

NONCLASSICAL PROBABILITY, CONVEX HULLS, AND DUTCH BOOKS

forthcoming in *Episteme*

Michał Gil Sanchez, Zalán Gyenis, Leszek Wroński*

Jagiellonian University, Kraków

April 25, 2022

Abstract

We report a solution to an open problem regarding the axiomatization of the convex hull of a type of nonclassical evaluations. We then investigate the meaning of this result for the larger context of the relation between rational credence functions and nonclassical probability. We claim that the notions of bets and Dutch Books typically employed in formal epistemology are of doubtful use outside the realm of classical logic, eventually proposing two novel ways of understanding Dutch Books in nonclassical settings.

1 INTRODUCTION

We would like to believe true propositions and avoid believing false ones. In formal epistemology it is typical to represent an agent's belief state by means of a credence function which assigns real numbers—usually taken from the $[0, 1]$ real segment—to propositions. Ideally, then, we would want our credences in true propositions to equal 1, and our credences in false propositions to equal 0. However, due to our cognitive and evidential limitations, leading to the typical human condition of imperfect information, we have to settle for something else. It is one of the basic tenets of formal epistemology that credences of a rational agent are weighted means of classical truth evaluations; this is the same as saying that they belong to the “convex hull” of classical evaluations, and, seen from yet another angle, it means that these credences satisfy the classical Kolmogorov probability axioms.

All this assumes, usually implicitly, that the underlying logic is classical. How does the situation change if this assumption is removed? At first glance, it might be intuitive to hold e.g. that if an agent knows, say, that some proposition A has the truth value $1/2$, their credence in A should be $1/2$. And more generally, just as classical probabilities are weighted means of classical evaluations, the term “nonclassical probability” can be taken to refer, to a first approximation, to a weighted mean of nonclassical ones.¹ Paris (2005) initiated the project of axiomatizing nonclassical probabilities, tying it also to the issue of Dutch-bookability: for

*Correspondence: leszek.wronski@uj.edu.pl.

¹This is the intended meaning of the term employed in a portion of the literature to which this paper aims to contribute, e.g. in Williams (2016) and Bradley (2017). Williams (2012a) uses the term “generalized probabilities”. In the paper which serves as a foundation for this research, Paris (2005) writes about “probabilities” and “analogs of probability functions for non-standard propositional logics”; Williams (2016) uses the term “nonclassical probabilities” to refer to such functions; earlier suggesting in (Williams, 2012b) that they are functions which are “nonclassically coherent”. The project of axiomatizing nonclassical probabilities thusly conceived should not be confused with those aiming to give some alternative formalization of the probability calculus, via Popper functions, Rényi axioms, etc.

a variety of nonclassical settings, credences satisfy Paris' axioms if and only if they are not Dutch-bookable (and thus avoid at least one source of irrationality).

However, as already mentioned, the above is only a first approximation of the issues to be discussed. This is because a proposition's truth value may by itself not be enough to fix the degree of belief an omniscient agent should invest in it. For example, various three-valued logics, employing the truth values of T , F and, say, O , may come with different interpretations of the "Other" value O . For some logics, if a proposition has the truth value O , then an omniscient, or even just a rational agent should better not have a credence in that proposition at all, and so their credence function should be appropriately gappy. According to other logics the credence should be defined in such a case, but should be set to 0; that is, for an omniscient agent a proposition with truth value O is as good as a false one. According to still other logics, in such cases the agent's credence should be 1; for still others, 0.5 (specific examples will be given later). These "credences an omniscient agent should invest in a proposition given its truth value" are called *cognitive loads*; it is typically assumed that each truth value has "its" cognitive load², and so each valuation, considered as a vector of truth values, generates a vector of cognitive loads. With this notion in hand, given a nonclassical (propositional) logic, instead of axiomatizing the convex hull of the valuations permitted by that logic, we can turn to axiomatizing the convex hulls of the corresponding vectors of cognitive loads (which we will call "cognitive evaluations" later on³). Members of that set are, after all, weighted means of credences of omniscient agents. It is in that vein that Paris' project has been continued by J.R.G. Williams and S. Bradley.

The paper is structured as follows. In the next Section we recall the basic results in this field which will come of use later. In Section 3 we report a solution to the problem (posed in Williams (2016)) of axiomatizing the convex hull of the set of cognitive evaluations of an intuitive 3-valued calculus called Symmetric Logic. We then turn to the issue of why one should be interested in such nonclassically probabilistic credences at all. As already mentioned, it is typical to use some sort of Dutch-book-related considerations to claim that such credences avoid at least one source of irrationality. In Section 4 we argue that, if such arguments are to be fruitful, they need to use different notions than the ones employed heretofore. We suggest two proposals going in that direction. Both of them call for modifying the formal details of what is to be called a Dutch Book: in the first case (labelled "real Dutch-bookability") this results from employing a different notion of credence, and in the second ("truth-value Dutch-bookability", Subsection 4.1.1) from using a different notion of bet altogether.

Shortly speaking, it is important to distinguish between the logical goal of axiomatizing certain convex hulls and the epistemological goal of arguing that satisfying those axioms is a matter of rationality. Section 3 reports on the former; Section 4 on the latter.

2 CONVEX HULLS OF EVALUATIONS AND NONCLASSICAL PROBABILITY

The project of axiomatizing convex hulls of nonclassical evaluations has its ultimate goal in describing credence functions which would be rational in nonclassical settings. Let us postpone the discussions of rationality to Section 4 and focus now on the aforementioned convex hulls.

We shall begin by stating Paris' initial result in the framework introduced in the recent handbook article by Williams [2016]. It requires a finite propositional language \mathcal{L} with the set of sentences $Sent_{\mathcal{L}}$ built using a set of connectives which includes \vee and \wedge (but possibly also other ones, including modalities). \mathbb{V} is a certain subset of the set of all functions from $Sent_{\mathcal{L}}$

²That is: gappy frameworks have, it seems, not been investigated yet in the literature on cognitive loads.

³The term has been introduced in Bradley (2017), and we intend to keep the intuitions behind the concept.

into a finite nonempty set of truth values TV . Williams introduces the term “cognitive load”:⁴ a cognitive load of a truth value is the supposed “ideal cognitive state” associated with it; in other words, it is the degree of belief an omniscient agent should invest in a proposition having that truth value. In the classical case, cognitive loads directly correspond to truth values 1 (true) and 0 (false), while in the general case the cognitive load function c is an arbitrary function from TV into $[0, 1]$. For any valuation V we can speak of “its” *cognitive evaluation* $c_V : Sent_{\mathcal{L}} \rightarrow [0, 1]$ defined as, for any $\varphi \in Sent_{\mathcal{L}}$, $c_V(\varphi) = c(V(\varphi))$. Williams’s idea is, in the context of some logic, to inquire about the convex combinations of something else than valuations. The reason for this is that two different logics, defined on the same language and having the same set TV of truth values, may give rise to exactly the same set of valuations. And yet, for example due to how the consequence relation differs between the two logics, the epistemic status of these valuations might be different.

Let us note that according to both Williams and Bradley the logics themselves are “cognitively loaded”, in that each truth value has “its” cognitive load (Williams, 2016, 255)⁵. For the purposes of stating the formal results we can therefore treat the cognitive load function to be a definitional element of the given logic. The logics are also “semantically driven”; we have our doubts as to what exactly this means⁶, but at the very least it seems to entail:

- first, that a logic expressed in a language carries with itself the information about what the possible valuations of the sentences of that language are (i.e., if a logic is semantically driven, it’s not something for which we could find different semantics: if you change the semantics, you change the logic);
- second, that truth values are categorical properties of sentences (as opposed to, say, uninterpreted formal devices used in achieving a different goal; see Field (2009));
- and third, that the logic involves a consequence relation defined so that whether a sentence entails another depends only on facts regarding valuations (and not, for example, on any syntactic considerations).

Before we state the starting result, let us note that we have made a conscious decision to depart from Paris’ original assumptions, to the effect that all languages under consideration here will include only truthfunctional operators. This is just for reasons of presentation. The main system under discussion, the Symmetric Logic, only includes operators of this kind anyway. But nonetheless, doing so allows us to simplify some statements of theorems based on Paris’ results. The languages in question have finitely many propositional variables and operators and the logics admit only finitely many possible valuations. Thanks to this we can speak of the function “ B ”⁷ as opposed to “every finite restriction of B ” being a convex combination of some vectors. More complicated variants of the relevant results, taking into the account also non-truthfunctional operators, can be provided after the inspection of Paris (2005).

Without further ado, here’s the generalized version of Paris’ theorem:

2.1 THEOREM (PARIS (2005), THEOREM 5 GENERALIZED)⁸ Fix a sentential language \mathcal{L} consisting of a finite set of propositional variables P and a finite set of logical connectives which

⁴Let us note here that Williams (2012b) uses a slightly different language to refer to essentially the same notions. We will not introduce that language in the current paper.

⁵One might be tempted by the examples in Section 3 in Williams (2016) (which are the three logics discussed below) to think that the given logic’s cognitive load function is somehow “generated” by its consequence relation. That, however, does not seem to be the intended interpretation; also, the function $[\cdot]$, which in Williams (2012b) plays the role of the cognitive load function, is defined for numerous logics and seems to have no straightforward connection to their consequence relations.

⁶For example: Bradley writes (*op. cit.*, p. 88) “it is facts about ways the truth statuses could be distributed that determine the logic”. This seems a radically strong thesis, since on a reading which seems the most natural to us it equates all logics which share truth tables; for example, the three logics KL , LP and SL discussed below. (We of course do not wish to claim that Bradley had this in mind.)

⁷To be interpreted as a function specifying degrees of belief.

⁸The crucial insight about the ‘no drop’ consequence is due to Williams (2012b).

includes \vee and \wedge . Take $Sent_{\mathcal{L}}$ to be the set of all sentences of \mathcal{L} . Let a logic be given by $L = (\mathbb{V}, \models, c)$, where valuations $V \in \mathbb{V}$ are functions from $Sent_{\mathcal{L}}$ into a finite nonempty set of truth values TV , \models is a consequence relation, and the cognitive load function c is an arbitrary function from TV to $[0, 1]$. The logic's "cognitive evaluations" are all functions $c_V : Sent_{\mathcal{L}} \rightarrow [0, 1]$ defined as, for any $\varphi \in Sent_{\mathcal{L}}$, $c_V(\varphi) = c(V(\varphi))$.

Let B be a function from $Sent_{\mathcal{L}}$ to $[0, 1]$. Then, **if**:

(*) the image of c is $\{0, 1\}$;

(**) \vee and \wedge operate classically with respect to the cognitive loads; that is, for any cognitive evaluation c_V ,

$$c_V(\varphi \vee \psi) = 0 \text{ iff } c_V(\varphi) = c_V(\psi) = 0$$

and

$$c_V(\varphi \wedge \psi) = 1 \text{ iff } c_V(\varphi) = c_V(\psi) = 1;$$

(***) the consequence operation satisfies the 'no drop' condition on the cognitive evaluations, that is,

$$\varphi \models \psi \text{ iff for any valuation } V, c_V(\varphi) \leq c_V(\psi);$$

then the following are equivalent.

- (A) B is a convex combination of the cognitive evaluations from $\{c_V \mid V \in \mathbb{V}\}$;
- (B) B satisfies the axioms below.

($\mathcal{L}1$) If $\varphi \models \psi$ then $B(\varphi) = 1$, and if $\psi \models \varphi$ then $B(\psi) = 0$,

($\mathcal{L}2$) If $\varphi \models \psi$ then $B(\varphi) \leq B(\psi)$,

($\mathcal{L}3$) $B(\varphi \vee \psi) + B(\varphi \wedge \psi) = B(\varphi) + B(\psi)$. ■

Already in three-valued cases the situation becomes non-trivial. If the set of truth values is, say, $\{T, O, F\}$ (for "true", "other", and "false"), the ideal cognitive state associated with O can be one of a number of things. If, for example, the logic dictates that O be read as "half-true", then the cognitive load of O can be naturally taken to be $1/2$. (This will be main case under discussion in the current paper.) For other logics, as we will see below, the ideal degree of belief invested in a proposition which has the truth value O might be 0 (if believing such a proposition should, according to the logic, be avoided), or 1 (if it is as belief-worthy as a true proposition). And, correspondingly, what rational credences are should depend on which logic governs the possible worlds: even if the set of valuations may be exactly the same.⁹

If we use the term " L -probabilities" for the elements of the convex hull of L 's cognitive evaluations, Theorem 2.1 says in effect that **if L satisfies (*)-(***), then L -probabilities are axiomatized by ($\mathcal{L}1$)-($\mathcal{L}3$)**. The general problem is **to give axiomatizations of M -probabilities for logics M which do not satisfy at least one of the conditions (*)-(***),** and we will consider

⁹A sidenote on generalization: Theorem 2.1 does not carry over to infinite sets P of propositional variables, not even in the classical propositional logic case. For if P is infinite, then the Lindenbaum-Tarski algebra of the classical propositional logic L is the countably generated free Boolean algebra \mathcal{B} . Each evaluation $V \in \mathbb{V}$ corresponds to an ultrafilter of \mathcal{B} . It is enough to show that there exists a (probability) function $B : \mathcal{B} \rightarrow [0, 1]$ that satisfies the axioms $\mathcal{L}1 - \mathcal{L}3$, but still B is not a convex combination of the cognitive evaluations. To this effect, take a countable partition $\{\varphi_i : i \in \mathbb{N}\}$ of \mathcal{B} and let B be an arbitrary probability function such that $B(\varphi_i) > 0$ for each $i \in \mathbb{N}$. By way of contradiction assume that

$$B = r_1 c_{V_1} + \dots + r_n c_{V_n}$$

for cognitive evaluations c_{V_k} and convex coefficients r_k . As each $c_{V_k}(\varphi_i)$ is either 1 or 0, by the homomorphism property of the V_k 's and that the φ_i 's form a partition, there must exist some i such that $c_{V_k}(\varphi_i) = 0$ for every k . But then

$$0 \neq B(\varphi_i) = r_1 c_{V_1}(\varphi_i) + \dots + r_n c_{V_n}(\varphi_i) = 0$$

is a contradiction.

here a particular case in which it is just the condition (*) that is violated: that is, the logic in question admits more than two cognitive loads.

A straightforward application of (a version of) Theorem 2.1 is noted by Williams in the context of the well-known three-valued logics *KL* (Kleene’s “strong logic of indeterminacy”) and *LP* (Priest’s “logic of paradox”). In fact, the Strong Kleene truth-tables used by them serve as the basis of probably the least complicated examples of the issue under discussion.¹⁰ The three-valued logics *KL* and *LP* along with the Kleene truth tables are introduced and discussed in detail in Section 7.3 of (Priest, 2001).

Consider, then, a sentential language \mathcal{L} consisting of a non-empty finite set P of propositional variables and the three connectives \wedge , \vee and \neg . A *valuation* V assigns to each propositional variables one the three possible truth values: T , O , and F . It is then extended to a mapping $V: \text{Sent}_{\mathcal{L}} \rightarrow \{T, O, F\}$ by the rules given by the *Kleene truth tables* as follows:

| | | | | | | | | | | | |
|----------|-----|-----|-----|--------|-----|-----|-----|--------|-----|-----|-----|
| \wedge | T | O | F | \vee | T | O | F | \neg | T | O | F |
| T | T | O | F | T | T | T | T | T | T | O | F |
| O | O | O | F | O | T | O | O | O | F | O | T |
| F | F | F | F | F | T | O | F | F | F | O | T |

The logics *KL*, *LP* and the “Symmetric Logic” *SL* use these truth tables; however, they differ in how their consequence relation \models is defined:

KL: $\varphi \models_{KL} \psi$ iff for every evaluation V we have

$$\text{if } V(\varphi) = T, \text{ then } V(\psi) = T. \quad (1)$$

LP: $\varphi \models_{LP} \psi$ iff for every evaluation V we have

$$\text{if } V(\varphi) = T \text{ or } O, \text{ then } V(\psi) = T \text{ or } O. \quad (2)$$

SL: $\varphi \models_{SL} \psi$ iff for every evaluation V we have

$$\text{if } V(\varphi) = T, \text{ then } V(\psi) = T; \text{ and} \quad (3)$$

$$\text{if } V(\varphi) = O, \text{ then } V(\psi) = T \text{ or } O. \quad (4)$$

That is, \models_{SL} is \models_{KL} and \models_{LP} “taken together”: $\varphi \models_{SL} \psi$ iff ($\varphi \models_{KL} \psi$ and $\varphi \models_{LP} \psi$). *KL* has no tautologies, *a fortiori* the principle of excluded middle $\varphi \vee \neg\varphi$ also fails to be one. *LP* is a paraconsistent logic, where $\varphi \wedge \neg\varphi$ is not explosive, i.e., it does not entail everything. *SL* allows us to enjoy both of these features.

In the 2016 handbook article Williams claims that the following are the cognitive loads of the three logics:

| | | | |
|--|-----|-----|-----|
| Truth value: | T | O | F |
| The <i>KL</i> cognitive load function c^{KL} : | 1 | 0 | 0 |
| The <i>LP</i> cognitive load function c^{LP} : | 1 | 1 | 0 |
| The <i>SL</i> cognitive load function c^{SL} : | 1 | 1/2 | 0 |

Note that, indeed, if we grant this assumption, then the logics *KL* and *LP* satisfy the conditions (*)-(***) and thus Theorem 2.1 can be applied to them directly.¹¹ Williams notes

¹⁰These examples appeared already in Williams (2016), but we include them here for the paper’s completeness and because we will eventually voice some doubts about Williams’ interpretation of them.

¹¹We are not convinced that merely the fact that the intended interpretation of “has the truth value O ” is “is both true and false” shows that $c^{LP}(O)$ should be set to 1. However, we plan to revisit this point in a future study, using the complex notion of credence introduced here in Section 4.3.

Note also the standard fact that there will be many logics under the *KL* name due to the various possibilities in which the set of variables might be chosen, which determines the set of valuations. That is, for two different sets of valuations \mathbb{V}_1 and \mathbb{V}_2 satisfying the Kleene truth tables, $(\mathbb{V}_1, \models_{KL}, c^{KL})$ and $(\mathbb{V}_2, \models_{KL}, c^{KL})$ will be different Kleene Logics, to both of which Theorem 2.1 directly applies. (All this applies to *LP* and *SL* too, of course.)

further that “it is a matter of hard graft to see whether similar completeness results can be derived for settings that fail the Parisian conditions (one representative of which is our Symmetric logic)”. As already mentioned, Paris himself extends the result so that it applies to finitely-valued Łukasiewicz logics, while Mundici (2006) achieves a similar goal for the infinitely-valued version.¹² Bradley (2017) continues the “hard graft”, covering some examples involving languages with non-truthfunctional operators.¹³ Our main goal here is to investigate how such nonclassical axioms can be argued to be requirements of rationality using Dutch Book considerations, an idea frequently mentioned, but one that has not, it seems, been so far put under sufficient scrutiny. Our point of departure will be the relatively uncomplicated example of Symmetric Logic, to which we now turn.

3 AXIOMATIZING CONVEX HULLS OF SYMMETRIC LOGIC

We have recently put forward the following solution of the problem of axiomatizing SL -probabilities in (Gil Sanchez et al., ms):

3.1 THEOREM (GIL SANCHEZ ET AL., MS). Let a sentential language \mathcal{L} consist of a nonempty finite set P of propositional variables and the three connectives \wedge , \vee and \neg . Take $Sent_{\mathcal{L}}$ to be the set of all sentences of \mathcal{L} . Let the logic SL be given as $(\mathbb{V}, \models_{SL}, c^{SL})$, where valuations $V \in \mathbb{V}$ are given by the Kleene truth tables. Let B be a function from $Sent_{\mathcal{L}}$ to $[0, 1]$. The following are equivalent.

- (A) B is a convex combination of the cognitive evaluations c_V^{SL} for $V \in \mathbb{V}$.
- (B) B satisfies the axioms below.

(SL1) If $\varphi \models \psi$ then $B(\varphi) \leq B(\psi)$,

(SL2) $B(\neg\varphi) = 1 - B(\varphi)$,

(SL3) $B(\varphi \vee \psi) = B(\varphi) + B(\psi) - B(\varphi \wedge \psi)$,


(SL4) $B(\varphi) = B(\psi \wedge \varphi) + B(\neg\psi \wedge \varphi) - B(\varphi \wedge \neg\psi \wedge \varphi)$. ■

In other words, SL -probabilities are axiomatized by the conditions (SL1)-(SL4).¹⁴

¹²The reader might be interested in why this shouldn't be straightforward, at least in the three-valued version of the Łukasiewicz logic, since it shares the truth tables for \wedge , \vee and \neg with the three logics under discussion here. Note, however, that its language contains also \rightarrow as a non-derived connective.

¹³See, however, footnote 16 below: we think that lattice-theoretic approach leads Bradley to some unfortunately phrased conclusions already in the truthfunctional cases.

¹⁴For those yearning to nibble at a morsel of additional formalism: the new axiom (SL4) is related to the presence in the Lindenbaum-Tarski algebra for SL logics of join-irreducible elements, which appear already once two propositional variables are admitted in the language. If we consider the Lindenbaum-Tarski algebra for the Symmetric Logic with two propositional variables x and y , then this will be the principal ideal generated by $x \wedge \neg y$:



As we can see, (the equivalence class of) $x \wedge \neg y$ cannot be arrived at as a join of two different elements. Therefore, even if the value of a credence function were fixed on all the atoms (literally: on all the elements of all the atoms) of the Lindenbaum-Tarski algebra, it wouldn't be possible to use the additivity axiom (SL3) to calculate its value for $x \wedge \neg y$. It is in such cases in which (SL4) is employed. For the gruesome details of this please refer to the technical paper (Gil Sanchez et al., ms).






Figure 1: The leftmost picture displays the Lindenbaum-Tarski algebra of the single-variable version of *SL*. The others display the three valuations possible in this context.

For an extended proof of this Theorem, see (Gil Sanchez et al., ms). One aspect of *SL* we've made use of is that, of the three logics *KL*, *LP* and *SL*, it is only the Symmetric Logic that is algebraizable: that is, only in that case the relation of mutual entailment is a congruence, and so the logic has its Lindenbaum-Tarski algebra.¹⁵ (Also, for similar reasons as in the classical propositional logic case above, the condition that *P* is finite cannot be dropped without significant modifications. The papers (Paris, 2005) and (Williams, 2012b) contain an exhaustive list of similar results for other logics and the issue of compositionality is also discussed therein.)

Notice that if the language has just a single propositional variable, all credences satisfying (SL1)-(SL4) are convex combinations of the three valuations displayed in Figure 1; that is, they are credences of this form:



The following is an example of such a credence:



It is important to distinguish two goals: the logical one of axiomatizing certain convex hulls, and the epistemological one of arguing that credences satisfying those axioms, or, equivalently, belonging to those convex hulls, are in some sense rational. With regard to the cognitive evaluations of Symmetric Logic, we've just described the achievement of the first of those two goals. We now turn to the second one.¹⁶

¹⁵If a logic is given semantically, with the entailment relation \models , then its Lindenbaum-Tarski algebra is a partition of the set of its sentences into equivalence classes (inside each of which, for every φ and ψ , $\varphi \models \psi$ and $\psi \models \varphi$) on which the operators of the logic's language behave "nicely". Take $[\varphi]$ to be the equivalence class of φ and assume the language has (only) the operators \wedge , \vee and \neg . In the Lindenbaum-Tarski algebra we require that, if $\alpha \in [\varphi]$ and $\beta \in [\psi]$:

- $\neg\alpha \in [\neg\varphi]$;
- $\alpha \wedge \beta \in [\varphi \wedge \psi]$; and
- $\alpha \vee \beta \in [\varphi \vee \psi]$.

Take, however, the two formulas $\theta_1 = p \wedge \neg p \wedge q$ and $\theta_2 = p \wedge \neg p \wedge \neg q$. They entail each other according to both \models_{KL} and \models_{LP} (so $\theta_1 \in [\theta_2]$ in both cases), but there are valuations in which their truth values differ (which can be used to show that, again in both cases, $\neg\theta_1 \notin [\neg\theta_2]$). Therefore no Lindenbaum-Tarski algebras for *KL* or *LP* exist. Their existence in the case of *SL* is exactly what allows us to draw the diagrams in this paper!

¹⁶This footnote is aimed at those readers who are familiar with Bradley (2017). The just provided simple examples of elements of the convex hull of cognitive evaluations of *SL* directly relate to one of the points made in that paper.

4 CONVEX HULLS, DUTCH BOOKS AND RATIONALITY

It is one thing to axiomatize the notion of a convex hull of cognitive loads in the context of some logic. It is another one to argue that these axioms should be satisfied by rational credence functions; in fact, it is achieving this goal that for epistemologists is probably the main allure of these considerations. Let us now turn to this task.

The chief reason stated in Paris (2005) for investigating axiomatizations of convex hulls of sets of valuations was that it is exactly the elements of such convex hulls that are not Dutch-bookable with respect to those valuations.¹⁷ One assumption was that the valuations are $\{1, 0\}$ -valued; however, Paris himself noted that this restriction could “clearly be relaxed”. Section 3 of (Williams, 2012a) contains a somewhat informal proof of a generalization of Paris’ result to finite “non-classical truth value distributions of truth values”¹⁸. Once cognitive loads are introduced, matters become a little bit more subtle. Assume possible worlds are governed by a logic with three truth values: *True*, *Other*, and *False*. Suppose a bet is bought for a proposition A with the stake S ; and that it turns out that the truth value of A is *Other*. What portion of S should be paid out—what Williams (2016) calls the “pragmatic load” of the *Other* value—depends on how *Other* is to be interpreted; it can be argued, e.g., that if the logic is KL , then the buyer should get nothing, if it is the LP , (s)he should receive the full S , and if it is the SL , then (s)he should end up with one half of S . In Section 5 of the 2016 handbook paper Williams points out that the demands *belong to the convex hull of cognitive evaluations* and *don’t be Dutch-bookable* might be inconsistent if cognitive loads differ from the pragmatic ones. To discuss the Paris-motivated connection between Dutch Books and convex hulls we thus assume that these two types of loads coincide.

Some comment regarding cognitive and pragmatic loads is in order. Williams (2016) takes them, in the context of some logic, to be properties of truth values¹⁹. This allows a hypothetical situation in which two different truth values share their cognitive load (say, 1) but differ in their pragmatic load (say, 1 vs. $1/2$). In such a case, even though there is no difference in the credence an omniscient agent should invest in a proposition depending on which of the two truth values it has, one of them is “worth more” in a betting situation. Similarly, truth values with differing cognitive loads might share their pragmatic load. One could certainly contemplate a different approach, in which one of the two types of loads was fundamental to the other; for example, pragmatic loads could be taken to “generate” cognitive loads, on the assumption that what omniscient agents care about can be reduced

Bradley aims to apply the Krein-Milman theorem, which says—under some assumptions about compactness and closedness—that the convex hull of the extremal elements of some convex set is that same convex set. His goal is to extend a theorem by Choquet which describes the extremal elements of the set of monotonic functions on a distributive lattice which satisfy the additivity axiom (SL3). The result he gives as his Theorem 3 is that, if we restrict our attention to functions which assign 1 to the top and 0 to the bottom element of the lattice, those extremal elements are indicator functions of ultrafilters.

Bradley states that in “nonclassical case[s] there can obviously be evaluations that aren’t indicator functions of ultrafilters (...) But (...) those evaluations will be non-extremal elements: they will be in the convex hull of the indicator functions of ultrafilters. So, as long as the indicator functions of ultrafilters are among the admissible evaluations, the convex hull of the evaluations will be equal to the convex hull of the indicator functions of ultrafilters” (p. 95). We believe it is important to note that this can be false e.g. if functions which assign something else than 0 to the bottom element of the lattice are considered. Take the single-variable version of SL : the three extremal elements of the set of evaluations are displayed in Figure 1. And while, indeed, the indicator functions of the two ultrafilters are there, the convex hull of the evaluations is decidedly *not* equal to the convex hull of the indicator functions of ultrafilters: for example, the “give-everything- $1/2$ ” evaluation displayed in the middle of Figure 1 cannot be obtained as a weighted mean of the two indicator functions of ultrafilters; one reason is that it has something else than 0 as the value of the bottom element. Therefore, despite the air of generality conveyed by Bradley’s ultimate conclusion that “as far as non-classical probability goes, it is only the lattice structure (encoded in the ultrafilters) that matters to what counts as probabilistically coherent”, we should take note that the claim has to be taken as holding only in the specific conditions in which the assumptions of Bradley’s Theorem 3 are satisfied.

¹⁷For introduction to Dutch Book arguments, see Vineberg (2016) and Pettigrew (2020).

¹⁸The result is not formally stated anywhere in (Williams, 2012a).

¹⁹That is: truth values as the term is used in the current paper. Williams (2016) writes about “truth statuses” here, explicitly using the term “truth value” so that its reference varies; see e.g. p. 262: “the ‘truth value’ of a sentence refers to the *pragmatic loading* of the relevant truth status, whereas in the previous results it referred to the *cognitive loading* of the truth statuses”. This leeway allows Williams to speak of “convex combinations of truth values”, even when what he calls “truth statuses” are not numerical in nature.

to betting profits. Since our project is to begin investigating how Dutch Books could be made to actually work in nonclassical settings, we shall put these issues aside here; matters will already be nontrivial if we stick to the identification of cognitive loads and pragmatic loads for the time being. And so, even if we will eventually argue that in nonclassical contexts the notion of Dutch Book needs to be modified, we shall now state the following generalization of Paris’ result using just the cognitive loads.

Assume, again, that we are given a propositional language \mathcal{L} with finitely many propositional variables and finitely many operators which include \vee and \wedge .

4.1 DEFINITION (DUTCH BOOK). In the context of a logic with the set of valuations \mathbb{V} , and the set of cognitive evaluations $\{c_V | V \in \mathbb{V}\}$, a function $B: \text{Sent}_{\mathcal{L}} \rightarrow [0, 1]$ *permits a Dutch Book* iff there are $\vartheta_1, \dots, \vartheta_n \in \text{Sent}_{\mathcal{L}}$ and $s_1, \dots, s_n \in \mathbb{R}$ such that for all elements of $\{c_V | V \in \mathbb{V}\}$ we have

$$\sum_{i=1}^n s_i (c_V(\vartheta_i) - B(\vartheta_i)) < 0. \quad (5)$$

4.2 THEOREM (VARIANT OF THEOREM 2 OF PARIS (2005)). In the context of a logic with the set of valuations \mathbb{V} , and the set of cognitive evaluations $\{c_V | V \in \mathbb{V}\}$, B does not permit a Dutch Book if and only if B is a convex combinations of the elements of $\{c_V | V \in \mathbb{V}\}$. ■

The shape of Formula (5) should be familiar to any reader of formal epistemology. The usual interpretation of the terms involved is as follows (assume, for brevity, that you can read “ B ” as “an agent with the credence function B ” whenever you feel it would be natural):

- s_i – the prize (stake) associated with the bet for ϑ_i ;
- $s_i \cdot B(\vartheta_i)$ – the cost of that bet;²⁰
- $s_i \cdot c_V(\vartheta_i)$ – the payout if ϑ_i has the truth value $V(\vartheta_i)$;²¹

and, crucially, it is understood that according to B , $s_i \cdot B(\vartheta_i)$ is the *fair price* for a bet with such a payout.²² A Dutch Book against B , then, is a collection of bets all of which B considers to be fairly priced but which taken together inevitably lead to loss.

Dutch Books are to be a sign of irrationality. Assume, then, that possible worlds are governed by the Symmetric Logic; for simplicity, suppose the language only has two variables, x and y . What’s wrong with not satisfying the (SL1)-(SL4) axioms? Since the first three conditions have been widely discussed in the literature²³, let us consider a B which violates (SL4) and construct a Dutch Book against it.

For convenience, let us label the sentence $x \wedge \neg x \wedge y \wedge \neg y$ —whose cell of the Lindenbaum-Tarski algebra for our logic is its bottom element—as \perp . Consider a B such that $B(x) = 1/3$, $B(y \wedge x) = B(\neg y \wedge x) = 1/4$ and $B(\perp) = 1/12$. Note that it does not follow from this assignment that any of the first three SL axioms be violated; however, (SL4) fails, since

$$1/3 = B(x) < B(y \wedge x) + B(\neg y \wedge x) - B(\perp) = 5/12.$$

To create a Dutch Book against this B we set the values for use in Formula (5) as per Table 1. Those values, according to the interpretation given above, mean that B considers it fair to:

- sell the bet for x for 4;

²⁰For a positive s_i , “how much B would pay for the bet had (s)he wanted to buy it”; for a negative s_i , “how much B would receive for the bet had (s)he wanted to sell it”.

²¹For a positive s_i , “how much B receives if she bought the bet and ϑ_i has the truth value $V(\vartheta_i)$ ”; for a negative s_i , “how much B has to pay if she sold the bet and ϑ_i has the truth value $V(\vartheta_i)$ ”. Remember also that we identify the pragmatic loads with the cognitive ones.

²²Why it should—or should not—be called so will be discussed below.

²³Although actual Dutch Books illustrating the violations of them are typically presented under classical assumptions.

| i | ϑ_i | s_i | $B(\vartheta_i)$ |
|-----|-------------------|-------|------------------|
| 1 | x | -12 | 1/3 |
| 2 | $y \wedge x$ | 12 | 1/4 |
| 3 | $\neg y \wedge x$ | 12 | 1/4 |
| 4 | \perp | -12 | 1/12 |

Table 1: The values for use in Formula (5), to create a Dutch Book against the B defined in the text.

| k | $c_{V_k}(x)$ | $c_{V_k}(y)$ | $c_{V_k}(y \wedge x)$ | $c_{V_k}(\neg y \wedge x)$ | $c_{V_k}(\perp)$ | Profit from bets | Total profit |
|-----|--------------|--------------|-----------------------|----------------------------|------------------|------------------|--------------|
| 1 | 1 | 1 | 1 | 0 | 0 | 0 | -1 |
| 2 | 1 | 1/2 | 1/2 | 1/2 | 0 | 0 | -1 |
| 3 | 1 | 0 | 0 | 1 | 0 | 0 | -1 |
| 4 | 0 | 1 | 0 | 0 | 0 | 0 | -1 |
| 5 | 0 | 1/2 | 0 | 0 | 0 | 0 | -1 |
| 6 | 0 | 0 | 0 | 0 | 0 | 0 | -1 |
| 7 | 1/2 | 1 | 1/2 | 0 | 0 | 0 | -1 |
| 8 | 1/2 | 1/2 | 1/2 | 1/2 | 1/2 | 0 | -1 |
| 9 | 1/2 | 0 | 0 | 1/2 | 0 | 0 | -1 |

Table 2: The payout table for the bets against a B violating the axiom (SL4) as discussed in the text.

- sell the bet for \perp for 1;
- buy the bet for $y \wedge x$ for 3;
- buy the bet for $\neg y \wedge x$ for 3.

Should all these bets go through, B would suffer a prior loss: -1 . The payouts from the bets in various worlds are given in Table 2. As we can see, the whole situation is a Dutch Book against B : no matter what happens, B ends up losing 1.

However, we find it highly doubtful that the existence of this Dutch Book is a sign of B 's irrationality. Bluntly put, why should B even care about *this* particular sets of bets? Why should we stipulate that the values from Table 1 encode numbers which B would take to be *fair prices* for the corresponding bets? For example, $B(x) = 1/3$. Why would B consider it to be fair to sell the bet for x —which pays 12 if the cognitive load of x is 1, pays 6 if it is $1/2$, and pays 0 otherwise—at the price of 4, that is, at one-third of the highest possible prize?

This assumption, which lies at the foundation of applying the Dutch Book idea in non-classical settings, seems not to have received sufficient scrutiny in the relevant literature. Williams (2012a) writes “[a]s standard in Dutch Book arguments, we assume that the fair price for an individual bet with unit prize for an individual with belief state b is specified by the degree of belief that b assigns to the proposition bet upon” (p. 817). He continues that he’s “not interested in whether the argument works”, and would rather investigate the “geometric structure of the argument”. However, from an epistemological point of view, in order to transform the formal insights into some conclusions regarding norms of rationality, we should definitely be interested whether the argument works. And on the most prevalent, indeed canonical, way of cashing out the notion of “fair price”, it just does *not* work.

The most common way of thinking about “fair price” in Dutch Book contexts, recall,²⁴ is to use the notion of expected value, and to think of bets which are fair according to B as those which B expects to favour neither buyer nor seller; that is, those which according to B have the expected value 0. Classically, to calculate the expected value of a bet for A from the perspective of B , we use the credence in that A is true, that is, $B(A)$, and the credence in that A is false, which we equate with $B(\neg A)$.²⁵ Since these alternatives exhaust the available

²⁴Consult Vineberg (2016) as a departure point.

²⁵This is problematic on its own; see Hedden (2013), Wroński and Godziszewski (2017), and Pettigrew (2021). (For starters, when arguing for probabilism it is a mistake to assume that $B(\neg A) = 1 - B(A)$; this should be a conclusion, not an assumption.)

options, we can calculate B 's expected profit from the given bet, assuming the price and the prize are provided. However, it should be clear that already in a three-valued setting this approach will not work. If the payout—whether 'cognitive loads' are considered or not—depends on the truth value of A , then we are lacking the required information about the credence in that A obtains the 'third' truth value. The values of B for A and $\neg A$, even assuming that the latter denotes the degree of belief in that A is false, are simply not enough to calculate B 's expected profit from the bet.

We will now propose a fix thanks to which we can, in nonclassical settings, consider Dutch Book arguments that indeed work.²⁶ It involves a modification of the usual notion of credence. On our proposal it will assign to propositions not single numbers, but rather vectors of numbers: as many as there are truth values according to the logic which governs the space of possible worlds. The approach will be similar to the one used in Janda (2016) in the context of accuracy measures. It will turn out that some results obtained using the usual notion of credence—for example, the axiomatization of the convex hulls of evaluations of SL —can be transformed so that their variants hold also when the new notion is used.

4.1 CREDENCES AS COMPLEX ATTITUDES

The idea is to treat credence in a proposition as a complex attitude, with as many dimensions as there are truth values. If, say, propositions can be True, False, or Half-True, one's credence in A is a triple of numbers: degrees of belief in *that A is True*, in *that A is False*, and in *that A is Half-True*.

More generally, assume a logic L is given in a language \mathcal{L} with valuations assigning to sentences elements from a finite set of truth values TV .²⁷ In such a context, credences are functions $B : TV \times Sent_{\mathcal{L}} \rightarrow [0, 1]$. We will use the expression $B_*(A)$ to denote $B(*, A)$; it is to be read as "credence in that A has the truth value $*$ ". If there are n truth values, we could equivalently be talking either about B or about n functions from $Sent_{\mathcal{L}}$ to $[0, 1]$; for example, in the case where TV has three elements—whatever they are—it might be convenient to speak about a credence function by referring to the indexed set of functions $\{B_1(\cdot), B_{1/2}(\cdot), B_0(\cdot)\}$.²⁸

While the ideas here are meant to be general, so that a variety of nonclassical logics and corresponding notions of credence can be considered, all examples illustrating our points will be given using the Symmetric Logic. As the Reader is well aware, it has three truth values, and three cognitive loads; while we noted that, under the assumptions that pragmatic loads are identical to cognitive loads, it is the cognitive loads which determine payouts from bets, in this particular case we can assume without loss of generality that it is the truth values that are doing the job. We will thus dispense with the notion of cognitive loads for the time being (we will come back to it in the general Definition 4.5). Let us also think of valuations as possible worlds, not because we wish to put any deep philosophy behind this choice, but so that we can conveniently speak of propositions having a certain truth value "at" a valuation V , and of agents "having profits" or "sustaining losses" at various V 's.

For convenience we assume that for any proposition all credences in that proposition having one of the various truth values sum up to 1; nothing formally important hangs on that, but without this assumption some formulas below would have to be more complicated. The idea is that there's some quantity of credence that's distributed among the possible options (the proposition in question having the various truth values); the convenient assumption

²⁶At least, that work as well as can be expected from a Dutch Book argument.

²⁷We would like to reiterate that in general we place no restrictions on what truth values are.

²⁸Note, again, that this way of writing is not meant to imply that the three truth values *are* 1, $1/2$, and 0. Below this notation is used so that $B_1(A)$ is to be understood as "credence in that A has the truth value associated with the pragmatic load 1", since it is the pragmatic loads—which, recall, determine the payout from a bet for a proposition enjoying a truth value with that pragmatic load—that are crucial for Dutch Book considerations. However, the convenient context of the Symmetric Logic allows us to equate truth values, cognitive loads, and pragmatic loads, which simplifies presentation. This will all be explained shortly.

amounts to a normalization of that quantity. We will write it out explicitly, since it does give us a philosophical bonus at one point, which some might find to be debatable:

Convenient Assumption (CA). For any $A \in \text{Sent}_{\mathcal{L}}$, $\sum_{* \in TV} B_*(A) = 1$.

Let us see how we can connect our modified concept of credence with the notion of bet assumed in Definition 4.1 and Theorem 4.2 to obtain examples of Dutch Book arguments in non-classical settings which actually work.

Following, for now, the lead of Williams (2012a), let us generalize the idea that a bet for A with the prize S pays out S if A is true and nothing if A is false. In other words, a bet for A pays out the portion of S given by the truth value of A .²⁹ That is, for a valuation V , a bet for A with the prize S pays out $V(A) \cdot S$ at V . For an agent who buys such a bet at cost C , then, the profit in V is $w(A) \cdot S - C$. We're after capturing the essence of Dutch-bookability, that is *exploitability* via *fair bets*: a Dutch Book against an agent (a credence function) is a collection of fair bets (that is, bets which are fair according to the agent's credence function) which ultimately lead to inevitable loss on part of any agent which would partake in all of them.

Which bets does a credence function consider fair? Given a proposition and a prize, a fair bet is that which has a *fair price*. We have already mentioned that, especially non nonclassical contexts, the formal epistemology literature has largely avoided extensive commentary on this issue. Typically, a price is considered to be fair according to B if under that price B expects the bet to favour neither buyer nor seller: according to B , the bet has the expected value 0.³⁰

Let us try to formulate a definition of a Dutch Book, suitable for nonclassical contexts, based on this notion of a fair price. How to cash it out formally? In our opinion the following idea is natural: given a prize S and proposition A , the expected value of a bet for A which costs C according to the credence function $B = \{B_*\}_{* \in TV}$ is

$$\sum_{* \in TV} B_*(A) (* \cdot S - C).$$

We're not the first to use the word 'expected' in this way: what we're doing here is essentially the same thing as e.g. what is proposed in Section 2 of Leitgeb and Pettigrew (2010), where the authors define 'expected inaccuracy', with the expectation calculated from the perspective of functions which are not assumed to be probabilities.³¹ (Our $B_*(A)$'s sum up to 1 for each A , so the formula *looks* like the classical expected value, but this is just because of our Convenient Assumption, made only for the purpose of simplifying the formulas involved.)

As already mentioned, we will be using *SL* to illustrate the proposed ideas. Assume, then, that the set of truth values is $TV = \{1, 1/2, 0\}$, and that a bet for A with prize S pays out $* \cdot S$ where $*$ is A 's truth value. It is then immediate to note that if $B(\cdot) = \{B_1(\cdot), B_{1/2}(\cdot), B_0(\cdot)\}$, then B 's fair price for a bet for A with prize S is $(B_1(A) + 0.5 \cdot B_{1/2}(A)) \cdot S$.³² We can thus now put forward the following definition of what it means for a credence to be Dutch-bookable, assuming it is the Symmetric Logic that governs the possible worlds.

²⁹Recall: for purposes of presentation only, since our primary topic until Definition 4.5 is the Symmetric Logic, we are identifying cognitive loads with truth values until we start discussing that definition. And in general, in this paper, we identify pragmatic loads with cognitive ones. So, to sum up: what actually determines the payout from a bet for A , given its prize S , is the pragmatic load of A 's truth value (which, in general, might be any object whatsoever). This load we assume in the current article to be identical to the value's cognitive load. And in the particular case of *SL* there is no harm in identifying it also with the truth value itself.

³⁰Cf. (Howson and Urbach, 2006, p. 54, our emphasis): "The condition of equal (and hence zero) risk is, of course, equal to that of equal (and hence zero) expected gain (...); thus *fair odds are those also that confer equal, meaning zero, advantage on each side of the bet*". At a fair price "you are indifferent between buying and selling the bet, and thus you see no advantage to either side" (Hájek, 2008, p. 795).

³¹Cf. p. 214 of Leitgeb and Pettigrew (2010): "while probability theory is the usual context in which expectations are defined, there is no objection in principle to extending the definition to cover the case of belief functions that may not be probability measures."

³²Without (CA), this formula would have to involve a fraction: $((B_1(A) + 0.5 \cdot B_{1/2}(A)) \cdot S) / ((B_1(A) + B_{1/2}(A) + B_0(A)))$. The situation is similar whenever fair prices are encountered below.

Table 3: A 'complex' credence function such that $B_1(\cdot) + 0.5 \cdot B_{1/2}(\cdot)$ gives the values of the credence function defined in Table 1, against which a Dutch Book exists as evidenced by Table 2.

| function \ Proposition | Proposition | | | |
|------------------------|-------------|--------------|-------------------|---------|
| | x | $y \wedge x$ | $\neg y \wedge x$ | \perp |
| $B_1(\cdot)$ | 9/36 | 6/36 | 6/36 | 2/36 |
| $B_{1/2}(\cdot)$ | 6/36 | 6/36 | 6/36 | 2/36 |
| $B_0(\cdot)$ | 21/36 | 24/36 | 24/36 | 32/36 |

4.3 DEFINITION (*SL-REALLY-DUTCH-BOOKABLE*). $B = \{B_1(\cdot), B_{1/2}(\cdot), B_0(\cdot)\}$ is **SL-really-Dutch-bookable** if there are $\vartheta_1, \dots, \vartheta_n \in \text{Sent}_{\mathcal{L}}$ and $s_1, \dots, s_n \in \mathbb{R}$ such that for all $v \in \mathbb{V}$ we have

$$\sum_{i=1}^n s_i \left(V(\vartheta_i) - (B_1(\vartheta_i) + 0.5 \cdot B_{1/2}(\vartheta_i)) \right) < 0. \quad (6)$$

That is, assuming that *SL* governs the possible worlds, B is really-Dutch-bookable if there is a series of bets B considers to be fair which inevitably lead to B 's loss: and thus the main intuition behind the notion of a Dutch Book is indeed captured.

It turns out we can use the result reported earlier, the axiomatization of convex hulls of evaluations of *SL* logic, to precisely specify which credences are not *SL*-really-Dutch-bookable:

4.4 FACT. $\{B_1(\cdot), B_{1/2}(\cdot), B_0(\cdot)\}$ is not *SL*-really-Dutch-bookable iff $B_1(\cdot) + 0.5 \cdot B_{1/2}(\cdot) : \text{Sent}_{\mathcal{L}} \rightarrow [0, 1]$ satisfies (SL1)-(SL4).

Proof. Define an 'old-style' credence function $b : \text{Sent}_{\mathcal{L}} \rightarrow [0, 1]$ as follows: $b(\varphi) := B_1(\varphi) + 0.5 \cdot B_{1/2}(\varphi)$. Then proceed through the following equivalences: $B_1(\cdot) + 0.5 \cdot B_{1/2}(\cdot)$ satisfies (SL1)-(SL4) iff b satisfies (SL1)-(SL4) iff b is not Dutch-bookable (in the sense of Definition 4.1, due to Theorem 4.2) iff there are no $\vartheta_1, \dots, \vartheta_n \in \text{Sent}_{\mathcal{L}}$ and $s_1, \dots, s_n \in \mathbb{R}$ such that for all $V \in \mathbb{V}$ (5) is satisfied (with b in place of B) iff there are no $\vartheta_1, \dots, \vartheta_n \in \text{Sent}_{\mathcal{L}}$ and $s_1, \dots, s_n \in \mathbb{R}$ such that for all $V \in \mathbb{V}$ (6) is satisfied iff $\{B_1(\cdot), B_{1/2}(\cdot), B_0(\cdot)\}$ is not *SL*-really-Dutch-bookable. ■

We thus have an example of a nonclassical setting and a Dutch Book argument that actually works in it: a function $B = \{B_1(\cdot), B_{1/2}(\cdot), B_0(\cdot)\}$ is not susceptible to an *SL*-real-Dutch Book only if $B_1(\cdot) + 0.5 \cdot B_{1/2}(\cdot)$ satisfies (SL1)-(SL4); otherwise there is a set of bets B considers to be fair which inevitably leads to B 's loss.

By way of illustration, let us see an example of a real Dutch Book in the *SL* setting. We will recreate the previous example in the new, 'complex' setting. Consider a credence $\{B_1, B_{1/2}, B_0\}$ with the values as given in Table 3. It is routine to check that (CA) is satisfied and that

$$\begin{aligned} 1/3 = B_1(x) + 0.5 \cdot B_{1/2}(x) &< B_1(y \wedge x) + 0.5 \cdot B_{1/2}(y \wedge x) + B_1(\neg y \wedge x) + \\ &+ 0.5 \cdot B_{1/2}(\neg y \wedge x) - B_1(\perp) - 0.5 \cdot B_{1/2}(\perp) = 5/12, \end{aligned}$$

that is, $B_1 + 0.5 \cdot B_{1/2}$ does not satisfy (SL4). The Dutch Book presented before shows that $\{B_1, B_{1/2}, B_0\}$ is really-Dutch-bookable; in Table 1 it suffices to substitute $B_1(\vartheta_i) + 0.5 \cdot B_{1/2}(\vartheta_i)$ for $B(\vartheta_i)$.

Having illustrated the idea behind real Dutch-bookability in the case of *SL*, let us give the general definition, involving cognitive loads (but sticking with the assumption that they are to be identified with pragmatic loads and continuing to assume (CA)). Suppose, then, that a logic L is given in a language \mathcal{L} as (\mathbb{V}, \models, c) with the valuations obtaining values in a finite set TV . Suppose that for each $* \in TV$, $c(*) \in [0, 1]$. If L is to be considered as governing the possible worlds, then credence functions B should be considered as being of

the form $B = \{B_*(\cdot)\}_{* \in TV}$. Then B 's expected profit from a bet for A which costs C and pays off $c(V(A)) \cdot S$ at world V is

$$\sum_{* \in TV} B_*(A)(c(*) \cdot S - C) \quad (7)$$

The fair price of such a bet from the perspective of such a B is the unique C which makes the expression (7) equal 0, that is, $C = \sum_{* \in TV} B_*(A) \cdot c(*) \cdot S$. (Note that this does indeed give back the classical “the fair price is the proportion of the prize which corresponds to the degree of belief” idea once enough assumptions are in place.)

With this in hand we can formulate the following general definition.

4.5 DEFINITION (*L*-REALLY-DUTCH-BOOKABLE). Suppose a logic L is given in a language \mathcal{L} as (\mathbb{V}, \models, c) with the valuations obtaining values in a finite set TV . Suppose that for each $* \in TV$, $c(*) \in [0, 1]$, and that credence functions B are of the form $B = \{B_*(\cdot)\}_{* \in TV}$.

B is *L*-**really-Dutch-bookable** if there are $\vartheta_1, \dots, \vartheta_n \in \text{Sent}_{\mathcal{L}}$ and $s_1, \dots, s_n \in \mathbb{R}$ such that for all $v \in \mathbb{V}$ we have

$$\sum_{i=1}^n s_i \left(c(V(\vartheta_i)) - \sum_{* \in TV} B_*(\vartheta_i) \cdot c(*) \right) < 0.$$

Note that if L is the classical logic, with $TV = \{True, False\}$ and $c(True) = 1$, $c(False) = 0$, then the above gives us the ‘typical’ notion of a Dutch Book. For a credence function $B = \{B_1(\cdot), B_0(\cdot)\}$, the expected value of a bet for A which costs C and whose prize is S is equal to $B_1(A) \cdot (S - C) - B_0(A) \cdot C$ (so, just like it should be, credence in $\neg A$ is not involved) and a variant of the norm of Probabilism can be recovered.

We hope *L*-real-Dutch-bookability will be studied for various logics L and that it will be possible to establish more connections between this notion and axiomatizations of sets of the given logic’s cognitive evaluations—an example of which is our Fact 4.4.

In the future it might also be fruitful to consider ‘complex’ cognitive loads. We’ve already mentioned our uneasiness with Williams’ proposal that in the case of *LP* the cognitive load of the truth value *Other* should be set to 1 just because its intended interpretation is “is both true and false”. Perhaps a more natural reading would be to say that, assuming *LP*, the credence an omniscient agent should assign to a proposition A with the truth value *Other* is $\{B_1(A) = 1, B_{1/2}(A) = 0, B_0(A) = 1\}$. We leave this topic for future research.

In most discussions of credential norms we find in formal epistemology, we can point to three aspects which are in a sense ‘classical’: the credence functions (considered as assigning single numbers to propositions), the (sometimes implicitly) classically conceived semantics, and the notion of a bet which is usually used in the literature. This paper contributes to the discussion in which the starting assumption is that the second element should be varied: the semantics under consideration may be nonclassical. In this Section we have so far been discussing varying the first element, so that we end up using some nonclassical version of the credence notion. However, one might not be happy with us keeping the notion of bet as employed by Williams. For example, if more than two cognitive loads are involved, a bet will have more than two possible payouts; some may find it unfortunate that it is no longer apparent what counts as winning or losing such a bet.³³ And one might have the intuition that a bet for A should be understood as a bet for *that A is true*, which should have two outcomes: it should be won if A is true and lost otherwise. We will now propose a formal generalization of this intuition, even though we believe that the presented notion of *L*-real-Dutch-bookability may be fruitfully used to deliver valid arguments for norms of rationality in various nonclassical settings. We will illustrate the idea using the classical and Symmetric logics, where cognitive loads (which we assume to be identical to pragmatic loads) directly correspond to truth values; in the following subsection we thus forego any mention of cognitive and pragmatic loads, assuming that is the truth values which straightforwardly determine the payouts.

³³We’d be happy with saying that “losing” and “winning” refer to the special case when a bet has only two possible payouts. Alternatively, “losing” may refer to receiving nothing and “winning”, which may then be a matter of degree, to any other outcome.

4.1.1 Truth-value bets

Traditionally, then, we seem to assume that a bet for A is a bet for *that A is true*. Let us once again tear off the classical shackles and stipulate that for a “truth-value bet” four things are needed: a proposition A , a truth value $*$, a prize S , and a cost C . We can then speak of a truth-value bet for *that A has the truth value $*$* . Note that no matter how many truth values there are, uncovering A ’s truth value leads to one of just two possible outcomes: either the bettor was right, or (s)he was wrong. It is thus natural to assume that such a bet, which costs C , pays out S if A has the truth value $*$, and pays out 0 otherwise. Where TV is a nonempty and finite set of truth values, the expected value of a truth-value bet for *that A has the truth value $*$* according to the credence $B = \{B_*(\cdot)\}_{* \in TV}$ is

$$B_*(A)(S - C) + \sum_{\# \in TV, \# \neq *} B_{\#}(A)(-C). \quad (8)$$

Sticking with the idea that the fair price of a truth-value bet is the cost C for which the above expression equals 0, we note that regardless of what TV consists of, **a fair price for the truth-value bet for that A has the truth value $*$ is $B_*(A)$** . A **truth-value Dutch Book** against a credence B is then a set of truth-value bets which are fair from the perspective of B , but which lead to B ’s inevitable loss.³⁴

4.6 DEFINITION (L -TRUTH-VALUE-DUTCH-BOOKABLE). Suppose a logic L is given in a language \mathcal{L} with the valuations obtaining values in a finite set $TV \subsetneq [0, 1]$. Assume credence functions B are of the form $B = \{B_*(\cdot)\}_{* \in TV}$.

B is **L -truth-value-Dutch-bookable** if there are $\vartheta_1, \dots, \vartheta_n \in \text{Sent}_{\mathcal{L}}$, $*_1, \dots, *_n \in TV$, and $s_1, \dots, s_n \in \mathbb{R}$ such that for all $V \in \mathbb{V}$ we have

$$\sum_{i: V(\vartheta_i) = *_i} s_i(1 - B_*(\vartheta_i)) + \sum_{i: V(\vartheta_i) \neq *_i} s_i(-B_*(\vartheta_i)) < 0. \quad (9)$$

The left sum in (9) refers to the profit from those bets belonging to the Dutch Book which are won at V , and the right one to the profit from those which are lost.

We will now consider two examples. For brevity, when we speak of (just) “a bet for A ”, we mean the bet as considered before the current subsection; when we speak of “a bet for *that A has truth-value $*$* ”, or “a bet for *that A is True*” (which is shorthand for “a bet for *that A has the truth-value 1*”), etc., we obviously mean a truth-value bet.

Let us first consider classical logic CL with $TV = \{1, 0\}$. Then a bet for *that A is False* with prize S and cost C is equivalent to the bet for *that A is True* with prize $-S$ and cost C . Therefore CL -truth-value-Dutch-bookability is reducible to “ CL -truth-value-Dutch-bookability exclusively via bets on that propositions are True”. This, in turn, is equivalent to Dutch-bookability (in the typical sense) of B_1 as a classically conceived credence. So: $B = \{B_1, B_0\}$ is not CL -truth-value-Dutch-bookable iff B_1 is a classical probability function.

Consider, now, the case of Symmetric Logic SL . Assume thus that $TV = \{1, 1/2, 0\}$. A bet for A with the fair price $(B_1(A) + 0.5 \cdot B_{1/2}(A)) \cdot S$ has the same payout table as the following pair of bets taken together:

- a bet for *that A is true* with prize S and cost $B_1(A) \cdot S$; and
- a bet for *that A has truth value $1/2$* with prize $0,5 \cdot S$ and cost $0,5 \cdot B_{1/2}(A) \cdot S$,

both of which are truth-value bets which are fair according to $\{B_1, B_{1/2}, B_0\}$.

Using this insight we can represent the previous Dutch Book as a truth-value Dutch Book. Consider the credence function displayed in Table 3 (p. 13). In Table 4, the column labelled

³⁴Recall that for ease of presentation the assumption of this subsection is that pragmatic loads coincide with truth values, and so if you’d like to use the concept of betting proposed here in the context of e.g. KL and LP , the definition would need to be suitably modified. Recall also that we continue assuming (CA).

| i | ϑ_i | $*_i$ | s_i | $B_*(\vartheta_i)$ |
|-----|-------------------|-------|-------|--------------------|
| 1 | x | 1 | -12 | 9/36 |
| 2 | x | 1/2 | -6 | 6/36 |
| 3 | $y \wedge x$ | 1 | 12 | 6/36 |
| 4 | $y \wedge x$ | 1/2 | 6 | 6/36 |
| 5 | $\neg y \wedge x$ | 1 | 12 | 6/36 |
| 6 | $\neg y \wedge x$ | 1/2 | 6 | 6/36 |
| 7 | \perp | 1 | -12 | 2/36 |
| 8 | \perp | 1/2 | -6 | 2/36 |

Table 4: A set of bets showing that the credence from Table 3 is *SL*-truth-value-Dutch-bookable.

' $*_i$ ' contains truth values. Each row i of that table defines a truth-value bet for *that* ϑ_i *has the truth value* $*_i$ with the prize s_i , the fair price of which is $B_*(\vartheta_i) \cdot s_i$. It is routine to check that participating in all of these bets yields the loss of 1. It is also routine to check that, again, Table 2 shows that this loss persists no matter what happens: therefore the whole situation is a truth-value Dutch Book against the credence defined in Table 3.

A straightforward generalization of this insight leads to the conclusion that **if a credence function is *SL*-really-Dutch-bookable, it is *SL*-truth-value-Dutch-bookable.**


However, the converse is not true. Assume the language has a single propositional variable x and consider the following credence B (which satisfies (CA)):

| | x | $\neg x$ | $x \vee \neg x$ | \perp |
|------------------|-----|----------|-----------------|---------|
| $B_1(\cdot)$ | 3/4 | 0 | 1/2 | 1/4 |
| $B_{1/2}(\cdot)$ | 0 | 1/2 | 1/2 | 0 |
| $B_0(\cdot)$ | 1/4 | 1/2 | 0 | 3/4 |

That this function is *SL*-truth-value-Dutch-bookable can be easily seen just from inspection of the top row of that table. Assume unitary prizes. From the perspective of the displayed credence the fair price for the bet for *that* x *is true* is 3/4, and the fair price for bet for *that* $x \vee \neg x$ *is true* is 1/2. Buying the first and selling the second bet establishes prior loss of -1/4. In each of the three possible worlds B ends up with a loss:

- if $V(x) = 1$, then both bets are won, and so the ultimate loss is the same as the prior loss;
- if $V(x) = 1/2$, then both bets are lost, and so the ultimate loss is the same as the prior loss;
- if $V(x) = 0$, then $V(x \vee \neg x) = 1$, and so the ultimate loss is -5/4.

Therefore B is *SL*-truth-value-Dutch-bookable. However, notice that $B_1 + 0.5 \cdot B_{1/2}$ is the following credence:



which we have already seen on p. 7 as a convex combination of two possible worlds (with weights 0.5). By the result of the current paper, we know that $B_1 + 0.5 \cdot B_{1/2}$ satisfies (SL1)-(SL4), and so is not *SL*-really-Dutch-Bookable.

We can therefore state the following Fact:

4.7 FACT. *SL*-real-Dutch-bookability implies *SL*-truth-value-Dutch-bookability; however, the converse does not hold in general.

In the context of truth-value bets we have discussed the classical and Symmetric logics. This is because in these contexts it is easy to think that truth values directly determine bet payouts. Generalizing Definition 4.6 so that it involved cognitive loads poses no formal problems. However, it's not evident for us what kind of betting we would then be modeling: while we find betting for that a proposition has a certain truth value to be somewhat intuitive, we would have to start thinking about betting that a certain proposition has this or that cognitive load, which we are reluctant to do without giving the matter more thought. We leave, then, investigating truth-value-Dutch-bookability in general for future research.

5 CONCLUSIONS

Continuing the “hard graft” proposed by Williams, we have reported a result concerning a logic which does not satisfy Paris’ conditions: the axiomatization of the convex hull of cognitive evaluations of Symmetric Logic (*SL*). We have then argued that it would be a mistake to claim that it is exactly the credences satisfying those axioms that are rational if *SL* governs the possible worlds, on the basis that the aforementioned convex hull coincides with the set of un-Dutch-bookable credences. That is, we have pointed out that the usual notion of a Dutch Book does not transfer immediately to nonclassical settings. To enable rigorous Dutch-Book-based arguments in such contexts, we have offered the notion of *L*-real-Dutch-bookability: what it means for a credence to be really Dutch-bookable on the assumption that the possible worlds are governed by a logic *L*. We have pointed out that in the case of Symmetric Logic the axiomatization of the convex hulls of the set of evaluation can in fact inform us about real-Dutch-bookability, too (Fact 4.4). Our hope is that *L*-real-Dutch-bookability can be fruitfully studied for various logics *L*.

Lastly, we have offered a modified—but also, we hope, intuitive—notation of bet, “truth-value bet”, which leads to another concept of Dutch Book. We have shown that in the case of *SL* this concept is weaker than the previous one. Its behaviour with regard to different logics is an open matter.

Acknowledgments

This research was supported by the grant 2019/34/E/HS1/00044, ‘Epistemic inaccuracy and foundational issues in formal epistemology’, of the National Science Centre (Poland) (recipient: Leszek Wroński).

We would like to thank Zoe Cocchiaro, Pavel Janda, Joanna Luc and Tomasz Placek for their insightful comments on earlier versions of the paper.

We would also like to express our gratitude towards the anonymous reviewer for the offered suggestions, which have clearly improved the paper.

REFERENCES

- Bradley, S. (2017). Nonclassical Probability and Convex Hulls. *Erkenntnis*, 82(1):87–101.
- Field, H. (2009). What is the normative role of logic? *Proceedings of the Aristotelean Society, Supplementary Volumes*, 83(1):251–268.
- Gil Sanchez, M., Gyenis, Z., and Wroński, L. (ms). Probability and symmetric logic. *Submitted*.
- Hájek, A. (2008). Arguments for-or against-probabilism. *British Journal for the Philosophy of Science*, 59:793–819.
- Hedden, B. (2013). Incoherence without exploitability. *Noûs*, 47(3):482–495.

- Howson, C. and Urbach, P. (2006). *Scientific Reasoning. The Bayesian Approach*, third edition. Open Court.
- Janda, P. (2016). Measuring inaccuracy of uncertain doxastic states in many-valued logical systems. *Journal of Applied Logic*, 14:95–112.
- Leitgeb, H. and Pettigrew, R. (2010). An Objective Justification of Bayesianism I: Measuring Inaccuracy. *Philosophy of Science*, 77(2):201–235.
- Mundici, D. (2006). Bookmaking over infinite-valued events. *International Journal of Approximate Reasoning*, 43:223–240.
- Paris, J. B. (2005). A note on the Dutch Book method. *Proceedings of the 2nd International Symposium on Imprecise Probabilities and their Applications, Ithaca, New York, 2001*, pp. 301-306. (The paper seems to have grown since its publication; see e.g. <http://www.stats.org.uk/why-probability/Dutch-book/Paris2002.pdf>)
- Pettigrew, R. (2020). *Dutch Book Arguments*. Cambridge University Press.
- Pettigrew, R. (2021). On the Expected Utility Objection to the Dutch Book Argument for Probabilism. *Noûs*, 55(1):23–38.
- Priest, G. (2001). *An Introduction to Non-Classical Logic*. Cambridge University Press.
- Vineberg, S. (2016). Dutch Book Arguments. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy (Spring 2016 Edition)*. URL = <https://plato.stanford.edu/archives/spr2016/entries/dutch-book/>.
- Williams, J. R. G. (2012a). Generalized Probabilism: Dutch Books and Accuracy Domination. *Journal of Philosophical Logic*, 41:811–840.
- Williams, J. R. G. (2012b). Gradational Accuracy and Nonclassical Semantics. *The Review of Symbolic Logic*, 5(4):513–537.
- Williams, J. R. G. (2016). Probability and non-classical logic. In Hájek, A. and Hitchcock, C. R., editors, *The Oxford Handbook of Probability and Philosophy*, pages 248–276. Oxford University Press.
- Wroński, L. and Godziszewski, M. T. (2017). The Stubborn Non-probabilist - 'Negation Incoherence' and a New Way to Block the Dutch Book Argument. In Baltag, A., Seligman, J., and Yamada, T., editors, *Logic, Rationality, and Interaction. Proceedings of the 6th International Workshop, LORI 2017, Sapporo, Japan, September 11-14, 2017*, pages 256–267. Springer.