An exercise in a non-classical semantics for reasoning with incompleteness and inconsistencies

Ivan Varzinczak

CRIL, Université d'Artois & CNRS, France
CAIR, Computer Science Division, Stellenbosch University, South Africa
varzinczak@cril.fr

Abstract. Reasoning in the presence of inconsistencies and in the absence of complete knowledge has long been a major challenge in artificial intelligence. In this paper, we revisit the classical semantics of propositional logic by generalising the notion of world (valuation) so that it allows for propositions to be both true and false, and also for their truth values not to be defined. We do so by adopting neither a many-valued stance nor the philosophical view that there are 'real' contradictions. Moreover, we show that satisfaction of complex sentences can still be defined in a compositional way. Armed with our semantic framework, we define some basic notions of semantic entailment generalising the classical one and analyse their logical properties. We believe our definitions can serve as a springboard to investigate more refined forms of non-classical entailment that can meet a variety of applications in knowledge representation and reasoning.

Keywords: Logic · Knowledge representation · Non-classical semantics.

1 Introduction

The problem of dealing with information that is either contradictory or incomplete (or even both) has long been a major challenge in human reasoning. With the advent of artificial intelligence (AI), such a problem has transferred to AI-based applications and has become one of the main topics of investigation of many areas at the intersection of AI and others.

Classical logic (and its many variants) is at the heart of knowledge representation and the formalisation of reasoning in AI. Alas, the classical semantics is naturally hostile to inconsistencies and does not cope well with lack of information. This has often forced applications of classical logic into resorting to 'workarounds' or limiting its scope.

In this paper, we make the first steps in the study of a generalised semantics for propositional logic in which contradictions and incompleteness are admitted at the very basics of the semantic framework. We do so by extending the notion of world (valuation) so that it allows for propositions to be both true and false, and also for their truth values not to be defined. With that as the basis, we define a truth-functional interpretation of complex sentences and show that it is possible to reason in the presence of inconsistencies or in the absence of truth values without resorting to a dialetheist [15, 16] or a many-valued stance [2, 9].

Of particular interest to us is the notion of entailment under the proposed semantics. In that respect, we also put forward some basic notions of semantic entailment generalising the classical one and analyse their logical properties. We point out which properties and reasoning patterns usually considered in classical logic are preserved and which ones are violated by the new definitions.

We believe that our constructions and preliminary results can serve as a springboard with which to investigate further the role of inconsistencies and incompleteness as basic logical notions and their applications in knowledge representation. The paper concludes by mentioning the next steps in this direction.

2 Basic semantic framework

Let $\mathcal{P}$ be a finite set of propositional *atoms*. We use $p, q, \ldots$ as meta-variables for atoms. Propositional sentences are denoted by $\alpha, \beta, \ldots$, and are recursively defined in the usual way:

$$\alpha \ ::= \top \mid \bot \mid \mathcal{P} \mid \neg\alpha \mid \alpha \wedge \alpha \mid \alpha \vee \alpha \mid \alpha \rightarrow \alpha$$

(As usually done, we see the bi-conditional $\alpha \leftrightarrow \beta$ as an abbreviation for the sentence $(\alpha \rightarrow \beta) \wedge (\beta \rightarrow \alpha)$.) We use $\mathcal{L}$ to denote the set of all propositional sentences built up according to the above grammar.

When it comes to the semantics, our point of departure is the thesis, endorsed by some in the past [17], that *functionality* and *totality* in the assignment of truth values is at the source of many of the criticisms voiced against classical reasoning. Therefore, here we shall forego both properties in the definitions of valuation (or world) and satisfaction below.

A (propositional) *world* is a relation on $\mathcal{P} \times \{0, 1\}$, where 1 represents truth and 0 falsity. $\mathcal{U} \stackrel{\text{def}}{=} 2^{\mathcal{P} \times \{0,1\}}$ is the *universe* (set of all worlds). We use $w, v, u, \ldots$, possibly with primes, to denote worlds. Whenever it eases the presentation, we shall represent valuations as sequences of atoms (e.g., p) and barred atoms (e.g., $\bar{p}$), with the understanding that the presence of p in w abbreviates $(p, 1) \in w$, while the presence of a barred atom $\bar{p}$ in w abbreviates $(p, 0) \in w$. Given $\mathcal{P}$, it is easy to verify that $|\mathcal{U}| = 2^{2 \times |\mathcal{P}|}$.

A *classical* world is a world that is a total function. A *partial* world is a world that is a partial function. An *absurd* world is a world w for which there is $p \in \mathcal{P}$ s.t. $w(p) = \{0, 1\}$. (Notice that there are some absurd worlds not assigning any truth value to at least one atom, i.e., there are absurd worlds that are extensions of partial ones.) A *possible* world is a world that is not absurd. A *non-classical* world is a world that is not classical. We can partition $\mathcal{U}$ into a set of classical worlds, denoted $\mathcal{U}_{\text{cl}}$, a set of partial worlds $\mathcal{U}_{\text{pa}}$, and a set of absurd worlds $\mathcal{U}_{\text{ab}}$, and such that $\mathcal{U} = \mathcal{U}_{\text{cl}} \cup \mathcal{U}_{\text{pa}} \cup \mathcal{U}_{\text{ab}}$. Moreover, with $\mathcal{U}_{\text{p}} \stackrel{\text{def}}{=} \mathcal{U}_{\text{cl}} \cup \mathcal{U}_{\text{pa}}$ we denote the set of possible worlds, and with $\mathcal{U}_{\text{nc}} \stackrel{\text{def}}{=} \mathcal{U}_{\text{pa}} \cup \mathcal{U}_{\text{ab}}$ the non-classical ones.

The idea behind our notions of absurd and partial worlds is certainly not new. For instance, they have been explored by Rescher and Brandom [17], even though their technical construction and proposed semantics for the propositional connectives is different from ours (see below).

The definition of satisfaction of a sentence $\alpha \in \mathcal{L}$ by a given world must be redefined w.r.t. the classical tradition because the standard, classical, notion of satisfaction is an 'all-or-nothing' notion. By that we mean $\Vdash$, although called satisfaction *relation*, is usually defined as a (recursive) *function*, which does not allow for a sentence $\alpha \in \mathcal{L}$ to be true and false at w at the same time, or to be just unknown at a given world. Therefore, just as valuations (worlds) in our setting are no longer total functions, which allows for propositions to be true and false simultaneously, or even completely unknown, so will the satisfaction relation be. The crux of the matter become then how to define $\Vdash$ in such a way as to allow α and $\neg\alpha$ to be true at a given w and how to express this in terms of the subsentences and propositions therein. Of course, if we want to give an intuitionistic flavour to our logic, $\Vdash$ should also be such that the values of both α and $\neg\alpha$, for some $\alpha \in \mathcal{L}$, are unknown.

Previous attempts to redefine satisfaction in order to allow for sentences and their negations to be true at the same time have stumbled upon the *exportation principle* [5, 6]: the notion of satisfaction, which sits at the meta-level, gets 'contaminated' by inconsistencies at the object level, resulting in a contradictory sentence being satisfied if and only if it is not satisfied (a contradiction 'exported' to the meta-level). Here we aim at avoiding precisely that by ensuring the meta-language we work with remains as much as possible consistent. (We shall come back to this matter at the end of the present section.)

The *satisfaction relation* is a binary relation $\Vdash \subseteq (\mathcal{U} \times \mathcal{L}) \times \{\mathbf{0}, \mathbf{1}\}$, where $\mathbf{1}$ and $\mathbf{0}$ are 'meta-truth' values standing for, respectively, 'is true' and 'is false' at world w. The intuition is that α is *true* at w if $\mathbf{1} \in \Vdash (w, \alpha)$, and *false* at w if $\mathbf{0} \in \Vdash (w, \alpha)$. Of course, to be false does not mean not to be true, and the other way round. (Notice that $\{\mathbf{0}, \mathbf{1}\}$ and $\{0, 1\}$ are not to be conflated, in the same vein as 'valuation' and 'satisfaction' are not synonyms.[1]) In that respect, a given $\alpha \in \mathcal{L}$ can have a truth value at a given world $w \in \mathcal{U}$ ($\mathbf{0}$, $\mathbf{1}$, or both, which, importantly, is *not* an extra truth value) or *none* at all (when $\Vdash (w, \alpha) = \emptyset$, which is not an extra truth value either). Before we provide the definition of satisfaction of a sentence in terms of its subsentences, we discuss its expected behaviour w.r.t. the sentence's main connective. For the sake of readability, in Tables 1–4, we represent $\{\mathbf{0}\}$, $\{\mathbf{1}\}$ and $\{\mathbf{0}, \mathbf{1}\}$ as, respectively, $\mathbf{0}$, $\mathbf{1}$ and $\mathbf{01}$, and the lack of truth value as $\emptyset$—which, again, is not an extra truth value. So, in the referred tables, $\mathbf{0}$ and $\mathbf{1}$ are read as usual, $\emptyset$ is read as 'has no value', and $\mathbf{01}$ reads 'is true and false', i.e., both truth values apply. (The reader not convinced by some of the entries in Tables 1–4 is invited to hold on until we state some of the validities holding under the notion of satisfaction we are about to define.)

[1] We could, in principle, also have used $\{0, 1\}$ in the definition of satisfaction, but here we shall adopt the (possibly superfluous) stance that truth of a fact within the 'actual' world and truth of a sentence at given worlds are notions sitting at different levels, or, at the very least, are notions of subtly different kinds.

Satisfaction of $\neg\alpha$ *at* w*:* If α is just true or false at a given world, then its negation should behave as usual, i.e., as a 'toggle' function. By the same principle (applied twice), if α happens to be true and false, then its negation should be false and true. For the odd case, namely when α has no (known) value, a legitimate question to ask is 'can we know more about the negation of a fact than we know about the fact itself?' A cautious answer would be 'no'. Table 1 summarises these considerations.

α	$\neg\alpha$
$\emptyset$	$\emptyset$
1	**0**
0	**1**
01	**01**

Table 1. Semantics of $\neg$.

Satisfaction of $\alpha \wedge \beta$ *at* w*:* As usual, a conjunction is true if both conjuncts are true, and false if at least one of them is false. A sentence being allowed more than one truth value means this rule can be triggered twice, with the conjunction being true and false at the same time. Not knowing the truth value of a conjunct may cast doubt on that of the conjunction; but if one of the conjuncts is surely false, so should their conjunction be. Table 2 captures all the possibilities.

α	β	$\alpha \wedge \beta$
$\emptyset$	$\emptyset$	$\emptyset$
$\emptyset$	**0**	**0**
$\emptyset$	**1**	$\emptyset$
$\emptyset$	**01**	**0**
0	$\emptyset$	**0**
0	**0**	**0**
0	**1**	**0**
0	**01**	**0**
1	$\emptyset$	$\emptyset$
1	**0**	**0**
1	**1**	**1**
1	**01**	**01**
01	$\emptyset$	**0**
01	**0**	**0**
01	**1**	**01**
01	**01**	**01**

Table 2. Semantics of $\wedge$.

Satisfaction of $\alpha \vee \beta$ at w: For $\vee$, the truth of at least one of the disjuncts is enough to enforce that of their disjunction, and that even if one does not know the truth value of all disjuncts. If both disjuncts are false, so should be their disjunction. And, once more, the last two rules may apply simultaneously. Table 3 summarises the behaviour of disjunction within our framework.

α	β	$\alpha \vee \beta$
$\emptyset$	$\emptyset$	$\emptyset$
$\emptyset$	**0**	$\emptyset$
$\emptyset$	**1**	**1**
$\emptyset$	**01**	**1**
0	$\emptyset$	$\emptyset$
0	**0**	**0**
0	**1**	**1**
0	**01**	**01**
1	$\emptyset$	**1**
1	**0**	**1**
1	**1**	**1**
1	**01**	**1**
01	$\emptyset$	**1**
01	**0**	**01**
01	**1**	**1**
01	**01**	**01**

Table 3. Semantics of $\vee$.

Satisfaction of $\alpha \to \beta$ at w: Traditionally, a conditional is deemed true if its antecedent is false or its consequent is true, and is false if the antecedent is true while the consequent is false. Allowing for a sentence to be true and false at the same time carries along a double application of this principle. Nevertheless, not knowing the value of the antecedent or that of the consequent may cast doubt on the application of the above principle: it should only be applied when a *sufficient* condition is met. For instance, if α is unknown and β is true and false, all we know is that β is true, and therefore the conditional is true. Similarly, if α is true and false, and β is unknown, all we know is that α is false, which is enough for the conditional to be verified. This discussion is summarised by Table 4.

In the light of the above considerations, satisfaction of sentences from $\mathcal{L}$ at worlds from $\mathcal{U}$ can be defined recursively as specified below.

Definition 1 (Satisfaction). *The* ***satisfaction relation*** *on sentences of $\mathcal{L}$ is defined in a compositional way as follows:*

- $\Vdash (w, \top) = \{\mathbf{1}\}$;
- $\Vdash (w, \bot) = \{\mathbf{0}\}$;

α	β	$\alpha \to \beta$
$\emptyset$	$\emptyset$	$\emptyset$
$\emptyset$	**0**	$\emptyset$
$\emptyset$	**1**	**1**
$\emptyset$	**01**	**1**
0	$\emptyset$	**1**
0	**0**	**1**
0	**1**	**1**
0	**01**	**1**
1	$\emptyset$	$\emptyset$
1	**0**	**0**
1	**1**	**1**
1	**01**	**01**
01	$\emptyset$	**1**
01	**0**	**01**
01	**1**	**1**
01	**01**	**01**

Table 4. Semantics of $\to$.

- $((w, p), \mathbf{1}) \in \Vdash$ *iff* $(p, 1) \in w$*;*
- $((w, p), \mathbf{0}) \in \Vdash$ *iff* $(p, 0) \in w$*;*
- $((w, \neg\alpha), \mathbf{1}) \in \Vdash$ *iff* $((w, \alpha), \mathbf{0}) \in \Vdash$*;*
- $((w, \neg\alpha), \mathbf{0}) \in \Vdash$ *iff* $((w, \alpha), \mathbf{1}) \in \Vdash$*;*
- $((w, \alpha \wedge \beta), \mathbf{1}) \in \Vdash$ *iff both* $((w, \alpha), \mathbf{1}) \in \Vdash$ *and* $((w, \beta), \mathbf{1}) \in \Vdash$*;*
- $((w, \alpha \wedge \beta), \mathbf{0}) \in \Vdash$ *iff either* $((w, \alpha), \mathbf{0}) \in \Vdash$ *or* $((w, \beta), \mathbf{0}) \in \Vdash$*;*
- $((w, \alpha \vee \beta), \mathbf{1}) \in \Vdash$ *iff either* $((w, \alpha), \mathbf{1}) \in \Vdash$ *or* $((w, \beta), \mathbf{1}) \in \Vdash$*;*
- $((w, \alpha \vee \beta), \mathbf{0}) \in \Vdash$ *iff both* $((w, \alpha), \mathbf{0}) \in \Vdash$ *and* $((w, \beta), \mathbf{0}) \in \Vdash$*;*
- $((w, \alpha \to \beta), \mathbf{1}) \in \Vdash$ *iff* $((w, \alpha), \mathbf{0}) \in \Vdash$ *or* $((w, \beta), \mathbf{1}) \in \Vdash$*;*
- $((w, \alpha \to \beta), \mathbf{0}) \in \Vdash$ *iff* $((w, \alpha), \mathbf{1}) \in \Vdash$ *and* $((w, \beta), \mathbf{0}) \in \Vdash$*.*

Henceforth, we shall sometimes use $w \Vdash \alpha$ as an abbreviation for $((w, \alpha), \mathbf{1}) \in \Vdash$, and $w \nVdash \alpha$ otherwise, i.e., when either $\Vdash (w, \alpha) = \emptyset$ or $\Vdash (w, \alpha) = \{\mathbf{0}\}$.

It is easy to see that the *satisfiability problem* for $\mathcal{L}$, i.e., the problem of determining whether for a given $\alpha \in \mathcal{L}$ there is $w \in \mathcal{U}$ s.t. $w \Vdash \alpha$, is at least as hard as the satisfiability problem for classical propositional logic. As for the upper bound, we already know there are $2^{2 \times |\mathcal{P}|}$ valuations to be checked. Furthermore, notice that the size of a world in $\mathcal{U}$ is, in the worst case, double the size of a classical propositional valuation and therefore verifying that a world does indeed satisfy a sentence α can be done in polynomial time. This establishes satisfiability of $\mathcal{L}$-sentences as an NP-COMPLETE problem [4, 14].

One of the difficulties brought about by some approaches allowing for the notion of contradiction-bearing worlds is the 'contamination' of the meta-language with inconsistencies via the exportation principle [5]. In such frameworks, when assessing the truth of a sentence of the form $\alpha \wedge \neg\alpha$, the corresponding definition of satisfaction leads to the fact α is true *and* is not true (in the meta-language). Let us see how things go in the light of our definitions above. Assume that $w \Vdash \alpha \wedge \neg\alpha$, i.e., $((w, \alpha \wedge \neg\alpha), \mathbf{1}) \in \Vdash$. By the semantics for conjunction, we have both $((w, \alpha), \mathbf{1}) \in \Vdash$ and $((w, \neg\alpha), \mathbf{1}) \in \Vdash$, which, according to Table 2, happens only if $\Vdash (w, \alpha) = \Vdash (w, \neg\alpha) = \{\mathbf{0}, \mathbf{1}\}$. As far as one can tell, the latter is not an antinomy in the meta-language.

It is worth noting that, for every $\alpha, \beta \in \mathcal{L}$, the truth conditions for both $\neg(\alpha \wedge \beta)$ and $\neg\alpha \vee \neg\beta$ are the same, and so are those for $\neg(\alpha \vee \beta)$ and $\neg\alpha \wedge \neg\beta$. In that respect, the De Morgan laws are preserved under our non-standard semantics. It turns out the truth conditions for $\alpha \rightarrow \beta$ and $\neg\alpha \vee \beta$ also coincide, and therefore the connective for material implication is superfluous. If we admit material implication and the constant $\perp$, then it is negation that becomes superfluous, as the semantics for $\neg\alpha$ and $\alpha \rightarrow \perp$ coincide. Moreover, both $\top$ and $\perp$ can be expressed in terms of each other with the help of negation.

As it is usually done, we can talk of *validity*, i.e., truth at all worlds under consideration. Let $\alpha \in \mathcal{L}$; we say that α is a *classical validity* (alias, α is *classically valid*), denoted $\models \alpha$, if $w \Vdash \alpha$ for every $w \in \mathcal{U}_{\text{cl}}$. (Obviously, our notion of classical validity and that of tautology in classical propositional logic coincide.) We say α is a *partial validity* (alias, α is *partially valid*), denoted $\models_{\text{pa}} \alpha$, if $w \Vdash \alpha$ for every $w \in \mathcal{U}_{\text{pa}}$. As it turns out, partial validity is quite a stingy notion of validity: the only valid sentences in $\mathcal{U}_{\text{pa}}$ are $\top$ and those of the form $\alpha \vee \top$ (or $\alpha \rightarrow \top$). Lest this be seen as a drawback of allowing partial valuations in our framework, here we claim this is rather a reminder that partiality may have as consequence that even the principles of the underlying logic do not hold at worlds where the truth of some or all the propositions is unknown. Moreover, the presence of partial worlds automatically rules out some validities from classical logic that are usually seen as unjustifiable in non-classical circles. Among those are the law of *excluded middle* ($\alpha \vee \neg\alpha$ is a validity), the law of *non-contradiction* ($\neg(\alpha \wedge \neg\alpha)$), the principle of *double negation* ($\neg\neg\alpha \leftrightarrow \alpha$), and the *principle of explosion* ($(\alpha \wedge \neg\alpha) \rightarrow \beta$, for every $\beta \in \mathcal{L}$). Neither of these is a partial validity, as it can easily be checked. (We shall come back to these principles later on.) Finally, we can also define the notion of *absurd validity* (denoted $\models_{\text{ab}}$), which amounts to satisfaction by all absurd worlds. Some examples of absurd validities within our framework are $\top$ and $(p_1 \wedge \neg p_1) \vee \ldots \vee (p_n \wedge \neg p_n)$, for $p_i \in \mathcal{P}$, $i = 1, \ldots, n$, with $n = |\mathcal{P}|$.

Before we carry on, let us consider the so-called 'paradoxes' of material implication [18], namely the sentences $\alpha \rightarrow (\beta \rightarrow \alpha)$, $(\alpha \rightarrow \beta) \vee (\beta \rightarrow \alpha)$, and $\alpha \rightarrow (\beta \vee \neg\beta)$, which are all classical validities. We already know that none of them is a partial validity, and it should not take too much effort to verify that they are not absurd validities either. Furthermore, according to our semantics, the three sentences above do not have the same meaning, i.e., their truth tables are pairwise different from each other. This means that our semantics can distinguish between these syntactically different sentences, which the classical semantics cannot do.

From the discussion above, one can see that not all classical tautologies are preserved in our semantic framework, which is just as intended. In that respect, our framework provides the semantic foundation for an *infra-classical* logic curtailing certain classical conclusions that are often perceived as problematic. In the next section, once we have defined a few forms of entailment, we shall also assess the validity and failure of some commonly considered rules of inference or reasoning patterns from classical logic.

3 Basic forms of entailment and their properties

Let $\alpha \in \mathcal{L}$. With $[\![\alpha]\!]_{\text{cl}} \stackrel{\text{def}}{=} \{w \in \mathcal{U}_{\text{cl}} \mid w \Vdash \alpha\}$ we denote the *classical models* of α; with $[\![\alpha]\!]_{\text{pa}} \stackrel{\text{def}}{=} \{w \in \mathcal{U}_{\text{pa}} \mid w \Vdash \alpha\}$ we denote the *partial models* of α, and with $[\![\alpha]\!]_{\text{ab}} \stackrel{\text{def}}{=} \{w \in \mathcal{U}_{\text{ab}} \mid w \Vdash \alpha\}$ we denote the *absurd models* of α. The *possible models* of α is the set $[\![\alpha]\!]_{\text{p}} \stackrel{\text{def}}{=} [\![\alpha]\!]_{\text{cl}} \cup [\![\alpha]\!]_{\text{pa}}$, whereas the *non-classical models* of α is the set $[\![\alpha]\!]_{\text{nc}} \stackrel{\text{def}}{=} [\![\alpha]\!]_{\text{pa}} \cup [\![\alpha]\!]_{\text{ab}}$. Finally, the *models* of α *tout court* is the set $[\![\alpha]\!] \stackrel{\text{def}}{=} [\![\alpha]\!]_{\text{cl}} \cup [\![\alpha]\!]_{\text{nc}}$.

The choice of which family of models one wants to work with gives rise to different notions of entailment or logical consequence. Below are those to which we shall give consideration in the present paper.[2]

Definition 2 (Classical entailment). *α **classically entails** β, denoted $\alpha \models \beta$, if $[\![\alpha]\!]_{\text{cl}} \subseteq [\![\beta]\!]$.*

Clearly, $\models$ here coincides with standard classical entailment. We shall use $\alpha \equiv \beta$ as an abbreviation for both $\alpha \models \beta$ and $\beta \models \alpha$.

Definition 3 (Possible entailment). *α **possibly entails** β, denoted $\alpha \models_{\text{p}} \beta$, if $[\![\alpha]\!]_{\text{p}} \subseteq [\![\beta]\!]$.*

Definition 4 (General entailment). *α **generally entails** β, denoted $\alpha \models_{\text{g}} \beta$, if $[\![\alpha]\!] \subseteq [\![\beta]\!]$.*

We use $\equiv_{\text{p}}$ and $\equiv_{\text{g}}$ to denote logical equivalence in the above class of models.

From the definitions at the beginning of the present section, one expects $\models_{\text{g}} \subseteq \models_{\text{p}} \subseteq \models$. This is indeed the case. In the remaining of the section, we look at specific properties of each form of non-classical entailment. Below, $\models_*$ denotes either of $\models_{\text{p}}$ or $\models_{\text{g}}$.

Possible entailment satisfies the following rule of Generalised Modus Ponens:

$$(\text{MP}) \quad \frac{\alpha \models_{\text{p}} \beta, \quad \alpha \models_{\text{p}} \beta \to \gamma}{\alpha \models_{\text{p}} \gamma}$$

[2] We do not rule out the remaining combinations; space and time constraints prevent us from assessing them here.

To see why, let $w \in [\![\alpha]\!]_{\mathrm{p}}$; then, by definition of $\models_{\mathrm{p}}$, we have both $w \in [\![\beta]\!]_{\mathrm{p}}$ and $w \in [\![\beta \to \gamma]\!]_{\mathrm{p}}$, i.e., $w \Vdash \beta$ and $w \Vdash \beta \to \gamma$. Then we have $((w, \beta), \mathbf{1}) \in \Vdash$ and $((w, \beta \to \gamma), \mathbf{1}) \in \Vdash$. Since $w \in \mathcal{U}_{\mathrm{p}}$, this only holds if $((w, \gamma), \mathbf{1}) \in \Vdash$ (cf. Table 4), i.e., $w \Vdash \gamma$, and therefore $w \in [\![\gamma]\!]$.

General entailment, on the other hand, fails (MP). Let $\alpha = \beta = p \wedge \neg p$, and let $\gamma = \bot$. It can easily be verified that $p \wedge \neg p \models_{\mathrm{g}} p \wedge \neg p$, $p \wedge \neg p \models_{\mathrm{g}} (p \wedge \neg p) \to \bot$, and $p \wedge \neg p \not\models_{\mathrm{g}} \bot$.

Neither possible nor general entailment satisfies the rule of Contraposition:

$$\text{(CP)} \quad \frac{\alpha \models_* \beta}{\neg\beta \models_* \neg\alpha}$$

For a counter-example, let $\alpha = p \wedge \neg p$ and $\beta = q$. We do have $p \wedge \neg p \models_{\mathrm{p}} q$, because $[\![p \wedge \neg p]\!]_{\mathrm{p}} = \emptyset$, but $\neg q \not\models_{\mathrm{p}} \neg(p \wedge \neg p)$, as there are partial worlds satisfying $\neg q$ which assign no value to p. The same counter-example applies to general entailment.

That both possible and general entailment satisfy the following Monotonicity and Cut rules can easily be verified and we omit the proofs here:

$$\text{(Mon)} \quad \frac{\alpha \models_* \beta, \quad \alpha' \models_* \alpha}{\alpha' \models_* \beta} \qquad\qquad \text{(Cut)} \quad \frac{\alpha \wedge \beta \models_* \gamma, \quad \alpha \models_* \beta}{\alpha \models_* \gamma}$$

Possible entailment also satisfies the so-called 'easy' half of the deduction theorem:

$$\text{(EHD)} \quad \frac{\alpha \models_{\mathrm{p}} \beta \to \gamma}{\alpha \wedge \beta \models_{\mathrm{p}} \gamma}$$

To witness, assume $\alpha \models_{\mathrm{p}} \beta \to \gamma$. By Mon, we have $\alpha \wedge \beta \models_{\mathrm{p}} \beta \to \gamma$. We also have $\alpha \wedge \beta \models_{\mathrm{p}} \beta$. By MP, we conclude $\alpha \wedge \beta \models_{\mathrm{p}} \gamma$.

To see that general entailment fails EHD, let again $\alpha = \beta = p \wedge \neg p$, and let $\gamma = \bot$. We have $p \wedge \neg p \models_{\mathrm{g}} (p \wedge \neg p) \to \bot$, but $p \wedge \neg p \wedge p \wedge \neg p \not\models_{\mathrm{g}} \bot$.

Both possible and general entailment fail the 'hard' half of the deduction theorem:

$$\text{(HHD)} \quad \frac{\alpha \wedge \beta \models_* \gamma}{\alpha \models_* \beta \to \gamma}$$

Indeed, we have $p \wedge q \models_{\mathrm{p}} q$, but for some w such that $w(p) = \{1\}$ and $w(q) = \emptyset$, we have $\Vdash (w, q \to q) = \emptyset$, and therefore $p \not\models_{\mathrm{p}} q \to q$. The case for $\models_{\mathrm{g}}$ is analogous.

The following Transitivity rule is a consequence of Monotonicity and is satisfied by both possible and general entailment:

$$\text{(Tran)} \quad \frac{\alpha \models_* \beta, \quad \beta \models_* \gamma}{\alpha \models_* \gamma}$$

Not surprisingly, possible and general entailment fail the First Disjunctive rule below, just as classical entailment does:

$$\text{(Disj1)} \quad \frac{\alpha \models_* \beta \vee \gamma}{\alpha \models_* \beta \ \text{ or } \ \alpha \models_* \gamma}$$

To witness, we have $p \rightarrow q \models_* \neg p \vee q$, but neither $p \rightarrow q \models_* \neg p$ nor $p \rightarrow q \models_* q$ holds.

The Second Disjunctive rule below is satisfied by both possible and general entailment, courtesy to the fact that, for every α, β, $[\![\alpha]\!] \subseteq [\![\alpha \vee \beta]\!]$.

$$\text{(Disj2)} \quad \frac{\alpha \vee \beta \models_* \gamma}{\alpha \models \gamma \ \text{ or } \ \beta \models_* \gamma}$$

Possible and general entailment satisfy the rule of Generalised Disjunction below:

$$\text{(GD)} \quad \frac{\alpha \models_* \beta, \quad \gamma \models_* \delta}{\alpha \vee \gamma \models_* \beta \vee \delta}$$

To see why, let $w \in [\![\alpha \vee \gamma]\!]_{\mathrm{p}}$, i.e., $w \Vdash \alpha \vee \gamma$, and then either $w \Vdash \alpha$ or $w \Vdash \gamma$, or both. If $w \Vdash \alpha$, then $w \Vdash \beta$, and therefore $w \Vdash \beta \vee \delta$ (cf. Table 3). If $w \Vdash \gamma$, then $w \Vdash \delta$, and hence $w \Vdash \beta \vee \delta$. The proof for $\models_g$ is analogous.

Both possible and general entailment fail the rule of Proof by Contradiction below:

$$\text{(PC)} \quad \frac{\alpha \wedge \neg\beta \models_* \gamma \wedge \neg\gamma}{\alpha \models_* \beta}$$

Indeed, we have $p \wedge \neg(q \vee \neg q) \models_{\mathrm{p}} q \wedge \neg q$. Nevertheless, there is $w \in \mathcal{U}_{\mathrm{p}}$ such that $w(p) = \{1\}$ and $w(q) = \emptyset$, from which follows $\Vdash (w, q \vee \neg q) = \emptyset$, and therefore $p \not\models_{\mathrm{p}} q \vee \neg q$. (For the case of $\models_{\mathrm{g}}$, just let $w \in \mathcal{U}_{\mathrm{ab}}$.)

The following rule of Proof by Cases is violated by both possible and general entailment:

$$\text{(D)} \quad \frac{\alpha \wedge \neg\beta \models \gamma, \quad \alpha \wedge \beta \models \gamma}{\alpha \models \gamma}$$

Indeed, we have $p \wedge \neg q \models_{\mathrm{p}} q \vee \neg q$ and $p \wedge q \models_{\mathrm{p}} q \vee \neg q$, but $p \not\models_{\mathrm{p}} q \vee \neg q$. The same can be shown for $\models_g$.

We conclude the present section with an assessment of a few more properties of classical reasoning. The first one is Disjunctive Syllogism: $(\neg\alpha \vee \beta) \wedge \alpha \models \beta$ (equivalently, $\beta \wedge \alpha \models \beta$). It is easy to verify that possible entailment satisfies both forms. General entailment, on the other hand, fails the first form and satisfies the second. For the latter, a quick check via Table 2 suffices. For the former, consider $w \in \mathcal{U}_{\mathrm{ab}}$ such that $w(p) = \{0, 1\}$ and $w(q) = \emptyset$. Then we have $w \Vdash (\neg p \vee q) \wedge p$, but $w \not\Vdash q$. So, once more, our absurd worlds allow for distinguishing two classically equivalent sentences.

We have seen that the semantics for both $\alpha \rightarrow \beta$ and $\neg\alpha \vee \beta$ coincide, which has as consequence that both possible and general entailment satisfy the Duns Scott law: $\neg\alpha \models_* \alpha \rightarrow \beta$. On the other hand, Reductio ad Absurdum, i.e., $\neg\alpha \rightarrow (\beta \wedge \neg\beta) \models \alpha$, is only satisfied by possible entailment: $\neg\alpha \rightarrow (\beta \wedge \neg\beta)$ is just $\alpha \vee (\beta \wedge \neg\beta)$, which possibly entails α. As a counter-example for general entailment, it suffices to verify that $\neg p \rightarrow (q \wedge \neg q) \not\models_g p$.

It is not hard to see that $\bot \models_* \alpha$, for every $\alpha \in \mathcal{L}$, but notice that we do *not* have $p \wedge \neg p \models_g \alpha$, for every $\alpha \in \mathcal{L}$. Indeed, $\bot \leftrightarrow (p \wedge \neg p)$ is not a validity in our framework, just as $\top \leftrightarrow (p \vee \neg p)$ is not one either. One of the consequences of the latter is that in our framework not all validities are omnigenerated. Nevertheless, it does hold that $\alpha \models_* \top$, for any $\alpha \in \mathcal{L}$, since $[\![\top]\!] = \mathcal{U}$.

4 Concluding Remarks

In this paper, we have revisited the semantics of classical propositional logic. We started by generalising the notion of propositional valuation to that of a world that may also admit inconsistencies, or lack of information, or both. We have seen that our definition of valuation remains suitable for a compositional interpretation of the truth value of a complex sentence, and that without appealing to either a dialetheist stance or the use of more than two truth values. In particular, we have seen that assuming a compositional semantics does not lead to difficulties brought about by the exportation principle, which is one of the limitations of previous approaches sharing our motivations. We have also seen that the adoption of a more general semantics, which brings in a higher number of possible states of affairs to consider, does not increase the computational complexity of the satisfiability problem for the underlying language. We have then explored some basic notions of entailment within our semantic framework and compared them against many of the properties or reasoning patterns that are usually considered in formal logic. Some of these are lost, as expected, while some are preserved.

Immediate next steps for further investigation include (*i*) an exploration of other definitions of semantic entailment, their properties and respective suitability (or not) for effective non-classical reasoning; (*ii*) a comparison with standard systems of paraconsistent logic and other existing non-classical logics; (*iii*) the identification of scenarios for potential applications of the framework here introduced, and (*iv*) the definition of a basic proof method, probably based on semantic tableaux [8], that can serve as the backbone of more elaborate proof systems for extensions of our semantic framework.

Further future work stemming from the basic definitions and results here put forward can branch in several fruitful directions. A non-exhaustive list includes: (*i*) investigating a generalisation of the satisfiability problem [4] and the adaption of existing approaches and optimised techniques for its solution; (*ii*) extending the Kripkean semantics of modal logics [7] to also allow for 'impossible' or 'incomplete' worlds, or the set-theoretic semantics of description logics [3] to capture 'incoherent' or 'partially-known' objects or individuals in formal ontologies, before considering a move to full first-order logic, and (*iii*) revisiting the areas of belief change [1, 10] and non-monotonic reasoning [13] in artificial intelligence, also benefitting from their semantic constructions in order to define more refined forms of entailment in our setting.

With regard to the last point above, extra structure may be added to $\mathcal{U}$, e.g. in the form of a *preference relation* or a *ranking function* [11, 12], in order to distinguish worlds according to their level of *logical plausibility*. For instance, absurd worlds can be deemed as the least plausible ones, and possible worlds can be further ranked given extra information (e.g. a knowledge base and its signature of relevant atomic propositions). The associated entailment relation becomes then parameterised by such levels and should give rise to a consequence relation with more interesting properties than those of the basic entailments we have seen.

Acknowledgements

I would like to thank the anonymous referees for their comments and helpful suggestions. This work was partially supported by the National Research Foundation (NRF) of South Africa.

References

1. Alchourrón, C., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. Journal of Symbolic Logic **50**, 510–530 (1985)
2. Arieli, O., Avron, A.: The value of the four values. Artificial Intelligence **102**, 97–141 (1998)
3. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P. (eds.): The Description Logic Handbook: Theory, Implementation and Applications. Cambridge University Press, 2 edn. (2007)
4. Ben-Ari, M.: Mathematical Logic for Computer Science. Springer-Verlag, 3 edn. (2012)
5. Berto, F., Jago, M.: Impossible Worlds. Oxford University Press (2019)
6. Berto, F., Jago, M.: Impossible Worlds. In: Zalta, E.N. (ed.) The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University, fall 2018 edn. (2018)
7. Chellas, B.: Modal logic: An introduction. Cambridge University Press (1980)
8. D'Agostino, M., Gabbay, D., Hähnle, R., Posegga, J. (eds.): Handbook of Tableau Methods. Kluwer Academic Publishers (1999)
9. Fitting, M.: Kleene's three valued logics and their children. Fundamenta Informaticae **20**(1), 113–131 (1994)
10. Hansson, S.: A Textbook of Belief Dynamics: Theory Change and Database Updating. Kluwer Academic Publishers (1999)
11. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. Artificial Intelligence **44**, 167–207 (1990)
12. Lehmann, D., Magidor, M.: What does a conditional knowledge base entail? Artificial Intelligence **55**, 1–60 (1992)
13. Makinson, D.: Bridges from Classical to Nonmonotonic Logic, Texts in Computing, vol. 5. King's College Publications (2005)
14. Papadimitriou, C.: Computational Complexity. Addison-Wesley (1994)
15. Priest, G.: An Introduction to Non-Classical Logic: From If to Is. Cambridge Introductions to Philosophy, Cambridge University Press, 2 edn. (2001)

16. Priest, G., Berto, F., Weber, Z.: Dialetheism. In: Zalta, E.N. (ed.) The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University, fall 2018 edn. (2018)
17. Rescher, N., Brandom, R.: The Logic of Inconsistency. A Study in Non-Standard Possible-Worlds Semantics and Ontology. APQ Library of Philosophy, Oxford, Basil Blackwell (1979)
18. Swart, H.: Logic: Mathematics, Language, Computer Science and Philosophy, vol. 1. Peter Lang (1993)