COSC410 Lecture 5
Belief change

Willem Labuschagne
University of Otago

2016

Introduction

We previously developed the following **Traffic System scenario**.

- The agent is a driver waiting to cross an intersection.
- The driver's representation language is $L_{\{p,q\}}$ where p stands for "The light is red" and q for "The oncoming car stops". The states of the Traffic System are given by $S = W_{\{p,q\}} = \{11, 10, 01, 00\}$.
- Assume that the waiting driver sees the light for cross-traffic is red.
- Semantically, the agent's observation provides enough information to reduce the possibilities to two states, 11 and 10, one of which must be the actual state of the system.
- This information can be expressed syntactically by the sentence p, since $Mod(p) = \{11, 10\}$.
- The waiting driver's difficulty is that she doesn't know whether the oncoming car is going to stop, because in one of the possible states, 11, it does and in the other, 10, it doesn't.
- The waiting driver has heuristic information that may be represented with the aid of a total preorder $\preccurlyeq$ given by

11	00
10	01

- $p \mid\!\sim q$ since the only maximal model of p is 11 and 11 satisfies q.
- Assume that the driver reasons according to the constraints embodied by $\mid\!\sim$ and therefore infers that the oncoming car is stopping. On the basis of this new belief q, she will be able to make a rational decision to drive across the intersection (rational, as opposed to merely hoping that everything will somehow be all right).
- The inference that delivers the conclusion q is defeasible, i.e. may turn out to be mistaken, because the driver does not know for sure that the actual state of the system is 11; the driver merely knows that the system would normally be in state 11, and has no reason to suspect that the system's current state is abnormal.

This brings us to the question: How should the driver change her beliefs about the world if the inference turns out to be mistaken? We shall consider two cases.

Case 1 *The driver hears police sirens that may be pursuing the oncoming car, and therefore becomes doubtful that the car will stop, so that she wishes to drop the belief q (while retaining as many of her other beliefs as possible).*

Case 2 *The driver actually sees the oncoming car speed up in order to run the red light, so that she wishes to add the new information $\neg q$ to her belief set.*

Belief sets

We are going to analyse this scenario from the perspective of the agent's beliefs. Beliefs are sentences expressing information possessed by the agent.

Recall that the driver formed the belief p by observing that the light for oncoming traffic had turned red, and then formed the belief q by making the defeasible inference $p \mid\!\sim q$. However, the driver's resulting belief set does not consist only of the two sentences p and q. Her beliefs must include all the bits of information carried by p and q. What does this mean?

If the agent believes p and believes q, then the agent rules out all states of the system in which p is false and all states in which q is false, i.e. the agent rules out $\{10, 01, 00\}$. This information can be broken down into various pieces corresponding to subsets of $\{10, 01, 00\}$. For instance, the information that rules out the state 00 is one small piece of the information that rules out all three states 00 and 01 and 10. There are 8 subsets that are ruled out:

$$\begin{array}{l}\{10, 01, 00\}\\ \{10, 01\}\\ \{10, 00\}\\ \{01, 00\}\\ \{10\}\\ \{01\}\\ \{00\}\\ \{\}.\end{array}$$

The agent's beliefs are the sentences that express each of these pieces of information. If we ignore equivalent sentences and pick just one sentence for each piece of information, then what we're looking for are sentences whose *nonmodels* match each of the subsets that is ruled out. (Alternatively, we could focus on the subsets of states that have been left as candidates for the actual state, i.e. those that are not ruled out. In that case we seek sentences having the subsets of candidate states as sets of models. We get the same sentences as beliefs either way.)

Thus the agent's belief set at this point in the scenario must be the set of sentences

$$\{p \wedge q,\ p \leftrightarrow q,\ q,\ p,\ \neg p \vee q,\ p \vee \neg q,\ p \vee q,\ p \vee \neg p\}$$

where each sentence has one of the above subsets as its set of nonmodels and the sentences are listed in the same order as the subsets. Strictly speaking, all the infinitely many sentences equivalent to $p \wedge q$ are also in the belief set, as are the sentences equivalent to $p \leftrightarrow q$, and those equivalent to q, and so forth. In describing the belief set, however, we list just one randomly selected sentence from each equivalence class, because what's important is the information being conveyed, not the syntactic shape of the sentence that carries the information.

By including all the pieces of information we ensure that the set of beliefs is closed under classical entailment $\models$, in other words that if φ is in the belief set and if $\varphi \models \psi$ then ψ is also in the belief set (or at least a sentence equivalent to ψ is in the belief set). After all, if $\varphi \models \psi$ then ψ rules out a subset of the states that φ rules out, so ψ conveys a part of the information conveyed by φ. We will take closure under classical entailment to be the characteristic property of belief sets, after we generalise it slightly to allow for a set of premisses instead of just one premiss.

Definition 1 *(Sets of premisses) Suppose $X \subseteq L_A$.*

- *We say that $X \models \beta$ iff $Mod(X) \subseteq Mod(\beta)$, i.e. iff it is the case that whenever x satisfies every sentence in X then x satisfies β.*
- *The set of consequences of X is $Cn(X) = \{\beta \mid X \models \beta\}$.*

For example in $L_{\{p,q\}}$, if $X = \{p, q\}$ then

$$Cn(\{p, q\}) = \{p \wedge q, p \leftrightarrow q, q, p, \neg p \vee q, p \vee \neg q, p \vee q, p \vee \neg p\}$$

where once again we do not explicitly list sentences equivalent to those mentioned.

Notation 1 *If X has only a single member, e.g. $X = \{\varphi\}$, then we write $Cn(\varphi)$ instead of $Cn(\{\varphi\})$ for ease of reading.*

Definition 2 *(Belief sets) A set K of sentences is a belief set iff $K = Cn(K)$.*

The idea is that a belief set K is closed under $\models$ in the sense that whenever β is a sentence such that $K \models \beta$, then β must already be a member of K.

So, in the **Traffic System scenario** we're considering, the agent begins by observing that the traffic light is red for the oncoming cross traffic, resulting in the belief set

$$K_0 = Cn(p) = \{p, p \vee q, p \vee \neg q, p \vee \neg p\}.$$

Notice that this set is closed under $\models$ as may be seen by inspecting the Lindenbaum-Tarski algebra you are diagramming for your assignment.

Next, through the mysterious process of everyday commonsense reasoning that we formalised using nonmonotonic logic, the agent adds belief q and ends up with the belief set

$$K_1 = Cn(p \wedge q) = \{p \wedge q,\ p \leftrightarrow q,\ q,\ p,\ \neg p \vee q,\ p \vee \neg q,\ p \vee q,\ p \vee \neg p\}.$$

K_1 is also closed under $\models$, as we'd expect .

How was K_1 built out of K_0? The process by which q was added to K_0 to give K_1 is called *expansion*, and is the first belief change operation we shall investigate below.

With K_1 we have now reached the point in the Traffic scenario where we have two different cases to consider, as described previously. Each of them causes the driver to doubt the belief q. In the first case, the doubt is just sufficient to make the driver want to retract the belief in q but not sufficient to convince the driver that $\neg q$ should be added in its place. In the second case, the driver actually wants to replace the belief q by the belief $\neg q$.

The two cases demand two different ways to change the agent's belief set, called *contraction* and *revision*. Contraction will be a way to remove an earlier belief, while revision will be a way to add a new belief. Neither operation turns out to be quite as simple as it might seem at first glance, and it is important to understand the relationship between them and *expansion*.

Expansion, Contraction, and Revision

Suppose K is a belief set and φ represents new information. The obvious way to incorporate the new information is to keep all the old beliefs in K, add to them the new belief φ, and then add in all the new consequences that now follow, so that the resulting set is closed under $\models$.

Remark 1 *Recall how the union of sets is defined as a way to add the contents of one set to another:* $X \cup Y = \{w \mid w \in X \text{ or } w \in Y\}$

Definition 3 *(Expansion)* $K + \varphi = Cn(K \cup \{\varphi\})$.

Let's return to the **Traffic System scenario** we're considering, and see what role expansion can play. First, the waiting driver has the belief set

$$\begin{aligned} K_0 &= Cn(p) \\ &= \{p, p \vee q, p \vee \neg q, p \vee \neg p\}. \end{aligned}$$

Next, the driver somehow acquires the new belief q. Perhaps she does some nonmonotonic reasoning and infers q, or perhaps her passenger, who is a more experienced driver, confidently assures her that q. The driver now expands her belief set to

$$\begin{aligned} K_0 + q &= Cn(Cn(p) \cup \{q\}) \\ &= Cn(\{p, p \vee q, p \vee \neg q, p \vee \neg p, q\}) \\ &= Cn(\{p, q\}) = Cn(p \wedge q) \\ &= \{p \wedge q, p \leftrightarrow q, q, p, \neg p \vee q, p \vee \neg q, p \vee q, p \vee \neg p\} \\ &= K_1. \end{aligned}$$

Thus expansion successfully produces the belief set K_1 from K_0. So far, so good!

Now let us consider each of the cases in which the driver is led to doubt q.

The first case (in which the driver heard sirens and suddenly doubted the belief q) involves removing information, not adding information, so expansion obviously can't help. We shall presently define a new kind of operation, called contraction, to cover this first case.

The second case does involve adding new information. We may imagine that our driver, as she prepares to cross the intersection, is shocked to see the oncoming car speed up and run the red light. With her own eyes, our pale and trembling driver has now seen that $\neg q$ is the case. So she wants to replace her belief q by $\neg q$. Should she expand her belief set by the new belief $\neg q$?

No, expanding with $\neg q$ would be a very bad idea. Her current belief set is $K_1 = K_0 + q = Cn(\{p \wedge q\}) = \{p \wedge q, p \leftrightarrow q, q, p, \neg p \vee q, p \vee \neg q, p \vee q, p \vee \neg p\}$. What is $K_1 + \neg q$? The moment we add $\neg q$ to K_1, the agent's beliefs become unsatisfiable, since the new belief set contains both q and $\neg q$. In other words $K_1 + \neg q = L_{\{p,q\}}$ because $\models$ is explosive, i.e. $K_1 + \neg q \models \beta$ for every $\beta \in L_{\{p,q\}}$. An unsatisfiable belief set means the agent believes everything. This is disaster. An agent who believes everything has no idea what is going on, i.e. is hopelessly deranged.

The example illustrates that if new information φ is inconsistent with a belief set K, then in order to incorporate φ some of the old information in K must be given up. Specifically, information contradicting φ must be sacrificed.

By a contradiction we understand a sentence of the form $\varphi \wedge \neg\varphi$, so the idea is that if φ is inconsistent with K then the information expressed by $\neg\varphi$ must be lurking somewhere in K, either in explicit form as the sentence $\neg\varphi$ or expressed in disguised form by other sentences. Giving up the information in $\neg\varphi$ usually means we have to remove from K a bunch of sentences that don't look like $\neg\varphi$; one can't just find the sentence $\neg\varphi$ and delete it.

To see that one can't simply delete a sentence, suppose $K = Cn(p) = \{p, p \vee q, p \vee \neg q, p \vee \neg p\}$ and we want to revise K so as to incorporate some new information $\varphi = \neg(p \vee \neg q)$. Clearly K does contain a sentence equivalent to $\neg\varphi$, namely the sentence $p \vee \neg q$. But if we simply to take out the old $p \vee \neg q$ and put in the new $\neg(p \vee \neg q)$, the resulting belief set still contains p, and since $p \vee \neg q$ is a classical consequence of p, and belief sets are closed under $\models$, the dratted $p \vee \neg q$ automatically slips back in, giving an unsatisfiable set $\{p, p \vee q, p \vee \neg q, \neg(p \vee \neg q), p \vee \neg p\}$. Closure has undone the removal we attempted to carry out. To really give up the information $p \vee \neg q$, we would need to remove p from K as well.

Thus it is more complicated to remove information from a belief set K than simply to delete a sentence from a list. It is even conceivable that there may be more than one way in which to remove from K the information that contradicts φ, and so the logicians Alchourrón, Gärdenfors, and Makinson in 1985 gave some very general constraints on how the removal should be done. Specifically, the *AGM postulates for contraction* give eight conditions that should be fulfilled by the set $K - \neg\varphi$ which results from the removal of the information contradicting φ. (Note that no specific operation for accomplishing the removal is given at this stage, just a general specification of how such an operation should behave.)

Similarly, the *AGM postulates for revision* describe eight conditions that should be fulfilled by the set $K * \varphi$ that results from the incorporation of the new information φ after removal of contradictory information.

Definition 4 *(AGM postulates for contraction) Given a belief set K, a contraction operation $-$ is an operation such that, for every sentence α, the set $K - \alpha$ (the result of contracting K by α) satisfies the following:*

1. $K - \alpha = Cn(K - \alpha)$
2. $K - \alpha \subseteq K$
3. $K - \alpha = K$ *if* $\alpha \notin K$
4. *If α is not a tautology,* $\alpha \notin K - \alpha$
5. *If* $\alpha \equiv \beta$, *then* $K - \alpha = K - \beta$
6. *If* $\alpha \in K$, *then* $(K - \alpha) + \alpha = K$ *(the Recovery Postulate)*
7. $(K - \alpha) \cap (K - \beta) \subseteq K - (\alpha \wedge \beta)$
8. *If* $\beta \notin K - (\alpha \wedge \beta)$ *then* $K - (\alpha \wedge \beta) \subseteq K - \beta$.

The first postulate says that $K - \alpha$ must again be a belief set.

The second ensures that contraction shrinks the belief set and doesn't expand it.

The third ensures that information is not taken away unnecessarily, in the sense that contracting will only change K if you contract by something actually in K.

The fourth postulate ensures that contraction by any sentence that actually contains a nonzero amount of semantic information will take that information away. (We can't get rid of beliefs that don't have any information, i.e. tautologies, because such tautologies exclude the set $\{\}$ of states, which is a subset of every excluded set of states, and therefore every belief set must contain a tautology.)

The fifth ensures that it is semantic information that is critical and that the contraction operation is not reacting to the syntactic form of α.

The sixth, known as the Recovery Postulate, requires the contraction of K by α to retain so much of the information of K that it is possible to recover the original K by putting α back (and adding in the classical consequences).

The last two postulates are included for technical reasons and are less intuitive.

Definition 5 *(AGM postulates for revision) Given a belief set K, a revision operation $*$ is an operation such that, for every sentence α, the set $K*\alpha$ (the result of revising K by α) satisfies:*

1. $K * \alpha = Cn(K * \alpha)$
2. $K * \alpha \subseteq K + \alpha$
3. $K + \alpha \subseteq K * \alpha$ *if* $\neg\alpha \notin K$
4. $\alpha \in K * \alpha$
5. *If* $\alpha \equiv \beta$*, then* $K * \alpha = K * \beta$
6. $K * \alpha = L_A$ *iff* α *is a contradiction*
7. $K * (\alpha \wedge \beta) \subseteq (K * \alpha) + \beta$
8. $(K * \alpha) + \beta \subseteq K * (\alpha \wedge \beta)$ *if* $\neg\beta \notin K * \alpha$.

The first revision postulate says that $K * \alpha$ must again be a belief set.

The second ensures that revision by α introduces no more information than conveyed by α.

The third ensures that if K contains no information contradicting α then the revised belief set keeps all the information in K and adds to it all the information of α.

The fourth postulate gives precedence to the information in α by ensuring that if any information has to be sacrificed to avoid inconsistency, it will be information in K and not α.

The fifth postulate stresses that semantic information counts, not syntactic form.

The sixth postulate expresses the idea that incorporating information by revision should result in a satisfiable belief set, but recognises that because of the fourth postulate, there is one case in which this goal cannot be achieved.

The last two postulates say that under certain conditions, successive revisions by α and then β give the same result as simultaneous revision by these sentences in the form of $\alpha \wedge \beta$.

How are expansion, contraction, and revision connected?

Theorem 1 *(The Levi identity)* $K * \varphi = (K - \neg\varphi) + \varphi$.

This tells us that we could arrive at the revised belief set $K * \varphi$ by first using contraction to remove from K any information contradicting φ and then using expansion to add φ to the result (which also adds in all the classical consequences). So, if we already have a contraction operation we could define a revision operation from it with the help of expansion. But this makes revision a clumsy two-step process, and we would prefer to revise more directly.

Implementing Contraction and Revision

We are now at the point where we have to face the million dollar question: How, exactly, should contraction and revision operations be defined?

A number of different proposals have been made. What fits nicely into the familiar framework of nonmonotonic logic is an approach first suggested by Grove in 1988, which makes use of preference relations. Grove showed that if we used preference relations that "faithfully" represented the belief set then, from those preference relations, there was a simple way to get all the contraction and revision operations that adhere to the AGM postulates. First let's say more precisely what we mean by the idea that a preference relation faithfully represents a belief set, and then we'll say how the operations should be defined using the preference relation.

In what follows, we will stick to total preorders on S as our preference relations, and we will stick to finite sets A and S so that every nonempty subset of S has at least one maximal element. These are the same restrictions we observed when introducing the rational consequence relations of nonmonotonic logic.

Now here's the idea of faithfulness. Recall that if $\leq$ is a preorder then we may write $s < t$ to mean that $s \leq t$ but not $t \leq s$. If we have a belief set K then we'll call a total preorder $\leq$ faithful to K if it places the models of K in the top layer and all nonmodels of K below them. (Note that faithfulness doesn't dictate how the nonmodels of K should be arranged relative to one another, so we will usually have several different preorders that are faithful to the same belief set K.)

Let's formalise the idea of faithfulness.

Definition 6 *(Faithfulness) Suppose that K is a belief set and that $\leq$ is a total preorder on S. Then $\leq$ faithfully represents K iff $\leq$ has the following two properties:*

- *for all $t \in Mod(K)$ and $s \notin Mod(K)$, we have $s < t$*
- *for all $t, t' \in Mod(K)$, we have $t \not< t'$.*

Examples of faithful total preorders:

- The total preorder

11	00
10	01

  faithfully represents the belief set $K = Cn(p \leftrightarrow q) = \{p \leftrightarrow q, \neg p \vee q, p \vee \neg q, p \vee \neg p\}$.

  But notice that this preorder does not faithfully represent the belief set $Cn(p \wedge q)$ because the state 00 in the top layer is not a model of $Cn(p \wedge q)$.

- The total preorder

11
10
00
01

  is faithful to $K = Cn(p \wedge q) = \{p \wedge q, p \leftrightarrow q, q, p, \neg p \vee q, p \vee \neg q, p \vee q, p \vee \neg p\}$.

  But notice that this preorder is not faithful to the belief set $Cn(p \leftrightarrow q)$ because 00 is a model of $Cn(p \leftrightarrow q)$ that is not in the top layer of the preorder.

Theorem 2 *(How to do contraction and revision) Consider a belief set K.*

A contraction operation for K satisfies the 8 contraction postulates iff there is some total preorder $\leq$ which is faithful to K and such that, for each α, $K - \alpha$ is the set of all sentences β such that $(Mod(K) \cup Max(\neg\alpha)) \subseteq Mod(\beta)$.

A revision operation for K satisfies the 8 revision postulates iff there is some total preorder $\leq$ which is faithful to K and such that, for each α, $K * \alpha$ is the set of all sentences β such that $Max(\alpha) \subseteq Mod(\beta)$.

We shall not prove the theorem, but it is important to appreciate what it says.

As far as defining contractions is concerned, the theorem tells us that if we have a belief set K and want to contract K by α, we need to have a faithful total preorder $\leq$ on S and then we may proceed in two stages: first we add to the models of K all the maximally preferred models of $\neg\alpha$, and then we take the set of all sentences satisfied by this enlarged set of models. The inclusion of models of $\neg\alpha$ will ensure that α is not one of the sentences in the contracted belief set. Note that in our finite set-up, once we have the desired set of models we can always find one sentence φ having exactly those models, and can then describe the contracted belief set as $Cn(\varphi)$.

As far as defining revisions is concerned, the theorem tells us that if we have a belief set K and learn reliable new information α, then we can get the revised belief set $K * \alpha$ by using a faithful total preorder $\leq$ on S to find $Max(\alpha)$ and then taking the set of all sentences satisfied by this set of models $Max(\alpha)$. Again, our finite set-up means that once we have the set of models we can find a single sentence φ having exactly those models and can thus describe the revised belief set as $Cn(\varphi)$.

Another way to think of the procedure by which we revise K by α is that we use the faithful total preorder $\leq$ on S to define the rational consequence relation $\mid\!\sim$ and then take $K * \alpha$ to be the set of all β such that $\alpha \mid\!\sim \beta$. Carrying this idea a step further, we see:

Corollary 1 *Suppose we have a belief set K and a total preorder $\preccurlyeq$ on S that is faithful to K. Then*

$$\beta \in K * \alpha \text{ iff } \alpha \mid\!\sim \beta.$$

This beautiful result tells us that nonmonotonic logic and AGM belief revision are two sides of the same coin. The beliefs an agent gets when revising by α are just the defeasible consequences of α under the appropriate rational consequence relation $\mid\!\sim$.

If we want to apply Theorem 2 to do contraction and revision, then we need to know what preorder accompanies the belief set. This preorder has to be faithful to the belief set, and so it usually can't be taken simply to be the preorder that represents the agent's heuristic information. There must be a way to also represent information obtained by observation, etc.

This area of belief change theory is still the subject of research. One influential approach, in case you are interested in reading more deeply, is that based on ordinal conditional functions, developed by Wolfgang Spohn. (Some references are given at the end.)

Contracting and revising in the Traffic System scenario

To illustrate how contraction and revision are done, we return to the driver waiting at the intersection. I will pick out one of the preorders faithful to the belief set K_1 and use it without claiming that it is a better choice than all other preorders faithful to K_1.

In the Traffic scenario, our driver initially formed the belief set

$$K_0 = \{p, p \vee q, p \vee \neg q, p \vee \neg p\}.$$

and then expanded with q to obtain the belief set

$$K_1 = K_0 + q = Cn(p \wedge q) = \{p \wedge q,\ p \leftrightarrow q,\ q,\ p,\ \neg p \vee q,\ p \vee \neg q,\ p \vee q,\ p \vee \neg p\}.$$

Suppose the total preorder faithful to K_1 is

11
10
00
01

Now recall Case 1, the situation in which the driver heard sirens and suddenly doubted the earlier inference that gave her the belief q. Because of the doubt, we don't want to add a new belief, we want to give up an old belief.

What do we get if we *contract* K_1 by q?

Well, $Mod(K_1) = \{11\}$ and $Max(\neg q) = \{10\}$ so $Mod(K_1 - q) = \{11\} \cup \{10\} = \{11, 10\}$. Looking around for a sentence with exactly these models we note that $Mod(p) = \{11, 10\}$ and thus:

$$K_1 - q = Cn(p).$$

This makes perfect sense. The sequence of events began with the driver first believing that p because she saw the light become red. Then she added the belief q because she defeasibly inferred that the oncoming car would stop. And then she contracted by q because she had reason to doubt her earlier inference, putting her back in the position of believing just p and its various classical consequences.

Note that the AGM postulates say nothing about the total preorder that is faithful to this new belief set $K_1 - q$. We have not covered the theoretical approaches to calculating the new preorder from the old, but obviously it would be important to have a way to produce a sensible preorder faithful to the new belief set, so that we can do further revisions and contractions as the need arises. The references on iterated revision at the end will tell you more about this.

What about Case 2? Suppose the driver with belief set $K_1 = Cn(p \wedge q)$ doesn't merely feel skeptical about the oncoming car stopping but actually sees the oncoming car run the red light. What do we get if we revise K_1 by $\neg q$, using the total preorder above?

Well, $Max(\neg q) = \{10\}$, and the sentence $p \wedge \neg q$ has exactly this set of models, so:

$$K_1 * \neg q = Cn(p \wedge \neg q).$$

Here revision by $\neg q$ has added in the new information $\neg q$, after first removing the contradictory information q. Again the result is as we would expect it to be. Notice also that the Levi identity holds, i.e. we could have obtained the revised belief set $K_1 * \neg q$ by first calculating the contraction $K_1 - q$ and then expanding by the new information $\neg q$ to get $(K_1 - q) + \neg q$. After all, $K_1 - q = Cn(p)$, so $(K_1 - q) + \neg q = Cn(p \wedge \neg q)$. But of course, we don't have to use the Levi identity because we have a direct way to get the revised set by using the preorder.

Notice that again we have said nothing about how one might get a new preorder faithful to the new belief set. There are various ways to do so, but general agreement has not yet been reached, and so this is still an area of research. (See the references on iterated revision for more about this.)

Finally, there is one aspect of belief sets that we have not explored. We all have some beliefs that we find harder to give up than other beliefs. If we were to take into account how *entrenched* various beliefs are, then this might affect the way contraction and revision are done. I have provided some references on entrenchment at the end, in case you may be interested.

Exercises

Quiz: The quiz question for lecture 6 will come from exercise 1 below.

1. Suppose we use $L_{\{p,q\}}$ and the set of states $S = W_{\{p,q\}} = \{11, 10, 01, 00\}$. Let the following be the total preorder faithful to the belief set K:

01
11
10
00

Taking for the belief set K the sentences satisfied by the states in the uppermost level of the refined preorder, work out the following belief sets:

K

$K - p$

$K - \neg p$

$K * p$

$K * \neg p$

$K - q$

$K - \neg q$

$K * q$

$K * \neg q$

$K - (p \wedge q)$

$K - \neg(p \wedge q)$

$K * (p \wedge q)$

$K * \neg(p \wedge q)$

$K - (p \wedge \neg q)$

$K * (p \wedge \neg q)$

$K - (p \vee q)$

$K * (p \vee q)$

$K - (p \to q)$

$K * (p \to q)$

Hint: To describe a belief set, you could of course list all its sentences, but this can be tedious. Find the set of models and use this to write the belief set in the form $Cn(\{\alpha\})$.

For further practice, pick your own total preorder on $S = W_{\{p,q\}} = \{11, 10, 01, 00\}$ and calculate the belief sets listed above. For example, pick two states to put into the top level, and place each of the remaining states into its own level lower down.

2. Consider the 3-Card System with $A = \{r_1, r_2, r_3, g_1, \ldots, b_3\}$ and $S = \{rgb, rbg, grb, gbr, brg, bgr\}$. Assume that the total preorder faithful to belief set K is given by the following:

rbg	gbr
rgb	grb
brg	bgr

Let K be the set of beliefs to which the ordering is faithful. Give, for each of the following, its set of models and then express the belief set in the form $Cn(\{\alpha\})$:

K

$K + r_1$

$K * r_1$

$K - b_2$

$K * g_2$

$K - g_3$

$K * g_3$

$K * r_3$

$K - (b_2 \to g_3)$

$K * (b_2 \to g_1)$

3. Suppose K is a belief set. Does contracting by a tautology leave K unchanged? Why?
4. Suppose K is a belief set. Does revising by a tautology leave K unchanged? Why?

Research

How good are the belief change operations that comply with the AGM postulates? This is an active area of research, and the operations of contraction and revision should be regarded as first approximations only. Below are some reasons why we do not consider them to be the final word. Selected research papers are given as a guide to further reading (although such additional reading is not required for exam purposes).

Should you decide to read some of the research on belief change, bear in mind that most researchers invert the direction of preference, so that more preferred states are lower down, whereas we have placed more preferred states higher up, as seems natural. The reason for the inverted direction of preference found in the literature is historical. In the early years of nonmonotonic logic, John McCarthy focused on how *abnormal* states were, not how normal they were, and thus greater abnormality placed the state higher up, i.e. more normal states were placed lower down. Whereas McCarthy concerned himself with how empty a glass is, some of us would rather think of how full it is.

Iterated belief revision

Although AGM belief revision is a step toward completing the picture of how agents ought to change their beliefs in the light of new information, it is incomplete. The approach describes how to produce, from a belief set K and a faithful ordering $\leq$, a new belief set $K * \varphi$ for any sentence φ, but no clue is given to suggest what the new order relation should be. It cannot be expected that the new ordering should be the same as $\leq$, because $\leq$ need not be faithful with regard to $K * \varphi$. The research area called *iterated belief revision* seeks to describe not only how the new belief set is to be produced but also how the new ordering (the *epistemic state*) is to be produced. For an introduction to research on iterated belief change, the following are some landmarks:

- Wolfgang Spohn: Ordinal conditional functions: A dynamic theory of epistemic states. In Harper and Skyrms (eds): *Causation in Decision, Belief Change, and Statistics* 105–134, Kluwer Academic Publishers 1988.
- Darwiche A & Pearl J: On the logic of iterated belief revision. *Artificial Intelligence* 89:1-29 1997.
- Chopra S, Ghose A & Meyer T: Iterated revision and the axiom of recovery: A unified treatment via epistemic states. In van Harmelen (ed): *ECAI2002: 15th European Conference on Artificial Intelligence*, pp 541-545, IOS Press 2002.
- Rott H: Shifting Priorities: Simple Representations for Twenty-Seven Iterated Theory Change Operators. In Makinson, Malinowski, and Wansing (eds): *Towards Mathematical Philosophy*, Trends in Logic, 28, pp 269–296, Springer 2009.

Variations on revision

Turning to a critical evaluation of revision, recall postulate 4 which says that revision by α is always successful, i.e. when the new information α is added to K, any contradictory information in K is sacrificed, so revision is biased in favour of the new information. This is fine if the new information is clearly correct, but what if the agent obtains the new information from a newspaper, so that the source is not absolutely reliable? Should the belief change operation not be more balanced when it looks at whether to sacrifice information in K?

Many alternative operations have been explored. One possible framework is described in Chopra, Ghose & Meyer (2003): Non-prioritized ranked belief change, *Journal of Philosophical Logic* 32:417-443.

This paper extends the semantics so that one can incorporate ways to add information that are more balanced than revision and that take the reliability of the information source into account. Quite fascinating.

Variations on contraction

Not everyone is convinced that the Recovery Postulate (the 6th contraction postulate) is correct. The Recovery Postulate says that $(K - \varphi) + \varphi = K$ as long as we start with a sentence $\varphi \in K$. A well-known researcher on belief change, Sven Ove Hansson, gives the following example to show why (he thinks) one might not always want this postulate to hold.

Suppose I read a book about Cleopatra, in which the claim is made that she had a son and a daughter. Suppose I then discover that the book is a work of fiction. According to Hansson, it would be quite reasonable for me to remove my belief that Cleopatra had children.

In other words, if we use the language with $A = \{p, q\}$ and let p stand for 'Cleopatra had a son' and q for 'Cleopatra had a daughter', then initially $K = Cn\{p \wedge q\} = \{p \wedge q,\ p,\ q,\ p \leftrightarrow q,\ p \vee q,\ p \vee \neg q,\ \neg p \vee q,\ p \vee \neg p\}$. Upon learning that the book was fictional, I change my belief set to $K - (p \vee q)$. For the moment, it is not clear exactly what the new belief set contains (no ordering on states has been given that I can use to construct the new belief set), but at the very least I know that the new belief set cannot contain $p \wedge q$, because if it did then it would also contain $p \vee q$, which is supposed to have been removed.

Now imagine that I go on to consult a history book and discover that Cleopatra did indeed have at least one child, although no details are given of whether there was more than one nor what the sex(es) may have been. This new information thus tells me simply that $p \vee q$.

If I now expand my belief set with the new information $p \vee q$, then by the Recovery Postulate the result should be K. In other words, having learned that Cleopatra had at least one child I should suddenly believe that she had both a son and a daughter (because $p \wedge q \in K$).

This would strike many people as unreasonable. Accordingly, research has been conducted on alternative forms of 'contraction' that do not have to satisfy the Recovery Postulate. The papers below look at alternative approaches, such as 'withdrawal' and 'removal' operations:

- Meyer T, Heidema J, Labuschagne W & Leenen L: Systematic withdrawal. *Journal of Philosophical Logic* **31**:415-443 2002.
- Booth R, Chopra S, Meyer T & Ghose A: A unifying semantics for belief change. In Ramon Lopez de Mantaras and Lorenza Saitta (eds): *ECAI2004: 16th European Conference on Artificial Intelligence*, pp 793-797, IOS Press 2004.

The Ghose-Goebel paradox

In the case of contraction, AGM postulate 3 says that if one contracts by a sentence α which is not in fact a belief inside K, then the effect of the contraction should be to do nothing, i.e. $K - \alpha = K$. But consider the Ghose-Goebel paradox. (We use $L_{\{p,q\}}$ with $S = \{11, 10, 01, 00\}$ to illustrate.)

Imagine an agent who starts off knowing nothing, i.e. having a belief set $K_0 = Cn(p \vee \neg p)$. Associate with this belief set the faithful total preorder $\preccurlyeq_0$ given by

$$\boxed{11 \quad 10 \quad 01 \quad 00}$$

which represents having zero information since no state is strictly preferred to any other state.

Next the agent learns that $p \to q$, and revision gives $K_1 = K_0 * (p \to q) = Cn(p \to q)$ which has to have associated with it the faithful total preorder $\preccurlyeq_1$ given by

11	01	00
	10	

since the maximal models of $p \to q$ (which happen to be all the models of $p \to q$ in the case of $\preccurlyeq_1$) must be in the top layer.

Now suppose the agent becomes skeptical about q and therefore contracts K_1 by q to get $K_2 = K_1 - q$. Calculating the models of K_2 we see that $Max(\neg q) = \{00\}$ so that $Mod(K_1) \cup \{00\} = Mod(K_1)$ and therefore nothing changes and $K_2 = K_1 - q = K_1$, just as the third AGM postulate for contraction demands, since q wasn't in K_1 to begin with. Since nothing has changed, we associate with K_2 the same total preorder as before, namely $\preccurlyeq_2$ given by

11	01	00
	10	

Finally, suppose the agent learns that p, so that revision gives $K_3 = K_2 * p = Cn(p \wedge q)$ since the maximal model of p in $\preccurlyeq_2$ is 11.

This has all been straightforward, so what's the paradox? Well, notice that $q \in K_3$. So the agent has ended up believing q. But recall that earlier the agent became skeptical of q, and nothing that happened afterwards had anything to do with making the agent feel less skeptical about q. After all, learning that p is the case is something independent of whether q is true or false.

The difficulty seems to be that because contracting K_1 by q didn't change the belief set, all memory of the skepticism about q was lost. How can one record the acquisition of such skepticism? Well, becoming skeptical about q is not the same thing as actually believing $\neg q$, which is why $K - q$ is something different from $K * \neg q$. But one clever idea is that when the agent becomes skeptical about a sentence α, then this should not only lead to a contraction by α but should also be recorded in a separate set of *disbeliefs*, and these disbeliefs ought then somehow to have an influence on what new beliefs the agent may acquire. The most thorough explanation of how this might work is given in the rather complicated but fascinating paper: Chopra, Ghose & Meyer (2003): Non-prioritized ranked belief change, *Journal of Philosophical Logic* 32:417-443.

Entrenchment

When an agent is forced to give up some of the information in its belief set, it is not equally ready to give up all beliefs — some beliefs are more entrenched than others. For example, an agent might hold more strongly to a belief based on what that agent saw with her own eyes than she would cling to a belief based on a journalist's report. A total preorder on the set of states induces an entrenchment ordering on the sentences of the language. For different approaches to entrenchment, see:

- Gärdenfors P & Makinson D: Revisions of knowledge systems using epistemic entrenchment. In Vardi MY (ed): *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge* pp83-95. Morgan Kaufmann 1988.
- Meyer T, Labuschagne W & Heidema J: Refined epistemic entrenchment. *Journal of Logic, Language, and Information* **9**:237-259 2000.