap [\![\psi]\!] \neq \emptyset$. Then since $\models \mathbf{K}_i((\neg\varphi \vee \neg\psi) \wedge \neg\varphi \leftrightarrow \neg\varphi)$, we have by $\Sigma 4$ that $\sigma(i,s,(\neg\varphi \vee \neg\psi) \wedge \neg\varphi) \cap [\![\varphi \wedge \psi]\!] \neq \emptyset$, and by $\Sigma 6$ that $\sigma(i,s,\neg\varphi \vee \neg\psi) \subseteq \sigma(i,s,(\neg\varphi \vee \neg\psi) \wedge \neg\varphi) = \sigma(i,s,\neg\varphi)$. Since by $\Sigma 1$, $\sigma(i,s,\neg\varphi \vee \neg\psi) \subseteq [\![\varphi \wedge \psi]\!]$, and given that $[\![\varphi \wedge \psi]\!] \subseteq [\![\psi]\!]$, it follows that $\sigma(i,s,\neg\varphi \vee \neg\psi) \subseteq \sigma(i,s,\neg\varphi) \cap [\![\psi]\!]$. Hence $\mathrm{B}_2(i,s) \subseteq \mathrm{B}_1(i,s)$.
- $\neg\varphi \notin \mathcal{B}(i,\mathcal{M},s)$. Note that in this case $\neg\varphi \vee \neg\psi \notin \mathcal{B}(i,\mathcal{M},s)$: for if $\neg\varphi \vee \neg\psi \in \mathcal{B}(i,\mathcal{M},s)$ and $\neg\varphi \notin \mathcal{B}(i,\mathcal{M},s)$, then from $B^\star 4$, $B^+ 2$ and $B^+ 3$ it follows that $\{\neg\varphi \vee \neg\psi, \varphi\} \subseteq \mathcal{B}^\star_\varphi(i,\mathcal{M},s)$. Since $\mathcal{B}^\star_\varphi(i,\mathcal{M},s)$ is deductively closed by $B^\star 1$ it follows that $\neg\psi \in \mathcal{B}^\star_\varphi(i,\mathcal{M},s)$ which contradicts the assumption that $\psi \notin \mathcal{B}^\star_\varphi(i,\mathcal{M},s)$. Hence $\neg\varphi \vee \neg\psi \notin \mathcal{B}(i,\mathcal{M},s)$. This implies that both $\mathrm{B}(i,s) \cap [\![\varphi]\!] \neq \emptyset$ and $\mathrm{B}(i,s) \cap [\![\varphi \wedge \psi]\!] \neq \emptyset$. Then it follows by $\Sigma 2$ that both $\sigma(i,s,\neg\varphi) \subseteq \mathrm{B}(i,s)$ and $\sigma(i,s,\neg\varphi \vee \neg\psi) \subseteq \mathrm{B}(i,s)$. Thus $\mathrm{B}_{11}(i,s) = \mathrm{B}(i,s) \cap [\![\varphi]\!]$, $\mathrm{B}_1(i,s) = (\mathrm{B}(i,s) \cap [\![\varphi]\!]) \cap [\![\psi]\!]$, $\mathrm{B}_{21}(i,s) = \mathrm{B}(i,s)$, and $\mathrm{B}_2(i,s) = \mathrm{B}(i,s) \cap [\![\varphi \wedge \psi]\!]$. Since for all $\mathcal{S}' \subseteq \mathcal{S}$ it holds that $(\mathcal{S}' \cap [\![\varphi]\!]) \cap [\![\psi]\!] = \mathcal{S}' \cap [\![\varphi \wedge \psi]\!]$, for all purely propositional formulae φ and ψ, it follows that $\mathrm{B}_1(i,s) = \mathrm{B}_2(i,s)$.

($B^\star 7$) $\mathcal{B}^\star_{\varphi\wedge\psi}(i,\mathcal{M},s) \subseteq \mathcal{B}^{\star +}_{\varphi\psi}(i,\mathcal{M},s)$.

($B^\star 8$) *If* $\neg\psi \notin \mathcal{B}^\star_\varphi(i,\mathcal{M},s)$, *then* $\mathcal{B}^{\star +}_{\varphi\psi}(i,\mathcal{M},s) \subseteq \mathcal{B}^\star_{\varphi\wedge\psi}(i,\mathcal{M},s)$.

Proof: Let $\mathcal{M}$ be some Kripke model with state s, and let i be some agent. Let σ be an arbitrary selection function for $\mathcal{M}$. Let φ be some propositional formula, and let $\mathcal{M}', s = \mathtt{r}(i, \mathtt{contract}\ \neg\varphi)(\mathcal{M}, s)$, let $\mathcal{M}'', s = \mathtt{r}(i, \mathtt{expand}\ \varphi)(\mathcal{M}', s) = \mathtt{r}(i, \mathtt{revise}\ \varphi)(\mathcal{M}, s)$, and let $\mathcal{M}''', s = \mathtt{r}(i, \mathtt{expand}\ \varphi)(\mathcal{M}, s)$.

($B^\star 1$) This postulate is shown in the same way as the corresponding postulates for belief expansion and contraction.

($B^\star 2$) Note that $\mathcal{B}^\star_\varphi(i,\mathcal{M},s) = \mathcal{B}(i,\mathcal{M}'',s) = \mathcal{B}^+_\varphi(i,\mathcal{M}',s)$. From B^+2 it follows that $\varphi \in \mathcal{B}^+_\varphi(i,\mathcal{M}',s)$, which suffices to conclude that the postulate is validated.

($B^\star 3$) By definition of $\mathtt{r}$ for $\mathtt{contract}$ and $\mathtt{expand}$ it follows that $\mathrm{B}''(i,s) = (\mathrm{B}(i,s) \cup \sigma(i,s,\neg\varphi)) \cap [\![\varphi]\!]$. Now if some formula $\psi \in \mathcal{B}^\star_\varphi(i,\mathcal{M},s)$, this means that $\mathcal{M}'', s'' \models \psi$ for all $s'' \in (\mathrm{B}(i,s) \cup \sigma(i,s,\neg\varphi)) \cap [\![\varphi]\!]$. Since ψ is purely propositional this implies that $\mathcal{M}, s'' \models \psi$ for all $s'' \in (\mathrm{B}(i,s) \cap [\![\varphi]\!])$. By definition of $\mathtt{r}(i, \mathtt{expand}\ \varphi)$ we have that $\mathrm{B}'''(i,s) = \mathrm{B}(i,s) \cap [\![\varphi]\!]$. But then $\mathcal{M}''', s''' \models \psi$ for all $s''' \in \mathrm{B}'''(i,s)$, and hence $\mathcal{M}''', s \models \mathbf{B}_i\psi$, which implies that $\psi \in \mathcal{B}^+_\varphi(i,\mathcal{M},s)$.

($B^\star 4$) If $\neg\varphi \notin \mathcal{B}(i,\mathcal{M},s)$, then $\mathrm{B}(i,s) \cap [\![\varphi]\!] \neq \emptyset$. From demand $\Sigma 2$ of selection functions, it follows that $\mathrm{B}'(i,s) = \mathrm{B}(i,s)$. Hence $\mathrm{B}''(i,s) = \mathrm{B}(i,s) \cap [\![\varphi]\!]$. Also $\mathrm{B}'''(i,s) = \mathrm{B}(i,s) \cap [\![\varphi]\!]$, and $\mathcal{M}'', s \models \mathbf{B}_i\psi$ iff $\mathcal{M}''', s \models \mathbf{B}_i\psi$ for all propositional formula ψ. Thus $\mathcal{B}^+_\varphi(i,\mathcal{M},s) = \mathcal{B}(i,\mathcal{M}''',s) = \mathcal{B}(i,\mathcal{M}'',s) = \mathcal{B}^\star_\varphi(i,\mathcal{M},s)$.

($B^\star 5$) We prove two implications.

- '$\Rightarrow$' Suppose $\mathcal{B}^\star_\varphi(i,\mathcal{M},s) = \mathcal{B}_\perp$. This implies that $\mathrm{B}''(i,s) = \emptyset$. Hence by definition of $\mathtt{r}$ for $\mathtt{contract}$ and $\mathtt{expand}$ this implies that $(\mathrm{B}(i,s) \cup \sigma(i,s,\neg\varphi)) \cap [\![\varphi]\!] = \emptyset$. In particular this implies that $\sigma(i,s,\neg\varphi) \cap [\![\varphi]\!] = \emptyset$, and since by $\Sigma 1$, $\sigma(i,s,\neg\varphi) \subseteq [\![\varphi]\!]$, we conclude that $\sigma(i,s,\neg\varphi) = \emptyset$. It follows by demand $\Sigma 3$ that $[s]_{\mathrm{R}(i)} \cap [\![\varphi]\!] = \emptyset$. This implies that $\mathcal{M}, s' \models \neg\varphi$ for all $s' \in [s]_{\mathrm{R}(i)}$ and thus $\mathcal{M}, s \models \mathbf{K}_i\neg\varphi$.
- '$\Leftarrow$' Suppose $\mathcal{M}, s \models \mathbf{K}_i\neg\varphi$. Then by demand $\Sigma 3$, $\sigma(i,s,\neg\varphi) = \emptyset$. Hence $\mathrm{B}(i,s) \subseteq [\![\neg\varphi]\!]$ and $\mathrm{B}'(i,s) \subseteq [\![\neg\varphi]\!]$. But then $\mathrm{B}''(i,s) = \mathrm{B}'(i,s) \cap [\![\varphi]\!] = \emptyset$, and hence $\mathcal{B}^\star_\varphi(i,\mathcal{M},s) = \mathcal{B}(i,\mathcal{M}'',s) = \mathcal{B}_\perp$.

($B^\star 6$) Suppose $\mathcal{M}, s \models \mathbf{K}_i(\varphi \leftrightarrow \psi)$. Then also $\mathcal{M}, s \models \mathbf{K}_i(\neg\varphi \leftrightarrow \neg\psi)$, and from demand $\Sigma 4$ it follows that $\sigma(i,s,\neg\varphi) = \sigma(i,s,\neg\psi)$. Hence $\mathtt{r}(i, \mathtt{contract}\ \varphi)(\mathcal{M}, s) = \mathtt{r}(i, \mathtt{contract}\ \psi)(\mathcal{M}, s) = \mathcal{M}', s$. Also $\mathrm{B}'(i,s) \cap [\![\varphi]\!] = \mathrm{B}'(i,s) \cap [\![\psi]\!]$, and $\mathtt{r}(i, \mathtt{expand}\ \varphi)(\mathcal{M}', s) = \mathtt{r}(i, \mathtt{expand}\ \psi)(\mathcal{M}', s)$. But then $\mathtt{r}(i, \mathtt{revise}\ \varphi)(\mathcal{M}, s) = \mathtt{r}(i, \mathtt{revise}\ \psi)(\mathcal{M}, s)$, and therefore $\mathcal{B}^\star_\varphi(i,\mathcal{M},s) = \mathcal{B}^\star_\psi(i,\mathcal{M},s)$.

($B^\star 7$) Assume that $\mathcal{M}_1, s = \mathtt{r}(i, \mathtt{revise}\ (\varphi \wedge \psi))(\mathcal{M}, s)$, and assume furthermore that $\mathcal{M}_2, s = \mathtt{r}(i, \mathtt{revise}\ \varphi; \mathtt{expand}\ \psi)(\mathcal{M}, s)$. From the definitions of $\mathtt{r}$ for revisions, contractions and expansions, it follows that $\mathrm{B}_1(i,s) = (\mathrm{B}(i,s) \cup \sigma(i,s,\neg\varphi \vee \neg\psi)) \cap [\![\varphi \wedge \psi]\!]$ and $\mathrm{B}_2(i,s) = (\mathrm{B}(i,s) \cup \sigma(i,s,\neg\varphi)) \cap [\![\varphi]\!] \cap [\![\psi]\!]$. Hence $\mathrm{B}_1(i,s) = (\mathrm{B}(i,s) \cap [\![\varphi \wedge \psi]\!]) \cup \sigma(i,s,\neg\varphi \vee \neg\psi)$ and $\mathrm{B}_2(i,s) = (\mathrm{B}(i,s) \cap [\![\varphi \wedge \psi]\!]) \cup (\sigma(i,s,\neg\varphi) \cap [\![\psi]\!])$. Hence, should

3.33. PROPOSITION. *For all agents i, and for all propositional formulae φ, ψ and ϑ we have:*

- $\models [\text{do}_i(\texttt{revise}\ \varphi)]\mathbf{B}_i\varphi$
- $\models [\text{do}_i(\texttt{revise}\ \varphi)]\mathbf{B}_i\vartheta \to [\text{do}_i(\texttt{expand}\ \varphi)]\mathbf{B}_i\vartheta$.
- $\models \neg\mathbf{B}_i\neg\varphi \to ([\text{do}_i(\texttt{expand}\ \varphi)]\mathbf{B}_i\vartheta \leftrightarrow [\text{do}_i(\texttt{revise}\ \varphi)]\mathbf{B}_i\vartheta)$.
- $\models \mathbf{K}_i\neg\varphi \leftrightarrow [\text{do}_i(\texttt{revise}\ \varphi)]\mathbf{B}_i\mathbf{ff}$.
- $\models \mathbf{K}_i(\varphi \leftrightarrow \psi) \to ([\text{do}_i(\texttt{revise}\ \varphi)]\mathbf{B}_i\vartheta \leftrightarrow [\text{do}_i(\texttt{revise}\ \psi)]\mathbf{B}_i\vartheta)$.
- $\models [\text{do}_i(\texttt{revise}\ (\varphi \wedge \psi))]\mathbf{B}_i\vartheta \to [\text{do}_i(\texttt{revise}\ \varphi; \texttt{expand}\ \psi)]\mathbf{B}_i\vartheta$.
- $\models \neg[\text{do}_i(\texttt{revise}\ \varphi)]\mathbf{B}_i\neg\psi \to$
  $([\text{do}_i(\texttt{revise}\ \varphi; \texttt{expand}\ \psi)]\mathbf{B}_i\vartheta \to [\text{do}_i(\texttt{revise}\ (\varphi \wedge \psi))]\mathbf{B}_i\vartheta)$.

PROOF: All clauses follow directly from Theorem 3.38. The first clause of Proposition 3.16 follows from $B^\star 2$ and so on until the last clause that follows from $B^\star 8$.
⊠

3.34. PROPOSITION. *For all agents i, for all propositional formulae φ, and for all formulae ψ we have:*

- $\models \langle\text{do}_i(\texttt{revise}\ \varphi)\rangle\mathbf{tt}$
- $\models \langle\text{do}_i(\texttt{revise}\ \varphi)\rangle\psi \leftrightarrow [\text{do}_i(\texttt{revise}\ \varphi)]\psi$
- $\models \langle\text{do}_i(\texttt{revise}\ \varphi; \texttt{revise}\ \varphi)\rangle\psi \leftrightarrow \langle\text{do}_i(\texttt{revise}\ \varphi)\rangle\psi$

PROOF: Directly from Def. 3.32.
⊠

3.35. PROPOSITION. *For all agents i, and for all propositional formulae φ and ψ we have:*

- $\models \mathbf{K}_i\neg\varphi \leftrightarrow [\text{do}_i(\texttt{revise}\ \varphi)]\mathbf{B}_i\mathbf{ff}$
- $\models \neg\mathbf{B}_i\neg\varphi \to ([\text{do}_i(\texttt{revise}\ \varphi)]\mathbf{B}_i\psi \leftrightarrow \mathbf{B}_i(\varphi \to \psi))$
- $\models \neg\mathbf{K}_i\neg\varphi \wedge \mathbf{B}_i\neg\varphi \to ([\text{do}_i(\texttt{revise}\ \varphi)]\mathbf{B}_i\psi \leftrightarrow \mathbf{K}_i(\varphi \to \psi))$ *if the definition of* r *for the* contract *action is for all models based on the AiG function.*

PROOF: The first clause follows directly from the first clause of Proposition 3.39. The second clause is proved by the same argument as given in the proof of Proposition 3.8. The third clause is analogous to the second clause: note that by the Deduction Theorem for propositional classical logic $\psi \in \text{Th}(\mathcal{K}(i, \mathcal{M}, s) \cup \{\varphi\})$ is equivalent with $\varphi \to \psi \in \text{Th}(\mathcal{K}(i, \mathcal{M}, s))$.
⊠

3.38. THEOREM. *Let $\mathcal{M}$ be some Kripke model. For all agents i, for all $s \in \mathcal{M}$ and for all formulae φ and ψ the following are true.*

($B^\star 1$) *$\mathcal{B}^\star_\varphi(i, \mathcal{M}, s)$ is a belief set.*
($B^\star 2$) *$\varphi \in \mathcal{B}^\star_\varphi(i, \mathcal{M}, s)$.*
($B^\star 3$) *$\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) \subseteq \mathcal{B}^+_\varphi(i, \mathcal{M}, s)$.*
($B^\star 4$) *If $\neg\varphi \notin \mathcal{B}(i, \mathcal{M}, s)$ then $\mathcal{B}^+_\varphi(i, \mathcal{M}, s) \subseteq \mathcal{B}^\star_\varphi(i, \mathcal{M}, s)$.*
($B^\star 5$) *$\mathcal{B}(i, \mathcal{M}, s) = \mathcal{B}_\perp$ if and only if $\mathcal{M}, s \models \mathbf{K}_i\neg\varphi$.*
($B^\star 6$) *If $\mathcal{M}, s \models \mathbf{K}_i(\varphi \leftrightarrow \psi)$ then $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) = \mathcal{B}^\star_\psi(i, \mathcal{M}, s)$.*

x	$\varsigma_0(x)$	x	$\varsigma_0(x)$
$\rho_1 : \bot$	$\{s_0\}$	$\rho_9 : \neg p \wedge \neg q$	$\{s_0\}$
$\rho_2 : p \wedge q$	$\{s_1, s_3\}$	$\rho_{10} : (\neg p \wedge \neg q) \vee (p \wedge q)$	$\{s_1, s_2\}$
$\rho_3 : p \wedge \neg q$	$\{s_0\}$	$\rho_{11} : \neg q$	$\{s_0\}$
$\rho_4 : p$	$\{s_2, s_3\}$	$\rho_{12} : \neg(\neg p \wedge q)$	$\{s_2\}$
$\rho_5 : \neg p \wedge q$	$\{s_0\}$	$\rho_{13} : \neg p$	$\{s_0\}$
$\rho_6 : q$	$\{s_1, s_3\}$	$\rho_{14} : \neg(p \wedge \neg q)$	$\{s_1\}$
$\rho_7 : (\neg p \wedge q) \vee (p \wedge \neg q)$	$\{s_0\}$	$\rho_{15} : \neg p \vee \neg q$	$\{s_0\}$
$\rho_8 : p \vee q$	$\{s_3\}$	$\rho_{16} : \top$	$\emptyset$

Define for all $s \in \mathcal{S}$, $\varsigma(1, s, \varphi) = \varsigma(\rho_j)$ for the unique ρ_j such that $\varphi \in [\rho_j]_{cl}$, where $[\rho_j]_{cl}$ denotes the equivalence class of ρ_j in classical propositional logic.

A.2. Lemma. *For ς as defined above we have:*

- *ς is an s-selection function for $\mathcal{M}$.*
- *ς is not a selection function for $\mathcal{M}$.*

Proof: We successively show both clauses.

- It is easily checked that ς validates the demands $S1$ through $S3$ as given in Def. 3.14, leaving only $S4$ to be shown. To prove $S4$ it needs to be shown that for any pair φ, ψ of formulae holds that
  $\varsigma(1, s, \varphi) \subseteq [\![\neg\psi]\!] \;\&\; \varsigma(1, s, \psi) \subseteq [\![\neg\varphi]\!] \Rightarrow$
  $\varsigma(1, s, \varphi) = \varsigma(1, s, \psi)$ (†)
  If at least one of φ and ψ is in any of $[\rho_{2j+1}]_{cl}$ for $j = 0 \ldots 7$ it is easily checked that (†) indeed holds. The case where one of φ and ψ is in $[\rho_{16}]_{cl}$ is trivial, leaving only the cases where both φ and ψ are in any of $[\rho_{2j}]_{cl}$ for $j = 1 \ldots 7$. We show the cases where $\varphi \in [\rho_2]_{cl}$ and $\psi \in [\rho_{2j}]_{cl}$ for $j = 2 \ldots 7$. All other cases can be checked in a similar fashion. So let $\varphi \in [\rho_2]_{cl}$; we distinguish the six cases for ψ:
  - $\psi \in [\rho_4]_{cl}$: since $\varsigma_0(\rho_2) \not\subseteq [\![\neg\rho_4]\!]$ this case goes through.
  - $\psi \in [\rho_6]_{cl}$: since $\varsigma_0(\rho_6) = \varsigma_0(\rho_2)$ this case goes through.
  - $\psi \in [\rho_8]_{cl}$: since $\varsigma_0(\rho_2) \not\subseteq [\![\neg\rho_8]\!]$ this case goes through.
  - $\psi \in [\rho_{10}]_{cl}$: since $\varsigma_0(\rho_2) \not\subseteq [\![\neg\rho_{10}]\!]$ this case goes through.
  - $\psi \in [\rho_{12}]_{cl}$: since $\varsigma_0(\rho_2) \not\subseteq [\![\neg\rho_{12}]\!]$ this case goes through.
  - $\psi \in [\rho_{14}]_{cl}$: since $\varsigma_0(\rho_2) \not\subseteq [\![\neg\rho_{14}]\!]$ this case goes through.

  Hence ς is indeed an s-selection function for $\mathcal{M}$.
- Note that $\varsigma(i, s, p \wedge q) = \{s_1, s_3\}$, and hence $\varsigma(i, s, p \wedge q) \cap [\![\neg p]\!] = \{s_3\} \neq \emptyset$. However $\varsigma(i, s, p) = \{s_2, s_3\} \not\subseteq \varsigma(i, s, p \wedge q)$, and hence ς does not validate demand $\Sigma 6$.

⊠

Defining the semantics for the `contract` action based on the function ς would yield that:

- $\mathcal{B}^-_{p\wedge q}(1, \mathcal{M}, s_0) = \mathrm{Th}(\{(p \wedge q) \vee (p \wedge \neg q) \vee (\neg p \wedge \neg q)\})$
- $\mathcal{B}^-_p(1, \mathcal{M}, s_0) = \mathrm{Th}(\{(p \wedge q) \vee (\neg p \wedge q) \vee (\neg p \wedge \neg q)\})$

Hence although $p \notin \mathcal{B}^-_{p\wedge q}(i, \mathcal{M}, s_0)$, $\mathcal{B}^-_p(i, \mathcal{M}, s_0) \not\subseteq \mathcal{B}^-_{p\wedge q}(1, \mathcal{M}, s_0)$, thereby contradicting postulate B^-8.

⊠

3.30. PROPOSITION. *For all models $\mathcal{M}$, and for all functions $\varsigma : \mathcal{A} \times \mathcal{S} \times \mathcal{L}_0 \to \wp(\mathcal{S})$ it holds that if ς is a selection function for $\mathcal{M}$ then ς is an s-selection function for $\mathcal{M}$.*

PROOF: Let some Kripke model $\mathcal{M}$ and function ς be given such that ς is a selection function for $\mathcal{M}$. We show that ς satisfies the demands for an s-selection function. The properties $S1$ through $S3$ follow directly from demands $\Sigma 1$ through $\Sigma 3$, leaving only $S4$ to be proved. Hence assume that $\neg\varphi$ is true at all states from $\varsigma(i,s,\psi)$ and $\neg\psi$ is true at all states from $\varsigma(i,s,\varphi)$. This implies that $\varsigma(i,s,\psi) \subseteq [\![\neg\varphi]\!]$ and $\varsigma(i,s,\varphi) \subseteq [\![\neg\psi]\!]$. From $\Sigma 1$ it follows that $\varsigma(i,s,\varphi) \subseteq [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]$, and hence $\varsigma(i,s,\varphi) \subseteq [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] \cap [\![\neg\psi]\!])$. Analogously it follows that $\varsigma(i,s,\psi) \subseteq [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] \cap [\![\neg\psi]\!]$. If either $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] = \emptyset$ or $[s]_{\mathrm{R}(i)} \cap [\![\neg\psi]\!] = \emptyset$, then both $\varsigma(i,s,\varphi) = \emptyset$ and $\varsigma(i,s,\psi) = \emptyset$ and hence $S4$ would be met. Hence assume that $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] \neq \emptyset$ and $[s]_{\mathrm{R}(i)} \cap [\![\neg\psi]\!] \neq \emptyset$. Then also $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi \vee \neg\psi]\!] \neq \emptyset$. From $\Sigma 3$ it follows that none of $\varsigma(i,s,\varphi)$, $\varsigma(i,s,\psi)$ and $\varsigma(i,s,\varphi \wedge \psi)$ is empty. By $\Sigma 5$ we have that $\varsigma(i,s,\varphi \wedge \psi) \subseteq \varsigma(i,s,\varphi) \cup \varsigma(i,s,\psi)$. Hence $\varsigma(i,s,\varphi \wedge \psi) \subseteq ([s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] \cap [\![\neg\psi]\!])$ (†). Then $\varsigma(i,s,\varphi \wedge \psi) \cap [\![\neg\varphi]\!] \neq \emptyset$, and hence by $\Sigma 6$ we have that $\varsigma(i,s,\varphi) \subseteq \varsigma(i,s,\varphi \wedge \psi)$ (‡). Analogously we have that $\varsigma(i,s,\psi) \subseteq \varsigma(i,s,\varphi \wedge \psi)$. Hence $\varsigma(i,s,\varphi \wedge \psi) = \varsigma(i,s,\varphi) \cup \varsigma(i,s,\psi)$.

Now $\models \mathbf{K}_i((\neg\varphi \vee \psi) \wedge \varphi \leftrightarrow \varphi \wedge \psi)$. Hence by $\Sigma 4$, $\varsigma(i,s,\varphi \wedge \psi) = \varsigma(i,s,(\neg\varphi \vee \psi) \wedge \varphi)$. From $\Sigma 5$ it follows that $\varsigma(i,s,\varphi \wedge \psi) \subseteq \varsigma(i,s,\neg\varphi \vee \psi) \cup \varsigma(i,s,\varphi)$. Since by $\Sigma 1$, $\varsigma(i,s,\neg\varphi \vee \psi) \subseteq [\![\varphi \wedge \neg\psi]\!]$, it follows from (†) that $\varsigma(i,s,\varphi \wedge \psi) \cap \varsigma(i,s,\neg\varphi \vee \psi) = \emptyset$. Hence $\varsigma(i,s,\varphi \wedge \psi) \subseteq \varsigma(i,s,\varphi)$. Combining this with (‡) yields that $\varsigma(i,s,\varphi \wedge \psi) = \varsigma(i,s,\varphi)$. By an analogous argument we conclude from $\models \mathbf{K}_i((\neg\psi \vee \varphi) \wedge \psi \leftrightarrow \varphi \wedge \psi)$ that $\varsigma(i,s,\varphi \wedge \psi) = \varsigma(i,s,\psi)$. Thus $\varsigma(i,s,\varphi) = \varsigma(i,s,\psi)$, which suffices to conclude that ς validates $S4$. Thus ς is an s-selection function.
⊠

3.31. PROPOSITION. *Some Kripke model $\mathcal{M}$ and function $\varsigma : \mathcal{A} \times \mathcal{S} \times \mathcal{L}_0 \to \wp(\mathcal{S})$ exists, such that ς is a s-selection function for $\mathcal{M}$, but not a selection function. Furthermore, when defining contractions based on the function ς, not all Gärdenfors postulates for belief contraction are validated.*

PROOF: Consider the language $\mathcal{L}$ based on the sets $\Pi = \{p, q\}$ $\mathcal{A} = \{1\}$, and At is arbitrary. Consider the Kripke model $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathtt{c} \rangle$ such that

- $\mathcal{S} = \{s_0, \ldots, s_3\}$,
- $\pi(p, s_j) = \mathbf{1}$ iff $j \in \{0, 1\}$, $\pi(q, s_j) = \mathbf{1}$ iff $j \in \{0, 2\}$,
- $\mathrm{R}(1) = \mathcal{S}^2$,
- $\mathrm{B}(1, s) = \{s_0\}$,
- $\mathtt{r}$ and $\mathtt{c}$ are arbitrary.

Define the function ς_0 as follows:

case of the proof of $\Sigma 5$, we conclude that $\sigma_a(i,s,\varphi\wedge\psi) = \sigma_a(i,s,\varphi)\cup\sigma_a(i,s,\psi)$.

- $\mathrm{B}(i,s)\cap[\![\neg\varphi\vee\neg\psi]\!]\neq\emptyset$. In this case $\sigma_a(i,s,\varphi\wedge\psi) = \mathrm{B}(i,s)\cap[\![\neg\varphi\vee\neg\psi]\!]$, and since $\sigma_a(i,s,\varphi\wedge\psi)\cap[\![\neg\varphi]\!]\neq\emptyset$, also $\mathrm{B}(i,s)\cap[\![\neg\varphi]\!]\neq\emptyset$. Hence $\sigma_a(i,s,\varphi) = \mathrm{B}(i,s)\cap[\![\neg\varphi]\!]$. Now if $\mathrm{B}(i,s)\cap[\![\neg\psi]\!]=\emptyset$ it follows by an identical argument as given in the third clause of the proof of $\Sigma 5$ that $\sigma_a(i,s,\varphi\wedge\psi)=\sigma_a(i,s,\varphi)$. If $\mathrm{B}(i,s)\cap[\![\neg\psi]\!]\neq\emptyset$ it follows by an identical argument as given in the fourth clause of the proof of $\Sigma 5$ that $\sigma_a(i,s,\varphi\wedge\psi)=\sigma_a(i,s,\varphi)\cup\sigma_a(i,s,\psi)$.

Since in both cases $\sigma_a(i,s,\varphi)\subseteq\sigma_a(i,s,\varphi\wedge\psi)$ we conclude that demand $\Sigma 6$ is validated.

⊠

3.27. Proposition. *For all agents i, and for all propositional formulae φ and ψ we have:*

- $\models \neg\mathbf{B}_i\varphi \rightarrow ([\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{B}_i\psi \leftrightarrow \mathbf{B}_i\psi)$
- $\models \mathbf{B}_i\varphi \rightarrow ([\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{B}_i\psi \leftrightarrow (\mathbf{B}_i\psi \wedge \mathbf{K}_i(\neg\varphi\rightarrow\psi)))$ *if the definition of* `r` *for the* `contract` *action is for all models based on the AiG function.*

Proof: Directly from Proposition 3.28.

⊠

3.28. Proposition. *For all Kripke models $\mathcal{M}$ with state s, for all agents i, and for all propositional formulae φ and ψ we have:*

- $\varphi\notin\mathcal{B}(i,\mathcal{M},s)\Rightarrow(\psi\in\mathcal{B}^-_\varphi(i,\mathcal{M},s)\Leftrightarrow\psi\in\mathcal{B}(i,\mathcal{M},s))$
- $\varphi\in\mathcal{B}(i,\mathcal{M},s)\Rightarrow$ $(\psi\in\mathcal{B}^-_\varphi(i,\mathcal{M},s)\Leftrightarrow\psi\in\mathcal{B}(i,\mathcal{M},s)\,\&\,(\neg\varphi\rightarrow\psi)\in\mathcal{K}(i,\mathcal{M},s))$ *if the definition of* `r` *for the* `contract` *action is based on the AiG function for $\mathcal{M}$.*

Proof: Let $\mathcal{M}$ be some Kripke model with state s, let i be some agent, and let φ,ψ be propositional formula.

- Suppose $\varphi\notin\mathcal{B}(i,\mathcal{M},s)$. By B^-3 we have that $\mathcal{B}^-_\varphi(i,\mathcal{M},s)=\mathcal{B}(i,\mathcal{M},s)$, and hence $\psi\in\mathcal{B}^-_\varphi(i,\mathcal{M},s)$ iff $\psi\in\mathcal{B}(i,\mathcal{M},s)$.
- Suppose $\varphi\in\mathcal{B}(i,\mathcal{M},s)$. Let $\mathcal{M}',s=\texttt{r}(i,\texttt{contract}\ \varphi)(\mathcal{M},s)$.

$$\begin{aligned}
&\psi\in\mathcal{B}^-_\varphi(i,\mathcal{M},s)\\
\Leftrightarrow\ &\psi\in\mathcal{B}(i,\mathcal{M}',s)\\
\Leftrightarrow\ &\mathcal{M}',s'\models\psi \text{ for all } s'\in\mathrm{B}'(i,s)\\
\Leftrightarrow\ &\mathcal{M}',s'\models\psi \text{ for all } s'\in\mathrm{B}(i,s)\cup\sigma_a(i,s,\varphi)\\
\Leftrightarrow\ &\mathcal{M}',s'\models\psi \text{ for all } s'\in\mathrm{B}(i,s) \text{ and } \mathcal{M}',s'\models\psi \text{ for all } s'\in\sigma_a(i,s,\varphi)\\
\Leftrightarrow\ &\mathcal{M},s'\models\psi \text{ for all } s'\in\mathrm{B}(i,s) \text{ and } \mathcal{M},s'\models\psi \text{ for all } s'\in[s]_{\mathrm{R}(i)}\cap[\![\neg\varphi]\!]\\
\Leftrightarrow\ &\mathcal{M},s\models\mathbf{B}_i\psi \text{ and } \mathcal{M},s\models\mathbf{K}_i(\neg\varphi\rightarrow\psi)\\
\Leftrightarrow\ &\psi\in\mathcal{B}(i,\mathcal{M},s) \text{ and } \neg\varphi\rightarrow\psi\in\mathcal{K}(i,\mathcal{M},s)
\end{aligned}$$

⊠

- $\varphi \notin \mathcal{B}(i, \mathcal{M}, s)$. In this case $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] \neq \emptyset$. Hence by $\Sigma 2$, $\sigma(i,s,\varphi) \subseteq \mathrm{B}(i,s)$. Then $S(\mathcal{B}(i,\mathcal{M},s)\bot\varphi) = \{\{\psi \in \mathcal{L}_0 \mid \forall s' \in \mathrm{B}(i,s)[\mathcal{M},s' \models \psi]\}\}$. Hence $\cap S(\mathcal{B}(i,\mathcal{M},s)\bot\varphi) = \{\psi \in \mathcal{L}_0 \mid \forall s' \in \mathrm{B}(i,s)[\mathcal{M},s \models \psi]\} = \mathcal{B}(i,\mathcal{M},s)$. If $\varphi \notin \mathcal{B}(i,\mathcal{M},s)$ it follows by B^-3 that $\mathcal{B}^-_\varphi(i,\mathcal{M},s) = \mathcal{B}(i,\mathcal{M},s)$ and thus $\mathcal{B}^-_\varphi(i,\mathcal{M},s) = \cap S(\mathcal{B}(i,\mathcal{M},s)\bot\varphi)$.
- $\varphi \in \mathcal{B}(i,\mathcal{M},s)$. Let $\rho \in \mathcal{B}^-_\varphi(i,\mathcal{M},s)$ be arbitrary. Then $\mathcal{M},s' \models \rho$ for all $s' \in \mathrm{B}(i,s) \cup \sigma(i,s,\varphi)$. But then $\rho \in \{\psi \in \mathcal{L}_0 \mid \forall s' \in \mathrm{B}(i,s) \cup \{s''\}[\mathcal{M},s' \models \psi]\}$ for all $s'' \in \sigma(i,s,\varphi)$. Hence $\rho \in \cap\{\{\psi \in \mathcal{L}_0 \mid \forall s' \in \mathrm{B}(i,s) \cup \{s''\}[\mathcal{M},s' \models \psi]\} \mid s'' \in \sigma(i,s,\varphi)\}$. Also for all $\rho \in \cap\{\{\psi \in \mathcal{L}_0 \mid \forall s' \in \mathrm{B}(i,s) \cup \{s''\}[\mathcal{M},s' \models \psi]\} \mid s'' \in \sigma(i,s,\varphi)\}$ it holds that $\mathcal{M},s' \models \rho$ for all $s' \in \mathrm{B}(i,s) \cup \sigma(i,s,\varphi)$, and thus $\mathcal{B}^-_\varphi(i,\mathcal{M},s) = \cap\{\{\psi \in \mathcal{L}_0 \mid \forall s' \in \mathrm{B}(i,s) \cup \{s''\}[\mathcal{M},s' \models \psi]\} \mid s'' \in \sigma(i,s,\varphi)\}$.

⊠

3.26. PROPOSITION. *The AiG function σ_a as given in Def. 3.25 is a selection function.*

PROOF: We successively show that the AiG function satisfies the demands for selection functions. So assume that σ_a is the AiG function for some model $\mathcal{M}$ with states s, s', and let i be some agent and φ and ψ be propositional formulae.

$\Sigma 0$. Suppose $s' \in [s]_{\mathrm{R}(i)}$. Then $\mathrm{B}(i,s) = \mathrm{B}(i,s')$ and $[s]_{\mathrm{R}(i)} = [s']_{\mathrm{R}(i)}$ and hence demand $\Sigma 0$ is met.

$\Sigma 1$. Since $\sigma_a(i,s,\varphi) = \mathrm{B}(i,s) \cap [\![\neg\varphi]\!] \subseteq [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]$ if $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] \neq \emptyset$, and $\sigma_a(i,s,\varphi) = [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]$ otherwise, demand $\Sigma 1$ is indeed met.

$\Sigma 2$. Since $\sigma_a(i,s,\varphi) = \mathrm{B}(i,s) \cap [\![\neg\varphi]\!] \subseteq \mathrm{B}(i,s)$ if $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] \neq \emptyset$, demand $\Sigma 2$ is obviously satisfied.

$\Sigma 3$. Demand $\Sigma 3$ follows directly from the definition of the AiG function.

$\Sigma 4$. If $\mathcal{M},s \models \mathbf{K}_i(\varphi \leftrightarrow \psi)$, then both $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] = \mathrm{B}(i,s) \cap [\![\neg\psi]\!]$ and $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] = [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]$, which suffices to conclude $\Sigma 4$.

$\Sigma 5$. We distinguish four cases:
  - $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] = \emptyset, \mathrm{B}(i,s) \cap [\![\neg\psi]\!] = \emptyset$. In this case also $\mathrm{B}(i,s) \cap [\![\neg\varphi \vee \neg\psi]\!] = \emptyset$. Hence $\sigma_a(i,s,\varphi \wedge \psi) = [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi \vee \neg\psi]\!] = [s]_{\mathrm{R}(i)} \cap ([\![\neg\varphi]\!] \cup [\![\neg\psi]\!]) = ([s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]) \cup ([s]_{\mathrm{R}(i)} \cap [\![\neg\psi]\!]) = \sigma_a(i,s,\varphi) \cup \sigma_a(i,s,\psi)$.
  - $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] = \emptyset, \mathrm{B}(i,s) \cap [\![\neg\psi]\!] \neq \emptyset$. In this case $\mathrm{B}(i,s) \cap [\![\neg\varphi \vee \neg\psi]\!] = \mathrm{B}(i,s) \cap [\![\neg\psi]\!]$. Since $\mathrm{B}(i,s) \cap [\![\neg\psi]\!] \neq \emptyset$, it follows that $\sigma_a(i,s,\varphi \wedge \psi) = \sigma_a(i,s,\psi)$.
  - $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] \neq \emptyset, \mathrm{B}(i,s) \cap [\![\neg\psi]\!] = \emptyset$. This case is completely analogous to the previous one.
  - $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] \neq \emptyset, \mathrm{B}(i,s) \cap [\![\neg\psi]\!] \neq \emptyset$. Then also $\mathrm{B}(i,s) \cap [\![\neg(\varphi \wedge \psi)]\!] \neq \emptyset$. In this case $\sigma_a(i,s,\varphi \wedge \psi) = \mathrm{B}(i,s) \cap [\![\neg\varphi \vee \neg\psi]\!] = \mathrm{B}(i,s) \cap ([\![\neg\varphi]\!] \cup [\![\neg\psi]\!]) = (\mathrm{B}(i,s) \cap [\![\neg\varphi]\!]) \cup (\mathrm{B}(i,s) \cap [\![\neg\psi]\!]) = \sigma_a(i,s,\varphi) \cup \sigma_a(i,s,\psi)$.

  Since in all four cases $\sigma_a(i,s,\varphi \wedge \psi) \subseteq \sigma_a(i,s,\varphi) \cup \sigma_a(i,s,\psi)$, we conclude that $\Sigma 5$ is validated.

$\Sigma 6$. We distinguish two cases:
  - $\mathrm{B}(i,s) \cap [\![\neg\varphi \vee \neg\psi]\!] = \emptyset$. In this case both $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] = \emptyset$ and $\mathrm{B}(i,s) \cap [\![\neg\psi]\!] = \emptyset$, and by an identical argument as given in the first

as in 3.15 but with σ replaced by an arbitrary function $\varsigma : \mathcal{A} \times \mathcal{S} \times \mathcal{L}_0 \to \wp(\mathcal{S})$. Then it holds that:

- *If ς satisfies $\Sigma 1$ then $B^{-}5$ is validated.*
- *If ς satisfies $\Sigma 2$ then $B^{-}3$ is validated.*
- *Given that ς satisfies $\Sigma 1$ it is the case that if ς satisfies $\Sigma 3$ then $B^{-}4$ is validated.*
- *If ς satisfies $\Sigma 4$ then $B^{-}6$ is validated.*
- *If ς satisfies $\Sigma 5$ then $B^{-}7$ is validated.*
- *If ς satisfies $\Sigma 6$ then $B^{-}8$ is validated.*

PROOF: The proof is a straightforward modification of the proof of Theorem 3.20 and therefore left to the reader.
⊠

3.23. PROPOSITION. *Let $\mathcal{M}$ be some Kripke model with state s, and let i be some agent. Assume that $\mathtt{r}(i, \mathtt{contract}\ \varphi)(\mathcal{M}, s)$ is defined as in 3.15 with σ replaced by an arbitrary function ς. Then if* `contract` *is to meet the demands presented in Theorem 3.20 it follows that:*

- $\varsigma(i, s, \varphi) \subseteq [s]_{\mathrm{R}(i)}$
- $\mathrm{B}(i, s) \cap [\![\neg\varphi]\!] = \emptyset \,\&\, [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] \neq \emptyset \Rightarrow \varsigma(i, s, \varphi) \cap [\![\neg\varphi]\!] \neq \emptyset$

PROOF: We successively show both cases. So let $\mathcal{M}$ be some Kripke model with state s, let i be some agent and let $\varsigma : \mathcal{A} \times \mathcal{S} \times \mathcal{L}_0 \to \wp(\mathcal{S})$ be some function.

- Since the result of a contraction is by $B^{-}1$ forced to be a well-defined Kripke model, it holds that in the model that results from a contraction the set of doxastic alternatives is contained in the set of epistemic alternatives. Hence the worlds that are to be added -these are the worlds from $\varsigma(i, s, \varphi)$- are to be a subset of the set of epistemic alternatives $[s]_{\mathrm{R}(i)}$.
- Assume that $\mathrm{B}(i, s) \cap [\![\neg\varphi]\!] = \emptyset$ and $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] \neq \emptyset$, for some $\varphi \in \mathcal{L}_0$. Then $\varphi \in \mathcal{B}(i, \mathcal{M}, s)$ and $\mathcal{M}, s \not\models \mathbf{K}_i\varphi$, hence from $B^{-}4$ it follows that in the model that results from the contraction, agent i does no longer believe φ in state s. But this implies that some world not supporting φ must have been added to the set of doxastic alternatives of the agent, and thus $\varsigma(i, s, \varphi) \cap [\![\neg\varphi]\!] \neq \emptyset$.

⊠

3.24. PROPOSITION. *Let some Kripke model $\mathcal{M}$ with state s, and agent i be given. Let σ be some selection function for $\mathcal{M}$. Define for propositional formulae φ:*

- $\mathcal{B}(i, \mathcal{M}, s)\perp\varphi = \mathcal{B}(i, \mathcal{M}, s)$ *if* $\varphi \in \mathcal{K}(i, \mathcal{M}, s)$
- $\mathcal{B}(i, \mathcal{M}, s)\perp\varphi = \{\{\psi \in \mathcal{L}_0 \mid \forall s' \in \mathrm{B}(i, s) \cup \{s''\}[\mathcal{M}, s' \models \psi]\} \mid s'' \in [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]\}$ *if* $\varphi \notin \mathcal{K}(i, \mathcal{M}, s)$
- $S(\mathcal{B}(i, \mathcal{M}, s)\perp\varphi) = \{\{\psi \in \mathcal{L}_0 \mid \forall s' \in \mathrm{B}(i, s) \cup \{s''\}[\mathcal{M}, s' \models \psi]\} \mid s'' \in \sigma(i, s, \varphi)\}$

Then $\mathcal{B}^{-}_{\varphi}(i, \mathcal{M}, s) = \cap S(\mathcal{B}(i, \mathcal{M}, s)\perp\varphi)$.

PROOF: Let $\mathcal{M}$ be some Kripke model, let i be some agent, and let φ be some propositional formula. We distinguish two cases.

(B^-3) If $\varphi \notin \mathcal{B}(i,\mathcal{M},s)$, then $\mathcal{M},s' \models \neg\varphi$ for some $s' \in \mathrm{B}(i,s)$. Then $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] \neq \emptyset$, and by $\Sigma 2$ it follows that $\sigma(i,s,\varphi) \subseteq \mathrm{B}(i,s)$. Thus $\mathrm{B}'(i,s) = \mathrm{B}(i,s)$ and hence $\mathcal{B}^-_\varphi(i,\mathcal{M},s) = \mathcal{B}(i,\mathcal{M},s)$.

(B^-4) If $\mathcal{M},s \not\models \mathbf{K}_i\varphi$, then $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] \neq \emptyset$. Hence by $\Sigma 3$, $\sigma(i,s,\varphi) \neq \emptyset$, and thus, by $\Sigma 1$, $\mathrm{B}'(i,s)$ contains some s' such that $\mathcal{M},s' \models \neg\varphi$. Since φ is propositional, then also $\mathcal{M}',s' \models \neg\varphi$, and hence $\mathcal{M}',s \not\models \mathbf{B}_i\varphi$. Thus $\varphi \notin \mathcal{B}^-_\varphi(i,\mathcal{M},s)$.

(B^-5) Suppose $\varphi \in \mathcal{B}(i,\mathcal{M},s)$. We distinguish two cases:

- $\mathcal{M},s \models \mathbf{K}_i\varphi$. Then $\sigma(i,s,\varphi) = \emptyset$ by $\Sigma 3$. Hence $\mathrm{B}'(i,s) = \mathrm{B}(i,s)$, and $\mathcal{B}^-_\varphi(i,\mathcal{M},s) = \mathcal{B}(i,\mathcal{M},s)$. Then $\mathcal{B}^{-+}_{\varphi\varphi}(i,\mathcal{M},s) = \mathcal{B}^+_\varphi(i,\mathcal{M},s)$. Now since $\varphi \in \mathcal{B}(i,\mathcal{M},s)$ it follows by B^+4 that $\mathcal{B}^{-+}_{\varphi\varphi}(i,\mathcal{M},s) = \mathcal{B}(i,\mathcal{M},s)$.
- $\mathcal{M},s \not\models \mathbf{K}_i\varphi$. Then $\mathrm{B}'(i,s) = \mathrm{B}(i,s) \cup \mathcal{S}'$ where $\mathcal{S}' = \sigma(i,s,\varphi)$. From $\Sigma 1$ it follows that $\mathcal{S}' \subseteq [\![\neg\varphi]\!]$. Let $\mathcal{M}'',s = \mathtt{r}(i,\mathtt{expand}\ \varphi)(\mathcal{M}',s)$. Then by Def. 3.5 it follows that $\mathrm{B}''(i,s) = \mathrm{B}'(i,s) \cap [\![\varphi]\!]$. Now since $\mathrm{B}'(i,s) = \mathrm{B}(i,s) \cup \mathcal{S}'$ and since $\varphi \in \mathcal{B}(i,\mathcal{M},s)$, and thus $\mathrm{B}(i,s) \cap [\![\varphi]\!] = \mathrm{B}(i,s)$, we have that $\mathrm{B}''(i,s) = \mathrm{B}(i,s)$. Hence $\mathcal{B}^{-+}_{\varphi\varphi}(i,\mathcal{M},s) = \mathcal{B}(i,\mathcal{M},s)$.

Since in both cases $\mathcal{B}(i,\mathcal{M},s) = \mathcal{B}^{-+}_{\varphi\varphi}(i,\mathcal{M},s)$, we conclude that postulate B^-5 is satisfied.

(B^-6) Suppose $\mathcal{M},s \models \mathbf{K}_i(\varphi \leftrightarrow \psi)$. Then from $\Sigma 4$ it follows that $\sigma(i,s,\varphi) = \sigma(i,s,\psi)$. Hence $\mathtt{r}(i,\mathtt{contract}\ \psi)(\mathcal{M},s) = \mathcal{M}',s$, and $\mathcal{B}^-_\varphi(i,\mathcal{M},s) = \mathcal{B}^-_\psi(i,\mathcal{M},s)$.

(B^-7) Let $\rho \in \mathcal{B}^-_\varphi(i,\mathcal{M},s) \cap \mathcal{B}^-_\psi(i,\mathcal{M},s)$. This implies that $\mathcal{M},s' \models \rho$ for all $s' \in \mathrm{B}(i,s) \cup \sigma(i,s,\varphi) \cup \sigma(i,s,\psi)$. Since by $\Sigma 5$, $\sigma(i,s,\varphi \wedge \psi) \subseteq \sigma(i,s,\varphi) \cup \sigma(i,s,\psi)$, it follows that $\mathcal{M},s' \models \rho$ for all $s' \in \mathrm{B}(i,s) \cup \sigma(i,s,\varphi \wedge \psi)$. But then $\rho \in \mathcal{B}^-_{\varphi\wedge\psi}(i,\mathcal{M},s)$. Since ρ was chosen arbitrarily it follows that $\mathcal{B}^-_\varphi(i,\mathcal{M},s) \cap \mathcal{B}^-_\psi(i,\mathcal{M},s) \subseteq \mathcal{B}^-_{\varphi\wedge\psi}(i,\mathcal{M},s)$.

(B^-8) Suppose $\varphi \notin \mathcal{B}^-_{\varphi\wedge\psi}(i,\mathcal{M},s)$. Let $\mathcal{M}'',s = \mathtt{r}(i,\mathtt{contract}\ (\varphi \wedge \psi))(\mathcal{M},s)$. We distinguish two cases:

- $\varphi \notin \mathcal{B}(i,\mathcal{M},s)$. Then also $\varphi\wedge\psi \notin \mathcal{B}(i,\mathcal{M},s)$, since $\mathcal{B}(i,\mathcal{M},s)$ is a belief set. From B^-3 it then follows that $\mathcal{B}^-_\varphi(i,\mathcal{M},s)$ and $\mathcal{B}^-_{\varphi\wedge\psi}(i,\mathcal{M},s)$ are both equal to $\mathcal{B}(i,\mathcal{M},s)$.
- $\varphi \in \mathcal{B}(i,\mathcal{M},s)$. Since $\varphi \notin \mathcal{B}^-_{\varphi\wedge\psi}(i,\mathcal{M},s)$, it follows that some $s' \in \sigma(i,s,\varphi\wedge\psi)$ exists such that $\mathcal{M}'',s' \models \neg\varphi$, and since φ is propositional, also $\mathcal{M},s' \models \neg\varphi$. But then $\sigma(i,s,\varphi \wedge \psi) \cap [\![\neg\varphi]\!] \neq \emptyset$, and by $\Sigma 6$ it follows that $\sigma(i,s,\varphi) \subseteq \sigma(i,s,\varphi \wedge \psi)$. Next, for all formulae $\rho \in \mathcal{B}^-_{\varphi\wedge\psi}(i,\mathcal{M},s)$, $\mathcal{M},s' \models \rho$ for all $s' \in \mathrm{B}(i,s) \cup \sigma(i,s,\varphi \wedge \psi)$. Since $\sigma(i,s,\varphi) \subseteq \sigma(i,s,\varphi \wedge \psi)$ it follows that $\mathcal{M},s' \models \rho$ for all $s' \in \mathrm{B}(i,s) \cup \sigma(i,s,\varphi)$, and hence $\mathcal{B}^-_{\varphi\wedge\psi}(i,\mathcal{M},s) \subseteq \mathcal{B}^-_\varphi(i,\mathcal{M},s)$.

Since in both cases $\mathcal{B}^-_{\varphi\wedge\psi}(i,\mathcal{M},s) \subseteq \mathcal{B}^-_\varphi(i,\mathcal{M},s)$, we conclude that postulate B^-8 is satisfied.

⊠

3.21. PROPOSITION. *Let $\mathcal{M}$ be some Kripke model with state s, let i be some agent and φ some propositional formula. Let* $\mathtt{r}(i,\mathtt{contract}\ \varphi)(\mathcal{M},s)$ *be defined*

- $\models [\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{B}_i\psi \to \mathbf{B}_i\psi$
- $\models \neg\mathbf{B}_i\varphi \to ([\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{B}_i\psi \leftrightarrow \mathbf{B}_i\psi)$
- $\models \neg\mathbf{K}_i\varphi \to [\mathrm{do}_i(\texttt{contract}\ \varphi)]\neg\mathbf{B}_i\varphi$
- $\models \mathbf{B}_i\varphi \to (\mathbf{B}_i\psi \to [\mathrm{do}_i(\texttt{contract}\ \varphi; \texttt{expand}\ \varphi)]\mathbf{B}_i\psi)$
- $\models \mathbf{K}_i(\varphi \leftrightarrow \psi) \to ([\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{B}_i\vartheta \leftrightarrow [\mathrm{do}_i(\texttt{contract}\ \psi)]\mathbf{B}_i\vartheta)$
- $\models ([\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{B}_i\vartheta \wedge [\mathrm{do}_i(\texttt{contract}\ \psi)]\mathbf{B}_i\vartheta) \to$
  $[\mathrm{do}_i(\texttt{contract}\ (\varphi \wedge \psi))]\mathbf{B}_i\vartheta$
- $\models [\mathrm{do}_i(\texttt{contract}\ (\varphi \wedge \psi))]\neg\mathbf{B}_i\varphi \to$
  $([\mathrm{do}_i(\texttt{contract}\ (\varphi \wedge \psi))]\mathbf{B}_i\vartheta \to [\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{B}_i\vartheta)$

Proof: All clauses follow directly from Theorem 3.20. The first clause of Proposition 3.16 follows from $B^{-}2$ and so on until the last clause that follows from $B^{-}8$.
⊠

3.17. Proposition. *For all agents i, for all propositional formulae φ, and for all formulae ψ we have:*

- $\models \langle \mathrm{do}_i(\texttt{contract}\ \varphi)\rangle\mathbf{tt}$
- $\models \langle \mathrm{do}_i(\texttt{contract}\ \varphi)\rangle\psi \leftrightarrow [\mathrm{do}_i(\texttt{contract}\ \varphi)]\psi$
- $\models \langle \mathrm{do}_i(\texttt{contract}\ \varphi; \texttt{contract}\ \varphi)\rangle\psi \leftrightarrow \langle \mathrm{do}_i(\texttt{contract}\ \varphi)\rangle\psi$

Proof: Directly from Def. 3.15.
⊠

3.20. Theorem. *Let $\mathcal{M}$ be some Kripke model. For all agents i, for all $s \in \mathcal{M}$ and for all formulae φ and ψ the following are true.*
($B^{-}1$) *$\mathcal{B}^{-}_{\varphi}(i, \mathcal{M}, s)$ is a belief set.*
($B^{-}2$) *$\mathcal{B}^{-}_{\varphi}(i, \mathcal{M}, s) \subseteq \mathcal{B}(i, \mathcal{M}, s)$.*
($B^{-}3$) *If $\varphi \notin \mathcal{B}(i, \mathcal{M}, s)$ then $\mathcal{B}^{-}_{\varphi}(i, \mathcal{M}, s) = \mathcal{B}(i, \mathcal{M}, s)$.*
($B^{-}4$) *If $\mathcal{M}, s \not\models \mathbf{K}_i\varphi$ then $\varphi \notin \mathcal{B}^{-}_{\varphi}(i, \mathcal{M}, s)$.*
($B^{-}5$) *If $\varphi \in \mathcal{B}(i, \mathcal{M}, s)$ then $\mathcal{B}(i, \mathcal{M}, s) \subseteq \mathcal{B}^{-+}_{\varphi\varphi}(i, \mathcal{M}, s)$.*
($B^{-}6$) *If $\mathcal{M}, s \models \mathbf{K}_i(\varphi \leftrightarrow \psi)$ then $\mathcal{B}^{-}_{\varphi}(i, \mathcal{M}, s) = \mathcal{B}^{-}_{\psi}(i, \mathcal{M}, s)$.*
($B^{-}7$) *$\mathcal{B}^{-}_{\varphi}(i, \mathcal{M}, s) \cap \mathcal{B}^{-}_{\psi}(i, \mathcal{M}, s) \subseteq \mathcal{B}^{-}_{\varphi\wedge\psi}(i, \mathcal{M}, s)$.*
($B^{-}8$) *If $\varphi \notin \mathcal{B}^{-}_{\varphi\wedge\psi}(i, \mathcal{M}, s)$ then $\mathcal{B}^{-}_{\varphi\wedge\psi}(i, \mathcal{M}, s) \subseteq \mathcal{B}^{-}_{\varphi}(i, \mathcal{M}, s)$.*

Proof: Let $\mathcal{M}$ be some Kripke model with state s, and let i be some agent. Let σ be an arbitrary selection function for $\mathcal{M}$. Let φ be some arbitrary propositional formula, and let $\mathcal{M}', s = \mathtt{r}(i, \texttt{contract}\ \varphi)(\mathcal{M}, s)$. We show that contractions based on σ satisfy the Gärdenfors postulates.

($B^{-}1$) This postulate is easily seen to be satisfied: in the case where $\mathrm{B}'(i, s) = \emptyset$, $\mathcal{B}(i, \mathcal{M}', s) = \mathcal{B}_{\perp}$, and otherwise $\mathcal{B}(i, \mathcal{M}', s)$ is consistent and deductively closed by definition of $\models$ for $\mathbf{B}_i\varphi$.

($B^{-}2$) By demand $\Sigma 1$ it follows that $\sigma(i, s, \varphi)$ yields a set of states from $\mathcal{M}$. It is easily seen that for propositional formulae it holds that if $\mathcal{M}, s' \models \varphi$ for all $s' \in \mathcal{S}'$ then for all $\mathcal{S}'' \subseteq \mathcal{S}'$, $\mathcal{M}, s'' \models \varphi$ for all $s'' \in \mathcal{S}''$. Now if $\psi \in \mathcal{B}^{-}_{\varphi}(i, \mathcal{M}, s)$, then $\mathcal{M}, s' \models \psi$ for all $s' \in \mathrm{B}(i, s) \cup \sigma(i, s, \varphi)$. Hence $\mathcal{M}, s' \models \psi$ for all $s' \in \mathrm{B}(i, s)$, and thus $\psi \in \mathcal{B}(i, \mathcal{M}, s)$.

($B^{+}1$) This postulate follows straightforwardly from the definition of $\texttt{r}$ for the $\texttt{expand}$ action. If the resulting model $\mathcal{M}'$ is such that $\mathrm{B}'(i,s) = \emptyset$ then $\mathcal{B}^{+}_{\varphi}(i,\mathcal{M},s) = \mathcal{B}_{\perp}$, and otherwise $\mathcal{B}^{+}_{\varphi}(i,\mathcal{M},s)$ is consistent and deductively closed by definition of $\models$ for belief formulae.

($B^{+}2$) Since φ is a propositional formula, it follows by definition of $\texttt{r}$ for $\texttt{expand}$ that $\mathcal{M}', s' \models \varphi$ for all $s' \in \mathrm{B}'(i,s)$. Hence $\mathcal{M}', s \models \mathbf{B}_i\varphi$ and thus $\varphi \in \mathcal{B}^{+}_{\varphi}(i,\mathcal{M},s)$.

($B^{+}3$) Let $\psi \in \mathcal{B}(i,\mathcal{M},s)$. Then $\mathcal{M}, s' \models \psi$ for all $\psi \in \mathrm{B}(i,s)$. Now $\mathrm{B}'(i,s) \subseteq \mathrm{B}(i,s)$ and since ψ is propositional it follows that $\mathcal{M}', s' \models \psi$ for all $s' \in \mathrm{B}'(i,s)$. Hence $\mathcal{M}', s \models \mathbf{B}_i\psi$ and thus $\psi \in \mathcal{B}^{+}_{\varphi}(i,\mathcal{M},s)$.

($B^{+}4$) Suppose $\varphi \in \mathcal{B}(i,\mathcal{M},s)$. Then $\mathcal{M}, s' \models \varphi$ for all $s' \in \mathrm{B}(i,s)$. Since φ is propositional it follows that $\mathrm{B}(i,s) \cap [\![\varphi]\!] = \mathrm{B}(i,s)$ and hence $\mathrm{B}'(i,s) = \mathrm{B}(i,s)$. Then $\mathcal{M}', s = \mathcal{M}, s$, and hence $\mathcal{B}^{+}_{\varphi}(i,\mathcal{M},s) = \mathcal{B}(i,\mathcal{M}',s) = \mathcal{B}(i,\mathcal{M},s)$.

($B^{+}5$) The proof that this postulate is validated is most easily given as a direct consequence of Proposition 3.13. For if $\mathcal{B}(i,\mathcal{M},s) \subseteq \mathcal{B}(i,\mathcal{M}',s')$ then $\mathrm{Th}(\mathcal{B}(i,\mathcal{M},s) \cap \{\varphi\}) \subseteq \mathrm{Th}(\mathcal{B}(i,\mathcal{M}',s') \cap \{\varphi\})$ and thus $\mathcal{B}^{+}_{\varphi}(i,\mathcal{M},s) \subseteq \mathcal{B}^{+}_{\varphi}(i,\mathcal{M}',s')$. Since the proof of Proposition 3.13 does not depend on $B^{+}5$, this postulate is validated.

($B^{+}6$) From $B^{+}2$, $B^{+}3$ and the fact that belief sets are deductively closed, it follows that $\mathrm{Th}(\mathcal{B}(i,\mathcal{M},s) \cup \{\varphi\}) \subseteq \mathcal{B}^{+}_{\varphi}(i,\mathcal{M},s)$. From Proposition 3.13, the proof of which does not depend on $B^{+}6$, it follows that $\mathcal{B}^{+}_{\varphi}(i,\mathcal{M},s)$ is indeed the smallest set that satisfies $B^{+}1$ through $B^{+}5$.

⊠

3.13. Proposition. *For all Kripke models $\mathcal{M}$ with states s, for all agents i, and for all propositional formulae φ we have:*

$$\mathcal{B}^{+}_{\varphi}(i,\mathcal{M},s) = \mathrm{Th}(\mathcal{B}(i,\mathcal{M},s) \cup \{\varphi\})$$

Proof: We prove that the two sets are equal by proving that each set is a subset of the other one. So let $\mathcal{M}$ be some Kripke model with state s, let i be some agent, φ be some propositional formula.

'$\supseteq$' This is shown by the argument given in the proof of $B^{+}6$: from $B^{+}2$, $B^{+}3$ and the fact that beliefs sets are deductively closed, it follows that $\mathrm{Th}(\mathcal{B}(i,\mathcal{M},s) \cup \{\varphi\}) \subseteq \mathcal{B}^{+}_{\varphi}(i,\mathcal{M},s)$.

'$\subseteq$' Suppose that the propositional formula $\psi \in \mathcal{B}^{+}_{\varphi}(i,\mathcal{M},s)$. If $\mathcal{M}', s = \texttt{r}(i, \texttt{expand}\ \varphi)(\mathcal{M},s)$ this implies that $\mathcal{M}', s' \models \psi$ for all $s' \in \mathrm{B}'(i,s)$. Since ψ is a propositional formula, and since $\mathrm{B}'(i,s) = \mathrm{B}(i,s) \cap [\![\varphi]\!]$, it follows that $\mathcal{M}, s' \models \psi$ for all $s' \in \mathrm{B}(i,s)$ such that $\mathcal{M}, s' \models \varphi$. Hence $\mathcal{M}, s' \models (\varphi \to \psi)$ for all $s' \in \mathrm{B}(i,s)$. Then $\varphi \to \psi \in \mathcal{B}(i,\mathcal{M},s)$ and thus $\psi \in \mathrm{Th}(\mathcal{B}(i,\mathcal{M},s) \cup \{\varphi\})$. Since ψ is arbitrary it follows that $\mathcal{B}^{+}_{\varphi}(i,\mathcal{M},s) \subseteq \mathrm{Th}(\mathcal{B}(i,\mathcal{M},s) \cup \{\varphi\})$.

⊠

3.16. Proposition. *For all agents i, and propositional formulae φ, ψ and ϑ we have:*

3.7. PROPOSITION. *For all agents i, for all propositional formulae φ, and for all formulae ψ we have:*

- $\models \langle \mathrm{do}_i(\texttt{expand}\ \varphi)\rangle \mathbf{tt}$
- $\models \langle \mathrm{do}_i(\texttt{expand}\ \varphi)\rangle \psi \leftrightarrow [\mathrm{do}_i(\texttt{expand}\ \varphi)]\psi$
- $\models \langle \mathrm{do}_i(\texttt{expand}\ \varphi; \texttt{expand}\ \varphi)\rangle \psi \leftrightarrow \langle \mathrm{do}_i(\texttt{expand}\ \varphi)\rangle \psi$

PROOF: Directly from Def. 3.5.
⊠

3.8. PROPOSITION. *For all agents i, for all propositional formulae φ and ψ we have:*

- $\models [\mathrm{do}_i(\texttt{expand}\ \varphi)]\mathbf{B}_i\psi \leftrightarrow \mathbf{B}_i(\varphi \to \psi)$

PROOF: Let $\mathcal{M}$ be some Kripke model with state s, let i be some agent and let φ and ψ be arbitrary propositional formulae. Then:

$$
\begin{array}{lll}
 & \mathcal{M}, s \models [\mathrm{do}_i(\texttt{expand}\ \varphi)]\mathbf{B}_i\psi & \\
\Leftrightarrow & \psi \in \mathcal{B}^+_\varphi(i, \mathcal{M}, s) & (\text{Definition of } \mathcal{B}^+_\varphi(i, \mathcal{M}, s)) \\
\Leftrightarrow & \psi \in \mathrm{Th}(\mathcal{B}(i, \mathcal{M}, s) \cup \{\varphi\}) & (\text{Proposition 3.13}) \\
\Leftrightarrow^* & (\varphi \to \psi) \in \mathrm{Th}(\mathcal{B}(i, \mathcal{M}, s)) & \\
\Leftrightarrow & \mathcal{M}, s \models \mathbf{B}_i(\varphi \to \psi) & (\text{Definition of } \mathcal{B}(i, \mathcal{M}, s))
\end{array}
$$

The starred equivalence holds due to the Deduction Theorem for propositional classical logic, which states that for sets of formulae Φ and formulae φ and ψ, $\psi \in \mathrm{Th}(\Phi \cup \{\varphi\})$ holds iff $(\varphi \to \psi) \in \mathrm{Th}(\Phi)$ holds.
⊠

3.9. COROLLARY. *For all agents i, for all propositional formulae φ and ψ we have:*

- $\models [\mathrm{do}_i(\texttt{expand}\ \varphi)]\mathbf{B}_i\mathbf{ff} \leftrightarrow \mathbf{B}_i\neg\varphi$

PROOF: The corollary follows directly from Proposition 3.8 and the observation that $\neg\varphi$ and $\varphi \to \mathbf{ff}$ are equivalent formulae.
⊠

3.12. THEOREM. *Let $\mathcal{M}$ and $\mathcal{M}'$ be Kripke models, with s in $\mathcal{M}$ and s' in $\mathcal{M}'$, and let i be some agent. The following is valid for all $\varphi \in \mathcal{L}_0$.*

(B^+1) *$\mathcal{B}^+_\varphi(i, \mathcal{M}, s)$ is a belief set, i.e., $\mathcal{B}^+_\varphi(i, \mathcal{M}, s)$ is either equal to $\mathcal{B}_\perp$ or it is consistent and deductively closed.*
(B^+2) *$\varphi \in \mathcal{B}^+_\varphi(i, \mathcal{M}, s)$.*
(B^+3) *$\mathcal{B}(i, \mathcal{M}, s) \subseteq \mathcal{B}^+_\varphi(i, \mathcal{M}, s)$.*
(B^+4) *If $\varphi \in \mathcal{B}(i, \mathcal{M}, s)$ then $\mathcal{B}^+_\varphi(i, \mathcal{M}, s) = \mathcal{B}(i, \mathcal{M}, s)$.*
(B^+5) *If $\mathcal{B}(i, \mathcal{M}, s) \subseteq \mathcal{B}(i, \mathcal{M}', s')$, then $\mathcal{B}^+_\varphi(i, \mathcal{M}, s) \subseteq \mathcal{B}^+_\varphi(i, \mathcal{M}', s')$.*
(B^+6) *$\mathcal{B}^+_\varphi(i, \mathcal{M}, s)$ is the smallest set that satisfies the postulates 1 – 5 as given above.*

PROOF: Let $\mathcal{M}$ be some Kripke model with state s, let i be some agent, φ be some propositional formula, and let

- To prove: $\models \neg\mathbf{B}_i\varphi \to \mathbf{B}_i\neg\mathbf{B}_i\varphi$. Let $\mathcal{M}$ be some Kripke model with state s, let i be some agent and let φ be some formula.
  $\mathcal{M}, s \models \neg\mathbf{B}_i\varphi$
  $\Leftrightarrow \exists s' \in \mathrm{B}(i,s)[\mathcal{M}, s' \not\models \varphi]$
  $\Rightarrow \forall s'' \in s\exists s' \in \mathrm{B}(i,s)[\mathcal{M}, s' \not\models \varphi]$
  $\Rightarrow \forall s'' \in \mathrm{B}(i,s)\exists s' \in \mathrm{B}(i,s)[\mathcal{M}, s' \not\models \varphi]$
  $\Leftrightarrow \mathcal{M}, s'' \models \neg\mathbf{B}_i\varphi$ for all $s'' \in \mathrm{B}(i,s)$
  $\Leftrightarrow \mathcal{M}, s \models \mathbf{B}_i\neg\mathbf{B}_i\varphi$
- To prove: $\models \neg\mathbf{B}_i\mathbf{ff} \to (\mathbf{B}_i\neg\mathbf{B}_i\varphi \to \neg\mathbf{B}_i\varphi)$. Let $\mathcal{M}$ be some Kripke model with state s, let i be some agent and let φ be some formula. Suppose $\mathcal{M}, s \models \neg\mathbf{B}_i\mathbf{ff}$, then $\mathrm{B}(i,s) \neq \emptyset$. Then:
  $\mathcal{M}, s \models \mathbf{B}_i\neg\mathbf{B}_i\varphi$
  $\Leftrightarrow \mathcal{M}, s' \models \neg\mathbf{B}_i\varphi$ for all $s' \in \mathrm{B}(i,s)$
  $\Leftrightarrow \forall s' \in \mathrm{B}(i,s)\exists s'' \in \mathrm{B}(i,s)[\mathcal{M}, s'' \not\models \varphi]$
  $\Leftrightarrow \exists s'' \in \mathrm{B}(i,s)[\mathcal{M}, s'' \models \neg\varphi]$ ('$\Rightarrow$': since $\mathrm{B}(i,s) \neq \emptyset$)
  $\Leftrightarrow \mathcal{M}, s \models \neg\mathbf{B}_i\varphi$

⊠

3.2. Proposition. *Let $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathtt{c}\rangle$ be some Kripke model with $s \in \mathcal{S}$, and let $\mathcal{M}' = \langle \mathcal{S}', \pi', \mathrm{R}', \mathrm{B}', \mathtt{r}', \mathtt{c}'\rangle$ be some Kripke model with $s' \in \mathcal{S}'$. Then it holds that:*

$$\forall p \in \Pi[\pi(p,s) = \pi'(p,s')] \Rightarrow \forall \psi \in \mathcal{L}_0[\mathcal{M}, s \models \psi \Leftrightarrow \mathcal{M}', s' \models \psi]$$

Proof: The proposition is shown by induction on the structure of ψ.

1. If $\psi \in \Pi$ then the proposition is trivial.
2. If $\psi = \neg\psi_1$ where ψ_1 is purely propositional then:
   $\mathcal{M}, s \models \neg\psi_1$
   $\Leftrightarrow \mathrm{not}(\mathcal{M}, s \models \psi_1)$
   $\Leftrightarrow \mathrm{not}(\mathcal{M}', s' \models \psi_1)$ (Induction Hypothesis)
   $\Leftrightarrow \mathcal{M}', s' \models \neg\psi_1$
3. If $\psi = \psi_1 \vee \psi_2$ where ψ_1, ψ_2 purely propositional then:
   $\mathcal{M}, s \models \psi_1 \vee \psi_2$
   $\Leftrightarrow \mathcal{M}, s \models \psi_1$ or $\mathcal{M}, s \models \psi_2$
   $\Leftrightarrow \mathcal{M}', s' \models \psi_1$ or $\mathcal{M}', s' \models \psi_2$ (Induction Hypothesis)
   $\Leftrightarrow \mathcal{M}', s' \models \psi_1 \vee \psi_2$

⊠

3.6. Proposition. *For all agents i, and for all propositional formulae φ and ψ we have:*

- $\models [\mathrm{do}_i(\mathtt{expand}\ \varphi)]\mathbf{B}_i\varphi$
- $\models \mathbf{B}_i\psi \to [\mathrm{do}_i(\mathtt{expand}\ \varphi)]\mathbf{B}_i\psi$
- $\models \mathbf{B}_i\varphi \to (\mathbf{B}_i\psi \leftrightarrow [\mathrm{do}_i(\mathtt{expand}\ \varphi)]\mathbf{B}_i\psi)$

Proof: All clauses follow directly from Theorem 3.12. The first clause of Proposition 3.6 follows from B^+2, the second clause follows from B^+3, and the third clause follows from B^+4.

⊠

the Katholieke Universiteit Nijmegen. Thanks are also due to the members of the Pst! colloquium, viz. Thomas Arts and Theo Huibers, for their very helpful comments and criticism on a draft version of this paper.

A. Appendix: the proofs

A.1. Remark. For reasons imposed by the demand for informational economy, in some of the proofs given below references are made to proofs that are given later on. In this cases the latter proof is either more elaborate or more instructive than the former one.

2.6. Proposition. *For all agents i and formulae φ and ψ we have:*

- $\models \mathbf{K}_i(\varphi \to \psi) \to (\mathbf{K}_i\varphi \to \mathbf{K}_i\psi)$
- $\models \mathbf{K}_i\varphi \to \varphi$
- $\models \mathbf{K}_i\varphi \to \mathbf{K}_i\mathbf{K}_i\varphi$
- $\models \neg\mathbf{K}_i\varphi \to \mathbf{K}_i\neg\mathbf{K}_i\varphi$
- $\models \varphi \Rightarrow \models \mathbf{K}_i\varphi$

Proof: The proof of this proposition is fairly standard, and can be found at various places in the literature [1, 16, 25].
⊠

2.7. Proposition. *For all agents i and formulae φ and ψ we have:*

- $\models \mathbf{B}_i(\varphi \to \psi) \to (\mathbf{B}_i\varphi \to \mathbf{B}_i\psi)$
- $\models \mathbf{B}_i\varphi \to \mathbf{B}_i\mathbf{B}_i\varphi$
- $\models \neg\mathbf{B}_i\varphi \to \mathbf{B}_i\neg\mathbf{B}_i\varphi$
- $\models \varphi \Rightarrow \models \mathbf{B}_i\varphi$

Proof: Again the proof of this proposition is standard [1, 16, 25].
⊠

2.8. Proposition. *For all agents i and formulae φ we have:*

- $\models \mathbf{K}_i\varphi \to \mathbf{B}_i\varphi$
- $\models \neg\mathbf{B}_i\mathbf{ff} \to (\mathbf{K}_i\neg\varphi \to \neg\mathbf{B}_i\varphi)$
- $\models \mathbf{K}_i\mathbf{K}_i\varphi \leftrightarrow \mathbf{K}_i\varphi$
- $\models \mathbf{K}_i\mathbf{B}_i\varphi \leftrightarrow \mathbf{B}_i\varphi$
- $\models \mathbf{K}_i\neg\mathbf{K}_i\varphi \leftrightarrow \neg\mathbf{K}_i\varphi$
- $\models \mathbf{K}_i\neg\mathbf{B}_i\varphi \leftrightarrow \neg\mathbf{B}_i\varphi$
- $\models \mathbf{K}_i\varphi \to \mathbf{B}_i\mathbf{K}_i\varphi$
- $\models \neg\mathbf{B}_i\mathbf{ff} \to (\mathbf{B}_i\mathbf{K}_i\varphi \to \mathbf{K}_i\varphi)$
- $\models \mathbf{B}_i\mathbf{B}_i\varphi \leftrightarrow \mathbf{B}_i\varphi$
- $\models \neg\mathbf{K}_i\varphi \to \mathbf{B}_i\neg\mathbf{K}_i\varphi$
- $\models \neg\mathbf{B}_i\mathbf{ff} \to (\mathbf{B}_i\neg\mathbf{K}_i\varphi \to \neg\mathbf{K}_i\varphi)$
- $\models \neg\mathbf{B}_i\varphi \to \mathbf{B}_i\neg\mathbf{B}_i\varphi$
- $\models \neg\mathbf{B}_i\mathbf{ff} \to (\mathbf{B}_i\neg\mathbf{B}_i\varphi \to \neg\mathbf{B}_i\varphi)$

Proof: All cases are easily shown. As an example we show the last two cases.

- $\mathbf{Cannot}_i(\alpha, \varphi) \equiv \mathbf{K}_i(\neg\langle \mathrm{do}_i(\alpha)\rangle\varphi \vee \neg\mathbf{A}_i\alpha)$.

Due to the specific properties of the belief-changing actions with respect to results, opportunities, and abilities, some remarkable validities can be shown.

3.43. PROPOSITION. *For all agents i, and for all propositional formulae φ we have:*

- $\models \mathbf{A}_i\texttt{expand}\ \varphi \rightarrow \mathbf{Can}_i(\texttt{expand}\ \varphi, \mathbf{B}_i\varphi \wedge \neg\mathbf{B}_i\mathbf{ff})$
- $\models \mathbf{A}_i\texttt{contract}\ \varphi \rightarrow \mathbf{Can}_i(\texttt{contract}\ \varphi, \neg\mathbf{B}_i\varphi)$
- $\models \mathbf{A}_i\texttt{contract}\ \neg\varphi \rightarrow \mathbf{Can}_i(\texttt{revise}\ \varphi, \mathbf{B}_i\varphi \wedge \neg\mathbf{B}_i\mathbf{ff})$

Intuitively, Proposition 3.43 states that any agent that has the ability to perform some belief-changing action *knows* that this action is feasible and behaves the way it should, i.e., the agent knows that this action is a correct and feasible plan to change its beliefs in the desired way.

4. Discussion

In this paper we defined actions that model three well-known changes of belief, viz. *expansions*, *contractions*, and *revisions*. We characterized the states of affairs that result from execution of these actions, the conditions that decide whether agents have the opportunity to perform these actions, and the capacities that agents should posses in order to be able to perform these actions. The action that models belief contractions is defined using the notion of selection functions. These are functions that select a subset of the set of epistemic alternatives of an agent that is to be added to its set of doxastic alternatives, in order to contract its set of beliefs. We prove that our kind of selection functions resemble both the selection functions proposed by Stalnaker [33] and partial meet contraction functions as defined by Gärdenfors [5]. The action that models belief revision is defined in terms of a contraction and an expansion in a way suggested by the Levi identity [21]. We showed that our belief-changing actions satisfy (semantic, agent-oriented variants of) the Gärdenfors postulates for belief expansions, belief contractions, and belief revisions, thereby supporting our claim that the formalization that we present is both an intuitively and philosophically acceptable one.

With regard to the topics covered in this paper, our future research will concentrate on agents acquiring information from multiple sources. For in this case, the reliability of the information and the source determines whether and how the beliefs of the agents should change. Actions that model belief revisions will obviously play an important part in situations like these. A general and important topic of future research will be the treatment of the motivational aspect of agents.

Acknowledgements

This research is partially supported by ESPRIT BRWG project No.8319 'MODELAGE', ESPRIT III BRA project No.6156 'DRUMS II', and the Vrije Universiteit Amsterdam; the third author is furthermore partially supported by

$$\begin{array}{lcl} \texttt{c}(i, \texttt{expand}\ \varphi)(\mathcal{M}, s) = \mathbf{1} & \Leftrightarrow & \mathcal{M}, s \models \neg \mathbf{B}_i \neg \varphi \\ \texttt{c}(i, \texttt{contract}\ \varphi)(\mathcal{M}, s) = \mathbf{1} & \Leftrightarrow & \mathcal{M}, s \models \neg \mathbf{K}_i \varphi \\ \texttt{c}(i, \texttt{revise}\ \varphi)(\mathcal{M}, s) & = & \texttt{c}(i, \texttt{contract}\ \neg\varphi; \texttt{expand}\ \varphi)(\mathcal{M}, s) \end{array}$$

The first clause of Def. 3.40 states that an agent is able to expand its set of beliefs with a formula if and only if it does not already believe the negation of the formula. The second clause formalizes the idea that an agent is able to remove some formula from its set of beliefs if and only if it does not consider the formula to be one of its principles. The ability for the `revise` action is defined through the Levi identity.

3.41. PROPOSITION. *For all agents i and for all propositional formulae φ we have:*

- $\models \mathbf{A}_i\texttt{expand}\ \varphi \leftrightarrow \mathbf{K}_i\mathbf{A}_i\texttt{expand}\ \varphi$
- $\models \mathbf{A}_i\texttt{expand}\ \varphi \rightarrow \langle \mathrm{do}_i(\texttt{expand}\ \varphi)\rangle \neg \mathbf{B}_i\mathbf{ff}$
- $\models \mathbf{A}_i\texttt{contract}\ \varphi \leftrightarrow \mathbf{K}_i\mathbf{A}_i\texttt{contract}\ \varphi$
- $\models \mathbf{A}_i\texttt{contract}\ \varphi \rightarrow \langle \mathrm{do}_i(\texttt{contract}\ \varphi)\rangle \neg \mathbf{B}_i\varphi$
- $\models \mathbf{A}_i\texttt{revise}\ \varphi \leftrightarrow \mathbf{A}_i\texttt{contract}\ \neg\varphi$
- $\models \mathbf{A}_i\texttt{revise}\ \varphi \rightarrow \langle \mathrm{do}_i(\texttt{revise}\ \varphi)\rangle (\mathbf{B}_i\varphi \wedge \neg \mathbf{B}_i\mathbf{ff})$

The first and third clause of Proposition 3.41 state that agents know of their ability to expand and contract their beliefs; a consequence of the fifth clause is that agents also know of their ability to revise their beliefs. The second, fourth and sixth clause formalize the idea that belief changes of which the agent is capable, behave the way they should, i.e., an expansion does not result in absurd belief sets, a contraction leads to disbelief in the contracted formula, and a revision results in the combination of these two.

3.4.1. The Can-predicate and the Cannot-predicate

Inspired by the concepts introduced by Moore [27], we defined our own version of the Can-predicate and the Cannot-predicate [11]. Intuitively these predicates formalize the knowledge and the reasoning of agents regarding the (in)correctness and (in)feasibility of their plans to achieve certain goals, and as such these predicates deal with important aspects of the agents' planning.

The definition of these predicates is based on the idea that an agent i knows that action α is a *correct* plan to achieve φ if and only if it knows that $\langle \mathrm{do}_i(\alpha)\rangle\varphi$ holds. Agent i knows that α is a *feasible* plan for φ if and only if it knows that it is able to do α, i.e., $\mathbf{K}_i\mathbf{A}_i\alpha$ holds. This intuition is formalized in the definition of the Can-predicate as we give it. We furthermore define the Cannot-predicate, which has as its intended meaning that the agent knows that it cannot reach some goal φ by performing some action α, since it knows that either the action does not lead to the desired goal or it is not capable of performing the action, i.e., the agent knows that the action is either an incorrect or an infeasible plan.

3.42. DEFINITION. For all agents i, action α and formula φ, the Can-predicate and the Cannot-predicate are defined as follows.

- $\mathbf{Can}_i(\alpha, \varphi) \equiv \mathbf{K}_i(\langle \mathrm{do}_i(\alpha)\rangle\varphi \wedge \mathbf{A}_i\alpha)$.

3.38. THEOREM. *Let $\mathcal{M}$ be some Kripke model. For all agents i, for all $s \in \mathcal{M}$ and for all formulae φ and ψ the following are true.*
($B^\star 1$) $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s)$ *is a belief set.*
($B^\star 2$) $\varphi \in \mathcal{B}^\star_\varphi(i, \mathcal{M}, s)$.
($B^\star 3$) $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) \subseteq \mathcal{B}^+_\varphi(i, \mathcal{M}, s)$.
($B^\star 4$) *If* $\neg\varphi \notin \mathcal{B}(i, \mathcal{M}, s)$ *then* $\mathcal{B}^+_\varphi(i, \mathcal{M}, s) \subseteq \mathcal{B}^\star_\varphi(i, \mathcal{M}, s)$.
($B^\star 5$) $\mathcal{B}(i, \mathcal{M}, s) = \mathcal{B}_\perp$ *if and only if* $\mathcal{M}, s \models \mathbf{K}_i \neg\varphi$.
($B^\star 6$) *If* $\mathcal{M}, s \models \mathbf{K}_i(\varphi \leftrightarrow \psi)$ *then* $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) = \mathcal{B}^\star_\psi(i, \mathcal{M}, s)$.
($B^\star 7$) $\mathcal{B}^\star_{\varphi \wedge \psi}(i, \mathcal{M}, s) \subseteq \mathcal{B}^{\star +}_{\varphi\psi}(i, \mathcal{M}, s)$.
($B^\star 8$) *If* $\neg\psi \notin \mathcal{B}^\star_\varphi(i, \mathcal{M}, s)$, *then* $\mathcal{B}^{\star +}_{\varphi\psi}(i, \mathcal{M}, s) \subseteq \mathcal{B}^\star_{\varphi \wedge \psi}(i, \mathcal{M}, s)$.

Also the result of Proposition 3.35 can be rephrased in terms that make it more in line with the framework of Gärdenfors.

3.39. PROPOSITION. *For all Kripke models $\mathcal{M}$ with state s, for all agents i, and for all propositional formulae φ we have:*

- *if* $\neg\varphi \in \mathcal{K}(i, \mathcal{M}, s)$ *then* $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) = \mathcal{B}_\perp$.
- *if* $\neg\varphi \notin \mathcal{B}(i, \mathcal{M}, s)$ *then* $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) = \mathrm{Th}(\mathcal{B}(i, \mathcal{M}, s) \cup \{\varphi\})$.
- *if* $\neg\varphi \in \mathcal{B}(i, \mathcal{M}, s) \setminus \mathcal{K}(i, \mathcal{M}, s)$ *then* $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) = \mathrm{Th}(\mathcal{K}(i, \mathcal{M}, s) \cup \{\varphi\})$ *if the definition of* `r` *for the* `contract` *action is based on the AiG function for $\mathcal{M}$.*

3.4. The ability to change one's mind

In the previous (sub)sections, we dealt with the formalization of the opportunity for and the result of the actions that model the belief changes of agents. Here we look at the *ability* of agents to change their beliefs.

For 'mental' actions, like testing (observing) and communicating, the abilities of agents are closely related to their (lack of) knowledge and/or belief. This observation seems to hold *a fortiori* for the abstract actions that cause agents to change their beliefs. For when testing and communicating, at least some interaction takes place, either with the real world in case of testing, or with other agents when communicating, whereas the changing of beliefs is a strictly mental, agent-internal, activity. Therefore, it seems natural to let the ability of an agent to change its beliefs be determined by its mental state only.

The intuitive idea behind the definitions as we present them, is that the ability to change one's beliefs can be used to guide the changes that the beliefs of an agent undergo. In particular, if an agent is able to change its beliefs in a certain way, then this change of belief should neither result in an absurd belief set nor cause no change at all. Another point of attention is given by the observation that the Levi identity should also be respected for abilities, i.e., an agent is capable of revising its beliefs with a formula φ if and only if it is able to contract its beliefs with $\neg\varphi$ and thereafter perform an expansion with φ.

3.40. DEFINITION. Let $\mathcal{M}$ be some Kripke model with state s, let i be some agent and let φ be some propositional formula. The capability function `c` is for the `expand`, `contract` and `revise` actions defined by:

3.35. PROPOSITION (CHARACTERIZATION OF REVISIONS). *For all agents i, and for all propositional formulae φ and ψ we have:*

- $\models \mathbf{K}_i \neg\varphi \leftrightarrow [\mathrm{do}_i(\texttt{revise }\varphi)]\mathbf{B}_i\mathbf{ff}$
- $\models \neg\mathbf{B}_i\neg\varphi \rightarrow ([\mathrm{do}_i(\texttt{revise }\varphi)]\mathbf{B}_i\psi \leftrightarrow \mathbf{B}_i(\varphi \rightarrow \psi))$
- $\models \neg\mathbf{K}_i\neg\varphi \wedge \mathbf{B}_i\neg\varphi \rightarrow ([\mathrm{do}_i(\texttt{revise }\varphi)]\mathbf{B}_i\psi \leftrightarrow \mathbf{K}_i(\varphi \rightarrow \psi))$ *if the definition of* `r` *for the* `contract` *action is for all models based on the AiG function.*

The first clause of Proposition 3.35 states that in cases where an agent *knows* the negation of some formula to be true, a revision with this formula results in absurd beliefs. The second clause states that in situations where the negation of some formula is not *believed*, revising beliefs with the formula amounts to an expansion with the formula. The last clause is possibly the most mysterious one. It states that in situations that are not of the kinds described in the first two clauses, formulae are believed after a revision with φ if it is *known* on beforehand that φ implies the formula, i.e., the belief set of the agent after a revision with φ consists of all those formulae that are known to be implied by φ.

3.3.1. Revisions and the Gärdenfors postulates

The Gärdenfors postulates for belief revision are given below. In these postulates K, φ and K^+_φ are assumed to have their usual connotation, and $K^\star_\varphi$ denotes the revision of K with the formula φ. The absurd belief set, consisting of all formulae from the language, is denoted by $K_\perp$.

3.36. DEFINITION. The Gärdenfors postulates for belief revision:
($G^\star 1$) $K^\star_\varphi$ is a belief set.
($G^\star 2$) $\varphi \in K^\star_\varphi$.
($G^\star 3$) $K^\star_\varphi \subseteq K^+_\varphi$.
($G^\star 4$) If $\neg\varphi \notin K$ then $K^+_\varphi \subseteq K^\star_\varphi$.
($G^\star 5$) $K^\star_\varphi = K_\perp$ if and only if $\vdash \neg\varphi$.
($G^\star 6$) If $\vdash \varphi \leftrightarrow \psi$ then $K^\star_\varphi = K^\star_\psi$.
($G^\star 7$) $K^\star_{\varphi\wedge\psi} \subseteq (K^\star_\varphi)^+_\psi$.
($G^\star 8$) If $\neg\psi \notin K^\star_\varphi$, then $(K^\star_\varphi)^+_\psi \subseteq K^\star_{\varphi\wedge\psi}$.

When defining revision through the Levi identity, starting from expansions and contractions that satisfy the appropriate postulates, the Gärdenfors postulates for belief revision are met [5, 21]. The same holds in our framework, i.e., our `revision` action behaves like a belief revision in the sense of Gärdenfors.

3.37. DEFINITION. The *revision* of $\mathcal{B}(i, \mathcal{M}, s)$ with a propositional formula φ, notation $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s)$ is defined by:

$$\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) = \{\psi \in \mathcal{L}_0 \mid \mathcal{M}, s \models [\mathrm{do}_i(\texttt{revise }\varphi)]\mathbf{B}_i\psi\}$$

The sequence of a revision with φ followed by an expansion with ψ of $\mathcal{B}(i, \mathcal{M}, s)$, notation $\mathcal{B}^{\star+}_{\varphi\psi}(i, \mathcal{M}, s)$, is defined by:

$$\mathcal{B}^{\star+}_{\varphi\psi}(i, \mathcal{M}, s) = \{\vartheta \in \mathcal{L}_0 \mid \mathcal{M}, s \models [\mathrm{do}_i(\texttt{revise }\varphi; \texttt{expand }\psi)]\mathbf{B}_i\vartheta\}$$

3.33. PROPOSITION. *For all agents i, and for all propositional formulae φ, ψ and ϑ we have:*

- $\models [\mathrm{do}_i(\mathtt{revise}\ \varphi)]\mathbf{B}_i\varphi$
- $\models [\mathrm{do}_i(\mathtt{revise}\ \varphi)]\mathbf{B}_i\vartheta \rightarrow [\mathrm{do}_i(\mathtt{expand}\ \varphi)]\mathbf{B}_i\vartheta$
- $\models \neg\mathbf{B}_i\neg\varphi \rightarrow ([\mathrm{do}_i(\mathtt{expand}\ \varphi)]\mathbf{B}_i\vartheta \leftrightarrow [\mathrm{do}_i(\mathtt{revise}\ \varphi)]\mathbf{B}_i\vartheta)$
- $\models \mathbf{K}_i\neg\varphi \leftrightarrow [\mathrm{do}_i(\mathtt{revise}\ \varphi)]\mathbf{B}_i\mathbf{ff}$
- $\models \mathbf{K}_i(\varphi \leftrightarrow \psi) \rightarrow ([\mathrm{do}_i(\mathtt{revise}\ \varphi)]\mathbf{B}_i\vartheta \leftrightarrow [\mathrm{do}_i(\mathtt{revise}\ \psi)]\mathbf{B}_i\vartheta)$
- $\models [\mathrm{do}_i(\mathtt{revise}\ (\varphi \wedge \psi))]\mathbf{B}_i\vartheta \rightarrow [\mathrm{do}_i(\mathtt{revise}\ \varphi; \mathtt{expand}\ \psi)]\mathbf{B}_i\vartheta$
- $\models \neg[\mathrm{do}_i(\mathtt{revise}\ \varphi)]\mathbf{B}_i\neg\psi \rightarrow$
  $([\mathrm{do}_i(\mathtt{revise}\ \varphi; \mathtt{expand}\ \psi)]\mathbf{B}_i\vartheta \rightarrow [\mathrm{do}_i(\mathtt{revise}\ (\varphi \wedge \psi))]\mathbf{B}_i\vartheta)$

The first clause of Proposition 3.33 states that agents believe φ as the result of revising their beliefs with φ. The second clause states that a revision with φ results in a belief set that is contained in the belief set that results from an expansion with φ, i.e., changing the belief set to incorporate φ in a consistent manner (if possible) -this is a revision with φ- results in a subset of the set of beliefs that results from straightforward inserting φ in the belief set -an expansion with φ. The third clause formalizes the idea that expansion is a special kind of revision: in cases where $\neg\varphi$ is not believed, expanding with φ and revising with φ amount to the same action. The left-to-right implication of the fourth clause states that if $\neg\varphi$ is known, i.e., $\neg\varphi$ is one of the formulae that the agent will never part from, then the revision with φ results in the absurd belief set, i.e., the agent believes **ff** as a result of revising with φ. The right-to-left implication of the fourth clause states that the absurd belief set will result only if a revision with a non-revisable formula is performed. The fifth clause states that revisions with formulae that are known to be equivalent have identical results. The sixth clause formalizes the idea that the revision with the conjunction $\varphi \wedge \psi$ results in a belief set that is a subset of the belief set that results from a revision with φ followed by an expansion with ψ. The last clause states that if a revision with φ does not result in $\neg\psi$ being believed, then the belief set that results from revising with $\varphi \wedge \psi$ is a superset of the belief set that results from a revision with φ followed by an expansion with ψ. As Gärdenfors remarks, these last two clauses provide for some sort of *minimal change* condition on revisions.

As was the case for expansions and contractions, revisions are realizable, deterministic and idempotent.

3.34. PROPOSITION. *For all agents i, for all propositional formulae φ, and for all formulae ψ we have:*

- $\models \langle\mathrm{do}_i(\mathtt{revise}\ \varphi)\rangle\mathbf{tt}$
- $\models \langle\mathrm{do}_i(\mathtt{revise}\ \varphi)\rangle\psi \leftrightarrow [\mathrm{do}_i(\mathtt{revise}\ \varphi)]\psi$
- $\models \langle\mathrm{do}_i(\mathtt{revise}\ \varphi; \mathtt{revise}\ \varphi)\rangle\psi \leftrightarrow \langle\mathrm{do}_i(\mathtt{revise}\ \varphi)\rangle\psi$

The belief sets resulting from application of the `revise` action can be characterized analogously to the characterization of `contract` as given in Proposition 3.27.

would result in selection functions, and *vice versa*. To investigate this relation in our framework we introduce the notion of *s-selection functions*. Basically these s-selection functions can be seen as conditional selection functions that are at some points adapted to make them more in line with our framework.

3.29. DEFINITION. Let some model $\mathcal{M}$ be given. A function $s : \mathcal{A} \times \mathcal{S} \times \wp(\mathcal{S})$ is an *s-selection function* for $\mathcal{M}$ if and only if it meets the following constraints for all $i \in \mathcal{A}, s \in \mathcal{S}$ and for all propositional formulae φ and ψ.

- $S1$. $\neg\varphi$ is true at all states from $s(i, s, \varphi)$
- $S2$. $s(i, s, \varphi)$ is empty only if s' is (epistemically) inaccessible from s for all worlds s' in which $\neg\varphi$ holds.
- $S3$. if $\mathrm{B}(i, s) \cap [\![\neg\varphi]\!] \neq \emptyset$ then $s(i, s, \varphi) \subseteq \mathrm{B}(i, s)$
- $S4$. if $\neg\varphi$ is true at all states from $s(i, s, \psi)$ and $\neg\psi$ is true at all states from $s(i, s, \varphi)$, then $s(i, s, \varphi) = s(i, s, \psi)$

It turns out that all selection functions for a given model, are also s-selection functions for the model. The converse does however not hold.

3.30. PROPOSITION. *For all models $\mathcal{M}$, and for all functions $\varsigma : \mathcal{A} \times \mathcal{S} \times \mathcal{L}_0 \rightarrow \wp(\mathcal{S})$ it holds that if ς is a selection function for $\mathcal{M}$ then ς is an s-selection function for $\mathcal{M}$.*

3.31. PROPOSITION. *Some Kripke model $\mathcal{M}$ and function $\varsigma : \mathcal{A} \times \mathcal{S} \times \mathcal{L}_0 \rightarrow \wp(\mathcal{S})$ exists, such that ς is an s-selection function for $\mathcal{M}$, but not a selection function. Furthermore, when defining contractions based on the function ς, not all Gärdenfors postulates for belief contraction are validated.*

3.3. The `revise` action

Having defined actions that model expansions and contractions, we now define actions that model revisions. A revision is a change of belief in which the belief status of some formula is reversed, i.e., for some formula $\neg\varphi$ that is believed on beforehand, φ is believed as the result of a revision with φ. In defining actions that model revisions we essentially use the *Levi identity* [21]. Levi showed that revisions can be defined in terms of contractions and expansions: a revision with some formula φ can be brought about as a contraction with $\neg\varphi$ followed by an expansion with φ. Given the definitions of contractions and expansions of the previous sections and the fact that the class of actions Ac that we consider is closed under sequential composition, the Levi identity provides for a means to define revisions as the sequential composition of a contraction and an expansion action.

3.32. DEFINITION. Let some model $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathtt{c} \rangle$ with $s \in \mathcal{S}$, and some agent i and propositional formula φ be given. We define:

$$\mathtt{r}(i, \mathtt{revise}\ \varphi)(\mathcal{M}, s) = \mathtt{r}(i, \mathtt{contract}\ \neg\varphi; \mathtt{expand}\ \varphi)(\mathcal{M}, s)$$

Definition 3.32 indeed provides for an intuitively acceptable formalization of belief revision.

The belief states following application of the `contract` action can completely be characterized in terms of *a priori* information, i.e., knowledge and belief, of the agent. In one of the clauses given below it is presupposed that `r` for the `contract` action is based on the AiG function, the other clause holds for general selection functions.

3.27. PROPOSITION (CHARACTERIZATION OF CONTRACTIONS). *For all agents i, and for all propositional formulae φ and ψ we have:*

- $\models \neg\mathbf{B}_i\varphi \rightarrow ([\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{B}_i\psi \leftrightarrow \mathbf{B}_i\psi)$
- $\models \mathbf{B}_i\varphi \rightarrow ([\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{B}_i\psi \leftrightarrow (\mathbf{B}_i\psi \wedge \mathbf{K}_i(\neg\varphi \rightarrow \psi)))$ *if the definition of* `r` *for the* `contract` *action is for all models based on the AiG function.*

Again the result of Proposition 3.27 can be rephrased to make it more in line with a result stated by Gärdenfors [5] for full meet contraction functions.

3.28. PROPOSITION. *For all Kripke models $\mathcal{M}$ with state s, for all agents i, and for all propositional formulae φ and ψ we have:*

- $\varphi \notin \mathcal{B}(i,\mathcal{M},s) \Rightarrow (\psi \in \mathcal{B}^-_\varphi(i,\mathcal{M},s) \Leftrightarrow \psi \in \mathcal{B}(i,\mathcal{M},s))$
- $\varphi \in \mathcal{B}(i,\mathcal{M},s) \Rightarrow$
  $(\psi \in \mathcal{B}^-_\varphi(i,\mathcal{M},s) \Leftrightarrow \psi \in \mathcal{B}(i,\mathcal{M},s)\,\&\,(\neg\varphi \rightarrow \psi) \in \mathcal{K}(i,\mathcal{M},s))$ *if the definition of* `r` *for the* `contract` *action is based on the AiG function for $\mathcal{M}$.*

Besides the relation between our kind of selection functions and the partial meet contraction functions of Gärdenfors, another interesting relation exists between our selection functions and those defined by Stalnaker [33]. Stalnaker uses selection functions (to avoid confusion we use the term *conditional selection functions* to refer to selection functions in the sense of Stalnaker) in the context of a Kripke style semantics for conditional logic. Given a Kripke model $\mathcal{M}$ and a state s in $\mathcal{M}$, a conditional selection f when applied to a pair (φ, s) yields the *most preferred* or *most reasonable* world given φ and s. Stalnaker gives four demands that a reasonable conditional selection function should meet:

1. φ is true at $f(\varphi, s)$.
2. $f(\varphi, s)$ is undefined only if s' is inaccessible from s for all worlds s' in which φ holds.
3. if φ is true at s then $f(\varphi, s) = s$.
4. if φ is true at $f(\psi, s)$ and ψ is true at $f(\varphi, s)$ then $f(\varphi, s) = f(\psi, s)$.

Intuitively, there seems to be at least some resemblance between Stalnaker's ideas underlying conditional selection function and the ideas underlying our selection functions. For although conditional selection functions aim at yielding a *single world* that *satisfies* a given formula and selection functions aim at yielding a *set of worlds* that *falsify* a given formula, both aim at yielding *reasonable* results. For conditional selection functions this reasonableness is enforced through the demands given above whereas for selection functions this is enforced through the demands $\Sigma 1$ through $\Sigma 6$ as given in Def. 3.14. One could ask whether imposing demands similar to those proposed by Stalnaker of maximal subsets of K that do not entail φ. More formal, $K^-_\varphi = \cap S(K \bot \varphi)$, where $K \bot \varphi$ is the set of belief sets K' that fail to imply φ and are maximal subsets of K, and S is a function that selects some of the elements of $K \bot \varphi$.

The intuitive idea behind our notion of selection functions resembles that of partial meet contraction functions. Whereas in partial meet contraction functions some of the maximal subsets not implying the contracted formula are selected, our selection function selects some of the maximal subsets of the belief set *given the model*. This can be seen as follows. Assume some Kripke model $\mathcal{M}$ with state s and agent i to be given. Assume furthermore that $\mathcal{M}, s \models \mathbf{B}_i\varphi \wedge \neg\mathbf{K}_i\varphi$ for some propositional formulae φ. Given this model, it is obvious that adding any of the worlds from $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]$ to the set of doxastic alternatives of i results in a model in which the agent no longer believes φ. Furthermore, under the constraints implied by Fig. 3, it is even so obvious that it is sufficient to add exactly one of the worlds from $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]$ in order to result in a model in which the agent no longer believes φ. From this point of view, a selection function selects some of the maximal subsets of a belief set, and the resulting belief set is the intersection of these maximal subsets, all with respect to the model $\mathcal{M}$. The following proposition formalizes these informal ideas.

3.24. PROPOSITION. *Let some Kripke model $\mathcal{M}$ with state s, and agent i be given. Let σ be some selection function for $\mathcal{M}$. Define for propositional formulae φ:*

- $\mathcal{B}(i, \mathcal{M}, s) \bot \varphi = \{\{\mathcal{B}(i, \mathcal{M}, s)\}\}$ *if* $\varphi \in \mathcal{K}(i, \mathcal{M}, s)$
- $\mathcal{B}(i, \mathcal{M}, s) \bot \varphi = \{\{\psi \in \mathcal{L}_0 \mid \forall s' \in \mathrm{B}(i, s) \cup \{s''\}[\mathcal{M}, s' \models \psi]\} \mid s'' \in [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]\}$ *if* $\varphi \notin \mathcal{K}(i, \mathcal{M}, s)$
- $S(\mathcal{B}(i, \mathcal{M}, s) \bot \varphi) = \{\{\psi \in \mathcal{L}_0 \mid \forall s' \in \mathrm{B}(i, s) \cup \{s''\}[\mathcal{M}, s' \models \psi]\} \mid s'' \in \sigma(i, s, \varphi)\}$

Then $\mathcal{B}^-_\varphi(i, \mathcal{M}, s) = \cap S(\mathcal{B}(i, \mathcal{M}, s) \bot \varphi)$.

The resemblance with the partial meet contraction functions suggest a concrete implementation of selection functions. The *full meet contraction function* of Gärdenfors defines the contraction of a belief set K with a formula φ to be the intersection of all maximal subsets of K that do not imply φ. In our framework this amounts to a selection function that simply adds all the possible worlds from the epistemic equivalence class that do not support the formula that is to be contracted, if the formula is accepted on beforehand. If the formula is not accepted on beforehand, no new worlds are added.

3.25. DEFINITION (All is Good). Let $\mathcal{M}$ be a Kripke model. The AiG function σ_a is for all $i \in \mathcal{A}$, $s \in \mathcal{M}$, and φ in $\mathcal{L}_0$ defined by:

- $\sigma_a(i, s, \varphi) = \mathrm{B}(i, s) \cap [\![\neg\varphi]\!]$ if $\mathrm{B}(i, s) \cap [\![\neg\varphi]\!] \neq \emptyset$.
- $\sigma_a(i, s, \varphi) = [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]$ otherwise.

3.26. PROPOSITION. *The AiG function σ_a as given in Def. 3.25 is a selection function.*

$\varsigma(1, s, \varphi) = \{t_0, u_0\}$ if $\mathcal{M}, t_0 \not\models \varphi$
$\varsigma(1, s, \varphi) = \{t_0\}$ if $\mathcal{M} \models \varphi$
$\varsigma(1, s, \varphi) = \{t \in \mathcal{T} \mid \mathcal{M}, t \not\models \varphi\} \cup \{t_0\}$ if $\mathcal{M}, t_0 \models \varphi \,\&\, \mathcal{M} \not\models \varphi \,\&\, \varphi \vdash_{cl} p$
$\varsigma(1, s, \varphi) = \{u \in \mathcal{U} \mid \mathcal{M}, u \not\models \varphi\} \cup \{t_0\}$ if $\mathcal{M}, t_0 \models \varphi \,\&\, \mathcal{M} \not\models \varphi \,\&\, \varphi \not\vdash_{cl} p$

The function ς is not a selection function for $\mathcal{M}$. In particular, ς does meet only one of the demands given in Def. 3.14. To see this, take some arbitrary $s \in \mathcal{S}$.

- Since ς yields identical results for all $s \in \mathcal{S}$, demand $\Sigma 0$ is met.
- Since $t_0 \in \varsigma(1, s, p)$ and $t_0 \notin [\![\neg p]\!]$, demand $\Sigma 1$ is not met.
- Since $u_0 \in \varsigma(1, s, \neg p)$ and $u_0 \notin \mathrm{B}(1, s)$, demand $\Sigma 2$ is not met.
- Although $\mathcal{S} \cap [\![\neg \mathbf{tt}]\!] = \emptyset$, $\varsigma(1, s, \mathbf{tt}) = \{t_0\}$, and hence demand $\Sigma 3$ is not met.
- Although $\mathcal{M}, t_0 \models \mathbf{K}_1(\neg q \leftrightarrow (p \wedge \neg q))$, $\varsigma(1, s, \neg q) = \{t_0, u_0, u_2\}$ whereas $\varsigma(1, s, p \wedge \neg q) = \{t_0, t_2\}$ and hence demand Σ_4 is not met.
- Since it is the case that $\varsigma(1, s, p \wedge q) = \{t_0, t_1, t_2\}$, $\varsigma(1, s, p) = \{t_0, t_2\}$ and $\varsigma(1, s, q) = \{t_0, u_1, u_2\}$, $\varsigma(1, s, p \wedge q) \not\subseteq \varsigma(1, s, p) \cup \varsigma(1, s, q)$, and hence demand $\Sigma 5$ is not met.
- Although $\varsigma(1, s, q \wedge p) \cap [\![\neg q]\!] = \{t_1, t_2\} \neq \emptyset$, $\varsigma(1, s, q) \not\subseteq \varsigma(1, s, q \wedge p)$ and hence demand $\Sigma 6$ is not met.

Hence ς is by no means a selection function for $\mathcal{M}$. It can however be checked that when defining $\mathtt{r}(1, \mathtt{contract}\ \varphi)(\mathcal{M}, s)$ based on the non-selection function ς, all the postulates given in Theorem 3.20 are validated[3].

Despite the negative results of Example 3.22, we can prove that when defining $\mathtt{r}$ for the $\mathtt{contract}$ action based on some function ς that adds doxastic alternatives (hence in accordance with Fig. 3), validation of the Gärdenfors postulates imposes some weak variants of the demands for selection functions on ς.

3.23. Proposition. *Let $\mathcal{M}$ be some Kripke model with state s, and let i be some agent. Assume that $\mathtt{r}(i, \mathtt{contract}\ \varphi)(\mathcal{M}, s)$ is defined as in 3.15 with σ replaced by an arbitrary function $\varsigma : \mathcal{A} \times \mathcal{S} \times \mathcal{L}_0 \to \wp(\mathcal{S})$. Then if* $\mathtt{contract}$ *is to meet the demands presented in Theorem 3.20 it follows that:*

- $\varsigma(i, s, \varphi) \subseteq [s]_{\mathrm{R}(i)}$
- $\mathrm{B}(i, s) \cap [\![\neg\varphi]\!] = \emptyset \,\&\, [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] \neq \emptyset \Rightarrow \varsigma(i, s, \varphi) \cap [\![\neg\varphi]\!] \neq \emptyset$

The fact that our approach using selection functions defines a contraction function which satisfies the Gärdenfors postulates is not as surprising as it might seem at first sight. Fact of the matter is that our selection functions can be seen as *model-based, knowledge-restricted partial meet contraction functions* as defined by Gärdenfors [5]. A partial meet contraction function performs contractions as follows. Given a belief set K and a formula φ that is to be contracted, a partial meet contraction function yields the intersection of a set

[3] One way to see this is by remarking that ς is a variant of the selection function σ_a as presented in Def. 3.25 in which it is used that $\mathcal{T}$ and $\mathcal{U}$ worlds are identical on the propositional level.

(B^-1) *$\mathcal{B}^-_\varphi(i, \mathcal{M}, s)$ is a belief set.*
(B^-2) *$\mathcal{B}^-_\varphi(i, \mathcal{M}, s) \subseteq \mathcal{B}(i, \mathcal{M}, s)$.*
(B^-3) *If $\varphi \notin \mathcal{B}(i, \mathcal{M}, s)$ then $\mathcal{B}^-_\varphi(i, \mathcal{M}, s) = \mathcal{B}(i, \mathcal{M}, s)$.*
(B^-4) *If $\mathcal{M}, s \not\models \mathbf{K}_i\varphi$ then $\varphi \notin \mathcal{B}^-_\varphi(i, \mathcal{M}, s)$.*
(B^-5) *If $\varphi \in \mathcal{B}(i, \mathcal{M}, s)$ then $\mathcal{B}(i, \mathcal{M}, s) \subseteq \mathcal{B}^{-+}_{\varphi\varphi}(i, \mathcal{M}, s)$.*
(B^-6) *If $\mathcal{M}, s \models \mathbf{K}_i(\varphi \leftrightarrow \psi)$ then $\mathcal{B}^-_\varphi(i, \mathcal{M}, s) = \mathcal{B}^-_\psi(i, \mathcal{M}, s)$.*
(B^-7) *$\mathcal{B}^-_\varphi(i, \mathcal{M}, s) \cap \mathcal{B}^-_\psi(i, \mathcal{M}, s) \subseteq \mathcal{B}^-_{\varphi\wedge\psi}(i, \mathcal{M}, s)$.*
(B^-8) *If $\varphi \notin \mathcal{B}^-_{\varphi\wedge\psi}(i, \mathcal{M}, s)$ then $\mathcal{B}^-_{\varphi\wedge\psi}(i, \mathcal{M}, s) \subseteq \mathcal{B}^-_\varphi(i, \mathcal{M}, s)$.*

3.2.2. Selection functions revisited

There is a more or less direct relation between the demands imposed on selection functions and the Gärdenfors postulates for contraction. This relation is given in the following proposition, leading to a refinement of the results obtained in Theorem 3.20.

3.21. Proposition. *Let $\mathcal{M}$ be some Kripke model with state s, let i be some agent and φ some propositional formula. Let $\mathtt{r}(i, \mathtt{contract}\ \varphi)(\mathcal{M}, s)$ be defined as in 3.15 but with σ replaced by an arbitrary function $\varsigma : \mathcal{A} \times \mathcal{S} \times \mathcal{L}_0 \to \wp(\mathcal{S})$. Then it holds that:*

- *If ς satisfies $\Sigma 1$ then B^-5 is validated.*
- *If ς satisfies $\Sigma 2$ then B^-3 is validated.*
- *Given that ς satisfies $\Sigma 1$ it is the case that if ς satisfies $\Sigma 3$ then B^-4 is validated.*
- *If ς satisfies $\Sigma 4$ then B^-6 is validated.*
- *If ς satisfies $\Sigma 5$ then B^-7 is validated.*
- *If ς satisfies $\Sigma 6$ then B^-8 is validated.*

The implications given in Proposition 3.21 cannot be generalized to equivalences. That is, the conditions imposed on selection functions are *sufficient* to bring about validation of the postulates for belief contraction but are not *necessary* to do so. The following example sheds some more light on this issue.

3.22. Example. Consider the single-agent language $\mathcal{L}$, based on $\Pi = \{p, q\}$ and $At = \{a\}$. Consider the Kripke model $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathsf{c} \rangle$ where

- $\mathcal{S} = \mathcal{T} \cup \mathcal{U}, \mathcal{T} = \{t_0, t_1, t_2\}, \mathcal{U} = \{u_0, u_1, u_2\}$,
- $\pi(p, t_j) = \mathbf{1}$ iff $\pi(p, u_j) = \mathbf{1}$ iff $j = 0$ or $j = 1$
  $\pi(q, t_j) = \mathbf{1}$ iff $\pi(q, u_j) = \mathbf{1}$ iff $j = 0$,
- $\mathrm{R}(1) = \mathcal{S}^2$,
- $\mathrm{B}(1, s) = \{t_0\}$ for all $s \in \mathcal{S}$,
- $\mathtt{r}$ is arbitrary,
- $\mathsf{c}(1, a)(t) = \mathbf{1}$ for all $t \in \mathcal{T}$, $\mathsf{c}(1, a)(u) = \mathbf{0}$ for all $t \in \mathcal{U}$.

Note that although the elements of $\mathcal{T}$ are copies of the elements of $\mathcal{U}$ (and vice versa) on the propositional level, they do not satisfy the same set of formula. For $\mathcal{M}, t \models \mathbf{A}_1 a$ for each $t \in \mathcal{T}$ whereas $\mathcal{M}, u \models \neg\mathbf{A}_1 a$ for all $u \in \mathcal{U}$. Define the function $\varsigma : \mathcal{A} \times \mathcal{S} \times \mathcal{L}_0 \to \wp(\mathcal{S})$ for all $s \in \mathcal{S}$ as follows:

believed both after a contraction with φ and after a contraction with ψ are believed after a contraction with $\varphi \wedge \psi$. The last clause states that if a contraction with $\varphi \wedge \psi$ results in φ not being believed, then in order to contract φ no more formulae need to be removed than those that were removed in order to contract $\varphi \wedge \psi$. This last clause is related to the property of *minimal change* for contractions.

As was the case for expansions, contractions are also realizable, deterministic, and idempotent.

3.17. PROPOSITION. *For all agents i, for all propositional formulae φ, and for all formulae ψ we have:*

- $\models \langle \mathrm{do}_i(\mathtt{contract}\ \varphi)\rangle \mathbf{tt}$
- $\models \langle \mathrm{do}_i(\mathtt{contract}\ \varphi)\rangle \psi \leftrightarrow [\mathrm{do}_i(\mathtt{contract}\ \varphi)]\psi$
- $\models \langle \mathrm{do}_i(\mathtt{contract}\ \varphi; \mathtt{contract}\ \varphi)\rangle \psi \leftrightarrow \langle \mathrm{do}_i(\mathtt{contract}\ \varphi)\rangle \psi$

3.2.1. Contractions and the Gärdenfors postulates

The Gärdenfors postulates for belief contraction are given below. In these postulates K, φ, ψ and K^+_φ are assumed to have their usual connotation, and K^-_φ denotes the contraction of K with the formula φ.

3.18. DEFINITION. The Gärdenfors postulates for belief contraction:
(G^-1) K^-_φ is a belief set.
(G^-2) $K^-_\varphi \subseteq K$.
(G^-3) If $\varphi \notin K$ then $K^-_\varphi = K$.
(G^-4) If $\not\vdash \varphi$ then $\varphi \notin K^-_\varphi$.
(G^-5) If $\varphi \in K$ then $K \subseteq (K^-_\varphi)^+_\varphi$.
(G^-6) If $\vdash \varphi \leftrightarrow \psi$ then $K^-_\varphi = K^-_\psi$.
(G^-7) $K^-_\varphi \cap K^-_\psi \subseteq K^-_{\varphi\wedge\psi}$.
(G^-8) If $\varphi \notin K^-_{\varphi\wedge\psi}$ then $K^-_{\varphi\wedge\psi} \subseteq K^-_\varphi$.

Using the definition of the `contract` action as given in 3.15, it is indeed the case that this action models contractions in the sense of Gärdenfors. As was the case for belief expansions, we have to modify the postulates for belief contraction somewhat to account for the agent-oriented, semantics based character of our framework.

3.19. DEFINITION. The *contraction* of $\mathcal{B}(i, \mathcal{M}, s)$ with a propositional formula φ, notation $\mathcal{B}^-_\varphi(i, \mathcal{M}, s)$ is defined by:

$$\mathcal{B}^-_\varphi(i, \mathcal{M}, s) = \{\psi \in \mathcal{L}_0 \mid \mathcal{M}, s \models [\mathrm{do}_i(\mathtt{contract}\ \varphi)]\mathbf{B}_i\psi\}$$

The sequence of a contraction with φ followed by an expansion with ψ of $\mathcal{B}(i, \mathcal{M}, s)$, notation $\mathcal{B}^{-+}_{\varphi\psi}(i, \mathcal{M}, s)$, is defined by:

$$\mathcal{B}^{-+}_{\varphi\psi}(i, \mathcal{M}, s) = \{\vartheta \in \mathcal{L}_0 \mid \mathcal{M}, s \models [\mathrm{do}_i(\mathtt{contract}\ \varphi; \mathtt{expand}\ \psi)]\mathbf{B}_i\vartheta\}$$

3.20. THEOREM. *Let $\mathcal{M}$ be some Kripke model. For all agents i, for all $s \in \mathcal{M}$ and for all formulae φ and ψ the following are true.*

for all $i \in \mathcal{A}, s, s' \in \mathcal{S}$ and for all propositional formulae φ and ψ.

Σ0. $\sigma(i,s,\varphi) = \sigma(i,s',\varphi)$ if $s' \in [s]_{\mathrm{R}(i)}$.

Σ1. $\sigma(i,s,\varphi) \subseteq [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]$.

Σ2. $\sigma(i,s,\varphi) \subseteq \mathrm{B}(i,s)$ if $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] \neq \emptyset$.

Σ3. $\sigma(i,s,\varphi) = \emptyset$ iff $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] = \emptyset$.

Σ4. if $\mathcal{M}, s \models \mathbf{K}_i(\varphi \leftrightarrow \psi)$ then $\sigma(i,s,\varphi) = \sigma(i,s,\psi)$.

Σ5. $\sigma(i,s,\varphi \wedge \psi) \subseteq \sigma(i,s,\varphi) \cup \sigma(i,s,\psi)$.

Σ6. if $\sigma(i,s,\varphi \wedge \psi) \cap [\![\neg\varphi]\!] \neq \emptyset$ then $\sigma(i,s,\varphi) \subseteq \sigma(i,s,\varphi \wedge \psi)$.

The definition of the `contract` action is based on the use of selection functions: a contraction is performed by adding exactly those worlds that are selected by the selection function to the set of doxastic alternatives of the agent.

3.15. Definition. Let some model $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathtt{c} \rangle$ with $s \in \mathcal{S}$, and an agent i and propositional formula φ be given. Let furthermore σ be an arbitrary but fixed selection function for $\mathcal{M}$. We define:

$$
\begin{aligned}
&\mathtt{r}(i, \mathtt{contract}\ \varphi)(\mathcal{M}, s) = \mathcal{M}', s \text{ where}\\
&\quad \mathcal{M}' = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}', \mathtt{r}, \mathtt{c} \rangle \text{ with}\\
&\qquad \mathrm{B}'(i', s') = \mathrm{B}(i', s') \text{ if } i' \neq i \text{ or } s' \notin [s]_{\mathrm{R}(i)}\\
&\qquad \mathrm{B}'(i, s') = \mathrm{B}(i, s') \cup \sigma(i, s, \varphi) \text{ for all } s' \in [s]_{\mathrm{R}(i)}
\end{aligned}
$$

Using selection functions to define the semantics for the `contract` action indeed results in an acceptable formalization of belief contraction.

3.16. Proposition. *For all agents i, and propositional formulae φ, ψ and ϑ we have:*

- $\models [\mathrm{do}_i(\mathtt{contract}\ \varphi)]\mathbf{B}_i\psi \to \mathbf{B}_i\psi$
- $\models \neg\mathbf{B}_i\varphi \to ([\mathrm{do}_i(\mathtt{contract}\ \varphi)]\mathbf{B}_i\psi \leftrightarrow \mathbf{B}_i\psi)$
- $\models \neg\mathbf{K}_i\varphi \to [\mathrm{do}_i(\mathtt{contract}\ \varphi)]\neg\mathbf{B}_i\varphi$
- $\models \mathbf{B}_i\varphi \to (\mathbf{B}_i\psi \to [\mathrm{do}_i(\mathtt{contract}\ \varphi; \mathtt{expand}\ \varphi)]\mathbf{B}_i\psi)$
- $\models \mathbf{K}_i(\varphi \leftrightarrow \psi) \to ([\mathrm{do}_i(\mathtt{contract}\ \varphi)]\mathbf{B}_i\vartheta \leftrightarrow [\mathrm{do}_i(\mathtt{contract}\ \psi)]\mathbf{B}_i\vartheta)$
- $\models ([\mathrm{do}_i(\mathtt{contract}\ \varphi)]\mathbf{B}_i\vartheta \wedge [\mathrm{do}_i(\mathtt{contract}\ \psi)]\mathbf{B}_i\vartheta) \to$ $[\mathrm{do}_i(\mathtt{contract}\ (\varphi \wedge \psi))]\mathbf{B}_i\vartheta$
- $\models [\mathrm{do}_i(\mathtt{contract}\ (\varphi \wedge \psi))]\neg\mathbf{B}_i\varphi \to$ $([\mathrm{do}_i(\mathtt{contract}\ (\varphi \wedge \psi))]\mathbf{B}_i\vartheta \to [\mathrm{do}_i(\mathtt{contract}\ \varphi)]\mathbf{B}_i\vartheta)$

The first clause of Proposition 3.16 states that a contraction results in a belief set that is contained in the belief set before the contraction. The second clause state that in situations in which φ is not believed, nothing changes as the result of contracting φ. Again this property reflects the criterion of informational economy. The third clause states that a contraction with a contractable formula, this is a formula that the agent is willing to part from, results in the agent not believing the contracted formula. The fourth clause states that all beliefs in the original belief set are recovered after a contraction with a formula followed by an expansion with the same formula. The fifth clause states that contractions with formulae that are known to be equivalent, result in identical belief sets. The sixth clause formalizes the idea that all formulae that are

3.2. The `contract` action

A belief contraction is the change of belief through which in general some formula that is believed on beforehand is no longer believed afterwards. As such, apparent beliefs that an agent has are turned into doubts as the result of a contraction. In terms of our framework this comes down to *expanding* the set of doxastic alternatives of an agent in order to encompass at least one state not satisfying the formula that is to be contracted. Our approach towards belief contraction is based on the following ideas.

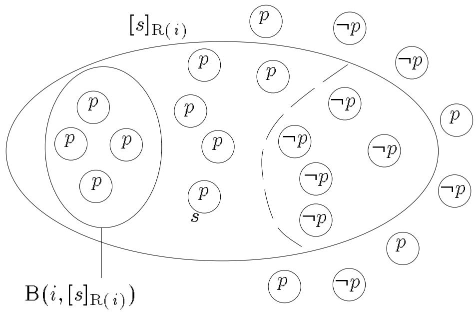

Figure 3: Contraction in Kripke models

In Fig. 3 agent i believes p since p holds in all its doxastic alternatives. When contracting p from the belief set of the agent, some of the $\neg p$ worlds are added to the set of doxastic alternatives of the agent. In order to end up with well-defined Kripke models, these worlds that are to be added, need to be in the set of epistemic alternatives of s. For in the Kripke models defined in 2.3, the set of doxastic alternatives for a given agent in a given state is contained in its set of epistemic alternatives in that state. Thus the worlds that are to be added to the set of doxastic alternatives of the agent are elements of the set of epistemic alternatives not supporting p. In Fig. 3 this is the rightmost part of $[s]_{R(i)}$ separated through the dotted line.

The problem with contractions defined in this way is that it is not straightforward as to decide which worlds need to be added. From the basic idea that knowledge -acting as the principles of agents- provides some sort of lower bound of the belief set of an agent, it is clear that in the case of a contraction with φ some states need to be added that are elements of the set of epistemic alternatives of the agent and do not support φ, but it is not clear exactly *which* elements of this set need to be chosen. That is, in Fig. 3 it is clear that the worlds to be chosen should be among the five rightmost worlds in $[s]_{R(i)}$, but it is not clear *which* of these worlds need to be added.

The approach that we propose is based on the use of so called *selection functions*. These are functions that (whenever possible) select a subset of the set of epistemic alternatives in such a way that the resulting `contract` action behaves in a reasonable, intuitively acceptably way.

3.14. Definition. Let some model $\mathcal{M}$ be given. A function $\sigma : \mathcal{A} \times \mathcal{S} \times \mathcal{L}_0 \rightarrow \wp(\mathcal{S})$ is a *selection function* for $\mathcal{M}$ if and only if it meets the following constraints

(G^+2) $\varphi \in K^+_\varphi$.
(G^+3) $K \subseteq K^+_\varphi$.
(G^+4) If $\varphi \in K$ then $K^+_\varphi = K$.
(G^+5) If $K \subseteq H$ then $K^+_\varphi \subseteq H^+_\varphi$.
(G^+6) For all belief sets K, and all sentences φ, K^+_φ is the smallest set that satisfies G^+1 – G^+5.

It turns out that our `expand` action can be seen as providing a belief expansion in the sense of Gärdenfors. To formulate the Gärdenfors postulates in our framework we introduce our own kind of belief sets. These belief sets are model-based and indexed with a particular agent. Furthermore the notion of knowledge sets, as providing the principles or prejudices that the agent will never part from, is defined below.

3.11. Definition. Let $\mathcal{M}$ be some Kripke model with $s \in \mathcal{M}$. The *belief set* of agent i in $\mathcal{M}, s$, notation $\mathcal{B}(i, \mathcal{M}, s)$, is defined by:

$$\mathcal{B}(i, \mathcal{M}, s) = \{\varphi \in \mathcal{L}_0 \mid \mathcal{M}, s \models \mathbf{B}_i\varphi\}$$

The *knowledge set* of agent i in $\mathcal{M}, s$, notation $\mathcal{K}(i, \mathcal{M}, s)$, is defined by:

$$\mathcal{K}(i, \mathcal{M}, s) = \{\varphi \in \mathcal{L}_0 \mid \mathcal{M}, s \models \mathbf{K}_i\varphi\}$$

The *expansion* of $\mathcal{B}(i, \mathcal{M}, s)$ with a formula $\varphi \in \mathcal{L}_0$, notation $\mathcal{B}^+_\varphi(i, \mathcal{M}, s)$, is defined by:

$$\mathcal{B}^+_\varphi(i, \mathcal{M}, s) = \{\psi \in \mathcal{L}_0 \mid \mathcal{M}, s \models [\mathrm{do}_i(\mathtt{expand}\ \varphi)]\mathbf{B}_i\psi\}$$

The unique absurd belief set $\mathcal{B}_\perp$ is defined to be $\mathcal{L}_0$.

3.12. Theorem. *Let $\mathcal{M}$ and $\mathcal{M}'$ be Kripke models, with s in $\mathcal{M}$ and s' in $\mathcal{M}'$, and let i be some agent. The following is valid for all $\varphi \in \mathcal{L}_0$.*
(B^+1) *$\mathcal{B}^+_\varphi(i, \mathcal{M}, s)$ is a belief set, i.e., $\mathcal{B}^+_\varphi(i, \mathcal{M}, s)$ is either equal to $\mathcal{B}_\perp$ or it is consistent and deductively closed.*
(B^+2) *$\varphi \in \mathcal{B}^+_\varphi(i, \mathcal{M}, s)$.*
(B^+3) *$\mathcal{B}(i, \mathcal{M}, s) \subseteq \mathcal{B}^+_\varphi(i, \mathcal{M}, s)$.*
(B^+4) *If $\varphi \in \mathcal{B}(i, \mathcal{M}, s)$ then $\mathcal{B}^+_\varphi(i, \mathcal{M}, s) = \mathcal{B}(i, \mathcal{M}, s)$.*
(B^+5) *If $\mathcal{B}(i, \mathcal{M}, s) \subseteq \mathcal{B}(i, \mathcal{M}', s')$, then $\mathcal{B}^+_\varphi(i, \mathcal{M}, s) \subseteq \mathcal{B}^+_\varphi(i, \mathcal{M}', s')$.*
(B^+6) *$\mathcal{B}^+_\varphi(i, \mathcal{M}, s)$ is the smallest set that satisfies the postulates B^+1 – B^+5 as given above.*

Gärdenfors shows that the postulates formulated in Def. 3.10 completely establish expansions [5]. Whereas the postulates for contraction and revision leave some degrees of freedom in defining contractions and revisions, G^+1 through G^+5 uniquely define expansions. Proposition 3.13 -a rephrasing of Proposition 3.8- states that the same holds in our framework, and furthermore, the unique definition that we end up with is identical to the one given by Gärdenfors.

3.13. Proposition. *For all Kripke models $\mathcal{M}$ with states s, for all agents i, and for all propositional formulae φ we have:*

$$\mathcal{B}^+_\varphi(i, \mathcal{M}, s) = \mathrm{Th}(\mathcal{B}(i, \mathcal{M}, s) \cup \{\varphi\})$$

in situations like these. Note that the first two clauses combined indicate that our definition of belief, and in particular the fact that we allow absurd belief sets, is a natural one. For an expansion with some formula φ in a situation in which $\neg\varphi$ is already believed, results in the agent believing both φ and $\neg\varphi$ and hence having inconsistent beliefs. The last clause states that in situations where some formula is already believed, nothing is changed as the result of an expansion with that formula. This latter property is suggested by the *criterion of informational economy* [5], which states that since information is in general not gratuitous unnecessary losses of information are to be avoided.

The `expand` action is furthermore *realizable*, *deterministic* and *idempotent*. Realizability of an action implies that agents have the opportunity to perform the action regardless of the situation, determinism of an action means that performing the action results in a unique state of affairs, and idempotence of an action implies that performing the action twice -or an arbitrary number of times- has the same effect as performing the action just once.

3.7. PROPOSITION. *For all agents i, for all propositional formulae φ, and for all formulae ψ we have:*

- $\models \langle \mathrm{do}_i(\texttt{expand}\ \varphi)\rangle \mathbf{tt}$
- $\models \langle \mathrm{do}_i(\texttt{expand}\ \varphi)\rangle\psi \leftrightarrow [\mathrm{do}_i(\texttt{expand}\ \varphi)]\psi$
- $\models \langle \mathrm{do}_i(\texttt{expand}\ \varphi; \texttt{expand}\ \varphi)\rangle\psi \leftrightarrow \langle \mathrm{do}_i(\texttt{expand}\ \varphi)\rangle\psi$

It turns out that expansions as formalized in Def. 3.5 can be completely characterized as follows.

3.8. PROPOSITION (CHARACTERIZATION OF EXPANSIONS). *For all agents i, for all propositional formulae φ and ψ we have:*

- $\models [\mathrm{do}_i(\texttt{expand}\ \varphi)]\mathbf{B}_i\psi \leftrightarrow \mathbf{B}_i(\varphi \rightarrow \psi)$

Proposition 3.8 states that some (propositional) formula ψ is believed after an expansion with φ if and only if the agent believes that φ implies ψ on beforehand. As a special case of Proposition 3.8 we can prove that an expansion with some formula results in the agent having absurd beliefs if and only if the agent believes the negation of the formula on beforehand.

3.9. COROLLARY. *For all agents i, for all propositional formulae φ and ψ we have:*

- $\models [\mathrm{do}_i(\texttt{expand}\ \varphi)]\mathbf{B}_i\mathbf{ff} \leftrightarrow \mathbf{B}_i\neg\varphi$

3.1.1. Expansions and the Gärdenfors postulates

As for contractions and revisions, Gärdenfors proposes some rationality postulates that describe belief expansions [5]. These postulates are given below, where K and H denote arbitrary belief sets, φ denotes some formula, and the expansion of K with φ is denoted by K^+_φ.

3.10. DEFINITION. The Gärdenfors postulates for belief expansion:
(G^+1) K^+_φ is a belief set.

p restricts the set of doxastic alternatives of the agent to those that support p. This is graphically depicted in Fig. 2.

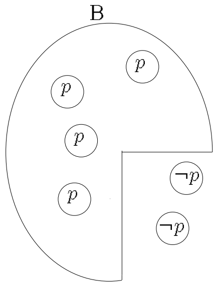

Figure 2: The set of doxastic alternatives after an expansion with p.

Note that in the resulting model it is indeed the case that the agent believes p, whereas it did not believe p on beforehand.

The definition of the $\mathtt{r}$ function for expansions is a direct formalization of the intuitive ideas given above: if some agent i performs an expansion with some formula φ in a world s in the model $\mathcal{M}$, the result of this will be that afterwards i has restricted its set of doxastic alternatives to those states that satisfy φ.

3.5. Definition. Let some model $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathtt{c} \rangle$ with $s \in \mathcal{S}$, and some agent i and propositional formula φ be given. We define:

$$
\begin{aligned}
&\mathtt{r}(i, \mathtt{expand}\ \varphi)(\mathcal{M}, s) = \mathcal{M}', s \text{ where}\\
&\quad \mathcal{M}' = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}', \mathtt{r}, \mathtt{c} \rangle \text{ with}\\
&\quad\ \mathrm{B}'(i', s') = \mathrm{B}(i', s') \text{ if } i' \neq i \text{ or } s' \notin [s]_{\mathrm{R}(i)}\\
&\quad\ \mathrm{B}'(i, s') = \mathrm{B}(i, s') \cap [\![\varphi]\!] \text{ if } s' \in [s]_{\mathrm{R}(i)}
\end{aligned}
$$

Definition 3.5 provides for an intuitively acceptable formalization of belief expansions as can be seen in the following proposition.

3.6. Proposition. *For all agents i, and for all propositional formulae φ and ψ we have:*

- $\models [\mathrm{do}_i(\mathtt{expand}\ \varphi)]\mathbf{B}_i\varphi$
- $\models \mathbf{B}_i\psi \to [\mathrm{do}_i(\mathtt{expand}\ \varphi)]\mathbf{B}_i\psi$
- $\models \mathbf{B}_i\varphi \to (\mathbf{B}_i\psi \leftrightarrow [\mathrm{do}_i(\mathtt{expand}\ \varphi)]\mathbf{B}_i\psi)$

The first clause of Proposition 3.6 states that an expansion with some formula results in the formula being believed. The second clause states that beliefs are persistent under expansions. In this clause the restriction to *propositional* formulae ψ is in general necessary. For consider a situation in which an agent does not believe φ and therefore believes that it does not believe φ (cf. clause 12 of Proposition 2.8). After expanding its beliefs with φ, the agent believes φ and, assuming that the resulting belief set is not the absurd one, it no longer believes that it does not believe φ. Hence not all beliefs of the agent persist

3.2. Proposition. *Let* $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathtt{c} \rangle$ *be some Kripke model with* $s \in \mathcal{S}$, *and let* $\mathcal{M}' = \langle \mathcal{S}', \pi', \mathrm{R}', \mathrm{B}', \mathtt{r}', \mathtt{c}' \rangle$ *be some Kripke model with* $s' \in \mathcal{S}'$. *Then it holds that:*

$$\forall p \in \Pi[\pi(p, s) = \pi'(p, s')] \Rightarrow \forall \psi \in \mathcal{L}_0[\mathcal{M}, s \models \psi \Leftrightarrow \mathcal{M}', s' \models \psi]$$

The result of Proposition 3.2 is used at various places in the proofs of the other propositions, theorems, and corollaries of this paper (see appendix A).

3.3. Convention. In the rest of this paper we follow the convention that whenever some model $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathtt{c} \rangle$ is clear from the context, $[\![\varphi]\!]$ denotes the set of states that satisfy φ, i.e., $[\![\varphi]\!] = \{s \in \mathcal{S} \mid \mathcal{M}, s \models \varphi\}$. The relation $\vdash_{cl} \subseteq \wp(\mathcal{L}_0) \times \mathcal{L}_0$ is the derivability relation of classical propositional logic. The function $\mathrm{Th} : \wp(\mathcal{L}_0) \to \wp(\mathcal{L}_0)$ that yields for every set Φ of propositional formula the set $\{\varphi \in \mathcal{L}_0 \mid \Phi \vdash_{cl} \varphi\}$ is the deductive closure operator associated with the derivability relation $\vdash_{cl}$. Note that both $\vdash_{cl}$ and Th work strictly on a propositional level; it is for example neither the case that $\{\mathbf{B}_i p\} \vdash_{cl} p \vee \neg p$ nor is it the case that $\{p\} \vdash_{cl} \mathbf{B}_i p \vee \neg \mathbf{B}_i p$.

3.1. The `expand` action

Informally, a belief expansion is an action that leads to a state of affairs in which some formula is believed. In our framework uncertainties of agents are formalized through the different doxastic alternatives that the agent has: if an agent does neither believe φ nor $\neg\varphi$ then it considers both doxastic alternatives supporting φ and doxastic alternatives supporting $\neg\varphi$ possible. Expanding the beliefs of the agent with φ then comes down to declaring all alternatives supporting $\neg\varphi$ to be 'doxastically impossible', i.e., on the ground of its beliefs the agent no longer considers these alternatives to be possible. Hence the *expansion* of the belief set of an agent is modelled through a *restriction* of its set of doxastic alternatives. The following example makes this point more clear.

3.4. Example. Consider a Kripke model $\mathcal{M}$ of which the relevant part is given in Fig. 1.

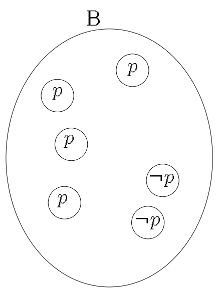

Figure 1: The set of doxastic alternatives before an expansion with p.

In the situation sketched in Fig. 1, the agent considers both worlds supporting p and worlds supporting $\neg p$ possible. An expansion with the formula

3. $\models \mathbf{K}_i\mathbf{K}_i\varphi \leftrightarrow \mathbf{K}_i\varphi$
4. $\models \mathbf{K}_i\mathbf{B}_i\varphi \leftrightarrow \mathbf{B}_i\varphi$
5. $\models \mathbf{K}_i\neg\mathbf{K}_i\varphi \leftrightarrow \neg\mathbf{K}_i\varphi$
6. $\models \mathbf{K}_i\neg\mathbf{B}_i\varphi \leftrightarrow \neg\mathbf{B}_i\varphi$
7. $\models \mathbf{K}_i\varphi \rightarrow \mathbf{B}_i\mathbf{K}_i\varphi$
8. $\models \neg\mathbf{B}_i\mathbf{ff} \rightarrow (\mathbf{B}_i\mathbf{K}_i\varphi \rightarrow \mathbf{K}_i\varphi)$
9. $\models \mathbf{B}_i\mathbf{B}_i\varphi \leftrightarrow \mathbf{B}_i\varphi$
10. $\models \neg\mathbf{K}_i\varphi \rightarrow \mathbf{B}_i\neg\mathbf{K}_i\varphi$
11. $\models \neg\mathbf{B}_i\mathbf{ff} \rightarrow (\mathbf{B}_i\neg\mathbf{K}_i\varphi \rightarrow \neg\mathbf{K}_i\varphi)$
12. $\models \neg\mathbf{B}_i\varphi \rightarrow \mathbf{B}_i\neg\mathbf{B}_i\varphi$
13. $\models \neg\mathbf{B}_i\mathbf{ff} \rightarrow (\mathbf{B}_i\neg\mathbf{B}_i\varphi \rightarrow \neg\mathbf{B}_i\varphi)$

3. Actions that change one's mind

As explained in Sect. 1, our approach towards belief changes in the agent-oriented, semantics based framework of Sect. 2 is based on the idea that belief changes are brought about by actions that the agents may perform. By performing belief-changing actions, agents expand, contract and revise their beliefs. We think of these belief-changing actions as working within the boundary set by the knowledge of the agents, i.e., knowledge is fixed within a given state of the model and does not change as the result of belief-changing actions. In this way knowledge can be seen as representing the beliefs that the agent will stick to against all odds.

From a syntactical point of view, the class of actions Ac is extended with three new, belief-changing, actions.

3.1. Definition. The class Ac of actions (and hence the language $\mathcal{L}$) as defined in 2.1 is extended as follows:

if $\varphi \in \mathcal{L}_0$ then $\texttt{expand}\ \varphi, \texttt{contract}\ \varphi, \texttt{revise}\ \varphi \in Ac$

The main reason underlying the restriction to propositional formulae in Def. 3.1 is the fact that changes of belief concerning non-propositional, and in particular doxastic formulae, are not very well understood. It is not at all clear what it means to revise the beliefs of some agent i with the formula $p \wedge \neg\mathbf{B}_i p$: does or doesn't the agent believe $p \wedge \neg\mathbf{B}_i p$ after its beliefs are revised with this formula?[2] On the other hand, changes of belief with propositional formulae are not only well understood and thoroughly investigated [5], but propositional formulae do also have some properties that are very useful in our semantic framework. The following proposition formalizes the most useful property of propositional formulae: the truth value of propositional formulae in a state of a model depends on the valuation for that state only.

[2] The natural language variant 'p, and i does not believe p' of the formula $p \wedge \neg\mathbf{B}_i p$ is considered by Thijsse ([34], pp. 131-132) to be a typical example of a non-contradictory sentence but a contradictory utterance. This implies that although the sentence in itself is consistent, it is not consistent to believe the sentence.

2.5. REMARK. With regard to the abilities of agents, the motivation for the choices made in Def. 2.4 is the following. The definition of $\mathtt{c}(i, \mathtt{confirm}\ \varphi)(s)$ expresses that an agent is able to get confirmation for a formula φ iff φ holds. Note that the definitions of $\mathtt{r}(i, \mathtt{confirm}\ \varphi)$ and $\mathtt{c}(i, \mathtt{confirm}\ \varphi)$ imply that in circumstances such that φ holds, the agents both have the opportunity and the ability to confirm φ. An agent is capable of performing a sequential composition $\alpha_1; \alpha_2$ iff it is capable of performing α_1 and it is capable of executing α_2 after it has performed α_1. An agent is capable of performing a conditional composition, if either it is able to get confirmation for the condition and thereafter perform the then-part, or it is able to confirm the negation of the condition and perform the else-part afterwards. An agent is capable of performing a repetitive composition $\mathtt{while}\ \varphi\ \mathtt{do}\ \alpha_1\ \mathtt{od}$ iff it is able to perform the action $(\mathtt{confirm}\ \varphi; \alpha_1)^k; \mathtt{confirm}\ \neg\varphi$ for some $k \in \mathbb{N}$.

When defining the R and B functions as in Def. 2.3, we end up with a notion of knowledge that satisfies an **S5** axiomatization, and a notion of belief that satisfies a **K45** axiomatization. The main difference between this approach towards knowledge and belief and the usual ones [10, 18, 23] is that our notion of belief does not satisfy the D-axiom $\neg(\mathbf{B}_i\varphi \wedge \mathbf{B}_i\neg\varphi)$. The reason for this is that the approach towards belief expansions [5] that we use in defining our expansion action presupposes the existence of inconsistent belief sets: expansions may result in the agent having inconsistent, or absurd, beliefs.

2.6. PROPOSITION (**S5** VALIDITIES FOR KNOWLEDGE). *For all agents i and formulae φ and ψ we have:*

- $\models \mathbf{K}_i(\varphi \to \psi) \to (\mathbf{K}_i\varphi \to \mathbf{K}_i\psi)$
- $\models \mathbf{K}_i\varphi \to \varphi$
- $\models \mathbf{K}_i\varphi \to \mathbf{K}_i\mathbf{K}_i\varphi$
- $\models \neg\mathbf{K}_i\varphi \to \mathbf{K}_i\neg\mathbf{K}_i\varphi$
- $\models \varphi \Rightarrow \models \mathbf{K}_i\varphi$

2.7. PROPOSITION (**K45** VALIDITIES FOR BELIEF). *For all agents i and formulae φ and ψ we have:*

- $\models \mathbf{B}_i(\varphi \to \psi) \to (\mathbf{B}_i\varphi \to \mathbf{B}_i\psi)$
- $\models \mathbf{B}_i\varphi \to \mathbf{B}_i\mathbf{B}_i\varphi$
- $\models \neg\mathbf{B}_i\varphi \to \mathbf{B}_i\neg\mathbf{B}_i\varphi$
- $\models \varphi \Rightarrow \models \mathbf{B}_i\varphi$

In proving validities concerning the relation between knowledge and belief, it has to be taken into account that in our approach agents may have inconsistent beliefs (this proviso is visible in the second, eight, eleventh and last clause of Proposition 2.8). Apart from this fact, knowledge and belief are related to each other as in the system of Kraus & Lehmann [10, 18].

2.8. PROPOSITION. *For all agents i and formulae φ we have:*

1. $\models \mathbf{K}_i\varphi \to \mathbf{B}_i\varphi$
2. $\models \neg\mathbf{B}_i\mathbf{ff} \to (\mathbf{K}_i\neg\varphi \to \neg\mathbf{B}_i\varphi)$

4. $\mathrm{B} : \mathcal{A} \times \mathcal{S} \to \wp(\mathcal{S})$ is a function that yields the set of doxastic alternatives for a given agent in a given state. To model the kind of belief that we like to model it is demanded that for all agents i, i' and for all possible worlds s and s' it holds that:
   - $\mathrm{B}(i,s) = \mathrm{B}(i,s')$ if $s' \in [s]_{\mathrm{R}(i)}$
   - $\mathrm{B}(i,s) \subseteq [s]_{\mathrm{R}(i)}$
5. $\mathtt{r} : \mathcal{A} \times At \to \mathcal{S} \to \wp(\mathcal{S})$ is such that $\mathtt{r}(i,a)(s)$ yields the (possibly empty) state transition in s caused by the event $\mathrm{do}_i(a)$. This function is such that for all atomic actions a it holds that $|\mathtt{r}(i,a)(s)| \leq 1$ for all i and s, i.e., these events are *deterministic*.
6. $\mathsf{c} : \mathcal{A} \times At \to \mathcal{S} \to \mathbf{bool}$ is the capability function such that $\mathsf{c}(i,a)(s)$ indicates whether the agent i is capable of performing the action a in s.

2.4. Definition. Let $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathsf{c} \rangle$ be some Kripke model from M. For propositional symbols, negated formulae, and conjunctions, $\mathcal{M}, s \models \varphi$ is inductively defined as usual. For the other clauses $\mathcal{M}, s \models \varphi$ is defined as follows:

$$\begin{array}{ll}
\mathcal{M}, s \models \mathbf{K}_i\varphi & \Leftrightarrow \forall s' \in \mathcal{S}[(s,s') \in \mathrm{R}(i) \Rightarrow \mathcal{M}, s' \models \varphi] \\
\mathcal{M}, s \models \mathbf{B}_i\varphi & \Leftrightarrow \forall s' \in \mathcal{S}[s' \in \mathrm{B}(i,s) \Rightarrow \mathcal{M}, s' \models \varphi] \\
\mathcal{M}, s \models \langle \mathrm{do}_i(\alpha) \rangle \varphi & \Leftrightarrow \exists \mathcal{M}', s'[\mathcal{M}', s' \in \mathtt{r}(i,\alpha)(\mathcal{M}, s) \,\&\, \mathcal{M}', s' \models \varphi] \\
\mathcal{M}, s \models \mathbf{A}_i\alpha & \Leftrightarrow \mathsf{c}(i,\alpha)(\mathcal{M}, s) = \mathbf{1}
\end{array}$$

where $\mathtt{r}$ and c are defined by:

$$\begin{array}{ll}
\mathtt{r} & : \ \mathcal{A} \times Ac \to (\mathsf{M} \times \mathcal{S}) \cup \mathcal{S} \to \wp(\mathsf{M} \times \mathcal{S}) \\
\mathtt{r}(i,a)(\mathcal{M}, s) & = \ \mathcal{M}, \mathtt{r}(i,a)(s) \\
\mathtt{r}(i, \mathtt{confirm}\,\varphi)(\mathcal{M}, s) & = \ \{(\mathcal{M}, s)\} \text{ if } \mathcal{M}, s \models \varphi \text{ and } \emptyset \text{ otherwise} \\
\mathtt{r}(i, \alpha_1; \alpha_2)(\mathcal{M}, s) & = \ \mathtt{r}(i,\alpha_2)(\mathtt{r}(i,\alpha_1)(\mathcal{M}, s)) \\
\mathtt{r}(i, \mathtt{if}\ \varphi\ \mathtt{then}\ \alpha_1 & = \ \mathtt{r}(i,\alpha_1)(\mathcal{M}, s) \text{ if } \mathcal{M}, s \models \varphi \text{ and} \\
\quad \mathtt{else}\ \alpha_2\ \mathtt{fi})(\mathcal{M}, s) & \quad \mathtt{r}(i,\alpha_2)(\mathcal{M}, s) \text{ otherwise} \\
\mathtt{r}(i, \mathtt{while}\ \varphi\ \mathtt{do}\ \alpha_1\ \mathtt{od})(\mathcal{M}, s) & = \ \{(\mathcal{M}', s') \mid \exists k \in \mathbb{N} \exists \mathcal{M}_0, s_0 \ldots \exists \mathcal{M}_k, s_k \\
 & \quad [\mathcal{M}_0, s_0 = \mathcal{M}, s \,\&\, \mathcal{M}_k, s_k = \mathcal{M}', s' \,\&\, \forall j < k \\
 & \quad [\mathcal{M}_{j+1}, s_{j+1} = \mathtt{r}(i, \mathtt{confirm}\,\varphi; \alpha_1)(\mathcal{M}_j, s_j)] \\
 & \quad \& \, \mathcal{M}', s' \models \neg\varphi]\} \\
\text{where } \mathtt{r}(i,\alpha)(\emptyset) & = \ \emptyset
\end{array}$$

and

$$\begin{array}{ll}
\mathsf{c} & : \ \mathcal{A} \times Ac \to (\mathsf{M} \times \mathcal{S}) \cup \mathcal{S} \to \mathbf{bool} \\
\mathsf{c}(i,a)(\mathcal{M}, s) & = \ \mathsf{c}(i,a)(s) \\
\mathsf{c}(i, \mathtt{confirm}\,\varphi)(\mathcal{M}, s) & = \ \mathbf{1} \text{ if } \mathcal{M}, s \models \varphi \text{ and } \mathbf{0} \text{ otherwise} \\
\mathsf{c}(i, \alpha_1; \alpha_2)(\mathcal{M}, s) & = \ \mathsf{c}(i,\alpha_1)(\mathcal{M}, s) \,\&\, \mathsf{c}(i,\alpha_2)(\mathtt{r}(i,\alpha_1)(\mathcal{M}, s)) \\
\mathsf{c}(i, \mathtt{if}\ \varphi\ \mathtt{then}\ \alpha_1 & = \ \mathsf{c}(i, \mathtt{confirm}\,\varphi; \alpha_1)(\mathcal{M}, s) \text{ or} \\
\quad \mathtt{else}\ \alpha_2\ \mathtt{fi})(\mathcal{M}, s) & \quad \mathsf{c}(i, \mathtt{confirm}\,\neg\varphi; \alpha_2)(\mathcal{M}, s) \\
\mathsf{c}(i, \mathtt{while}\ \varphi\ \mathtt{do}\ \alpha_1\ \mathtt{od})(\mathcal{M}, s) & = \ \mathbf{1} \text{ if } \exists k \in \mathbb{N}[\mathsf{c}(i, (\mathtt{confirm}\,\varphi; \alpha_1)^k; \\
 & \quad \mathtt{confirm}\,\neg\varphi)(\mathcal{M}, s) = \mathbf{1}] \\
 & \quad \text{and } \mathbf{0} \text{ otherwise} \\
\text{where } \mathsf{c}(i,\alpha)(\emptyset) & = \ \mathbf{1}.
\end{array}$$

Satisfiability and validity are defined as usual.

of affairs satisfying φ would result. Besides the possibility to formalize both opportunities and results when using dynamic logic, another advantage lies in the compatibility of epistemic and dynamic logic from a semantic point of view: the possible world semantics can be used to provide meaning both to epistemic and dynamic notions.

2.1. The formal definitions

2.1. Definition. Let a finite set $\mathcal{A} = \{1, \ldots, n\}$ of agents, and some denumerable sets Π of propositional symbols and At of atomic actions be given. The language $\mathcal{L}$ and the class of actions Ac are defined by mutual induction as follows.

1. $\mathcal{L}$ is the smallest superset of Π such that
   - if $\varphi, \psi \in \mathcal{L}$ then $\neg\varphi, \varphi \vee \psi \in \mathcal{L}$,
   - if $i \in \mathcal{A}$, $\alpha \in Ac$ and $\varphi \in \mathcal{L}$ then $\mathbf{K}_i\varphi, \mathbf{B}_i\varphi, \langle \mathrm{do}_i(\alpha)\rangle\varphi, \mathbf{A}_i\alpha \in \mathcal{L}$.
2. Ac is the smallest superset of At such that
   - if $\varphi \in \mathcal{L}$ then $\texttt{confirm}\ \varphi \in Ac$,
   - if $\alpha_1 \in Ac$ and $\alpha_2 \in Ac$ then $\alpha_1; \alpha_2 \in Ac$,
   - if $\varphi \in \mathcal{L}$ and $\alpha_1, \alpha_2 \in Ac$ then $\texttt{if}\ \varphi\ \texttt{then}\ \alpha_1\ \texttt{else}\ \alpha_2\ \texttt{fi} \in Ac$,
   - if $\varphi \in \mathcal{L}$ and $\alpha_1 \in Ac$ then $\texttt{while}\ \varphi\ \texttt{do}\ \alpha_1\ \texttt{od} \in Ac$.

The purely propositional fragment of $\mathcal{L}$ is denoted by $\mathcal{L}_0$.

The constructs $\wedge, \rightarrow, \leftrightarrow, \mathbf{tt}, \mathbf{ff}, \mathbf{M}_i\varphi$ and $[\mathrm{do}_i(\alpha)]\varphi$ are defined in the usual way. Other additional constructs are introduced by definitional abbreviation: `skip` is `confirm` $\mathbf{tt}$, `fail` is `confirm` $\mathbf{ff}$, α^0 is `skip`, and α^{n+1} is $\alpha; \alpha^n$.

2.2. Remark. The `confirm` action behaves essentially like the test actions in dynamic logic [6, 8]. As such this action differs substantially from tests as they are looked upon by humans: these genuine tests are usually assumed to contribute to the knowledge (or belief) of the agent that performs the test [24], whereas by performing `confirm` φ it is just verified (checked, confirmed) that φ holds. The meaning of the other actions in Ac is respectively: the atomic action, sequential composition, conditional composition, and repetitive composition. Furthermore, `skip` denotes the empty action, and `fail` denotes the never succeeding action.

In the following definitions it is assumed that some set $\mathbf{bool} = \{\mathbf{0}, \mathbf{1}\}$ of truth values is given.

2.3. Definition. The class M of Kripke models contains all tuples $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathtt{c}\rangle$ such that

1. $\mathcal{S}$ is a set of possible worlds, or states.
2. $\pi : \Pi \times \mathcal{S} \rightarrow \mathbf{bool}$ is a total function that assigns a truth value to propositional symbols in possible worlds.
3. $\mathrm{R} : \mathcal{A} \rightarrow \wp(\mathcal{S} \times \mathcal{S})$ is a function that yields the epistemic accessibility relations for a given agent. Since we assume to deal with **S5** models, it is demanded that $\mathrm{R}(i)$ is an equivalence relation for all i. For reasons of practical convenience we define $[s]_{\mathrm{R}(i)}$ to be $\{s' \in \mathcal{S} \mid (s, s') \in \mathrm{R}(i)\}$.

abilities, opportunities, and results; furthermore the formal definitions of our framework are given. We also elaborate on the notion of informative actions. In Sect. 3 we introduce actions that model expansions, revisions, and contractions. We define the states of affairs following these actions, conditions that need to be fulfilled for agents to have the opportunity to perform these actions, and (mental) capacities that agents should have to be capable of performing these actions. Furthermore we show that these actions satisfy (slightly adapted) versions of the Gärdenfors postulates. In Sect. 4 we summarize and discuss options for further research. The appendix contains the proofs of all theorems, propositions and corollaries given in this paper.

2. Knowledge, belief, abilities, opportunities, and results

As already stated in Sect. 1, for the moment we restrict ourselves to informational and action aspects of rational agents. At the informational level we consider both *knowledge and belief*. Formalizing these notions has been a subject of continuing research both in analytical philosophy and in AI [7, 9, 26]. In representing knowledge and belief we follow, both from a syntactical and a semantic point of view, the approach common in epistemic and doxastic logic: the formula $\mathbf{K}_i\varphi$ denotes the fact that agent i knows φ, and the formula $\mathbf{B}_i\varphi$ denotes the fact that agent i beliefs φ. For the semantics we use Kripke-style possible worlds models.

At the action level we consider *results*, *abilities* and *opportunities*. In defining the result of an action, we follow ideas of Von Wright [39], in which the state of affairs brought about by execution of the action is defined to be its result. An important aspect of any investigation of action is the relation that exists between ability and opportunity. In order to successfully complete an action, both the opportunity and the ability to perform the action are necessary. Although these notions are interconnected, they are surely not identical [17]: the abilities of agents can be seen as comprising mental and physical powers, moral capacities, and human and physical possibility, whereas the opportunity to perform actions is best described by the notion of circumstantial possibility. A nice example that illustrates the difference between ability and opportunity is that of a lion in a zoo [3]: although the lion will (ideally) never have the opportunity to eat a zebra, it certainly has the ability to do so. We propose that in order to make our formalization of rational agents, like for instance robots, as accurate and realistic as possible, abilities and opportunities need also be distinguished in AI environments. The abilities of the agents are formalized via the $\mathbf{A}_i$ operator; the formula $\mathbf{A}_i\alpha$ denotes the fact that agent i has the ability to do α. When using the definitions of opportunities and results as given above, the framework of (propositional) dynamic logic provides an excellent means to formalize these notions. Using events $\mathrm{do}_i(\alpha)$ to refer to the performance of the action α by the agent i, we consider the formulae $\langle \mathrm{do}_i(\alpha)\rangle\varphi$ and $[\mathrm{do}_i(\alpha)]\varphi$. In our deterministic framework, $\langle \mathrm{do}_i(\alpha)\rangle\varphi$ is the stronger of these formulae; it represents the fact that the agent i has the opportunity to do α and that doing α leads to φ. The formula $[\mathrm{do}_i(\alpha)]\varphi$ is noncommittal about the opportunity of the agent to do α but states that should the opportunity arise, only states aspects like knowledge and belief (for a survey see [4, 7, 25]), and *motivational*[1] aspects like commitments and obligations [2]. Recent developments include the work on agent-oriented programming [31, 35], the BDI-architecture for formalizing rational agents [28, 29, 30], logics for the specification and verification of multi-agent systems [36, 37], logics for agents with bounded rationality [14, 15], and cognitive robotics [19, 20] (for a survey on recent developments see [13, 38]).

In our research [11, 12, 22, 23, 24] we defined a *theorist* logic for rational agents, i.e., a logic that is used to *specify*, and to *reason about*, (various aspects of) the behaviour of rational agents. In our framework we concentrate on informational and action aspects, leaving motivational aspects (for the moment) out of consideration. In the basic architecture the *knowledge* and *abilities* of agents, as well as the *opportunities* for and the *results* of their actions are formalized. In this framework it can for instance be modelled that an agent knows that some action is a *correct* plan to achieve some goal since it knows that performing the action will lead to the goal, and that it knows that an action is a *feasible* plan since the agent knows of its ability to perform the action. In subsequent research we extended our framework with nondeterministic actions [12], epistemic tests [24], communicative actions [22] and actions that model default reasoning [23].

The aim of this paper is a formalization of belief-changing actions in the framework mentioned above. We consider three kinds of belief-changing actions: *expansions*, *contractions* and *revisions*. Informally speaking, expansions result in some formula being believed, contractions result in some formula no longer being believed, and revisions reverse the belief status of a formula, i.e., for some formula that is believed on beforehand the negation is believed afterwards. As for any action in our framework, we define the states of affairs resulting from execution of belief-changing actions, conditions that need to be satisfied in order for agents to have the opportunity to perform these actions, and capacities that the agents must posses in order to be capable of performing these actions. The resulting definitions provide for an intuitively acceptable formalization of expansions, contractions and revisions that furthermore satisfy the Gärdenfors postulates for belief changes [5]. Although already intrinsically interesting from a philosophical point of view, belief-changing actions are particularly important when formalizing rational agents that acquire information from multiple sources. For whenever some source provides reliable information that contradicts the information that an agent already has, the agent has to change its beliefs if it wants to incorporate this new information whilst keeping its set of beliefs consistent, and this is where belief-changing actions play a part. The formalization of this kind of multiple-source information acquisition is subject of further research.

1.1. Organization of the paper

The rest of the paper is organized as follows.

To sketch the context and the area of application of this research, we start in Sect. 2 with the (re)introduction of some of our ideas on knowledge, belief,

[1] The terms *informational* and *motivational* are both due to Shoham [32].

Actions that Make you Change your Mind
Belief Revision in an Agent-Oriented Setting

B. van Linder W. van der Hoek J.-J. Ch. Meyer
Utrecht University
Department of Computer Science
P.O. Box 80.089
3508 TB Utrecht
The Netherlands
Email: bernd@cs.ruu.nl

Abstract

In this paper we study the dynamics of belief from an agent-oriented, semantics-based point of view. In a formal framework used to specify and to reason about formal agents, we define actions that model three well-known changes of belief, viz. *expansions*, *contractions*, and *revisions*. We treat these belief changes as full-fledged actions by defining both the *opportunity for* and *result of* these actions, and the *ability* of agents to apply these belief-changing actions. In defining the result of the contraction action we introduce the concept of selection functions. These are special functions that select a subset of the set of states that is to be added to the set of doxastic alternatives of an agent, thereby contracting its set of beliefs. The action that models belief revisions is defined as the sequential composition of a contraction and an expansion. We show that these belief-changing actions are defined in an intuitively acceptable, reasonable way by proving that the Gärdenfors postulates for belief changes are validated. The ability of agents to apply belief-changing actions is defined in terms of their knowledge and belief. These definitions are such that actions that an agent is capable of performing lead to desired states of affairs. The resulting framework provides an intuitively acceptable yet simple formalization of expansions, contractions and revisions as actions in a dynamic/epistemic, agent-oriented framework.

Content Areas: Belief Revision, Knowledge Representation, Reasoning about Action.

1. Introduction

The formalization of rational agents is a topic of continuing interest in Artificial Intelligence. Research on this subject has held the limelight ever since the pioneering work of Moore [26, 27] in which knowledge and actions are considered. Over the years important contributions have been made on both *informational*