Conditional Beliefs in a Bipolar Framework

Jonathan Lawry and Trevor Martin

Department of Engineering Mathematics,
University of Bristol,
Bristol, UK
j.lawry@bris.ac.uk

Abstract. A framework for quantifying lower and upper bipolar belief is introduced, which incorporates aspects of stochastic and of semantic uncertainty as well as an indeterministic truth-model allowing for inherent linguistic vagueness at the propositional level. This is then extended to include lower and upper measures of conditional belief given information in the form of lower and upper truth-valuations. The properties of these measures are explored and their relationship with conditional belief in other uncertainty theories is highlighted.

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 classified as being absolutely *short* nor absolutely *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* suggesting that a non-Tarskian notion of truth may be required to capture this aspect of vagueness. A model of this kind with distinct, although related, valuations for absolute truth and absolute falsity exhibits, what Dubois and Prade [1], refer to as *symmetric bivariate unipolarity*, whereby judgments are made according to two distinct evaluations on unipolar scales i.e. distinct evaluations about the truth value of a sentence and that of its negation. In the current context, we have a strong and a weak evaluation criterion where the former corresponds to *absolute truth* and the latter *not absolute falsity.* As with many examples of this type of bipolarity there is then a natural duality between the two evaluation criteria in that a proposition is absolutely true if and only if its negation is absolutely false.

The development of formal models incorporating truth-gaps has potentially important applications in artificial intelligence systems. For example, allowing for borderline cases can help to mitigate the risks associated with making forecasts [15]. In this context, a bipolar framework can form the basis of a decision theoretic model to enable natural language generation systems, such as automatic weather forecasters, to decide between different assertions with different levels of semantic precision, so as to minimize risk and maximize performance [5]. In multi-agent systems where agents need to reach consensus concerning a set of

L.C. van der Gaag (Ed.): ECSQARU 2013, LNAI 7958, pp. 364–375, 2013.

propositions, the use of borderline cases can allow agents to adapt their beliefs so as to reach a compromise with others, whilst maintaining a certain level of internal consistency [6]. Furthermore, in multi-agent dialogues a bipolar approach can help to distinguish between strong and weak viewpoints in opinion formation [8]. Another application area of growing importance is in the representation of so-called flexible specifications for adaptive autonomous systems. The deployment of autonomous systems in complex dynamic environments tends to naturally result in a tension between the requirement that the system's behaviour conforms to a predefined specification, and the need for it to be sufficiently flexible so as to cope with severe uncertainty and unexpected scenarios. For example, it might find itself in situations not envisaged by its designers, where all available actions result in some violation of its specification. In such cases, a more flexible form of specification may allow for some constraints to be only borderline satisfied in certain conditions. Furthermore, the blurring of concept boundaries in the interpretation would then permit some aspect of gradedness, potentially allowing the system to choose between different suboptimal possibilities.

In all of the above application areas the adequate representation of epistemic uncertainty combined with bipolarity is also of central importance. 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 partial knowledge of language conventions resulting in agents being unsure about conceptual boundaries. Here we extend bipolar belief measures, recently proposed in [7], which combine probabilistic uncertainty with truth-gaps as represented in Kleene's strong three-valued logic [4]. More specifically, the main contribution of this paper is the introduction of natural measures of conditional belief within this framework. We then discuss their properties and relate and contrast these measures to existing definitions of conditional belief in the literature such as in Dempster-Shafer theory and fuzzy logic.

An outline of the paper is as follows: Section 2 introduces valuation pairs as a bipolar truth-model based on Kleene's three-valued logic. Section 3 defines bipolar belief pairs in terms of probability distributions over the set of valuation pairs and shows their relationship to lower and upper membership functions in interval-valued fuzzy logic. Extending this idea, section 4 proposed definitions for conditional belief pairs and investigates their properties. Finally, in section 5 we have conclusions and further discussion of potential applications of the framework.

2 Valuation Pairs

In this section, we introduce valuation pairs as a bipolar model of truth which allows for the explicit representation of borderline cases. Typical examples are declarative sentences containing vague adjectives e.g. *low*, *tall*, *fast* etc, although truth-gaps can of course result from other sources of vagueness such as from verbs and nouns. We now propose to model truth-gaps by replacing a single binary, true or false, valuation on propositions with distinct lower and upper valuations representing absolutely true and not absolutely false respectively. Borderline cases then correspond to those sentences in which the lower and upper valuation differ.

Let $\mathcal{L}$ be a language of propositional logic with connectives $\wedge$, $\vee$ and $\neg$ and propositional variables $P = \{p_1, \ldots, p_n\}$, and let $S\mathcal{L}$ denote the sentences of $\mathcal{L}$ as generated recursively from P by application of the connectives. A valuation pair on $S\mathcal{L}$ consists of two binary functions $\underline{v}$ and $\overline{v}$ representing lower and upper truth-values. The underlying idea is that $\underline{v}$ represents the strong criterion of *absolutely true* while $\overline{v}$ represents the weaker criteria of *not absolutely false*. In accordance with [11], we might think of a sentence being absolutely true as meaning that it can be uncontroversially asserted without any risk of censure, while being not absolutely false only means that it is acceptable to assert i.e. one can get away with such an assertion. 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 p='the suspect was short', even though this statement would not be viewed as being absolutely true. One possible bipolar model of the concept *short* exhibiting such truth-gaps could be as follows: Let h be the height of the suspect and suppose that *short* is defined in terms of lower and upper thresholds $\underline{h} \leq \overline{h}$ on heights. In this case p is *absolutely true* if $h \leq \underline{h}$, *absolutely false* if $h > \overline{h}$ and *borderline* if $\underline{h} < h \leq \overline{h}$ (see figure 1).

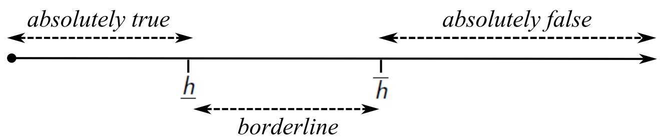

Fig. 1. A bipolar interpretation of the concept short

It is important to note that in this model truth-gaps corresponding to different lower and upper truth valuations are not the result of epistemic uncertainty concerning the state of the world but rather due to inherent flexibility in the underlying language conventions. In other words, a truth-gap (or middle truth-value in three-valued logic) does not represent an *uncertain* epistemic state. For example, given absolute certainty about suspect's height the proposition p may then be *known* to be borderline because of the inherent flexibility (or vagueness) in the definition of the concept short i.e. because $\underline{h} < h \leq \overline{h}$. The potential confusion resulting from applying many-valued logic to model epistemic uncertainty is highlighted by Dubois in [2]. In the sequel we shall emphasize the truth-value status of the intermediate case by using the term *borderline* rather than 'uncertain' or 'unknown' as originally suggested by Kleene [4].

Definition 1. *Kleene Valuation Pairs*
A Kleene valuation pair on $\mathcal{L}$ is a pair of functions $\boldsymbol{v} = (\underline{v}, \overline{v})$ where $\underline{v} : S\mathcal{L} \to \{0,1\}$ and $\overline{v} : S\mathcal{L} \to \{0,1\}$ such that $\underline{v} \leq \overline{v}$ and where $\forall \theta, \varphi \in S\mathcal{L}$ 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* $\overline{v}(\theta \vee \varphi) = \max(\overline{v}(\theta), \overline{v}(\varphi))$

We use $\mathbb{V}$ to denote the set of all Kleene valuation pairs on $\mathcal{L}$.

The link to three-valued logic is clear 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. From definition 1 we can then determine truth-tables for the connectives $\wedge$, $\vee$ and $\neg$ in terms of the truth-values $\{\mathbf{t}, \mathbf{b}, \mathbf{f}\}$ identical to those of Kleene's logic [4]. Shapiro [14] has recently proposed the use of Kleene's three-valued logic to model truth-gaps in vague predicates, arguing that Kleene's truth tables 'reflect the open-texture of vague predicates'. For example, if instead we were to adopt Lukasiewicz logic [10] this would mean that for two borderline propositional variables their conjunction would be absolutely false, even though neither conjunct was absolutely false. This would seem to be a totally unwarranted elimination of vagueness. One might of course consider a non-functional calculus for valuation pairs based, for example, on supervaluationist principles as explored in Lawry and Tang [5]. Another possibility would be to introduce many-valued logics with more than three truth-values. From the current perspective this would correspond to propositions being *borderline* to differing degrees. However, the representational utility of making such distinctions between borderline cases is not entirely clear, as is discussed in more details in [5].

Definition 2. *For $\theta, \varphi \in S\mathcal{L}$, θ and φ are equivalent, denoted $\theta \equiv \varphi$, if and only if $\forall \boldsymbol{v} \in \mathbb{V}$ $\boldsymbol{v}(\theta) = \boldsymbol{v}(\varphi)$.*

The following theorem identifies a number of well known equivalences from Kleene three-valued logic.

Theorem 1. *Important Equivalences [7]*
$\forall \theta, \varphi, \psi \in SL$ the following sentences are equivalent:

- *De Morgan's Laws:* $\neg(\theta \wedge \varphi) \equiv \neg\theta \vee \neg\varphi$ *and* $\neg(\theta \vee \varphi) \equiv \neg\theta \wedge \neg\varphi$
- *Double Negation:* $\neg(\neg\theta) \equiv \theta$
- *Idempotence:* $\theta \wedge \theta \equiv \theta$ *and* $\theta \vee \theta \equiv \theta$
- *Commutativity:* $\theta \vee \varphi \equiv \varphi \vee \theta$ *and* $\theta \wedge \varphi \equiv \varphi \wedge \theta$
- *Associativity:* $\theta \vee (\varphi \vee \psi) \equiv (\theta \vee \varphi) \vee \psi$ *and* $\theta \wedge (\varphi \wedge \psi) \equiv (\theta \wedge \varphi) \wedge \psi$
- *Distributivity:* $\theta \vee (\varphi \wedge \psi) \equiv (\theta \vee \varphi) \wedge (\theta \vee \psi)$ *and* $\theta \wedge (\varphi \vee \psi) \equiv (\theta \wedge \varphi) \vee (\theta \wedge \psi)$

Kleene valuation pairs do not completely satisfy the laws of non-contradiction and excluded middle in borderline cases. While it is the case that for any sentence $\varphi \in S\mathcal{L}$, $\underline{v}(\varphi \wedge \neg\varphi) = 0$ and $\overline{v}(\varphi \wedge \neg\varphi) = 1$ the same equalities do not necessarily hold for the corresponding upper and lower valuations respectively. In fact, any such partial failure of the laws of non-contradiction and excluded middle exactly correspond with φ being a borderline case, as the following result shows.

Theorem 2. *$\boldsymbol{v} \in \mathbb{V}$, $\boldsymbol{v}(\varphi) = \mathbf{b}$ if and only if $\overline{v}(\varphi \wedge \neg\varphi) = 1$ if and only if $\underline{v}(\varphi \vee \neg\varphi) = 0$.*

Proof. ($\Rightarrow$) Suppose $\boldsymbol{v}(\varphi) = \mathbf{b}$ then $\underline{v}(\varphi) = 0 \Rightarrow \overline{v}(\neg\varphi) = 1$ and since also $\overline{v}(\varphi) = 1 \Rightarrow \overline{v}(\varphi \wedge \neg\varphi) = 1$. ($\Leftarrow$) $\overline{v}(\varphi \wedge \neg\varphi) = 1 \Rightarrow \overline{v}(\varphi) = 1$ and also $\Rightarrow \overline{v}(\neg\varphi) = 1 \Rightarrow \underline{v}(\varphi) = 0$. Furthermore, by duality and de Morgan's law (theorem 2) it follows that $\underline{v}(\varphi \vee \neg\varphi) = 0$ if and only if $\overline{v}(\varphi \wedge \neg\varphi) = 1$ as required.

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 $\boldsymbol{v}_1$ is less semantically precise than $\boldsymbol{v}_2$ if they disagree only for some subset of sentences of $\mathcal{L}$, which being identified as either absolutely true or absolutely false by $\boldsymbol{v}_2$, are classified as borderline by $\boldsymbol{v}_1$.

Definition 3. *Semantic Precision*
$\boldsymbol{v}_1 \preceq \boldsymbol{v}_2$ *iff* $\forall\theta \in S\mathcal{L}$ $\underline{v_1}(\theta) \leq \underline{v_2}(\theta)$ *and* $\overline{v_1}(\theta) \geq \overline{v_2}(\theta)$.

Shapiro [14] proposed essentially the same ordering of interpretations which he refers to as *sharpening* i.e. $\boldsymbol{v}_1 \preceq \boldsymbol{v}_2$ means that $\boldsymbol{v}_2$ extends or sharpens $\boldsymbol{v}_1$. Here we shall refer to $\preceq$ as the *semantic precision* ordering on valuation pairs whereby, if $\boldsymbol{v}_1 \preceq \boldsymbol{v}_2$ then $\boldsymbol{v}_1$ tends to classify more sentences of $\mathcal{L}$ as *borderline* than $\boldsymbol{v}_2$. In other words, one might think of $\preceq$ as ordering valuation pairs according to their relative vagueness.

3 Belief Pairs

Within the proposed bipolar framework, uncertainty concerning the sentences of $\mathcal{L}$ effectively corresponds to uncertainty as to which is the correct Kleene valuation pair for $\mathcal{L}$. In practice, there are likely to be many different sources of this uncertainty, however one natural division of uncertainty types is as follows:

- *Semantic uncertainty* about the linguistic conventions defining concepts relevant to the sentences of $\mathcal{L}$. For example, an agent may be uncertain as to whether or not a proposition such as 'the suspect is short' is *absolutely true* or *not absolutely false* even if the suspect's height h is known precisely. For instance, this might manifest itself in terms of uncertainty about the exact values of the thresholds $\underline{h}$ and $\overline{h}$ (see figure 1). This uncertainty naturally arises from the distributed manner in which language is learnt through communications with other agents across a population of interacting agents.
- *Stochastic uncertainty* arising from a lack of knowledge concerning the state of the world. For example, being uncertain about the suspect's height h in the proposition 'the suspect is short'.

In general we view uncertainty as being epistemic in nature, resulting from a lack of knowledge concerning either, the state of the world to which propositions refer, or the linguistic conventions governing the assertability of propositions as part of communications. Viewing semantic uncertainty as being epistemic in nature requires that agents make the assumption that there is a correct underlying interpretation of the language $\mathcal{L}$, but about which they may be uncertain. This is a weaker version of the epistemic theory of vagueness as expounded by Timothy Williamson [16] referred to as the *epistemic stance* [9]. Williamson's theory assumes that for the extension of a vague concept there is a precise but unknown boundary between it and the extension of its negation. In contrast the epistemic stance corresponds to the more pragmatic view that individuals, when faced with decision problems about what to assert, find it useful as part of a decision making strategy to simply *assume* that there is an underlying correct interpretation of $\mathcal{L}$. In other words, when deciding what to assert agents behave as is the epistemic theory is correct. Another difference between the epistemic theory and our current approach is that the former assumes that the underlying truth model is classical while here we assume a bipolar model which can exhibit truth-gaps.

In the following definition we assume that uncertainty is quantified by a probability measure w on the set of Kleene valuation pairs $\mathbb{V}$.

Definition 4. *Kleene Belief Pairs [7]*
Let $\mathbb{V}$ be the set of all Kleene valuation pairs on $\mathcal{L}$ and let w be a probability distribution defined on $\mathbb{V}$ so that $w(\boldsymbol{v})$ is the agent's subjective belief that $\boldsymbol{v}$ is the true valuation pair for $\mathcal{L}$. Then $\boldsymbol{\mu} = (\underline{\mu}, \overline{\mu})$ is a Kleene belief pair where $\forall \theta \in S\mathcal{L}$, $\underline{\mu}(\theta) = w(\{\boldsymbol{v} \in \mathbb{V} : \underline{v}(\theta) = 1\})$ and $\overline{\mu}(\theta) = w(\{\boldsymbol{v} \in \mathbb{V} : \overline{v}(\theta) = 1\})$.

There is a clear rationality argument for defining belief measures in this manner when Kleene valuation pairs are the underlying truth model for $\mathcal{L}$. From a general result due to Paris [12], it follows that an agent can only avoid Dutch books where the outcomes of bets are dependent on lower (upper) Kleene valuations if their belief measures on $S\mathcal{L}$ correspond to lower (upper) belief measures as given in definition 4. This idea is explored in more detail in Lawry and Tang [5] in the context of lower and upper bets. The following theorem highlights a number of properties of Kleene belief pairs, including additivity. The latter property in particular distinguishes Kleene Belief pairs from Dempster-Shafer belief and plausibility measures [13] on $S\mathcal{L}$ which are not, in general, additive.

Theorem 3. *For all $\theta, \varphi \in S\mathcal{L}$, the following hold:*

- $\underline{\mu}(\theta) \leq \overline{\mu}(\theta)$
- $\underline{\mu}(\neg\theta) = 1 - \overline{\mu}(\theta)$ *and* $\overline{\mu}(\neg\theta) = 1 - \underline{\mu}(\theta)$.
- $\underline{\mu}(\theta \vee \varphi) = \underline{\mu}(\theta) + \underline{\mu}(\varphi) - \underline{\mu}(\theta \wedge \varphi)$ *and* $\overline{\mu}(\theta \vee \varphi) = \overline{\mu}(\theta) + \overline{\mu}(\varphi) - \overline{\mu}(\theta \wedge \varphi)$

It is also interesting to note that a special case of Kleene belief pairs has the same calculus as the interval (or type 2) fuzzy membership functions proposed by Zadeh [17]. This is the case of Kleene belief pairs in which there is only uncertainty about the level of semantic precision of the valuation pair. More formally we have the following result:

Theorem 4. *[7] Let w be a probability distribution on $\mathbb{V}$ for which $\{\boldsymbol{v} \in \mathbb{V} : w(\boldsymbol{v}) > 0\} = \{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k\}$ can be ordered such that $\boldsymbol{v}_1 \preceq \boldsymbol{v}_2 \ldots \preceq \boldsymbol{v}_k$. In this case $\boldsymbol{\mu}$ satisfies the following properties; $\forall \theta, \varphi \in S\mathcal{L}$,*

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

Example 1. Recall the example from section 2 concerning the proposition $p =$'the suspect is short', where the concept short is defined by two thresholds on height $0 \leq \underline{h} \leq \overline{h}$, so that an individual is absolutely *short* if their height is less than or equal to $\underline{h}$ and absolutely not short if their height is greater than $\overline{h}$. Hence, if the suspect's height is known to be h then an agent's beliefs about the interpretation of $\mathcal{L}$ can be modelled by a valuation pair $\boldsymbol{v}$ such that:

$$\underline{v}(p) = 1 \text{ if and only if } h \leq \underline{h} \text{ and } \overline{v}(p) = 1 \text{ if and only if } h \leq \overline{h}$$

We might further assume, perhaps reasonably in this case, that the agent's semantic uncertainty with regard to p is limited to uncertainty about the actual values of the thresholds $\underline{h}$ and $\overline{h}$. Further suppose that the agent's beliefs about these thresholds is represented by a joint probability density function f on $(\underline{h}, \overline{h})$ satisfying:

$$\int_0^\infty \int_{\underline{h}}^\infty f(\underline{h}, \overline{h}) \, d\overline{h} \, d\underline{h} = 1$$

Based on this the agent can define a lower measure of their belief in p, $\underline{\mu}(p)$, corresponding to the probability that the lower threshold $\underline{h} \geq h$ and similarly and upper measure, $\overline{\mu}(p)$, corresponding to the probability that the upper threshold $\overline{h} \geq h$ i.e.

$$\underline{\mu}(p) = \int_h^\infty \int_{\underline{h}}^\infty f(\underline{h}, \overline{h}) \, d\overline{h} \, d\underline{h} \text{ and } \overline{\mu}(p) = \int_h^\infty \int_0^{\overline{h}} f(\underline{h}, \overline{h}) \, d\underline{h} \, d\overline{h}$$

Now suppose that in this case the agent believes that $\underline{h}$ and $\overline{h}$ are independent variables both with triangular distributions centered around $130cm$ and $150cm$ respectively. More specifically; $f(\underline{h}, \overline{h}) = f_1(\underline{h}) \times f_2(\overline{h})$ where

$$f_1(\underline{h}) = \begin{cases} \frac{\underline{h}-120}{100} : \underline{h} \in [120, 130) \\ \frac{140-\underline{h}}{100} : \underline{h} \in [130, 140] \\ 0 : \text{ otherwise} \end{cases} \quad \text{and } f_2(\overline{h}) = \begin{cases} \frac{\overline{h}-140}{100} : \overline{h} \in [140, 150) \\ \frac{160-\overline{h}}{100} : \overline{h} \in [150, 160] \\ 0 : \text{ otherwise} \end{cases}$$

In this case the resulting values for $\underline{\mu}(p)$ and $\overline{\mu}(p)$ are shown in figure 2 as height h varies. Similarly, figure 3 shows the agent's belief that p is a borderline proposition, as quantified by $\overline{\mu}(p) - \underline{\mu}(p)$, for different values of h.

We can also consider the possibility that the agent is uncertain about the value of suspect's height. Suppose that the agent's knowledge about h is characterised

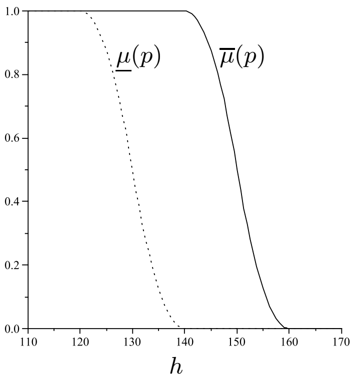

Fig. 2. Lower and upper belief values for proposition p as the suspect's height h varies

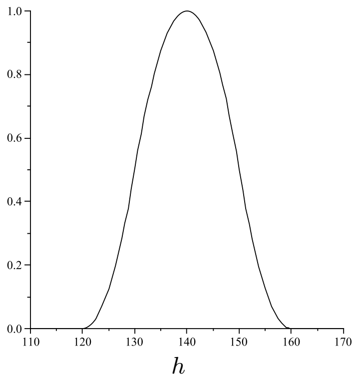

Fig. 3. Belief that p is a borderline proposition, given by $\overline{\mu}(p) - \underline{\mu}(p)$, as h varies

by a probability density function g and further suppose that h is taken to be independent of the thresholds $\underline{h}$ and $\overline{h}$. This additional uncertainty can then be included in the calculation of the lower and upper belief measures as follows:

$$\underline{\mu}(p) = \int_0^\infty \int_h^\infty \int_{\underline{h}}^\infty f(\underline{h}, \overline{h}) g(h) \; \mathrm{d}\overline{h} \; \mathrm{d}\underline{h} \; \mathrm{d}h \text{ and}$$

$$\overline{\mu}(p) = \int_0^\infty \int_h^\infty \int_0^{\overline{h}} f(\underline{h}, \overline{h}) g(h) \; \mathrm{d}\underline{h} \; \mathrm{d}\overline{h} \; \mathrm{d}h$$

For example, if g is a normal distribution with mean 140 and standard deviation 7 then $\underline{\mu}(p) = 0.1092$ and $\overline{\mu}(p) = 0.8908$.

4 Conditional Belief Pairs

In this section we propose a conditioning model by which agents can update their subjective belief pairs on the basis of new information concerning the absolute truth and absolute falsity of sentences in $\mathcal{L}$. In view of the inherently probabilistic nature of belief pairs, one obvious method is based on conditional probabilities. For this approach we assume that new knowledge takes the form of lower and upper valuation constraints, which it is then assumed that the *correct* valuation for $\mathcal{L}$ must satisfy. From the perspective of the above discussion on uncertainty in a bipolar context, we can think of such constraints as providing new information both about the state of the world and about the underlying interpretation of $\mathcal{L}$. This knowledge allows us to define conditional lower and upper belief measures by determining a posterior distribution on valuation pairs from the prior w, according to the standard definition of conditional probability

Definition 5. *Conditional Belief Pairs*
Suppose an agent obtains new knowledge regarding the assertability of sentences in $S\mathcal{L}$ in the form of a set K of constraints on lower and upper valuations of the following form:

$$K = \{\underline{v}(\theta_1) = 1, \ldots, \underline{v}(\theta_t) = 1, \overline{v}(\varphi_1) = 1, \ldots, \overline{v}(\varphi_s) = 1\}$$

Then we define lower and upper conditional belief measures conditional on K as follows:

$$\underline{\mu}(\theta|K) = \frac{w(\{\boldsymbol{v} \in \mathbb{V}(K) : \underline{v}(\theta) = 1\})}{w(\mathbb{V}(K))} \text{ and } \overline{\mu}(\theta|K) = \frac{w(\{\boldsymbol{v} \in \mathbb{V}(K) : \overline{v}(\theta) = 1\})}{w(\mathbb{V}(K))}$$

where $\mathbb{V}(K) \subseteq \mathbb{V}$ denotes the set of Kleene valuation pairs on $\mathcal{L}$ which satisfy the constraints K.

A possible source of knowledge constraints of the form given in definition 5, is from strong and weak assertion in agent dialogues [8]. For example, a witness might describe the suspect as ‘absolutely short’ or ‘definitely short’. Alternatively, they might only be prepared to say that the suspect was ‘possibly short’ or ‘short-ish’. The former might be regarded as strong assertions concerning the proposition p =‘the suspect is tall’ corresponding to the knowledge that $\boldsymbol{v}(p) = \mathbf{t}$. In contrast, the latter are weak assertions corresponding to $\boldsymbol{v}(p) \neq \mathbf{f}$ and $\boldsymbol{v}(p) = \mathbf{b}$ respectively. One can then envisage a knowledge base K as in definition 5, being derived from a dialogue with other agents consisting of such strong and weak assertions.

We now consider the special cases where $K = \{\underline{v}(\varphi) = 1\}$, $K = \{\overline{v}(\varphi) = 1\}$ and $K = \{\underline{v}(\varphi) = 0, \overline{v}(\varphi) = 1\}$ for some sentence $\varphi \in S\mathcal{L}$. Notice, that these correspond to the knowledge that $\boldsymbol{v}(\varphi) = \mathbf{t}$, $\boldsymbol{v}(\varphi) \neq \mathbf{f}$ and $\boldsymbol{v}(\varphi) = \mathbf{b}$ respectively.

Theorem 5

$$\underline{\mu}(\theta|\underline{v}(\varphi) = 1) = \frac{\underline{\mu}(\theta \wedge \varphi)}{\underline{\mu}(\varphi)} \text{ and } \overline{\mu}(\theta|\underline{v}(\varphi) = 1) = \frac{\overline{\mu}(\theta \vee \neg\varphi) - \overline{\mu}(\neg\varphi)}{1 - \overline{\mu}(\neg\varphi)}$$

Proof

$$\forall \theta, \varphi \in S\mathcal{L},\ \underline{\mu}(\theta|\underline{v}(\varphi) = 1) = \frac{w(\{\boldsymbol{v} \in \mathbb{V} : \underline{v}(\theta) = 1, \underline{v}(\varphi) = 1\})}{w(\{\boldsymbol{v} \in \mathbb{V} : \underline{v}(\varphi) = 1\})}$$

$$= \frac{w(\{\boldsymbol{v} \in \mathbb{V} : \underline{v}(\theta \wedge \varphi) = 1\})}{w(\{\boldsymbol{v} \in \mathbb{V} : \underline{v}(\varphi) = 1\})} \text{ by definition 1 } = \frac{\underline{\mu}(\theta \wedge \varphi)}{\underline{\mu}(\varphi)} \text{ by definition 4.}$$

In addition, by duality we have that:

$$\overline{\mu}(\theta|\underline{v}(\varphi) = 1) = 1 - \underline{\mu}(\neg\theta|\underline{\varphi} = 1) \text{by the above } = 1 - \frac{\underline{\mu}(\neg\theta \wedge \varphi)}{\underline{\mu}(\varphi)}$$

$$= \frac{\underline{\mu}(\varphi) - \underline{\mu}(\neg\theta \wedge \varphi)}{\underline{\mu}(\varphi)} = \frac{1 - \overline{\mu}(\neg\varphi) - 1 + \overline{\mu}(\neg(\neg\theta \wedge \varphi))}{1 - \overline{\mu}(\neg\varphi)}$$

$$= \frac{\overline{\mu}(\theta \vee \neg\varphi) - \overline{\mu}(\neg\varphi)}{1 - \overline{\mu}(\neg\varphi)} \text{ by de Morgan's law (theorem 1)}$$

Theorem 6

$$\underline{\mu}(\theta|\overline{v}(\varphi)=1)=\frac{\underline{\mu}(\theta\vee\neg\varphi)-\underline{\mu}(\neg\varphi)}{1-\underline{\mu}(\neg\varphi)} \text{ and } \overline{\mu}(\theta|\overline{v}(\varphi)=1)=\frac{\overline{\mu}(\theta\wedge\varphi)}{\overline{\mu}(\varphi)}$$

Proof. Similar to theorem 5.

Notice that the lower and upper conditions in theorem 6 have the same definition relative to the underlying belief measures as conditional belief and plausibility in Dempster-Shafer theory [13]. However, recall that Kleene belief pairs are not Dempster Shafer measures since, for example, they satisfy additivity (see theorem 3).

Theorem 7

$$\underline{\mu}(\theta|\boldsymbol{v}(\varphi)=\mathbf{b})=\frac{\underline{\mu}(\theta\vee\varphi\vee\neg\varphi)-\underline{\mu}(\varphi\vee\neg\varphi)}{1-\underline{\mu}(\varphi\vee\neg\varphi)} \text{ and } \overline{\mu}(\theta|\boldsymbol{v}(\varphi)=\mathbf{b})=\frac{\overline{\mu}(\theta\wedge\varphi\wedge\neg\varphi)}{\overline{\mu}(\varphi\wedge\neg\varphi)}$$

Proof

$$\overline{\mu}(\theta|\boldsymbol{v}(\varphi)=\mathbf{b})=\frac{w(\{\boldsymbol{v}:\boldsymbol{v}(\varphi)=\mathbf{b},\overline{v}(\theta)=1\})}{w(\{\boldsymbol{v}:\boldsymbol{v}(\varphi)=(0,1)\})}=\frac{w(\{\boldsymbol{v}:\overline{v}(\varphi\wedge\neg\varphi)=1,\overline{v}(\theta)=1\})}{w(\{\boldsymbol{v}:\boldsymbol{v}(\varphi)=(0,1)\})}$$

$$\text{by theorem 2 }=\frac{w(\{\boldsymbol{v}:\overline{v}(\theta\wedge\varphi\wedge\neg\varphi)=1\})}{w(\{\boldsymbol{v}:\overline{v}(\varphi\wedge\neg\varphi)=1\})}=\frac{\overline{\mu}(\theta\wedge\varphi\wedge\neg\varphi)}{\overline{\mu}(\varphi\wedge\neg\varphi)}$$

Also we have that,

$$\underline{\mu}(\theta|\boldsymbol{v}(\varphi)=\mathbf{b})=\frac{w(\{\boldsymbol{v}:\underline{v}(\theta)=1,\boldsymbol{v}(\varphi)=\mathbf{b}\})}{w(\{\boldsymbol{v}:\boldsymbol{v}=\mathbf{b}\})}=\frac{w(\{\boldsymbol{v}:\underline{v}(\theta)=1,\underline{v}(\varphi\vee\neg\varphi)=0\})}{w(\{\boldsymbol{v}:\underline{v}(\varphi\vee\neg\varphi)=0\})}$$

$$\text{by theorem 2}=\frac{w(\{\boldsymbol{v}:\underline{v}(\varphi\vee\neg\varphi)=0\})-w(\{\boldsymbol{v}:\underline{v}(\theta)=0,\underline{v}(\varphi\vee\neg\varphi)=0\})}{w(\{\boldsymbol{v}:\underline{v}(\varphi\vee\neg\varphi)=0\})}$$

$$=\frac{w(\{\boldsymbol{v}:\underline{v}(\varphi\vee\neg\varphi)=0\})-w(\{\boldsymbol{v}:\underline{v}(\theta\vee\varphi\vee\neg\varphi)=0\})}{w(\{\boldsymbol{v}:\underline{v}(\varphi\vee\neg\varphi)=0\})}$$

$$=\frac{1-\underline{\mu}(\varphi\vee\neg\varphi)-(1-\underline{\mu}(\theta\vee\varphi\vee\neg\varphi))}{1-\underline{\mu}(\varphi\vee\neg\varphi)}=\frac{\underline{\mu}(\theta\vee\varphi\vee\neg\varphi)-\underline{\mu}(\varphi\vee\neg\varphi)}{1-\underline{\mu}(\varphi\vee\neg\varphi)}$$

Corollary 1. *Let w be a probability distribution on $\mathbb{V}$ for which $\{\boldsymbol{v}\in\mathbb{V}:w(\boldsymbol{v})>0\}=\{\boldsymbol{v}_1,\ldots,\boldsymbol{v}_k\}$ can be ordered such that $\boldsymbol{v}_1\preceq\boldsymbol{v}_2\ldots\preceq\boldsymbol{v}_k$. Then for $\theta,\varphi\in SL$ it holds that:*

$$\underline{\mu}(\theta|\underline{v}(\varphi)=1)=\begin{cases}\frac{\underline{\mu}(\theta)}{\underline{\mu}(\varphi)} & :\underline{\mu}(\theta)\leq\underline{\mu}(\varphi)\\ 1 & : \textit{otherwise}\end{cases} \textit{ and}$$

$$\overline{\mu}(\theta|\underline{v}(\varphi)=1)=\begin{cases}\frac{\overline{\mu}(\theta)+\underline{\mu}(\varphi)-1}{\underline{\mu}(\varphi)} & :\overline{\mu}(\theta)+\underline{\mu}(\varphi)\geq 1\\ 0 & : \textit{otherwise}\end{cases}$$

$$\underline{\mu}(\theta|\overline{v}(\varphi)=1)=\begin{cases}\frac{\underline{\mu}(\theta)+\overline{\mu}(\varphi)-1}{\overline{\mu}(\varphi)} & :\underline{\mu}(\theta)+\overline{\mu}(\varphi)\geq 1\\ 0 & : \textit{otherwise}\end{cases} \textit{ and}$$

$$\overline{\mu}(\theta|\overline{v}(\varphi)=1)=\begin{cases}\frac{\overline{\mu}(\theta)}{\overline{\mu}(\varphi)} & :\overline{\mu}(\theta)\leq\overline{\mu}(\varphi)\\ 1 & : \textit{otherwise}\end{cases}$$

Proof. Follows immediately from theorem 4 and theorems 5 and 6.

Notice that in corollary 1 $\underline{\mu}(\theta|\underline{v}(\varphi) = 1)$ and $\overline{\mu}(\theta|\overline{v}(\varphi) = 1)$ correspond to the Goguen implication operator [3] applied to the lower and upper belief values of θ and φ respectively.

Example 2. Recall the proposition $p =$'the suspect is short' as described in example 1. Now consider an additional proposition $q =$'the suspect is very short' where the concept *very short* is defined by lower and upper height thresholds $\underline{h}'$ and $\overline{h}'$. Further suppose that these thresholds are dependent on the thresholds of *short*, according to $\underline{h}' = 0.9\underline{h}$ and $\overline{h}' = 0.9\overline{h}$. Further suppose that, as in example 1, the semantic and stochastic uncertainty is modelled by the joint distribution f on the threshold $\underline{h}$ and $\overline{h}$, and the distribution g on h respectively. Now suppose that the agent learns that the suspect is *borderline very short*. How does this change their level of belief that the suspect is short? In other words, what are the values of the conditional beliefs $\underline{\mu}(p|\boldsymbol{v}(q) = \mathbf{b})$ and $\overline{\mu}(p|\boldsymbol{v}(q) = \mathbf{b})$? Notice that given the above definition of $\overline{h}'$ then it follows that $h \leq \overline{h}'$ implies that $h \leq \overline{h}$ and hence $\overline{\mu}(p|\boldsymbol{v}(q) = \mathbf{b}) = 1$. Now in this example $w(\{\boldsymbol{v} : \boldsymbol{v}(q) = b\})$ corresponds to the probability that $\underline{h}' \leq h \leq \overline{h}'$ or alternatively that $\overline{h} \geq \frac{h}{0.9}$ and $\underline{h} \leq \frac{h}{0.9}$. Hence, we have that:

$$w(\{\boldsymbol{v} : \boldsymbol{v}(q) = b\}) = \int_0^\infty \int_0^{\frac{h}{0.9}} \int_{\frac{h}{0.9}}^\infty f(\underline{h}, \overline{h}) g(h) \, \mathrm{d}\overline{h} \, \mathrm{d}\underline{h} \, \mathrm{d}h = 0.2625$$

Similarly we have that:

$$w(\{\boldsymbol{v} : \underline{v}(p) = 1, \boldsymbol{v}(q) = b\}) = \int_0^\infty \int_h^{\frac{h}{0.9}} \int_{\frac{h}{0.9}}^\infty f(\underline{h}, \overline{h}) g(h) \, \mathrm{d}\overline{h} \, \mathrm{d}\underline{h} \, \mathrm{d}h = 0.0888$$

Hence,

$$\underline{\mu}(p|\boldsymbol{v}(q) = \mathbf{b}) = \frac{0.0888}{0.2625} = 0.3383$$

In comparison with the values obtained in example 1 we see that both $\underline{\mu}(p|\boldsymbol{v}(q) = b) > \underline{\mu}(p)$ and $\overline{\mu}(p|\boldsymbol{v}(q) = b) > \overline{\mu}(p)$. Clearly then, learning that q is a borderline case is informative when trying to determine the truth value of p. This emphasises the difference in terms of conditioning between the two distinct interpretations of truth-gaps (or middle truth-values) either as being borderline cases due to inherent vagueness or as representing epistemic ignorance. Indeed, if all we were to learn was that the truth value of q was *unknown* then this would tell us nothing about the truth-value of p, and therefore conditioning would not result in any change to belief values.

5 Conclusion and Discussion

In this paper we have proposed definitions for lower and upper conditional belief pairs, extending the framework introduced in [7]. The properties of these measures has been investigated and the relationship to conditional belief in existing uncertainty theories has been highlighted.

The belief pairs framework, incorporating the conditional measures proposed in this paper, is sufficiently rich to capture aspects of both stochastic and semantic uncertainty together with indeterminism in the underlying truth model. This can permit the definition of more flexible rules and specifications for intelligent autonomous systems, as well as providing an enhanced model of decision making in the presence of both uncertainty and conceptual vagueness. For example, one can envisage flexible requirements concerning the relationship between a pair of propositions p and q which include the requirement that p must be *absolutely true* in those circumstances in which q is only *borderline true*. Furthermore, in the presence of significant uncertainty probabilistic requirements may be more appropriate in the form of constraints on lower and upper condition beliefs e.g. $\underline{\mu}(p|\boldsymbol{v}(q) = b) \geq \alpha$ for a suitable confidence level α. Future work will aim to explore the application of the belief pairs framework to the formal representation of flexible specifications and their verification.

Acknowledgements. This work is partially funded by EPSRC grant EP/J01205X/1.

References

1. Dubois, D., Prade, H.: An Introduction to Bipolar Representations of Information and Preference. Int. Journal of Intelligent Systems 23(8), 866–877 (2008)
2. Dubois, D.: On Ignorance and Contradiction Considered as Truth-Values. Logic Journal of the IGPL 16(2), 195–216 (2008)
3. Goguen, J.A.: The Logic of Inexact Concepts. Synthese 19, 325–373 (1968)
4. Kleene, S.C.: Introduction to Metamathematics. D. Van Nostrand Company Inc., Princeton (1952)
5. Lawry, J., Tang, Y.: On Truth-gaps, Bipolar Belief and the Assertability of Vague Propositions. Artificial Intelligence 191-192, 20–41 (2012)
6. Lawry, J., Dubois, D.: A Bipolar Framework for Combining Beliefs about Vague Propositions. In: Proceedings of 13th International Conference on Principles of Knowledge Representation and Reasoning, pp. 530–540 (2012)
7. Lawry, J., González-Rodríguez, I.: A Bipolar Model of Assertability and Belief. International Journal of Approximate Reasoning 52, 76–91 (2011)
8. Lawry, J.: Imprecise Bipolar Belief Measures Based on Partial Knowledge from Agent Dialogues. In: Deshpande, A., Hunter, A. (eds.) SUM 2010. LNCS, vol. 6379, pp. 205–218. Springer, Heidelberg (2010)
9. Lawry, J.: Appropriateness Measures: An Uncertainty Model for Vague Concepts. Synthese 161(2), 255–269 (2008)
10. Lukasiewicz, J.: O logice trojwartosciowej (On three-valued logic). Ruch Filozoficzny 5, 170–171 (1920)
11. Parikh, R.: Vague Predicates and Language Games. Theoria XI(27), 97–107 (1996)
12. Paris, J.B.: A Note on the Dutch Book Method. In: Proceedings of ISIPTA 2001, Ithaca, New York (2001)
13. Shafer, G.: A Mathematical Theory of Evidence. Princeton University Press (1976)
14. Shapiro, S.: Vagueness in Context. Oxford University Press (2006)
15. van Deemter, K.: Utility and Language Generation: The Case of Vagueness. Journal of Philosophical Logic 38, 607–632 (2009)
16. Williamson, T.: Vagueness. Routledge (1994)
17. Zadeh, L.A.: The Concept of a Linguistic Variable and its Application to Approximate Reasoning: I. Information Sciences 8, 199–249 (1975)