



Lawry, J., & Tang, Y. (2012). On truth-gaps, bipolar belief and the assertability of vague propositions. *Artificial Intelligence*, 191-192, 20-41. <https://doi.org/10.1016/j.artint.2012.07.004>

Early version, also known as pre-print

Link to published version (if available):
[10.1016/j.artint.2012.07.004](https://doi.org/10.1016/j.artint.2012.07.004)

[Link to publication record in Explore Bristol Research](#)
PDF-document

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/pure/about/ebr-terms>

On Truth-gaps, Bipolar Belief and the Assertability of Vague Propositions

Jonathan Lawry,
Intelligent Systems Laboratory,
University of Bristol,
Bristol BS8 1TR, UK
email: j.lawry@bris.ac.uk

Yongchuan Tang,
College of Computer Science
Zhejiang University,
Hangzhou 310027, PR China
email: tyongchuan@gmail.com

Abstract

This paper proposes an integrated approach to indeterminacy and epistemic uncertainty in order to model an intelligent agent's decision making about the assertability of vague statements. Initially, valuation pairs are introduced as a model of truth-gaps for propositional logic sentences. These take the form of lower and upper truth-valuations representing *absolutely true* and *not absolutely false* respectively. In particular, we consider valuation pairs based on supervaluationist principles and also on Kleene's three valued logic. The relationship between Kleene valuation pairs and supervaluation pairs is then explored in some detail with particular reference to a natural ordering on semantic precision. In the second part of the paper we extend this approach by proposing bipolar belief pairs as an integrated model combining epistemic uncertainty and indeterminacy. These comprise of lower and upper belief measures on propositional sentences, defined by a probability distribution on a finite set of possible valuation pairs. The properties of these measures are investigated together with their relationship to different types of uncertainty measure. Finally, we apply bipolar belief measures in a preliminary decision theoretic study so as to begin to understand how the use of vague expressions can help to mitigate the risk associated with making forecasts or promises. This then has potential applications to natural language generation systems.

Keywords Vagueness, truth-gaps, valuation pairs, semantic uncertainty, integrated uncertainty, bipolar belief measures.

1 Introduction

A defining feature of vague concepts is that they admit borderline cases which neither definitely satisfy the concept nor its negation. For example, there are some height values which would neither be definitely classified as being *short* nor *not short*. For propositions involving vague concepts this naturally results in truth-gaps. In other words, there are cases in which a proposition is neither *absolutely true* nor *absolutely false*. If Ethel's height lies in a certain intermediate range then the proposition 'Ethel is short' may be inherently borderline. Such truth-gaps suggests that a non-Tarskian notion of truth may be required

to capture this aspect of vagueness even in a simple propositional framework. There has been a number of different possibilities proposed in the literature for this alternative model of truth including three valued logics and supervaluations. In the sequel we will discuss and relate two different models of truth-gaps, supervaluationism and Kleene's strong three valued logic, in the context of a new framework for bipolar valuations. In particular, we will investigate propositional truth-models taking the form of a lower and an upper truth valuation on the sentences of the language. The underlying idea is that, given such a *valuation pair*, the lower truth valuation represents the strong criterion of being *absolutely true*, while the upper valuation represents the weaker criterion of being *not absolutely false*. In this context, borderline statements are those for which there is a difference between the lower and upper valuations (i.e. a truth-gap).

A different perspective on borderline cases relates to the notion of assertability. Here we take the view that concept definitions are to a large extent determined by linguistic convention, and according to such conventions a statement may or may not be assertable given a particular state of the world. Interestingly, a case can be made that assertability is inherently bipolar, a phenomenon which manifests itself in a distinction between those propositions which convention would deem definitely assertable, and those which convention would not classify as incorrect, or perhaps even dishonest, to assert. Parikh [27] observes that:

Certain sentences are assertible in the sense that we might ourselves assert them and other cases of sentences which are non-assertible in the sense that we ourselves (and many others) would reproach someone who used them. But there will also be the intermediate kind of sentences, where we might allow their use.

For example, consider a witness in a court of law describing a suspect as being *short*. Depending on the actual height of the suspect this statement may be deemed as clearly true or clearly false, in which latter case the witness could be accused of perjury. However, there will also be an intermediate height range for which, while there may be doubt and differing opinions concerning the use of the description *short*, it would not be deemed as definitely inappropriate and hence the witness would not be viewed as committing perjury. In other words, for certain height values of the suspect, it may be acceptable to assert the statement 'the suspect was short', even though this statement would not be viewed as being absolutely true. Clearly there is a natural connection between this bipolar aspect of assertability and the idea of truth-gaps for borderline cases outlined above. If a statement θ is absolutely true, a judgment which is of course dependent both on the state of the world and on how linguistic convention defines the relevant concepts, then θ would be *definitely assertable*. On the other hand, provided that θ is not absolutely false then θ would be deemed *acceptable to assert*. The bipolarity of assertability would seem to be

a special case of what Dubois and Prade [7] refer to as *symmetric bivariate unipolarity*, whereby judgments are made according to two distinct evaluations on unipolar scales i.e. distinct evaluations about the assertability of a sentence and its negation. In the current context, we have a strong and a weak evaluation criterion where the former corresponds to definite assertability and the latter to acceptable assertability. As with many examples of this type of bipolarity there is a natural duality between the two evaluation criterion in that a proposition is definitely assertable if and only if it is not acceptable to assert its negation.

The adequate representation of epistemic uncertainty is of central importance in any effective model of belief. Typically we think of uncertainty as arising because of insufficient information about the state of the world. However, in the presence of vagueness there may also be *semantic uncertainty* due to our having only partial knowledge of language conventions. For example, consider the proposition ‘Ethel is short’. Here an agent with certain knowledge of Ethel’s height may still be uncertain as to the truth of this proposition due to uncertainty about the conventions governing the definition of the concept short. Such uncertainty may naturally arise from the distributed manner in which language is learnt across a population of communicating agents. Semantic uncertainty often occurs in conjunction with a lack of knowledge concerning the underlying state of the world. In our example, the agent may also be uncertain as to the precise value of Ethel’s height. In the sequel then we propose an integrated model of semantic and stochastic uncertainty in the context of language conventions which admit borderline cases. Here we view truth as a function of both the state of the world, e.g. Ethel’s height, and language convention, e.g. the interpretation of the concept short in terms of height values. An integrated model of epistemic uncertainty and truth-gaps can then take the form of a probability distribution on the cross product of the set of possible world states and the set of possible language conventions. Furthermore, if a convention maps each state of the world to a valuation pair, then this naturally results in a probability distribution on possible valuation pairs. Given such a distribution we can immediately define lower and upper measures by evaluating the probabilities of those valuation pairs in which a given sentence is absolutely true and of those in which it is not absolutely false respectively. We refer to these lower and upper measures on the sentences of the language as a *bipolar belief pair*.

We argue that valuation pairs are one of the most straight-forward representations of truth-gaps in natural language propositions. Hence, by taking a probability distribution over a set of possible valuation pairs for the language we generate a very natural integrated model of belief for propositions and sentences which involve vague concepts and about which there is inherent uncertainty. As such, the proposed framework provides an ideal platform from which we can begin to explore issues concerning the utility of vagueness in communication. Certainly there are many potential applications of such a study including in natural language generation [39], consensus modelling [24] and multi-agent dialogues

[22]. In this paper we shall focus mainly on the first of these application areas.

A fundamental open problem in natural language generation is that of understanding why individuals often choose to make vague assertions rather than semantically similar crisp (non-vague) ones. In particular, what are the practical advantages of such a decision from the perspective of an asserting agent? One approach is to apply decision theory so that utility values are associated with different expressions quantifying the benefit or gain resulting from their assertion given a particular state of the world and in the context of particular language conventions. Given such a formalism the problem is then that of understanding why, in certain situations, a vague assertion may result in a higher expected utility than for a similar crisp assertion. One type of situation where vague assertions may be preferable in this way, is referred to by van Deemter [39] as *future contingencies*. Here we assume that agents are playing what may, to some extent, be seen as a non-cooperative language game in which different agents have different, and possibly conflicting, goals and objectives. Within communication environments of this kind there may be a risk associated with making forecasts or promises which turn out to be wrong or which cannot be kept. This would then be reflected in the utility values for different assertions within a particular context. Mitigating this risk is certainly of practical importance in AI systems such as for automatic weather forecasting or medical diagnosis. In [39] it is suggested that for predictions or promises vague assertions may be lower risk than crisp ones. In this paper we shall apply the bipolar belief framework outlined above to carry out a very preliminary study of future contingencies, from a decision theoretic perspective. The aim will be to demonstrate the potential of our approach for modelling assertion behaviour, as well as providing some small insight into how truth-gaps may be exploited so as to minimize risk and maximize gain in complex multi-agent dialogues.

An outline of the paper is as follows: In following section we describe several approaches to modelling truth-gaps and emphasise how these differ from theories of epistemic uncertainty. Section 3 then introduces valuation pairs and provides axiomatic characterisations both of supervaluation pairs and of Kleene valuation pairs. In this section we also investigate the relationship between Kleene and supervaluation pairs as well as discussing the functionality and truth-functionality of different classes of valuation pairs. Section 4 proposes bipolar belief measures as an integrated model combining epistemic uncertainty and truth-gaps. We discuss the motivation and justification for defining lower and upper belief measures generated from probability distributions on valuation pairs, and investigate the properties of such belief pairs especially relating to a semantic precision ordering on valuation pairs. This section will also include an exploration of the relationship between classes of bipolar belief measures and other types of uncertainty measures, such as possibility measures and interval-valued fuzzy sets. Section 5 outlines a very preliminary decision theoretic analysis of the use of vague statements in order to mitigate the risk associated with making predictions or promises. Finally, section 6 gives conclusions and

proposes some future directions of study.

2 Models for Truth Gaps

In this section we outline two general approaches to the modelling of truth-gaps in propositional logic. The first is that of three-valued logics in which an additional truth-value is included in the underlying truth-model in order to represent borderline cases. A fully truth-functional calculus is then defined for the truth values $\mathbb{T} = \{\mathbf{t}, \mathbf{b}, \mathbf{f}\}$ where \mathbf{t} denotes *absolutely true*, \mathbf{b} denotes *borderline* and \mathbf{f} denotes *absolutely false*. Different logics are then characterised by different three-valued truth tables for the different connectives. Well known examples of three-valued logics include Lukasiewicz logic [26] and Kleene’s weak and strong logics [16]. For example, table 1 show the truth tables for the connectives \neg (negation), \wedge (conjunction) and \vee (disjunction) for Kleene’s strong logic.

	\neg
\mathbf{t}	\mathbf{f}
\mathbf{b}	\mathbf{b}
\mathbf{f}	\mathbf{t}

\wedge	\mathbf{t}	\mathbf{b}	\mathbf{f}
\mathbf{t}	\mathbf{t}	\mathbf{b}	\mathbf{f}
\mathbf{b}	\mathbf{b}	\mathbf{b}	\mathbf{f}
\mathbf{f}	\mathbf{f}	\mathbf{f}	\mathbf{f}

\vee	\mathbf{t}	\mathbf{b}	\mathbf{f}
\mathbf{t}	\mathbf{t}	\mathbf{t}	\mathbf{t}
\mathbf{b}	\mathbf{t}	\mathbf{b}	\mathbf{b}
\mathbf{f}	\mathbf{t}	\mathbf{b}	\mathbf{f}

Table 1: Truth tables from Kleene’s strong three-valued logic.

By referring to the middle truth-value as *borderline* we have consciously adopted somewhat non-standard terminology. Both Lukasiewicz and Kleene use terminology suggestive of an epistemic interpretation of \mathbf{b} corresponding to either *possibly true* or *unknown*¹. This would seem to assume an underlying binary truth-model about which there is only partial or incomplete knowledge. However, this view is problematic since, as highlighted by Dubois [6], three-valued logics do not provide an adequate model of ignorance or partial knowledge. For example, even if there is uncertainty concerning the *binary truth-value* of a proposition p then the contradiction $p \wedge \neg p$ would still be *known* to be false since it is false in all binary valuations. However, for Kleene’s strong logic we can see from table 1 that if p has the middle truth value then so has $p \wedge \neg p$.

In contrast to the epistemic view, here we are interested in the application of three-valued logic to model truth-gaps for propositions involving vague concepts. From this perspective a proposition may be *known* to be borderline. For example, given absolute certainty about Ethel’s height then we may also be certain that the proposition $p =$ ‘Ethel is tall’ is a borderline case. The intermediate truth-value of p results then from the inherent flexibility in the concept tall and is not the result of a lack of knowledge about any underlying binary truth-model. Indeed, from this perspective there is no underlying binary truth-model since truth valuations on vague propositions are inherently three-valued.

¹Shapiro [35] uses the term *indeterminate* which seems closer to our interpretation.

A second model of truth-gaps is *supervaluationism* as proposed by Fine [10] (see also Williamson [44] for an exposition)². In this approach it is assumed that vague concepts have different admissible precise interpretations referred to as precisifications. For example, the concept tall may admit a range of admissible threshold values on height, each defining a different precisification of the concept. The simplest formulation of supervaluationism in a propositional logic framework is in terms of a set of admissible classical (Tarski) valuations. In this case a sentence is said to be *supertrue* if it is true in every admissible classical valuation and *superfalse* if it is false in every admissible valuation (i.e. its negation is supertrue). Borderline cases are then those sentences which are neither supertrue nor superfalse. Notice, however, that unlike three-valued logic supervaluationism preserves classical tautologies, contradictions and equivalences. For example, even though proposition p and its negation $\neg p$ may both be borderline cases, $p \vee \neg p$ and $p \wedge \neg p$ are always supertrue and superfalse respectively. Furthermore, if two propositional logic sentences are classically equivalent then they are either both supertrue, both superfalse or both borderline.

The relationship between supervaluationism and models of partial or incomplete knowledge is perhaps less straightforward than in the case of three-valued logic. In one sense there clearly is an underlying binary truth-model since supervaluations are defined in terms of sets of classical valuations. Indeed, sets of valuations have been proposed as a way of representing incomplete knowledge of an agent's beliefs. An example of such a model would be *Boolean (two-valued) possibility theory* [9], [3] in which the *necessity* or *possibility* of a sentence is decided on the basis of whether or not it is respectively a direct consequence of, or consistent with a set of sentences K representing the agent's incomplete beliefs about the world. In this context, K can be characterised by the set of valuations for which every sentence in K is true, and hence a sentence is necessarily true if it is true in every such valuation and possibly true if it is true in some of these valuations. However, in supervaluationism we are using sets of valuations in a completely different way so as to capture vagueness rather than partial knowledge. To explain the difference between these two distinct uses of sets of valuations we shall employ the notions of *conjunctive and disjunctive sets* [42], [5].

A conjunctive set is an inclusive collection of elements representing a conjunctive property, whereas a disjunctive set is an exclusive collection of elements each representing one of a number of possibilities only one of which can actually be realised. For example, the set of languages spoken by Ethel is a conjunctive set since each element is equally admissible, while the set $\{1, \dots, 6\}$ representing the possible outcomes of a throw of a die is a disjunctive set since, while all outcomes are possible, only one outcome can be the actual score.

²The term supervaluationism was originally introduced by van Fraassen [41] with regard to the truth-gaps which result when some of the terms in a predicate logic language do not have referents in a particular interpretation.

We argue then that the sets of valuations in supervaluationism are conjunctive, while those in Boolean possibility theory are disjunctive. More specifically, in supervaluationism a set of valuations identifies a conjunction of interpretations *all of which* are admissible precisifications, whereas in Boolean possibility theory the set of valuations correspond to possible epistemic states, *only one of which* is the true state of the world. Interestingly, this view of supervaluationism is what Smith [36] refers to as *plurivaluationism*, according to which a language has multiple valid interpretations.

In the following sections we introduce valuation pairs as a general model of truth-gaps for sentences from a propositional language, and which allows for a unified treatment of supervaluationism and three-valued logics. Such a treatment will enable us to make a close comparison between the different truth-models, and here we shall focus in particular on the relationship between supervaluationism and Kleene's strong logic.

3 Valuation Pairs

Valuation pairs provide a direct, and relatively straightforward, model of truth-gaps and indeterminism based on two primitive truth-states; *absolutely true* and *not absolutely false*. These lower and upper truth-states then drive (or are perhaps motivated by) the type of bipolarity of assertability identified by Parikh [27] as being fundamental to the use of vague concepts in natural language. In this section we give the general definition of valuation pairs and identify a number of sub-classes of valuations pairs for further study. Initially, however, we recall some basic concepts from propositional logic.

Let \mathcal{L} be a language of propositional logic with connectives \wedge , \vee and \neg and a finite set of propositional variables $P = \{p_1, \dots, p_n\}$. Let $S\mathcal{L}$ denote the sentences of \mathcal{L} generated by recursive application of the connectives to P . We also define $S\mathcal{L}^+$ as the sentences of \mathcal{L} generated recursively from the propositional variables by application of the connectives \wedge and \vee only, and similarly we let $S\mathcal{L}^-$ denote the sentences of \mathcal{L} generated recursively from the negated propositional variables by application of \wedge and \vee only. Here we can think of $S\mathcal{L}^+$ as the set of *entirely positive* sentences and $S\mathcal{L}^-$ as the set of *entirely negative* sentences of \mathcal{L} respectively.

Classical valuations are based on Tarski's binary truth-model as follows:

Definition 1. *Classical Valuations*

A classical valuation v is a function $v : S\mathcal{L} \rightarrow \{0, 1\}$ satisfying $\forall \theta, \varphi \in S\mathcal{L}$,

- $v(\neg\theta) = 1 - v(\theta)$
- $v(\theta \wedge \varphi) = \min(v(\theta), v(\varphi))$
- $v(\theta \vee \varphi) = \max(v(\theta), v(\varphi))$

A classical valuation is characterised by a subset $F \subseteq P$ so that:

$$v_F(p_i) = 1 \text{ iff } p_i \in F$$

On the basis of classical valuations we can then define the standard propositional logic notions of entailment and equivalence according to which $\theta \models \varphi$ (φ follows from θ) if for all classical valuations $v(\theta) = 1$ implies that $v(\varphi) = 1$, and $\theta \equiv \varphi$ (θ and φ are equivalent) if for all classical valuations $v(\theta) = v(\varphi)$. A sentence θ is then a classical tautology (denoted $\models \theta$) if $v(\theta) = 1$ in all classical valuations, while θ is a classical contradiction $\models \neg\theta$ if $v(\theta) = 0$ for all classical valuations.

Definition 2. *Atoms of \mathcal{L}*

The atoms of \mathcal{L} are those sentences of the form:

$$\alpha = \bigwedge_{i=1}^n \pm p_i \text{ where } + p_i \text{ denotes the propositional } p_i \text{ and } - p_i \text{ denotes the negated propositional variable } \neg p_i$$

For atom $\alpha \in S\mathcal{L}$ let $F_\alpha = \{p_i : \alpha \models p_i\}$. Also any subset $F \subseteq P$ characterises an atom so that:

$$\alpha_F = \left(\bigwedge_{p_i \in F} p_i \right) \wedge \left(\bigwedge_{p_i \notin F} \neg p_i \right)$$

By the Disjunctive Normal Form theorem it holds that $\forall \theta \in S\mathcal{L}$,

$$\theta \equiv \bigvee_{F \in \mathcal{P}_\theta} \alpha_F \text{ where } \mathcal{P}_\theta = \{F : v_F(\theta) = 1\}$$

Definition 3. *Valuation Pairs*

A valuation pair is a pair of functions $\vec{v} = (\underline{v}, \bar{v})$ where $\underline{v} : S\mathcal{L} \rightarrow \{0, 1\}$ and $\bar{v} : S\mathcal{L} \rightarrow \{0, 1\}$ such that $\underline{v} \leq \bar{v}$. Furthermore, $\forall \theta, \varphi \in S\mathcal{L}$, if $\underline{v}(\theta) = \bar{v}(\theta) = \alpha$ and $\underline{v}(\varphi) = \bar{v}(\varphi) = \beta$ then $\underline{v}(\neg\theta) = \bar{v}(\neg\theta) = 1 - \alpha$, $\underline{v}(\theta \wedge \varphi) = \bar{v}(\theta \wedge \varphi) = \min(\alpha, \beta)$ and $\underline{v}(\theta \vee \varphi) = \bar{v}(\theta \vee \varphi) = \max(\alpha, \beta)$.

Intuitively, \underline{v} and \bar{v} are valuation functions for the truth-states *absolutely true* and *not absolutely false* respectively. The requirement that $\underline{v} \leq \bar{v}$ formally defines \underline{v} as a stronger criterion than \bar{v} . Hence, a sentence $\theta \in S\mathcal{L}$ is absolutely true only if it is not absolutely false. However, the converse does not necessarily hold in which case the difference between \bar{v} and \underline{v} is due to inherent vagueness in the sentences of \mathcal{L} . So if $\bar{v}(\theta) - \underline{v}(\theta) = 1$ then θ is a borderline case which is neither absolutely true nor absolutely false. In other words, θ is a sentence which while not definitely assertable is nonetheless acceptable to assert. In the case that we restrict ourselves to sentences of \mathcal{L} which are not viewed as being borderline it is assumed that valuation pairs behave like classical valuations as given in definition 1.

Indeed, from the perspective of definition 3, classical valuations can be viewed as special cases of valuation pairs in which $\underline{v} = \bar{v}$.

In the following definition we propose supervaluation pairs as a valuation pair model of supervaluationism. A supervaluation pair is defined by a set of classical valuations which, when adopting the set of propositional variables characterisation of classical valuations outlined above, is represented by a subset \mathcal{P} of 2^P (i.e. the power set of P). In this context $\underline{v}(\theta) = 1$ if and only if sentence θ is true in every admissible classical valuation in \mathcal{P} , while $\bar{v}(\theta) = 1$ if and only if θ is true in at least one classical valuation in \mathcal{P} . Clearly then we are identifying the truth-state *absolutely true* with supertruth and that of *not absolutely false* with non-superfalseness. Upper valuations of the latter kind have been investigated by Hyde [14] who refers to them as *subvaluations*.

Definition 4. *Supervaluation Pairs*

Let $\emptyset \neq \mathcal{P} \subseteq 2^P$ then a supervaluation pair $\vec{v} = (\underline{v}, \bar{v})$ generated by \mathcal{P} is given by: $\forall \theta \in \mathcal{L}$
 $\underline{v}(\theta) = \min\{v_F(\theta) : F \in \mathcal{P}\}$ and $\bar{v}(\theta) = \max\{v_F(\theta) : F \in \mathcal{P}\}$

In the sequel we will show that supervaluation pairs retain many of the properties of classical valuations including classical equivalence, so that if $\theta \equiv \varphi$ then $\vec{v}(\theta) = \vec{v}(\varphi)$ for all supervaluation pairs, and classical tautologies, so that if $\models \theta$ then $\vec{v}(\theta) = (1, 1)$ for all supervaluation pairs. Furthermore, supervaluation pairs satisfy the natural *duality property* that for sentence θ , $\underline{v}(\neg\theta) = 1 - \bar{v}(\theta)$ and $\bar{v}(\neg\theta) = 1 - \underline{v}(\theta)$. In other words, a sentence is absolutely true if and only if its negation is absolutely false. An obvious consequence of this and the fact that supervaluation pairs preserve classical tautologies is that classical contradictions are also preserved.

A supervaluation pair $\vec{v} = (\underline{v}, \bar{v})$ naturally identifies lower and upper sets of propositional variables of the form $\underline{F} = \{p_i \in P : \underline{v}(p_i) = 1\} \subseteq \bar{F} = \{p_i \in P : \bar{v}(p_i) = 1\}$. Notice immediately that for any $F \in \mathcal{P}$ it must hold that $\underline{F} \subseteq F \subseteq \bar{F}$, however, it does not necessarily follow from definition 4 that either \underline{F} or \bar{F} are themselves in \mathcal{P} . In the next definition we identify *bounded supervaluation pairs* as the class of supervaluation pairs for which $\mathcal{P} \supseteq \{\underline{F}, \bar{F}\}$. The failure to include \underline{F} or \bar{F} in \mathcal{P} has the following interesting consequences. If $\underline{F} \notin \mathcal{P}$ then $\vec{v}(\neg p_i) = (0, 1)$ for $p_i \in \bar{F} - \underline{F}$ while $\vec{v}(\bigwedge_{p_i \in \bar{F} - \underline{F}} \neg p_i) = (0, 0)$. In other words, even though the negation of every propositional variable in $\bar{F} - \underline{F}$ is a borderline case their conjunction is absolutely false. Similarly, if $\bar{F} \notin \mathcal{P}$, then $\vec{v}(\bigwedge_{p_i \in \bar{F} - \underline{F}} p_i) = (0, 0)$. That is, the conjunction of all borderline propositional variables in \mathcal{L} is absolutely false. Hence, for bounded supervaluation pairs we are restricting ourselves to sets of admissible classical valuations which include a valuation in which all borderline propositional variables are true, and also one in which they are all false.

Definition 5. *Bounded Supervaluation Pairs*

A bounded supervaluation pair is a supervaluation pair where $\exists \underline{F}, \bar{F} \in \mathcal{P}$ such that $\underline{F} \subseteq \bar{F}$ and $\forall F \in \mathcal{P}$, $\underline{F} \subseteq F \subseteq \bar{F}$

Complete bounded supervaluation pairs are a more restricted class of bounded supervaluation pairs intended to capture the situation in which an agent's truth-model is generated solely by them considering propositional variables, and by identifying those propositional variables which are absolutely true and those which are not absolutely false. The agent then simply identifies the set of admissible valuations as being all those classical valuations consistent with these lower and upper valuations of the propositional variables.

Definition 6. *Complete Bounded Supervaluation Pairs*

A complete bounded supervaluation pair is a bounded supervaluation pair where $\mathcal{P} = \{F : \underline{F} \subseteq F \subseteq \overline{F}\}$.

An alternative approach to defining valuation pairs is via compositional rules for each of the connectives. In particular, *Kleene valuation pairs* are defined in terms of an explicit duality rule for negation together with min and max combination operators for conjunction and disjunction respectively. Kleene valuation pairs have already been defined in [23], where some of their basic properties were described.

Definition 7. *Kleene Valuation Pairs*

A Kleene valuation pair is a valuation pair $\vec{v} = (\underline{v}, \overline{v})$ such that $\forall \theta, \varphi \in S\mathcal{L}$ the following hold:

- $\underline{v}(\neg\theta) = 1 - \overline{v}(\theta)$ and $\overline{v}(\neg\theta) = 1 - \underline{v}(\theta)$
- $\underline{v}(\theta \wedge \varphi) = \min(\underline{v}(\theta), \underline{v}(\varphi))$ and $\overline{v}(\theta \wedge \varphi) = \min(\overline{v}(\theta), \overline{v}(\varphi))$
- $\underline{v}(\theta \vee \varphi) = \max(\underline{v}(\theta), \underline{v}(\varphi))$ and $\overline{v}(\theta \vee \varphi) = \max(\overline{v}(\theta), \overline{v}(\varphi))$

There is a natural link between Kleene valuation pairs and Kleene's strong three-valued logic [16]. This can be seen clearly when we view the three possible values of a valuation pair for a sentence as truth-values i.e. $\mathbf{t} = (1, 1)$ as absolutely true, $\mathbf{b} = (0, 1)$ as borderline and $\mathbf{f} = (0, 0)$ as absolutely false. If we then view a valuation pair as a function $\vec{v} : S\mathcal{L} \rightarrow \mathbb{T}$ where $\mathbb{T} = \{(0, 0), (0, 1), (1, 1)\}$ then Kleene valuation pairs as given in definition 7, satisfy the truth-tables for Kleene's strong three-valued logic as described in section 2, table 1.

In contrast to supervaluation pairs, Kleene valuation pairs do not preserve classical tautologies and contradictions. Indeed from a well-known property of Kleene's strong logic there are no sentences θ for which $\vec{v}(\theta) = (1, 1)$ for all Kleene valuation pairs. However, the laws of excluded middle and non contradiction are at least partially preserved since $\theta \vee \neg\theta$ is not absolutely false and $\theta \wedge \neg\theta$ is not absolutely true for all Kleene valuation pairs. Furthermore, Kleene valuation pairs do not preserve all classical equivalences. For example, consider a Kleene valuation pair for which $\vec{v}(p_1) = (0, 0)$ and $\vec{v}(p_2) = (0, 1)$. In this case $\vec{v}(p_1 \vee (p_2 \wedge \neg p_2)) = (0, 1) \neq (0, 0) = \vec{v}(p_1)$. However, a number of classical equivalences are preserved [23] including, de Morgan's laws, double negation, idempotence, commutativity, associativity and distributivity.

Notice, that classical valuation pairs can be viewed as special cases of both Kleene and supervaluation pairs in which $\underline{v} = \bar{v}$. Of course many other classes of valuation pairs can also be defined including, for example, valuation pairs related to Lukasiewicz three-valued logic³. However, for the scope of this paper we will focus on supervaluation pairs and Kleene valuation pairs, so as to explore the relationship between them. In the following section we will investigate axiomatic justifications for both these families of valuation pairs in terms of a set of ‘desirable’ properties for lower and upper valuations. Initially, however, we introduce notation for sets of valuation pairs of different classes.

Definition 8. *Classes of Valuation Pairs*

- Let \mathbb{V} denote the set of all valuation pairs on \mathcal{L} .
- Let \mathbb{V}_c denote the set of all classical valuations on \mathcal{L} .
- Let \mathbb{V}_s denote the set of all supervaluation pairs on \mathcal{L} .
- Let $\mathbb{V}_{bs} \subseteq \mathbb{V}_s$ denote the set of all bounded supervaluation pairs on \mathcal{L} .
- Let $\mathbb{V}_{cbs} \subseteq \mathbb{V}_{bs}$ denote the set of complete bounded supervaluation pairs on \mathcal{L} .
- Let \mathbb{V}_k denote the set of Kleene valuation pairs on \mathcal{L} .

3.1 Axioms for Valuation Pairs

Here we introduce a number of elementary properties which it might be considered desirable for a valuation pair to satisfy. The aim is to provide an axiomatic characterisation of some of the classes of valuation pairs introduced in the previous sub-section, in terms of a small number of basic requirements for truth-models of indeterminacy.

VP1 (Duality) $\forall \theta \in S\mathcal{L}, \underline{v}(\neg\theta) = 1 - \bar{v}(\theta)$ and $\bar{v}(\neg\theta) = 1 - \underline{v}(\theta)$.

VP2 (Tautology Preservation) If $\models \theta$ (i.e. θ is a classical tautology) then $\vec{v}(\theta) = (1, 1)$.

VP3 (Equivalence) If $\theta \equiv \varphi$ (i.e. θ and φ are classically equivalent) then $\vec{v}(\theta) = \vec{v}(\varphi)$

VP4 (Maximum Upper) $\forall \theta, \varphi \in S\mathcal{L}, \bar{v}(\theta \vee \varphi) = \max(\bar{v}(\theta), \bar{v}(\varphi))$

VP5 (Maximum Lower) $\forall \theta, \varphi \in S\mathcal{L}, \underline{v}(\theta \vee \varphi) = \max(\underline{v}(\theta), \underline{v}(\varphi))$

VP1 is simply the duality property for negation already discussed above, in which it is required that a sentence is absolutely true if and only if its negation is absolutely

³Lukasiewicz valuation pairs can be defined such that: $\underline{v}(\neg\theta) = 1 - \bar{v}(\theta)$ and $\bar{v}(\neg\theta) = 1 - \underline{v}(\theta)$, $\underline{v}(\theta \wedge \varphi) = \min(\underline{v}(\theta), \underline{v}(\varphi))$ and $\bar{v}(\theta \wedge \varphi) = \max(\min(\underline{v}(\theta), \bar{v}(\varphi)), \min(\bar{v}(\theta), \underline{v}(\varphi)))$ and $\underline{v}(\theta \vee \varphi) = \min(\max(\underline{v}(\theta), \bar{v}(\varphi)), \max(\bar{v}(\theta), \underline{v}(\varphi)))$ and $\bar{v}(\theta \vee \varphi) = \max(\bar{v}(\theta), \bar{v}(\varphi))$

false. VP2 and VP3 require the preservation of classical tautologies and equivalences. Also, notice that when taken together with VP1, VP2 also implies the preservation of classical contradictions. The fundamental idea here is that valuation pairs are simply a generalisation of classical valuations so as to admit indeterminacy, and there is no reason that such a generalisation should result in a change to the underlying logical equivalences between sentences of the language. Furthermore, if a sentence is universally true across all interpretations of the language in which the underlying concepts are assumed to be crisp, then it should also be *absolutely true* in any interpretation which admits borderline cases. This claim has been vigorously contested by proponents of many-valued and infinitely-valued (fuzzy) logic who argue that there is nothing particularly special about the tautologies and equivalences of classical logic requiring them to be accorded special status above those of any other internally consistent logic. Indeed, they argue that admitting additional truth-values fundamentally changes the semantics of the language so that we should not be surprised if the resulting logic has different tautologies and equivalences.

VP4 requires that a disjunction should be not absolutely false when at least one of its disjuncts is not absolutely false. Furthermore, this should be the only manner in which a disjunction is not absolutely false. In other words, a disjunction should be absolutely false *exactly when* both of its disjuncts are absolutely false. Together with VP1, VP4 seems perhaps the least controversial of the properties. VP5 again relates to disjunctions of sentences requiring that a disjunction is absolutely true *exactly when* at least one of its disjuncts is absolutely true. However, there is an inherent tension between this requirement and VP2. To see this suppose that we have a valuation pair for which $\vec{v}(p) = (0, 1)$ and, as would be required by VP1, also $\vec{v}(\neg p) = (0, 1)$. If in this case VP5 holds then $\underline{v}(p \vee \neg p) = 0$ contradicting VP2. Notice there is no problem here with the requirement that a disjunction with at least one absolutely true disjunct is itself absolutely true, but rather with the requirement that this is *the only way* in which a disjunction can be absolutely true. In other words, there is no tension between VP2 and the weaker requirement that $\underline{v}(\theta \vee \varphi) \geq \max(\underline{v}(\theta), \underline{v}(\varphi))$. Finally, notice that assuming VP1 holds together with de Morgan's laws then VP4 and VP5 imply dual constraints on conjunctions of the form $\underline{v}(\theta \wedge \varphi) = \min(\underline{v}(\theta), \underline{v}(\varphi))$ and $\bar{v}(\theta \wedge \varphi) = \min(\bar{v}(\theta), \bar{v}(\varphi))$ respectively.

The following results show that VP1 through to VP4 provide a characterisation of supervaluation pairs.

Theorem 9. *If $\vec{v} \in \mathbb{V}$ satisfies axioms VP1 through to VP4 then $\exists \emptyset \neq \mathcal{P} \subseteq 2^P$ such that $\forall \theta \in S\mathcal{L}$, $\underline{v}(\theta) = \min\{v_F(\theta) : F \in \mathcal{P}\}$ and $\bar{v}(\theta) = \max\{v_F(\theta) : F \in \mathcal{P}\}$ i.e. \vec{v} is a supervaluation pair.*

Proof. Let $\mathcal{P} = \{F_{\alpha_j} : \bar{v}(\alpha_j) = 1\}$. Notice that $\mathcal{P} \neq \emptyset$ since if $\mathcal{P} = \emptyset$ then by the definition of \mathcal{P} we have that $\forall \alpha_j, \bar{v}(\alpha_j) = 0 \Rightarrow$ by definition of valuation pairs $\underline{v}(\alpha_j) = 0$. Therefore, by VP1 we for any atom α_j that $\bar{v}(\neg \alpha_j) = 1$. However, $\neg \alpha_j \equiv \bigvee_{k \neq j} \alpha_k$ and therefore by

VP3 and VP4 we have that

$$\bar{v}(-\alpha_j) = \bar{v}\left(\bigvee_{k \neq j} \alpha_k\right) = \max\{\bar{v}(\alpha_k) : k \neq j\} = 0$$

This is a contradiction.

Now suppose θ is a tautology then by VP2 it holds that $\vec{v}(\theta) = (1, 1)$. Also $\forall F \subseteq P$ it holds that $v_F(\theta) = 1$. Therefore, $\min\{v_F(\theta) : F \in \mathcal{P}\} = \max\{v_F(\theta) : F \in \mathcal{P}\} = 1$ as required. Now suppose θ is not a tautology so that $\mathcal{P}_\theta \neq \emptyset$ and $\mathcal{P}_{-\theta} \neq \emptyset$. We consider three cases:

$$\underline{\vec{v}(\theta) = (0, 0)}$$

$$\vec{v}(\theta) = (0, 0) \Rightarrow (\text{ by VP3 }) \vec{v}\left(\bigvee_{F \in \mathcal{P}_\theta} \alpha_F\right) = (0, 0) \Rightarrow \bar{v}\left(\bigvee_{F \in \mathcal{P}_\theta} \alpha_F\right) = 0$$

$$\Rightarrow (\text{ by VP4 }) \max\{\bar{v}(\alpha_F) : F \in \mathcal{P}_\theta\} = 0 \Rightarrow \forall F \in \mathcal{P}_\theta, \bar{v}(\alpha_F) = 0$$

and hence by definition of \mathcal{P} it follows that $\mathcal{P}_\theta \cap \mathcal{P} = \emptyset \Rightarrow \forall F \in \mathcal{P}, v_F(\theta) = 0$

$$\Rightarrow \min\{v_F(\theta) : F \in \mathcal{P}\} = \max\{v_F(\theta) : F \in \mathcal{P}\} = 0$$

as required.

$$\underline{\vec{v}(\theta) = (0, 1)}$$

Suppose that $\forall F \in \mathcal{P} v_F(\theta) = 1 \Rightarrow \mathcal{P} \subseteq \mathcal{P}_\theta \Rightarrow \forall F \in \mathcal{P}_\theta^c (= \mathcal{P}_{-\theta}), \bar{v}(\alpha_F) = 0$ (by the definition of \mathcal{P}). Now by VP3 and VP4

$$\bar{v}(-\theta) = \bar{v}\left(\bigvee_{F \in \mathcal{P}_{-\theta}} \alpha_F\right) = \max\{\bar{v}(\alpha_F) : F \in \mathcal{P}_{-\theta}\} = 0 \Rightarrow \underline{v}(\theta) = 1 \text{ by axiom VP1}$$

This a contradiction and hence $\exists F \in \mathcal{P}, v_F(\theta) = 0 \Rightarrow \min\{v_F(\theta) : F \in \mathcal{P}\} = 0$ as required.

Alternatively suppose that $\forall F \in \mathcal{P}, v_F(\theta) = 0 \Rightarrow \forall F \in \mathcal{P}, v_F(-\theta) = 1 \Rightarrow \mathcal{P} \subseteq \mathcal{P}_{-\theta} \Rightarrow \forall F \in \mathcal{P}_\theta, \bar{v}(\alpha_F) = 0$ by definition of \mathcal{P} . Now by VP3 and VP4

$$\bar{v}(\theta) = \bar{v}\left(\bigvee_{F \in \mathcal{P}_\theta} \alpha_F\right) = \max\{\bar{v}(\alpha_F) : F \in \mathcal{P}_\theta\} = 0$$

This is a contradiction and therefore $\exists F \in \mathcal{P}, v_F(\theta) = 1 \Rightarrow \max\{v_F(\theta) : F \in \mathcal{P}\} = 1$ as required.

$$\underline{\vec{v}(\theta) = (1, 1)}$$

Suppose $\exists F \in \mathcal{P}$ such that $v_F(\theta) = 0 \Rightarrow \exists F \in \mathcal{P}$ such that $\alpha_F \not\models \theta$ (i.e. $F \in \mathcal{P}_{-\theta}$) $\Rightarrow \exists F \in \mathcal{P} \cap \mathcal{P}_{-\theta} \Rightarrow$ by the definition of \mathcal{P} that $\exists F \in \mathcal{P}_{-\theta}$ such that $\bar{v}(\alpha_F) = 1 \Rightarrow$ by VP3 and VP4 that

$$\bar{v}(-\theta) = \max\{\bar{v}(\alpha_F) : F \in \mathcal{P}_{-\theta}\} = 1$$

$\Rightarrow \underline{v}(\theta) = 0$ by VP1 which is a contradiction. Hence $\forall F \in \mathcal{P}, v_F(\theta) = 1 \Rightarrow$

$$\min\{v_F(\theta) : F \in \mathcal{P}\} = \max\{v_F(\theta) : F \in \mathcal{P}\} = 1$$

as required. □

Corollary 10. *A valuation pair \vec{v} satisfies VP1 to VP4 if and only if \vec{v} is a supervaluation pair*

Proof. From theorem 9 it only remains to show that supervaluation pairs satisfy VP1 to VP4. This follows trivially from the definition of supervaluation pairs. \square

From our earlier discussions it is clear that supervaluation pairs do not, in general, satisfy VP5. This is because they do satisfy VP2 and that there are clearly examples of supervaluation pairs for which $\vec{v}(p) = \vec{v}(\neg p) = (0, 1)$ for some propositional variable p (e.g. take $\mathcal{P} = \{\{p\}, \emptyset\}$).

In the following results we show that, taken together with de Morgan's laws, VP1, VP4 and VP5 provide a characterisation of Kleene valuation pairs.

Theorem 11. *If $\vec{v} \in \mathbb{V}$ satisfies VP1, VP4, VP5 and de Morgan's laws then \vec{v} is a Kleene valuation pair.*

Proof. It is sufficient to show that $\forall \theta, \varphi \in S\mathcal{L}$, $\underline{v}(\theta \wedge \varphi) = \min(\underline{v}(\theta), \underline{v}(\varphi))$ and $\bar{v}(\theta \wedge \varphi) = \min(\bar{v}(\theta), \bar{v}(\varphi))$. By VP1 $\bar{v}(\theta \wedge \varphi) = 1 - \underline{v}(\neg(\theta \wedge \varphi)) = 1 - \underline{v}(\neg\theta \vee \neg\varphi)$ by de Morgan's laws $= 1 - \max(\underline{v}(\neg\theta), \underline{v}(\neg\varphi))$ by VP5 $= 1 - \max(1 - \underline{v}(\theta), 1 - \underline{v}(\varphi))$ by VP1 $= \min(\underline{v}(\theta), \underline{v}(\varphi))$. The result for $\underline{v}(\theta \wedge \varphi)$ follows similarly. \square

Corollary 12. *A valuation pair $\vec{v} \in \mathbb{V}$ satisfies VP1, VP4, VP5 and de Morgan's laws if and only if \vec{v} is a Kleene valuation pair.*

Proof. From theorem 11 and definition 7 we need only show that Kleene valuation pairs satisfy de Morgan's laws. This is proved in [23]. \square

We now define *semantic precision* as a natural partial ordering on \mathbb{V} . This concerns the situation in which one valuation pair admits more borderline cases than another but where otherwise their truth valuations agree. More formally, valuation pair \vec{v}_1 is less semantically precise than \vec{v}_2 if they disagree only for some set of sentences of \mathcal{L} , which being identified as either absolutely true or absolutely false by \vec{v}_2 , are classified as being borderline cases by \vec{v}_1 . In other words, \vec{v}_1 is less semantically precise than \vec{v}_2 if all of the absolutely true and absolutely false valuations of \vec{v}_1 are preserved by \vec{v}_2 .

Definition 13. *Semantic Precision*

$\forall \vec{v}_1, \vec{v}_2 \in \mathbb{V}$, $\vec{v}_1 \preceq \vec{v}_2$ if and only if $\forall \theta \in S\mathcal{L}$, $\underline{v}_1(\theta) \leq \underline{v}_2(\theta)$ and $\bar{v}_1(\theta) \geq \bar{v}_2(\theta)$.

We might now think of the disagreement between two valuation pairs related by \preceq as intuitively being less than that of two unrelated valuation pairs. Indeed, if $\vec{v}_1 \preceq \vec{v}_2$ then another perspective is to think of \vec{v}_2 being generated as a precisification (in the general sense) of \vec{v}_1 by, perhaps even arbitrarily, classifying some of the borderline cases

identified by \vec{v}_1 as either being absolutely true or absolutely false. The following two results show alternative characterisations of semantic precision both for Kleene valuation pairs and supervaluation pairs. In the latter case notice that $\vec{v}_1 \preceq \vec{v}_2$ exactly when the first interpretation is a precisification of the second in the sense of Fine [10]. Indeed, Shapiro [35] proposes a version of this ordering for Kleene's strong logic which he refers to as *sharpening*, where $\vec{v}_1 \preceq \vec{v}_2$ means that \vec{v}_2 extends or sharpens \vec{v}_1 .

Theorem 14. [23] $\forall \vec{v}_1, \vec{v}_2 \in \mathbb{V}_k, \vec{v}_1 \preceq \vec{v}_2$ if and only if $\{p_i \in P : \underline{v}_1(p_i) = 1\} \subseteq \{p_i \in P : \underline{v}_2(p_i) = 1\}$ and $\{p_i \in P : \bar{v}_1(p_i) = 1\} \supseteq \{p_i \in P : \bar{v}_2(p_i) = 1\}$.

Theorem 15. $\forall \vec{v}_1, \vec{v}_2 \in \mathbb{V}_s, \vec{v}_1 \preceq \vec{v}_2$ if and only if $\mathcal{P}_1 \supseteq \mathcal{P}_2$ where \mathcal{P}_1 and \mathcal{P}_2 are the sets of admissible valuations for \vec{v}_1 and \vec{v}_2 respectively.

For theorem 14 the proof follows trivially by induction on the complexity of sentences in SL , while theorem 15 follows immediately from the definition of supervaluation pairs.

3.2 Relating Kleene and Supervaluation Pairs

In this section we investigate the relationship between Kleene valuation pairs and bounded supervaluation pairs. In particular, we will show that for every bounded supervaluation pair there is a less semantically precise Kleene valuation pair with the same values on the sets of entirely positive and entirely negative sentences of \mathcal{L} . The following two initial lemmas prove useful nestedness properties for classical valuations when restricted to sentences in SL^+ and SL^- .

Lemma 16. If $\theta \in SL^+$ and $F \subseteq F' \subseteq P$ then $v_F(\theta) = 1$ implies that $v_{F'}(\theta) = 1$

Proof. Let $SL^{+,0} = P$ and $SL^{+,n} = SL^{+,n-1} \cup \{\theta \wedge \varphi, \theta \vee \varphi : \theta, \varphi \in SL^{+,n-1}\}$. Then by induction on n .

For $p_i \in P, v_F(p_i) = 1 \Rightarrow p_i \in F \Rightarrow p_i \in F' \Rightarrow v_{F'}(p_i) = 1$. Now suppose $\Psi \in SL^{+,n+1}$ then either $\Psi \in SL^{+,n}$ in which case the result follows trivially or $\exists \theta, \varphi \in SL^{+,n}$ and one of the following holds:

- $\Psi = \theta \wedge \varphi$: In this case $v_F(\theta \wedge \varphi) = 1 \Rightarrow v_F(\theta) = 1$ and $v_F(\varphi) = 1 \Rightarrow$ (By induction) $v_{F'}(\theta) = 1$ and $v_{F'}(\varphi) = 1 \Rightarrow v_{F'}(\theta \wedge \varphi) = 1$ as required.
- $\Psi = \theta \vee \varphi$: In this case $v_F(\theta \vee \varphi) = 1 \Rightarrow v_F(\theta) = 1$ or $v_F(\varphi) = 1 \Rightarrow$ (By induction) $v_{F'}(\theta) = 1$ or $v_{F'}(\varphi) = 1 \Rightarrow v_{F'}(\theta \vee \varphi) = 1$ as required.

□

Lemma 17. If $\theta \in SL^-$ and $F' \subseteq F \subseteq P$ then $v_F(\theta) = 1$ implies that $v_{F'}(\theta) = 1$

Proof. The proof mirrors that of lemma 16 but with induction on the complexity of sentences in $S\mathcal{L}^-$ instead of $S\mathcal{L}^+$. \square

The next result shows that for any bounded supervaluation pair its value on strictly positive or strictly negative sentences is determined entirely from its lower and upper admissible valuations. In other words, it has the same value on $S\mathcal{L}^+ \cup S\mathcal{L}^-$ as the bounded supervaluation pair which admits only these lower and upper valuations.

Theorem 18. *Let $\vec{v} \in \mathbb{V}_{bs}$ be a bounded supervaluation pair with lower and upper admissible valuations \underline{F} and \overline{F} and let \vec{v}' be the bounded supervaluation pair for which $\mathcal{P}' = \{\underline{F}, \overline{F}\}$ then $\forall \theta \in S\mathcal{L}^+ \cup S\mathcal{L}^-, \vec{v}(\theta) = \vec{v}'(\theta)$*

Proof. For $\theta \in S\mathcal{L}^+$: If $\vec{v}(\theta) = 1 \Rightarrow \exists F \in \mathcal{P}$ such that $v_F(\theta) = 1 \Rightarrow v_{\overline{F}}(\theta) = 1$ by lemma 16 $\Rightarrow \vec{v}'(\theta) = 1$. Also, if $\vec{v}'(\theta) = 1 \Rightarrow v_{\overline{F}}(\theta) = 1 \Rightarrow \exists F \in \mathcal{P}$ such that $v_F(\theta) = 1 \Rightarrow \vec{v}(\theta) = 1$

If $\underline{v}(\theta) = 1 \Rightarrow \forall F \in \mathcal{P}, v_F(\theta) = 1 \Rightarrow v_{\underline{F}}(\theta) = 1$ and $v_{\overline{F}}(\theta) = 1 \Rightarrow \underline{v}'(\theta) = 1$. Also, if $\underline{v}'(\theta) = 1 \Rightarrow v_{\underline{F}}(\theta) = 1 \Rightarrow$ by lemma 16 $\forall F \in \mathcal{P}, v_F(\theta) = 1 \Rightarrow \underline{v}(\theta) = 1$

For $\theta \in S\mathcal{L}^-$ the result follows similarly from lemma 17 by swapping the roles of \underline{F} and \overline{F} in the above argument. \square

Notice, however, it is not generally the case that for $\theta \in S\mathcal{L} - (S\mathcal{L}^+ \cup S\mathcal{L}^-)$, $\vec{v}(\theta) = \vec{v}'(\theta)$ as can be seen from the following example:

Example 19. *Let $\mathcal{P} = \{\{p_1\}, \{p_1, p_2\}, \{p_1, p_2, p_3\}\}$ so that $\mathcal{P}' = \{\{p_1\}, \{p_1, p_2, p_3\}\}$. In this case $\vec{v}(p_2 \wedge \neg p_3) = 1$ while $\vec{v}'(p_2 \wedge \neg p_3) = 0$*

We now show that for entirely positive and entirely negative sentences, bounded supervaluation pairs satisfy property VP5 i.e. that a disjunction is absolutely true exactly when at least one of its disjuncts is absolutely true. Furthermore, in this case we also have the dual result that a conjunction is not absolutely false exactly when both of its conjuncts are not absolutely false.

Theorem 20. *Let $\vec{v} \in \mathbb{V}_{bs}$ be a bounded supervaluation pair then $\forall \theta, \varphi \in S\mathcal{L}^+$ and $\forall \theta, \varphi \in S\mathcal{L}^-$ it holds that:*

$$\vec{v}(\theta \wedge \varphi) = \min(\vec{v}(\theta), \vec{v}(\varphi)) \text{ and } \underline{v}(\theta \vee \varphi) = \max(\underline{v}(\theta), \underline{v}(\varphi))$$

Proof. As before let $\mathcal{P}' = \{\underline{F}, \overline{F}\}$ then for $\theta, \varphi \in S\mathcal{L}^+$, $\vec{v}(\theta \wedge \varphi) = 1$ iff $\vec{v}'(\theta \wedge \varphi) = 1$ by theorem 18 iff $v_{\overline{F}}(\theta \wedge \varphi) = 1$ by lemma 16 iff $v_{\overline{F}}(\theta) = 1$ and $v_{\overline{F}}(\varphi) = 1$ iff $\vec{v}'(\theta) = 1$ and $\vec{v}'(\varphi) = 1$ by lemma 16 iff $\vec{v}(\theta) = 1$ and $\vec{v}(\varphi) = 1$ by theorem 18 as required.

Similarly, $\underline{v}(\theta \vee \varphi) = 1$ iff $\underline{v}'(\theta \vee \varphi) = 1$ by theorem 18, iff $v_{\underline{F}}(\theta \vee \varphi) = 1$ by lemma 16 iff $v_{\underline{F}}(\theta) = 1$ or $v_{\underline{F}}(\varphi) = 1$ iff $\underline{v}'(\theta) = 1$ or $\underline{v}'(\varphi) = 1$ by lemma 16 iff $\underline{v}(\theta) = 1$ or $\underline{v}(\varphi) = 1$ by theorem 18 as required.

For $\theta, \varphi \in S\mathcal{L}^-$ the result follows similarly from lemma 17 by swapping the roles of \underline{F} and \overline{F} in the above argument. \square

Theorem 21. *Let $\vec{v}_{bs} \in \mathbb{V}_{bs}$ be a bounded supervaluation pair then there exists a unique Kleene valuation pair \vec{v}_k such that $\vec{v}_k \preceq \vec{v}_{bs}$ and $\forall \theta \in S\mathcal{L}^+ \cup S\mathcal{L}^-, \vec{v}_k(\theta) = \vec{v}_{bs}(\theta)$.*

Proof. We define $\vec{v}_k \in \mathbb{V}_k$ such that $\forall p_i \in P$ let $\underline{v}_k(p_i) = 1$ iff $p_i \in \underline{F}$ and $\overline{v}_k(p_i) = 1$ iff $p_i \in \overline{F}$, where \underline{F} and \overline{F} are the lower and upper sets of admissible valuations for \vec{v}_{bs} .

Firstly, we show that $\forall \theta \in S\mathcal{L}^+ \cup S\mathcal{L}^-, \vec{v}_k(\theta) = \vec{v}_{bs}(\theta)$. Here we prove this result for sentences in $S\mathcal{L}^+$, and the result for $S\mathcal{L}^-$ follows similarly. The proof is by induction as follows. For $p_i \in P$ the result follows trivially. Now suppose $\Psi \in S\mathcal{L}^{+,n+1}$ then either $\Psi \in S\mathcal{L}^{+,n}$ in which case the result follows trivially or $\exists \theta, \varphi \in S\mathcal{L}^{+,n}$ and one of the following holds:

- $\Psi = \theta \wedge \varphi$: $\underline{v}_k(\theta \wedge \varphi) = \min(\underline{v}_k(\theta), \underline{v}_k(\varphi))$ by definition of Kleene valuation pairs = $\min(\underline{v}_{bs}(\theta), \underline{v}_{bs}(\varphi))$ by the inductive hypothesis = $\underline{v}_{bs}(\theta \wedge \varphi)$ by a basic property of supervaluation pairs. Also, $\overline{v}_k(\theta \wedge \varphi) = \min(\overline{v}_k(\theta), \overline{v}_k(\varphi))$ by definition of Kleene valuation pairs = $\min(\overline{v}_{bs}(\theta), \overline{v}_{bs}(\varphi))$ by the inductive hypothesis = $\overline{v}_{bs}(\theta \wedge \varphi)$ by theorem 20.
- $\Psi = \theta \vee \varphi$: $\underline{v}_k(\theta \vee \varphi) = \max(\underline{v}_k(\theta), \underline{v}_k(\varphi))$ by definition of Kleene valuation pairs = $\max(\underline{v}_{bs}(\theta), \underline{v}_{bs}(\varphi))$ by the inductive hypothesis = $\underline{v}_{bs}(\theta \vee \varphi)$ by theorem 20. Also, $\overline{v}_k(\theta \vee \varphi) = \max(\overline{v}_k(\theta), \overline{v}_k(\varphi))$ by definition of Kleene valuation pairs = $\max(\overline{v}_{bs}(\theta), \overline{v}_{bs}(\varphi))$ by the inductive hypothesis = $\overline{v}_{bs}(\theta \vee \varphi)$ by a basic property of supervaluation pairs.

To see that \vec{v}_k defined as above is the only Kleene valuation agreeing with \vec{v}_{bs} for all sentences in $S\mathcal{L}^+ \cup S\mathcal{L}^-$ note that if $\forall p_i \in P, \vec{v}_k(p_i) = \vec{v}_{bs}(p_i)$ then $\underline{v}_k(p_i) = 1$ iff $\underline{v}_{bs}(p_i) = 1$ iff $p_i \in \underline{F}$ and $\overline{v}_k(p_i) = 1$ iff $\overline{v}_{bs}(p_i) = 1$ iff $p_i \in \overline{F}$.

We now show that $\vec{v}_k \preceq \vec{v}_{bs}$, for \vec{v}_k defined as above. The proof is by induction as follows: If $\Psi \in S\mathcal{L}^0$ then $\Psi = p_i \in P$. Clearly, by the definition of $\vec{v}_k, \vec{v}_k(p_i) = \vec{v}_{bs}(p_i)$ and so the result follows trivially in this case. Now suppose $\Psi \in S\mathcal{L}^{n+1}$ then either $\Psi \in S\mathcal{L}^n$ in which case the result follows trivially or $\exists \theta, \varphi \in S\mathcal{L}^n$ such that one of the following holds:

- $\Psi = \theta \wedge \varphi$: In this case, if $\underline{v}_k(\theta \wedge \varphi) = 1$ then $\min(\underline{v}_k(\theta), \underline{v}_k(\varphi)) = 1$ which implies that $\underline{v}_k(\theta) = 1$ and $\underline{v}_k(\varphi) = 1$. Hence, by induction, $\underline{v}_{bs}(\theta) = 1$ and $\underline{v}_{bs}(\varphi) = 1$ which implies that $\underline{v}_{bs}(\theta \wedge \varphi) = 1$ as required. Also, if $\overline{v}_{bs}(\theta \wedge \varphi) = 1$ then $\overline{v}_{bs}(\theta) = 1$ and $\overline{v}_{bs}(\varphi) = 1$ which implies, by induction, that $\overline{v}_k(\theta) = 1$ and $\overline{v}_k(\varphi) = 1$. Hence, $\min(\overline{v}_k(\theta), \overline{v}_k(\varphi)) = \overline{v}_k(\theta \wedge \varphi) = 1$ as required.

- $\Psi = \theta \vee \varphi$: In this case, if $\underline{v}_k(\theta \vee \varphi) = 1$ then $\max(\underline{v}_k(\theta), \underline{v}_k(\varphi)) = 1$ which implies that $\underline{v}_k(\theta) = 1$ or $\underline{v}_k(\varphi) = 1$. Hence, by induction, $\underline{v}_{bs}(\theta) = 1$ or $\underline{v}_{bs}(\varphi) = 1$ which implies that $\underline{v}_{bs}(\theta \vee \varphi) = 1$ as required. Also, if $\bar{v}_{bs}(\theta \vee \varphi) = 1$ then $\bar{v}_{bs}(\theta) = 1$ or $\bar{v}_{bs}(\varphi) = 1$ which implies, by induction, that $\bar{v}_k(\theta) = 1$ or $\bar{v}_k(\varphi) = 1$. Hence, $\max(\bar{v}_k(\theta), \bar{v}_k(\varphi)) = \bar{v}_k(\theta \vee \varphi) = 1$ as required.
- $\Psi = \neg\theta$: In this case, if $\underline{v}_k(\neg\theta) = 1$ then $\bar{v}_k(\theta) = 0$, which implies, by induction, that $\bar{v}_{bs}(\theta) = 0$. Hence, $\underline{v}_{bs}(\neg\theta) = 1$ as required. Also, if $\bar{v}_{bs}(\neg\theta) = 1$ then $\underline{v}_{bs}(\theta) = 0$ which implies, by induction, that $\underline{v}_k(\theta) = 0$. Hence, $\bar{v}_k(\neg\theta) = 1$ as required.

□

There are a number of different possible viewpoints from which we can interpret theorem 21. One perspective might be that in some sense bounded supervaluation pairs provide a better model of truth-gaps than Kleene valuation pairs, perhaps on the basis of the axiomatic characterisations given in section 3.1. In this case we could view the Kleene valuation pair \vec{v}_k identified in theorem 21 as a less semantically precise approximation of \vec{v}_{bs} , agreeing with \vec{v}_{bs} on an important subclass of sentences. Such an approximation might be advantageous on computational grounds because Kleene valuation pairs are truth-functional, while bounded supervaluation pairs are not even functional in a weaker sense. We will consider this issue in more detail in the following section.

An alternative perspective would be to simply view Kleene valuation pair \vec{v}_k as a more vague interpretation of \mathcal{L} for which theorem 21 identifies a set of bounded supervaluation pairs as different precisifications. More formally, given Kleene valuation pair $\vec{v}_k \in \mathbb{V}_k$ for which $\underline{F} = \{p_i : \underline{v}_k(p_i) = 1\}$ and $\bar{F} = \{p_i : \bar{v}_k(p_i) = 1\}$ then we can identify the class of bounded supervaluation pairs with lower and upper admissible valuations \underline{F} and \bar{F} , given by:

$$\llbracket \underline{F}, \bar{F} \rrbracket = \{\vec{v}_{bs} \in \mathbb{V}_{bs} : \{p_i : \underline{v}_{bs}(p_i) = 1\} = \underline{F}, \{p_i : \bar{v}_{bs}(p_i) = 1\} = \bar{F}\}$$

Now theorem 21 identifies the same Kleene valuation pair \vec{v}_k for each bounded supervaluation pair in $\llbracket \underline{F}, \bar{F} \rrbracket$. Hence, we might think of $\llbracket \underline{F}, \bar{F} \rrbracket$ as a natural class of precisifications of \vec{v}_k .

Interestingly, the class $\llbracket \underline{F}, \bar{F} \rrbracket$ also provides us with some insight into the relationship between bounded supervaluation pairs and complete bounded supervaluation pairs. Notice that as $\underline{F} \subseteq \bar{F}$ range across all possible such pairs of subsets of P , then the sets $\llbracket \underline{F}, \bar{F} \rrbracket$ form a partition of \mathbb{V}_{bs} . Also, each complete bounded supervaluation pair $\vec{v}_{cbs} \in \mathbb{V}_{cbs}$ is a member of exactly one of these sets. In fact, if \vec{v}_{cbs} has lower and upper admissible valuations \underline{F} and \bar{F} then $\vec{v}_{cbs} \in \llbracket \underline{F}, \bar{F} \rrbracket$ and by theorem 15 $\forall \vec{v}_{bs} \in \llbracket \underline{F}, \bar{F} \rrbracket, \vec{v}_{cbs} \preceq \vec{v}_{bs}$. Hence, $\forall \theta \in S\mathcal{L}$,

$$\underline{v}_{cbs}(\theta) = \min\{\underline{v}_{bs}(\theta) : \vec{v}_{bs} \in \llbracket \underline{F}, \bar{F} \rrbracket\} \text{ and } \bar{v}_{cbs}(\theta) = \max\{\bar{v}_{bs}(\theta) : \vec{v}_{bs} \in \llbracket \underline{F}, \bar{F} \rrbracket\}$$

Also notice by theorem 18 that \vec{v}_{obs} agrees with all the valuation pairs in $\llbracket \underline{F}, \overline{F} \rrbracket$ on the sentences in $S\mathcal{L}^+ \cup S\mathcal{L}^-$. Hence, in summary, we have that for every partition set of bounded supervaluation pairs $\llbracket \underline{F}, \overline{F} \rrbracket$ there is both a unique Kleene valuation pair and a unique complete bounded supervaluation pair which are less semantically precise than all the valuation pairs in this set, but which agree with all of them on the sentences of $S\mathcal{L}^+ \cup S\mathcal{L}^-$.

In the following section we consider the representational power of supervaluation pairs, bounded and complete bounded supervaluation pairs, and Kleene valuation pairs, with particular regard to penumbral connections.

3.3 Penumbral Connections

Fine [10] highlights the capacity to capture penumbral connections as a significant advantage of supervaluationist theories of vagueness over many-valued logics. Penumbral connections are defined to be those ‘logical relations [that] hold between indefinite sentences’ [10]. In particular, given a set of borderline literals, penumbral connections may ensure that certain conjunctions or disjunctions of these literals are nonetheless absolutely true or absolutely false. For example, suppose we have a sequence of heights $h_1 < h_2 < \dots < h_n$ where only h_1 is classed as being absolutely not tall, and only h_n is absolutely tall, so that all other heights are borderline cases of tall. However, despite their borderline status, if we were to learn that h_i was indeed tall for any $i \in \{2, \dots, n-1\}$ we would immediately infer that h_{i+1} was also tall, simply because $h_{i+1} > h_i$. More formally, let p_i denote the proposition ‘a person of height h_i is tall’ then we can capture the above penumbral connections by the supervaluation pair where $\mathcal{P} = \{\{p_2, \dots, p_n\}, \dots, \{p_i, \dots, p_n\}, \dots, \{p_{n-1}, p_n\}, \{p_n\}\}$. In this case, $\vec{v}(p_1) = (0, 0)$, $\vec{v}(p_n) = (1, 1)$ and $\vec{v}(p_i) = (0, 1)$ for $i = 2, \dots, n-1$. Furthermore, $\vec{v}(p_i \rightarrow p_{i+1}) = (1, 1)$ ⁴ for $i = 1, \dots, n-1$. Notice that this is a bounded supervaluation pair with $\underline{F} = \{p_n\}$ and $\overline{F} = \{p_2, \dots, p_n\}$ but it is not a complete bounded supervaluation pair. Indeed, the complete bounded supervaluation pair generated by \underline{F} and \overline{F} as above, would lose the semantic information about the penumbral connections between p_1, \dots, p_n since it would, by definition, contain other admissible valuations consistent with $p_i \wedge \neg p_j$ for $j > i$.

Another type of penumbral connections concerns two propositional variables which are exclusive and exhaustive, but where both are borderline cases. Fine [10] uses the example of colours red and pink, suggesting that a given object on the borderline of these two colours must nonetheless be described as one or the other but not both. Suppose then that we have two borderline propositions p_1 and p_2 , for which $p_1 \wedge p_2$ is absolutely false and $p_1 \vee p_2$ is absolutely true. In order for a supervaluation pair to capture this semantic information we would require the following; 1) $\forall F \in \mathcal{P}$ either $p_1 \in F$ or $p_2 \in F$

⁴Here $p_i \rightarrow p_j$ is used as shorthand for $\neg p_i \vee p_j$.

but $\{p_1, p_2\} \not\subseteq F$, and 2) $\exists F, F' \in \mathcal{P}$ such that $p_1 \in F$ and $p_2 \in F'$. However, a set of admissible valuations with these properties cannot define a bounded supervaluation pair. To see this notice that taking $\overline{F} = \{p_i \in P : \bar{v}(p_i) = 1\}$, it follows that $\{p_1, p_2\} \subseteq \overline{F}$ and therefore $\overline{F} \notin \mathcal{P}$.

The above examples provide some illustration of the loss of representational power resulting from restricting general supervaluation pairs to either bounded or complete bounded supervaluation pairs. Whilst corollary 10 means that any penumbral connections resulting from classical equivalences, tautologies or contradictions can be captured by all supervaluation pairs, restricting ourselves to bounded and then complete bounded supervaluation pairs successively limits the type of penumbral connections which can be captured. Kleene valuation pairs have even less capacity for representing the type of penumbral connections we have outlined here. Indeed, even connections represented by classical equivalences, tautologies and contradictions cannot in general be captured. For example, there is no Kleene valuation pair where $\vec{v}(\theta) = \vec{v}(\neg\theta) = (0, 1)$ but where $\vec{v}(\theta \wedge \neg\theta) = (0, 0)$ and $\vec{v}(\theta \vee \neg\theta) = (1, 1)$. In the next section, however, we will argue that, despite the reduction in capacity for representing penumbral connections, there are some computational advantage of restricting ourselves to bounded or complete bounded supervaluation pairs, or indeed even to Kleene valuation pairs

3.4 Functionality and Truth-Functionality of Valuation Pairs

Truth-functionality or compositionality is a fundamental aspect of most classical formal systems. The truth value of any compound expression, it is assumed, can be determined from the truth-values of its components by means of a recursive process dependent of the logical structure of that expression. This provides a computationally feasible method for determining truth-values with a computational cost bounded by the number of independent occurrences of connectives in the given formula. Furthermore, an agent is only required to identify truth-values for the primitives in the language; in this case the propositional variables.

In this section we consider the issue of truth-functionality for valuation pairs together with a weaker form of functionality. For the latter, rather than assuming that valuations are determined from mappings based directly around the connectives of the language, we require only that, for each sentence θ , a function can be identified which determines the valuation of θ on the basis of the valuations for the propositional variables only. Such an assumption restricts the necessity of direct truth evaluations to the propositional variables which can at least potentially limit an agent's storage (or memory) requirements. However, as we shall see the computational cost of evaluating a sentence can, in the worst case, remain high.

The following are formal definitions of both functionality and truth-functionality for a

class $\mathbb{V}' \subseteq \mathbb{V}$ of valuation pairs.

Definition 22. *Functionality*

Let $\mathbb{T} = \{(0,0), (0,1), (1,1)\}$. Then a class of valuation pairs $\mathbb{V}' \subseteq \mathbb{V}$ is functional if $\forall \theta \in S\mathcal{L}$ there exists a function $f_\theta : \mathbb{T}^n \rightarrow \mathbb{T}$ such that:

$$\forall \vec{v} \in \mathbb{V}', \vec{v}(\theta) = f_\theta(\vec{v}(p_1), \dots, \vec{v}(p_n))$$

Definition 23. *Truth-Functionality*

A class of valuation pairs $\mathbb{V}' \subseteq \mathbb{V}$ is truth-functional if there are functions $f_\wedge : \mathbb{T}^2 \rightarrow \mathbb{T}$, $f_\vee : \mathbb{T}^2 \rightarrow \mathbb{T}$ and $f_- : \mathbb{T} \rightarrow \mathbb{T}$ such that $\forall \theta, \varphi \in S\mathcal{L}$ and $\forall \vec{v} \in \mathbb{V}'$:

$$\vec{v}(-\theta) = f_-(\vec{v}(\theta)), \vec{v}(\theta \wedge \varphi) = f_\wedge(\vec{v}(\theta), \vec{v}(\varphi)), \text{ and } \vec{v}(\theta \vee \varphi) = f_\vee(\vec{v}(\theta), \vec{v}(\varphi))$$

Naturally any class of valuation pairs which are truth-functional are also functional. For example, Kleene valuation pairs are truth-functional. However, the following example shows that in general supervaluation pairs are not functional.

Example 24. Let $P = \{p_1, p_2, p_3\}$ and \vec{v}_1 and \vec{v}_2 be supervaluation pairs with sets of admissible valuations $\mathcal{P}_1 = \{\{p_1\}, \{p_2\}, \{p_3\}\}$ and $\mathcal{P}_2 = \{\{p_1, p_2\}, \{p_3\}\}$ respectively. Now for both \vec{v}_1 and \vec{v}_2 we have that $\vec{v}_1(p_1) = \vec{v}_2(p_1) = (0, 1)$, $\vec{v}_1(p_2) = \vec{v}_2(p_2) = (0, 1)$ and $\vec{v}_1(p_3) = \vec{v}_2(p_3) = (0, 1)$. But $\vec{v}_1(p_1 \wedge p_2) = (0, 0)$ while $\vec{v}_2(p_1 \wedge p_2) = (0, 1)$.

In fact it is not even the case that the restricted class of bounded supervaluation pairs are functional as can be seen from the following example:

Example 25. Let $P = \{p_1, p_2, p_3\}$ and \vec{v}_1 and \vec{v}_2 be bounded supervaluation pairs with sets of admissible valuations $\mathcal{P}_1 = \{\{p_1\}, \{p_1, p_2\}, \{p_1, p_2, p_3\}\}$ and $\mathcal{P}_2 = \{\{p_1\}, \{p_1, p_3\}, \{p_1, p_2, p_3\}\}$ respectively. For both \vec{v}_1 and \vec{v}_2 we have that $\vec{v}_1(p_1) = \vec{v}_2(p_1) = (1, 1)$, $\vec{v}_1(p_2) = \vec{v}_2(p_2) = (0, 1)$ and $\vec{v}_1(p_3) = \vec{v}_2(p_3) = (0, 1)$. However, for $\theta = p_1 \wedge p_2 \wedge \neg p_3$ we have that $\vec{v}_1(\theta) = (0, 1)$ while $\vec{v}_2(\theta) = (0, 0)$.

On the other hand, complete bounded supervaluation pairs are functional as we can see from the following argument: If $\vec{v} \in \mathbb{V}_{cbs}$ then if we know $\vec{v}(p_i)$ for $i = 1, \dots, n$ then we can determine the lower and upper admissible valuations according to:

$$\underline{F} = \{p_i : \underline{v}(p_i) = 1\} \text{ and } \overline{F} = \{p_i : \overline{v}(p_i) = 1\}$$

and consequently we can define for any $\theta \in S\mathcal{L}$, $f_\theta : \mathbb{T}^n \rightarrow \mathbb{T}$ such that:

$$f_\theta(\vec{v}(p_1), \dots, \vec{v}(p_n)) = (\min\{v_F(\theta) : \underline{F} \subseteq F \subseteq \overline{F}\}, \max\{v_F(\theta) : \underline{F} \subseteq F \subseteq \overline{F}\})$$

However, as the next result shows, no class of supervaluation pairs which strictly extends classical valuations is truth-functional.

Theorem 26. *Let $\mathbb{V}_c \subset \mathbb{V}' \subseteq \mathbb{V}_s$ then the class of supervaluation pairs \mathbb{V}' is not truth-functional.*

Proof. Since $\mathbb{V}' \supset \mathbb{V}_c$ then there exists $\vec{v} \in \mathbb{V}'$ with admissible valuations $\mathcal{P} \supseteq \{F, F'\}$ where $F \neq F'$. Now w.l.o.g we can assume that $\exists p_i \in F$ such that $p_i \notin F'$. Hence, $\vec{v}(p_i) = (0, 1)$, $\vec{v}(\neg p_i) = (0, 1)$ but $\vec{v}(p_i \wedge p_i) = (0, 1)$ while $\vec{v}(p_i \wedge \neg p_i) = (0, 0)$. Hence, \mathbb{V}' is not truth-functional. \square

Given all of the above we can now consider the computational cost of truth evaluations for different classes of valuation pairs, from the perspective of the functionality and truth-functionality properties. If a sentence θ has m distinct occurrences of the connectives then for any truth-functional class \mathbb{V}' the computational cost of evaluating $\vec{v}(\theta)$ is $\mathcal{O}(m)$ ⁵. For non truth-functional \mathbb{V}' this cost can be much higher. For example, if $\vec{v} \in \mathbb{V}_s$ then the cost of evaluating $\vec{v}(\theta)$ is, in the worst case, $\mathcal{O}(m|\mathcal{P}|)$ i.e. the combined cost of evaluating the truth value of θ for every admissible classical valuation in \mathcal{P} .⁶ Indeed, this is also true for complete bounded supervaluation pairs, even though their functionality means that \mathcal{P} does not need to be stored directly but can be determined from the valuations on propositional variables. However, these are worst case costs and in many situations evaluating valuation pairs can be much more computationally feasible. For instance, notice that from theorem 20 it follows that for sentences restricted to $S\mathcal{L}^+$ or $S\mathcal{L}^-$ bounded supervaluation pairs are effectively truth-functional with $f_\wedge = \min$ and $f_\vee = \max$. Also, recall from theorem 21 that any bounded supervaluation pair can be ‘approximated’ by a semantically less precise but fully truth-functional Kleene valuation pair.

4 Bipolar Belief Measures

Recall, as discussed in section 2, that there is a clear distinction between indeterminism and epistemic uncertainty. For vague concepts, truth-gaps are an integral part of the underlying truth-model and do not result from lack of knowledge about any of the relevant attributes involved in the concept definitions. In particular, borderline cases do not model uncertainty, and indeed an agent can be certain that a particular proposition is a borderline case. For example, a given height may definitely be a borderline case of the predicate *tall*. Similarly, learning that the proposition ‘Bill is tall’ is borderline is informative since it suggests that Bill’s height is a borderline case of the concept tall and hence provides information about its actual value. On the other hand, learning that the truth-value of ‘Bill is tall’ is uncertain or unknown provides no information about Bill’s height. Hence, from

⁵This is of course assuming no cost for evaluating the connective functions f_\neg , f_\wedge and f_\vee

⁶This worst case occurs when determining the valuation pair values of a sentence θ which is either absolutely true or absolutely false. In such a case the truth value of θ must be calculated for every admissible valuation in \mathcal{P} , with a cost for each valuation dependent on the number of connectives in θ .

this perspective, valuation pairs provide a truth-model incorporating truth-gaps and not a three-valued representation of uncertainty. This is in contrast to a significant proportion of the literature on three-valued logic which tends to use the third truth-value to represent ‘unknown’ [6].

In view of the above distinction, it is natural to propose a combined model incorporating both indeterminism and epistemic uncertainty in order to represent an agent’s subjective beliefs concerning sentences which involve vague concepts. In the sequel we will argue that this model should be characterised by a probability distribution over a set of possible valuations. This will then result in lower and upper (bipolar) measures of subjective belief on the sentences of \mathcal{L} , as outlined below. Now, in practice, there are many different sources of epistemic uncertainty that will influence an agent’s belief about $S\mathcal{L}$. However, here we suggest that one natural division of uncertainty types is as follows:

- *Semantic Uncertainty*: This takes the form of uncertainty about the linguistic conventions governing the assertability of sentences of \mathcal{L} . In other words, uncertainty about what is the *correct* interpretation of \mathcal{L} . For example, an agent may be uncertain as to whether or not a proposition such as ‘Ethel is tall’ is definitely or acceptably assertable even if they know Ethel’s height precisely. One would expect this type of uncertainty to be a natural consequence of the distributed manner in which an agent learns language through communications with other individuals across a population of interacting agents.
- *Possible Worlds Uncertainty*: This type of uncertainty arises from a lack of knowledge concerning the current state of the world and in particular about the referents of sentences in \mathcal{L} . For example, an agent may not know Ethel’s height precisely and hence be uncertain about the truth value of the proposition ‘Ethel is tall’. Another way of viewing this is that the agent has identified a set of possible worlds, perhaps corresponding to different possible values for Ethel’s height, but where he/she is uncertain as to which of these correspond to the actual world.

The second of these uncertainty types would seem to be consistent with the conventional understanding of epistemic uncertainty. In contrast, the treatment of semantic uncertainty as a form of epistemic uncertainty perhaps requires some justification. In our view semantic uncertainty is another important aspect of vagueness in language, but one which is best understood in terms of a lack of knowledge concerning the underlying conventions of the language. By adopting this view we are, in effect, assuming that each agent believes in the existence of a coherent set of rules governing the assertability of the sentences of \mathcal{L} , which they should adhere to in order to communicate effectively with other agents. In other words, agents make the assumption that there is a *correct* interpretation of the language \mathcal{L} ⁷. This would seem to bring us close to the epistemic theory of vagueness as

⁷although there is no requirement that this should be a classical interpretation

expounded by Timothy Williamson [44]. The epistemic theory of vagueness assumes that for the extension of a vague concept there is a precise but unknown dividing boundary between it and the extension of its negation. For example, consider the set of heights which are classified as being tall, then according to the epistemic theory there is a precise but unknown height threshold ϵ for which all heights less than ϵ are not tall and all those greater than or equal to ϵ are tall.

While there are marked similarities between the epistemic theory of vagueness and the notion of semantic uncertainty outlined above, there are also some important differences. In the first instance the epistemic view assumes that the underlying interpretation of the language is classical, while in the sequel we will assume that it may admit truth-gaps by, for example, taking the form of a valuation pair. Perhaps more importantly though, the epistemic theory would seem to assume the existence of some *objectively correct* interpretation of \mathcal{L} which is not necessarily correlated with language use.⁸ Instead, we adopt the more pragmatic view that individuals, when faced with decision problems about assertions, find it useful as part of decision making strategy to simply *assume* that there is a correct interpretation of \mathcal{L} . In other words, when deciding what can be asserted agents behave *as if* the epistemic theory is correct. In earlier work we have referred to this strategic assumption across a population of agents as the *epistemic stance* [19], a concise statement of which is as follows:

Each individual agent in the population assumes the existence of a correct set of language conventions, governing what can appropriately (or truthfully) be asserted given a particular state of the world.

Of course, in practice such conventions would not be imposed by some outside authority. Indeed, they may not exist at all in a single centralised form. Rather they are represented as a distributed body of knowledge concerning the interpretation of concepts in various cases, shared across a population of agents, and emerging as the result of interactions and communications between individual agents all adopting the epistemic stance. The idea is that the learning processes of individual agents, all sharing the fundamental aim of understanding how words can be appropriately used to communicate information, will eventually converge to some degree on a set of shared conventions.⁹ The very process of convergence would then to some extent vindicate the epistemic stance from the perspective of individual agents.

⁸While this is true of the epistemic theory in general, it is perhaps worth noting that Williamson does assume that the correctness of one particular interpretation *is* determined by the usage of speakers. Many thanks to one of the anonymous referees for this observation.

⁹This kind of convergent process has been shown to take place in various agent-based simulations of language games. See, for example, Steels and Belpaeme [37].

4.1 An Integrated Probabilistic Uncertainty Model

Here we propose an integrated treatment of both the types of epistemic uncertainty described above, but with the additional assumption that interpretations of \mathcal{L} may admit truth-gaps i.e. that they correspond to valuation pairs. Let \mathcal{S} denote the different possible states of the world (or alternatively the set of possible worlds). A language convention c is then viewed as a function which maps from states of the world to truth valuations on \mathcal{L} . Hence, in the current context, a language convention is a function $c : \mathcal{S} \rightarrow \mathbb{V}$ from states of the world to valuation pairs. Let \mathcal{C} denote the set of such conventions. Naturally, further restrictions of this definition may then be considered such as $c : \mathcal{S} \rightarrow \mathbb{V}_s$ or $c : \mathcal{S} \rightarrow \mathbb{V}_k$ depending on the class of valuation pairs considered as providing the appropriate form of truth-model for \mathcal{L} . For example, a convention c would be associated with a particular pair of lower and upper thresholds $\underline{\epsilon} \leq \bar{\epsilon}$ on height for the concept short. Whereas a state of the world $\sigma \in \mathcal{S}$ would be associated with a particular height value h for Ethel. The proposition ‘Ethel is short’ would then be absolutely true for those convention state pairs (c, σ) for which $h \leq \underline{\epsilon}$ and not necessarily false in those for which $h \leq \bar{\epsilon}$. From this perspective, sentences of the language relate to the objective state of the world, but this relationship is based on linguistic conventions which may be vague in the sense of admitting borderline cases. Hence the truth (or falsity) of propositions depends on both the state of the world and the linguistic conventions according to which the relevant concepts are defined relative to the different possible states of the world \mathcal{S} . A model of epistemic uncertainty integrating both the uncertainty types can therefore be defined in terms of a probability measure P on subsets of the cross product space $\mathcal{C} \times \mathcal{S}$ of linguistic conventions with world states. Such a model naturally generates a probability distribution on valuation pairs given by:

$$w(\vec{v}) = P(\{(c, \sigma) : c(\sigma) = \vec{v}\})$$

From this we can naturally define lower and upper measures $\underline{\mu} \leq \bar{\mu}$ such that for $\theta \in S\mathcal{L}$, $\underline{\mu}(\theta)$ is the probability of a valuation pair in which θ is absolutely true, while $\bar{\mu}(\theta)$ is the probability of a valuation pair in which θ is not absolutely false. That is:

$$\underline{\mu}(\theta) = w(\{\vec{v} : \underline{v}(\theta) = 1\}) \text{ and } \bar{\mu}(\theta) = w(\{\vec{v} : \bar{v}(\theta) = 1\})$$

This formulation would seem to have significant overlap with that of Lassiter [18], who has proposed a model of linguistic vagueness defined on the cross-product space of possible worlds with possible languages. However, Lassiter’s approach only refers to uncertainty about ‘precise languages’ suggesting that he is assuming a classical underlying truth-model. In this respect, his work is strongly related to probabilistic interpretations of fuzzy sets (see [8] for an overview), in which a fuzzy concept is defined by a probability distribution on possible crisp set definitions. On the other hand, we argue that a bipolar model combining epistemic uncertainty with truth-gaps provides extra flexibility in decision making as well as helping to explain why it may sometimes be optimal for an agent to choose to assert a

vague statement rather than a semantically similar crisp one. In section 5 we will expand on this claim, however, initially we formally develop the notion of bipolar belief pairs.

4.2 Belief Pairs from Probability Distributions on Valuation Pairs

In this section we consider bipolar belief measures on $S\mathcal{L}$ as generated from probability distributions defined on valuation pairs. The motivation for this approach comes from the idea outlined in section 4.1 that epistemic uncertainty has both a semantic and stochastic element, naturally resulting in probability distributions on valuation pairs when these provide the possible valuations of the language \mathcal{L} . In particular, we will focus on Kleene and supervaluation pairs and investigate some of the resulting bipolar uncertainty measures on $S\mathcal{L}$ and their properties.

Definition 27. Supervaluation Belief Pair

Let $w : \mathbb{V}_s \rightarrow [0, 1]$ be a probability distribution on \mathbb{V}_s (i.e. the set of all supervaluation pairs on \mathcal{L}) then the Supervaluation Belief Pair (SBP) generated by w is a pair of lower and upper measures $\vec{\mu}_{s,w} = (\underline{\mu}_{s,w}, \bar{\mu}_{s,w})$ ¹⁰ defined by: $\forall \theta \in S\mathcal{L}$

$$\underline{\mu}_{s,w}(\theta) = w(\{\vec{v} \in \mathbb{V}_s : \underline{v}(\theta) = 1\}) \text{ and } \bar{\mu}_{s,w}(\theta) = w(\{\vec{v} \in \mathbb{V}_s : \bar{v}(\theta) = 1\})$$

We refer to SBPs where w is non-zero only on \mathbb{V}_{bs} as Bounded Supervaluation Belief Pairs (BSBP) which are then denoted by $\vec{\mu}_{bs,w} = (\underline{\mu}_{bs,w}, \bar{\mu}_{bs,w})$. Similarly, BSBPs where w is non-zero only on \mathbb{V}_{cbs} are referred to as Complete Bounded Supervaluation Belief Pairs (CBSBP) and denoted by $\vec{\mu}_{cbs,w} = (\underline{\mu}_{cbs,w}, \bar{\mu}_{cbs,w})$.

Notice that for SBPs we can define a mass function on sets of valuations $m : 2^{2^P} \rightarrow [0, 1]$ where $m(\emptyset) = 0$ and $\sum_{\mathcal{P} \subseteq 2^P} m(\mathcal{P}) = 1$ such that $\forall \vec{v} \in \mathbb{V}_s$,

$$m(\{F_{\alpha_j} : \bar{v}(\alpha_j) = 1\}) = w(\vec{v})$$

Notice that in this case:

$$\underline{\mu}_{s,w}(\theta) = \sum_{\mathcal{P} \subseteq \mathcal{P}_\theta} m(\mathcal{P}) \text{ and } \bar{\mu}_{s,w}(\theta) = \sum_{\mathcal{P} \cap \mathcal{P}_\theta \neq \emptyset} m(\mathcal{P})$$

Hence, $\underline{\mu}_{s,w}$ and $\bar{\mu}_{s,w}$ are Dempster-Shafer belief and plausibility measures on $S\mathcal{L}$ respectively [34]. Notice, in the case that $w : \mathbb{V}_s \rightarrow [0, 1]$ is non-zero on $\{\vec{v}_1, \dots, \vec{v}_r\}$ and where \vec{v}_i is generated by the set of admissible valuations \mathcal{P}_i which can be ordered such that $\mathcal{P}_i \supseteq \mathcal{P}_{i+1}$ where $i = 1, \dots, r-1$, then $\underline{\mu}_{s,w}$ and $\bar{\mu}_{s,w}$ are necessity and possibility measures respectively and hence: $\forall \theta, \varphi \in S\mathcal{L}$,

$$\underline{\mu}_{s,w}(\theta \wedge \varphi) = \min(\underline{\mu}_{s,w}(\theta), \underline{\mu}_{s,w}(\varphi)) \text{ and } \bar{\mu}_{s,w}(\theta \vee \varphi) = \max(\bar{\mu}_{s,w}(\theta), \bar{\mu}_{s,w}(\varphi))$$

¹⁰As with valuation pairs the notation $\vec{\mu} = (\underline{\mu}, \bar{\mu})$ refers to values of a lower and an upper belief measure and not an interval of possible values for a single belief measure.

Notice from theorem 15 that $\underline{\mu}_{s,w}$ and $\bar{\mu}_{s,w}$ are necessity and possibility measures exactly when w is non-zero only on a set of supervaluation pairs which can be totally ordered according to the semantic precision relation. In effect, in this case the agent’s uncertainty is restricted to a lack of knowledge about exactly how semantically precise the correct interpretation of \mathcal{L} should be.

We now define belief pairs for the case when the underlying truth-model is that of Kleene valuation pairs.

Definition 28. *Kleene Belief Pair*

Let $w : \mathbb{V}_k \rightarrow [0, 1]$ be a probability distribution on \mathbb{V}_k (i.e. the set of all Kleene valuation pairs on \mathcal{L}) then the Kleene Belief Pair (KBP) generated by w is a pair of lower and upper measures $\vec{\mu}_{k,w} = (\underline{\mu}_{k,w}, \bar{\mu}_{k,w})$ defined by: $\forall \theta \in S\mathcal{L}$

$$\underline{\mu}_{k,w}(\theta) = w(\{\vec{v} \in \mathbb{V}_k : \underline{v}(\theta) = 1\}) \text{ and } \bar{\mu}_{k,w}(\theta) = w(\{\vec{v} \in \mathbb{V}_k : \bar{v}(\theta) = 1\})$$

Unsurprisingly, the properties of belief pairs are fundamentally constrained by the properties of the underlying valuation pairs. For example, it follows immediately that belief pairs which are defined in terms of valuation pairs satisfying VP1, including both SBP and KBP, also satisfy duality in that $\forall \theta \in S\mathcal{L}, \underline{\mu}(-\theta) = 1 - \bar{\mu}(\theta)$, and vice-versa. In the sequel we will investigate the properties of supervaluation belief pairs and Kleene belief pairs in the light of the results given in section 3.2 relating Kleene and bounded supervaluation pairs in the context of the semantic precision relation. Initially, however, we propose an operational justification for the definition of belief pairs in terms of probability distributions on valuation pairs by introducing an extension of de Finetti’s betting semantics for subjective probability.

4.3 Betting Semantics for Belief Pairs

Inspired by the operationalism movement in physics de Finetti [4] proposed betting behaviour as an operational semantics for subjective probability. In this section we employ a generalisation of de Finetti’s result due to Paris [28] in order to give a betting interpretation of belief pairs as defined in section 4.2. This will require a short initial digression in order to outline de Finetti’s basic idea.

The motivation behind de Finetti’s approach was to provide a mechanism for eliciting a number in $[0, 1]$ quantifying an agent’s belief in a sentence θ by requiring them to decide whether or not to accept a series of bets involving θ . It can then be shown that in order for the agent to avoid situations of sure loss, referred to as Dutch books, an agent’s quantitative beliefs must then conform to the axioms of probability theory. More formally, we can identify a bet by triple (s, α, θ) where $\theta \in S\mathcal{L}$ is the statement on which we are betting, $s \in \mathbb{R}$ is the stake which can be either positive or negative, and $\alpha \in [0, 1]$ is the odds. Accepting or buying this bet then means agreeing to the following terms:

- pay $s \times \alpha$
- for which you will receive $\mathcal{L}s$ if θ is *true* and receive $\mathcal{L}0$ if θ is *false*.

Note that if $s < 0$ you are effectively selling the bet instead of buying it. Now supposing that the classical valuation $v \in \mathbb{V}_c$ is representative of the true state of the world combined with the correct language convention then the gain from bet (s, α, θ) is $s(v(\theta) - \alpha)$. Consequently, if language conventions are such that only classical valuation occur for any state of the world then a Dutch book can be defined as a set of bets $(s_i, \alpha_i, \theta_i) : i = 1, \dots, m$ such that $\forall v \in \mathbb{V}_c, \sum_{i=1}^m s_i(v(\theta_i) - \alpha_i) < 0$.

Given this we are now able to define a *rational agent* as follows: Let $\mu : S\mathcal{L} \rightarrow [0, 1]$ be a function representing an agent's quantitative beliefs about the sentences of \mathcal{L} . Suppose μ is determined so that for all sentences θ , $\mu(\theta)$ is the value for which the agent will accept the bet $(s, \mu(\theta), \theta)$ for any stake $s \in \mathbb{R}$ i.e. $\mu(\theta)$ represents the odds for which the agent will bet on θ no matter what the value of the stake. In this case an agent will accept any set of bets of the form $(s_i, \mu(\theta_i), \theta_i) : i = 1, \dots, m$ and consequently it is desirable that no such sequence should be a Dutch book. Hence, a rational agent is one for whom their belief function μ is defined in such a way that no set of sentences of this form is a Dutch book. de Finetti's famous theorem shows that, according to this definition, an agent is rational if and only if their belief is a probability measure on $S\mathcal{L}$.

Now the above argument assumes that the underlying truth model is classical. We now consider the consequences of allowing truth-gaps by assuming that the set of possible truth models is a finite set of valuation pairs $\mathbb{V}' \subseteq \mathbb{V}$ which satisfy the duality principle VP1. In this case it is natural to define both lower and upper bets where the lower bet is dependent on the sentence involved being absolutely true while the upper bet requires only that it is not absolutely false. Hence, let $(s, \alpha, \theta)_*$ denote the lower bet:

- pay $s \times \alpha$
- receive $\mathcal{L}s$ if θ is *absolutely true* and receive $\mathcal{L}0$ if θ is *not absolutely true*

while $(s, \alpha, \theta)^*$ denotes the upper bet:

- pay $s \times \alpha$
- receive $\mathcal{L}s$ if θ is *not absolutely false* and receive $\mathcal{L}0$ if θ is *absolutely false*

If $\vec{v} \in \mathbb{V}'$ is the valuation pair representing the true state of the world combined with the correct language interpretation then the gain from $(s, \alpha, \theta)_*$ is $s(\underline{v}(\theta) - \alpha)$ while that of $(s, \alpha, \theta)^*$ is $s(\vec{v}(\theta) - \alpha)$. An obvious consequence of this is that if $s \geq 0$ then the gain from $(s, \alpha, \theta)^*$ is always greater than or equal to that of $(s, \alpha, \theta)_*$ while if $s < 0$ then the situation is reversed. Consequently a rational agent should always prefer to buy upper

bets and sell lower bets whenever possible. Furthermore, the duality property VP1 implies a strong relationship between lower and upper bets as follows: Notice, that the gain for lower bet $(s, \alpha, \neg\theta)_*$ is $s(\underline{v}(\neg\theta) - \alpha) = s(1 - \bar{v}(\theta) - \alpha) = -s(\bar{v}(\theta) - (1 - \alpha))$ which is the same as the gain for the upper bet $(-s, 1 - \alpha, \theta)^*$. Consequently, these two bets should be viewed as equivalent since a rational agent should be willing to accept $(s, \alpha, \neg\theta)_*$ if and only if they are also willing to accept $(-s, 1 - \alpha, \theta)^*$. In other words, for any stake a rational agent should be equally willing to buy a lower bet on θ at a given odds and sell an upper bet on $\neg\theta$ at one minus those odds, and vice-versa.

We now adopt an operational model in which an agent's beliefs are quantified by a belief pair $\bar{\mu}$ where for any sentence θ , $\underline{\mu}(\theta)$ corresponds to the odds at which the agent is prepared to accept the lower bet $(s, \alpha, \theta)_*$ for any stake $s \in \mathbb{R}$, while $\bar{\mu}(\theta)$ is the odds at which the agent will accept the upper bet $(s, \alpha, \theta)^*$ for any stake $s \in \mathbb{R}$. Notice immediately from the above duality argument that the two bets $(s, \underline{\mu}(\neg\theta), \neg\theta)_*$ and $(-s, 1 - \underline{\mu}(\neg\theta), \theta)$ are equally acceptable for any $s \in \mathbb{R}$ and hence $\bar{\mu}(\theta) = 1 - \underline{\mu}(\neg\theta)$.

Now in this extended context a Dutch book is a combined set of lower and upper bets $(s_i, \alpha_i, \theta_i)_* : i = 1, \dots, r$ and $(s_j, \alpha_j, \theta_j)^* : j = r + 1, \dots, m$ for $m \geq r$ such that $\forall \vec{v} \in \mathbb{V}'$,

$$\sum_{i=1}^r s_i(\underline{v}(\theta_i) - \alpha_i) + \sum_{j=r+1}^m s_j(\bar{v}(\theta_j) - \alpha_j) < 0$$

However, from the duality relationship described above we can restrict our attention to Dutch books involving only lower bets since the set $(s_i, \alpha_i, \theta_i)_* : i = 1, \dots, r$ and $(s_j, \alpha_j, \theta_j)^* : j = r + 1, \dots, m$ is a Dutch book if and only if the set $(s_i, \alpha_i, \theta_i)_* : i = 1, \dots, r$ and $(-s_j, 1 - \alpha_j, \neg\theta_j)_* : j = r + 1, \dots, m$ is a Dutch book.

Theorem 29. *An agent is rational (in the sense of avoiding Dutch books comprised of lower and upper bets) if and only if $\forall \theta \in S\mathcal{L}$,*

$$\underline{\mu}(\theta) = w(\{\vec{v} \in \mathbb{V}' : \underline{v}(\theta) = 1\}) \text{ and } \bar{\mu}(\theta) = w(\{\vec{v} \in \mathbb{V}' : \bar{v}(\theta) = 1\})$$

where w is a probability distribution on \mathbb{V}' .

Proof. Paris [28] proved the following: Let \mathbb{B} be a finite set of binary functions on $S\mathcal{L}$. Here we are assuming that \mathbb{B} is the set of possible truth-models for \mathcal{L} , so that if $b \in \mathbb{B}$ is the true valuation then the gain from the bet (s, α, θ) is $s(b(\theta) - \alpha)$. Also, let $\mu : S\mathcal{L} \rightarrow [0, 1]$ denote the odds for which the bet (s, α, θ) is acceptable for all $s \in \mathbb{R}$. Then there is no set of bets $(s_i, b(\theta_i), \theta_i) : i = 1, \dots, m$ for which $\forall b \in \mathbb{B}$,

$$\sum_{i=1}^m s_i(b(\theta_i) - \mu(\theta_i)) < 0$$

if and only if μ is a convex linear combinations of functions in \mathbb{B} i.e. $\forall \theta \in S\mathcal{L}$, $\mu(\theta) = w(\{b \in \mathbb{B} : b(\theta) = 1\})$ where w is a probability distribution on \mathbb{B} .

Now from the above we know that a belief pair $\vec{\mu}$ avoids Dutch books involving lower and upper bets if and only if the lower measure $\underline{\mu}$ avoids Dutch books comprised only of lower bets. Hence, the result follows immediately from Paris' theorem by taking $\mathbb{B} = \{\underline{v} : \vec{v} \in \mathbb{V}'\}$ and from the duality between $\underline{\mu}$ and $\bar{\mu}$. \square

Clearly then theorem 29 provides an added justification for definitions 27 and 28 by taking $\mathbb{V}' = \mathbb{V}_s$ and $\mathbb{V}' = \mathbb{V}_k$ respectively. That is, if the underlying truth models of \mathcal{L} are supervaluation pairs or Kleene valuation pairs, then a rational agent should respectively adopt supervaluation belief pairs or Kleene belief pairs as their subjective belief model in order to avoid Dutch books comprised of sets of lower and upper bets.

4.4 Belief Pairs and Semantic Precision

We now investigate the properties of belief pairs relative to the semantic precision partial ordering on valuation pairs. The aim is to highlight certain classes of belief pairs which provide computationally efficient models under particularly restricted uncertainty conditions. We begin by extending theorem 21 so as to show a natural relationship between bounded supervaluation belief pairs and Kleene belief pairs.

Theorem 30. *For any BSBP $\vec{\mu}_{bs,w}$ there is a KBP $\vec{\mu}_{k,w'}$ such that $\forall \theta \in S\mathcal{L}^+ \cup S\mathcal{L}^-$, $\vec{\mu}_{bs,w}(\theta) = \vec{\mu}_{k,w'}(\theta)$ and $\forall \theta \in S\mathcal{L}$, $\underline{\mu}_{k,w'}(\theta) \leq \underline{\mu}_{bs,w}(\theta)$ and $\bar{\mu}_{k,w'}(\theta) \geq \bar{\mu}_{bs,w}(\theta)$.*

Proof. By theorem 21 there exists a function $g : \mathbb{V}_{bs} \rightarrow \mathbb{V}_k$ such that $g(\vec{v}_{bs}) = \vec{v}_k$ where $\underline{v}_k(p_i) = 1$ iff $p_i \in \underline{F}$ and $\bar{v}_k(p_i) = 1$ iff $p_i \in \bar{F}$ for $\underline{F} \subseteq \bar{F}$ the lower and upper admissible valuations of \vec{v}_{bs} and where $\forall \theta \in S\mathcal{L}^+ \cup S\mathcal{L}^-$, and $\forall \vec{v} \in \mathbb{V}_{bs}$, \vec{v} and $g(\vec{v})$ agree on θ .

Now given BSBP $\vec{\mu}_{bs,w}$ generated by probability distribution w on \mathbb{V}_{bs} we define the distribution w' on \mathbb{V}_k as follows

$$w'(\vec{v}') = w(\{\vec{v} \in \mathbb{V}_{bs} : g(\vec{v}) = \vec{v}'\})$$

Then $\forall \theta \in S\mathcal{L}^+ \cup S\mathcal{L}^-$ we have that:

$$\begin{aligned} \underline{\mu}_{bs,w}(\theta) &= w(\{\vec{v} \in \mathbb{V}_{bs} : \underline{v}(\theta) = 1\}) = w(\{\vec{v} \in \mathbb{V}_{bs} : g(\vec{v}) = \vec{v}' \text{ and } \underline{v}'(\theta) = 1\}) \\ &= w'(\{\vec{v}' \in \mathbb{V}_k : \underline{v}'(\theta) = 1\}) = \underline{\mu}_{k,w'}(\theta) \end{aligned}$$

The result follows similarly for $\bar{\mu}_{bs,w}(\theta)$.

Furthermore, from theorem 21 we have that if $g(\vec{v}) = \vec{v}'$ then $\vec{v}' \preceq \vec{v}$. Hence,

$$\begin{aligned} \underline{\mu}_{k,w'}(\theta) &= w'(\{\vec{v}' \in \mathbb{V}_k : \underline{v}'(\theta) = 1\}) = w(\{\vec{v} \in \mathbb{V}_{bs} : g(\vec{v}) = \vec{v}', \underline{v}'(\theta) = 1\}) \\ &\leq w(\{\vec{v} \in \mathbb{V}_{bs} : \underline{v}(\theta) = 1\}) = \underline{\mu}_{bs,w}(\theta) \end{aligned}$$

The proof that $\bar{\mu}_{k,w'}(\theta) \geq \bar{\mu}_{bs,w}(\theta)$ follows similarly. \square

For BSBP the following result identifies nestedness constraints on the set of valuation pairs with non-zero probability, resulting in a belief pair which is truth-functional on $S\mathcal{L}^+$ and $S\mathcal{L}^-$.

Theorem 31. *Let $\vec{\mu}_{bs,w}$ be a BSBP where w is non-zero on the set of bounded supervaluation pairs $\{\vec{v}_1, \dots, \vec{v}_r\}$ such that for \vec{v}_i and \vec{v}_j where $i \neq j$ either $\underline{F}_i \subseteq \underline{F}_j$ or $\underline{F}_j \subseteq \underline{F}_i$ and either $\overline{F}_i \subseteq \overline{F}_j$ or $\overline{F}_j \subseteq \overline{F}_i$. In this case: $\forall \theta, \varphi \in S\mathcal{L}^+$ and $\forall \theta, \varphi \in S\mathcal{L}^-$.*

$$\begin{aligned} \underline{\mu}_{bs,w}(\theta \wedge \varphi) &= \min(\underline{\mu}_{bs,w}(\theta), \underline{\mu}_{bs,w}(\varphi)) \text{ and } \overline{\mu}_{bs,w}(\theta \wedge \varphi) = \min(\overline{\mu}_{bs,w}(\theta), \overline{\mu}_{bs,w}(\varphi)) \\ \underline{\mu}_{bs,w}(\theta \vee \varphi) &= \max(\underline{\mu}_{bs,w}(\theta), \underline{\mu}_{bs,w}(\varphi)) \text{ and } \overline{\mu}_{bs,w}(\theta \vee \varphi) = \max(\overline{\mu}_{bs,w}(\theta), \overline{\mu}_{bs,w}(\varphi)) \end{aligned}$$

Proof. Here we only prove the result for sentences in $S\mathcal{L}^+$. The result follows similarly for sentences in $S\mathcal{L}^-$. Given the conditions of the theorem we can assume w.l.o.g. that the set of bounded supervaluation pairs $\{\vec{v}_1, \dots, \vec{v}_r\}$ are ordered such that $\underline{F}_1 \subseteq \dots \subseteq \underline{F}_r$. Now by lemma 16 it holds that $\forall \theta \in S\mathcal{L}^+$, $\underline{v}_i(\theta) = 1$ if and only if $v_{\underline{F}_i}(\theta) = 1$. Also, by lemma 16 it follows that if $v_{\underline{F}_i}(\theta) = 1$ then $v_{\underline{F}_j}(\theta) = 1$ for $j = i + 1, \dots, r$. Hence, $\exists t \leq r$ such that:

$$\{\vec{v}_i : \underline{v}_i(\theta) = 1\} = \{\vec{v}_t, \dots, \vec{v}_r\}$$

Similarly, for $\varphi \in S\mathcal{L}^+$, $\exists t' \leq r$ such that

$$\{\vec{v}_i : \underline{v}_i(\varphi) = 1\} = \{\vec{v}_{t'}, \dots, \vec{v}_r\}$$

Hence, since from a basic property of supervaluation pairs $\underline{v}_i(\theta \wedge \varphi) = \min(\underline{v}_i(\theta), \underline{v}_i(\varphi))$ it follows that

$$\{\vec{v}_i : \underline{v}_i(\theta \wedge \varphi) = 1\} = \{\vec{v}_{\max(t,t')}, \dots, \vec{v}_r\}$$

Consequently,

$$\begin{aligned} \underline{\mu}_{bs,w}(\theta) &= \sum_{j=t}^r w(\vec{v}_j), \quad \underline{\mu}_{bs,w}(\varphi) = \sum_{j=t'}^r w(\vec{v}_j) \text{ and } \underline{\mu}_{bs,w}(\theta \wedge \varphi) = \sum_{j=\max(t,t')}^r w(\vec{v}_j) \\ &= \min(\underline{\mu}_{bs,w}(\theta), \underline{\mu}_{bs,w}(\varphi)) \end{aligned}$$

Also, by theorem 20 it follows that $\underline{v}_i(\theta \vee \varphi) = \max(\underline{v}_i(\theta), \underline{v}_i(\varphi))$ and hence,

$$\underline{\mu}_{bs,w}(\theta \vee \varphi) = \sum_{j=\min(t,t')}^r w(\vec{v}_j) = \max(\underline{\mu}_{bs,w}(\theta), \underline{\mu}_{bs,w}(\varphi))$$

The required results for $\overline{\mu}_{bs,w}$ can be proved in a similar manner since we can reorder the relevant valuation pairs so that $\{\vec{v}_1, \dots, \vec{v}_r\} = \{\vec{v}'_1, \dots, \vec{v}'_r\}$ where $\overline{F}'_1 \subseteq \dots \subseteq \overline{F}'_r$. \square

Corollary 32. Let $\vec{\mu}_{bs,w}$ be a BSBP for which $\{\vec{v} \in \mathbb{V}_{bs} : w(\vec{v}) > 0\} = \{\vec{v}_1, \dots, \vec{v}_r\}$ can be ordered such that $\vec{v}_1 \preceq \vec{v}_2 \preceq \dots \preceq \vec{v}_r$ then the following holds:

$$\forall \theta, \varphi \in S\mathcal{L}, \underline{\mu}_{bs,w}(\theta \wedge \varphi) = \min(\underline{\mu}_{bs,w}(\theta), \underline{\mu}_{bs,w}(\varphi)) \text{ and}$$

$$\overline{\mu}_{bs,w}(\theta \vee \varphi) = \max(\overline{\mu}_{bs,w}(\theta), \overline{\mu}_{bs,w}(\varphi))$$

and

$$\forall \theta, \varphi \in S\mathcal{L}^+ \text{ and } \forall \theta, \varphi \in S\mathcal{L}^-, \underline{\mu}_{bs,w}(\theta \vee \varphi) = \max(\underline{\mu}_{bs,w}(\theta), \underline{\mu}_{bs,w}(\varphi)) \text{ and}$$

$$\overline{\mu}_{bs,w}(\theta \wedge \varphi) = \min(\overline{\mu}_{bs,w}(\theta), \overline{\mu}_{bs,w}(\varphi))$$

Proof. Since $\vec{v}_1 \preceq \dots \preceq \vec{v}_r$ it follows from theorem 15 that $\mathcal{P}_1 \supseteq \dots \supseteq \mathcal{P}_r$. Hence $\underline{\mu}_{bs,w}$ and $\overline{\mu}_{bs,w}$ are necessity and possibility measures on $S\mathcal{L}$ respectively. Also, it trivially holds that $\underline{F}_i \subseteq \underline{F}_{i+1}$ and $\overline{F}_i \supseteq \overline{F}_{i+1}$ for $i = 1, \dots, r-1$. The result then follows immediately from theorem 31. \square

Corollary 32 identifies a special case of the conditions of theorem 31 for BSBPs generated from a probability distribution which is non-zero only on a set of valuation pairs that can be totally ordered by the semantic precision relation. Recall from section 4.2 that in this case a SBP corresponds to a necessity and possibility measure on $S\mathcal{L}$. Combined with theorem 31 this results in uncertainty measures that are partially functional in general, and fully truth-functional on $S\mathcal{L}^+$ and $S\mathcal{L}^-$.

Corollary 32 has close parallels with a related result for KBP originally given in [23] as follows:

Theorem 33. Let $\vec{\mu}_{k,w}$ be a KBP for which $\{\vec{v} \in \mathbb{V}_k : w(\vec{v}) > 0\} = \{\vec{v}_1, \dots, \vec{v}_r\}$ can be ordered such that $\vec{v}_1 \preceq \vec{v}_2 \preceq \dots \preceq \vec{v}_r$. In this case: $\forall \theta, \varphi \in S\mathcal{L}$

$$\underline{\mu}_{k,w}(\theta \wedge \varphi) = \min(\underline{\mu}_{k,w}(\theta), \underline{\mu}_{k,w}(\varphi)) \text{ and } \overline{\mu}_{k,w}(\theta \wedge \varphi) = \min(\overline{\mu}_{k,w}(\theta), \overline{\mu}_{k,w}(\varphi))$$

$$\underline{\mu}_{k,w}(\theta \vee \varphi) = \max(\underline{\mu}_{k,w}(\theta), \underline{\mu}_{k,w}(\varphi)) \text{ and } \overline{\mu}_{k,w}(\theta \vee \varphi) = \max(\overline{\mu}_{k,w}(\theta), \overline{\mu}_{k,w}(\varphi))$$

Theorem 33 is an even stronger result than corollary 32 in that for KBPs, if an agent's uncertainty concerns only the semantic precision of the correct interpretation of \mathcal{L} , then the resulting belief pair is fully truth-functional. Indeed, in such a case KBPs satisfy the definition of lower and upper truth values (membership values) in interval fuzzy logic (set theory) as proposed independently by Zadeh [43], Grattan-Guinness [13], Jahn [15] and Sambuc [33]. Furthermore, given the isomorphic relationship between interval fuzzy logic and Atanassov's intuitionistic fuzzy logic [1], theorem 33 also allows for an interpretation of the latter in terms of KBPs. More specifically, we can interpret intuitionistic fuzzy truth and falsity degrees in a sentence θ as corresponding to the probability that

θ is absolutely true and the probability that θ is absolutely false respectively, based on an underlying truth-model of Kleene valuation pairs about which the only uncertainty relates to semantic precision. In this case we can think of these measures as analogous to necessity and possibility measures, but where the truth-model is that of Kleene rather than supervaluation pairs.

5 Decision Making about Assertions

A fundamental challenge in natural language communications is to understand why we often choose to make vague assertions when we have many potentially suitable crisp (semantically precise) alternatives at our disposal. If the aim were simply to convey information then we might suppose that the use of vague statements would only result in reduced efficiency. Indeed some formal results in the game theoretic literature seem to support this view [25]. However, van Deemter [39], [40] identifies a number of plausible reasons for the utility of vagueness in natural language. In [39] a decision theoretic approach is proposed according to which different utterances have different rewards or costs (negative rewards), and where the choice of assertion in a particular context is driven by utility. Such an approach has similarities with the assertability risk models of Giles [11] and Kyburg [17]. In this section we outline a very preliminary application of belief pairs to assertion decisions, focusing on a class of communication scenarios referred to by van Deemter [39] as *future contingencies*. Here we consider a communications model in which different agents have different, and to some degree conflicting, goals and objectives. An agent is then required to make a statement, perhaps in the form of a prediction or commitment, about some future state of the world. While there are potential rewards for making such statements, there may also be the risk of significant costs if they prove to be incorrect or if promises are seen to be broken. In such scenarios, van Deemter argues, the utility of making a vague assertion may be potentially higher than that for making a crisp assertion, primarily because the risk of incurring costs is significantly reduced. Future contingency scenarios are important in many AI applications involving natural language generation (see [31] for an overview) such as automatic weather forecast generation [12] or in automatic medical diagnosis systems [29].

To illustrate the type of problems involved consider the example of a politician who is considering making a pledge at an upcoming British General Election. She would like to promise as high a reduction in the budget deficit as possible (achieved by cutting public expenditure) resulting in as small an increase as possible in unemployment. She judges that a promise of this kind would be popular with voters and would increase the probability of her (or her political party) winning the election. One possibility would be to decide on a crisp assertion and promise say ‘at least a 40% reduction in the deficit and no more than a 2% increase in unemployment’. However, she may then be worried that she will find it

difficult to meet this commitment exactly and subsequently be punished by the voters at the following election in five years time. If on the other hand she promises ‘a significant reduction in the debt with only a minor increase in unemployment’ then she may still be able to curry favour with the voters at the current general election while reducing the risk of being accused of breaking promises at the subsequent election. To quote van Deemter [39] ‘a precise promise is easier to break than a vague one’. The notion of valuation pairs and truth-gaps may provide some insight into why this is indeed the case. The voters in our example may interpret the politician as promising that their assertion will be absolutely true by some future date, perhaps by the subsequent election, and reward them accordingly. However, they will not then significantly punish them unless the assertion turns out to be absolutely false. Hence, in future contingency scenarios the truth-gap associated with a vague assertion provides the agent with a window of opportunity in which they can maximize their chance of reward and minimize their risk of loss. Furthermore, in practice almost all assertion decisions of this kind will be made in the presence of significant epistemic uncertainty. Consequently, by combining a model of truth-gaps with that of epistemic uncertainty, bipolar belief pairs may well provide a useful framework for the analysis of future contingency and similar assertion decision scenarios. Indeed, for the above example there is a clear need for an integrated model of this kind. Since phrases such as ‘significant reduction’ and ‘minor increase’ are vague we would certainly expect linguistic conventions to be based around truth-models which explicitly represent borderline cases. For example, the truth valuation of the statement that there will be only ‘a minor increase in unemployment’ might be determined by lower and upper thresholds on the overall percentage increase in unemployment, along the lines described in section 4.1 for the concept short. In addition, the inherent semantic uncertainty associated with the interpretation of such natural language descriptions would mean that the politician will be uncertain as to the exact nature of the language convention which should be applied in this case. This semantic uncertainty could then manifest itself as uncertainty regarding the values of the relevant lower and upper thresholds for phrases such as ‘minor increase in unemployment’. Furthermore, since the politician is making a prediction about the future, she will be uncertain as to what will be the actual state of the world. For example, she will be unsure about what will be the actual percentage changes to unemployment and to the deficit. The bipolar model outlined in section 4 would then allow the politician to capture all these different aspects influencing the likely truth or falsity of their predictions by lower and upper belief measures on the sentences concerned.

We now propose a simple utility decision model based on bipolar belief measures in order to explore the potential advantages of asserting a vague sentence over a crisp one. Suppose that, at time t_0 , an agent must choose between assertions θ and φ concerning the state of the world at some future time t_1 , where φ is crisp and θ is vague. Assume that the reward for making either assertion at time t_0 is $x > 0$. At time t_1 there will be a

further reward $y > 0$ if the assertion made at t_0 is judged to be absolutely true given the actual state of the world at t_1 . Similarly, there will be a cost $-z$ ($z > 0$) if the assertion made is judged to be absolutely false given the state of the world at t_1 . Furthermore, in the case of vague assertion θ , there will be a small cost $-w$ ($0 \leq w \leq z$) if θ is judged to be a borderline statement at t_1 .

Now suppose that at time t_0 the agents' beliefs about the state of the world at t_1 is represented by a bipolar belief pair $\vec{\mu}$ ¹¹. Let U_θ and U_φ denote the total rewards (which may be negative) for making assertions θ and φ respectively. Then based on $\vec{\mu}$, the agent can evaluate their expected reward for making either assertion θ or φ at time t_0 as follows:

$$E(U_\theta) = x + \underline{\mu}(\theta)y - \underline{\mu}(\neg\theta)z - (\bar{\mu}(\theta) - \underline{\mu}(\theta))w = x - z + \underline{\mu}(\theta)(y + w) + \bar{\mu}(\theta)(z - w)$$

$$E(U_\varphi) = x + \mu(\varphi)y - \mu(\neg\varphi)z = x - z + \mu(\varphi)(y + z)$$

where, since φ is a crisp statement, $\mu(\varphi) = \bar{\mu}(\varphi) = \underline{\mu}(\varphi)$. Given this formulation we can now investigate under what circumstances the agent would expect a higher reward for asserting the vague statement θ than for asserting the precise statement φ . Initially, we make the simplifying assumption that θ being a borderline case is effectively cost neutral so that $w = 0$. Furthermore, we assume that θ is selected such that $\underline{\mu}(\theta) < \mu(\varphi) < \bar{\mu}(\theta)$ ¹². Now let $\alpha = \frac{y}{z}$ be the ratio of the reward for the assertion being absolutely true at t_1 against the cost of the assertion be absolutely false at t_1 . Then we have that:

$$E(U_\theta) \geq E(U_\varphi) \text{ if and only if } \alpha \leq \frac{\bar{\mu}(\theta) - \mu(\varphi)}{\mu(\varphi) - \underline{\mu}(\theta)}$$

Hence, we have an upper bound on the ratio of reward y over cost z for which $E(U_\theta) \geq E(U_\varphi)$, so that the higher the cost of making an assertion which is absolutely false at t_1 relative to the reward of making an assertion which is absolutely true at t_1 , the more likely it is that the agent will be better off choosing a vague over a crisp assertion. In addition, as the belief value of φ approaches to the lower belief value for θ then there is an increasing range of ratio values for which the agent will be better off choosing the vague over the crisp assertion. Furthermore, notice that if $\mu(\varphi) \leq \frac{\underline{\mu}(\theta) + \bar{\mu}(\theta)}{2}$ then $E(U_\theta) \geq E(U_\varphi)$ for all $\alpha \in [0, 1]$ (i.e. for all cases where $y \leq z$). A particular instance of this is when $\vec{\mu}(\theta) = (0, 1)$ and $\mu(\varphi) \leq 0.5$ and hence, if the cost of making an assertion which is absolutely false at t_1 is higher than the reward for making an assertion which is absolutely true at t_1 , then a rational agent should choose a statement which they are certain is borderline over a crisp statement for which their belief value is at most 0.5.

¹¹For this analysis it is not important whether $\vec{\mu}$ is a SBP or a KBP

¹²A justification of this assumption is as follows: Suppose that θ is selected by the agent so that it is a less semantically precise statement than φ . Then in the framework of valuation pairs this would mean the following: If $\vec{\mu}$ is generated by probability distribution w where $\{\vec{v} : w(\vec{v}) > 0\} = \{\vec{v}_1, \dots, \vec{v}_r\}$ then since φ is crisp $\underline{v}_i(\varphi) = \bar{v}_i(\varphi) = v_i(\varphi)$ for $i = 1, \dots, r$ and since θ is less semantically precise than φ then $\underline{v}_i(\theta) \leq v_i(\varphi) \leq \bar{v}_i(\theta)$ for $i = 1, \dots, r$, where in some cases these inequalities are strict. From this it follows immediately that $\underline{\mu}(\theta) < \mu(\theta) < \bar{\mu}(\theta)$.

These result must, of course, be understood in the light of the modelling assumption we have made. In particular, we have assumed that there are the same rewards and costs, x, y and z , for both crisp and vague assertions. We might perhaps justify this assumption by supposing that the two assertions are sufficiently semantically similar that there would be no distinction made between them in the allocation of rewards and costs. However, it is equally possible that other agents may prefer crisp assertions over vague ones. For example, voters may be more likely to vote for a politician who makes crisp promises because such assertions are easier to evaluate at a future date than similar but vague commitments. One way of modelling this would be to suppose different rewards x_1 and x_2 for asserting θ and φ at t_0 respectively, and where $x_2 \geq x_1$. In this case the expected utilities from asserting θ and φ are given by:

$$E(U_\theta) = x_1 - z + \underline{\mu}(\theta)(y + w) + \bar{\mu}(\theta)(z - w) \text{ and } E(U_\varphi) = x_2 - z + \mu(\varphi)(y + z)$$

From this we obtain that $E(U_\theta) \geq E(U_\varphi)$ if and only if $\alpha \leq \frac{\bar{\mu}(\theta) - \mu(\varphi) - \beta}{\mu(\varphi) - \underline{\mu}(\theta)}$ and $\beta \leq \bar{\mu}(\theta) - \mu(\varphi)$ where $\beta = \frac{x_2 - x_1}{z}$ is the ratio of the difference between the two rewards at t_0 over the cost at t_1 of making an absolutely false assertion. Consequently, the larger the cost of making an absolutely false assertion relative to the different rewards x_1, x_2 and y , the more likely it is that the agent will be better off making a vague rather than a crisp assertion.

The above is of course a highly simplified analysis with many assumptions which are probably unrealistic. Despite these limitations we nonetheless suggest that it does take some small steps towards demonstrating the potential applicability of the bipolar belief framework for modelling decision problems involving the choice of vague over crisp assertions. A decision theoretic analysis of this kind could have significant applications in Artificial Intelligence. Indeed, some natural language generation systems already include the use of vague words. For example, the weather forecasting system FOG developed by Goldberg et al. [12] includes terms such as westerly or southerly to describe wind direction. A more developed analysis along the lines of that outline above, raises the possibility of providing such systems with the capability of choosing between vague and crisp assertions on the basis of a range of factors including context and uncertainty levels.

6 Conclusions

In this paper we have introduced valuation pairs as a truth-model which captures borderline cases relating to vague propositional statements. Different classes of valuation pairs have been discussed and the relationships between them have been investigated. In particular, we have provided axiomatic characterisations of supervaluation pairs and Kleene valuation pairs. Furthermore, we have exposed the close relationship between Kleene valuation pairs and bounded supervaluation pairs.

We have then proposed an integrated approach to truth-gaps and epistemic uncertainty

in the form of lower and upper belief measures on the sentences of the language. In this context the lower measure of a sentence θ is the probability that θ is absolutely true, while the upper measure is the probability that θ is not absolutely false. The definition of these two measures in terms of an underlying probability distribution on valuation pairs can be justified by assuming that the truth-model for the language is functionally dependent on both the state of the world and conventions for language use, and where epistemic uncertainty is modelled by a probability distribution on world state and language convention pairs. Further justification is provided by a betting argument according to which a rational agent must define their lower and upper measures so as to avoid accepting Dutch books consisting of sets of lower and upper bets. Within this bipolar belief framework we have outlined a possible decision theoretic argument as to why, in some circumstances, it might be optimal for an agent to choose to assert a vague statement over a crisp one when making forecasts or promises relating to some future state of the world about which they are uncertain.

The model of truth-gaps provided by valuation pairs assumes that the notions of absolutely true and absolutely false are primitives and consequently our framework lacks some of the expressiveness of other approaches especially with regard to higher order vagueness. For example, many formalisations of supervaluationism are based on modal logic (e.g. Bennett [2]) and which therefore allow the explicit representation of higher order statements such as ‘absolutely true θ ’ and ‘borderline true that absolutely true θ ’ etc. Indeed, the importance of capturing higher-order borderline cases in this manner is largely taken for granted by theorists of vagueness. However, we argue that there are a number of justifications for a more straightforward approach which avoids representing higher-order statements as part of the language.

From a practical perspective the valuation pair framework is sufficiently general to encompass a range of different theoretical approaches to truth-gaps. This enables direct comparisons between, for instance, supervaluationism and three-valued logic of the kind we have outlined in this paper. Furthermore, the valuation pair approach can be naturally extended so as to provide an integrated model of truth-gaps and epistemic uncertainty based on lower and upper measures as described in section 4. From a more theoretical perspective we argue that the idea of higher-order truth-gaps is actually quite problematic and that it is far from clear that valuations of the form *borderline true that absolutely true* are even meaningful. Consider, for example, the concept of *short* as defined on the scale of heights. Suppose that instead of the two threshold values $\underline{\epsilon}$ and $\bar{\epsilon}$ suggested in section 4.1, we have a series of increasing thresholds representing borderline cases, borderline borderline cases and so on. The problem with this interpretation of the concept is that, as pointed out by Sainsbury [32], it naturally collapses to the original two-borderline case. This is because, given the original intuition behind borderline cases as being those which are neither absolutely true nor absolutely false, we must be able to partition the

set of heights as follows: There is a set of lower heights to which the concept *short* is absolutely and totally applicable. Similarly, there is an upper set of heights to which *short* is absolutely and totally not applicable. The union of all other sets of heights, no matter what level of borderline they constitute, can simply be viewed as a type of borderline i.e. heights which for which short is neither absolutely applicable nor absolutely not applicable. Another perspective on this is given by Raffman [30] who suggests that second order statement such as ‘absolutely true θ ’ and ‘absolutely false θ ’ simply do not admit truth-gaps. In other words, such statements admit only a binary truth-model. Indeed, Raffman argues, that given our intuitive understanding of ‘absolutely true θ ’ and ‘absolutely false θ ’ as meaning that θ is unquestionably true or unquestionably false respectively, then a statement such as ‘borderline absolutely true θ ’ can have no other truth value but false. Of course, none of this means that an agent may not be uncertain as to the exact boundaries of a concept. However, we would argue that this is a form of semantic uncertainty about the conventions of language use and consequently should be modelled probabilistically as described in section 4, rather than being incorporated into the underlying truth-model by means of modal operators.

It is certainly the case that a purely propositional framework, as described in this paper, has definite limitations. Future work should therefore explore extending these ideas to predicate logic. One initial approach might be to develop random set and prototype theory models of the kind introduced in [20], so as to incorporate both lower and upper threshold distances from prototypes. The idea here is that the interpretation of a predicate R would be defined by a set of prototypes PR in some underlying metric space, together with lower and upper thresholds, $\underline{\epsilon}$ and $\bar{\epsilon}$, on the distance of elements from PR . For an element of the space x with distance less than $\underline{\epsilon}$ from PR , $R(x)$ would be absolutely true, while for an element x with distance from PR greater than $\bar{\epsilon}$, $R(x)$ would be absolutely false. All other elements would then be borderline cases of R . Semantic uncertainty about the exact values of $\underline{\epsilon}$ and $\bar{\epsilon}$ could then be modelled by a joint probability distribution on the two thresholds. An elementary model of this kind has already been proposed in [21] and in [38], but there is scope to develop a much more general version of this prototype theory based framework to capture borderline cases and semantic uncertainty in predicate definitions.

As we have already indicated in section 5 the bipolar framework outlined in this paper has potential application to natural language generation systems. Other areas of Artificial Intelligence where our approach may have applications include consensus modelling and multi-agent dialogues. For the former, truth-gaps may provide a means by which an agent can adapt their beliefs so as to reach consensus with other agents whilst maintaining a certain level of internal consistency. Indeed truth-gaps are fundamental here, in allowing for a weaker notion of consistency according to which different valuations may still be considered consistent. Furthermore, consensus modelling provides an example of an appli-

cation in which the interests of different agents are more or less aligned. More specifically, agents share a common goal of, where possible, reaching a consensus about their beliefs. In multi-agent dialogues, lower and upper valuations can provide a model of strong and weak assertions, with semantic precision providing a structured criterion according to which different viewpoints can be clustered together. Preliminary studies have been carried out in both application areas for Kleene valuation pairs (see [24] and [22]). There are plans to extend these to more in depth studies, including developing the proposed methods within the general valuation pairs framework.

Acknowledgements

The authors would like to thank the anonymous referees for their insightful comments and suggestions.

References

- [1] K. Atanassov, (1986), ‘Intuitionistic Fuzzy Sets’, *Fuzzy Sets and Systems*, Vol. 20, pp87-96.
- [2] B. Bennett, (1998), ‘Modal Semantics for Knowledge Bases Dealing with Vague Concepts’, in (Eds. A.G. Cohn, L. Schubert and S. Shapiro) *Principles of Knowledge Representation and Reasoning: Proceedings in KR-98*, pp234-244.
- [3] G. de Cooman, (1995), ‘Towards a Possibilistic Logic’, in *Fuzzy Set Theory and Advanced Mathematical Applications*, (Ed. D. Ruan), pp89-133.
- [4] B. de Finetti, (1931), ‘Sul significato soggettivo della probabilità’, *Fundamenta Mathematicae*, Vol. 17, pp298-329.
- [5] D. Dubois, H. Prade, (2011), ‘Gradualness, Uncertainty and Bipolarity: Making Sense of Fuzzy Sets’, *Fuzzy Sets and Systems*, In Press.
- [6] D. Dubois, (2008), ‘On Ignorance and Contradiction Considered as Truth-values’, *Logic Journal of the IGPL*, Vol. 16, No. 2, pp195-216.
- [7] D. Dubois, H. Prade, (2008), ‘An Introduction to Bipolar Representations of Information and Preference’, *International Journal of Intelligent Systems*, Vol. 23, pp866-877.
- [8] D. Dubois, H. Prade, (1997), ‘The Three Semantics of Fuzzy Sets’, *Fuzzy Sets and Systems*, Vol. 90, pp141-150.

- [9] D. Dubois, H. Prade, (2001), ‘Possibility theory, Probability Theory and Multiple-valued Logics: A Clarification’, *Annals of Mathematics and Artificial Intelligence*, Vol. 32, pp35-66.
- [10] K. Fine, (1975), ‘Vagueness, Truth and Logic’, *Synthese*, Vol. 30, pp265-300.
- [11] R. Giles, (1988), ‘The Concept of Grade of Membership’, *Fuzzy Sets and Systems*, Vol. 25, pp297-323.
- [12] E. Goldberg, N. Driedger, R.I. Kittredge, (1994), ‘Using Natural-Language Processing to Produce Weather Forecasts’, *IEEE Expert: Intelligent Systems and their Applications*, Vol. 9, No.2, pp45-53.
- [13] I. Grattan-Guinness, (1975), ‘Fuzzy Membership Mapped onto Interval and Many-Valued Quantities’, *Z. Math. Logik Grundlehren Math*, Vol. 22, pp149-160.
- [14] D. Hyde, (1997), ‘From Heaps and Gaps to Heaps of Gluts’, *Mind*, Vol. 106, pp641-660.
- [15] K.U. Jahn, (1975), ‘Intervall-wertige Mengen’, *Math. Nach.*, Vol. 68, pp115-132.
- [16] S.C. Kleene, (1952), *Introduction to Metamathematics* D. Van Nostrand Company Inc., Princeton, New Jersey.
- [17] A. Kyburg, (2000), ‘When Vague Sentences Inform: A Model of Assertability’, *Synthese*, Vol. 124, pp175-192.
- [18] D. Lassiter, (2011), ‘Vagueness as Probabilistic Linguistic Knowledge’, *Lecture Notes in Artificial Intelligence*, Vol. 6517, pp127-150.
- [19] J. Lawry, (2008), ‘Appropriateness Measures: An uncertainty model for vague concepts’, *Synthese*, Vol. 161, No. 2, pp255-269.
- [20] J. Lawry, Y. Tang, (2009), ‘Uncertainty modelling for vague concepts: A prototype theory approach’, *Artificial Intelligence*, Vol. 173, pp1539-1558.
- [21] J. Lawry, (2010), ‘A Random Set and Prototype Theory Interpretation on Intuitionistic Fuzzy Sets’, in (Eds E. Hüllermeier, R. Kruse, F. Hoffman) *Information Processing and Management of Uncertainty in Knowledge-Based Systems: Theory and Methods*, pp618-628.
- [22] J. Lawry, (2010), ‘Imprecise Bipolar Belief Measures Based on Partial Knowledge from Agent Dialogues’, *Lecture Notes in Computer Science*, Vol. 6379, pp205-218.
- [23] J. Lawry, I. Gonzalez-Rodriguez, (2011), ‘A Bipolar Model of Assertability and Belief’, *International Journal of Approximate Reasoning*, Vol. 52, pp76-91.

- [24] J. Lawry, D. Dubois, (2012), ‘A Bipolar Framework for Combining Beliefs about Vague Propositions’, *Proceedings of 13th International Conference on Principles of Knowledge Representation and Reasoning*, pp530-540.
- [25] B. L. Lipman, (2006), ‘Why is language vague?’, *Working Paper*, Department of Economics, Boston University.
- [26] J. Lukasiewicz, (1920), ‘O logice trojwartosciowej (On three-valued logic).’, *Ruch filozoficzny*, Vol 5, pp170-171.
- [27] R. Parikh, (1996), ‘Vague Predicates and Language Games’, *Theoria (Spain)*, Vol. XI(27), pp97-107.
- [28] J. B. Paris, (2001), ‘A Note on the Dutch Book Method’, *Proceedings of the Second International Symposium on Imprecise Probabilities and their Applications*, Ithaca, New York.
- [29] F. Portet, E. Reiter, A. Gatt, J. Hunter, S. Sripada, Y. Freer, C. Sykes, (2009), ‘Automatic Generation of Textual Summaries from Neonatal Intensive Care Data’, *Artificial Intelligence*, Vol. 173, No. 7-8, pp789-816.
- [30] D. Raffman, (2010), ‘Demoting Higher-Order Vagueness’ in (Eds. R. Dietz, S. Moruzzi) *Cuts and Clouds: Vagueness, its Nature and its Logic*, OUP.
- [31] E. Reiter, R. Dale, (2006), *Building Natural Language Generation Systems*, Cambridge University Press.
- [32] R.M. Sainsbury, (1997), ‘Concepts without Boundaries’, in (Eds. R. Keefe, P. Smith) *Vagueness: A Reader*, MIT press.
- [33] R. Sambuc, (1975), ‘Fonctions-floues. Application a l’aide au diagnostic en pathologie thyroïdienne’, *PhD Thesis, Univ. Marseille, France*.
- [34] G. Shafer, (1976), *A Mathematical Theory of Evidence*, Princeton University Press.
- [35] S. Shapiro, (2006), *Vagueness in Context*, Oxford University Press.
- [36] N.J.J. Smith, (2008), *Vagueness and Degrees of Truth*, Oxford University Press.
- [37] L. Steels, T. Belpaeme, (2005), ‘Coordinating Perceptually Grounded Categories Through Language’, *Behavioural and Brain Sciences*, Vol. 28, pp469-529.
- [38] Y. Tang, J. Lawry, (2011), ‘Bipolar Semantic Cells: An Interval Model of Linguistic Labels’, *Lecture Notes in Computer Science*, Vol. 7027, pp60-71.
- [39] K. van Deemter, (2009), ‘Utility and Language Generation: The Case of Vagueness’, *Journal of Philosophical Logic*, Vol. 38, pp607-632.

- [40] K. van Deemter, (2010), *Not Exactly: In Praise of Vagueness*, Oxford University Press.
- [41] B.C. van Fraassen, (1966), 'Singular Terms, Truth-Value Gaps, and Free Logic', *The Journal of Philosophy*, Vol. 63, No. 17, pp481-495.
- [42] R. Yager, (1987), 'Set-Based Representations of Conjunctive and Disjunctive Knowledge', *Information Sciences*, Vol. 41, pp1-22.
- [43] L.A. Zadeh, (1975), 'The Concept of a Linguistic Variable and its Application to Approximate Reasoning:I', *Information Sciences*, Vol. 8, pp199-249.
- [44] T. Williamson, (1994), *Vagueness*, Routledge.