

# Chapter 2

## A General Framework for Handling Inconsistency

This chapter proposes a general framework for handling inconsistency under any monotonic logic. Basically, our approach to reason with an inconsistent knowledge base (KB) is a process which follows three steps:

1. Determining consistent “subbases”;
2. Selecting among all the subbases the ones that are *preferred*;
3. Applying entailment on the preferred subbases.

Throughout the rest of this chapter, we assume that we have an arbitrary, but fixed *monotonic* logic  $(\mathcal{L}, \text{CN})$ .

The basic idea behind our framework is to construct what we call *options*, and then to define a preference relation on these options. The *preferred options* are intended to support the conclusions to be drawn from the inconsistent knowledge base. Intuitively, an option is a set of formulas that is both consistent and closed w.r.t. consequence in logic  $(\mathcal{L}, \text{CN})$ .

**Definition 2.1 (Options).** An *option* is any set  $\mathcal{O}$  of elements of  $\mathcal{L}$  such that:

- $\mathcal{O}$  is consistent.
- $\mathcal{O}$  is closed, i.e.,  $\mathcal{O} = \text{CN}(\mathcal{O})$ .

We use  $\text{Opt}(\mathcal{L})$  to denote the set of all options that can be built from  $(\mathcal{L}, \text{CN})$ .

Note that the empty set is not necessarily an option. This depends on the value of  $\text{CN}(\emptyset)$  in the considered logic  $(\mathcal{L}, \text{CN})$ . For instance, in propositional logic, it is clear that  $\emptyset$  is not an option since all the tautologies will be inferred from it. Indeed, it is easy to see that  $\emptyset$  is an option iff  $\text{CN}(\emptyset) = \emptyset$ .

Clearly, for each consistent subset  $X$  of  $\mathcal{L}$ , it holds that  $\text{CN}(X)$  is an option (as  $\text{CN}(X)$  is consistent and Idempotence axiom entails that  $\text{CN}(X)$  is closed). Since we are considering generic logic, we can show that options do not always exist.

**Proposition 2.1.** *The set  $\text{Opt}(\mathcal{L}) = \emptyset$  iff*

1.  $\forall \psi \in \mathcal{L}, \text{CN}(\{\psi\})$  is inconsistent, and
2.  $\text{CN}(\emptyset) \neq \emptyset$ .

*Proof.* ( $\Rightarrow$ ) Let us assume that  $\text{Opt}(\mathcal{L}) = \emptyset$ . Let us also assume by contradiction that  $\exists \psi \in \mathcal{L}$  such that  $\text{CN}(\{\psi\})$  is consistent. Since  $\text{CN}(\{\psi\})$  is closed by the Idempotence axiom,  $\text{CN}(\{\psi\})$  is an option, which is a contradiction.

Assume now that  $\text{CN}(\emptyset) = \emptyset$ . This means that  $\emptyset$  is an option since it is closed and consistent ( $\text{CN}(\emptyset) \neq \mathcal{L}$ ), which is a contradiction.

( $\Leftarrow$ ) Let us assume that (i)  $\forall \psi \in \mathcal{L}$ ,  $\text{CN}(\{\psi\})$  is inconsistent, and (ii)  $\text{CN}(\emptyset) \neq \emptyset$ . Assume also by contradiction that  $\text{Opt}(\mathcal{L}) \neq \emptyset$  and let  $\mathcal{O} \in \text{Opt}(\mathcal{L})$ . There are two cases:

Case 1:  $\mathcal{O} = \emptyset$ . Consequently,  $\text{CN}(\emptyset) = \emptyset$ , which contradicts assumption (ii).

Case 2:  $\mathcal{O} \neq \emptyset$ . Since  $\mathcal{O}$  is consistent,  $\exists \psi \in \mathcal{O}$  s.t.  $\{\psi\}$  is consistent and thus  $\text{CN}(\{\psi\})$  is consistent. This contradicts assumption (i).  $\square$

So far, we have defined the concept of option for any logic  $(\mathcal{L}, \text{CN})$  in a way that is independent of a knowledge base. We now show how to associate a set of options with an inconsistent knowledge base.

In most approaches for handling inconsistency, the maximal consistent subsets of a given inconsistent knowledge base have an important role. This may induce one to think of determining the options of a knowledge base as the closure of its maximal consistent subsets. However, this approach has the side effect of dropping entire formulas, whereas more fine-grained approaches could be adopted in order to preserve more information of the original knowledge base. This is shown in the following example.

*Example 2.1.* Consider the propositional knowledge base  $\mathcal{K} = \{(a \wedge b); \neg b\}$ . There are two maximal consistent subsets, namely  $MCS_1 = \{a \wedge b\}$  and  $MCS_2 = \{\neg b\}$ . However, one could argue that  $MCS_2$  is too weak, since we could have included  $a$  by “weakening” the formula  $(a \wedge b)$  instead of dropping it altogether.

The “maximal consistent subset” approach, as well as the one suggested in the previous example, can be seen as a particular case of a more general approach, where one considers consistent “relaxations” (or *weakenings*) of a given inconsistent knowledge base. The ways in which such weakenings are determined might be different, as the following examples show.

*Example 2.2.* Consider again the temporal knowledge base of Example 1.2. An intuitive way to “weaken” the knowledge base might consist of replacing the  $\bigcirc$  (*next moment in time*) connective with the  $\diamond$  (*sometime in the future*) connective. So, for instance,  $\bigcirc \text{processed} \leftarrow \text{received}$  might be replaced by  $\diamond \text{processed} \leftarrow \text{received}$ , thus saying that if *received* is true at time  $t$ , then *processed* is true at some subsequent time  $t' \geq t$  (not necessarily at time  $t + 1$ ). This would lead to a consistent knowledge base, whose closure is clearly an option. Likewise, we might weaken only (1.6), obtaining another consistent knowledge base whose closure is an option.

*Example 2.3.* Consider the probabilistic knowledge base of Example 1.3. A reasonable way to make a probabilistic formula  $\phi : [\ell, u]$  weaker, might be to replace it with another formula  $\phi : [\ell', u']$  where  $[\ell, u] \subseteq [\ell', u']$ .

The preceding examples suggest that a flexible way to determine the options of a given knowledge base should be provided, since what is considered reasonable to be an option might depend on the logic and the application domain at hand, and, more importantly, it should depend on the user's preferences. The basic idea is to consider *weakenings* of a given knowledge base  $\mathcal{K}$  whose closures yield options. For instance, as said before, weakenings might be subsets of the knowledge base. Although such a weakening mechanism is general enough to be applicable to many logics, more tailored mechanisms could be defined for specific logics. For instance, the two reasonable approaches illustrated in Examples 2.2 and 2.3 above cannot be captured by considering subsets of the original knowledge bases; as another example, let us reconsider Example 2.1: by looking at subsets of the knowledge base, it is not possible to get an option containing both  $a$  and  $\neg b$ . We formally introduce the notion of *weakening* as follows.

**Definition 2.2.** Given an element  $\psi$  of  $\mathcal{L}$ ,

$$\text{weakening}(\psi) = \begin{cases} \text{CN}(\{\psi\}) & \text{if } \psi \text{ is consistent} \\ \emptyset & \text{otherwise} \end{cases}$$

**Definition 2.3.** Given a knowledge base  $\mathcal{K}$ ,

$$\text{weakening}(\mathcal{K}) = \{\mathcal{K}' \subseteq \mathcal{L} \mid \forall \psi' \in \mathcal{K}' (\exists \psi \in \mathcal{K} . \psi' \in \text{weakening}(\psi))\}$$

According to the preceding definitions, to weaken a knowledge base intuitively means to weaken formulas in it; to weaken a formula  $\psi$  means to take some formulas in  $\text{CN}(\{\psi\})$  if  $\psi$  is consistent, or to otherwise drop  $\psi$  altogether (note that a consistent formula could also be dropped). The set  $\text{weakening}(\mathcal{K})$  can be computed by first finding  $\text{weakening}(\psi)$  for all  $\psi \in \mathcal{K}$  and then returning the subsets of  $\bigcup_{\psi \in \mathcal{K}} \text{weakening}(\psi)$ . It is easy to see that if  $\mathcal{K}' \in \text{weakening}(\mathcal{K})$ , then  $\text{CN}(\mathcal{K}') \subseteq \text{CN}(\mathcal{K})$ .

Observe that although a knowledge base in  $\text{weakening}(\mathcal{K})$  does not contain any inconsistent formulas, it could be inconsistent.

**Definition 2.4.** A *weakening mechanism* is a function  $\mathcal{W} : 2^{\mathcal{L}} \rightarrow 2^{2^{\mathcal{L}}}$  such that  $\mathcal{W}(\mathcal{K}) \subseteq \text{weakening}(\mathcal{K})$  for any  $\mathcal{K} \in 2^{\mathcal{L}}$ .

The preceding definition says that a weakening mechanism is a function that maps a knowledge base into knowledge bases that are *weaker* in some sense. For instance, an example of a weakening mechanism is  $\mathcal{W}(\mathcal{K}) = \text{weakening}(\mathcal{K})$ . This returns *all* the weaker knowledge bases associated with  $\mathcal{K}$ . We use  $\mathcal{W}_{\text{all}}$  to denote this weakening mechanism.

We now define the set of options for a given knowledge base (w.r.t. a selected weakening mechanism).

**Definition 2.5.** Let  $\mathcal{K}$  be a knowledge base in logic  $(\mathcal{L}, \text{CN})$  and  $\mathcal{W}$  a weakening mechanism. We say that an option  $\mathcal{O} \in \text{Opt}(\mathcal{L})$  is an option for  $\mathcal{K}$  (w.r.t.  $\mathcal{W}$ ) iff there exists  $\mathcal{K}'$  in  $\mathcal{W}(\mathcal{K})$  such that  $\mathcal{O} = \text{CN}(\mathcal{K}')$ .

Thus, an option for  $\mathcal{K}$  is the closure of some weakening  $\mathcal{K}'$  of  $\mathcal{K}$ . Clearly,  $\mathcal{K}'$  must be consistent because  $\mathcal{O}$  is consistent (by virtue of being an option) and because  $\mathcal{O} = \text{CN}(\mathcal{K}')$ . In other words, the options for  $\mathcal{K}$  are the closure of consistent weakenings of  $\mathcal{K}$ . We use  $\text{Opt}(\mathcal{K}, \mathcal{W})$  to denote the set of options for  $\mathcal{K}$  under the weakening mechanism  $\mathcal{W}$ . Whenever  $\mathcal{W}$  is clear from the context, we simply write  $\text{Opt}(\mathcal{K})$  instead of  $\text{Opt}(\mathcal{K}, \mathcal{W})$ .

Note that if we restrict  $\mathcal{W}(\mathcal{K})$  to be  $\{\mathcal{K}' \mid \mathcal{K}' \subseteq \mathcal{K}\}$ , Definition 2.5 corresponds to that presented in Subrahmanian and Amgoud (2007) (we will refer to such a weakening mechanism as  $\mathcal{W}_{\subseteq}$ ). Moreover, observe that every option for a knowledge base w.r.t. this weakening mechanism is also an option for the knowledge base when  $\mathcal{W}_{\text{all}}$  is adopted, that is, the options obtained in the former case are a subset of those obtained in the latter case.

*Example 2.4.* Consider again the knowledge base of Example 2.1 and let  $\mathcal{W}_{\text{all}}$  be the adopted weakening mechanism. Our framework is flexible enough to allow the set  $\text{CN}(\{a, \neg b\})$  to be an option for  $\mathcal{K}$ . This weakening mechanism preserves more information from the original knowledge base than the classical “maximal consistent subsets” approach.

In Chap. 4 we will consider specific monotonic logics and show more tailored weakening mechanisms.

The framework for reasoning about inconsistency has three components: the set of all options for a given knowledge base, a preference relation between options, and an inference mechanism.

**Definition 2.6 (General framework).** A general framework for reasoning about inconsistency in a knowledge base  $\mathcal{K}$  is a triple  $(\text{Opt}(\mathcal{K}, \mathcal{W}), \succeq, \sim)$  such that:

- $\text{Opt}(\mathcal{K}, \mathcal{W})$  is the set of options for  $\mathcal{K}$  w.r.t. the weakening mechanism  $\mathcal{W}$ .
- $\succeq \subseteq \text{Opt}(\mathcal{K}, \mathcal{W}) \times \text{Opt}(\mathcal{K}, \mathcal{W})$ .  $\succeq$  is a partial (or total) preorder (i.e., it is reflexive and transitive).
- $\sim : 2^{\text{Opt}(\mathcal{K}, \mathcal{W})} \rightarrow \text{Opt}(\mathcal{L})$ .

The second important concept of the general framework above is the preference relation  $\succeq$  among options. Indeed,  $\mathcal{O}_1 \succeq \mathcal{O}_2$  means that the option  $\mathcal{O}_1$  is at least as preferred as  $\mathcal{O}_2$ . This relation captures the idea that some options are better than others because, for instance, the user has decided that this is the case, or because those preferred options satisfy the requirements imposed by the developer of a conflict management system. For instance, in Example 1.1, the user chooses certain options (e.g., the options where the salary is minimal or where the salary is maximal based on his needs). From the partial preorder  $\succeq$  we can derive the strict partial order  $\succ$  (i.e., it is irreflexive and transitive) over  $\text{Opt}(\mathcal{K}, \mathcal{W})$  as follows: for any  $\mathcal{O}_1, \mathcal{O}_2 \in \text{Opt}(\mathcal{K}, \mathcal{W})$  we say  $\mathcal{O}_1 \succ \mathcal{O}_2$  iff  $\mathcal{O}_1 \succeq \mathcal{O}_2$  and  $\mathcal{O}_2 \not\succeq \mathcal{O}_1$ . Intuitively,  $\mathcal{O}_1 \succ \mathcal{O}_2$  means that  $\mathcal{O}_1$  is *strictly* preferable to  $\mathcal{O}_2$ . The set of *preferred* options in  $\text{Opt}(\mathcal{K}, \mathcal{W})$  determined by  $\succeq$  is  $\text{Opt}^{\succeq}(\mathcal{K}, \mathcal{W}) = \{\mathcal{O} \mid \mathcal{O} \in \text{Opt}(\mathcal{K}, \mathcal{W}) \wedge \nexists \mathcal{O}' \in \text{Opt}(\mathcal{K}, \mathcal{W}) \text{ with } \mathcal{O}' \succ \mathcal{O}\}$ . Whenever  $\mathcal{W}$  is clear from the context, we simply write  $\text{Opt}^{\succeq}(\mathcal{K})$  instead of  $\text{Opt}^{\succeq}(\mathcal{K}, \mathcal{W})$ .

In the following three examples, we come back to the example theories of Chap. 1 to show how our framework can handle them.

*Example 2.5.* Let us consider again the knowledge base  $S$  of Example 1.1. Consider the options  $\mathcal{O}_1 = \text{CN}(\{(1.1), (1.3)\})$ ,  $\mathcal{O}_2 = \text{CN}(\{(1.1), (1.2)\})$ ,  $\mathcal{O}_3 = \text{CN}(\{(1.2), (1.3)\})$ , and let us say that these three options are strictly preferable to all other options for  $S$ ; then, we have to determine the preferred options among these three. Different criteria might be used to determine the preferred options:

- Suppose the *score*  $sc(\mathcal{O}_i)$  of option  $\mathcal{O}_i$  is the sum of the elements in the multiset  $\{S \mid \text{salary}(\text{John}, S) \in \mathcal{O}_i\}$ . In this case, the score of  $\mathcal{O}_1$  is 50K, that of  $\mathcal{O}_2$  is 110K, and that of  $\mathcal{O}_3$  is 60K. We could now say that  $\mathcal{O}_i \succeq \mathcal{O}_j$  iff  $sc(\mathcal{O}_i) \leq sc(\mathcal{O}_j)$ . In this case, the only preferred option is  $\mathcal{O}_1$ , which corresponds to the bank manager's viewpoint.
- On the other hand, suppose we say that  $\mathcal{O}_i \succeq \mathcal{O}_j$  iff  $sc(\mathcal{O}_i) \geq sc(\mathcal{O}_j)$ . In this case, the only preferred option is  $\mathcal{O}_2$ ; this corresponds to the view that the rule saying everyone has only one salary is wrong (perhaps the database has John being paid out of two projects simultaneously and 50K of his salary is charged to one project and 60K to another).
- Now consider the case where we change our scoring method and say that  $sc(\mathcal{O}_i) = \min\{S \mid \text{salary}(\text{John}, S) \in \mathcal{O}_i\}$ . In this case,  $sc(\mathcal{O}_1) = 50\text{K}$ ,  $sc(\mathcal{O}_2) = 50\text{K}$ ,  $sc(\mathcal{O}_3) = 60\text{K}$ . Let us suppose that the preference relation says that  $\mathcal{O}_i \succeq \mathcal{O}_j$  iff  $sc(\mathcal{O}_i) \geq sc(\mathcal{O}_j)$ . Then, the only preferred option is  $\mathcal{O}_3$ , which corresponds exactly to the tax agency's viewpoint.

*Example 2.6.* Let us consider the temporal logic theory  $T$  of Example 1.2. We may choose to consider just three options for determining the preferred ones:  $\mathcal{O}_1 = \text{CN}(\{(1.4), (1.5)\})$ ,  $\mathcal{O}_2 = \text{CN}(\{(1.4), (1.6)\})$ ,  $\mathcal{O}_3 = \text{CN}(\{(1.5), (1.6)\})$ . Suppose now that we can associate a numeric score with each formula in  $T$ , describing the reliability of the source that provided the formula. Let us say these scores are 3, 1, and 2 for formulas (1.4), (1.5) and (1.6), respectively, and the weight of an option  $\mathcal{O}_i$  is the sum of the scores of the formulas in  $T \cap \mathcal{O}_i$ . We might say  $\mathcal{O}_i \succeq \mathcal{O}_j$  iff the score of  $\mathcal{O}_i$  is greater than or equal to the score of  $\mathcal{O}_j$ . In this case, the only preferred option is  $\mathcal{O}_2$ .

*Example 2.7.* Consider the probabilistic logic theory  $P$  of Example 1.3. Suppose that in order to determine the preferred options, we consider only options that assign a single non-empty probability interval to  $p$ , namely options of the form  $\text{CN}(\{p : [\ell, u]\})$ . For two atoms  $A_1 = p : [\ell_1, u_1]$  and  $A_2 = p : [\ell_2, u_2]$ , let  $\text{diff}(A_1, A_2) = \text{abs}(\ell_1 - \ell_2) + \text{abs}(u_1 - u_2)$ . Let us say that the score of an option  $\mathcal{O} = \text{CN}(\{A\})$ , denoted by  $\text{score}(\mathcal{O})$ , is given by  $\sum_{A' \in P} \text{diff}(A, A')$ . Suppose we say that  $\mathcal{O}_i \succeq \mathcal{O}_j$  iff  $\text{score}(\mathcal{O}_i) \leq \text{score}(\mathcal{O}_j)$ . Intuitively, this means that we are preferring options that change the lower and upper bounds in  $P$  as little as possible. In this case,  $\text{CN}(\{p : [0.41, 0.43]\})$  is a preferred option.

Thus, we see that our general framework for managing inconsistency is very powerful – it can be used to handle inconsistencies in different ways based upon

how the preference relation between options is defined. In Chap. 4, we will consider more logics and illustrate more examples showing how the proposed framework is suitable for handling inconsistency in a flexible way.

The following definition introduces a preference criterion where an option is preferable to another if and only if the latter is a weakening of the former.

**Definition 2.7.** Consider a knowledge base  $\mathcal{K}$  and a weakening mechanism  $\mathcal{W}$ . Let  $\mathcal{O}_1, \mathcal{O}_2 \in \text{Opt}(\mathcal{K}, \mathcal{W})$ . We say  $\mathcal{O}_1 \succeq_W \mathcal{O}_2$  iff  $\mathcal{O}_2 \in \text{weakening}(\mathcal{O}_1)$ .

**Proposition 2.2.** Consider a knowledge base  $\mathcal{K}$  and a weakening mechanism  $\mathcal{W}$ . Let  $\mathcal{O}_1, \mathcal{O}_2 \in \text{Opt}(\mathcal{K}, \mathcal{W})$ .  $\mathcal{O}_1 \succeq_W \mathcal{O}_2$  iff  $\mathcal{O}_1 \supseteq \mathcal{O}_2$ .

*Proof.* ( $\Rightarrow$ ) Let  $\psi_2 \in \mathcal{O}_2$ . By definition of  $\succeq_W$ , there exists  $\psi_1 \in \mathcal{O}_1$  s.t.  $\psi_2 \in \text{weakening}(\psi_1)$ ; that is  $\psi_2 \in \text{CN}(\{\psi_1\})$ . Since  $\{\psi_1\} \subseteq \mathcal{O}_1$ , it follows that  $\text{CN}(\{\psi_1\}) \subseteq \mathcal{O}_1$  (by Monotonicity and the fact that  $\mathcal{O}_1$  is closed). Hence,  $\psi_2 \in \mathcal{O}_1$ .

( $\Leftarrow$ ) Let  $\psi_2 \in \mathcal{O}_2$ . Clearly,  $\psi_2 \in \text{weakening}(\psi_2)$ , since  $\psi_2$  is consistent and  $\psi_2 \in \text{CN}(\{\psi_2\})$  (Expansion axiom). As  $\psi_2 \in \mathcal{O}_1$ , the condition expressed in Definition 2.3 trivially holds and  $\mathcal{O}_1 \succeq_W \mathcal{O}_2$ .  $\square$

The following corollary states that  $\succeq_W$  is indeed a preorder (in particular, a partial order).

**Corollary 2.1.** Consider a knowledge base  $\mathcal{K}$  and a weakening mechanism  $\mathcal{W}$ .  $\succeq_W$  is a partial order over  $\text{Opt}(\mathcal{K}, \mathcal{W})$ .

*Proof.* Straightforward from Proposition 2.2.  $\square$

If the user's preferences are expressed according to  $\succeq_W$ , then the preferred options are the least weak or, in other words, in view of Proposition 2.2, they are the maximal ones under set inclusion.

The third component of the framework is a mechanism for selecting the inferences to be drawn from the knowledge base. In our framework, the set of inferences is itself an option. Thus, it should be consistent. This requirement is of great importance, since it ensures that the framework delivers *safe* conclusions. Note that this inference mechanism returns an option of the language from the set of options for a given knowledge base. The set of inferences is generally computed from the preferred options. Different mechanisms can be defined for selecting the inferences to be drawn. Here is an example of such a mechanism.

**Definition 2.8 (Universal Consequences).** Let  $(\text{Opt}(\mathcal{K}, \mathcal{W}), \succeq, \sim)$  be a framework. A formula  $\psi \in \mathcal{L}$  is a universal consequence of  $\mathcal{K}$  iff  $(\forall \mathcal{O} \in \text{Opt}^{\succeq}(\mathcal{K}, \mathcal{W})) \psi \in \mathcal{O}$ .

We can show that the set of inferences made using the universal criterion is itself an option of  $\mathcal{K}$ , and thus the universal criterion is a valid mechanism of inference. Moreover, it is included in every preferred option.

**Proposition 2.3.** *Let  $\langle \text{Opt}(\mathcal{K}, \mathcal{W}), \succeq, |\sim \rangle$  be a framework. The set  $\{\psi \mid \psi \text{ is a universal consequence of } \mathcal{K}\}$  is an option in  $\text{Opt}(\mathcal{L})$ .*

*Proof.* Let  $\mathcal{C} = \{\psi \mid \psi \text{ is a universal consequence of } \mathcal{K}\}$ . As each  $\mathcal{O}_i \in \text{Opt}^{\succeq}(\mathcal{K}, \mathcal{W})$  is an option,  $\mathcal{O}_i$  is consistent. Thus,  $\mathcal{C}$  (which is a subset of every  $\mathcal{O}_i$ ) is also consistent. Moreover, since  $\mathcal{C} \subseteq \mathcal{O}_i$ , thus  $\text{CN}(\mathcal{C}) \subseteq \mathcal{O}_i$  (by Monotonicity and Idempotence axioms),  $\forall \mathcal{O}_i \in \text{Opt}^{\succeq}(\mathcal{K}, \mathcal{W})$ . Consequently,  $\text{CN}(\mathcal{C}) \subseteq \mathcal{C}$  (according to the above definition of universal consequences). In particular,  $\text{CN}(\mathcal{C}) = \mathcal{C}$  because of the expansion axiom. Thus,  $\mathcal{C}$  is closed and consistent, and is therefore an option in  $\text{Opt}(\mathcal{L})$ .  $\square$

However, the following criterion

$$\mathcal{K} \sim \psi \text{ iff } \exists \mathcal{O} \in \text{Opt}^{\succeq}(\mathcal{K}, \mathcal{W}) \text{ such that } \psi \in \mathcal{O}$$

is not a valid inference mechanism since the set of consequences returned by it may be inconsistent, thus, it is not an option.