Noname manuscript No.
(will be inserted by the editor)

Believing in Default Rules: Inclusive Default Reasoning

Forthcoming in *Synthese*

Frederik J. Andersen & Rasmus K. Rendsvig

Abstract This paper argues for the reasonableness of an *inclusive* conception of default reasoning. The inclusive conception allows untriggered default rules to influence beliefs: Since a default "from φ, infer ψ" is a defeasible inference rule, it by default warrants a belief in the material implication $\varphi \to \psi$, even if φ is not believed. Such inferences are not allowed in standard default logic of the Reiter tradition, but are reasonable by analogy to the Deduction Theorem for classical logic. Our main contribution is a formal framework for inclusive default reasoning. The framework has a solid philosophical foundation, it draws conclusions non-trivially different from non-inclusive frameworks, and it exhibits a host of benchmark properties deemed desirable in the literature—e.g., that extensions always exist and are consistent.

Keywords: Rational Beliefs; Non-Monotonic Reasoning; Default Reasoning; Default Logic; Prioritized Default Logic; Packard's Lexicographic Lifting Rule.

1 Introduction

Default reasoning follows patterns like *in the absence of reasons to the contrary, from φ, conclude ψ*. Such inference patterns are widely used in everyday thinking and are established research topics in epistemology and computer science. Topics of interest in *default logic* include, e.g., what beliefs an ideal reasoner should hold given an acceptable set of *default rules* (or simply *defaults*), how some defaults defeat others, and what sets of defaults may reasonably be held jointly. A default rule may be thought of as a *defeasible generalization*, where learning the premise (by default) warrants a belief in the conclusion.

The first system of default logic was proposed and developed by Raymond Reiter (1980).[1] The purpose of Reiter's seminal work was to formalize reasoning with default assumptions, to which end he used defaults of the form

$$A : C/B \tag{1}$$

read by Horty (2007a) as "if A belongs to the agent's stock of beliefs, and C is consistent with these beliefs, then the agent should believe B as well" (p. 386). In

Address(es) of author(s) should be given

[1] Reiter's default logic was one of many non-monotonic reasoning frameworks developed in the late 1970 and 80s, with early papers collected and presented in (Ginsberg, 1987), including (Reiter, 1980) and other default logic approaches (Reiter and Criscuolo, 1981; Etherington and Reiter, 1983; Touretzky, 1986; Poole, 1988), but also, e.g., circumscription (McCarthy, 1980) and modal logical approaches (McDermott and Doyle, 1982; Moore, 1985). Later, AGM belief revision theory has been proposed as a non-monotonic system (Makinson and Gärdenfors, 1991). For overviews, see, e.g., (Ginsberg, 1987; Antoniou, 1999; Delgrande et al., 2004; Antonelli, 2005; Koons, 2017; Strasser and Antonelli, 2019).

a Reiter default like (1), A is called the *prerequisite*, C is the *justification* and B is the *consequent*.[2] A Reiter default in which the justification is logically equivalent to the consequent is called *normal*. Throughout this paper, we focus on normal defaults $A : B/B$, which we write '$A \rightsquigarrow B$'.[3]

As reasoning with defaults may be non-monotonic, classical logic does not suffice[4] as a guide for what conclusions to draw given some background information and a set of defaults (jointly called a *default theory*). Thus, a main task in default logic is to specify what conclusions are reasonable—to find the so-called *extensions* of a given default theory. Such extensions are often considered as rational fixed points that may be understood as cognitive equilibria of an ideal reasoner, and so may be equated with rational beliefs held on the basis of the default theory.

1.1 Two Guiding Examples

As a first very simple example of default reasoning—which involves no conflicts between defaults, no defeat, and no non-monotonic reasoning—consider the following adaption of Horty's example *Wedding of a Distant Relative* from (2012) (see also Section 4.3.3):

Two Guest Wedding. A relative is to be married, and only you and your aunt Petunia are invited as guests. You must RSVP and have no reasons not to attend. Moreover, you like spending time with Aunt Petunia, so if she goes, it gives you a reason to show as well. If Petunia cannot go, this also provides a reason for you to go, as the wedding will otherwise sadly be held without guests. If decided, what do you RSVP?

The case involves two defaults: "from Petunia attends, conclude I have a reason to attend", and "from Petunia does not attend, conclude I have a reason to attend". The structure of the example is simple for illustrative purposes. It could be part of arbitrarily complex and convoluted cases where the decision to attend, not attend, or remain agnostic, may result in different cascades of triggered defaults. Hence, the decision made in Two Guest Wedding may have severe consequences in more complex default reasoning scenarios.

Beliefs About Matches. Next, consider another simple example where an agent is reasoning about lighting a match, initially only with the propositions "the match

[2] A Reiter default may have multiple justifications, i.e., be of the form $\varphi : \psi_1, \ldots, \psi_n \setminus \chi$ with the reading "if φ is derived, and $\psi_1, \ldots, \psi_n$ are separately consistent with what is derived, then infer χ" (Brewka and Eiter, 2000), or "if φ is known and consistent with assumptions $\psi_1, \ldots, \psi_n$, then conclude χ" (Antoniou, 1999). Additionally, the formulas are normally allowed to be first-order.

[3] We still treat normal defaults as (defeasible) inference rules, and not formulas, but find the notation '$\rightsquigarrow$' easy to read.

[4] Monotonic logics—such as classical logic—satisfy that for any formula φ from the language L, if $\varphi \in L$ is a consequence of a set of formulas $\Gamma \subseteq L$ and if $\Gamma \subseteq \Delta \subseteq L$, then φ is also a consequence of Δ. That is, adding premises does not remove conclusions. Non-monotonic logics lack this property, and allow conclusions to be withdrawn in the light of new information.

is struck" and "the match ignites".[5] A *default* one may accept regarding these propositions is: "If the match is struck, it ignites". Given no background information at all, e.g., concerning the quality of the matches at hand and the potentially relevant weather conditions, what should the agent believe about the relationship between match striking and ignition? We return with our considerations below.

1.2 Two Conceptions of Default Reasoning

The paradigm of Reiter holds a near-monopoly on default logic. We can see this illustrated, for example, by the fact that a standard "sanity check" result about a new default logic framework is that its extensions are *refinements* of *Reiter extensions* (see below). An assumption of the Reiter paradigm is that untriggered defaults never influence beliefs. We argue that dropping this assumption may be well-motivated and lead to non-trivial differences in analyses of non-monotonic reasoning.

To delineate default logics with or without the assumption that only triggered defaults influence beliefs, we will refer to the *exclusive* conception versus the *inclusive* conception of default logic. According to the exclusive conception, untriggered defaults are *excluded* from influencing beliefs, while under the inclusive conception, they are not. These two conceptions of default reasoning can lead to very different conclusions in cases such as Two Guest Wedding and Beliefs About Matches.[6]

Under the exclusive conception, only triggered defaults may inform beliefs, i.e., only defaults with satisfied prerequisites (antecedents) may be used to extend beliefs beyond the given background information. In the Two Guest Wedding, the agent has no background information about Petunia's attendance. Hence, neither default is triggered, and so neither may influence the agent's beliefs about whether to attend. Under the exclusive conception, the agent thus *dispenses their decision about attending.*

Contrary to this, the decision-theoretic *Sure-Thing Principle* (Savage, 1954) holds that the agent should attend: The agent would take the action if they knew P (Petunia attends), and would take the action if they knew $\neg P$ (Petunia does not attend). Hence, even when the agent has no knowledge about P, they should attend, A—since A is a *sure thing.*

Under the inclusive conception, the agent attends the wedding, *because it is a sure thing.* The two untriggered defaults form a basis for beliefs $\{P \to A, \neg P \to A\}$ of which A is a logical consequence.

We present the finer details of both conceptions and analyses below, and here only reiterate that the difference observed in the Two Guest Wedding is non-negligible. One aspect is that the difference in decision could in more complex scenarios cascade, e.g., if A serves as a prerequisite in other defaults. Another more general perspective is

[5] We thank an anonymous reviewer for this example.

[6] Note however that a few exclusive accounts have "inclusive tendencies" (like input/output logic, cf. Section 5) in the sense that they can give the same solution to the Two Guest Wedding as our inclusive account.

that only the inclusive conception allows the agent to make use of *conditional proofs* in their belief formation—as done in the Sure-Thing Principle and the Deduction Theorem from classical logic that we return to for motivation.

As for the Beliefs About Matches-example, we can illustrate the difference between the inclusive and exclusive conceptions of default reasoning as follows. Consider an agent reasoning about lighting a match, initially only with the propositions S and I, which read “the match is struck” and “the match ignites”, respectively. Again, a default one may accept regarding these propositions is $(S \rightsquigarrow I)$, i.e., if the match is struck, it ignites. Given the background information that *the match is struck* both inclusive and exclusive accounts uncontroversially hold that the agent should believe proposition I, i.e., that the match ignites. The two conceptions differ when there is no background information, e.g., in the case where the match stays in the match box. Here, the default is not triggered, so exclusive accounts hold that the agent believes only the logical validities, while the inclusive conception holds that the agent additionally believes the material implication $(S \to I)$, i.e., if the match is struck, then it ignites. Hence, given no background information, an inclusive agent holds more beliefs about their world, rooted in the defaults they accept, than their exclusive counterpart. With all matches in the box, if asked about what they believe about the relationship between match striking and ignition, an exclusive agent will reply that they hold no particular beliefs about this relationship, except those logically valid, while an inclusive agent will reply that they believe that if one of the matches is struck, it will ignite.

We—the authors—think that these inclusive beliefs portray an epistemically bolder attitude towards the world. Even when untriggered, the defaults accepted by the agent still effect the agent’s beliefs. This may in turn effect their plans, e.g., about how to light a fire, as they can build a plan based upon the belief that striking the match will ignite it (which again may influence the actions they subsequently take). Under exclusive accounts of default reasoning, the accepted yet untriggered default goes to waste, epistemically speaking: The agent will hold no beliefs about match striking and ignition, and so have a relatively smaller epistemic base for planning.

However, the inclusive conception also holds that the agent believes the logical consequences of $(S \to I)$, and these may lead to false beliefs, which may naturally lead to objections to the conception. To exemplify, extend the case with the proposition M, which reads “the match is moist”—again with no background information available. A logical consequence of $(S \to I)$ is $(S \wedge M) \to I$, so the inclusive agent will also falsely believe that if the match is struck while moist, it will ignite. The exclusive agent will be spared this false belief: They will still believe only the logical validities. Holding such a false belief is of course unfortunate—but has noone ever rationally stricken a wet match hoping for fire? Prior to striking the moist match, the inclusive agent has no information that indicates that the match would not ignite. In making the attempt, they acted in accordance with their beliefs, and learned that they were wrong. Having made the attempt, the agent may revise their accepted defaults, and accept instead that striking dry matches ignite them and striking moist matches do not: $(S \wedge \neg M) \rightsquigarrow I$ and $(S \wedge M) \rightsquigarrow \neg I$. Given these accepted defaults and no background information, the inclusive agent will hold non-valid beliefs about their world, while the exclusive agent will not. Again, this will give the inclusive agent a basis for planning their actions, while the exclusive agent will be in the dark. Of course, the new inclusive beliefs may be again false in certain contexts: If the moist match is struck while lightning strikes, too, the match may ignite anyway. Another occasion for revising your defaults.

In sum, inclusive agents incorporate their untriggered defaults in their beliefs, and are thus more opinionated epistemically than exclusive agents in this regard. This may be beneficial as untriggered defaults lead to actionable beliefs, but also risky, as these beliefs may be false. Exclusive agents do not run this risk because they ignore the information contained in theoretically accepted defaults until they are forced by triggering background information to take it into account.

1.3 The Exclusive Conception

The exclusive conception is standard in the literature on default logic, going back to (Reiter, 1980). To illustrate the exclusive nature of Reiter's default logic, consider this simplified[7] rendition of his fixed point definition of *extensions*—the possible belief sets held on the basis of some given background information and a set of defaults:

Definition 1 (Reiter Extensions) Let $\Delta = (W, D)$ be a *default theory* with background information W and defaults D. For any set of formulas S, let $\Gamma(S)$ be the smallest set of formulas that satisfy

1. $W \subseteq \Gamma(S)$
2. $Th(\Gamma(S)) = \Gamma(S)$, where Th denotes logical closure, and
3. If the default "if φ, then ψ" is in D, and $\varphi \in \Gamma(S)$ while $\neg\psi \notin S$, then $\psi \in \Gamma(S)$.

A set of formulas E is a *Reiter extension* of Δ iff $\Gamma(E) = E$.

Applying the definition to Two Guest Wedding as our default theory, the unique Reiter extension is $E = Th(\emptyset)$, i.e., the logical closure of the empty set. Hence, all Reiter extensions of Two Guest Wedding exclude their defaults from influencing beliefs.

That Reiter's default logic is exclusive entails that all later augmentations of his original framework are exclusive, too. Of specific interest to this paper are frameworks that augment default theories (W, D) to *prioritized default theories* $(W, D, \leq)$, where $\leq$ is a priority (or preference) order on D, determining which defaults should be favored over others in case of conflict, or used to determine which Reiter extension to prefer when multiple exist. The predominant systems for prioritized default reasoning

[7] Simplified as we have not yet introduced a formal syntax, and to fit normal defaults. See (Reiter, 1980, Def. 1) for full details and the general case.

all suggest notions of extensions that are *refinements* of Reiter extensions; in the sense that they suggest preferred extensions that form a subset of the Reiter extensions for the same default theory. Hence, such frameworks inherit the exclusive trait from Reiter's default logic. This is the case for the prioritized frameworks of Marek and Truszczyński (1993), Baader and Hollunder (1993; 1995), Brewka (1994a; 1994b), Rintanen (1995; 1998), Brewka and Eiter (2000), Delgrande and Schaub (2000a), and Horty (2007b; 2007a; 2012).

1.4 The Inclusive Conception

Contrary to the predominant exclusive analysis, we submit that untriggered defaults may often reasonably inform beliefs, and that the inclusive conception leads to an intuitive analysis of cases like Two Guest Wedding and Beliefs About Matches.

Default rules that occur in a (prioritized) default theory are rules for drawing conclusions from premises; rules the agent finds *plausible*, i.e., the very fact that the default is present in the default theory *presupposes* the agent's readiness to believe its conclusion given the truth of its premise (unless this leads to conflict etc.). This sounds like deduction and it's supposed to: Default rules may be viewed as accepted inference rules, differing from classical inference rules like Modus Ponens only by being defeasible.

Now, treat default rules on par with classical inference rules, whenever this does not lead to conflict etc. For classical logic, we have the following

Deduction Theorem. For any set of formulas Σ and any formulas φ and ψ, if $\Sigma \cup \{\varphi\} \vdash \psi$, then $\Sigma \vdash \varphi \rightarrow \psi$.

Inspired by the Deduction Theorem for cases where $\Sigma = \emptyset$, we can formulate a

Minimal Default Deduction Assumption. For any default theory (W, D) and any consistent formulas φ, ψ, if $W = \emptyset$ and $D = \{$"from φ, infer ψ"$\}$, then $(\varphi \rightarrow \psi)$ is in all extensions of (W, D).

This assumption is in line with the Deduction Theorem in the sense that both allow notions of conditional proofs: If we from an assumed hypothesis/assumed piece of background information φ can prove/defeasibly conclude ψ, then this proves/defeasibly entails $(\varphi \rightarrow \psi)$.

Notice, however, that the minimal assumption is not quite as general as it was just stated. It allows conditional proofs only in near trivial, single rule default theories, about which it states that the default rule should be treated on par with classical inference rules by having its "material implication counterpart" included among the "theorems" of the default theory. If one thinks of default rules as standard inference rules (only defeasible), then the assumption is reasonable: As the default rule is in the default theory, the agent finds that it warrants its conclusion, unless the rule is defeated—which the consistency of φ and ψ ensures it isn't. Further, if the agent had background information $\{\varphi\}$, then they would infer ψ. Hence, in line with the Deduction Theorem, from the background information $W = \emptyset$, the agent may reasonably infer $(\varphi \rightarrow \psi)$. Thus, through thinking of default rules as inference rules and applying classical logical reasoning, the agent may endorse the material implication that the rule expresses.

We take satisfying the Minimal Default Deduction Assumption (properly stated for the framework at hand) as a **defining characteristic of an inclusive default logic** framework. To illustrate with an important case, Reiter's default logic does not satisfy this assumption. For the default theory $\Delta = (W, D)$ with $W = \emptyset$ and $D = \{A \rightsquigarrow B\}$ where A and B are consistent and A is non-valid, the unique Reiter extension is $E = Th(\emptyset)$. Reiter's framework is therefore exclusive, together with its many augmentations. The only inclusive framework we know of is the non-prioritized framework of Poole (1988), which we'll discuss in detail in Section 5 below. For now we simply take notice of Poole's original motivation for his *inclusive* framework:

> [...] [T]here seems to be two approaches to solving this problem of nonmonotonicity [...] The first is to claim that there is obviously something wrong with (say) classical logic and so there is a need to define a new logic to handle nonmonotonic reasoning [...] An alternative is to say that there is nothing wrong with classical logic; we should not expect reasoning to be just deduction from our knowledge. The proposal in this paper follows this second approach. (Poole, 1988, pp. 27-28)

According to Poole, Reiter's original default logic (1980) fits within the first of these mentioned approaches, while we—just as Poole—find nothing obviously "wrong" with golden standard monotonic logics, like classical logic, and thus see our own framework for non-monotonic reasoning as subsumed under the second approach. Rather than being a problem with logic, nonmonotonicity is a problem of how logic is properly used. In the present paper we stay close to classical logic in our reasoning about specific cases of interest. By analogy to the Deduction Theorem for classical logic, we use the Minimal Default Deduction Assumption as a defining characteristic of inclusive default logic frameworks because (1) its minimality makes it easy to find the appropriate phrasing of the condition for differing frameworks, thus making it easy to classify a framework as inclusive or exclusive, and (2) because the treatment of simple cases it describes, e.g., in the Two Guest Wedding-example, seems guiding for whether a framework follows the exclusive or the inclusive conception.

While the Minimal Default Deduction Assumption does not strictly speaking partition the theoretical space into exclusive and inclusive frameworks (the lack of a unified metatheory for all the relevant non-monotonic frameworks makes it difficult to phrase a requirement that applies to them all), our alignment with Poole is still a historically motivated way of distinguishing the framework we are proposing from those of the dominant Reiter tradition, among others. In approaching the problem of nonmonotonicity, we see our inclusive conception of default reasoning as closer in heritage to Poole's original work than that of Reiter. A substantial difference between Poole and Reiter is that Poole does not work with a set of default rules, but a set of possible hypotheses, which may be arbitrary consistent formulas. Poole's main thesis is that: *If one allows for hypothetical reasoning, there is no need to define a new logic to handle non-monotonic reasoning.* We find the implied simplicity intriguing and here submit a prioritized inclusive framework capable of non-monotonic reasoning with less divergence from classical logic than prioritized exclusive frameworks.

The inclusive conception provides an intuitive complimentary alternative to the predominant exclusive conception, and since the literature on prioritized default logics focuses on the exclusive conception, this paper develops an inclusive default logic framework for prioritized default theories.

We motivate the development of our inclusive framework by a close comparison to the exclusive framework due to John F. Horty, focusing on his presentation from the 2012 monograph *Reasons as Defaults*. We use *Reasons as Defaults* as a basis for two main reasons. The first is that the framework is mature, well-developed, and philosophically relevant: (Horty, 2012) is based on previous papers (Horty, 2007a,b), was well-received in both *Analysis* and *Mind* (Saka, 2014; Chrisman, 2015), is used without alterations in later philosophical applications (Horty, 2016), and is considered among the standard views of the logic of reasons (Bonevac, 2018). Second, the monograph presents an exclusive framework in a neatly modular fashion with individual definitions and numerous examples thoroughly discussed. This provides an opportunity to relate our main concepts with Horty's concepts (including their philosophical motivations), while also comparing the frameworks' analyses in concrete cases.

1.5 Roadmap

The paper proceeds as follows. Section 2 presents Horty's framework, first in overview, and then formally, throughout including Horty's intuitions. Section 3 presents the framework developed in this paper, referred to as the *inclusive model*. In defining the model, we make extensive use of a result from social choice theory by Packard (1981). The result concerns how one may lift an order on a set to an order on the power set of that set. Along the way, we'll discuss the assumptions behind Packard's result in relation to the inclusive model, and Horty's framework and intuitions. We further show that according to the inclusive model, for any default theory, beliefs exist and are consistent, a success property not satisfied by all default logic frameworks. We also establish that the inclusive model is indeed inclusive. Section 4 constitutes a sanity check and comparison to Horty's framework: We apply the inclusive model to multiple cases, including several from *Reasons as Defaults*. The section illustrates some differences and similarities between the exclusive and inclusive conceptions of default reasoning and serves as a basis for comparing with the intuitions of (Horty, 2012). Section 5 relates the inclusive model to the literature more broadly. This is done mainly by discussing various benchmark properties that are deemed desirable in other works on (prioritized) default logic, and showing how the inclusive model fares with respect to these. Finally, Section 6 contains some concluding remarks.

2 The Default Logic of *Reasons as Defaults*

To define the rational belief set(s) of an agent in an informational context, Horty's framework involves a host of notions defined in the following subsections. A rough outline of the framework is as follows.

The main aim is to establish an agent's *full belief set* given a context, i.e., a *default theory.* A default theory represents the initial data an idealized agent uses as a basis for reasoning (Horty, 2012, p. 22).

Horty works with prioritized default theories, each containing a set of *hard background information* W, a set of *defaults* D, and an *order* $<$ on the defaults. How the context is arrived at is not in question, only what to believe on the background of it is.

As the defaults D may produce conflicts (inconsistency), as some defaults may defeat others through priority, and as untriggered defaults should be excluded from influencing beliefs, the agent must select a reasonable subset of D on which to base their beliefs: They must find a *proper scenario.* Defining proper scenarios, i.e., scenarios that are also *rationally acceptable* (e.g., it should not allow us to conclude a contradiction from true premises), is the main task of the framework.

Finally, the rational belief set(s) are determined: A set of formulas is a rational belief set if it is the set of logical consequences of the combination of the background information and a proper scenario. As a default theory may admit multiple proper scenarios, each may also admit multiple rational belief sets, called *extensions.*

In the following, we present the formal details of the *Reasons as Defaults* framework, with a running commentary on interpretation.

Remark 1 The definitions below are labeled with references to (Horty, 2012). The labels are meant as conjunctions, so for example (Def. 7, p. 17, Def. 9) specifies a definition which is based on Horty's Definition 7, content from page 17 and Definition 9. We use notation that slightly differs from Horty's (e.g., the symbol '$\rightsquigarrow$' used in default rules) and introduce a few sets (e.g., $\mathcal{D}$ as the set of all default rules), but make no alterations to concepts defined in (Horty, 2012).

2.1 Language and Default Rules

Definition 2 (pp. 15–18) Throughout, fix a countable set of atomic propositions Φ and a language $\mathcal{L}$ given by

$$\varphi := p \mid \top \mid \neg\psi \mid \psi \rightarrow \psi'$$

where the symbol '$\to$' denotes material implication.[8] The remaining Boolean connectives are defined as usual. For $\Gamma \subseteq \mathcal{L}$, $\varphi \in \mathcal{L}$, write $\Gamma \vdash \varphi$ when φ is classically deducible from Γ. Denote the logical closure of Γ by $Th(\Gamma) := \{\varphi : \Gamma \vdash \varphi\}$.

Where $\varphi, \psi \in \mathcal{L}$, a **default rule**[9] is any expression of the form

$$(\varphi \rightsquigarrow \psi)$$

Denote the set of all default rules by $\mathcal{D}$ with typical elements δ, δ'. For a default rule $\delta = (\varphi \rightsquigarrow \psi)$ or a set of default rules $D \subseteq \mathcal{D}$, let

$$\begin{aligned} Premise(\delta) &:= \varphi & \qquad Premises(D) &:= \{Premise(\delta) : \delta \in D\}. \\ Conclusion(\delta) &:= \psi & \qquad Conclusions(D) &:= \{Conclusion(\delta) : \delta \in D\}. \end{aligned}$$

Intuitively, a default rule may be thought of as a *defeasible generalization*. A classic example is (*Tweety is a bird* $\rightsquigarrow$ *Tweety can fly*). By default, learning the premise warrants a belief in the conclusion, but additionally learning that Tweety is a penguin delegitimizes it. Hence, the rule is defeasible. As additional information may invalidate the conclusion, the inference is an example of non-monotonic reasoning. Horty interprets defaults as providing reasons for conclusions.[10]

2.2 Default Theories

The next definition specifies the core notions of the framework: a (fixed priority) default theory represents the initial data that an idealized agent can use as a basis for reasoning (Horty, 2012, p.22). Ensuing definitions provide refinements.

Definition 3 (Def. 1, p.22) A **(fixed priority) default theory** is a tuple $\Delta = (W, D, <)$ where

$W \subseteq \mathcal{L}$ is a set of *background information*,

$D \subseteq \mathcal{D}$ is a set of *available default rules*, and

$<$ is a strict partial *priority order* on D (i.e., $<$ is transitive and irreflexive).

A **scenario** based on Δ is a subset $S \subseteq D$.

Intuitively, "[...] a scenario is supposed to represent the particular subset of available defaults that have actually been selected by the reasoning agent as providing sufficient support for their conclusions—the particular subset of defaults, that is, to be used by the agent in extending the initial information from W to a full belief set, which we can then speak of as the belief set that is generated by the scenario" (Horty, 2012, p. 23).

[8] This notation—featuring vertical bars—is a common way of presenting a formal grammar in, for example, computer science (cf. Backus–Naur form). Note also that parts of sections 2 and 4 are based on (Andersen, 2024a,b).

[9] Default rules are not expressible in $\mathcal{L}$, and, as in (Horty, 2012), '$\rightsquigarrow$' cannot be nested.

[10] Horty (2007a, p. 368) writes: "Where A and B are formulas from the background language, we then let $A \rightsquigarrow B$ represent the *default rule* that allows us to conclude B, by default, whenever it has been established that A. It is most useful, I believe, to think of default rules as providing *reasons* for conclusions." In the quote, we have replaced Horty's notation '$\to$' with the present '$\rightsquigarrow$'.

Concerning the requirements on the relation $<$, Horty argues that transitivity is a natural requirement, that the relation should be irreflexive (i.e., strict) so that "no default can ever have a higher priority than itself" (*ibid.*, p. 20), and that the relation should not be strongly connected[11] as—though this would help to resolve conflicts between defaults—the requirement would be unreasonable, because: (1) some defaults are simply incommensurable, and (2) some defaults may have equal priorities.

Remark 2 Reason (2) contrasts with the choice of a *strict* order, and suggests using a preorder $\leq$ instead—which we do in Section 3. The order is then *reflexive* instead of irreflexive, with the also natural interpretation that every default is comparable to itself, and to itself it has *the same* priority. As a preorder, it may still be partial, in accordance with Horty's intuitive examples (*ibid.*, p. 20).

Remark 3 Horty's fixed priority default theories may be seen as a generalization of *normal* Reiter default theories, i.e., Reiter default theories (W, D) where all defaults are normal, with (W, D) represented by $\Delta = (W, D, \emptyset)$, cf. (Horty, 2007a).

2.3 Proper Scenarios

Horty remarks that belief sets based on arbitrary scenarios are unsatisfactory (*ibid.*, p. 23). Satisfactory belief sets are obtained only from *proper scenarios.* The definition of a proper scenario requires the auxiliary notions of *triggered*, *conflicted* and *defeated* defaults.

Definition 4 (Def. 2, p. 25, Def. 3, p. 27, Def. 4, p. 29) Let $S \subseteq D$ be a scenario based on $\Delta = (W, D, <)$. Define

$$\begin{aligned}
Triggered(\Delta, S) &= \{\delta \in D\colon W \cup Conclusions(S) \vdash Premise(\delta)\}.\\
Conflicted(\Delta, S) &= \{\delta \in D\colon W \cup Conclusions(S) \vdash \neg Conclusion(\delta)\}.\\
Defeated(\Delta, S) &= \{\delta \in D\colon \exists \delta' \in Triggered(\Delta, S) \text{ such that}\\
&\qquad \delta < \delta' \text{ and } W \cup \{Conclusion(\delta')\} \vdash \neg Conclusion(\delta)\}.
\end{aligned}$$

Using these three notions, Horty presents two definitions of a proper scenario. The first definition relies on the notion of a *binding default.* It is preliminary, but used throughout the book. The second is presented in his Appendix A.1 to handle certain problem cases.[12] We state the definitions in turn.

Definition 5 (Def. 5, p. 30) Let $S \subseteq D$ be a scenario based on $\Delta = (W, D, <)$. Define

$$Binding(\Delta, S) = (Triggered(\Delta, S) - Conflicted(\Delta, S)) - Defeated(\Delta, S).$$

[11] The priority order should not be assumed *connex,* that for any defaults δ, δ', either $\delta < \delta'$ or $\delta' < \delta$.

[12] Horty exemplifies: Let $\delta = \varphi \rightsquigarrow \varphi$ and $\Delta = (W, D, <)$ with $W = \emptyset$, $D = \{\delta\}$ and $< = \emptyset$. Then $S = \{\delta\}$ is stable as δ is triggered, and neither conflicted nor defeated. Yet the belief set $E = Th(\{\varphi\})$ is not grounded in the background information.

A scenario $S \subseteq D$ based on $\Delta = (W, D, <)$ is **stable** iff $S = Binding(\Delta, S)$. The scenario S is **proper**$_1$ iff it is stable.

The second definition is stronger, in that it implies stability, cf. Horty's Theorem 1 (*ibid.*, p. 223). It is based on the notion of an *approximating sequence*:

Definition 6 (Def. 26, Def. 27, pp. 222–223) Let $S \subseteq D$ be a scenario based on $\Delta = (W, D, <)$. Then $(S_n)_{n \in \mathbb{N}} = S_0, S_1, S_2, ...$ is an **approximating sequence** based on Δ and constrained by S iff

$$S_0 = \emptyset,$$

$$S_{i+1} = \{\delta : \delta \in Triggered(\Delta, S_i), \delta \notin Conflicted(\Delta, S), \delta \notin Defeated(\Delta, S)\}$$

The scenario S is **proper**$_2$ iff $S = \bigcup_{i \geqslant 0} S_i$ for some approximating sequence $(S_n)_{n \in \mathbb{N}}$.

Definitions 5 and 6 indicate that Horty's framework is exclusive, by requiring that a proper scenario contains only triggered defaults. We return to this below.

2.4 Extensions, Beliefs, and Exclusivity

Finally, Horty defines extensions of default theories:

Definition 7 (Def. 8, p. 32) The set E is an *extension* of $\Delta = (W, D, <)$ if there is some proper$_{\{1,2\}}$ scenario S such that

$$E = Th(W \cup Conclusion(S)).$$

This concludes the formal framework.[13]

Horty *does not* directly associate extensions with beliefs, cf. his discussion on pp. 34–40: A default theory Δ may have multiple or no extensions, and identifying the Δ-beliefs with *the* extension of Δ is therefore not well-defined. Horty discusses both multiple and lacking extensions, but he does not give a solution. As lacking extensions will not play a role in this paper, we ignore that problem. For multiple extensions, we conform our terminology to what we consider the least committal of Horty's three proposals: We interpret every extension of a default theory as a possible equilibrium state that an ideal reasoner might arrive at—as a *possible belief state*.

In relation to the exclusive/inclusive distinction, a result of Horty's, that relates his framework to Reiter's, shows that Horty's framework is exclusive:

[13] Horty revises the definition of defeat in Chapter 8, but writes "[...] [T]his preliminary definition [Def. 4] is adequate for a wide variety of ordinary examples, and in order to avoid unnecessary complication, we will rely on it as our official definition throughout the bulk of this book." (Horty, 2012, p.30). The revision affects the definitions of binding defaults and of approximating sequences, resulting in the two additional definitions of proper scenarios, but neither affects the result presented below. For completeness, we include the revised definition:
Definition (Def. 21, p.196)**.** Let $S \subseteq D$ be a scenario based on $\Delta = (W, D, <)$. Define $Defeated(\Delta, S) = \{\delta \in D\colon$ there is a set $D' \subseteq Triggered(\Delta, S)$ such that (1) $\delta < D'$, and (2) there is a set $S' \subseteq S$ such that (*a*) $S' < D'$, (*b*) $W \cup Conclusion((S - S') \cup D')$ is consistent, and (*c*) $W \cup Conclusion((S - S') \cup D') \vdash \neg Conclusion(\delta)\}$.

Theorem 1 (Thm. 5, p.232) *Let $\Delta = (W, D, <)$ be a default theory with D finite. Then any extension E of $(W, D, <)$ is a Reiter extension of (W, D).*

Hence, as Reiter extensions do not satisfy the Minimal Default Deduction Assumption, neither do Horty's.[14] [15]

3 A Framework for Inclusive Default Reasoning

In this section, we present a default reasoning framework for the inclusive conception. For ease of reference, we refer to the framework as the *inclusive model.*

The rudimentary idea is as follows: given a (preordered) default theory $\Delta = (W, D, \leq)$, we lift the priority (pre)order $\leq$ to a (pre)order $\preceq$ on the power set of defaults, 2^D, from which we define proper* scenarios as the maximal elements. To stay close to Horty's intuitions, the definition of $\preceq$ is motivated by his conception of defeat.

The order $\preceq$ satisfies that a single higher priority default out-prioritizes any number of lower priority ones (i.e., if for all $\delta_b \in B$, $\delta > \delta_b$, then $\{\delta\} \succ B$), while conservatively extending the subset relation. Hence, D will always be $\preceq$-maximal in 2^D. As D may often be unreasonable—e.g., by leading to inconsistent beliefs—we use two conditions to prune 2^D, resulting in sane maximal elements. These conditions may be seen as counterparts to Horty's notions of conflict and triggering.

Finally, we define proper* belief sets as those belief sets obtained from proper* scenarios, i.e., the $\preceq$-maximal sets of defaults of the pruned power set.

Two important properties of the inclusive model are that it is indeed inclusive, in the sense that it satisfies the Minimal Default Deduction Assumption, and that it ensures that belief sets always exist.

To induce the priority order $\preceq$, we make use of a result from social choice theory by Packard (1981), which requires a bit of work. To ease the presentation, we therefore first re-define belief sets, discuss power set pruning, and define proper* scenarios and proper* belief sets.

3.1 Priority Preorders and Finite Sets of Defaults

To apply the results from (Packard, 1981), we must make two concessions, in the form of two assumptions held throughout the remainder of the paper. In defining default

[14] The exclusivity may also be seen directly. Let $\Delta = (W, D, <)$ be a prioritized default theory with $W = \emptyset$, $D = \{(\varphi \rightsquigarrow \psi)\}$ and $<= \emptyset$. There are two possible scenarios of Δ: $S_0 = \emptyset$ and $S_1 = D$. Here, S_0 is trivially both proper$_1$ and proper$_2$, and provides the extension $E = Th(\emptyset)$. Hence, the Minimal Default Deduction Assumption is not satisfied.

[15] Horty (2012) also draws other relations to Reiter's framework, showing for instance that for any default theory with an empty priority order, his framework produces exactly the Reiter extensions (Thm. 4, p. 229), and when there are non-trivial priorities on infinite sets of defaults, his extensions may not all be Reiter extensions (p. 230).

theories $\Delta = (W, D, <)$, Horty states that $<$ should be a strict partial order, i.e., irreflexive, transitive and possibly not connected (connex). Beyond stating that it is natural that no default has a higher priority than itself, no argument is provided as to why the ordering must be *strict*, but it is argued from incommensurability that it should be *partial*. In the construction below, we apply a theorem by Packard (1981) detailing how to lift a *non-strict* and *total* order $\leq$ on a *finite* set X to a non-strict total order $\preceq$ on the power set 2^X. In a compromise between Horty and Packard, we retain partiality from Horty, but assume Packard's non-strictness and finiteness. Thus, we henceforth work with the following

Definition 8 A **default theory** is a tuple $\Delta = (W, D, \leq)$ with $W \subseteq \mathcal{L}$ a set of background information, D a finite set of defaults, and $\leq$ a preorder on D.

A preorder is a reflexive and transitive relation. For defaults $x, y \in D$, when $x \leq y$ and $y \geq x$, we write $x = y$ to mean that the defaults have *equal priority*. Hence, we do not think that a default can have higher priority than itself, but that it has equal priority with itself. With $\leq$ a preorder, defaults may still be incomparable, so we find the preorder assumption unproblematic. While the finiteness assumption is of course limiting, it is irrelevant to the examples discussed in (Horty, 2012).

With $\leq$ assumed a preorder, the order $\preceq$ we will define in Section 3.5 on the power set 2^D will also be a preorder, allowing both incomparability and equality. For $X, Y \in D$, we write $X \prec Y$ when $X \preceq Y$ and not $X \succeq Y$, and write $X \approx Y$ when $X \preceq Y$ and $Y \preceq X$. If $X \approx Y$, we say that X and Y have *equal priority*. The maximal elements of a subset of the power set $\mathcal{X} \subseteq 2^D$ is

$$\max_{\preceq} \mathcal{X} := \{A \in \mathcal{X} \colon \text{ for all } B \in \mathcal{X}, B \succeq A \text{ implies } B \approx A\}.$$

3.2 Belief Sets

When Horty uses extensions $E = Th(W \cup Conclusions(S))$ to define belief sets, this presupposes that the scenario S is suitably conflict-free and that all its defaults are triggered. Contrary to this, under the inclusive conception, we do want untriggered defaults to sometimes influence beliefs. To avoid requiring that defaults must be triggered to influence beliefs, we avoid making use of the $Conclusions(\cdot)$ map, and instead rely on *Modus Ponens* and the minimality of logical closure to ensure that no default conclusions, unwarranted by the background information, are believed. As such, the definition of belief sets in use here is closer in spirit to Poole's definition of extensions, treating defaults as Poole's *possible hypotheses* (Poole, 1988), to which we compare in Sec. 5.1.

Definition 9 For a default theory $\Delta = (W, D, \leq)$ and scenario $S \subseteq D$, define the **material implication counterpart** of S as

$$S^{\rightarrow} := \{(\varphi \rightarrow \psi) \in \mathcal{L} \colon (\varphi \rightsquigarrow \psi) \in S\}.$$

Define the **belief set** given W and S as

$$\mathsf{B}(W,S) := Th(W \cup S^{\rightarrow}).$$

A belief set need not be very reasonable. For that, we want restrictions on S, ultimately such that it is proper*—i.e., $\preceq$-maximal in the pruned power set.

3.3 Power Set Pruning

As mentioned above, the $\preceq$-maximal element of 2^D will always be D itself, which may lead to inconsistent or otherwise unreasonable beliefs. However, belief sets are only of interest when consistent. Hence, we want to prune the power set down to contain only scenarios that ensure consistency:

Definition 10 For default theory $\Delta = (W, D, \leq)$, define the set of **Δ-consistent scenarios** by

$$\mathsf{Con}_{\Delta} := \{S \subseteq D \colon \mathsf{B}(W,S) \text{ is consistent}\}.$$

Seeking $\preceq$-maximal elements only in Con_{Δ} is one way to prune 2^D, and by definition, then, for any $S \in \mathsf{Con}_{\Delta}$, the belief set $\mathsf{B}(W,S)$ is consistent. By ensuring consistency, this pruning ensures the core functionality of Horty's concept of conflict. Horty (2012, p. 27) writes that a default is *conflicted* whenever the reasoner in question is "...already committed to the negation of its conclusion."

Consistency with the background information is the bare minimum. A slightly stronger requirement for proper* scenarios is that they also yield consistent conclusions to any consistent combination of their premises. This stronger requirement rules out, e.g., $\{p \rightsquigarrow q, p \rightsquigarrow \neg q\}$ as a proper* scenario for a default theory with $W = \{\neg p\}$, while consistency alone does not. An additional pruning of the power set, which seems in line with intuitions underlying Horty's notion of conflict,[16] is the pruning to scenarios that are also *coherent*:

Definition 11 For default theory $\Delta = (W, D, \leq)$, call $S \subseteq D$ **coherent** when for all $S' \subseteq S$, if $Premises(S')$ is consistent, then $Conclusions(S')$ is consistent. Let

$$\mathsf{Coh}_{\Delta} := \{S \subseteq D \colon S \text{ is coherent}\}, \text{ and}$$
$$\mathsf{CC}_{\Delta} := \mathsf{Con}_{\Delta} \cap \mathsf{Coh}_{\Delta}.$$

The final pruning condition concerns discerning defaults that are triggered from those that are not. Horty (2012, p. 25) writes: "The *triggered* defaults are supposed to represent those that are applicable in the context of a particular scenario; they are defined as the defaults whose premises are entailed by that scenario—those defaults, that is, whose premises follow from the initial information belonging to the underlying default

[16] Requiring only consistency is fully compatible with the remaining approach and main results, though the analysis of some examples will change.

theory together with the conclusions of the defaults already endorsed." Inspired by this, we say that:

Definition 12 For default theory $\Delta = (W, D, \leq)$, call $S \subseteq D$ **triggered** if for all defaults $(\varphi \rightsquigarrow \psi) \in S$, $\varphi \in Th(W \cup S^{\rightarrow})$, and let

$$\mathsf{Tr}_\Delta := \{S \subseteq D \colon S \text{ is triggered}\}.$$

We will *not* be invoking triggering as a pruning requirement in the same fashion as we will consistency and coherence. If we prune the power set down to contain only triggered scenarios, then we again produce an exclusive framework. Instead, we only require that a proper* scenario must be a superset of some triggered scenario (which may possibly be empty).

3.4 Proper* Scenarios and Proper* Belief Sets

The preorder $\preceq$ will be used to select the highest prioritized sets of defaults in two steps. First, we prune the power set down to the set of consistent, coherent, and triggered scenarios. Among these, we select the maximal elements to form foundations $\mathcal{F}_\Delta$ for the proper* scenarios. Second, we then re-inflate these foundations to consistent and coherent scenarios $\mathcal{S}_\Delta$. Among these re-inflations, the maximal elements constitute the proper* scenarios, which in turn result in "rational" belief sets $\mathcal{B}_\Delta$:

Definition 13 Given $\Delta = (W, D, \leq)$, let $\preceq$ be the preorder induced on 2^D. Let

$$\mathcal{F}_\Delta := \max_{\preceq}(\mathsf{CC}_\Delta \cap \mathsf{Tr}_\Delta),$$

and define the set of **proper* scenarios** for Δ to be

$$\mathcal{S}_\Delta := \max_{\preceq}(\{S \in \mathsf{CC}_\Delta \colon S \supseteq T \text{ for some } T \in \mathcal{F}_\Delta\}).$$

Finally, call

$$\mathcal{B}_\Delta := \{\mathsf{B}(W, S) \colon S \in \mathcal{S}_\Delta\}$$

the set of **proper* belief sets** for Δ.

That $\mathcal{S}_\Delta$ and $\mathcal{B}_\Delta$ are well-defined and non-empty follows from Corollary 1 shown in the end of the section.

3.5 Defining $\preceq$: Packard's Characterization through Priority and Defeat

What criteria should we impose for the relationship between $\leq$ on D and $\preceq$ on the power set 2^D to stay faithful to concepts of priority and defeat? Initially, we stipulate that $\preceq$ should be an *extension* of $\leq$, i.e., should satisfy

Extension : For all $x, y \in D$, if $x \leq y$ and $\{x\}, \{y\} \in 2^D$, then $\{x\} \preceq \{y\}$.

How should we extend $\preceq$ beyond the singletons?

3.5.1 A Criterion from Defeat

Recall from Horty's framework that

$$\begin{aligned} Defeated(\Delta, S) = \{\delta \in D\colon \exists\delta' \in Triggered(\Delta, S) \text{ such that} \\ \delta < \delta' \text{ and } W \cup \{Conclusion(\delta')\} \vdash \neg Conclusion(\delta)\}. \end{aligned}$$

Inspecting this definition, we see that a single (triggered) high-priority default takes priority over any set of (triggered) lower-priority defaults. Hence, if a single high-priority default (together with the background information) entails the negation of the conclusion of *every* low-priority default in some set, then that entire set is defeated. As a criterion on the preorder $\preceq$, this translates to the following:

Single Dominance : For $\{a\}, B \in 2^D$ such that $\{a\} \cup B \notin \mathsf{CC}_\Delta$, if $a > y$ for all $y \in B$, then $\{a\} \succ B$.

The definition of defeat also implies the following stronger condition, generalizing Single Dominance from a singleton $\{a\}$ to any non-empty set A:

General Dominance : For $A, B \in 2^D, A \neq \emptyset$, such that $A \cup B \notin \mathsf{CC}_\Delta$, if $\forall x \in A, \forall y \in B, x > y$, then $A \succ B$.

General Dominance ranks only very specific sets of defaults, viz., those that are not jointly coherent or consistent, and where all elements in one defeats all those in the other.

What about sets that are not mutually defeating? Say, for example, that defaults a, b and c are jointly coherent and consistent with the background information while $a > b > c$. How should we consider the relationship between $\{a\}$ and $\{b, c\}$? General Dominance does not apply as $\{a\} \cup \{b, c\} \in \mathsf{CC}_\Delta$.

Here, we turn to a condition suggested by Packard (1981) in discussing plausibility orders on finite, consistent sets of statements. The condition states that if two sets are *disjoint* and all elements of the first are ranked higher than those of the second, then the first set is ranked higher than the second set:

Composition : For $A \neq \emptyset, A \cap B = \emptyset$, if $\forall x \in A, \forall y \in B, \{x\} \succ \{y\}$, then $A \succ B$.

Composition combines aspects of Extension and General Dominance to allow the dominance aspect of defeat to also rank compatible sets of defaults—such as $\{a\}$ and $\{b, c\}$. However, it does not imply rank between non-disjoint subsets, which we turn to below.

The effect of Composition is to ensure that many low-ranked elements do not add up to outrank fewer highly ranked elements. Additionally, Composition implies that any set of defaults is better than none: if A is non-empty, then $A \succ \emptyset$. Conceptually, this means that no defaults have downright negative priority.[17] Beyond this, we find that Composition captures the core role of the priority order in relation to defeat.

[17] One may object that not every set of defaults is better than none, as there are defaults that are of downright negative priority. We agree. We ignore this problem here, for two reasons. First, it can be

3.5.2 Ordering Overlapping Sets of Defaults

Where Composition is a condition on disjoint sets, Packard additionally considers a condition for overlapping sets. It states that whenever two sets ($X = A \cup C$ and $Y = B \cup C$) have a non-empty intersection (C), then the ordering is independent of their overlap:

Independence : For $A \cap C = B \cap C = \emptyset$, $A \cup C \preceq B \cup C$ iff $A \preceq B$.

We find Independence plausible for governing the priorities between default rules, but see no direct arguments for or against it in (Horty, 2012). It may be seen as a generalization of the desired benchmark property *Principle I* of Brewka and Eiter (2000), discussed in Section 5.

3.5.3 Packard's Characterization Result

Packard considers Extension, Composition, and Independence in relation to total plausibility orders on finite, consistent sets of statements, and shows that the criteria jointly characterize a unique total order.[18] In Section 3.6, we use Packard's Theorem to provide the preorder $\preceq$ used in defining proper* scenarios. Subsequent analyses will not use the explicit definition of $\preceq^L$ mentioned in Packard's Theorem, but it may be found in the Appendix on page 40. In short, $\preceq^L$ is lexicographic with $\leq$-priority as main component and number of highest, 2nd highest... etc. defaults as secondary. Section 3.6.2 shows properties of the preorder version used in analyses.

Theorem 2 (Packard) *Let $(X, \leq)$ be a finite, totally ordered set. Then $\preceq^L$ is the unique total order on 2^X that satisfies Extension, Composition, and Independence with respect to $\leq$.*

Definition 14 For any $(X, \leq)$ finite, totally ordered set, call $(2^X, \preceq^L)$ the **Packard order** of $(X, \leq)$.

3.6 Proper* Scenarios from Priority Preorders

Given a default theory $\Delta = (W, D, \leq)$, Packard's Theorem may be applied to select a scenario used for belief formation, but only if $\leq$ is total—which is not generally assumed. However, Packard's result can be exploited to obtain a natural preorder on 2^D from a priority preorder.

solved: Heiner and Packard (1983) generalize Packard (1981)'s constructions to situations involving downright implausible statements. Their generalization requires additional definitions. Second, as nothing forces an agent to pick a set with negative-priority defaults instead of picking the set with those removed, we do not think that ordering also in accordance with negative priorities is of utmost importance. Yet, we thought the issue deserved a remark.

[18] The formulation below is based more directly on (Heiner and Packard, 1984), as Packard (1981) did not explicitly consider Extension.

3.6.1 Packard Preorders: Preorder on 2^X from Preorder on X

The idea behind inducing a preorder $(2^X, \preceq)$ from a preorder $(X, \leq)$ is that the latter can be seen as a family of total orders $(X_i, \leq_i)_{i\in I}$, on each of which we can apply Packard's Theorem, thereby inducing a family of total Packard orders $(2^{X_i}, \preceq_i)_{i\in I}$. Taking the union of this family, we obtain a preorder on 2^X. However, this preorder does not generally respect Independence. Therefore, *we additionally require that the partial order extends the subset relation.* Beyond salvaging Independence, this requirement is in line with Composition's implication that $A \succ \emptyset$ for any non-empty set A, i.e., that no defaults have downright negative priority. Intuitively, then, adding a default to A should, if anything, raise its priority. Thereby, we obtain the smallest preorder that satisfies the conditions Extension, Composition, and Independence.

Definition 15 Let $(X, \leq)$ be a finite, preordered set, and let $(X_i, \leq_i)_{i\in I}$ be the finite family of maximal chains[19] of $(X, \leq)$, with Packard orders $(2^{X_i}, \preceq_i)_{i\in I}$. With $\sqsubseteq$ the subset relation on the power set of X, call

$$(2^X, \preceq) \text{ with } \preceq := \sqsubseteq \cup \bigcup_{i\in I} \preceq_i$$

the **Packard preorder** of $(X, \leq)$.

Theorem 3 *Let $(X, \leq)$ be a finite preorder. Then $(2^X, \preceq)$, the Packard preorder of $(X, \leq)$, is the smallest preorder on 2^X that satisfies Extension, Composition, and Independence with respect to $\leq$.*

In proving Theorem 3, we use the following Lemma:

Lemma 1 *Let $(X, \leq)$ be a finite preorder with maximal chains $(X_i, \leq_i)_{i\in I}$ with Packard orders $(2^{X_i}, \preceq_i)$. Then if $A, B \subseteq X_i \cap X_j$ for some $i, j \in I$, then $A \preceq_i B$ iff $A \preceq_j B$.*

The proofs of both are found in Appendix 2.

3.6.2 Properties of Packard Preorders

As the Packard preorder $(2^X, \preceq)$ is induced by the preorder $(X, \leq)$, it is itself often not total. By extending the subset relation, it does however produce natural chains through 2^X. Items 1 and 2 of the Lemma below are adapted from (Packard, 1981):

Lemma 2 *Let $(2^X, \preceq)$ be the Packard preorder of $(X, \leq)$ and let $A, B, C, D \in 2^X$. Then*

1. *For any non-empty A, $\emptyset \prec A$.*
2. *If A is a proper subset of B, then $A \prec B$.*
3. *If for all $x \in A$, $x < b$ for some $b \in B$, then $A \prec B$.*

[19] A chain is a totally ordered subset of a partially ordered set; a chain is maximal if it is not a proper subset of any other chain. With $(X_i, \leq_i)_{i\in I}$ being the finite family of maximal chains of $(X, \leq)$, it follows that $(\cup_{i\in I} X_i, \cup_{i\in I} \leq_i) = (X, \leq)$.

Again, the proof may be found in Appendix 2.

Remark 4 Packard preorders are not in general as rich as Packard orders. Beyond those of Lemma 2, (Packard, 1981) gives three properties of the Packard order, which do not hold in the preorder case:

1. If $A \succeq C$ and $B \succeq D$, then $A \cup B \succeq C \cup D$, where $A \cap B = C \cap D = \emptyset$;
2. If $A \succ C$ and $B \succeq D$, then $A \cup B \succ C \cup D$, where $A \cap B = C \cap D = \emptyset$;
3. If $A \approx C$ and $B \approx D$, then $A \cup B \approx C \cup D$, where $A \cap B = C \cap D = \emptyset$;

Counterexamples to 1., 2., and 3., for Packard preorders may be simple. For instance, the Packard preorder $(2^X, \preceq)$ for $X = \{a, b\}$ with a and b incomparable under the preorder $\leq$ fails to satisfy 1. for $A = \{a\}$, $B = \{b\}$ and $C = D = \emptyset$.

The failure of these three properties is a consequence of using preorders on defaults. If one desires a higher degree of comparability between sets of defaults while only requiring defaults preordered, one must enforce requirements in addition to those characterizing Packard preorders, but one should be wary to not force inconsistencies with Extension, Composition, and Independence. In this vein, there is a large literature of established possibility and impossibility results, among others counting the cited (Packard, 1981) and (Heiner and Packard, 1983). See (Barberà et al., 2004) for an overview.

3.6.3 Proper Scenarios and Proper* Belief Sets for Priority Preorders*

Applying Theorem 3 to default theories, we obtain the following

Corollary 1 *Let a default theory $\Delta = (W, D, \leq)$ be given, and let $(2^D, \preceq)$ be the Packard preorder of $(D, \leq)$. Then*

$$\preceq \cap (\mathsf{CC}_\Delta \cap \mathsf{Tr}_\Delta)^2 \text{ and } \preceq \cap (\mathsf{CC}_\Delta)^2$$

are the smallest preorders on $\mathsf{CC}_\Delta \cap \mathsf{Tr}_\Delta$ and CC_Δ, respectively, that satisfy Extension, Composition, and Independence with respect to $\leq$.

As Packard's Theorem implied for total priority orders, Corollary 1 implies that for any default theory $\Delta = (W, D, \leq)$ with a priority preorder, the sets

$$\mathcal{F}_\Delta = \max_{\prec} \mathsf{CC}_\Delta \cap \mathsf{Tr}_\Delta,$$

$$\mathcal{S}_\Delta = \max_{\preceq} \{S \in \mathsf{CC}_\Delta \colon S \supseteq T \text{ for some } T \in \mathcal{F}_\Delta\}, \text{ and}$$

$$\mathcal{B}_\Delta = \{\mathsf{B}(W, S) \colon S \in \mathcal{S}_\Delta\}$$

used to define proper* scenarios ($\mathcal{S}_\Delta$) and proper* belief sets ($\mathcal{B}_\Delta$) are well-defined and non-empty. Hence, we obtain the following success property:

Proposition 1 *Beliefs are well-defined and consistent: For any default theory $\Delta = (W, D, \leq)$, the set of proper* belief states $\mathcal{B}_\Delta$ is non-empty and every belief set $\mathsf{B}(W, S) \in \mathcal{B}_\Delta$ is consistent.*

Proof $\mathcal{B}_\Delta$ is non-empty as $\mathcal{S}_\Delta$ is non-empty, which it is as $\mathcal{F}_\Delta$ is non-empty, which it is as $\mathsf{CC}_\Delta \cap \mathsf{Tr}_\Delta$ is finite and non-empty. $\mathsf{CC}_\Delta \cap \mathsf{Tr}_\Delta$ is finite as D is assumed finite, and it is non-empty as W is assumed consistent: Hence at least $\emptyset \in \mathsf{CC}_\Delta$—and $\emptyset \in \mathsf{Tr}_\Delta$ trivially. Finally, $\mathsf{B}(W, S)$ is consistent by definition.

Hence, for any default theory with consistent background information, the agent will have a non-empty set of possible, resulting proper* belief sets.

3.6.4 Inclusiveness

Finally, before turning to comparisons with Horty's framework, we show that the introduced framework is indeed inclusive, i.e., it satisfies the Minimal Default Deduction Assumption:

Proposition 2 *For any default theory $\Delta = (W, D, \leq)$ and any consistent formulas φ, ψ, if $W = \emptyset$ and $D = \{\varphi \rightsquigarrow \psi\}$, then $(\varphi \to \psi)$ is in any extension of Δ.*

Proof Let Δ be as described. Note first that $\emptyset \prec D$. Further, $\mathsf{CC}_\Delta = \mathsf{Con}_\Delta \cap \mathsf{Coh}_\Delta = \{\emptyset, D\}$ and $\mathsf{Tr}_\Delta = \{\emptyset\}$. Hence $\mathcal{F}_\Delta = \max_\preceq(\mathsf{CC}_\Delta \cap \mathsf{Tr}_\Delta) = \{\emptyset\}$, and

$$\begin{aligned}\mathcal{S}_\Delta &= \max_\preceq(\{S \in \mathsf{CC}_\Delta \colon S \supseteq T \text{ for some } T \in \mathcal{F}_\Delta\}) \\ &= \{D\}\end{aligned}$$

Finally, $\mathcal{B}_\Delta = \{\mathsf{B}(W, S) \colon S \in \mathcal{S}_\Delta\} = \{Th(D)^\to\}$, so $(\varphi \to \psi)$ is in every extension of Δ.

4 Examples and Comparison

This section compares the inclusive model with Horty's exclusive framework in their treatment of selected examples. Following analyses of the Two Guest Wedding and an example involving reasoning with necessary and sufficient conditions, where the frameworks yield importantly differing conclusions, we finally compare using multiple examples from *Reasons as Defaults.*

4.1 Formal Two Guest Wedding

Figure 1 presents a diagrammatic representation of the Two Guest Wedding example from the introduction (where you and Aunt Petunia are invited as the only guests to a relative's wedding, and are asked to RSVP).

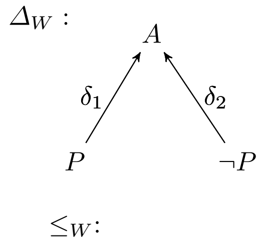

Propositions:
A: I attend. P: Petunia attends.

Δ_W consistent and coherent scenarios:
$S_0 = \emptyset$
$S_1 = \{\delta_1\}$, $S_2 = \{\delta_2\}$, $S_3 = \{\delta_1, \delta_2\}$

Fig. 1 Two Guest Wedding. *Left*: A diagrammatic illustration of the default theory $\Delta_W = (W_W, D_W, \leq_W)$. Circled propositions constitute the set of background information W_W: There are none such, as W_W is empty. A δ-labeled arrow from one formula φ to another ψ means the default $\delta = (\varphi \rightsquigarrow \psi)$ is among the available defaults D_W. The order $\leq_W$, we write out explicitly, omitting reflexive loops and links obtainable by transitive closure—nothing is specified here, as $\leq_W$ contains only reflexive loops. *Right*: Supplementary information.

Horty's conclusion. Horty's framework prescribes the unique extension $Th(\emptyset)$. This follows as $S_0 = \emptyset$ is the only proper$_1$ or proper$_2$ scenario in Δ_W. To see this, recall Definition 5 stating that a scenario $S \subseteq D$ is proper$_1$ iff it is stable in $\Delta = (W, D, <)$ iff $S = Binding(\Delta, S)$ with $Binding(\Delta, S) = (Triggered(\Delta, S) - Conflicted(\Delta, S)) - Defeated(\Delta, S)$, while it is proper$_2$ iff $S = \bigcup_{i \geqslant 0} S_i$ for some approximating sequence $(S_n)_{n \in \mathbb{N}}$ given by

$$S_0 = \emptyset,$$
$$S_{i+1} = \{\delta : \delta \in Triggered(\Delta, S_i), \delta \notin Conflicted(\Delta, S), \delta \notin Defeated(\Delta, S)\}\,.$$

That S_0 is the only proper$_1$ or proper$_2$ scenario follows as $Triggered(\Delta_W, S_k) = \emptyset$ for $k = 0, ..., 3$. Hence, on Horty's exclusive account, the agent ends up believing only the logical validities given by the unique extension

$$Th(\emptyset \cup Conclusion(S_0)) = Th(\emptyset).$$

Inclusive analysis. In contrast, on the inclusive analysis, the agent ends up with the unique belief set $\mathsf{B}(\{\delta_1, \delta_2\}) = Th(\{P \to A, \neg P \to A\})$, which entails A.

To see this, we first identify the ($\preceq$-)maximal sets of consistent, coherent, and triggered scenarios, $\max_{\preceq} \mathsf{CC}_{\Delta_W} \cap \mathsf{Tr}_{\Delta_W}$. The consistent and coherent scenarios are $\mathsf{CC}_{\Delta_\mathsf{W}} = \{S_0, S_1, S_2, S_3\}$ while the triggered are $Tr_{\Delta_W} = \{S_0\}$. So, $\mathcal{F}_{\Delta_W} = \max_{\prec} \mathsf{CC}_{\Delta_W} \cap \mathsf{Tr}_{\Delta_W} = \{S_0\}$. Second, we re-inflate the maximal consistent, coherent, and triggered scenarios to consistent and coherent scenarios, and find the maximal elements

$$\mathcal{S}_{\Delta_W} = \max_{\preceq}\{S \in \mathsf{CC}_{\Delta_W} : S \supseteq T \text{ for some } T \in \mathcal{F}_{\Delta_W}\} = \{S_3\}.$$

Here, S_3 is maximal as it is a proper superset of S_0, S_1 and S_2. That $\mathcal{S}_{\Delta_W}$ contains a unique element entails that there is a unique proper* belief set, which is

$$\mathsf{B}(\{\delta_1, \delta_2\}) = Th(\{P \to A, \neg P \to A\}).$$

Hence, we conclude that no matter whether P is the case or not, one should believe that A, which contrasts Horty's result.

Discussion. Due to its simplicity, Two Guest Wedding highlights the difference between the exclusive and inclusive conceptions of default reasoning (as represented by Horty's framework and the inclusive model, respectively).

According to Horty's exclusive framework only defaults with satisfied prerequisites (antecedents) can be used to extend beliefs beyond the given background information, i.e., only *triggered* defaults may inform beliefs, while untriggered defaults are *excluded* from serving this purpose. In the Two Guest Wedding, the exclusive conception entails that the agent *dispenses its decision* about attending: Without any background information about Petunia's attendance, neither default is triggered, and hence neither influence the agent's beliefs about whether to attend.

Under the inclusive model, however, defaults with unsatisfied prerequisites may be used to extend beliefs beyond the given background information, i.e., *untriggered* defaults may be used to this end. In the Two Guest Wedding, the inclusive conception entails that the agent *decides to attend*: Both defaults are allowed to inform beliefs, the logical closure of which implies that the agent believes it has reason to attend the wedding.

4.2 Reasoning with Necessary and Sufficient Conditions

We will now turn to a simple example involving Newtonian mechanics. The example aims to illustrate how the inclusive model and Horty's exclusive framework can diverge when it comes to reasoning with necessary and sufficient conditions.

Think of a simple physical setup where a body is under the influence of two opposing forces, keeping it at rest. To frame a simple image, consider a vase at rest on a pillar in a system so isolated that the pillar's force (which opposes gravity) is a necessary and sufficient condition for the vase to not drop.

Given this setup, let the two atomic propositions V and P and their negations have the following readings:

P:	The pillar stands.	$\neg P$:	The pillar falls.
V:	The vase stays.	$\neg V$:	The vase drops.

As background information, take $W_N = \{P\}$. From the atoms, we may form a host of default rules, for example the following:

	Defaults in D_N:	**Readings:**
δ_1	$P \rightsquigarrow V$	If the pillar stands, then the vase stays.
δ_2	$\neg P \rightsquigarrow \neg V$	If the pillar falls, then the vase drops.
δ_3	$P \rightsquigarrow \neg V$	If the pillar stands, then the vase drops.
δ_4	$\neg P \rightsquigarrow V$	If the pillar falls, then the vase stays.

Of these, learning the premise of the first two (by default) warrants a belief in the conclusion, while this is not the case for the latter two. These warrants are justified by an appeal to Newtonian mechanics: Given the setup, the first two are in line with Newtonian physics' first law of motion—the law of inertia—that an object is at rest if, and only if, the sum of forces exerted on it is zero. The latter two are not. Therefore, a default theory for the case could include δ_1 and δ_2, but not δ_3 and δ_4. Figure 2 illustrates one such theory, analyzed below.

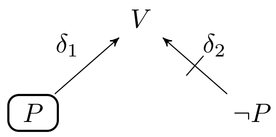

$\leq_N$: $\delta_1 =_N \delta_2$

Propositions:
P: The pillar stands. $\neg P$: The pillar falls.
V: The vase stays. $\neg V$: The vase drops.

Default rules:
$\delta_1 : \; P \rightsquigarrow V$
$\delta_2 : \neg P \rightsquigarrow \neg V$

Consistent and coherent scenarios:
$\emptyset, \quad \{\delta_1\}, \quad \{\delta_2\}, \quad \{\delta_1, \delta_2\}$

Fig. 2 Newtonian Mechanics. *Left:* A diagrammatic illustration of the default theory $\Delta_N = (W_N, D_N, \leq_N)$. Circled propositions constitute the set of background information W_N. A δ-labeled arrow from φ to ψ means the default $(\varphi \rightsquigarrow \psi)$ is in D_N, while a *crossed* δ'-labeled from φ to ψ means the default $\delta' = (\varphi \rightsquigarrow \neg\psi)$ is available. The order $\leq_N$ is stated directly, reflexivity omitted. *Right:* Supplementary information.

Horty's conclusion. Horty's exclusive framework prescribes that the Δ_N should entail the beliefs $E = Th\{P, V\} = Th\{P, P \to V\}$, as $S_1 = \{\delta_1\}$ is the only proper$_1$ (or proper$_2$) scenario.

Inclusive model. Contrary to Horty's conclusion, the inclusive model prescribes beliefs $\mathsf{B}(W_N, \{\delta_1, \delta_2\}) = Th(\{P, (P \to V), (\neg P \to \neg V)\})$.

To see this, identify the maximal sets of consistent, coherent, and triggered scenarios, $\max_{\preceq} \mathsf{CC}_{\Delta_N} \cap \mathsf{Tr}_{\Delta_N}$. The set CC_{Δ_N} is the full power set of D_N, listed in Figure 2. Of the sets in CC_{Δ_N}, only $\{\delta_1\}$ is triggered, so trivially $\max_{\preceq} \mathsf{CC}_{\Delta_N} \cap \mathsf{Tr}_{\Delta_N} = \{\{\delta_1\}\}$.

Second, we re-inflate the maximal consistent, coherent, and triggered scenarios to consistent and coherent scenarios, and find the maximal elements. The set $\{\delta_1\}$ can be inflated to $\{\delta_1, \delta_2\}$, and by Lemma 2.2, $\{\delta_1\} \prec \{\delta_1, \delta_2\}$. Hence the set of proper* scenarios is

$$\mathcal{S}_\Delta = \max_{\preceq}\{S \in \mathsf{CC}_\Delta \colon S \supseteq T \text{ for some } T \in \max_{\preceq} \mathsf{CC}_{\Delta_N} \cap \mathsf{Tr}_{\Delta_N}\} = \{\{\delta_1, \delta_2\}\},$$

entailing that there is a unique proper* belief set

$$\mathsf{B}(W_N, \{\delta_1, \delta_2\}) = Th(\{P, (P \to V), (\neg P \to \neg V)\}).$$

Discussion. The difference in the conclusions of the two frameworks may be summarized by Horty's only prescribing that one direction of (this specific instance of) the law of inertia should be believed, while the inclusive model prescribes that both directions should be believed.

This is a non-trivial difference that may influence the actions of the agent. Consider for example an agent with the goal that the vase drops, who can choose to knock over the pillar. Under the inclusive model, the agent can form a plan they believe will attain the goal: Knock over the pillar. Under the exclusive reading, this is not the case, as the agent doesn't believe that knocking over the pillar is a *sufficient* condition for the vase to drop.

4.3 Examples from *Reasons as Defaults*

To compare Horty's exclusive framework with the inclusive model, we analyze six examples from *Reasons as Defaults*. In the first five, the two frameworks reach the same conclusions. In the sixth, the frameworks differ, and this is discussed.

Each example is presented with the *strict* order $<$ on defaults given by Horty, and is followed first by Horty's conclusions, and then analyzed under the inclusive model, using instead a preorder $\leq$ on defaults, where $\leq$ is always the reflexive closure of $<$.

4.3.1 The Tweety Triangle

The first example, a classic example of non-monotonic reasoning, concerns the bird Tweety and its ability to fly. That Tweety is a bird is a reason to conclude that Tweety can fly. But if Tweety is also a penguin, then the reason to think that Tweety can fly is defeated. This illustrates how default rules can be thought of as defeasible generalizations.

The setup of the Tweety Triangle example is summarized in Figure 3. Given this setup, what are the proper scenarios, and, in extension, the potential rational belief sets?

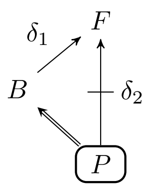

$\delta_1 < \delta_2$

Propositions:
B : Tweety is a bird. F : Tweety flies.
P : Tweety is a penguin.

All scenarios:
$\emptyset$,
$S_1 = \{\delta_1\}$, $S_2 = \{\delta_2\}$,
$S_3 = \{\delta_1, \delta_2\}$

Fig. 3 The Tweety Triangle, cf. Horty (2012, p. 23-25, 32-33). The double arrow shows that $(P \rightarrow B)$ is in the background information.

Horty's conclusion. Horty finds it intuitive that the agent should only endorse the default $\delta_2 = (P \rightsquigarrow \neg F)$, and thus only the scenario $S_2 = \{\delta_2\}$—which is also the unique proper$_{\{1,2\}}$ scenario (Horty, 2012, p. 23–25, 32–33). The agent thus reaches the conclusion that Tweety cannot fly with the belief set

$$Th(\{P, P \to B, \neg F\}).$$

By logical closure, this set includes $(P \to \neg F)$.

Inclusive analysis. The inclusive model reaches the same conclusion as Horty's. First, we identify the maximal sets of consistent, coherent and triggered scenarios. All scenarios are triggered, but S_3 is incoherent, so by Extension it follows that $\delta_1 < \delta_2$, which implies that $\max_{\preceq} \mathsf{CC}_\Delta \cap \mathsf{Tr}_\Delta = \{\{\delta_2\}\}$. Second, as we cannot consistently and coherently inflate $\{\delta_2\}$, we obtain that

$$\mathcal{S}_\Delta = \max_{\preceq}\{S \in \mathsf{CC}_\Delta \colon S \supseteq T \text{ for some } T \in \max_{\preceq} \mathsf{CC}_\Delta \cap \mathsf{Tr}_\Delta\} = \{\{\delta_2\}\},$$

entailing that there is a unique proper* belief set

$$\mathsf{B}(W, \{\delta_2\}) = Th(\{P, P \to B, P \to \neg F, \neg F\}),$$

in agreement with Horty's framework.

Exemplifying non-monotonic reasoning. To illustrate that the inclusive model indeed performs non-monotonic reasoning, consider the Tweety Triangle, but without the background information that Tweety is a penguin. Hence, let $\nabla = (W', D, \leq)$ be as in Figure 3, but with only $W' = \{B, P \to B\}$. In this case, only the scenario $\{\delta_1\}$ is triggered, entailing that $\mathcal{S}_\nabla = \{\{\delta_1\}\}$, implying the proper* belief set $\mathsf{B}(W', \{\delta_1\}) = Th(\{B, P \to B, B \to F, F\})$. The agent thus believes F. This belief was not held in Δ which contained additional background information, thereby illustrating non-monotonic reasoning.

4.3.2 The Nixon Diamond

Another well-known example concerns conflicting information about the former US president Nixon. The so-called "Nixon Diamond" is summarized in Figure 4. This example illustrates the difficulties of drawing an unambiguous conclusion when one is presented with conflicting information: That Nixon is a Quaker constitutes a reason to believe that he is a pacifist, while Nixon being Republican provides a reason to believe that he is not.

Horty's conclusion. Horty's framework reflects the seemingly insolvable conflict between the two defaults δ_1 and δ_2, as the Nixon Diamond has exactly two proper$_{\{1,2\}}$ scenarios, namely $S_1 = \{\delta_1\}$ and $S_2 = \{\delta_2\}$, with extensions $E_1 = Th(\{Q, R, P\})$ and $E_2 = Th(\{Q, R, \neg P\})$, respectively. Hence, the framework does not specify a unique belief set, but leaves us with an open-ended conclusion.[20]

[20] Horty (2012, p. 34–37) discusses three possible ways to deal with multiple extensions put forth in the literature. For further philosophical treatment, see (Horty, 2002).

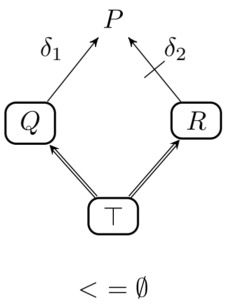

Propositions:
Q : Nixon is a Quaker.
R : Nixon is Republican.
P : Nixon is a pacifist.

All scenarios:
$\emptyset$,
$S_1 = \{\delta_1\}$, $S_2 = \{\delta_2\}$,
$S_3 = \{\delta_1, \delta_2\}$

Fig. 4 **Nixon Diamond**, cf. (Horty, 2012, p. 26–28, 34–37). The bottom $\top$-node and the double arrows are included to retain the traditional diamond shape, but are superfluous when circling the background information.

Inclusive analysis. The inclusive model reaches the same open-ended verdict as Horty's framework. All scenarios are triggered, but as S_3 is incoherent and $\leq$ is empty, $\max_{\preceq} \mathsf{CC}_{\Delta} \cap \mathsf{Tr}_{\Delta} = \{\{\delta_1\}, \{\delta_2\}\}$. Neither can be inflated consistently and coherently, so we obtain two proper* scenarios: $\mathcal{S}_{\Delta} = \{\{\delta_1\}, \{\delta_2\}\}$. Hence, the Nixon Diamond allows the two proper* belief sets

$$\mathsf{B}(W, \{\delta_1\}) = Th(W \cup \{\delta_1\}^{\rightarrow}) = Th(\{Q, R, Q \rightarrow P\}), \text{ and}$$
$$\mathsf{B}(W, \{\delta_2\}) = Th(W \cup \{\delta_2\}^{\rightarrow}) = Th(\{Q, R, R \rightarrow \neg P\}),$$

leaving us with an open-ended conclusion.

4.3.3 The Wedding of a Distant Relative

To discuss amalgamation of reasons, Horty introduces an example concerning a wedding (Horty, 2012, p. 59–61): You are invited to the wedding of a distant relative. It takes place at a busy time, making it inconvenient for you to participate. However, two aunts that you hold dear, Olive and Petunia, are also invited. You prefer to go if, and only if, both your aunts participate. Two formalizations of the example are presented in Figure 5.

Horty's conclusion. To model the example, Horty considers two default theories, Δ_1 and Δ_2 (Figure 5). In Δ_1, the inconvenience of the trip outranks all reasons for going, so the unique proper scenario relative to Δ_1 is $S_3 = \{I \rightsquigarrow \neg A\}$, entailing beliefs $E = Th\{W \cup \{\neg A\}\}$. Hence, you end up not going. According to Horty (2012, p. 60), this is counterintuitive, as the reasons are not amalgamated: Intuitively, the joy of seeing *both* aunts at the wedding should outrank the inconvenience. However, such amalgamation is not accounted for in Δ_1. Δ_2, on the other hand, takes amalgamation into account (by the conjunctive antecedent of δ_4) and as a consequence has $S_4 = \{((O \wedge P) \rightsquigarrow A)\}$ as unique proper scenario. In Δ_2, then, you will attend the wedding as both aunts do so, thus reaching the verdict Horty finds intuitive.

Δ_1 :

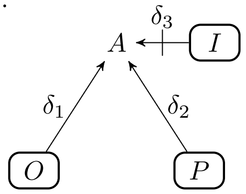

$\delta_1, \delta_2 < \delta_3$

Δ_2 :

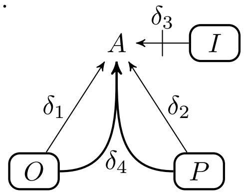

$\delta_1, \delta_2 < \delta_3 < \delta_4$

Propositions:
I: It's inconvenient. O: Olive attends.
A: I attend. P: Petunia attends.

Δ_1 **Consistent and coherent scenarios:**
$\emptyset$
$S_1 = \{\delta_1\}$, $S_2 = \{\delta_2\}$, $S_3 = \{\delta_3\}$,
$S_4 = \{\delta_1, \delta_2\}$

Δ_2 **Consistent and coherent scenarios:**
$\emptyset, S_1, \ldots, S_4$, $S_5 = \{\delta_4\}$,
$S_6 = \{\delta_1, \delta_4\}$, $S_7 = \{\delta_2, \delta_4\}$,
$S_8 = \{\delta_1, \delta_2, \delta_4\}$

Fig. 5 The Wedding of a Distant Relative, cf. Horty (2012, p. 59–61). The converging double arrow in Δ_2 illustrates that the default $\delta_4 = (O \wedge P \rightsquigarrow A)$ with amalgamated reasons is available.

Inclusive analysis. The inclusive model provides the same results as Horty's, in both cases. In both cases, all scenarios are triggered, which entails that proper* scenarios are simply the maximal consistent and coherent ones, i.e., $\mathcal{S}_{\Delta_i} = \max_{\preceq} \mathsf{CC}_{\Delta_i}$, for $i \in \{1, 2\}$.

For Δ_1, $\mathsf{CC}_{\Delta_1} = \{\emptyset, S_1, \ldots, S_4\}$, cf. Figure 5. The $\preceq$-maximal element of $\{\emptyset, S_1, \ldots, S_4\}$ is S_3: both $S_4 \succ S_1$ and $S_4 \succ S_2$ as S_4 is a superset of both S_1 and S_2 (cf. Lemma 2.2), and $S_3 \succ S_4$ as $\delta_3 > x$ for all $x \in S_4$ (cf. Lemma 2.3). The unique proper* belief set is therefore $\mathsf{B}(W, S_3) = Th(\{I, O, P, I \to \neg A\}) = Th(\{I, O, P, \neg A\})$, as in Horty's analysis.

For Δ_2, $\mathsf{CC}_{\Delta_2} = \{\emptyset, S_1, \ldots, S_8\}$, cf. Figure 5. The $\preceq$-maximal element of $\{\emptyset, S_1, \ldots, S_8\}$ is S_8, which is a superset of S_1, S_2, and $S_4, \ldots, S_7$, and $\delta_4 > x$ for all $x \in S_3$. Hence $\mathsf{B}(W, S_8) = Th(\{I, O, P, A\})$, in line with Horty's conclusion.

Remark 5 Horty remarks (2012, p. 60) that including the default $\delta_4 = ((O \wedge P) \rightsquigarrow A)$ does not solve the problem of amalgamation, but merely relocates it, as the prioritization of δ_4 relative to δ_1 and δ_2 is handled manually. The same carries over to the inclusive model: By brute force, the problem is solved in this instance, but it is not solved in general.

4.3.4 A Control Scenario Anomaly

In relating his framework to that of Brewka (1994a; 1994b), Horty analyzes the abstract example summarized in Figure 6, which presents the two (of 16) possible scenarios favored by respectively Brewka (S_1) and Horty (S_2).

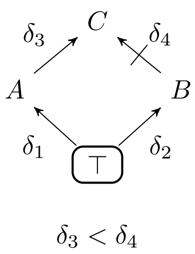

Propositions:
A, B, C, read abstractly.

Scenarios discussed by Brewka and Horty:
$S_1 = \{\delta_1, \delta_2, \delta_3\}, \quad S_2 = \{\delta_1, \delta_2, \delta_4\}$

Fig. 6 A Control Scenario Anomaly, cf. Horty (2012, p. 200–201).

Horty's conclusion. Horty argues that the default theory $\varDelta$ in Figure 6 illustrates a deficiency in the framework of Brewka, as it identifies S_1 as the proper extension of $\varDelta$. Horty favors a different conclusion, both intuitively and formally, viz., that the proper scenario is S_2, as δ_4 has a higher priority than δ_3.

Inclusive analysis. The inclusive model reaches Horty's favored conclusion. All scenarios are triggered, which entails that proper* scenarios are simply the maximal consistent and coherent ones, i.e., $\mathcal{S}_\varDelta = \max_{\preceq} \mathsf{CC}_\varDelta$. The inconsistent or incoherent scenarios are the supersets of $\{\delta_3, \delta_4\}$, so $S_1, S_2 \in \mathsf{CC}_\varDelta$. As every other scenario in $\mathsf{CC}_\varDelta$ is a strict subset of one of these, they are all non-maximal (cf. Lemma 2.2). Finally, $S_1 \prec S_2$: As $\delta_3 < \delta_4$, we have $\{\delta_3\} \prec \{\delta_4\}$ by Extension, and as $S_1 = \{\delta_1, \delta_2\} \cup \{\delta_3\}$ and $S_2 = \{\delta_1, \delta_2\} \cup \{\delta_4\}$ it follows by Independence that $S_1 \prec S_2$. Hence S_2 is the unique proper* scenario, with resulting proper* belief set $\mathsf{B}(\emptyset, S_4^{\rightarrow}) = Th(\{A, B, B \rightarrow \neg C\})$.

4.3.5 Order Puzzle

A contested example is the *Order Puzzle,* cf. Figure 7. It divides opinions, with Brewka (1994a; 1994b) advocating one solution (S_5),[21] Horty another (S_6), and Delgrande and Schaub (2000a) arguing that the underlying default theory is meaningless (Horty, 2012, p. 202).

The puzzle may be interpreted as a set of orders given to an underling by three superiors, 1, 2, and 3, the higher outranking the lower, with proper scenarios seen as the sets of orders one may follow to avoid being court-martialled for failing to follow orders while not following higher-ranking orders to justify the disobedience. See (Horty, 2012, pp. 201–206) for an extended discussion of this and other interpretations.

[21] Brewka's approach differs from Horty's by using the priority ordering to control the order of application of defaults, so higher priority defaults are applied/satisfied before lower priority defaults are considered.

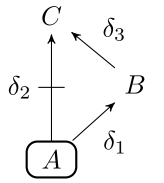

$\delta_1 < \delta_2 < \delta_3$

Propositions:
A, B, C, read abstractly.

All scenarios:
$\emptyset$,
$S_1 = \{\delta_1\}$, $S_2 = \{\delta_2\}$, $S_3 = \{\delta_3\}$,
$S_4 = \{\delta_1,, \delta_2\}$, $S_5 = \{\delta_1, \delta_3\}$, $S_6 = \{\delta_2, \delta_3\}$,
$S_7 = \{\delta_1, \delta_2, \delta_3\}$.

Fig. 7 The Order Puzzle, cf. Horty (2012, p. 201–206).

Horty's conclusion. $S_5 = \{\delta_1, \delta_3\}$ is the unique proper scenario in Horty's framework. As $\delta_1 < \delta_2 < \delta_3$, it follows that the defeating default of δ_1, i.e., δ_2, is itself defeated by δ_3. The background information $W = \{A\}$ triggers δ_1, which in turn triggers δ_3, leaving the agent with the belief set $Th(\{A, B, C\})$.

Inclusive analysis. The inclusive model again reaches the same conclusion as Horty's framework. The consistent, coherent, and triggered scenarios are $\{\emptyset, S_1, S_2, S_4, S_5\}$, of which the $\preceq$-maximal element is S_5. S_5 cannot be consistent and coherently inflated, so it is itself the unique proper* scenario, resulting in the proper* belief set $Th(\{A, B, C\})$.

4.3.6 Inappropriate Equilibria

Horty also provides an example where he finds his framework reaches an inappropriate conclusion. Again, an underling receives orders from three superiors, 1, 2, and 3, the higher outranking the lower. 1 orders the underling to see to it that A (δ_1), 2 to see to it that B (δ_2), and 3 to see to it that $\neg B$, conditional on A (δ_3).[22] The setup is summarized in Figure 8.

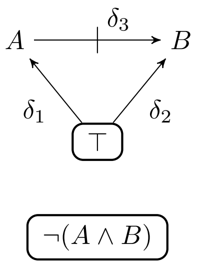

$\delta_1 < \delta_2 < \delta_3$

Propositions:
A, B, are unspecified orders from superiors.

All scenarios:
$\emptyset$,
$S_1 = \{\delta_1\},^\dagger$ $S_2 = \{\delta_2\},^\dagger$ $S_3 = \{\delta_3\}$,
$S_4 = \{\delta_1, \delta_2\}$, $S_5 = \{\delta_1, \delta_3\},^\dagger$ $S_6 = \{\delta_2, \delta_3\}$
$S_7 = \{\delta_1, \delta_2, \delta_3\}$

Fig. 8 Inappropriate Equilibrium, cf. Horty (2012, p. 206–207). The background information includes the constraint that $\neg(A \wedge B)$. The consistent, coherent and triggered scenarios are marked by a dagger (†).

[22] Horty remarks that 3's order is odd, as the background information already ensures it, but that nothing prevents a superior officer from giving weird orders.

Horty's conclusion. Horty's framework holds both S_2 and S_5 as proper scenarios. In $S_2 = \{\delta_2\}$, the underling obeys 2's order, while justifiably disobeying 1. 3's order is not triggered as it is conditional on A, which is incompatible with B. In S_5, the underling obeys 1 and 3, but justifiably disobeys 2.

Horty finds it counterintuitive that S_5 is formally deemed proper: "From an intuitive standpoint, it seems almost as if the defaults have been considered in the wrong order. The initial conflict, one wants to say, lies between [...] δ_1 and [...] δ_2. This conflict should of course be resolved in favor of [2], in which case [...] δ_3 is never even triggered [...]"(Horty, 2012, p.207). Thus, Horty finds it intuitive that matters be settled favoring only S_2 and obtaining $Th(W \cup \{B\})$, but formally also S_5—concluding $Th(W \cup \{A, \neg B\})$—is proper.

Inclusive analysis. The inclusive model agrees with neither Horty's framework nor his intuitions in this case. The only consistent, coherent, and triggered scenarios are $\emptyset, S_1 = \{\delta_1\}, S_2 = \{\delta_2\}$ and $S_5 = \{\delta_1, \delta_3\}$. Of these, only S_5 is $\preceq$-maximal. It cannot be consistent and coherently inflated, so it is also the unique proper* scenario. This is in contrast with Horty's framework and intuitions.

Discussion. As cited above, Horty remarks that in deeming S_5 proper "... it seems almost as if the defaults have been considered in the wrong order." This indicates an intuitive reading of the example that lies closer to other frameworks in the literature (see Section 5 below), where the priority between defaults determines the order in which they are to be applied, with lower-priority defaults applied after higher-priority ones. Contrary to this approach, both Horty's framework and the inclusive model are *holistic*: They evaluate the scenarios in full—and from a holistic perspective, we do find S_5 to be the unique reasonable scenario.

An anonymous reviewer notes that some may find S_6 the most intuitively reasonable. We, too, see an argument for this: S_6 is indeed the highest ordered scenario consistent with the background information. However, S_6 is not coherent: Were we to learn A, S_6 would lead to inconsistent beliefs. As mentioned in Sec. 3.3, our requirement of coherence seems in line with intuitions underlying Horty's notion of conflict, yet it may be dropped from the power set method without affecting the paper's main results, only the analyses in the examples. We must leave a systematic comparison with the effects of omitting the coherence requirement for future work.

4.4 Summing Up

The two frameworks are in thorough disagreement about the conclusions to be drawn in Two Guest Wedding and Newtonian Mechanics. This is by design: The examples are designed exactly to show differences between the exclusive and the inclusive conception by obtaining defaults that will remain untriggered. Hence, the disagreement between the inclusive model and Horty's framework is intentional.

Concerning the examples from (Horty, 2012), they all exhibit the trait that all defaults in their default theories may be triggered either directly by the background information or by conclusions by other triggered defaults. Therefore, we would not expect to see large differences in the analyses offered by exclusive and inclusive conceptions, and indeed, large differences are not found: The inclusive model reaches Horty's conclusions in all cases but one. Disagreement occurs in Inappropriate Equilibria, where the inclusive model disagrees with Horty's stated intuitions. Given the holistic approach to default reasoning taken by both Horty's framework and the inclusive model, we find this disagreement reasonable.

5 Benchmarks and Relations to the Literature

With the above section, we have illustrated that the inclusive model stays in line with most intuitions from *Reasons as Defaults*. To round off, we now situate the inclusive model in the literature more broadly. The main insights are propositions showing that the inclusive model satisfies some benchmark desiderata suggested in the literature. The inclusive model fails one principle from (Brewka and Eiter, 2000), but satisfies a weakened version, which we argue is actually better in line with the inclusive conception of default reasoning.

5.1 Relations to Poole's Logical Framework for Default Reasoning

A benchmark often discussed in the literature and already touched on above is whether a given framework is a conservative extension/refinement of Reiter's model—in the sense that a framework is a refinement of Reiter's if the extensions of a prioritized default theory found by the framework are Reiter extensions (as Horty's is when finite sets of defaults are considered, cf. Theorem 1). As mentioned in the introduction, if a default framework refines on Reiter's, then it cannot be inclusive.

As also mentioned in the introduction, the non-prioritized framework of Poole (1988) is inclusive. Here, we briefly present Poole's framework (slightly adapted to the assumptions and notation of the present paper), and show that the inclusive model offers a refinement. A main difference between Poole and Reiter's approaches is that Poole works not with a set of default rules, but a set of possible hypotheses, which may be arbitrary consistent formulas.

Definition 16 Let W and H be sets of consistent formulas, with the *background information* W assumed consistent and H called the *possible hypotheses.*

A **Poole scenario** of (W, H) is a set $S \cup W$ such that $S \subseteq H$ and $S \cup W$ is consistent.

A **Poole extension** of (W, H) is the set of logical consequences of a maximal (with respect to set inclusion) Poole scenario of (W, H).

Poole's framework is inclusive: With $W = \emptyset$, φ and ψ consistent, and $H = \{(\varphi \to \psi)\}$, the unique Poole extension of (W, H) is $Th(H)$.

In fact, Poole's framework is "more" inclusive than the inclusive model, as the latter restricts extensions (proper* belief sets) to be based on coherent scenarios. For a more exact comparison, we think of the set H of possible hypotheses as restricted to only material implications,[23] so we may consider H as the set of material implication counterparts of a set of defaults D, i.e.,

$$H = D^{\to} = \{(\varphi \to \psi) \in \mathcal{L} \colon (\varphi \rightsquigarrow \psi) \in D\}.$$

Taking this perspective, it may be seen that the inclusive model is *not* a refinement of Poole's framework, i.e., it suggests extensions that differ from those suggested by Poole's model, even when priorities are irrelevant.

To see this, let $W = \emptyset$ and $H = \{(p \to q), (p \to \neg q)\}$, for which the unique Poole extension is $Th(H)$. In the inclusive model, $Th(H)$ is not an extension (proper* belief set) of the default theory with equal priorities $\Delta = (W = \emptyset, D = \{(p \rightsquigarrow q), (p \rightsquigarrow \neg q)\}, \leq = D^2)$ because the scenario $S = D$ is not coherent. In the inclusive model, Δ has two proper* belief sets (extensions), namely $Th(\{(p \rightsquigarrow q)\})$ and $Th(\{(p \rightsquigarrow \neg q)\})$.

The difference in extensions between the inclusive model and Poole's framework seemingly stems only from the pruning to coherent scenarios in the inclusive model. When no coherence issues can arise, the inclusive model is a conservative extension of Poole's framework:

Proposition 3 *Let $\Delta = (W, D, \leq)$ be a default theory where all defaults have equal priority (i.e., $\leq = D^2$) and where $D \in \mathsf{Coh}_\Delta$. Let $H = D^{\to}$. Then:*

If B is a proper belief set of Δ, then B is a Poole extension of (W, H).*

Proof Assume $B = \mathsf{B}(W, S) \in \mathcal{B}_\Delta$. Then B is the logical closure of the Poole scenario $W \cup S^{\to}$. $W \cup S^{\to}$ is also a maximal (with respect to set inclusion) Poole scenario: $\mathsf{B}(W, S) \in \mathcal{B}_\Delta$ implies $S \in \mathcal{S}_\Delta$, and as all defaults have equal priority, the Packard preorder on D is simply the set inclusion order, entailing that we can add no further $\delta \in D$ to S without losing consistency.

We leave it as an open question whether omitting the coherence requirement from the inclusive model would produce a reasonable prioritized variant of Poole's framework.

5.2 Constrained Input/Output Logic

A second class of frameworks with inclusive tendencies is the (prioritized) *input/output logic(s)* of (Makinson and Van Der Torre, 2000, 2001; Boella and van der Torre, 2008;

[23] This is innocent as every set of formulas A is equivalent with the set of material implications $\{(\top \to \varphi) \colon \varphi \in A\}$.

Parent, 2011; Tucker, 2018).[24] To ease the comparison to these, we present them using notation assimilated to that of the present paper, and focus on logical input/outputs. See the cited papers for more general perspectives.

In the terminology of input/output logic, an expression $(\varphi \rightsquigarrow \psi)$, $\varphi, \psi \in \mathcal{L}$, is a *conditional norm*, with φ called the *input* (representing a condition) and ψ the *output* (representing what the norm holds obligatory).

Given a *generator* (a set of defaults) $D \subseteq \mathcal{L}^2$ and an *input set* $W \subseteq \mathcal{L}$, the *output set* of W under D is

$$D(W) = \{\psi \in \mathcal{L} \colon (\varphi \rightsquigarrow \psi) \in D \text{ for some } \varphi \in W\}$$

Makinson and Van Der Torre (2000) present four methods for calculating the output of a generator D given an input set W. We follow the definition found in (Parent, 2011).

Let $W \subset \mathcal{L}$ be an input set (background information) and D a set of generators (defaults). Let $D(W)$ denote the "image" of W under D, so $D(W) = \{\psi \in \mathcal{L} \colon (\varphi \rightsquigarrow \psi) \in D$ for some $\varphi \in W\}$. Call a set $V \subseteq \mathcal{L}$ a *complete* set if it is either maximally consistent or if $V = \mathcal{L}$. Define the four outputs as

$$\begin{aligned}
out_1(D,W) &= Th(D(Th(W)))\\
out_2(D,W) &= \bigcap\{Th(D(V)) \colon W \subseteq V, V \text{ is complete}\}\\
out_3(D,W) &= \bigcap\{Th(D(W')) \colon W \subseteq W' \supseteq Th(W') \supseteq D(W')\}\\
out_4(D,W) &= \bigcap\{Th(D(V)) \colon W \subseteq V \supseteq D(V), V \text{ is complete}\}
\end{aligned}$$

As noted by Makinson and Van Der Torre (2000) only out_3 and out_4 allow for reusing outputs as new inputs, and only out_2 and out_4 allow for a type of reasoning by cases. This is exemplified below.

Before moving to prioritized input/output logic, we note that out_1–out_4 each define an *exclusive* framework, in the sense that neither satisfy the Minimal Default Deduction Assumption (in its natural translation to input/output logic, i.e., where inputs take the place of background information, generators take the place of defaults, and outputs take the place of extension(s)). To see this, consider the input set $W = \emptyset$ and a singleton generator $D = \{\varphi \rightsquigarrow \psi\}$ for any consistent formulas $\varphi, \psi \in \mathcal{L}$. Assume further that $(\varphi \to \psi)$ and φ are *not* logical validities. Then, the Minimal Default Deduction Assumption (in its natural translation to input/output logic) requires that $(\varphi \to \psi)$ is in the output of D and W. However, $out_1(D, \emptyset)$–$out_4(D, \emptyset)$ all contain all and only the logical validities of $\mathcal{L}$. Hence no conservative extensions of these input/output logics are inclusive.

Building on out_1–out_4 (Parent, 2011), let's explore input/output logic with *priorities* by considering a preorder $\leq$ on D, and a set W of background information constraining deliberation. Parent and van der Torre (2013, pp. 520-522) present a

[24] We thank an anonymous reviewer for pointing us to this strand of literature.

process of deliberation that proceeds in two steps. To illustrate the process, we follow (Tucker, 2018, Sec. 2) and consider Tucker's (2018) variant of Horty's Order Puzzle. With A as a set of input-formulas such that $a, b, c, d \in \mathcal{L}$:

$\delta_1 = (a \rightsquigarrow b)$,
$\delta_2 = (a \rightsquigarrow c)$,
$\delta_3 = (b \rightsquigarrow d)$,
$D = \{\delta_1, \delta_2, \delta_3\}$,
$\delta_3 > \delta_2 > \delta_1$,
$A = \{a\}$,
$W = \{\neg(c \wedge d)\}$.

Step 1: Identify every maximal subset D' of D whose output, given A, is consistent with W; the D' is called the maxfamily(D, A, W). In the above version of the Order Puzzle, there are three members of the maxfamily(D, A, W):

$$D_1 = \{\delta_1, \delta_2\} \quad D_2 = \{\delta_1, \delta_3\} \quad D_3 = \{\delta_2, \delta_3\}.$$

Assuming that $out \in \{out_3, out_4\}$, as Tucker (2018, p. 949) does, we get outputs $out(D_1) = Th(b, c)$, $out(D_2) = Th(b, d)$, and $out(D_3) = Th(c)$.

Step 2: Remove those members of the maxfamily that are not maximal according to $\geq$. The remaining set of the preferred members of the maxfamily is called the preffamily(D, A, W). As in our inclusive model, $\geq$ is then lifted from elements of D to subsets of D. The resulting order on the power set 2^D is denoted $\succeq$ and defined by

$$D_i \succeq D_j \text{ iff for every } \delta_j \in D_j - D_i \text{ there is a } \delta_i \in D_i - D_j \text{ such that } \delta_i \geq \delta_j \quad (2)$$

with maximality relative to $\varDelta$ defined as

A set $D_i \in 2^{\varDelta}$ is $\succeq$-maximal in $\varDelta$ iff for all $D_j \in 2^{\varDelta}, D_j \succeq D_i$ implies $D_i \succeq D_j$.

Under these definitions, $D_j \subset D_i$ implies $D_j \prec D_i$. In the concrete case of the Order Puzzle, we have that D_3 is the only $\succeq$-maximal member of the maxfamily, and so it is the sole member of the preffamily.[25]

Now, while some variants of input/output logic allow for reusing outputs as new inputs, and some allow for a type of reasoning by cases, the input/output lifting operation still differs importantly from the one presented in the present paper (as part of our inclusive model): The preorder $\preceq$ on 2^D defined in (2) is not lexicographic—as Packard preorders are. Where the Packard preorder also takes into consideration the *number* of reasons presented, the input/output preorder $\preceq$ does not.

To see the lack of lexicography, consider the input/output preorder $\preceq$ on the power set of $D = \{\delta_1, \delta_2, \delta_3\}$ given $\leq$ under which δ_1,δ_2, and δ_3 are equal. For $\{\delta_1, \delta_2\}$ and $\{\delta_3\}$, we get $\{\delta_1, \delta_2\} \approx \{\delta_3\}$. That is, more good reasons are not preferred over fewer equally good reasons.

[25] Note that the basic lifting operations found in all of (Boella and van der Torre, 2008; Parent, 2011; Parent and van der Torre, 2013; Tucker, 2018) are mere notational variants of the one found in (Brass, 1991). This is explicitly acknowledged by the various authors.

As a curious similarity, it may be noted that the definition in (2) has strong similarities to the three properties of the Packard order stated in Lemma 2: (2) resembles 3. and implies both 1. and 2. Yet, as we have just seen, the orders differ in important aspects.

5.3 Benchmarks for Non-Prioritized Defaults

Though with a topic in common, many of the frameworks found in this section, and also in later proposals, rely on differing formalizations and definitions to achieve a variety of goals. Therefore, one should perhaps not expect to find *the* correct account of non-monotonic reasoning, but rather several systems with individual strengths and weaknesses (Makinson and Gärdenfors, 1991; Antoniou, 1999; Makinson, 2005). Despite—or perhaps due to—this, the literature contains a number of properties argued desirable for non-monotonic systems. One such desiderata is

Existence of Extensions : Every default theory has at least one extension.

This is not satisfied by Reiter's default logic, unless restricted to normal defaults only (Reiter, 1980).[26] The inclusive model in general satisfies the existence of extensions by Proposition 1, stating that for any default theory $\varDelta$, the set of possible belief sets $\mathcal{B}_{\varDelta}$ is non-empty.

One argument for Existence of Extensions presented by Antoniou (1999) is that the *nonexistence* of extensions in the case of Reiter's default logic violates *semi-monotonicity*, which states that while being non-monotonic in the set of background information, default theories should be monotonic in the set of available defaults—i.e., the addition of new defaults to a theory should facilitate more, not less, conclusions.[27]

Semi-Monotonicity : Let $\varDelta = (W, D)$ and $\nabla = (W, D')$ be *non-prioritized* default theories with $D \subseteq D'$. Then for every extension E of $\varDelta$, there exists an extension E' of ∇ such that $E \subseteq E'$.

Semi-monotonicity is not satisfied in general for the prioritized default theories of the inclusive model, as priorities may be used to design natural violations.[28] However, for the special case emulating non-prioritized default theories, the following is immediate.

Proposition 4 *Let $\varDelta = (W, D, \leq)$ and $\nabla = (W, D', \sqsubseteq)$ be default theories with $D \subseteq D'$, both with all defaults given equal priority. Then for every belief set $B \in \mathcal{B}_{\varDelta}$, there is a belief set $B' \in \mathcal{B}_{\nabla}$ such that $B \subseteq B'$.*

[26] This limits the expressive power of the logic, cf. (Antoniou, 1999), which also discusses alternative solutions, including Łukaszewicz' *Justified Default Logic* (Łukaszewicz, 1988) and Schaub's *Constrained Default Logic* (Schaub, 1992).

[27] Brewka (1991) argues against, and Brewka and Eiter (2000) do not see existence of extensions as essential. We return to the latter when discussing prioritized default theories below.

[28] For a violation, consider for example $W = \emptyset$, $D = \{\top \rightsquigarrow p\}$ and $D' = \{\top \rightsquigarrow p, \top \rightsquigarrow \neg p\}$ with $(\top \rightsquigarrow p) < (\top \rightsquigarrow \neg p)$. Then for $\varDelta = (\emptyset, D, D^2)$, $\mathcal{B}_{\varDelta} = \{Th(p)\}$, while for $\nabla = (\emptyset, D', \leq)$, $\mathcal{B}_{\nabla} = \{Th(\neg p)\}$.

Below we discuss a somewhat related requirement for prioritized defaults, namely Principle II of Brewka and Eiter (2000).

5.4 Benchmarks for Non-Monotonic Consequence Relations

The desirable properties of non-monotonic consequence relations has been a topic of discussion in default logic—Antonelli (2005) provides an overview. Both Antoniou (1999) and Antonelli (2005) define *semantic* consequence relations between (non-prioritized) default theories and formulas. Concerning their variant, Antonelli's main result is that it satisfies (the semantic variant of) three desiderata by Gabbay (1985), also argued for in (Gabbay et al., 1994; Stalnaker, 1994; Antonelli, 2005). Fitting Antoniou and Antonelli's definitions to the present, we obtain the following

Definition 17 Let $\Delta = (W, D, \leq)$ be a default theory. Then for any formula $\varphi \in \mathcal{L}$, say that φ is a **defeasible consequence** of Δ, written $\Delta \Vdash \varphi$, iff $\varphi \in B$ for all belief sets $B \in \mathcal{B}_\Delta$.

By additionally fitting Antonelli's desired semantic properties, we obtain a set of benchmark properties for the defined consequence relation $\Vdash$. The relation satisfies the property if the description holds for every default theory $\Delta = (W, D, \leq)$:

Reflexivity : If $\varphi \in W$, then $\Delta \Vdash \varphi$.
Cautious Monotonicity : If $\Delta \Vdash \varphi$ and $\Delta \Vdash \psi$, then $(W \cup \{\varphi\}, D, \leq) \Vdash \psi$.
Cut : If $\Delta \Vdash \varphi$ and $(W \cup \{\varphi\}, D, \leq) \Vdash \psi$, then $\Delta \Vdash \psi$.

The conjunction of Cautious Monotonicity and Cut is called *Cumulativity* in (Antoniou, 1999), described as capturing standard mathematical usage of lemmas: With φ interpreted as a lemma provable from W, Cumulativity states that we can prove the same from W and $W \cup \{\varphi\}$. (This property does not hold for Reiter's default logic, see (Antoniou, 1999) for a counterexample due to Makinson (1994)).

Of these, the consequence relation $\Vdash$ satisfies only Reflexivity. To establish this we assume that $\varphi \in W$, and let $B \in \mathcal{B}_\Delta$. By definition, $B = Th(W \cup S^{\rightarrow})$ for some set $S \subseteq D$, so $Th(W) \subseteq B$, so in particular $\varphi \in B$. As B was arbitrary, $\Delta \Vdash \varphi$.

Counterexamples to both Cautious Monotonicity and Cut can be found.[29] For a counterexample to Cautious Monotonicity: Let $\Delta = (W, D, \leq)$ with $W = \emptyset$, $D = \{\top \rightsquigarrow c, a \vee c \rightsquigarrow \neg c\}$ and $\leq = \emptyset$. Thus, we have $\Delta \Vdash c$, $\Delta \Vdash a \vee c$, and $(\{a \vee c\}, D, \leq) \nVdash c$.

For a counterexample to Cut: Let $\Delta = (W, D, \leq)$ with $W = \emptyset$, $D = \{\top \rightsquigarrow c, a \rightsquigarrow b, b \rightsquigarrow \neg a\}$ and $b \rightsquigarrow \neg a > a \rightsquigarrow b$ and $b \rightsquigarrow \neg a > \top \rightsquigarrow a$. Thus, we have $\Delta \Vdash b$, $(\{b\}, D, \leq) \Vdash \neg a$ and $\Delta \nVdash \neg a$.

[29] We thank an anonymous reviewer for providing these counterexamples.

5.5 Benchmarks for Prioritized Defaults

Using a priority ordering on defaults to guide the choice of extension in case multiple defaults are in conflict was not first introduced by Horty (2007b; 2007a; 2012), but has been studied in multiple prior works, over which (Delgrande et al., 2004) offers a survey and comparison. In short, Delgrande and Schaub (2000a; 2000b) compile priorities into standard default rules, building on an idea of Reiter and Criscuolo (1981), also developed differently by Etherington and Reiter (1983). Rintanen (1995, 1998) induces an order on extensions based on the order of defaults generating them, with preferred extensions being those maximal in the induced order. In a more uniform branch of literature, Marek and Truszczyński (1993), Brewka (1994a,b), Baader and Hollunder (1995), and Brewka and Eiter (2000), construct a sequence of extensions by using priorities to guide the defaults' order of application, with higher priority defaults applied before lower priority ones. See (Horty, 2007a, Sec. 4.2) for a comparison to this branch.

Brewka and Eiter (2000) present two principles they find intuitive for prioritized defaults, cast as general benchmarks for prioritization logics, and both satisfied by their system. They compare with related literature by arguing that the frameworks of (Marek and Truszczyński, 1993), (Brewka, 1994a), and (Baader and Hollunder, 1995), fail to satisfy the first principle, while that of (Rintanen, 1995) fails to satisfy the second (see also (Delgrande et al., 2004)). The two benchmark principles are

Principle I. Let B_1 and B_2 be two extensions of a prioritized default theory Δ, generated by the defaults $R \cup \{\delta_1\}$ and $R \cup \{\delta_2\}$, where $\delta_1, \delta_2 \notin R$, respectively. If $\delta_1$is preferred (strictly) over δ_2, then B_2 is not a *preferred extension* of Δ.

Principle II. Let E be a *preferred extension* of a prioritized default theory $\Delta = (D, W, <)$, and let δ be a default such that the prerequisite of δ is not in E. Then E is a *preferred extension* of $\nabla = (W, D \cup \{\delta\}, \sqsubset)$ whenever $\sqsubset$ agrees with $<$ on priorities among defaults in D.

Informally, Principle I essentially states that priorities between defaults should be reflected in the preference of extensions, while Principle II states that adding an untriggered default should not influence the preference of extensions.

To compare with the presented, we take a 'preferred extension' to correspond to a proper* belief set. Under this assumption, we see that

Proposition 5 *The presented inclusive model satisfies Principle I.*

Proof See Appendix 2.

By virtue of being inclusive, the inclusive model does not satisfy Principle II as Principle II is in contradiction with the Minimal Default Deduction Assumption. For example, let both φ and ψ be consistent, let $\delta_1 = (\varphi \rightsquigarrow \psi)$ and let $\Delta = (\emptyset, \emptyset, \emptyset)$ and $\nabla = (\emptyset, \{\delta_1\}, \{(\delta_1, \delta_1)\})$ be prioritized default theories. Then $Th(\emptyset)$ is an extension of Δ, but the Minimal Default Deduction Assumption requires that all extensions of ∇ contain $(\varphi \to \psi)$, contradicting Principle II. Hence, in general, Principle II is not well-suited for the inclusive conception of default reasoning.

The inclusive model does satisfies a weakened version of Principle II, essentially capturing that adding an untriggered default to a default theory Δ will not *defeat* any beliefs held based on Δ. The weakening consists in potentially allowing for *new* beliefs to be added based on this untriggered default, in contrast to Principle II, which requires that adding an untriggered default has no effect at all.

Proposition 6 *Let $B \in \mathcal{B}_\Delta$ for some default theory $\Delta = (D, W, \leq)$, and let δ be a default for which $Premise(\delta)$ is not in B. If $\nabla = (W, D \cup \{\delta\}, \sqsubseteq)$ and $\sqsubseteq$ agrees with $\leq$ on the priorities among defaults in D, then there exists a proper* belief set $B' \in \mathcal{B}_\nabla$ such that $B \subseteq B'$.*

Proof See Appendix 2.

Beyond Principles I and II, there are two points of comparison on which we would like to remark. First, the inclusive model assumes that formulas in defaults stem from a propositional language, whereas Brewka and Eiter assume a first-order language. The present approach would work without alteration if default rules were taken to consist of first-order sentences (assuming a finite set of defaults).[30] Second, in some cases, the approach of Brewka and Eiter entails that no preferred extension exists: i.e., the approach violates the Existence of Extensions benchmark discussed above. We agree with Brewka and Eiter that this violation is "less desirable" (p. 38). As stated above, the inclusive model does not face this shortcoming: as Proposition 1 shows, in the inclusive model, proper* belief sets always exist and are consistent.

6 Concluding Remarks

In this paper, we have introduced the distinction between an exclusive and inclusive conception of default reasoning, and argued why the inclusive conception is a meaningful and important form of reasoning. We showed that Reiter's default logic and its conservative extensions—including Horty's—are exclusive, while Poole's and our *inclusive model* are inclusive.

Our main contribution lies in the development of the formal framework for prioritized inclusive default reasoning. This development was philosophically guided by intuitions underlying Horty's exclusive framework, to which we have compared throughout. Beyond the comparison to Horty's framework, we have showed that the inclusive model is a conservative extension of neither the framework of Reiter (1980) nor Poole (1988), despite the latter being inclusive. We have further showed that the inclusive

[30] This finiteness assumption means that the default theory used in the proof of Brewka and Eiter's Proposition 18 is not representable in the present setting.

model satisfies a number of desirable benchmark properties from the literature, including Existence of Extensions, Semi-Monotonicity, Reflexivity, as well as Brewka and Eiter (2000)'s Principle I. As their Principle II is inherently exclusive, there is a straightforward explanation why the inclusive model does not satisfy it, but only satisfies a version of Principle II suitably weakened to the inclusive conception of default reasoning.

Several topics are now open for future work. A first topic is mainly conceptual, and relates to the general functioning of the inclusive model as presented. As we stated in the analysis of the Inappropriate Equilibria example (Sec. 4.3.6), the inclusive model—like Horty's—is holistic in its treatment of default priorities, in the sense that a scenario with overall higher priority will outweigh a scenario with lower priority, even if the defaults triggered immediately by the background information have a higher priority in a low priority scenario than the immediately triggered defaults do in the high priority scenario. Whether this holism is desirable may depend on the case at hand, and it would therefore be interesting to develop a "stepwise", non-holistic variant of the inclusive model, in the spirit of, e.g., Marek and Truszczyński (1993), Brewka (1994a,b), Baader and Hollunder (1995), and Brewka and Eiter (2000), who all use priorities to guide the defaults' order of application.

A second, related topic, is to also include adding priorities to the background information itself, so it may obtain a structure representing the entrenchment of beliefs (see, e.g., Baltag and Smets (2008)). Here, different configurations of the relationship between the two priority orders would seemingly produce a (hopefully) systematic span of default logics.

Finally, what the computational complexity of finding proper* scenarios is, what concrete algorithms would be suitable for doing so, and how these compare to similar results from other existing frameworks, are open questions.

Appendix 1: Defining Packard's Lexicographic Order

Let $\leq$ be a total order on a finite set X (of statements). Partition X according to $\leq$-equality: let $\boldsymbol{x} = \{y \in X : y \leq x \text{ and } x \leq y\}$ for all $x \in X$. Then $\boldsymbol{X} = \{\boldsymbol{x} : x \in X\}$ has finite cardinality $n \in \mathbb{N}$. Enumerate $\boldsymbol{X}$ according to $\leq$ such that

$$\boldsymbol{x}_i = \min(X \setminus \cup_{j<i} \boldsymbol{x}_j).$$

Hence $\boldsymbol{x}_1$ contains all the smallest (worst) elements of X according to $\leq$, $\boldsymbol{x}_2$ all the second-smallest, etc., and $\boldsymbol{x}_n$ all the largest (best).[31]

[31] The use of bold in our notation should not be ignored here. For it is a typographic device used to distinguish between two different mathematical objects. We could have used $[x]$ for $\boldsymbol{x}$, which is also frequently used for equivalence classes. Note further that we use the backslash notation for *set difference*: $A \backslash B$ is the set of elements in A such that they are not also in B. We could just as well have symbolized set difference using a minus sign.

Associate each subset $Y \subseteq X$ with an n-tuple

$$\overline{Y} = (|Y \cap \boldsymbol{x}_n|, |Y \cap \boldsymbol{x}_{n-1}|, ..., |Y \cap \boldsymbol{x}_1|)$$

Each $\overline{Y}$ then encodes how many best, second-best, etc. elements Y has. E.g., if $\boldsymbol{X}$ contains 4 equivalence classes and $\overline{Y} = (3, 4, 0, 1)$ then $\overline{Y}$ contains three of the best elements, four the second best, zero second-worst, and one worst element.

Each $\overline{Y}$ is an element of $\mathbb{N}^{|\boldsymbol{X}|}$. Let $\leq^L$ be the lexicographic order on $\mathbb{N}^{|\boldsymbol{X}|}$, i.e.,

$$\begin{aligned} \overline{Y} \leq^L \overline{Z} \text{ iff either } & \overline{Y}_i = \overline{Z}_i \text{ for all } i \leq n \\ \text{or } & \overline{Y}_i < \overline{Z}_i \text{ for some } i \leq n \text{ and } \overline{Y}_j = \overline{Z}_j \text{ for all } j < i. \end{aligned}$$

That is, $\overline{Y}$ is equal-to-or-worse-than $\overline{Z}$ if $\overline{Y}$ and $\overline{Z}$ are equal on all coordinates, or if there is some coordinate i on which $\overline{Y}$ is strictly worse than $\overline{Z}$ while on all more important coordinates j, $\overline{Y}$ and $\overline{Z}$ are equal.

Finally, transfer the lexicographic order $\leq^L$ on $\mathbb{N}^{|\boldsymbol{X}|}$ to the subsets of X: Define the lexicographic order $\preceq^L$ on the power set 2^X of X by

$$Y \preceq^L Z \text{ iff } \overline{Y} \leq^L \overline{Z}.$$

Then $\preceq^L$ is a total order on 2^X. It orders the sets $Y, Z \subseteq X$ in accordance with the plausibility of the statements they contain, lexicographically. Hence $Z \succeq Y$ if Z holds more of the most plausible statements than Y, or if they hold equally many, then Z holds more of the second-most plausible statements, or if they hold equally many ... etc.

Appendix 2: Proofs

Proofs

Lemma 1 *Let $(X, \leq)$ be a finite preorder with maximal chains $(X_i, \leq_i)_{i \in I}$ with Packard orders $(2^{X_i}, \preceq_i)$. Then if $A, B \subseteq X_i \cap X_j$ for some $i, j \in I$, then $A \preceq_i B$ iff $A \preceq_j B$.*

Proof Assume that $A, B \subseteq X_i \cap X_j$. Then for all $a \in A, b \in B$, $a \leq_i b$ iff $a \leq_j b$ as both $\leq_i$ and $\leq_j$ are maximal chains. Assume that $A \preceq_i B$. By the definition of $\preceq_i$, this depends only on the elements in A, the elements in B, and on $\leq_i \cap (A \times B)$. As $\leq_i$ agrees with $\leq_j$ on all elements in $A \cap B$, $(\leq_j \cap (A \times B)) = (\leq_i \cap (A \times B))$. Hence, $A \preceq_i B$ implies $A \preceq_j B$. That the converse holds follows by a symmetrical argument.

Theorem 3 *Let $(X, \leq)$ be a finite preorder. Then $(2^X, \preceq)$, the Packard preorder of $(X, \leq)$, is the smallest preorder on 2^X that satisfies Extension, Composition, and Independence with respect to $\leq$.*

Proof ***$\preceq$ is a preorder:*** The relation $\preceq$ is reflexive as it includes $\sqsubseteq$. For transitivity, assume that $A \preceq B \preceq C$. If $A \subseteq B \subseteq C$, then $A \preceq C$ as $\sqsubseteq$ is included. There are three further cases. i) If $A \subseteq B \preceq_i C$ for some $i \in I$, then also $A \preceq_i B$, as Packard orders respects subset inclusion, cf. (Packard, 1981, p. 416, item 5). Hence $A \preceq_i C$, so $A \preceq C$. ii) The case for $A \preceq_i B \subseteq C$ is similar to i). iii) If $A \preceq_i B \preceq_j C$, then—as $(X_i, \leq_i)_{i \in I}$, on which $(2^{X_i}, \preceq_i)_{i \in I}$ is build, is the family of all maximal chains—there is some $k \in I$ such that $A \preceq_k B \preceq_k C$. Hence $A \preceq_k C$, so $A \preceq C$.

$\preceq$ ***satisfies the three properties:*** That $\preceq$ satisfies Extension holds by the following argument. Assume that $x \leq y$ for some $x, y \in X$. Then there is some maximal chain $\leq_k$ such that $x \leq_k y$. Hence, by Packard's Theorem, $\{x\} \preceq_i \{y\}$. By definition $\{x\} \preceq \{y\}$ as required.[32]

That $\preceq$ satisfies Composition holds by the following argument. Assume that $A \neq \emptyset, A \cap B = \emptyset$, now if $\forall x \in A, \forall y \in B, \{x\} \prec \{y\}$, then, as each $\preceq_i$ satisfies Composition by Packard's Theorem, we have that $A \prec B$ by Lemma 1 as required.

That $\preceq$ satisfies Independence holds by the following argument. (Left-to-right direction). Assume that $A \cap B = B \cap C = \emptyset$ and $A \cup B \preceq B \cup C$. To show: $A \preceq B$. Suppose $A \cup C \subseteq B \cup C$. Then $A \subseteq B$, and so $A \preceq B$ as required. Suppose $A \cup B \preceq_i B \cup C$ for some i. Then $A \preceq_i B$ by Packard' s theorem and hence $A \preceq B$ as required. (Right-to-left direction). Assume that $A \cap B = B \cap C = \emptyset$ and $A \preceq B$. To show: $A \cup C \preceq B \cup C$. If $A \subseteq B$, we immediately have the desired result. So suppose instead that $A \preceq_i B$ for some i. Then $A \cup C \preceq_i B \cup C$ by Packart's theorem, and so $A \cup C \preceq B \cup C$.

$\preceq$ ***is the smallest:*** Assume $\precsim$ is a preorder on 2^X that satisfies Extension, Composition and Independence with respect to $\leq$. Then $\preceq$ is a subset of $\precsim$, as 1) each $\preceq_i$ is a subset of $\precsim$ and 2) $\sqsubseteq$ is a subset of $\precsim$.

1) Each $\preceq_i$ is a subset of $\precsim$ as $\precsim$ satisfies Extension, Composition and Independence with respect to $\leq$, subsets of $\precsim$ must, by Packard's Theorem, form total orders that satisfy them with respect to the maximal chains $(X_i, \leq_i)_{i \in I}$. Again by Packard's Theorem, these total orders are unique, and by assumption $(2^{X_i}, \preceq_i)_{i \in I}$. Hence each $\preceq_i$ is a subset of $\precsim$.

2) $\sqsubseteq$ is a subset of $\precsim$: Assume that for some $D, C \in 2^X$, $D \subseteq C$. If $D = C$, then $D \precsim C$ as $\precsim$ is a preorder. Assume $D \subset C$. Then $C = B \cup D$ for some non-empty B such that $B \cap D = \emptyset$. Trivially, for $A = \emptyset$, $D = A \cup D$ with $A \cap D = \emptyset$. Then Independence states that $A \cup D \precsim B \cup D$ iff $A \precsim B$, which is the case by Composition. Hence $D \precsim C$.

Lemma 2 *Let $(2^X, \preceq)$ be the Packard preorder of $(X, \leq)$ and let $A, B, C, D \in 2^X$. Then*

1. *For any non-empty A, $\emptyset \prec A$.*
2. *If A is a proper subset of B, then $A \prec B$.*
3. *If for all $x \in A$, $x < b$ for some $b \in B$, then $A \prec B$.*

[32] We thank an anonymous reviewer for supplying our argumentation on this point.

Proof 1. Follows directly from Composition. 2. Assume $A \subset B$. Then $A \preceq B$ as $\preceq$ includes the subset relation. Also, not $B \preceq A$: As $A \subset B$, $B = A \cup C$ for some non-empty C with $A \cap C = \emptyset$, while trivially, $A = A \cup \emptyset$. By Independence, $A \cup C \preceq A \cup \emptyset$ iff $C \preceq \emptyset$. But by 1., $\emptyset \prec C$. Hence, $A \prec B$. 3. If for all $x \in A$, $x < b$, then $\{x\} \prec \{b\}$ as $\preceq$ extends $\leq$. By Composition, $A \prec \{b\}$. As $\{b\} \subseteq B$, $\{b\} \preceq B$, so $A \prec B$.

Proposition 5 *The presented inclusive model satisfies Principle I.*

Proof Let $\Delta = (W, D, \leq)$ with $R \cup \{\delta_1, \delta_2\} \subseteq D$ and $\delta_2 < \delta_1$. Let $B_1 = \mathsf{B}(W, R \cup \{\delta_1\})$ and $B_2 = \mathsf{B}(W, R \cup \{\delta_2\})$. From $\delta_2 < \delta_1$ and Extension, we get that $\{\delta_2\} \prec \{\delta_1\}$. Hence, from Independence, it follows that $R \cup \{\delta_2\} \prec R \cup \{\delta_1\}$. Hence $R \cup \{\delta_2\} \notin \mathcal{S}_\Delta$, so $B_2 \notin \mathcal{B}_\Delta$.

Proposition 6 *Let $B \in \mathcal{B}_\Delta$ for some default theory $\Delta = (D, W, \leq)$, and let δ be a default for which $Premise(\delta)$ is not in B. If $\nabla = (W, D \cup \{\delta\}, \sqsubseteq)$ and $\sqsubseteq$ agrees with $\leq$ on the priorities among defaults in D, then there exists a proper* belief set $B' \in \mathcal{B}_\nabla$ such that $B \subseteq B'$.*

Proof For $\Delta = (D, W, \leq)$, assume $B = \mathsf{B}(W, S) \in \mathcal{B}_\Delta$ and that $\mathsf{B}(W, S)$ does not entail φ. Let $\delta = (\varphi \rightsquigarrow \psi)$ and let $\nabla = (D', W, \sqsubseteq)$ with $D' = D \cup \{\delta\}$ and such that $\leq\, = \,\sqsubseteq \cap D$. If $S' = S \cup \{\delta\}$ is inconsistent or incoherent, then it is irrelevant to the fact that $S \in \mathcal{S}_\nabla$. Hence also $B \in \mathcal{B}_\nabla$, establishing the desired. Hence assume $S' \in \mathsf{CC}_\nabla$. The scenario S is a superset of some $T \in \max_\preceq(\mathsf{CC}_\Delta \cap \mathsf{Tr}_\Delta)$. By assumption, δ is not triggered in ∇, so $\mathsf{Tr}_\nabla = \mathsf{Tr}_\Delta$. Hence, $T \in \max_\preceq(\mathsf{CC}_\Delta \cap \mathsf{Tr}_\Delta)$ implies $T \in \max_\preceq(\mathsf{CC}_\nabla \cap \mathsf{Tr}_\nabla) = \mathcal{F}_\nabla$. As $S' \succ S$ cf. Lemma 2.2, S' is in $\mathcal{S}_\nabla = \max_\preceq(\{S \in \mathsf{CC}_\nabla : S \supseteq T$ for some $T \in F_\nabla\})$. I.e., S' is proper* in ∇. Hence $B' = \mathsf{B}(W, S') \in \mathcal{B}_\nabla$. As $S \subseteq S'$, $\mathsf{B}(W, S) \subseteq \mathsf{B}(W, S')$, as desired.

References

Andersen, F. J. (2024a). *Logical Disagreement: An Epistemological Study.* PhD thesis, The University of St Andrews.

Andersen, F. J. (2024b). Modeling Deep Disagreement in Default Logic. *Australasian Journal of Logic*, 21:47–63.

Antonelli, A. (2005). *Grounded Consequence for Defeasible Logic.* Cambridge University Press.

Antoniou, G. (1999). A Tutorial on Default Logics. *ACM Computing Surveys (CSUR)*, 31(4):337–359.

Baader, F. and Hollunder, B. (1993). How to Prefer More Specific Defaults in Terminological Default Logic. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI 1993)*, pages 669–674.

Baader, F. and Hollunder, B. (1995). Priorities on Defaults with Prerequisites, and their Application in Treating Specificity in Terminological Default Logic. *Journal of Automated Reasoning*, 15(1):41–68.

Baltag, A. and Smets, S. (2008). A Qualitative Theory of Dynamic Interactive Belief Revision. In Bonanno, G., van der Hoek, W., and Wooldridge, M., editors, *Logic and the Foundations of Game and Decision Theory (LOFT 7)*, number Loft 7 in Texts in Logic and Games, Vol. 3, pages 9–58. Amsterdam University Press. Republished in *Readings in Formal Epistemology*, H. Arló-Costa, J. van Benthem, V. F. Hendricks (Eds.), Springer Graduate Texts in Philosophy, Springer, 2016.

Barberà, S., Bossert, W., and Pattanaik, P. K. (2004). Ranking Sets of Objects. In Barberà, S., Hammond, P. J., and Seidl, C., editors, *Handbook of Utility Theory. Volume 2: Extensions*, pages 893–977. Kluwer Academic Publishers.

Boella, G. and van der Torre, L. (2008). Institutions with a hierarchy of authorities in distributed dynamic environments. *Artificial Intelligence and Law*, 16:53–71.

Bonevac, D. (2018). Defaulting on Reasons. *Noûs*, 52(2):229–259.

Brass, S. (1991). Deduction with supernormal defaults. In Brewka, G. et al., editors, *Nonmonotonic and Inductive Logics, volume 659 of Lecture Notes in Computer Science*, pages 153–174. Springer.

Brewka, G. (1991). *Nonmonotonic Reasoning: Logical Foundations of Commonsense*. Cambridge University Press.

Brewka, G. (1994a). Adding Priorities and Specificity to Default Logic. In *Logics in Artificial Intelligence: European Workshop JELIA '94*, pages 247–260. Springer.

Brewka, G. (1994b). Reasoning about Priorities in Default Logic. In *AAAI*, volume 1994, pages 940–945.

Brewka, G. and Eiter, T. (2000). Prioritizing Default Logic. In Hölldobler, S., editor, *Intellectics and Computational Logic: Papers in Honor of Wolfgang Bibel*, volume 19 of *Applied Logic Series*, pages 27–45. Kluwer Academic Publishers.

Chrisman, M. (2015). Reasons as Defaults by John F. Horty. *Mind*, 124(495):919–924.

Delgrande, J. P. and Schaub, T. (2000a). Expressing preferences in default logic. *Artificial Intelligence*, 123(1-2):41–87.

Delgrande, J. P. and Schaub, T. (2000b). The Role of Default Logic in Knowledge Representation. In *Logic-Based Artificial Intelligence*, pages 107–126. Springer.

Delgrande, J. P., Schaub, T., Tompits, H., , and Wang, K. (2004). A Classification and Survey of Preference Handling Approaches in Nonmonotonic Reasoning. *Computational Intelligence*, 20:308–334.

Etherington, D. and Reiter, R. (1983). On Inheritance Hierarchies with Exceptions. In *Proceedings of the Third National Conference on Artificial Intelligence (AAAI-83)*, pages 104–108.

Gabbay, D. (1985). Theoretical Foundations for Nonmonotonic Reasoning in Expert Systems. In Apt, K., editor, *Logics and Models of Concurrent Systems*, pages 439–459. Springer.

Gabbay, D. M., Hogger, C. J., and Robinson, J. A. (1994). *Handbook of Logic in Artificial Intelligence and Logic Programming: Volume 3: Logic Programming*. Oxford University Press.

Ginsberg, M., editor (1987). *Readings in Nonmonotonic Reasoning.* Morgan Kaufmann Publishers.

Heiner, R. A. and Packard, D. J. (1983). More on Plausibility Orderings. *Synthese*, 55:333–337.

Heiner, R. A. and Packard, D. J. (1984). A Uniqueness Result for Extending Orders; with Application to Collective Choice as Inconsistency Resolution. *Journal of Economic Theory*, 32(1):180–184.

Horty, J. F. (2002). Skepticism and Floating Conclusions. *Artificial Intelligence*, 135(1-2):55–72.

Horty, J. F. (2007a). Defaults with Priorities. *Journal of Philosophical Logic*, 36(4):367–413.

Horty, J. F. (2007b). Reasons as Defaults. *Philosophers' Imprint*, 7:1–28.

Horty, J. F. (2012). *Reasons as Defaults.* Oxford University Press.

Horty, J. F. (2016). Reasoning with Precedents as Constrained Natural Reasoning. In Lord, E. and Maguire, B., editors, *Weighing Reasons.* Oxford University Press.

Koons, R. (2017). Defeasible reasoning. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy.* Metaphysics Research Lab, Stanford University, winter 2017 edition.

Łukaszewicz, W. (1988). Considerations on default logic: An alternative approach 1. *Computational intelligence*, 4(1):1–16.

Makinson, D. (1994). General Patterns in Non-Monotonic Reasoning. In Gabbay, D. M., Hogger, C. J., and Robinson, J. A., editors, *Handbook of Logic in Artificial Intelligence and Logic Programming: Nonmonotonic Reasoning and Uncertain Reasoning (vol. 3)*, pages 35–110. Oxford University Press.

Makinson, D. (2005). *Bridges from Classical to Nonmonotonic Logic.* King's College Publications.

Makinson, D. and Gärdenfors, P. (1991). Relations between the Logic of Theory Change and Nonmonotonic Logic. In *The Logic of Theory Change*, pages 183–205. Springer.

Makinson, D. and Van Der Torre, L. (2000). Input/output logics. *Journal of philosophical logic*, 29:383–408.

Makinson, D. and Van Der Torre, L. (2001). Constraints for input/output logics. *Journal of philosophical logic*, 30:155–185.

Marek, V. and Truszczyński, M. (1993). *Nonmonotonic Logic: Context-Dependent Reasoning.* Springer.

McCarthy, J. (1980). Circumscription – A Form of Non-Monotonic Reasoning. *Artificial intelligence*, 13(1-2):27–39.

McDermott, D. and Doyle, J. (1982). Non-Monotonic Logic 1. *Artificial Intelligence*, 13:41–72.

Moore, R. C. (1985). Semantical Considerations on Nonmonotonic Logic. *Artificial intelligence*, 25(1):75–94.

Packard, D. J. (1981). Plausibility Orderings and Social Choice. *Synthese*, 49:415—-418.

Parent, X. (2011). Moral particularism in the light of deontic logic. *Artificial Intelligence and Law*, 19:75–98.

Parent, X. and van der Torre, L. (2013). Input/output logic. In Gabbay, D. et al., editors, *Handbook of deontic logic and normative systems*, pages 499–544. College Publications.

Poole, D. (1988). A Logical Framework for Default Reasoning. *Artificial Intelligence*, 36:27–47.

Reiter, R. (1980). A Logic for Default Reasoning. *Artificial Intelligence*, 13:81–132.

Reiter, R. and Criscuolo, G. (1981). On Interacting Defaults. In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence (IJCAI-81)*, pages 270–276.

Rintanen, J. (1995). On Specificity in Default Logic. In Mellish, C., editor, *IJCAI'95*, pages 1474–1479. AAAI Press.

Rintanen, J. (1998). Lexicographic Priorities in Default Logic. *Artificial Intelligence*, 106(2):221–265.

Saka, P. (2014). Reasons as Defaults by John F. Horty. *Analysis*, 74(2):358–360.

Savage, L. J. (1954). *The foundations of statistics*. John Wiley & Sons Inc.

Schaub, T. (1992). On Constrained Default Theories. In *ECAI*, pages 304–308.

Stalnaker, R. (1994). What is a Nonmonotonic Consequence Relation? *Fundamenta Informaticae*, 21(1, 2):7–21.

Strasser, C. and Antonelli, A. (2019). Non-monotonic Logic. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2019 edition.

Touretzky, D. S. (1986). *The Mathematics of Inheritance Systems*. Morgan Kaufmann Publishers.

Tucker, D. (2018). Variable priorities and exclusionary reasons in input/output logic. *Journal of Philosophical Logic*, 47(6):947–964.