Conditional quantification, or poor man's probability

Michiel van Lambalgen*
ILLC
University of Amsterdam
Plantage Muidergracht 24
NL 1018 TV Amsterdam
michiell@wins.uva.nl

June 25, 1999

1 Introduction

The purpose of this note is to take a few steps toward a first order logic of incomplete information, based on the idea that such a logic may be obtained by importing an essential ingredient of probability theory, conditional expectation, into the logic. The reason that we focus on conditional expectation is that it is often used to model a kind of incomplete knowledge that is of interest in itself: incomplete knowledge about the value of a variable, where the 'degree of incompleteness' is given by some algebraic structure. Conditional expectation often comes into play when what can be described as a change of granularity is involved. Formal details will be given below, but a typical example is this. Let Ω be a sample space equipped with a second countable Hausdorff topology, and let $X : \Omega \longrightarrow \mathbb{R}$ be a random variable. Think of X as representing some measurement apparatus. If $\mathcal{B}$ is the Borel σ-algebra on Ω, then the fact that singletons are in $\mathcal{B}$ represents the assumption that an outcome can be observed with arbitrary precision. In practice, however, it may be impossible to observe an outcome $X(\omega)$ with arbitrary precision. For instance, the best we can do may be to locate $X(\omega)$ in some element of a partition of $\mathbb{R}$ into intervals of length ϵ. Let $\mathcal{B}'$ be the σ-algebra generated by those intervals. Then $\mathcal{G} = X^{-1}\mathcal{B}'$ is strictly contained in $\mathcal{B}$. The conditional expectation $\mathbf{E}(X|\mathcal{G})$ is in a sense X viewed from the perspective of $\mathcal{G}$; the values of X are averaged over an element

*I am indebted to Samson Abramsky, Alan Bundy, Jaap van der Does, Rosalie Iemhoff, Alekos Kechris, Daniele Turi and Marco Vervoort for helpful comments. Support by the Netherlands Organisation for Scientific research (NWO) under grant PGS 22 262 is gratefully acknowledged. This paper was begun in Amsterdam and finished at HCRC in Edinburgh. I thank Keith Stenning for his friendship and hospitality, and the EPSRC for financial support.

in the generated partition on Ω. This is brought out by the defining condition of $\mathbf{E}$: $\mathbf{E}(X|\mathcal{G})$ is a $\mathcal{G}$-measurable random variable such that for A in $\mathcal{G}$,

$$\int_A \mathbf{E}(X|\mathcal{G})dP \;=\; \int_A XdP.$$

Incomplete information is not always a negative qualification. In some cases one wants to expressly *reduce* the information present in a signal in order to make that information more useful. This happens for instance when noise is suppressed in an audio signal. Moving to a coarser granularity can thus be beneficial as well. For our present purposes it is of interest that this process of filtering is also modelled by means of conditional expectation.

These two aspects of changes in granularity occur as well in qualitative reasoning with incomplete information. In an interesting programmatic paper, Hobbs [8] emphasises the positive aspects of switching to coarser granularities:

> Our ability to conceptualize the world at different granularities and to switch among these granularities is fundamental to our intelligence and flexibility. It enables us to map the complexities of the world around us into simple theories that are computationally tractable to reason in.

A few years earlier, Marr, in his book *Vision*, emphasised the same point while sketching a program for semantics (Marr [12, p. 357–8])

> I expect that at the heart of our understanding of intelligence will lie at least one and probably several important principles abour organizing and representing knowledge that in some sense capture what is important about our intellectual capabilities. . . . The perception of an event or an object must include the simultaneous computation of several different descriptions of it that capture different aspects of the use, purpose or circumstances of the event or object. . . . The various descriptions include coarse versions as well as fine ones. These coarse descriptions are a vital link in choosing appropriate overall scenarios . . . and in correctly establishing the roles played by the objects and actions that caused those scenarios to be chosen.

In fact, natural language even has adverbial expressions, such as ‘actually’, ‘really’, which indicate such shifts. Here is an example due to Asher and Vieu

> The point of this pencil is actually an irregular surface with several peaks.

It would be useful to have a formal apparatus capturing the process of simultaneously viewing the world at different scales; the claim of this paper is that one way to approach this problem is via a suitable notion of generalised quantification.

We are thus motivated by the following analogy. Moving to a larger grain size is analogous to taking conditional expectation, and it will be shown that this process can be captured by a quantifier which shares many formal properties with conditional expectation. We will investigate several ways of defining grain size with their associated conditioning structures and quantifiers. An important (and largely open) question that arises here is: when and why can we be confident that a result (for instance a plan) obtained in a coarse world also applies to the real, or at least a finer world? This leads to the study of a type of preservation theorems which are analogous to the martingale convergence theorems of probability theory. The ideas sketched above have been applied to the development of a logic of visual perception in a series of papers [21], [19] and [20]; the present article intends to provide some theoretical complements.

To further clarify the point of departure of the present article it is useful to contrast it with the aims of Keisler's probability logic (cf. Keisler [10]). Keisler presents an infinitary axiomatisation of such concepts as probability measure, integral and conditional expectation, and proves the completeness of the axiomatisation with respect to the usual probability spaces. By contrast, we look at what conditional expectation is used for in probability theory, try to identify situations of a more qualitative nature where approximately the same use is called for, and define qualitative analogues of conditional expectation which can perform these functions. There appears to be a certain logical and philosophical interest in this procedure, because it touches upon the old problem of the relation between logic and probability. It is our contention that logic and probability are alike in that their fundamental concepts, quantification and conditional expectation, are basically the same. Indeed, it has been observed several times that algebraically speaking the existential quantifier and conditional expectation are in the same class of operators on a Boolean algebra, namely the hemimorphisms[1] α satisfying the *averaging identity* $\alpha(p \wedge \alpha q) = \alpha p \wedge \alpha q$. (See, for example, Wright [24],[25] and the references given therein.) Furthermore Ellerman and Rota [4] showed that existential quantification $\exists x$ *is* conditional expectation with respect to the algebra determined by projection along x and a suitably generalised notion of measure, which is such that even on uncountable models, every subset of the domain has positive measure.

What remains to be done is to further exploit the analogy and to generalise existential quantification in such a way that the very special algebra used to define $\exists x$ (namely, the algebra of the sets of assignments obtained by projection along x) can be replaced by arbitrary Boolean algebras, or even more general structures. In view of the analogy with conditional expectation, the new quantifiers will be called *conditional* quantifiers. It will be very useful to think of these quantifiers as *resource-bounded* quantifiers: the resource is some algebraic structure which determines the granularity of the available information with respect to which one quantifies. For example, put in these terms quantification $\exists x \varphi$ means that we have no information about x, but full information about the other variables.

The remainder of the paper is organised as follows. We shall first, in sec-

[1] A hemimorphism on a lattice L with $\mathbf{1}$ is a mapping $\alpha : L \longrightarrow L$ satisfying $\alpha(a \vee b) = \alpha(a) \vee \alpha(b)$ and $\alpha\mathbf{1} = \mathbf{1}$.

tion 2, briefly rehearse theory and applications of conditional expectation in so far as relevant for our purposes. We then look at several examples of shifts of granularity, and show that they call for different types of resources (section 3). In section 4 the resources considered are Boolean algebras, and we try to retain as many of the properties of the existential quantifier as possible. Although we prove an infinitary completeness theorem here, the results are not very satisfactory, and in the next section (6) we consider, along with Boolean resources, a modified notion of quantifier, namely a quantifier $\exists$ which lacks the property $\varphi \leq \exists\varphi$. These quantifiers, which bear a stronger resemblance to conditional quantification then the existential quantifier itself, have a smoother theory, although at some points we need the continuum hypothesis to get things off the ground. In section 8 we consider resources which do not have the structure of a Boolean algebra, notably Lawvere's co-Heyting algebras (Lawvere [11]). These algebras naturally arise in situations where there is a distinction between positive and negative information. Again we supply an infinitary axiomatisation with a completeness theorem. Lastly, we return to our original motivation, conditional quantification as capturing shifts of granularity, which leads us to define a logical analogue of the *martingales* familiar from probability theory. We show that there exists a natural relation between nonmonotonic reasoning and martingale convergence, and we close with a conjecture on the precise form of martingale convergence in the logical setting.

2 Conditional expectation

In this section we give a very brief introduction to the fundamental properties of conditional expectation. The reader wishing to know more is advised to consult any introductory treatise on measure theoretic probability, a beautiful specimen of which is Williams' *Probability with martingales* [23]. We begin with a simple example. Suppose we have a variable X on a sample space Ω which takes values 0 and 1 both with probability $\frac{1}{2}$. Let A_i be the subset of Ω on which X takes value i, then $P(A_i) = \frac{1}{2}$. Suppose furthermore that we cannot measure X directly, but can only measure $X+Y$, where Y is some small perturbation. Let $B_i \subseteq A_i$ be such that $P(B_0) = \epsilon \neq P(B_1) = \delta$, where $\epsilon, \delta \ll \frac{1}{2}$; Y takes value $\frac{1}{2}$ on B_0, value $-\frac{1}{2}$ on B_1 and 0 elsewhere. Then $X+Y$ takes value 0 on $A_0 - B_0$, value 1 on $A_1 - B_1$ and value $\frac{1}{2}$ on $B_0 \cup B_1$. The smallest Boolean algebra $\mathcal{B}$ which contains $A_0 - B_0$, $A_1 - B_1$ and $B_0 \cup B_1$ does not contain A_0 and A_1. The algebra $\mathcal{B}$ represents the situation that we have incomplete information about X, and precisely codes the kind of information that we do have available about X, namely $X+Y$, since $\mathcal{B}$ is the smallest algebra with respect to which $X+Y$ is measurable. If we can measure only $X+Y$, not X, this implies that all information about X is represented by the function $\mathbf{E}(X|\mathcal{B})$, which takes value 0 on $A_0 - B_0$, value $\frac{1}{2} - \delta$ on $A_1 - B_1$, and value δ on $B_0 \cup B_1$. It follows that we have no information about the event $X > 0$ (since $A_0 \notin \mathcal{B}$, but our best estimate for the probability of this event is given by $P(\mathbf{E}(X|\mathcal{B}) > 0) = \frac{1}{2} + \epsilon$.

The next example is more realistic example, with continuous noise. Suppose we have a random variable X on a sample space Ω (measurable with respect to a σ-algebra $\mathcal{B}$) which we want to measure; for the sake of definiteness, assume X is distributed as the Gaussian $N(0, \sigma^2)$. Due to noise, we can only observe $X + c.\xi$, where ξ is independent of X, ξ is distributed as the Gaussian $N(0, 1)$ and $c \in I\!R$. $X + c.\xi$ can be taken to be measurable with respect to a σ-algebra $\mathcal{G}$ which is not the same as $\mathcal{B}$. For example, the event $\{X < 0\} \in \mathcal{B}$ is not in $\mathcal{G}$ if $\mathcal{G}$ is the smallest σ-algebra with respect to which $X + c.\xi$ is measurable. Hence we can only determine properties of X as filtered through $\mathcal{G}$; this is represented by the conditional expectation $\mathbf{E}(X|\mathcal{G})$. It is not possible to determine precisely which sample point ω has been chosen. We can only ask whether $\omega \in A$ for $A \in \mathcal{G}$. Then the expected value of X given such information, namely $\int_A X dP$, should equal $\int_A \mathbf{E}(X|\mathcal{G})dP$; from the point of view of $\mathcal{G}$, no other questions about X can be answered. This leads to the following

Definition 1 *Let $\mathcal{B}$ be a σ-algebra on Ω. A conditional expectation of a $\mathcal{B}$-measurable random variable $X : \Omega \longrightarrow I\!R$ with respect to the sub-σ-algebra $\mathcal{G}$ and the probability measure P is a function $\mathbf{E}(X|\mathcal{G}) : \Omega \longrightarrow I\!R$ satisfying*

$$\int_A \mathbf{E}(X|\mathcal{G})dP \;=\; \int_A X dP.$$

The Radon-Nikodym theorem is used to show that a function $\mathbf{E}(X|\mathcal{G})(\omega)$ with these properties exists. However, it is not unique, in the sense that for given X two versions of $\mathbf{E}(X|\mathcal{G})$ only agree almost everywhere[2]. In the following list of properties of $\mathbf{E}(X|\mathcal{G})$, the (in)equalities must therefore be understood as holding almost everywhere:

1. $\mathbf{E}(1_\emptyset|\mathcal{G}) = 0$, $\mathbf{E}(1_\Omega|\mathcal{G}) = 1$;
2. $X \subseteq Y$ implies $\mathbf{E}(1_X|\mathcal{G}) \leq \mathbf{E}(1_Y|\mathcal{G})$; and
3. $\mathbf{E}(X \cdot \mathbf{E}(Y|\mathcal{G})|\mathcal{G}) = \mathbf{E}(X|\mathcal{G}) \cdot \mathbf{E}(Y|\mathcal{G})$.

All three properties are strongly suggestive of quantifier properties, as has often been remarked in the literature; see, for example, Birkhoff's survey paper [3]. However, conditional expectation as constructed above by means of the Radon-Nikodym theorem is not yet quite analogous to quantification, since the construction is not uniform. For example, for each X, Y there is a nullset N such that off N, 3 holds; but there need not be a nullset which does the job for every X, Y. The required uniformity is given by the following definition. For ease of exposition, we formulate it in terms of conditional *probability* $\mathbf{P}(\cdot|\mathcal{G})$ determined by

$$\mathbf{P}(A|\mathcal{G}) \;=\; \mathbf{E}(1_A|\mathcal{G}).$$

Definition 2 *A* regular conditional probability *is a function* $\mathbf{P}(F|\mathcal{G})(\omega) : \mathcal{B} \times \Omega \longrightarrow [0, 1]$ *satisfying*

[2] The analogues of conditional expectation in section 6 suffer from this non-uniqueness as well.

1. *for each* $F \in \mathcal{B}$, $\mathbf{P}(F|\mathcal{G})(\cdot) = \mathbf{E}(1_F|\mathcal{G})$ *a.s.*
2. *for almost every* ω, *the map* $F \mapsto \mathbf{P}(F|\mathcal{G})(\omega)$ *is a probability measure on* $\mathcal{B}$.

The last condition ensures that the additivity and monotonicity properties hold in a uniform manner, in the sense that there exists a single nullset such that outside this nullset these properties hold. Regular conditional probabilities do not always exist, but we need not inquire into the conditions for existence here. It is important to observe that also in this case the function $\mathbf{P}(F \mid \mathcal{G})(\omega)$ is not completely determined by $\mathcal{G}$, and that essential reference is made to nullsets.

In probability one often considers families of conditional expectation operators. Since this suggests new ways of looking at quantifiers, we include a brief description. Let T be a *directed* set, i.e. a set partially ordered by a relation $\geq$ such that for $s, t \in T$, there exists $r \in T$ with $r \geq s, t$. A family of algebras $\{\mathcal{B}_s\}_{s \in T}$ is called a *net* (sometimes *filtration*) if $s \geq t$ implies $\mathcal{B}_s \supseteq \mathcal{B}_t$. The intuitive idea is that $\mathcal{B}_s$ contains all the information at hand at 'stage' s (which could be a timepoint, or a location in space *etc.*). A set of random variables $\{X_s\}_{s \in T}$ is called *adapted* (w.r.t. $\{\mathcal{B}_s\}_{s \in T}$) if each X_s is measurable with respect to $\mathcal{B}_s$. The set $\{X_s\}_{s \in T}$ is called a *martingale* if for each s, t such that $s \geq t$, $\mathbf{E}(X_s|\mathcal{B}_t) = X_t$. For our purposes the following fundamental theorem is of interest.

Theorem 1 *Let X be a random variable measurable with respect to* $\mathcal{B} = \bigcup_{s \in T} \mathcal{B}_s$.

1. *The family of random variables* $\{\mathbf{E}(X \mid \mathcal{B}_s)\}_{s \in T}$ *is adapted w.r.t.* $\{\mathcal{B}_s\}_{s \in T}$ *and is a martingale.*
2. $\lim_{s \in T} \mathbf{E}(X \mid \mathcal{B}_s) = \mathbf{E}(X \mid \mathcal{B})$ *a.s.. (In many interesting cases, but not always, one has in addition that* $X = \mathbf{E}(X \mid \mathcal{B})$ *a.s..)*

The preceding material suggests that the family of quantifiers considered by first order logic, $\{\exists y \mid y \text{ a variable}\}$, may be only one possibility among many interesting families of quantifiers. In section 9 we shall study one example in some detail, namely the logical analogue of a martingale. Another interesting example is furnished by the logical analogue of a Markov random field, a family of quantifiers that was implicitly used in the proof theory for generalised quantifiers presented in Alechina and van Lambalgen [1]. In slightly more formal detail: it can be shown that every filter quantifier is definable using a family of resource-bounded quantifiers which is the qualitative analogue of a Markov random field. This however is best left for a separate paper.

3 Examples: partiality and granularity

We proceed to give a few examples of conditional quantifiers, motivated by the idea of 'changes in granularity'. These examples play an important role in the logic of visual perception developed in [21] and [19], and the reader is encouraged to consult the papers cited for fuller information on their use. Here, we shall discuss the motivating examples mostly with reference to Hobbs' paper. We will return to vision in section 7, after we have developed the necessary technical apparatus.

3.1 Restricted signature

Let $\mathcal{L}$ be a first order language, $\mathcal{L}'$ a sublanguage of $\mathcal{L}$ containing only the 'relevant' predicates of $\mathcal{L}$. A concrete case is furnished by visual occlusion: if I am looking at the back of someone's head, a predicate such as 'mouth' is not applicable, hence not part of my model of the situation, or, equivalently, not part of the language with which I describe that situation. Given a model $\mathcal{M}$ ('the real world'), we may then define an equivalence relation E on $\mathcal{M}$ by $E(a,b)$ iff for all formulas $\varphi(x)$ in the language $\mathcal{L}'$, $\mathcal{M} \models \varphi(a) \iff \mathcal{M} \models \varphi(b)$.

Hobbs proposes to define a model $\mathcal{N}$ with domain the E-equivalence classes of elements of $\mathcal{M}$, such that (for monadic predicates) $\mathcal{N} \models A([a])$ iff $\mathcal{M} \models A(a)$. By the choice of E, this is welldefined. However, apart from the fact that this does not generalise to predicates of higher arity, the truth definition does not capture the intuitive motivation. The trouble lies in the direction from left to right: if $\mathcal{N} \models A([a])$ then $\mathcal{M} \models A(a)$. When planning a hike, a curve on the map (represented by a suitable equivalence class) be a trail from x to y *from the perspective of the map's scale*, without being an actual trail (e.g. it may have been overgrown in part). A more intuitive truth definition is obtained by putting $\mathcal{N} \models A([a])$ iff $\exists b(E(a,b) \& \mathcal{M} \models A(b))$. Since E is an equivalence relation this is again welldefined, but the definition now better captures the uncertain nature of the inference from $\mathcal{N}$ to $\mathcal{M}$.

In order to treat predicates of arbitrary arity, we can define an equivalence relation R on assignments by $R(f,g)$ iff for all formulas φ in the language $\mathcal{L}'$, $\mathcal{M} \models \varphi[f] \iff \mathcal{M} \models \varphi[g]$. R generates a quantifier[3] $\exists_R$ by

$$\mathcal{M} \models \exists_R[f] \iff \exists g(R(f,g) \ \& \ \mathcal{M} \models \varphi[g]).$$

If A is a predicate of $\mathcal{L}$, $\exists_R A$ represents the set of tuples of R-equivalence classes corresponding to tuples satisfying A. $\exists_R$ thus defines a model $\mathcal{N}$ which can be viewed as a coarsening of $\mathcal{M}$[4].

Let $\mathcal{B}$ be the algebra of $\mathcal{L}$-definable sets of assignments, and let $\mathcal{G}$ be the subalgebra of $\mathcal{B}$ determined by $\mathcal{L}'$. We will see later that $\exists_R$ is in a sense a quantifier conditional on $\mathcal{G}$. It will also become clear later that $\exists_R$ only has the right properties for a resource-bounded quantifier when $\mathcal{N}$ is finite, but at this stage $\exists_R$ is useful to fix ideas. The important point to remember is that the 'best estimate' of a predicate is represented by a quantifier. We could also use two quantifiers $\exists$ and $\forall$ corresponding to best upper and best lower estimate, but in the Boolean case $\exists$ and $\forall$ will be dual.

[3] Quantifiers of this type (which satisfy the $S5$ axioms) were introduced by Halmos [6].

[4] The reader may observe that Hobbs's truth definition actually corresponds to using the dual quantifier $\forall_R$ when interpreting predicates on $\mathcal{N}$.

3.2 Homomorphic image

The map $a \mapsto [a]$ constructed above is an example of a homomorphism:

Definition 3 *A mapping $h : \mathcal{M} \longrightarrow \mathcal{N}$ is a homomorphism if it is a surjective map such that $\mathcal{M} \models A(a_1, \ldots, a_n)$ implies $\mathcal{N} \models A(h(a_1), \ldots, h(a_n))$.*

Note that, in the previous example, both truth definitions for the predicates give rise to homomorphisms. Concrete examples of homomorphisms are furnished by projections of vector spaces onto spaces of lower dimension. For instance when planning a trip from x to y, the (two-dimensional[5]) road between x and y is reduced to a one-dimensional curve. For another example, we may consider object recognition: to identify some complex shape such as a human body (which requires a very large number of parameters for a complete specification), it is often sufficient to use a rough model, specifiable by few parameters (cf. Marr [12]). For more examples in this vein, some borrowed from Herskovits' semantics for spatial prepositions [7], see van Lambalgen [20].

In the obvious way h generates a projection π on the assignment spaces of $\mathcal{M}$ and $\mathcal{N}$, $\pi : \mathcal{F}_{\mathcal{M}} \longrightarrow \mathcal{F}_{\mathcal{N}}$. Define an equivalence relation R on assignments by $R(f, g)$ iff $\pi(f) = \pi(g)$; R defines a quantifier $\exists_R$ as above.

Let us now compare the model $\mathcal{N}$ and the quantifier $\exists_R$; both in some sense represent a coarsening of the model $\mathcal{M}$, and if h is a homomorphism we have $\mathcal{N} \models A(h(a))$ if $\mathcal{M} \models \exists_R A(a)$. For arbitrary formulas φ there is no such pleasant relation between truth on $\mathcal{N}$ and $\exists_R \varphi$, so we must ask which 'estimate' of φ is best: $\{f \mid \mathcal{N} \models \varphi[\pi(f)]\}$ or $\{f \mid \mathcal{M} \models \exists_R \varphi[f]\}$. (Observe that both sets are R-invariant, so they actually determine sets of equivalence classes, consistent with the intuition that individual assignments are not accessible.)

The answer has to do with the role of knowledge. Given that $\mathcal{N} \models A(h(a))$, we may be ignorant whether $\mathcal{M} \models A(a)$; but we may know for certain that $\mathcal{M} \models \forall x(A(x) \rightarrow \neg B(x))$. Suppose the model $\mathcal{N}$ is such that $A_{\mathcal{N}}$ and $B_{\mathcal{N}}$ overlap and $B_{\mathcal{M}} = \pi^{-1}(B_{\mathcal{N}})$; this is consistent with $h : \mathcal{M} \longrightarrow \mathcal{N}$ being a homomorphism. In this case $A_{\mathcal{N}}$ is evidently too rough as an estimate for $A_{\mathcal{M}}$, since we already know that the intersection $A_{\mathcal{N}} \cap B_{\mathcal{N}}$ cannot be part of $A_{\mathcal{M}}$. On the other hand, $\exists_R$ does take account of such knowledge, and can be shown to have the property that $\mathcal{M} \models \forall x(A(x) \rightarrow \neg B(x))$ implies $\mathcal{M} \models \forall x(\exists_R A(x) \rightarrow \neg B(x))$ (for this particular choice of B).

We now reformulate the property of being a 'best estimator' abstractly, borrowing a notation from probability theory. Let $\mathcal{G}'$ be the algebra of first order definable subsets of $\mathcal{F}_{\mathcal{N}}$, $\mathcal{G} = \{\pi^{-1}(C) \mid C \in \mathcal{G}'\}$. We would now like to have a quantifier $\exists(\cdot \mid \mathcal{G})$ which gives us '$\mathcal{M}$ seen from the point of view of $\mathcal{N}$'. Desirable properties are, for instance, monotonicity, distribution over disjunction, and the property that, if $C \in \mathcal{G}$ and, on $\mathcal{M}$, $\varphi \subseteq C$ (as sets of assignments), then $\exists(\varphi|\mathcal{G}) \subseteq C$.

Formally, the quantifier $\exists(\cdot|\mathcal{G})$ should be a map from sets of assignments to sets of assignments, with range in $\mathcal{G}$, satisfying properties such as monotonicity

[5]This is already an idealisation, i.e. a projection.

and distribution over $\vee$. At this stage it is not yet clear that this is at all possible, we just have the vague analogy with conditional expectation to guide us.

3.3 Combining information

In the preceding examples the resource was determined by a single model, and one might ask what the advantage is of the resource-bounded quantifier over the model. Although there is a definite advantage, for instance when the resource is only a proper subalgebra of the algebra determined by a model, it is when we want to combine information from several sources, that resource-bounded quantification really comes into its own. For example, when planning a trip, we combine information from a road atlas with information about road blocks. Both sources can be viewed as homomorphic images, say $\mathcal{N}$ and $\mathcal{K}$, of the real world $\mathcal{M}$, which lead to algebras $\mathcal{G}_{\mathcal{N}}$ and $\mathcal{G}_{\mathcal{K}}$ on $\mathcal{M}$. If $\mathcal{G}$ is the smallest algebra containing both $\mathcal{G}_{\mathcal{N}}$ and $\mathcal{G}_{\mathcal{K}}$, and if $\exists x\varphi(x,a,b)$ is the sentence 'there is a path from A to B', then the formula $\exists(\exists x\varphi(x,a,b)|\mathcal{G})$ should give us the best estimate of the truth of $\exists x\varphi(x,a,b)$ given the information in $\mathcal{N}$ and $\mathcal{K}$. It should be possible to construct a model from $\exists(\cdot|\mathcal{G})$ which represents our best knowledge, but the usefulness of $\exists(\cdot|\mathcal{G})$ resides in the possibility to combine information from a 'fixed' source (the road atlas) with temporary information.

3.4 Partial homomorphisms

A more realistic example of coarse graining is obtained when a only few elements of the domain are approximated, and the rest is not taken into consideration. This happens for instance in vision, when we combine the effects of the approximative nature of perception with the necessarily restricted perceptual field. Hobbs [8, p. 2] treats this case by throwing in the element EE (for 'everything else') in the approximating model, and mapping all elements not considered on EE. There exists another option, which will be seen to lead to a quantifier.

Definition 4 *A partial homomorphism $h : \mathcal{M} \hookrightarrow \mathcal{N}$ is a surjective homomorphism from a submodel $\mathcal{M}' \subseteq \mathcal{M}$ to $\mathcal{N}$.*

If h is a partial homomorphism, the projection $\pi : \mathcal{F}_{\mathcal{M}} \hookrightarrow \mathcal{F}_{\mathcal{N}}$, defined by $\pi(f)(x) = h(f(x))$ is likewise a partial function.

The quantifier $\exists$ associated to such a partial homomorphism is different from the $S5$-type quantifiers considered above; that is, it does satisfy the 'desirable properties' of section 3.2, but it will lack the properties $\varphi \leq \exists\varphi$ and $\exists\mathbf{1} = \mathbf{1}$. In case we have a partial homomorphism h which arises from restricting a *total* homomorphism h' to a domain D (say, the visual field), the quantifier $\exists$ can be made to satisfy $\exists\mathbf{1} = \mathbf{1}$ as well, which leads to the very useful concept of an *average*, introduced by Wright [25]. It will be studied in detail in section 6.

3.5 Non-Boolean resources

So far we took the resources to be Boolean algebras. In some of the applications of [21] and [19] it turned out be useful to have resources with weaker logical properties, in order to model various kinds of partiality. For example, in practice there will often be an assymmetry between a predicate A and its negation. Here is the pertinent definition:

Definition 5 *A co-Heyting algebra is a bounded distributive lattice L with a subtraction operation $\backslash : L \times L \longrightarrow L$ with the property*

$$x \backslash y \leq \Longleftrightarrow x \leq y \vee z.$$

Definition 6 *A complete co-Heyting algebra is a lattice in which arbitrary infima exist subject to the following distributive law*

$$x \vee \bigwedge_i y_i = \bigwedge_i (x \vee y_i).$$

A concrete example of a complete co-Heyting algebra is obtained when we take the lattice of sets of assignments on a model $\mathcal{N}$ defined by positive formulas, and close under $\bigwedge$. It will be seen in section 8 that this choice of resource changes the logic drastically.

As a last example, we consider an even weaker structure, also motivated by the idea of partiality. If we think of the quantifier $\exists$ applied to the predicate A as giving an estimate of A on the basis of restricted information, then we would like our estimate to be somehow informative. In the present set-up, $\exists A$ is always defined, but it may sometimes return the value $\mathbf{1}$. In treating Barwise's example of a non-inference due to partiality, 'Whitehead saw Russell. Russell winked. Therefore Whitehead saw Russell wink.', in [19], it was found useful to allow the possibility that $\exists A$ is undefined, when it would otherwise have returned the value $\mathbf{1}$. This leads to the following structure:

Definition 7 *A* pseudolattice *L is a partially ordered set in which meets and joins of finite non-empty sets exist. A pseudolattice L is a pseudo co-Heyting lattice if it is closed under arbitrary non-empty meets, such that the following distributive law holds:*

$$a \vee \bigwedge_{i \in I} b_i = \bigwedge_{i \in I} (a \vee b_i).$$

Note that in lattices in which arbitrary (i.e. also empty) finite meets and joins exists, top and bottom can be defined by $\mathbf{1} = \bigwedge \emptyset$ and $\mathbf{0} = \bigvee \emptyset$; but a pseudo co-Heyting lattice may lack top and/or bottom. Hence, if $\mathcal{H}$ is pseudo co-Heyting, and if $\exists$ is a quantifier conditional on $\mathcal{H}$ (in a sense to be made precise later), then $\exists A$ need not be defined, and $\exists \mathbf{0}$ can be non-empty.

It will be seen that, in order to construct the quantifiers corresponding to these examples, the conditioning algebra must be chosen carefully, in particular with regard to completeness properties; but we hope that the general idea is clear. In closing this section, we remark that common to all examples given is that we start from a model $\mathcal{N}$ and a 'filter' through which to view $\mathcal{N}$ (be it a Boolean algebra or a weaker structure); these two ingredients then define a conditional quantifier. It is also possible to proceed in the opposite direction: starting from a conditional quantifier, to define a class of models and resources which could have given rise to this quantifier. (We are really concerned with classes of models here, since as we have seen, a resource may filter out quite a lot of information about a model.) The technical development of this idea is rather lengthy however, so it will not be included here.

4 Conditioning on Boolean algebras

4.1 Preliminaries

Let $\mathcal{B}$ be a Boolean algebra of sets of assignments, $\mathcal{G}$ a subalgebra of $\mathcal{B}$. As a concrete case, let us refer back to section 3.2. Let $h : \mathcal{M} \longrightarrow \mathcal{N}$ be a homomorphism, then h generates a projection π on the assignment spaces of $\mathcal{M}$ and $\mathcal{N}$, $\pi : \mathcal{F}_{\mathcal{M}} \longrightarrow \mathcal{F}_{\mathcal{N}}$, by putting $\pi(f)(x) = h(f(x))$. Now put $\mathcal{G} = \{\pi^{-1}\{g \mid \mathcal{N} \models \psi[g]\} \mid \psi$ a first order formula$\}$. We want to construct quantifiers as mappings $\exists : \mathcal{B} \longrightarrow \mathcal{G}$ which somehow give the best estimate given the information present in $\mathcal{G}$; this generalises both first order logic and conditional expectation in probability theory. The analogy suggests that one should think of the conditioning algebra $\mathcal{G}$ as a *resource*, so that the generalised notion of quantification sought for is *resource bounded* quantification.

Above, we have met several examples of possible resources, determined by models which were in a sense coarser than the ground models. Hobbs' remark that change of granularity allows us to reason in computationally tractable theories, can be provisionally be translated into the requirement that the resource is specifiable in a 'simple' way. A prime example is of the course the resource derived from a finite model. Other possibilities suggest themselves, for instance the case where the resource is obtained by applying negation, countable conjunction and countable disjunction to predicates. For instance, if the domain of the ground model consists of the reals, then the resource, which is thus a Borel σ-algebra, can in general be said to be computationally simpler than the algebra of first order definable sets[6]. This case will be our main example below. Of course, this is not exactly real life computability, but it may give an idea of what could possibly be done here.

Considering the resources, several questions immediately arise

1. What are the possible resources? E.g. Boolean algebras or also other structures? Do these structures have to be complete in some sense?

[6] This of course depends on the signature; due to quantifier elimination it is not true with only addition and multiplication.

2. How does one define a quantifier given a resource and how does this choice, together with the choice of a resource, influence the logical properties of the resulting quantifier?

3. Is the quantifier uniquely determined by the resource?

In this section we suppose that the resource is given by a Boolean algebra $\mathcal{G}$ and we define the quantifier by means of a standard adjointness condition; as a consequence the logic of the conditional quantifier will be classical, and the quantifier will satisfy $S5$-like properties. (Other resources will be considered in section 8.) Since we think of the resource as the available information, this means that the quantifier should take values in $\mathcal{G}$ (compare the discussion of conditional expectation in section 2). It turns out that this is in general impossible without moving to some kind of completion of $\mathcal{G}$. If the answer to question 3 is affirmative, then the result of applying the quantifier resource-bounded by $\mathcal{G}$ to the formula φ is in a loose sense computable from $\mathcal{G}$ and φ. As will become clear in the course of the paper, this apparently can be achieved only when the resource satisfies a logic much weaker than classical logic.

Definition 8 *Let $\mathcal{M}$ be a model, $\mathcal{F}$ the set of assignments $f : VAR \longrightarrow |\mathcal{M}|$, $\mathcal{B}$ the algebra of first order definable sets of assignments, $\mathcal{G}$ a subalgebra of $\mathcal{B}$. A quantifier* conditional *on $\mathcal{G}$ is a function $\exists(\cdot|\mathcal{G}) : \mathcal{B} \longrightarrow \mathcal{G}$ satisfying the* adjointness condition[7]

$$\textit{for all } A \in \mathcal{G}, \textit{ for all } \varphi \in \mathcal{B},\ \varphi \leq A \textit{ iff } \exists(\varphi|\mathcal{G}) \leq A.$$

In this definition we have tacitly introduced the notational convention that definable sets of assignments are represented by the formulas defining them. It is easy to see that $\exists(\cdot|\mathcal{G})$, when it exists, is unique, and must satisfy

$$\exists(\varphi|\mathcal{G}) = \bigwedge\{A \in \mathcal{G} \mid \varphi \leq A\}.$$

The inclusion from left to right follows from the adjointness condition and we have equality because $\exists(\varphi|\mathcal{G}) \in \mathcal{G}$ by hypothesis. It is also clear that the existential quantifier $\exists x$ is a conditional quantifier: let $\mathcal{G}_x$ be the algebra of sets of assignments definable by formulas which do not contain x free, then $\exists x$ is the quantifier conditional on $\mathcal{G}_x$.
It is instructive to rewrite the Galois condition in order to emphasise the analogy between conditional quantification and conditional expectation.

Lemma 1 *A mapping $\exists(\cdot|\mathcal{G}) : \mathcal{B} \longrightarrow \mathcal{G}$ is a conditional quantifier if and only if the following conditions hold*

1. $\exists(\mathbf{0}|\mathcal{G}) = \mathbf{0}$, $\exists(\mathbf{1}|\mathcal{G}) = \mathbf{1}$

2. $\varphi \leq \exists(\varphi|\mathcal{G})$

[7] In the interest of stylistic variation, adjointness will sometimes be called 'Galois condition'.

3. $\varphi \leq \psi$ *implies* $\exists(\varphi|\mathcal{G}) \leq \exists(\psi|\mathcal{G})$

4. $\exists(\varphi \vee \psi|\mathcal{G}) = \exists(\varphi|\mathcal{G}) \vee \exists(\psi|\mathcal{G})$

5. *(a)* $\exists(\varphi \wedge \psi|\mathcal{G}) = \exists(\varphi|\mathcal{G}) \wedge \psi)$ *for* $\psi \in \mathcal{G}$

(b) $\exists(\varphi \wedge \exists(\psi|\mathcal{G}) \mid \mathcal{G}) = \exists(\varphi|\mathcal{G}) \wedge \exists(\psi|\mathcal{G})$.

The equivalent properties 5a and 5b will collectively be referred to as the Frobenius property.

Proof. Suppose $\exists(\cdot|\mathcal{G})$ is a conditional quantifier. All properties are trivial except 5a and 5b. The latter property follows from the former since $\exists(\cdot|\mathcal{G}) \in \mathcal{G}$. The inclusion from left to right in 5a follows from 2 (i.e. Galois from right to left), and Galois from left to right. The other direction is more interesting. We have to show $\exists(\varphi|\mathcal{G}) \wedge \psi \leq \exists(\varphi \wedge \psi|\mathcal{G})$; since $\mathcal{G}$ has an implication $\rightarrow$ which is also the implication of $\mathcal{B}$, this follows if we can show $\exists(\varphi|\mathcal{G}) \leq \psi \rightarrow \exists(\varphi \wedge \psi|\mathcal{G})$. Now the *r.h.s.* is in $\mathcal{G}$ so it suffices to show that $\varphi \leq \psi \rightarrow \exists(\varphi \wedge \psi|\mathcal{G})$, but this follows from 2 and the properties of $\rightarrow$.
Now suppose that the conditional quantifier satisfies properties 1 to 5b. The right to left direction of Galois follows from 2. Suppose $\varphi \leq A$, where $A \in \mathcal{G}$. By 3, $\exists(\varphi|\mathcal{G}) \leq \exists(A|\mathcal{G}) = A$, where the last equality holds because of 5a. □

Observe that all these properties correspond to properties of conditional expectation, except for 2. This is also the property wich makes it hard to achieve our desired goal, as will be seen in a moment, and in section 6 we shall consider quantifiers which lack it, the so called averages. Note also that these properties are true of the existential quantifier.

4.2 Existence of a conditional quantifier

In section 3 we gave a few examples of conditioning algebras; do there exist conditional quantifiers relative to these algebras? More formally, we start from a model consisting of a pair $\langle \mathcal{M}, \mathcal{G} \rangle$, where $\mathcal{G}$ is the Boolean algebra which represents the resource, and we are asking whether there exists an algebra $\mathcal{B}$, which extends the algebra of first order definable sets of assignments, and a conditional quantifier $\exists : \mathcal{B} \longrightarrow \mathcal{G}$. If this is so, we say that the model $\langle \mathcal{M}, \mathcal{G} \rangle$ can be extended with the conditional quantifier $\exists : \mathcal{B} \longrightarrow \mathcal{G}$. In this weak form of the existence question, we ask for the smallest extension of the algebra of first order definable sets of assignments on which a quantifier conditional on $\mathcal{G}$ can be defined. Much stronger versions of the question can be asked, for instance we could demand that $\mathcal{B}$ is closed under first order quantifiers, or under other conditional quantifiers; in sum, we could ask for the existence of a *family* of conditional quantifiers having certain interactions; we will return to this below. Even in this weak form, the answer is no in general, since it is dependent upon completeness properties of $\mathcal{G}$. However, we shall first consider a case where the answer is affirmative. Suppose that in example 3.2 we take the model $\mathcal{N}$ to be finite. This case is of importance in the logic of vision of [21] and [19], which mostly considers models that can be written as inverse limits of inverse systems of finite models (cf. section 9 below).

Lemma 2 *With the notation of section 3.2, if $\mathcal{N}$ is finite, then $\langle \mathcal{M}, \mathcal{G} \rangle$ can be extended with a conditional quantifier $\exists(\cdot|\mathcal{G})$.*

Proof. Suppose $\varphi(\overline{x}) \leq A \in \mathcal{G}$, then we may assume A is of the form $\pi^{-1}\psi(\overline{x}, \overline{y})$, ψ first order. Let u be a variable occurring in $\overline{x}$, and let $\overline{u}$ be obtained from $\overline{y}$ by substituting every variable with u. Then $\varphi(\overline{x}) \leq \pi^{-1}\psi(\overline{x}, \overline{u}) \leq \pi^{-1}\psi(\overline{x}, \overline{y})$. Since $\mathcal{N}$ is finite, it follows that $\bigwedge\{A \in \mathcal{G} \mid \varphi \leq A\}$ is actually a finite meet, which is thus in $\mathcal{G}$. □

The next example, motivated by section 3.1, gives a construction of a Boolean algebra $\mathcal{G}$ such that conditioning with respect to $\mathcal{G}$ cannot be defined. Suppose we have a model $\mathcal{M}$ with infinite signature $\{C, A_1, A_2, \ldots\}$, where all predicates are unary. Let $\mathcal{M}$ satisfy

1. $\forall n \forall y (C(y) \to A_n(y)$
2. $\forall n \forall y (A_{n+1}(y) \to A_n(y))$
3. $\forall n \exists y (\neg A_{n+1}(y) \wedge A_n(y))$
4. $\forall n \exists y (\neg C(y) \wedge A_n(y))$.

Let $\mathcal{G}$ be the Boolean algebra generated by the A_n, and let $\mathcal{M}$ be an ω-saturated modelfor 1–4. Under these conditions, $\exists(C(x,y)|\mathcal{G})$ would have to be equal to $\bigwedge_n A_n(y)$, but this is not in $\mathcal{G}$.

Hence $\mathcal{G}$ must have some completeness properties. We would like to follow probability theory in being able to specify the resource $\mathcal{G}$ independently; ideally, the meaning of the quantifier $\exists(\cdot|\mathcal{G})$ should then be completely determined by $\mathcal{G}$. In the above examples, especially in 3.2, $\mathcal{G}$ should represent 'the' information present in a model. However, there is no unique choice here; it all depends on the logic in which we describe the model. Even if we fix the logic to be classical, there are still various possibilities for making the logic infinitary. The easy way out is to allow all infima and suprema.

Suppose $\mathcal{G}_0$ is a Boolean algebra of first order definable sets of assignments. Define an equivalence relation $R_{\mathcal{G}_0}$ by

$$R_{\mathcal{G}_0}(f,g) \text{ iff } \forall \psi \in \mathcal{G}_0 (f \in \psi \Longrightarrow g \in \psi).$$

Let $\mathcal{G}$ be the algebra of $R_{\mathcal{G}_0}$-invariant sets of assignments, then $\mathcal{G}$ is effectively the power set on the set of equivalence classes $\{[f]_{R_{\mathcal{G}_0}} \mid f \in \mathcal{F}\}$. In this case we have

Lemma 3 *Let $\mathcal{M}$ be a model, $\mathcal{F}$ its set of assignments, $\mathcal{G}$ the $R_{\mathcal{G}_0}$-invariant algebra contained in $\wp(\mathcal{F})$. Then the conditional quantifier $\exists(\cdot|\mathcal{G}) : \wp(\mathcal{F}) \longrightarrow \mathcal{G}$ exists.*

Proof. Define $\exists(A|\mathcal{G}) = \{f \in \mathcal{F} \mid \exists g \in \mathcal{F}(R_{\mathcal{G}_0}(f,g) \wedge g \in A)\}$. $\exists(A|\mathcal{G}) = A$ for $A \in \mathcal{G} \subseteq \mathcal{B}$, hence the range of $\exists(\cdot|\mathcal{G})$ is $\mathcal{G}$. Since $R_{\mathcal{G}_0}$ is an equivalence relation, $\exists(\cdot|\mathcal{G})$ satisfies the Galois condition with respect to $\mathcal{G}$. □

In fact there is an intimate relation between conditional quantifiers and equivalence relations, as the following argument, essentially due to Halmos, shows. Strictly speaking, we cannot yet state this result since we have not introduced a formal language for conditional quantifiers. The only thing we need from the next section is, however, that for the resources $\mathcal{G}$ that are of interest, for each formula φ, $\exists(\varphi|\mathcal{G})$ is a formula with finitely many free variables.

We remind the reader that if a model $\mathcal{M}$ is ω_1-saturated, then for any set $\{\varphi_i \mid i \in \omega\}$ of formulas (in arbitrary free variables): if for each n, $\{\varphi_i \mid i \in n\}$ is satisfiable, then so is $\{\varphi_i \mid i \in \omega\}$. For obvious reasons this property will be referred to as compactness.

Theorem 2 *Suppose we have a model $\mathcal{M}$ on which is defined a conditional quantifier $\exists(\cdot|\mathcal{G})$, such that $\mathcal{M}$ is ω_1-saturated with respect to the language containing the conditional quantifier. Then there exists an equivalence relation R such that for all formulas φ in the language, $\exists(\varphi|\mathcal{G}) = \{f \in \mathcal{F} \mid \exists g \in \mathcal{F}(R(f,g) \wedge g \in \varphi)\}$.*

Proof. Given $\exists(\cdot|\mathcal{G})$, define R by $R(f,g)$ iff for all C in the range of $\exists(\cdot|\mathcal{G})$, $f \in C$ iff $g \in C$. Then $\exists g \in \mathcal{F}(R(f,g) \wedge g \in \varphi)$ implies $f \in \exists(\varphi|\mathcal{G})$ because $\exists(\varphi|\mathcal{G})$ is R-invariant. Conversely, if $\forall g(g \in \varphi \rightarrow \neg R(f,g))$, for each $g \in \varphi$ choose C_g in the range of $\exists(\cdot|\mathcal{G})$ such that $g \in C_g, f \notin C_g$. This can be done because the range of $\exists(\cdot|\mathcal{G})$ is an algebra (this will be formally proved in the proof of theorem 5). It follows that $\varphi \subseteq \bigcup_g C_g$ and by ω_1-saturation φ is also covered by a finite union, which is therefore in the range of $\exists(\cdot|\mathcal{G})$. By the Galois property, it follows that $f \notin \exists(\varphi|\mathcal{G})$. □

Recall that part of our motivation is that we want to construct the conditional quantifier in such a way that $\exists(\varphi|\mathcal{G})$ is in a sense computable from φ and $\mathcal{G}$. The case of interest to us is where $\mathcal{G}$ is the collection of $R_{\mathcal{G}_0}$-invariant sets of assignments, for a Boolean algebra $\mathcal{G}_0$ of first order formulas.

Theorem 3 *Let the model $\mathcal{M}$ be ω_1-saturated, $\mathcal{G}$, $\mathcal{G}_0$ as above. Then for first order φ, $\exists(\varphi|\mathcal{G}) = \bigwedge\{\psi \in \mathcal{G}_0 \mid \varphi \leq \psi\}$.*

Proof. Let $[f]$ denote the $R_{\mathcal{G}_0}$ equivalence class of an assignment. For arbitrary formulas φ we then have

$$\exists(\varphi|\mathcal{G}) = \bigcup\{[f] \mid \varphi \cap [f] \neq \emptyset\}.$$

The inclusion from left to right follows because the r.h.s. is an invariant set which covers φ, hence by Galois also covers $\exists(\varphi|\mathcal{G})$. Conversely, if $g \notin \exists(\varphi|\mathcal{G})$, then since $\exists(\varphi|\mathcal{G})$ is invariant, $[g] \cap \exists(\varphi|\mathcal{G}) = \emptyset$, and it follows that $[g] \cap \varphi = \emptyset$. By construction, each $[f]$ is of the form $\bigcap\{C \in \mathcal{G}_0 \mid f \in C\}$. Now let φ be first order and suppose there exists an assignment g such that for all $\psi \in \mathcal{G}_0$

such that $\varphi \leq \psi$, $g \in \psi$, but $[g] \cap \varphi = \emptyset$. Since $[g] = \bigcap\{\theta \in \mathcal{G}_0 \mid g \in \theta\}$, by ω_1-saturation there exists $\theta \in \mathcal{G}_0$ such that $g \in \theta$ and $\varphi \leq \neg\theta$. Since $\neg\theta \in \mathcal{G}_0$ we have $g \in \neg\theta$, a contradiction. □

Hence, although $\mathcal{G}$ is very large (for example, it strictly includes the σ-algebra generated by $\mathcal{G}_0$), in order to 'compute' $\exists(\varphi|\mathcal{G})$ we need only a countable set of elements of lesser complexity than $\exists(\varphi|\mathcal{G})$ itself. Furthermore, when working with a single conditional quantifier, this is basically all we need. In the following, $\overline{(\neg)\theta}$ stands for a conjunction of a finite but arbitrary number of formulas, some of which may be negated.

Lemma 4 *Every formula in the language obtained by adding a single conditional quantifier $\exists(\cdot|\mathcal{G})$ to a first order language without $\forall$ or $\exists$ is equivalent to a formula of the form: a disjunction of conjuncts $\overline{\psi} \wedge \overline{(\neg)\exists(\varphi|\mathcal{G})}$.*

PROOF OF LEMMA 4 If θ is a formula of the language, write θ in disjunctive normal form. $\exists(\cdot|\mathcal{G})$ distributes over disjunction, hence we have to show that each conjunct of θ is of the required form.
By lemma 1, part 5b, every quantifier $\exists(\cdot|\mathcal{G})$ occurring within the scope of another such quantifier can be pulled out. □

Speaking topologically, we have the following situation. Let the first order formulas determine a basis for a totally disconnected topology τ on the set $\mathcal{F}$ of assignments of an ω_1-saturated model $\mathcal{M}$. τ is compact, and if we identify assignments not distinguished by a formula, also Hausdorff. It follows that $(\mathcal{F}, \tau)$ is isomorphic to a standard Borel space. We have just shown that the set of formulas constructed from a single conditional quantifier determines sets of assignments which are Boolean combinations of open and closed sets. Sets are becoming more complicated when we add a second conditional quantifier; e.g. $\exists(\neg\exists(\varphi|\mathcal{G})|\mathcal{H})$ will in general be F_σ. This suggests that, in order to accomodate families of conditional quantifiers, the domain of each of them should be the Borel σ-algebra generated by τ. However, in view of the form of the definition of conditional quantifiers, one may suspect that adding one more conditional quantifier may then lead to analytic, non-Borel sets of the form $\exists(\neg\exists(\neg\exists(\varphi|\mathcal{G})|\mathcal{H})|\mathcal{K})$. In general one has the following result, kindly communicated to me by A. Kechris:

Lemma 5 *Let X be a standard Borel space, and let $E \subseteq X \times X$ be the Borel equivalence relation determined by a countable collection $\mathcal{G}_0$ of Borel sets, via*

$$Eab \iff \forall C \in \mathcal{G}_0(a \in C \leftrightarrow b \in C).$$

Then the following are equivalent

1. *for some Borel set $A \subseteq X$, the set $\{a \in A \mid \exists b(Eab \;\&\; b \in A)\}$ is complete analytic*
2. *E has uncountably many uncountable classes.*

It is easy to cook up an example along the lines of 3.1 which satisfies the second condition. It follows that we cannot simply take the domain of a conditional quantifier to be the Borel σ-algebra generated by τ. This is somewhat out of keeping with our general aim of using a logic which is as simple as possible to describe the resource. Below we shall tackle this problem in two ways. In the next section, we show that at least as far as completeness is concerned, the complexity of the denotation of conditionally quantified formulas can be kept small; this is achieved by resorting to a type of nonstandard model first introduced by Harvey Friedman, the so called Borel models. We then move on to slightly change the logical properties of the conditional quantifier, dropping the property that $\varphi \leq \exists(\varphi|\mathcal{G})$, and show that it has the effect of lowering complexity of denotation. In this case the Borel σ-algebra suffices as a domain of the quantifier. Conditional quantifiers will make a further appearance in section 9, where they will be used to provide simple examples of interesting families of quantifiers, martingales.

5 The logic of conditional quantifiers

In the literature on algebraic logic, what we call conditional quantifiers have been studied mainly by means of representation theory: given a mapping $\exists : \mathcal{B} \longrightarrow \mathcal{G}$, satisfying certain properties, where $\mathcal{G}$ is a subalgebra of the Boolean algebra $\mathcal{B}$, find out what corresponds to $\exists$ on the Stone spaces associated to $\mathcal{B}$ and $\mathcal{G}$. For our purposes this is not quite sufficient. Most importantly, representation theory does not tell us for which $\mathcal{G}$ the desired mappings $\exists$ exist. Hence, *given* an abstract quantifier, we can find a representation; but this gives us no clue how to construct such quantifiers *as conditional on a given algebra.* In this respect, the Halmos-Wright theory of quantifiers differs from its probababilistic counterpart. A related point is that the algebraic theory in itself gives us little idea how to construct a semantics for these quantifiers, in particular it is unclear how they interact with the ordinary quantifiers. We shall thus proceed in a traditional manner and develop logics for conditional quantifiers, with proof rules, semantics and completeness results relating the two.

Even so, there are several ways to do this. Staying close to probability theory would suggest adding conditional quantifiers to $L_{\omega_1\omega}$ or some fragment thereof, as in Keisler's probability logics [10]. The other option is to add conditional quantifiers to first order logic, even though the resource $\mathcal{G}$ need not itself be specifiable as a set of first order formulas. We shall choose the latter option, which is roughly analogous to the addition of generalised quantifiers to first order logic.

5.1 Syntax

Definition 9 *We assume given a supply of names $\mathcal{B}, \mathcal{G}, \mathcal{G}_s, \ldots$ for Boolean algebras. For any such name $\mathcal{G}$, $\exists(\cdot \mid \mathcal{G})$ denotes a conditional quantifier. The language $\mathcal{L}(CQ(\mathcal{G}))$ consists of a first order language together with a single* *quantifier symbol* $\exists(\cdot \mid \mathcal{G})$ *with the formation rule*

If φ *is a formula, then so is* $\exists(\varphi \mid \mathcal{G})$.

If we want to consider a language with many conditional quantifiers $\{\mathcal{G}_s \mid s \in T\}$*, we write* $\mathcal{L}(CQ(\{\mathcal{G}_s \mid s \in T\}))$.

We now encounter a technical subtlety. Normally a quantifier binds a variable, or perhaps several variables, but what are the free variables of $\exists(\varphi|\mathcal{G})$? The Galois condition forces $\exists(\varphi|\mathcal{G})$ to mean $\bigwedge\{\psi \in \mathcal{G} \mid \varphi \leq \psi\}$, from which it follows that the free variables of $\exists(\varphi|\mathcal{G})$ are all those occurring in some $\psi \in \mathcal{G}$ with $\varphi \leq \psi$. What those variables are very much depends on the interpretation of $\mathcal{G}$ in a model. Syntactically we can thus say only that the set of free variables of $\exists(\varphi|\mathcal{G})$ coincide with the set of free variables of the language. Below, we shall put a reasonable restriction on $\mathcal{G}$ which allows us to identify the free variables of $\exists(\varphi|\mathcal{G})$ with those of φ. Also, as we shall see later the effect of a conditional quantifier is best described as 'partially binding' variables; see section 9.

5.2 Semantics

Definition 10 *A* conditional quantification model *for* $\mathcal{L}(CQ(\mathcal{H}))$ *is a structure* $\langle \mathcal{M}, \mathcal{G}, \mathcal{G}' \rangle$*, where* $\mathcal{G}$ *is a Boolean algebra of sets of assignments and* $\mathcal{G}'$ *is a subalgebra of* $\mathcal{G}$*. The interpretation of* $\exists(\varphi|\mathcal{H})$ *on* $\langle \mathcal{M}, \mathcal{G}, \mathcal{G}' \rangle$ *is given by*

1. *The name* $\mathcal{H}$ *is interpreted by the algebra* $\mathcal{G}$.
2. $\langle \mathcal{M}, \mathcal{G}, \mathcal{G}' \rangle \models \exists(\varphi|\mathcal{H})[f]$ *iff for all* $A \in \mathcal{G}'$, $\varphi \leq A$ *implies* $f \in A$.

The quantifier $\exists(\varphi|\mathcal{H})$ *thus interpreted on the structure* $\langle \mathcal{M}, \mathcal{G}, \mathcal{G}' \rangle$ *is called* adapted *if the range of* $\exists(\cdot \mid \mathcal{H})$ *equals* $\mathcal{G}'$*. Henceforth we do not distinguish between names for algebras and their interpretations, and we shall take the second argument of* $\exists(\cdot|\cdot)$ *to be the algebra itself. We shall also sometimes suppress reference to the algebras when considering a structure.*

Similarly, a model for $\mathcal{L}(CQ(\{\mathcal{G}_s \mid s \in T\}))$ *is of the form* $\langle \mathcal{M}, \{\mathcal{G}_s \mid s \in T\}, \{\mathcal{G}'_s \mid s \in T\} \rangle$*, where each* $\mathcal{G}'_s \subseteq \mathcal{G}_s$.

Intuitively, $\mathcal{G}$ represents the resource, and $\mathcal{G}'$ represents the range of the conditional quantifier, which could be smaller, for instance when the language is countable and $\mathcal{G}$ is uncountable. Adaptedness is the analogue of measurability in our context. The distinction between $\mathcal{G}$ and $\mathcal{G}'$ is necessitated by our decision to add conditional quantifiers to first order logic; this implies that the domain of $\exists(\cdot|\mathcal{G})$ in general will not contain all of $\mathcal{G}$, and as a consequence condition 2 can only be formulated for the range of $\exists(\cdot|\mathcal{G})$, not for $\mathcal{G}$ itself. This is so because condition 2 corresponds directly to the Galois condition, formulated as a proof rule; and since in the proof rules only formulas of the language can occur, a proof of the completeness theorem will usually say little about sets in $\mathcal{G} - \mathcal{G}'$. The price we have to pay for this is that the truth condition 2 is not inductive, since to verify that $\exists(\varphi|\mathcal{G})$ satisfies it, we have to know $\mathcal{G}'$, hence $\exists(\varphi|\mathcal{G})$ itself. However, when the structure upon which the quantifier conditions is not a Boolean algebra but a co-Heyting algebra, the situation is much better; cf. section 8.

5.3 Formal system and completeness

As we have seen, without restrictions on $\mathcal{G}$, a formula of the form $\exists(\cdot|\mathcal{G})$ may contain infinitely many variables. We first briefly study a fragment of $\mathcal{L}(CQ)\mathcal{G}))$ where the first order quantifiers $\forall, \exists$ are not allowed to have the conditional quantifier $\exists(\cdot|\mathcal{G})$ within their scope. This fragment will be called $\mathcal{L}(CQ)\mathcal{G}))^-$. The fragment is axiomatised by a sequent calculus $\mathcal{A}(CQ(\mathcal{G}))^-$ consisting of the ordinary structural rules and rules for the connectives, the rules for $\forall, \exists$ with the proviso that they are not applicable to formulas containing a conditional quantifier, together with the following rules for conditional quantifiers

$$(lCQ^-)\ \frac{\Gamma, \varphi \Longrightarrow \Delta}{\Gamma, \exists(\varphi|\mathcal{G}) \Longrightarrow \Delta}$$

where formulas in Γ, Δ are of the form $\exists(\theta|\mathcal{G})$;

and

$$(rCQ^-)\ \frac{\Gamma \Longrightarrow \varphi, \Delta}{\Gamma \Longrightarrow \exists(\varphi|\mathcal{G}), \Delta}$$

(without restrictions on Γ, Δ).

Lemma 6 *$\mathcal{A}(CQ(\mathcal{G}))^-$ is sound with respect to adapted conditional quantification models.*

Proof. We check lCQ. Suppose $\mathcal{M} \models \Gamma, \exists(\varphi|\mathcal{G})[f]$, where Γ is a set of conditionally quantified formulas. The interpretation of $\exists(\varphi|\mathcal{G})$ entails that for all $A \in \mathcal{G}'$ such that $\varphi \leq A$, $f \in A$. The premiss of the rule gives $\varphi \leq \bigwedge \Gamma \to \bigvee \Delta$, and since the model is adapted and $\mathcal{G}'$ is a Boolean algebra, the *r.h.s.* is in $\mathcal{G}'$. It follows that $f \in \bigwedge \Gamma \to \bigvee \Delta$, whence $f \in \bigvee \Delta$. The soundness of rCQ is trivial. □

Theorem 4 *$\mathcal{A}(CQ(\mathcal{G}))^-$ is complete, i.e. every set Γ of sentences in $\mathcal{L}(CQ)\mathcal{G}))^-$ consistent with $\mathcal{A}(CQ(\mathcal{G}))^-$ has an adapted conditional quantification model.*

The proof is based on the idea that each formula of the form $\exists(\varphi|\mathcal{G})$ can be treated as a separate predicate of infinite arity, together with the normal form lemma 4, which we restate for convenience:

Lemma 7 *Every $\mathcal{L}(CQ)\mathcal{G}))^-$ - formula is equivalent to a formula of the form: a disjunction of conjuncts $\overline{\psi} \wedge \overline{(\neg)\exists(\varphi|\mathcal{G})}$.*

The proof of theorem 4 is then completed by a Henkin construction. □

However, the lack of ordinary quantifier rules somewhat complicates the usual Henkin construction, so we will move to the more interesting system where ordinary quantifiers may have scope over conditional quantifiers. In order to get a result for the full language $\mathcal{L}(CQ)\mathcal{G}))$ we have to restrict the algebras under consideration to the ones that are analogous to those considered in examples 3.1 and 3.2. The distinguishing feature of these algebras is given by the following

Definition 11 *An algebra $\mathcal{G}$ of sets of assignments is* homogeneous *if it is closed under substitution, where a substitution is an operation $S_r : \mathcal{G} \longrightarrow \mathcal{G}$, determined by a function $r : \omega \longrightarrow \omega$, satisfying $S_r(A) = \{f \mid \exists g \in A\ (g(x_i) = f(x_{r(i)}))\}$. A homogeneous conditional quantification model is a conditional quantification model $\langle \mathcal{M}, \mathcal{G}, \mathcal{G}' \rangle$ where both $\mathcal{G}$ and $\mathcal{G}'$ are homogeneous.*

The definition is as convoluted as it is because of the following subtlety: if we substitute x_1 for x_2 in $\theta(x_1, x_2)$, we want the result to be two-dimensional, not one-dimensional. The advantage of homogeneous conditioning algebras $\mathcal{G}$ is that we may intuitively think of a formula $\exists(\varphi(\overline{x})|\mathcal{G})$, as having the free variables $\overline{x}$, for if there were $\theta \in \mathcal{G}, \varphi \leq \theta$ containing additional free variables they can be replaced by variables from $\overline{x}$. We may therefore write $\exists(\varphi(\overline{x})|\mathcal{G})$ as $\exists(\varphi|\mathcal{G})(\overline{x})$. Furthermore the result of substituting a variable z for y in $\exists(\varphi(y, \overline{x})|\mathcal{G})$ is now simply $\exists(\varphi(z, \overline{x})|\mathcal{G})$. Formally,

Lemma 8 *Let $\langle \mathcal{M}, \mathcal{G}, \mathcal{G}' \rangle$ be a homogeneous conditional quantification model. Then if x does not occur free in φ, $\langle \mathcal{M}, \mathcal{G}, \mathcal{G}' \rangle \models \varphi[f], g =_x f$ implies $\langle \mathcal{M}, \mathcal{G} \rangle \models \varphi[g]$.*

Proof. It suffices to prove the statement of the lemma for φ of the form $\exists(\psi(\overline{x})|\mathcal{G})$. Suppose $y \notin \overline{x}, f \in \exists(\psi(\overline{x})|\mathcal{G}),\ g =_y f$. Choose $B \in \mathcal{G}, \psi(\overline{x}) \leq B$, then $f \in B$. Let $r : VAR \longrightarrow \overline{x}$ be a mapping which is the identity on $\overline{x}$, and which maps the remaining free variables of B into $\overline{x}$. Then $\psi(\overline{x}) \leq S_r(B) \leq B$ and $S_r(B) \in \mathcal{G}$. It follows that $f \in S_r(B)$, whence $g \in S_r(B)$ and $g \in B$. $\square$

We now introduce the sequent calculus $\mathcal{A}(CQ(\mathcal{G}))$ appropriate to this situation. The previous rule (lCQ^-) is replaced by

$$(lCQ)\quad \frac{\Gamma, \varphi(\overline{x}) \Longrightarrow \Delta}{\Gamma, \exists(\varphi|\mathcal{G})(\overline{x}) \Longrightarrow \Delta}$$

where formulas in Δ are of the form $\exists(\theta|\mathcal{G}), \neg\exists(\theta|\mathcal{G})$, and where Γ does not contain $\overline{x}$ free.
(rCQ) is identical to (rCQ^-), and we allow the rules for $\exists, \forall$ to be applied to formulas containing a conditional quantifier.

It is instructive to compare (lCQ) to the $l\exists$ rule. To introduce $\exists x$ we require that neither Γ nor Δ contains x free, i.e. that both Γ and Δ contain only $\exists(\cdot|\mathcal{G}_x)$-formulas (where $\mathcal{G}_x$ is the algebra of formulas not containing x free). Hence $l\exists$ is an instance of (lCQ^-). The stronger rule (lCQ) makes clear that x actually plays a dual role: roughly speaking, introduction of $\forall x$ on $\varphi \to \bigvee \Delta$, and introduction of $\exists x$ on φ. These roles must be kept separate for general conditional quantifiers. Another difference between (lCQ^-) and (lCQ) is that in the later rule we do not allow formulas of the form $\exists(\theta|\mathcal{G})(\overline{x})$ to occur on Γ, but move them (with negations added) all to Δ. The reason for this technical; we later need an infinitary variant of (lCQ) which is only sound with the new stipulation.

As a consequence of lemma 8 we obtain

Lemma 9 $\mathcal{A}(CQ(\mathcal{G}))$ *is sound for adapted homogeneous conditional quantification models.*

PROOF. Lemma 8 guarantees that the unrestricted $l\exists$ and $r\forall$ are sound. Now consider (lCQ). Suppose $\mathcal{M} \models \Gamma, \exists(\varphi|\mathcal{G})(\overline{x})[f]$. Since $\mathcal{M}$ is adapted, we know that for all ψ in the range of the conditional quantifier: if $\mathcal{M} \models \forall\overline{x}(\varphi(\overline{x}) \rightarrow \psi)[f]$, then $\mathcal{M} \models \psi[f]$. The restriction on$\overline{x}$ now ensures that $\mathcal{M} \models \forall\overline{x}(\varphi(\overline{x}) \rightarrow \bigvee\Delta)[f]$. □

We assume that the role of adaptedness has been made sufficiently clear; henceforth we shall take all models of interest to be adapted.

Theorem 5 $\mathcal{A}(CQ(\mathcal{G}))$ *is complete for homogeneous conditional quantification models.*

PROOF. For each formula $\exists(\varphi(\overline{x})|\mathcal{G})$ of $\mathcal{L}(CQ(\mathcal{G}))$ introduce a predicate, also written $\exists(\varphi(\overline{x})|\mathcal{G})$, with free variables $\overline{x}$. Suppose the sequent $\Lambda \Longrightarrow \Delta$ is not derivable in $\mathcal{A}(CQ(\mathcal{G}))$. It follows from the rules that $\Lambda \cup \{\neg\delta \mid \delta \in \Delta\}$ is consistent in $\mathcal{A}(CQ(\mathcal{G}))$. Let Γ be any set of formulas of $\mathcal{L}(CQ(\mathcal{G}))$ consistent in $\mathcal{A}(CQ(\mathcal{G}))$; by the above translation we can think of Γ as a set of first order formulae. Add a countable set of Henkin witnesses to the language. Γ can be extended to a set T maximally consistent in $\mathcal{A}(CQ(\mathcal{G}))$, satisfying

1. for each φ in the language, $\varphi \in T$ or $\neg\varphi \in T$
2. for first order φ, if $\exists x\varphi(x) \in T$, then for some a, $\varphi(a) \in T$.

An equivalence relation $\sim$ is defined on the constants by putting $a \sim b$ iff $a = b \in T$, and if $[a]$ is the $\sim$-equivalence class of a, one puts $A([a_1]\dots[a_n])$ iff $A(a_1 \dots a_n) \in T$, where A is a standard predicate. This determines a first order model $\mathcal{M}$. For formulas of the form $\exists(\varphi|\mathcal{G})$ the truth condition is similar: $\exists(\varphi|\mathcal{G})([a_1]\dots[a_n])$ iff $\exists(\varphi|\mathcal{G})(a_1 \dots a_n) \in T$. By induction one proves the truth lemma for $\mathcal{M}$ and T:

$$\mathcal{M} \models \varphi([a^1]\dots[a^k]) \text{ iff } \varphi(a^1 \dots a^k) \in T,$$

where $\langle a^1 \dots a^k\rangle$ is any sequence of constants.
Define $\mathcal{G} = \mathcal{G}'$ to be $\{\varphi \mid \mathcal{M} \models \forall\overline{x}(\varphi(\overline{x}) \leftrightarrow \exists(\varphi|\mathcal{G})(\overline{x}))\}$. We then have to show that $\mathcal{G}$ is a Boolean algebra, and that $\langle\mathcal{M}, \mathcal{G}, \mathcal{G}\rangle$ is a homogeneous conditional quantification model. It is easy to see that the properties of lemma 1 can be derived in $\mathcal{A}(CQ(\mathcal{G}))$. By lemma 1, part 5b, $\mathcal{G}$ is closed under $\wedge$, and by part 4 of the same lemma $\mathcal{G}$ is closed under $\vee$, hence $\mathcal{G}$ is a lattice. By rule lCQ, $\exists(\neg\exists(\varphi|\mathcal{G})|\mathcal{G}) = \neg\exists(\varphi|\mathcal{G})$, so that $\mathcal{G}$ is also closed under Boolean negation. By lemma 1, part 5b, $\exists(\exists(\varphi|\mathcal{G})|\mathcal{G}) = \exists(\varphi|\mathcal{G})$, so that the range of $\exists(\cdot|\mathcal{G})$ equals $\mathcal{G}$. The defining condition of $\mathcal{G}$ ensures that it is homogeneous.
It remains to check that $\mathcal{M} \models \exists(\varphi|\mathcal{G})[a^1 \dots a^k]$ iff for all $\exists(\psi|\mathcal{G}) \in \mathcal{G}' = \mathcal{G}$: $\varphi \leq \exists(\psi|\mathcal{G})$ on $\mathcal{M}$ implies $\mathcal{M} \models \exists(\psi|\mathcal{G})[a^1 \dots a^k]$. The direction from right to left is trivial: T is consistent with rCQ, hence $\forall\overline{x}(\varphi(\overline{x}) \rightarrow \exists(\varphi|\mathcal{G})(\overline{x})) \in T$. For

the converse direction, choose an element of $A \in \mathcal{G}$ such that on $\mathcal{M}$, $\varphi \leq A$. By construction, A is of the form $\exists(\psi|\mathcal{G})$. We now observe that since T is consistent with lCQ, if $\forall \overline{x}(\varphi(\overline{x}) \to \exists(\psi|\mathcal{G})(\overline{x})) \in T$, also $\forall \overline{x}(\exists(\varphi|\mathcal{G})(\overline{x}) \to \exists(\psi|\mathcal{G})(\overline{x})) \in T$, whence it follows that $\mathcal{M} \models \exists(\varphi|\mathcal{G})[a^1 \ldots a^k]$. $\square$

This result is far from optimal in that we have little control over the algebra $\mathcal{G}$ constructed in the proof of the completeness theorem, contrary to the intuitive motivation. Consider the following simple example. We have a language with two unary predicates A and C, satisfying $\forall x(C(x) \to A(x))$. We want to think of the algebra generated by A as the resource. Then, intuitively, $\exists(C|\mathcal{G})$ should be equal to A. However, a model as constructed in the proof of the completeness theorem, satisfying the Galois conditions for $\exists(\cdot|\mathcal{G})$ plus the property $\forall x(C(x) \to A(x))$, need not satisfy $\exists(C|\mathcal{G}) = A$; in general, one will only get that $\exists(C|\mathcal{G})$ is contained in A. Hence, despite appearances, the $S5$ axioms are not complete for the resource-bounded quantifier $\exists(C|\mathcal{G})$; precisely the fact that $\mathcal{G}$ is the resource in $\exists(C|\mathcal{G})$ is left out.

To see what kind of additional axiom might be needed, let us complicate the situation slightly by supposing that we now have finitely many predicates $A_1, \ldots, A_n$ all satisfying $\forall x(C(x) \to A_i(x))$, such that the resource $\mathcal{G}$ is now determined by the A_i. Then clearly we should have $\exists(C|\mathcal{G}) = \bigwedge_{i \leq n} A_i$, and this is forced by requiring additionally that $\forall xy(\bigwedge_{i \leq n}(A_i(x) \leftrightarrow A_i(y)) \to (\exists(C|\mathcal{G})(x) \leftrightarrow \exists(C|\mathcal{G})(y)))$. Indeed, the Galois condition forces that $\exists(C|\mathcal{G}) \subseteq \bigwedge_{i \leq n} A_i$, and if the inclusion were proper, this could be used to generate a counterexample to the axiom. Note that the axiom is a way of expressing in logic that $\exists(\cdot|\mathcal{G})$ is measurable with respect to the algebra $\mathcal{G}$.

We are now ready to introduce the formal system meeting our specifications, with the associated standard models. By induction, using the fact that homomorphisms are surjective, one first proves

Lemma 10 *Suppose $h : \mathcal{M} \longrightarrow \mathcal{N}$ is a homomorphism, π the associated projection on assignments. For every predicate A in the signature of $\mathcal{M}$, add a new predicate A', which is interpreted on $\mathcal{M}$ by $A'_{\mathcal{M}} = h^{-1}A_{\mathcal{N}}$. If $\mathcal{G}_{\mathcal{N}}$ is the algebra of sets of assignments first order definable on $\mathcal{N}$, then $\pi^{-1}(\mathcal{G}_{\mathcal{N}})$ is the algebra of sets of assignments on $\mathcal{M}$ first order definable with respect to the language only containing the A'.*

Definition 12 *A* standard *conditional quantification model is a homogeneous conditional quantification model $\langle \mathcal{M}, \mathcal{G}_0, \mathcal{G}, \mathcal{G}' \rangle$ such that $\mathcal{G}_0$ is a countable algebra of first order formulas (or the sets of assignments defined by these formulas) and $\mathcal{G} = \sigma(\mathcal{G}_0)$. As before $\mathcal{G}' \subseteq \mathcal{G}$ is the range of the conditional quantifier.*

By lemma 10, the condition that $\mathcal{G}_0$ is a countable set of first order formulas captures all cases of interest. Note that $\sigma(\mathcal{G}_0)$ is homogeneous if the generating algebra $\mathcal{G}_0$ is homogeneous; it follows that we may think of a formula $\exists(\theta|\mathcal{G})$, where θ has arity n, as determining a set of n-tuples.

The formal system now consists of $\mathcal{A}(CQ(\mathcal{G}))$ plus a new rule $(M)_{\mathcal{G}_0}$, where $\mathcal{G}_0$ is a Boolean algebra of first order formulas, which says:

$$\frac{\Gamma \Longrightarrow \exists(\varphi|\mathcal{G})(\overline{x}) \quad \{\Gamma, \psi(\overline{x}) \Longrightarrow \psi(\overline{y}), \Delta \mid \psi \in \mathcal{G}_0\}}{\Gamma \Longrightarrow \exists(\varphi|\mathcal{G})(\overline{y}), \Delta}$$

Clearly $(M)_{\mathcal{G}_0}$ has infinitely many premisses when $\mathcal{G}_0$ is infinite, although each sequent itself is finite. The next lemma establishes a connection between the rule $(M_{\mathcal{G}_0})$ and measurability with respect to $\mathcal{G}_0$ in the setting of Borel models:

Lemma 11 (Blackwell; cf. Exercise 14.16 in [9]) *Let X be a standard Borel space. Let (A_n) be a countable collection of Borel subsets of X; define an equivalence relation R by $R(x,y)$ iff $\forall n(x \in A_n \Leftrightarrow y \in A_n)$. Then for any Borel set B, B is R-invariant iff $B \in \sigma(A_n)$.*

Proof. The direction from right to left is proved by induction, and does not need the assumption that X is a Borel space. For the converse direction, observe that since the collection (A_n) is countable, each equivalence class is determined by a sequence in 2^ω. The canonical projection $\pi : X \longrightarrow 2^\omega$ is Borel measurable if we equip 2^ω with the product topology. The Borel σ-algebra on 2^ω is isomorphic to $\sigma(A_n)$ via π. For each Borel set B, $\pi(B)$ is analytic. For invariant Borel sets we have in addition that $\pi(B^c) = \pi(B)^c$, so that by Suslin's theorem $\pi(B)$ is actually Borel. It follows that $B \in \sigma(A_n)$. □

Lemma 12 *$(M)_{\mathcal{G}_0}$ is sound with respect to standard conditional quantification models.*

Proof. Let $\mathcal{M}$ be a standard conditional quantification model, f an assignment such that $\mathcal{M} \models \Gamma[f]$, and for all $\delta \in \Delta, \mathcal{M} \not\models \delta[f]$. We then have that $\exists(\varphi|\mathcal{G}) \in \mathcal{G} = \sigma(\mathcal{G}_0)$ and $\mathcal{M} \models \exists(\varphi|\mathcal{G})(\overline{x})[f]$. Let $f(\overline{x}) = \overline{a}$, $f(\overline{y}) = \overline{b}$. The second premiss, combined with the homogeneity of $\mathcal{G}_0$ says that $\overline{a}, \overline{b}$ are $\mathcal{G}_0$-equivalent. The easy half of Blackwell's lemma then ensures that $\overline{b} \in \exists(\varphi|\mathcal{G})$.

□

The main ingredient of the completeness proof is provided by the following definition

Definition 13 (H. Friedman) *A first order model $\mathcal{M}$ is* totally Borel *if its domain is (homeomorphic to) $\mathbb{R}$ and all relations definable on $\mathcal{M}$ by means of parameters are Borel.*

Theorem 6 (H. Friedman; cf. Steinhorn [14], [15]) *A first order theory with an infinite model has a totally Borel model.*

The addition of the infinitary rule necessitates a modified notion of a maximally consistent theory. As before, $\mathcal{G}_0$ is a homogeneous algebra of first order formulas.

Definition 14 *A set of formulas Γ in $\mathcal{L}(CQ(\mathcal{G}))$ is* $\mathcal{G}_0$-saturated *if it is maximal with respect to satisfying the following conditions*

1. Γ *is consistent in* $\mathcal{A}(CQ(\mathcal{G})) + (M)_{\mathcal{G}_0}$.

2. $\varphi \vee \psi \in \Gamma$ *implies* $\varphi \in \Gamma$ *or* $\psi \in \Gamma$.

3. $\exists x\varphi(x) \in \Gamma$ *implies* $\varphi(a) \in \Gamma$ *for some Henkin witness* a.

4. *If for some terms* s, t, $\exists(\varphi|\mathcal{G})(t), \neg\exists(\varphi|\mathcal{G})(s) \in \Gamma$, *then there exists a formula* $\psi \in \mathcal{G}_0$ *such that* $\psi(t), \neg\psi(s) \in \Gamma$.

Lemma 13 *Every set of formulas* Γ *consistent in* $\mathcal{A}(CQ(\mathcal{G})) + (M)_{\mathcal{G}_0}$ *has a* $\mathcal{G}_0$*-saturated extension.*

Proof. We do the two nontrivial cases. Enumerate the formulas in the extended language such that each formula occurs infinitely often. Suppose we have reached a stage T_k where the formula to be treated is $\exists(\varphi|\mathcal{G})(s) \in T_k$. Suppose for some t, $\neg\exists(\varphi|\mathcal{G})(t) \in T_k$. If for all $\psi \in \mathcal{G}_0$ such that $T_k \cup \{\psi(s)\}$ consistent, $T_k \cup \{\psi(s)\} \vdash \psi(t)$, then rule $(M)_{\mathcal{G}_0}$ would give us $T_k \vdash \exists(\varphi|\mathcal{G})(t)$, a contradiction. Hence we may choose ψ such that $T_k \cup \{\psi(s), \psi(t)\}$ is consistent. Do this for each t such that $\neg\exists(\varphi|\mathcal{G})(t) \in T_k$, and let the result be T_{k+1}.
We next show that witnesses for existential formulas can be added such that the resulting theory is still consistent with $(M)_{\mathcal{G}_0}$. So suppose we have constructed a theory T_n, consistent in $\mathcal{A}(CQ(\mathcal{G})) + (M)_{\mathcal{G}_0}$, such that $\exists x\varphi(x) \in T_n$; we want to construct T_{n+1} such that for some witness a, $\varphi(a) \in T_{n+1}$, while T_{n+1} is still consistent with $(M)_{\mathcal{G}_0}$. Assume for a term b and a formula θ, $\neg\exists(\theta|\mathcal{G})(b) \in T_n$. We have a contradiction with $(M)_{\mathcal{G}_0}$ if for all (new) witnesses a, on the one hand $T_n, \varphi(a) \vdash \exists(\theta|\mathcal{G})(a)$, and on the other hand, for all $\psi \in \mathcal{G}_0$, $T_n, \varphi(a) \vdash (\psi(a) \leftrightarrow \psi(b))$. However, rule $(M)_{\mathcal{G}_0}$ then implies that for all new a, $T_n, \varphi(a) \vdash \exists(\theta|\mathcal{G})(b)$, and it follows that $T_n \vdash \neg\exists x\varphi(x)$, a contradiction. □

Putting all these ingredients together we now have

Theorem 7 $\mathcal{A}(CQ(\mathcal{G}))+(M)_{\mathcal{G}_0}$ *is complete with respect to standard conditional quantification models.*

Proof. Let Γ be consistent in $\mathcal{A}(CQ(\mathcal{G})) + (M)_{\mathcal{G}_0}$. As before, think of conditionally quantified formulas as being predicates. We may assume that Γ has an infinite model, for if not, the argument given just after theorem 5 does the job. Extend Γ to Γ^* with built-in Skolem functions and extend Γ^* to a $\mathcal{G}_0$-saturated theory T. Let $\mathcal{M}$ be a totally Borel model for T. Define $\mathcal{G}'$ to be $\{\varphi \mid \mathcal{M} \models \forall\overline{x}(\varphi(\overline{x}) \leftrightarrow \exists(\varphi|\mathcal{G})(\overline{x}))\}$. Let $\mathcal{G}_0 = \{\varphi \mid \varphi \text{ first order}, \mathcal{M} \models \varphi \equiv \exists(\varphi|\mathcal{G})\}$. It remains to be shown that $\mathcal{G}' \subseteq \sigma(\mathcal{G}_0)$. Define the equivalence relation R by $R(f, g)$ iff for all $\psi \in \mathcal{G}_0, \mathcal{M} \models \psi[f] \Leftrightarrow \mathcal{M} \models \psi[g]$. R is a Borel equivalence relation on $I\!R^\omega$ and by Blackwell's lemma it suffices to show that each $\exists(\varphi|\mathcal{G})$ is invariant under R (since $\exists(\varphi|\mathcal{G})$ is a Borel subset of $I\!R^\omega$).
For notational convenience we shall assume that $\exists(\varphi|\mathcal{G})(x)$ only contains the free variable x. The proof of Theorem 6 shows that each element of the universe of $\mathcal{M}$ is of the form $t(i_1, \ldots, i_n)$, where t is a Skolem term and $i_1, \ldots, i_n$ are indiscernibles from a linear order of indiscernibles homeomorphic to the irrationals. Now suppose $\mathcal{M} \models \exists(\varphi|\mathcal{G})(x)[f]$, where $f(x) = t(\overline{\imath})$, for t a Skolem

term. Let $R(f, g)$ with $g(x) = s(\overline{j})$. We then have, for all $\psi \in \mathcal{G}_0$, $\mathcal{M} \models \psi(t(\overline{i}))$ implies $\mathcal{M} \models \psi(s(\overline{j}))$. By construction this is equivalent to: for all $\psi \in \mathcal{G}_0$, $T \vdash \psi(t(\overline{u}))$ implies $T \vdash \psi(s(\overline{v}))$. Since $\exists(\varphi|\mathcal{G})(t(\overline{u})) \in T$, by rule $(M)_{\mathcal{G}_0}$ this implies $\exists(\varphi|\mathcal{G})(s(\overline{v})) \in T$, whence $\mathcal{M} \models \exists(\varphi|\mathcal{G})(s(\overline{j}))$. $\square$

A much simpler argument also shows that the rules given (i.e. including the infinitary rule) are also complete with respect to models of the form $\langle \mathcal{M}, \mathcal{G}_0, \mathcal{G} \rangle$, $\mathcal{G}$ the algebra of all $\mathcal{G}_0$-invariant sets. However, the point of the argument just given is that it forces the denotations of the formulas to be in an algebra which is effectively described.

The argument goes over without change to families of conditional quantifiers. Say the language is now $\mathcal{L}(CQ(\{\mathcal{G}_s \mid s \in T\}))$; this requires the obvious extension of $\mathcal{A}(CQ(\mathcal{G})) + (M)_{\mathcal{G}_0}$ to $\mathcal{A}(CQ(\mathcal{G})) + \{(M)_{\mathcal{G}_{s0}} \mid s \in T\}$. We then have

Corollary 1 *$\mathcal{A}(CQ(\mathcal{G})) + \{(M)_{\mathcal{G}_{s0}} \mid s \in T\}$ is complete with respect to standard homogeneous conditional quantification models of the form $\langle \mathcal{M}, \{\mathcal{G}_{s0} \mid s \in T\}, \{\mathcal{G}_s \mid s \in T\}, \{\mathcal{G}'_s \mid s \in T\}\rangle$, where each $\mathcal{G}'_s \subseteq \mathcal{G}_s$ and $\mathcal{G} = \sigma(\mathcal{G}_{s0})$.*

6 Averages

One of the themes of our investigation is the relation between the resource chosen and the logical properties of the conditional quantifier. We have seen that it is in general impossible to construct a conditional quantifier satisfying the $S5$ properties, which is conditional on a resource which is the σ-algebra generated by a set of first order formulae. Here, we show that, if we drop the requirement that $\varphi \leq \exists(\varphi \mid \mathcal{G})$, then conditioning on a σ-algebra is possible. We also obtain a very intuitive correspondence between this form of conditional quantification and conditional expectation, whenever the latter is defined. There is a price to pay, however; the construction requires the continuum hypothesis, and the models have to be expanded by means of an ideal. We first give an algebraic definition of the quantifier of interest.

Definition 15 (Wright [25]) *Let $\mathcal{B}$ be a Boolean algebra. An* average *is a mapping $\exists : \mathcal{B} \longrightarrow \mathcal{B}$, satisfying*

1. $\exists \mathbf{0} = \mathbf{0}, \exists \mathbf{1} = \mathbf{1}$
2. *$\exists$ commutes with $\vee$*
3. *for all A in the range of $\exists$, for all $B \in \mathcal{B}$,*
   (a) $B \leq A \Longrightarrow \exists B \leq A$
   (b) $B \wedge A = \mathbf{0} \Longrightarrow \exists B \wedge A = \mathbf{0}$.

Equivalently,

Lemma 14 (Wright [25]) *An average is a mapping $\exists : \mathcal{B} \longrightarrow \mathcal{B}$, satisfying*

1. $\exists\mathbf{0} = \mathbf{0}, \exists\mathbf{1} = \mathbf{1}$

2. $\exists$ *commutes with* $\vee$

3. $\exists(A \wedge \exists B) = \exists A \wedge \exists B$.

An example of an average on a Boolean algebra $\mathcal{B}$ is obtained when there exists a subalgebra $\mathcal{G}$ which is a retract of $\mathcal{B}$, i.e. when there exists a surjective homomomorphism $h : \mathcal{B} \longrightarrow \mathcal{G}$ which is the identity on $\mathcal{G}$. Referring to example 3.1, we obtain an average when a predicate in $\mathcal{L}$ but not in $\mathcal{L}'$ is defined by means of a formula in $\mathcal{L}'$. An extreme case of an average is a homomorphism $h : \mathcal{B} \longrightarrow \mathbf{2}$, i.e. an ultrafilter. A more typical example of an average $\exists$ arises is obtained when we have a probability measure P on $\mathcal{B}$, so that we may define, for all $a \in \mathcal{B}$: $\exists a = \mathbf{1}$ iff $P(a) > 0$. This example suggests that an average can also be characterised by means of a modified Galois condition, involving an ideal.

Lemma 15 (Wright [25]) *If $\exists$ is an average with range $\mathcal{G}$, the set $I = \{N \mid N \leq B \wedge \neg\exists B$, for some $B \in \mathcal{B}\}$ is an ideal. Moreover, $I \cap \mathcal{G} = \{\mathbf{0}\}$.*

Proof. By the Frobenius property, $\exists(B \wedge \neg\exists B) = \exists B \wedge \neg\exists B = \mathbf{0}$, and the result follows since $\exists$ commutes with $\vee$. For the second part of the assertion, assume $N \in I \cap \mathcal{G}$, then $\exists N = N$ and there exists B such that $N \leq B \wedge \neg\exists B$. Hence $N = \exists N \leq \exists(B \wedge \neg\exists B) = \mathbf{0}$. □

Lemma 16 *A mapping $\exists : \mathcal{B} \longrightarrow \mathcal{B}$ is an average iff there exists an ideal I on $\mathcal{B}$ such that $I \cap range(\exists) = \{\mathbf{0}\}$ and, for all $B \in \mathcal{B}$, all A in the range of $\exists$:*

$$B \wedge \neg A \in I \Longleftrightarrow \exists B \leq A.$$

Proofsketch. If $\exists$ is an average, $B \wedge \neg A \in I$ implies $\exists(B \wedge \neg A) = \exists B \wedge \neg A = \mathbf{0}$, whence $\exists B \leq A$. If $\exists B \leq A$, then $B \wedge \neg A \in I$ since $B \wedge \neg\exists B \in I$. □

The latter characterisation will be used in our semantics for averages, now considered as quantifiers added to a first order language. It suggests that in this case the resource consists of two parts: an algebra and an ideal of negligible sets.

We are now in a position to elaborate the example sketched in section 3.4. Consider a partial homomorphism $h : \mathcal{M} \longrightarrow \mathcal{N}$ with domain D; I is then the prime ideal determined by $\mathcal{F} - D^\omega$. Let $\mathcal{G}$ be the algebra on $\mathcal{M}$ determined by $\mathcal{N}$ and h; for the moment we waive issues about completeness. Define a quantifier $\exists : \mathcal{B} \longrightarrow \mathcal{G}$ by $\exists\varphi = \bigwedge\{A \in \mathcal{G} \mid \varphi \leq_I A\}$, then the range of $\exists$ will be a subalgebra of $\mathcal{G}$ and it will satisfy all properties of an average except $\exists\mathbf{1} = \mathbf{1}$; in fact we will have $\exists D^\omega = \exists\mathbf{1} = D^\omega$. For the purposes of the logic of vision outlined in [19] this is not satisfactory, since it provides the perceptual field with sharp boundaries; whereas the logic should allow for the possibility that one sees an object which does not occur in the perceptual field, so that $\exists D^\omega = \exists\mathbf{1} \neq D^\omega$. It is then technically convenient to have $\exists$ satisfy $\exists\mathbf{1} = \mathbf{1}$.

This requires changing the set-up; we now must assume that h derives from a total homomorphism h' so that $h = h'|D$ and $\mathcal{G}$ is generated from h' rather than from h. The desideratum that the perceptual field does not have sharp boundaries translates into the requirement that no nonempty part of D^ω is in $\mathcal{G}$, or, what amounts to the same, that $I \cap \mathcal{G} = \{\mathbf{0}\}$. Under these conditions, $\exists$ is an average. As an interesting additional consequence we obtain that $\exists$ is in fact surjective on $\mathcal{G}$, for if $A, B \in \mathcal{G}$, then $A \leq_I B$ iff $A \leq B$. This example clearly shows that the two resources defining an average, concretely: approximate vision and restricted perceptual field, should be independent.

6.1 Existence of averages

Wright [25][p. 446] gives a representation theorem for averages in terms of hypersections on the Stone spaces of $\mathcal{B}$ and $\mathcal{G}$. This gives us little information about the kind of resource that can actually occur for an average, so we shall approach the matter from the other end, by fixing the resource.

Definition 16 *An ideal I in a Boolean algebra $\mathcal{B}$ is ω_1-saturated if for any family of pairwise disjoint sets $\{B_\alpha | \alpha < \omega_1\} \subseteq \mathcal{B} - I$, there exist distinct B_α, B_γ such that $B_\alpha \cap B_\gamma \notin I$.*

Lemma 17 *Let $\mathcal{B}$ be a Boolean σ-algebra. Then a σ-ideal $I \subseteq \mathcal{B}$ is ω_1-saturated iff there does not exist an uncountable set $A \subseteq \mathcal{B} - I$ of pairwise disjoint sets.*

PROOF. The direction from left to right is trivial. For the other direction, suppose for all distinct $B_\alpha, B_\gamma, B_\alpha \cap B_\gamma \in I$. Define a pairwise disjoint family $\{A_\alpha | \alpha < \omega_1\}$ by putting

$$A_0 = B_0$$
$$A_\alpha = B_\alpha \cap (\bigcup_{\gamma<\alpha} B_\gamma)^c.$$

By hypothesis, $B_\alpha \cap \bigcup_{\gamma<\alpha} B_\gamma \in I$, whence $A_\alpha \notin I$. It follows that the family A_α must be countable, a contradiction. □

Corollary 2 *Let (X, τ) be a second-countable topological space, $\mathcal{B}$ the Borel σ-algebra on X.*

1. *If μ is a σ-additive measure on $\mathcal{B}$, then $I = \{N \mid \mu(N) = 0\}$ is an ω_1-saturated σ-ideal.*

2. *$I = \{N \mid N \text{is meagre}\}$ is an ω_1-saturated σ-ideal.*

PROOF. 1 is immediate by Lemma 17.
For 2, let $\{B_\alpha\}_{\alpha<\omega_1}$ be an uncountable pairwise disjoint subset of $\mathcal{B} - I$. Let O_α be an open set congruent to B_α modulo I. The O_α can at most overlap on a meagre set and hence must be pairwise disjoint by the Baire category theorem. This contradicts the fact that the space is second-countable. □

Recall that a collection of sets is a π-*system* if it is closed under finite intersection.

Theorem 8 *$(ZFC + CH)$ Let $\mathcal{B}$ be the σ-algebra generated by a class of first order definable sets of assignments, let $\mathcal{G}_0$ be a countable subalgebra of $\mathcal{B}$ determined by a set of first order formulas with generated σ-algebra $\sigma(\mathcal{G}_0)$, and let I be an ω_1-saturated σ-ideal. Furthermore, suppose that for some π-system Δ generating $\mathcal{G}_0$, $I \cap \Delta = \{\emptyset\}$. Then there exists an average $\exists(\cdot|\mathcal{G}_0, I) : \mathcal{B} \longrightarrow \sigma(\mathcal{G}_0)$, satisfying*

$$B \subseteq_I A \text{ iff } \exists(B|\mathcal{G}_0, I) \subseteq A,$$

for $B \in \mathcal{B}$, A in the range of $\exists(\cdot|\mathcal{G}_0, I)$, and such that $\mathcal{G}_0$ is contained in the range.

Proof. The quotient σ-algebra $\mathcal{B}/I$ satisfies the countable chain by ω_1-saturation of I, hence is complete; analogously for the subalgebra $\sigma(\mathcal{G}_0)/I$. Let π be the canonical projection $\mathcal{B} \longrightarrow \mathcal{B}/I$. Define a map $F : \mathcal{B} \longrightarrow \sigma(\mathcal{G}_0)/I$ by $F(B) = \bigwedge\{A \in \sigma(\mathcal{G}_0)/I | \pi(B) \leq A\}$.

We now choose representatives from the elements of $\sigma(\mathcal{G}_0)/I$ such that the algebraic relations are preserved and $\mathcal{G}_0$ is included in the set of representatives.

In general, let $\mathcal{G}$ be a Borel σ-algebra generated by a topology τ, and let I be a σ-ideal disjoint from τ. A *lifting* $\rho : \mathcal{G} \longrightarrow \mathcal{G}$ is a map which satisfies

1. $\rho(A) = A \pmod I$
2. $A = B \pmod I$ implies $\rho(A) = \rho(B)$
3. $\rho(\emptyset) = \emptyset$
4. $\rho(A \cap B) = \rho(A) \cap \rho(B)$
5. $\rho(A \cup B) = \rho(A) \cup \rho(B)$.

A lifting is *strong* if we have in addition

(vi) for $U \in \tau$, $U \subseteq \rho(U)$.

The existence of liftings can be proved in ZFC; but assuming the continuum hypothesis, a strong lifting ρ exists for $\mathcal{G}$ if τ is second countable and Hausdorff; see Graf [5][p. 425].

In particular this is true for our case, where $\mathcal{G} = \sigma(\mathcal{G}_0)$, for as in the proof of lemma 11, we may view $\mathcal{G}$ as the Borel σ-algebra on 2^ω. Moreover, since $\mathcal{G}_0$ consists of clopen sets, we may use (vi) to conclude that ρ is actually the identity on $\mathcal{G}_0$.

Note that we may take ρ to be a mapping on $\sigma(\mathcal{G}_0)/I$ and define $\exists(B|\mathcal{G}_0, I) = \rho(F(B))$. Then the range of $\exists(B|\mathcal{G}_0, I)$ is a subalgebra of $\sigma(\mathcal{G}_0)$ containing $\mathcal{G}_0$ and for first order φ, $\exists(\varphi|\mathcal{G}_0, I) \subseteq \bigwedge\{C \in \mathcal{G}_0 | \varphi \subseteq C\}$.

Lastly, we show that the Galois condition holds. The direction from right to left is trivial, since by construction $B \subseteq_I \exists(B|\mathcal{G}_0, I)$. For the converse direction, if $B \subseteq_I A$, where A is in the range of $\exists(\cdot|\mathcal{G}_0, I)$ then $\pi(B) \leq \pi(A)$, whence

$F(B) \leq F(A)$. It follows that $\rho(F(B)) \subseteq \rho(F(A))$, but since A is in the range of $\exists(\cdot|\mathcal{G}_0, I)$, $\rho(F(A)) = A$. $\Box$

In logical terms, the meaning of the result is the following. Suppose we drop ordinary $\forall, \exists$ altogether and work in the fragment of $L_{\omega_1\omega}$ not containing quantifiers, as in Keisler's probability logic [10][8]. Let $L_{\omega_1\omega}(Av(\mathcal{G}))$ denote this fragment with the average $\exists(\cdot|\mathcal{G}, I)$ added.

Corollary 3 *(ZFC + CH) Let $\mathcal{M}$ be a model, $\mathcal{G} = \sigma(\mathcal{G}_0$ a countably generated σ-algebra, I an ω_1-saturated σ-ideal on $\mathcal{G}$ such that $I \cap \mathcal{G}_0 = \{\emptyset\}$. Then $\mathcal{M}$ can be extended to a model of $L_{\omega_1\omega}(Av(\mathcal{G}))$ with $\exists(\cdot|\mathcal{G}, I)$ interpreted by means of $\mathcal{G}, I$.*

$\Box$

The use of a lifting in the proof of theorem 8 is unavoidable, as can be seen as follows. Suppose we are given an average $\alpha : \mathcal{B} \longrightarrow \mathcal{G}$ with range $\mathcal{G}$; let I be the ideal $\{a \in \mathcal{B}|\alpha a = \mathbf{0}\}$. Let $alg(\mathcal{G} \cup I)$ be the Boolean algebra generated by $\mathcal{G} \cup I$, then $\mathcal{G} \cap I = \{\emptyset\}$ and $\mathcal{G} \simeq alg(\mathcal{G} \cup I/I)$. Let π be the canonical projection $\pi : \mathcal{B} \longrightarrow \mathcal{B}/I$, then we can define a conditional quantifier $\exists : \mathcal{B}/I \longrightarrow alg(\mathcal{G} \cup I)/I$ by $\exists \pi a = \pi \alpha a$. Furthermore, the range of α is obtained by applying a lifting to the range of $\exists$. This parallels the steps in our proof. Furthermore, we want the lifting to be strng in order to ensure that the generating algebra $\mathcal{G}_0$ is in the range; I am not aware of any method of doing this without the continuum hypothesis.

Measures provide a source for the required ideals. Assume for the following that $\mathcal{M}$ is ω_1-saturated. Let $\mathcal{F}$ be the space of assignments, with its compact topology of first order definable subsets. Let $\mathcal{C}(\mathcal{F})$ denote the set of continuous functions on $\mathcal{F}$, and let $\mathcal{M}(\mathcal{F})$ be the space of probability measures on $\mathcal{F}$. $\mathcal{M}(\mathcal{F})$ is compact in the weak $\star$-topology, the smallest topology making each of the maps $\mu \longrightarrow \int F d\mu$, $F \in \mathcal{C}(\mathcal{F})$, continuous. It follows that there exist many non-atomic measures on $\mathcal{F}$. Indeed, take any sequence (f_n) of elements of $\mathcal{F}$, and define a measure μ_n by

$$\mu_n(B) = \frac{1}{n} \sum_{i=1}^{n} \chi_B(f_i),$$

where χ_B is the characteristic function of B. By compactness, the sequence (μ_n) will have a limit point in $\mathcal{M}(\mathcal{F})$. Given a $\mu \in \mathcal{M}(\mathcal{F})$, define $I = \{B|\mu(B) = 0\}$, then I is an ω_1-saturated σ-ideal. Given a π-system Δ generating $\mathcal{G}_0$, we can always construct a measure μ which assigns positive measure to each element of Δ (by taking a suitable infinite convex combination) and in this case we will have $\Delta \cap I = \{\emptyset\}$.

We are now in a position to establish a connection between averages and conditional expectation.

[8]Although we do not need Keisler's restriction to admissible fragments.

Corollary 4 *(ZFC + CH) Let $\mathcal{G} = \sigma(\mathcal{G}_0)$, let μ be a probability measure on $\mathcal{B}$, the σ-algebra generated by the first order definable sets of assignments, and let $\mathbf{E}_\mu(\cdot|\mathcal{G})$ denote conditional expectation with respect to μ. Let I be the set of μ-nullsets. Then for all $B \in \mathcal{B}$, for μ-a.a. $f \in \mathcal{F}$:*

$$f \in \exists(B|\mathcal{G}, I) \Longleftrightarrow \mathbf{E}_\mu(B|\mathcal{G})(f) > 0.$$

Proof. From right to left: suppose $\mu\{f|\mathbf{E}_\mu(B|\mathcal{G})(f) > 0 \,\&\, f \notin \exists(B|\mathcal{G}, I)\} > 0$, then

$$\int \mathbf{E}_\mu(B|\mathcal{G}) \cdot \mathbf{1}_{\exists(B|\mathcal{G},I)^c} d\mu > 0,$$

whence

$$\int \mathbf{1}_B \cdot \mathbf{1}_{\exists(B|\mathcal{G},I)^c} d\mu > 0,$$

a contradiction. For the converse direction, we use the Galois property (and CH). Let ρ be the strong lifting used to construct $\exists(\cdot|\mathcal{G}, I)$. Put

$$A = \rho([\{f \mid \mathbf{E}_\mu(B|\mathcal{G})(f) > 0\}]),$$

then, since A is in the range of the average, it will follow that $\exists(B|\mathcal{G}, I) \subseteq A$ if $B \subseteq A$ mod I.
But we have $\int_{A^c} \mathbf{1}_B d\mu = \int_{A^c} \mathbf{E}_\mu(B|\mathcal{G}) d\mu = 0$, and we are done. □

6.1.1 Uniqueness

One of the questions we posed at the beginning was: is the quantifier unique given the resource? In this case the resource consists of an algebra together with an ideal. If the algebra were complete, the average would be definable as $\exists(\varphi|\mathcal{G}, I) = \bigwedge\{\psi \in \mathcal{G} \mid \varphi \leq_I \psi\}$, so that it is unique. In the case of countable completeness, the situation is different. Suppose we have averages $\exists_0, \exists_1$ both with respect to the resource $\mathcal{G}, I$; let $\exists_0$ have range $\mathcal{G}_0 \subseteq \mathcal{G}$ and let $\exists_1$ have range $\mathcal{G}_1 \subseteq \mathcal{G}$. An easy computation then shows that for all b in the domain of $\exists_0, \exists_1$, $\exists_0 b \triangle \exists_1 b \in I$; uniqueness holds only modulo the ideal. This means that also here $\exists(\varphi|\mathcal{G}, I)$ is not computable from φ and the resource $\mathcal{G}$. At most the equivalence class containing $\exists(\varphi|\mathcal{G}, I)$ can be determined from knowledge of $\mathcal{G}$, φ and I, but the actual choice depends upon CH.

6.2 Logic of averages

6.2.1 Syntax

The syntax of the logic is the same as that for conditional quantifiers, except that we write an average as $\exists(\cdot|\mathcal{G}, I)$, thus explicitly mentioning the second component of the resource. To mark the difference the language of the average-quantifier will be denoted $\mathcal{L}(Av)$.

From the rules of the sequent calculus $\mathcal{A}(CQ(\mathcal{G})) + (M)_{\mathcal{G}_0}$ we keep only the left introduction rule. Note that this takes care of both parts of the modified

Galois condition in 15, due to the presence of the side formulas Γ, Δ. Dropping the right introduction rule has the consequence that we have to force distribution over $\vee$ by means of the axiom

$$\exists(\varphi \vee \psi|\mathcal{G}, I) \Longrightarrow \exists(\varphi|\mathcal{G}, I), \exists(\psi|\mathcal{G}, I).$$

The resulting system is called $\mathcal{A}(Av) + (M)$.

6.2.2 Semantics

Motivated by the discussion in section 6.1, models for averages will take the following form

Definition 17 *A* standard average model *is a structure of the form $\langle \mathcal{M}, \mathcal{G}, \mathcal{G}_0, \mathcal{G}', I\rangle$, where $\mathcal{M}$ is a first order model, $\mathcal{G}_0$ is a homogeneous algebra of first order definable sets of assignments, $\mathcal{G} = \sigma(\mathcal{G}_0)$, $\mathcal{G}'$ is a Boolean algebra contained in $\mathcal{G}$, and I is an ideal such that $I \cap \mathcal{G}' = \{\mathbf{0}\}$. The average $\exists(\cdot|\mathcal{G}, I)$ is interpreted on this structure as follows*

1. *the range of $\exists(\cdot|\mathcal{G}, I)$ is $\mathcal{G}'$*
2. *$\mathcal{M} \models \exists(\varphi|\mathcal{G}, I)[f]$ iff for all $A \in \mathcal{G}', \varphi \leq_I A \iff f \in A$ (where $\varphi \leq_I A$ means that $\varphi \wedge \neg A \in I$).*

The soundness of $\mathcal{A}(Av) + (M)_{\mathcal{G}_0}$ with respect to standard average models is proved as for standard conditional quantification models.

6.2.3 Completeness

Theorem 9 *$\mathcal{A}(Av) + (M)_{\mathcal{G}_0}$ is complete with respect to standard average models.*

Proof. The proof is almost the same is that of theorem 7, but of course we now have to take care of the ideal. Let I be the ideal of $\mathcal{L}(Av)$-definable sets of assignments generated by $\{\varphi \wedge \neg\exists(\varphi|\mathcal{G}, I) \mid \varphi \text{ a formula of } \mathcal{L}(Av)\}$. We must show that $I \cap \mathcal{G}' = \{\mathbf{0}\}$. An element in $I \cap \mathcal{G}'$ is a formula ψ such that for a set $\{\varphi_n \wedge \neg\exists(\varphi_n|\mathcal{G}, I)|n \leq k\}$, $\psi \leq \bigvee_{n \leq k}(\varphi_n \wedge \neg\exists(\varphi_n|\mathcal{G}, I))$. Applying $\exists(\cdot|\mathcal{G}, I))$ to the *r.h.s* yields $\mathbf{0}$ in virtue of the Frobenius property, hence the *r.h.s.* is in I. It follows that $\exists(\psi|\mathcal{G}, I) = \mathbf{0}$. Since $\psi \in \mathcal{G}'$, $\exists(\psi|\mathcal{G}, I) = \psi$, whence $\psi = \mathbf{0}$.

It is now easy to see that $\exists(\varphi|\mathcal{G}, I) = \bigwedge\{A \in \mathcal{G}' \mid \varphi \leq_I A\}$. The direction from right to left follows because by definition of I, $\varphi \leq_I \exists(\varphi|\mathcal{G}, I)$. For the converse direction, observe that if $\varphi \leq_I A$, then there are formulae $\psi_n,\ n \leq k$ such that $\varphi \wedge \neg A \leq \bigvee_{n \leq k} \psi_n \wedge \exists(\psi_n|\mathcal{G}, I)$, whence we have $\exists(\varphi \wedge \neg A|\mathcal{G}, I) \leq \bigvee_{n \leq k} \exists(\psi_n \wedge \exists(\psi_n|\mathcal{G}, I) \mid \mathcal{G}, I) = \mathbf{0}$ by Frobenius, so that $\exists(\varphi|\mathcal{G}, I) \leq A$. □

The ideal constructed in the course of the completeness proof apparently lacks the strong properties demanded by the existence theorem 8. One way to force those properties would be to show that every theory consistent with $\mathcal{A}(Av) + (M)_{\mathcal{G}_0}$ can be extended to a theory consistent with Friedman's logic for 'almost all' (cf. Steinhorn [14]) together with a set of axioms of the form $Qx\neg(\varphi(x) \wedge \neg\exists(\varphi|\mathcal{G}, I)(x))$, where Q is the 'almost all' quantifier. We have not been able to show, however, that such an extension is always possible.

7 Interlude: illusory conjunctions

Now that we have treated conditional quantifiers and averages in some detail, we briefly pause to give an application of these quantifiers to the logic of vision outlined in [21], [19] and [20]. The main thrust of those papers is, that if one takes seriously the idea that seeing is filtering of information, then a plausible logic of perception and of perception reports can be obtained if the required filtering is performed by various quantifiers of the conditional type. Here, we extend the argument by applying the logic of conditional quantifiers to a psychological phenomenon called 'illusory conjunction', for which see Treisman and Gelade [17], Treisman and Schmidt [18] and Treisman [16].

First a brief remark on the relation between conditional quantification and vision. It seems profitable to split the meaning of the sentence 'I see a φ' into two components, the first semantic, the second pragmatic. The semantic component is 'With the present approximation what I perceive $[x]$ is identified as a φ'[9], whereas the pragmatic component can be rendered as 'I expect this identification to remain true for every more refined approximation'. Some of the theory of the pragmatic component will be studied in section 9. The semantic component is formalised as $\exists(\varphi(x)|\mathcal{G})$, where $\mathcal{G}$ is the filtering structure corresponding to 'the present approximation'. The reasons for this formalisation, which derive from the importance of filtering to perception as discussed in Marr [12], are extensively discussed in the papers cited, especially [20], and we will not dwell upon this here. But note a consequence of this choice: there is a difference between $\exists(\varphi|\mathcal{G}) \wedge \exists(\psi|\mathcal{G})$ and $\exists(\varphi \wedge \psi|\mathcal{G})$, the latter being in general the smaller set. It was left open in [21] and [19] whether this difference has any visual meaning. (Incidentally, situation semantics faces this problem too, as was noted by Barwise in [2].) This is where Treisman's work comes in, which we shall now briefly describe.

The phenomenon to be explained, illusory conjunctions, is that features (such as color or shape) which are perceived in a unattended part of the visual field, may combine in a random manner. Thus, if the perceptual field is a display containing only pink X's and green T's, one may still report seeing a pink T when experimental conditions are such that serially scrutinising each letter is impossible. The pink T is an illusory conjunction in the sense that it is spuriously combining features from two objects which do exist. The theory proposed to explain the phenomenon is that features are detected in parallel and fast, but in such a manner that they are not yet bound to a location; it requires attention, which consists in serially scrutinising each object or location to conjoin the features into a unitary whole. Experimentally, the distinction between searching for features and searching for an object shows up as follows.

[9] Here, 'perceive' is taken in a passive, pre-attentive, sense.

Suppose we have a display which is filled with green X's and brown T's and which may or may not contain the target. In the feature condition, one asks subjects to search for a target which is either blue or an S; in the object condition, the target is a green T. In both cases one has to search for two features simultaneously, but they have to be conjoined only in the latter condition. The experimental findings are that in the feature condition search time is independent of display size, whereas in the object condition search time increases linearly with display size (where display size is measured by the number of items the display contains). This strongly suggests that search in the first condition is parallel, whereas it is serial in the second condition. Parallel search points to the existence of many specialised feature detectors, each devoted to a small part of the visual field, while serial search is a consequence of focussing attention on one item at a time.

What has all this to do with logic? Our claim is that viewing visual perception as filtering leaves room for illusory conjunctions; furthermore, that, also as a matter of logic, restricting attention to one object, as formalised by a suitable ideal, eliminates the illusory conjunctions. Formally, an illusory conjunction of predicates A and B is a situation where (relative to a visual field) $\forall x(A(x) \to \neg B(x))$, but $\exists x(\exists(A(x)|\mathcal{G}) \wedge \exists(B(x)|\mathcal{G}))$; the existential quantifier is the correct formalisation here, because we are asking only for the existence of a target, not for its location. Such situations can be constructed easily, not only for conditional quantifiers, but also for averages. (A concrete instance of a filter here is given by the experiment reported in Treissman and Gelade [17] which shows that judgments of presence or absence of a feature are independent of judgments of spatial location of that feature.) Now recall from section 6 that we may model the effect of a restricted perceptual field by means of averages; if the set of objects D represents the perceptual field, the ideal I is given by the subsets of $(D^c)^\omega$. Similarly we may go one step further and use I to represent the field of attention; if this consists of the object d, we may put $I = \{F \subseteq \mathcal{F} \mid f(x) \neq d\}$. Let $\exists(\cdot|\mathcal{G}, I)$ be the associated average. Then we have generally

$$\exists(\varphi(x)|\mathcal{G}, I) \wedge \exists(\psi(x)|\mathcal{G}, I) = \exists(\varphi(x) \wedge \psi(x)|\mathcal{G}, I).$$

There are two cases. First suppose $\varphi(d) \wedge \psi(d)$. By definition of an average, $\varphi(x) \wedge \psi(x) \leq_I \exists(\varphi(x) \wedge \psi(x)|\mathcal{G}, I)$, so that for our definition of I, $f \in \varphi(x) \wedge \psi(x) \wedge \neg\exists(\varphi(x) \wedge \psi(x)|\mathcal{G}, I)$ implies $f(x) \neq d$. It follows that $\varphi(x) \leq_I \exists(\varphi(x) \wedge \psi(x)|\mathcal{G}, I)$, or equivalently $\varphi(x) \wedge \neg\exists(\varphi(x) \wedge \psi(x)|\mathcal{G}, I) \in I$: for if $f \in \varphi(x)$ and $f(x) = d$, then $f \in \psi(x)$ and by the preceding observation $f \in \exists(\varphi(x) \wedge \psi(x)|\mathcal{G}, I)$. It follows that $\exists(\varphi(x)|\mathcal{G}, I) \leq \exists(\varphi(x) \wedge \psi(x)|\mathcal{G}, I)$, whence also $\exists(\varphi(x)|\mathcal{G}, I) \wedge \exists(\psi(x)|\mathcal{G}, I) \leq \exists(\varphi(x) \wedge \psi(x)|\mathcal{G}, I)$. In the second case, $\varphi(d) \wedge \neg\psi(d)$. Now we must have $\psi(x) \in I$, so that $\exists(\psi(x)|\mathcal{G}, I) = \mathbf{0}$ and again $\exists(\varphi(x)|\mathcal{G}, I) \wedge \exists(\psi(x)|\mathcal{G}, I) \leq \exists(\varphi(x) \wedge \psi(x)|\mathcal{G}, I)$. The argument depends upon the fact that I contains all subsets of $\{f \mid f(x) \neq d\}$. It would fail, as it should, if I allowed attention to be divided between at least two objects.

Summarising: if one takes seriously the point of view that perception involves filtering, or information reduction, then illusory conjunctions can be expected.

As a matter of logic, the only way to eliminate these conjunctions is by focussing attention on one object at a time.

8 Other conditioning structures

One theme of this investigation is the relation between the strength of the resource and the logical properties of the quantifier bounded by this resource. Here is one easy result, already used implicitly in section 6, which shows how fixing the logic also fixes the structure of the resource.

Lemma 18 (Wright [25]) *Let $\exists$ be an average defined on a Boolean algebra $\mathcal{B}$. Then its range $\mathcal{G}$ is a Boolean algebra.*

Proof. It suffices to show that the complement of $\mathcal{B}$ is the complement of $\mathcal{G}$. Choose $\exists a \in \mathcal{G}$, and let $(\exists a)'$ be its complement in $\mathcal{B}$. We are done if we can show that $(\exists a)' \in \mathcal{G}$, or, equivalently, $\exists((\exists a)') = (\exists a)'$. We have $\mathbf{1} = \exists(\exists a \vee (\exists a)') = \exists\exists a \vee \exists((\exists a)') = \exists a \vee \exists((\exists a)')$, whence $(\exists a)' \leq \exists((\exists a)')$.
On the other hand, $\mathbf{0} = \exists a \wedge (\exists a)'$, whence $\exists((\exists a)') \wedge \exists a = \mathbf{0}$, i.e. $\exists((\exists a)') \leq (\exists a)'$. □

Hence if we have a resource $\mathcal{H} \subseteq \mathcal{B}$ which does not contain a nontrivial Boolean subalgebra, it will be impossible to define an average $\exists : \mathcal{B} \longrightarrow \mathcal{H}$ with nontrivial range.

One may also ask several converse questions. Suppose we fix (i) the structure of the resource, (ii) the desired properties of the quantifier conditional on this resource and (iii) the basic first order logic, what is the resulting logic of the conditional quantifier? This demand for a complete characterisation of course presupposes that (i) - (iii) are consistent. Ideally, the complete axiomatisation would be obtained by an operation on the algebras given by (i) and (iii), but this appears to be difficult in general. However, we already proved one result along these lines, namely lemma 4, which can be restated as follows

Lemma 19 *Suppose $\exists$ is a quantifier on a classical first order language satisfying the axioms for an average, such that the range of $\exists$ (i.e. the resource) satisfies classical logic. Then the first order language extended with $\exists$ satisfies classical logic.*

Proof. Lemma 4 shows that the Lindenbaum algebra of the logic can be obtained as the direct product of two Boolean algebras, namely the resource and the first order logic. □

So much for the easy results. In the rest of the section we consider examples of resources weaker than Boolean algebras.

8.1 Positive information

In section 3 we gave an example where the resource $\mathcal{G}$ is the co-Heyting algebra generated by positive formulas. This example was motivated by the consideration that there sometimes exists an asymmetry between positive and negative information, where only the former is 'effective'. A concrete case of this occurs in perception, in particular in the interaction between knowledge and perception. Let the predicate A denote a set of $2D$ objects, B the predicate 'has mass', $\mathcal{G}_s$ a collection of tests comparing images to $3D$ templates at level of accuracy s; e.g. as in Marr [12], Chapter 5. Then even though $A \cap B = \emptyset$, one may have $\exists(A|\mathcal{G}_s) \cap B \neq \emptyset$; in fact, this is just the visual illusion which allows us to enjoy (?) television. Formally, what is at issue here is that we do not always want to require adjointness with respect to B^c.

We now investigate this case in slightly more detail. As will be seen, it is the only example where the conditioning structure naturally has the right completeness properties and where in addition $\exists(\varphi|\mathcal{G})$ can be said to be computable from φ and $\mathcal{G}$.

Let $\mathcal{M}$ be a model, $\mathcal{G}_0$ the distributive lattice of (sets of assignments on $\mathcal{M}$ definable by) positive formulas. Let $\mathcal{G}$ be the lattice obtained by closing $\mathcal{G}_0$ under countable conjunctions, then since $\mathcal{G}$ satisfies the infinite distributive law $\bigwedge_{i\in I}(a_i \vee b) = \bigwedge_{i\in I} a_i \vee b$, $\mathcal{G}$ is actually a co-Heyting algebra. We may now define a quantifier $\exists(\cdot|\mathcal{G})$ by putting $\exists(\varphi|\mathcal{G}) = \bigwedge\{\psi \in \mathcal{G} \mid \varphi \subseteq \psi\}$, for φ in the language $\mathcal{L}(CQ(\mathcal{G}))$. Although the definition is similar to what we have seen before, with this choice of resource it stands to reason that

$$\varphi \wedge \exists(\psi|\mathcal{G}) = \mathbf{0} \Longrightarrow \exists(\varphi|\mathcal{G}) \wedge \exists(\psi|\mathcal{G}) = \mathbf{0}$$

must fail, because it expresses that the estimate $\exists(\varphi|\mathcal{G})$ is affected by the negative information $\neg\exists(\psi|\mathcal{G})$, which need not be in $\mathcal{G}$. For suitable $\mathcal{G}$ this can indeed be shown, and it follows that the Frobenius property fails. The question of interest is then whether there exists a complete axiomatisation of conditional quantifiers with co-Heyting algebras as range.

Definition 18 *The language $\mathcal{L}(CQ(\text{co-}H))$ has propositional constants $\top, \bot$, connectives $\wedge, \vee, \neg$ and quantifiers $\exists, \forall, \exists(\cdot|\mathcal{G})$. The formation rules are as usual, except that $\neg\varphi$ is defined only for formulas φ not containing a conditional quantifier.*

Definition 19 *A* standard co-Heyting model *is a structure $\langle \mathcal{M}, \mathcal{G}_0, \mathcal{G}, \rangle$, where $\mathcal{M}$ is a first order model and $\mathcal{G}$ is the complete co-Heyting algebra of sets of assignments generated by a homogeneous set of first order formulas $\mathcal{G}_0$. The interpretation of $\backslash$ on $\mathcal{M}$ should satisfy $\mathcal{M} \models \varphi\backslash\psi[f]$ iff for all θ, if $\mathcal{M} \models \varphi[f] \Rightarrow \mathcal{M} \models \psi \vee \theta[f]$, then $\mathcal{M} \models \theta[f]$.*
The interpretation of $\exists(\varphi|\mathcal{H})$ on $\langle \mathcal{M}, \mathcal{G}_0, \mathcal{G}, \rangle$ is given by

1. *The name $\mathcal{H}$ is interpreted by the co-Heyting algebra $\mathcal{G}$.*
2. *$\langle \mathcal{M}, \mathcal{G}_0, \mathcal{G}, \rangle \models \exists(\varphi|\mathcal{H})[f]$ iff for all $A \in \mathcal{G}$, $\varphi \leq A$ implies $f \in A$. .*

Observe that in this case we do not make a distinction between the resource and the range of the quantifier; this choice will be justified by the completeness theorem below. Standard co-Heyting models can easily be constructed starting from an ω_1-saturated first order model $\mathcal{M}$ and a homogeneous set of first order formulas $\mathcal{G}_0$. Let $\mathcal{H}$ be the complete co-Heyting algebra generated by the sets of assignments definable on $\mathcal{M}$ by first order formulas, $\mathcal{G}$ the complete co-Heyting subalgebra generated by $\mathcal{G}_0$. By ω_1-saturation, $\mathcal{H}$ is closed under $\exists$. We may therefore define $\setminus$ explicitly by As before, the proof calculus consists of a basic system together with an infinitary measurability rule.

The basic sequent calculus $\mathcal{A}(CQ(\text{co-H}))$ comprises

1. The classical rules for $\wedge, \vee, \forall, \exists, \top, \bot$.

2. The classical rules for $\neg$, applicable only to formulas not containing $\exists(\cdot|\mathcal{G})$.

3. The classical structural rules, including cut.

4. $(lCQ(\text{co-H}))$ $\dfrac{\Gamma, \varphi(x) \Longrightarrow \Delta}{\Gamma, \exists(\varphi|\mathcal{G})(x) \Longrightarrow \Delta}$,

   where x does not occur free in Γ and Δ consists of $\exists(\cdot|\mathcal{G})$-quantified formulas, or conjunctions of these. The deviant condition on Δ is necessary to enforce that the range of the conditional quantifier will be a lattice. Note that in this case we definitely could not allow formulas of the form v to be present in Γ, as this would make the Frobenius property derivable.

5. $(rCQ(\text{co-H}))$ $\dfrac{\Gamma \Longrightarrow \varphi, \Delta}{\Gamma \Longrightarrow \exists(\varphi|\mathcal{G}), \Delta}$.

Trivially one has

Lemma 20 *$\mathcal{A}(CQ(co\text{-}H))$ is sound with respect to standard co-Heyting models.*

The basic sequent calculus leads to a weak completeness theorem, in the sense that the resource will be a distributive lattice. However, in order for the resource to be a complete co-Heyting algebra, we need an infinitary rule.

Definition 20 *For a set $\mathcal{G}_0$ of first order formulas, $(McoH)_{G_0}$ is the following infinitary rule*

$$\frac{\{\Gamma, \exists\overline{x}(\theta(\overline{x}) \wedge \neg\psi(\overline{x})) \Longrightarrow \exists\overline{x}(\varphi(\overline{x}) \wedge \neg\psi(\overline{x})) \mid \psi \in \mathcal{G}_0\}}{\Gamma, \exists(\theta|\mathcal{G}) \Longrightarrow \exists(\varphi|\mathcal{G})}.$$

The idea behind the rule is simple: given the intended interpretation of $\exists(\cdot|\mathcal{G})$, if for all $\psi \in \mathcal{G}_0$, $\varphi \leq \psi$ implies $\theta \leq \psi$, then $\exists(\theta|\mathcal{G}) \leq \exists(\varphi|\mathcal{G})$. The complicated formulation arises from the fact that we do not have implication in our language. We thus have

Lemma 21 *$(McoH)_{G_0}$ is sound with respect to standard co-Heyting models.*

We are now in a position to prove

Theorem 10 $\mathcal{A}(CQ(co\text{-}H)) + (McoH)_{G_0}$ *is complete with respect to standard co-Heyting models.*

PROOF. Suppose $\Lambda \Longrightarrow \Delta$ is not derivable. Construct a maximal filter F containing Λ such that $\bigvee \Delta \notin F$. Enumerate the formulas such that each formula occurs infinitely often; as usual, we construct F in stages F_n ensuring the various closure properties; but since there is no negation applicable to all formulas, we have to increase Δ as well. We do the case corresponding to rule $(McoH)_{G_0}$. Suppose F_n, Δ_n have been constructed such that $F_n \Longrightarrow \Delta_n$ is not derivable, and suppose that the formula to be treated is $\exists(\theta|\mathcal{G}) \in F_n$. Consider a formula $\exists(\varphi|\mathcal{G})$ with number smaller than that of $\exists(\theta|\mathcal{G})$. If

$$F_n, \exists(\varphi|\mathcal{G}) \Longrightarrow \Delta_n$$

is not derivable, put $F_{n+1} := F_n \cup \{\exists(\varphi|\mathcal{G})\}$ and $\Delta_{n+1}; = \Delta_n$. If

$$F_n, \exists(\varphi|\mathcal{G}) \Longrightarrow \Delta_n$$

is derivable, then there must exist $\psi \in \mathcal{G}_0$ such that

$$F_n, \exists\overline{x}(\theta \wedge \neg\psi) \Longrightarrow \exists\overline{x}(\varphi \wedge \neg\psi), \Delta_n$$

is not derivable. Indeed, if all these sequents were derivable, then one would have the following derivation

$$\cfrac{\cfrac{\{F_n, \exists\overline{x}(\theta \wedge \neg\psi) \Longrightarrow \exists\overline{x}(\varphi \wedge \neg\psi), \Delta_n \mid \psi \in \mathcal{G}_0\}}{F_n, \exists(\theta|\mathcal{G}) \Longrightarrow \exists(\varphi|\mathcal{G}), \Delta_n}}{F_n \Longrightarrow \exists(\varphi|\mathcal{G}), \Delta_n} \qquad ,$$

from which it follows that $F_n \Longrightarrow \Delta_n$ by cut; a contradiction. For the ψ found above, define

$$F_{n+1} := F_n \cup \{\exists\overline{x}(\theta \wedge \neg\psi)\},\ \Delta_{n+1} := \Delta_n \cup \{\exists\overline{x}(\varphi \wedge \neg\psi)\},$$

and proceed.
One then shows that for any such maximal F, for all first order τ, either $\tau \in F$ or $\neg\tau \in F$, and that F is closed under Henkin witnesses. Let $\mathcal{M}$ be the model constructed from F. Let $\mathcal{G}_0$ be as in the infinitary rule, and let $\mathcal{G}$ be the complete co-Heyting algebra generated by $\mathcal{G}_0$. To ensure that the model $\mathcal{M}$ satisfies $\exists(\varphi|\mathcal{G}) = \bigwedge\{\psi \in \mathcal{G}_0 \mid \mathcal{M} \models \forall x(\varphi(x) \to \psi(x))\}$, it suffices to show that we can take $\mathcal{M}$ to omit the type $\Sigma = \{\psi(\overline{x}) \in \mathcal{G}_0 \mid \exists\overline{x}(\varphi(\overline{x}) \wedge \neg\psi(\overline{x})) \notin F\} \cup \{\neg\exists(\varphi|\mathcal{G})\}$. But the rule $(McoH)_{G_0}$ implies that Σ is locally omitted by F, and we are done.

□

8.2 Bi-Heyting algebras

Our ostensible motivation in studying co-Heyting algebras as resources was that they lead to the failure of the property $p \wedge \exists q = \mathbf{0} \;\; \Rightarrow \;\; \exists p \wedge \exists q = \mathbf{0}$, so that the estimate $\exists p$ of p need not be influenced by negative information. It thus seemed natural to take the co-Heyting algebra generated by positive first order formulas. However, one has to be careful here, for if we close the finitary positive formulas under both $\bigwedge$ and $\bigvee$, we obtain a complete bi-Heyting algebra:

Definition 21 *A bi-Heyting algebra is an algebra that is both Heyting and co-Heyting. A complete bi-Heyting lattice is a lattice closed under arbitrary suprema and infima such that the following distributive laws hold*

$$x \vee \bigwedge_i y_i = \bigwedge_i (x \vee y_i), \;\; x \wedge \bigvee_i y_i = \bigvee_i (x \wedge y_i).$$

An interesting feature of bi-Heyting algebras is that they have two negations: ordinary intuitionistic negation $\neg a = a \to \bot$, and the *supplement* $\sim a = \top \backslash a$. The latter satisfies $a \vee \sim a = \top$, but not necessarily $a \wedge \sim a = \bot$.

Lemma 22 *A complete lattice satisfying the distributive laws above is a bi-Heyting algebra.*

Proof. Define $a \to b$ by $\bigvee\{c \mid a \wedge c \leq b\}$ and $a \backslash b$ by $\bigwedge\{c \mid a \leq b \vee c\}$. The distributive laws ensure that these operations have the right properties. $\Box$

Now if the resource is bi-Heyting, so is the resulting logic of the conditional quantifier, hence we also have an implication here. However, implication is sufficient to derive the Frobenius property, whence also the undesired influence of negative information. Positiveness must therefore be taken in a rather strict sense.

We remark in passing that there exist more natural examples of bi-Heyting algebras than the one just given. Suppose we allow partial predicates on models, i.e. predicates A which have positive and negative extensions A^+ and A^- such that $A^+ \cap A^- = \emptyset$, but not necessarily $A^+ \cup A^- = |\mathcal{M}|$. The logic of partial predicates is provided by Kleene's three-valued logic. Now using strong Kleene negation gives us implication, whereas the weak negation also gives us subtraction, thus yielding a bi-Heyting algebra.
Furthermore, in a topos the set of subobjects of a given object forms a complete bi-Heyting algebra, so that here the theory of conditional quantifiers works very smoothly. However, it would go against our motivation to fix one particular type of resource for its technical elegance, instead of considering a variety of resources reflecting degrees of incomplete knowledge.

9 Martingales, nonmonotonicity and an open problem

In logic a variable is either free or bound, with nothing in between. One idea which probability theory may have to offer to logic is that there is actually a continuous scale between free and bound variables; for in probability we have not only the random variable $X : \Omega \longrightarrow I\!R$ and an integral $\int XdP = \mathbf{E}(X|\{\emptyset, \Omega\})$ which is a real number, but also conditional expectations $\mathbf{E}(X|\mathcal{G})$, which are nonconstant functions if $\mathcal{G} \neq \{\emptyset, \Omega\}$.

Now suppose we have a formula $\varphi(x)$ with only the variable x free. Existentially binding x means applying the conditional quantifier $\exists(\cdot|\{\mathbf{0}, \mathbf{1}\})$ to φ; semantically this expresses that we have no information about x. Replacing $\{\mathbf{0}, \mathbf{1}\}$ by an algebra of formulas in some of which x occurs free now expresses that we have *some* information about x. If we replace $\{\mathbf{0}, \mathbf{1}\}$ by the algebra $\mathcal{G}$ of all formulas containing x free, we indicate by this that we have full information about x, and accordingly $\exists(\varphi(x)|\mathcal{G}) = \varphi(x)$. Summarising, we could say that we have replaced the two possibilities 'free/bound' by a continuum of possibilities, where 'less information about x' corresponds to 'x is bound stronger'. This continuum of possibilities suggests that we consider notions of limit that might be applicable here. The most useful one appears to be the logical analogue of the martingales that were introduced in section 2. In the following we consider conditional quantifiers, hence averages are for the moment excluded.

Definition 22 *Let $\langle T, \geq \rangle$ be a directed set, i.e. $\geq$ is a partial order on T such that for all $s, t \in T$ there exists $r \in T$ such that $r \geq s, t$. A family of conditional quantifiers $\{\exists(\cdot|\mathcal{G}_s) \mid s \in T\}$ is a* martingale *if for all formulas φ, $\exists(\varphi|\mathcal{G}_s) \leq \exists(\varphi|\mathcal{G}_t)$.*

With this definition one trivially has

Lemma 23 *1. $\{\exists(\cdot|\mathcal{G}_s) \mid s \in T\}$ is a martingale if $s \geq t$ implies $\mathcal{G}_s \supseteq \mathcal{G}_t$.*

2. If $\{\exists(\cdot|\mathcal{G}_s) \mid s \in T\}$ is a martingale, then $\exists(\exists(\varphi|\mathcal{G}_s)|\mathcal{G}_t) = \exists(\varphi|\mathcal{G}_t)$.

□

Observe that 2 is the defining condition for martingales in probability theory. By analogy, we may now ask whether there exists a suitable notion of limit here. Now clearly, if $\{\exists(\cdot|\mathcal{G}_s) \mid s \in T\}$ is a martingale, $\lim_{s\in T} \exists(\varphi|\mathcal{G}_s)(f) = 1$ iff $\exists s \forall t \geq s \exists(\varphi|\mathcal{G}_t)(f) = 1$ iff $\forall t \exists(\varphi|\mathcal{G}_t)(f) = 1$, so that the limit is simple intersection. The direction from right to left is trivial; for the converse direction, choose arbitrary r, then for $s \geq s_0 r$, $\exists(\varphi|\mathcal{G}_s)(f) = 1$, so that by the martingale property $\exists(\varphi|\mathcal{G}_r)(f) = 1$. Two questions now arise:

1. Is $\lim_{s\in T} \exists(\varphi|\mathcal{G}_s)$ itself a conditional quantifier?
2. When is $\lim_{s\in T} \exists(\varphi|\mathcal{G}_s)$ equal to φ?

These questions are not easy to answer. For the case that the quantifiers are generated by equivalence relations, the first question has a positive answer, at least on suitably saturated models. This is the content of Wright's theorem 3 in [24].[10] For general models the question is open. On the other hand, if the

[10] The saturation comes in because Wright proves his result via Stone representation theory for Boolean algebras.

resources are co-Heyting algebras, then the answer is affirmative. The second question is even more interesting, and to explain why, we will make an excursion to nonmonotonic logic.

In Reiter's version of default logic ([13]), a *default* is a rule of the form

$$\alpha : \beta_1, \ldots, \beta_n / \omega$$

where α is the *prerequisite* of the rule, $\beta_1, \ldots, \beta_n$ are its *justifications*, and ω is its *consequent*. The customary interpretation of the rule is: 'if α has been derived from the bakcground knowledge and $\beta_1, \ldots, \beta_n$ are consistent with what has been derived, conclude ω'.

A *normal* default is one in which there is a single justification which is identical to the consequent; this is the kind of default of interest to us. Normally, defaults are used to express rules with exceptions, such as 'Birds fly', formalised as

$$B(a) : F(a)/F(a),$$

for every constant a. A default theory consists of a set of facts and a set of default rules. The facts ('Tweety is a bird') are taken to be specific and reliable information, and the defaults represent general information.

We will use default rules to express a slightly different concept. For the following discussion a martingale $\{\exists(\cdot|\mathcal{G}_s) \mid s \in T\}$ is assumed to be given. We want to consider a default rule of the form

$$\exists(A(x)|\mathcal{G}_s) : \exists(A(x)|\mathcal{G}_t)/\exists(A(x)|\mathcal{G}_t) \text{ (for } t > s),$$

which says that if evidence at stage s derives $\exists(A(x)|\mathcal{G}_s)$, and if $\exists(A(x)|\mathcal{G}_t)$ is consistent with the evidence (where $t > s$), then assume $\exists(A(x)|\mathcal{G}_t)$.[11] Note the subtle difference with Reiter's default rules: here the specific information consists of *observational* judgments, which are always approximate, hence defeasible. The default rules now express the expectation that the judgment will continue to be true in more refined approximations. This 'evidential' interpretation also leads to a formal difference.

We assumed that observations on variables are always performed with finite accuracy. At any given time, there will be a maximum accuracy s with which observations can be performed. Intuitively this suggests that for $t > s$, $\exists(A(x)|\mathcal{G}_t)$ will be consistent with the evidence; for if not, we would seem to have evidence at accuracy level t, which is impossible. This intuitive argument can be given formal status when we consider quantifiers conditional on Boolean algebras, but in general not when the quantifier is conditional on a co-Heyting algebra. Here is the argument.

Suppose the information about X at stage s is summarised by $\exists(\psi|\mathcal{G}_s)$ where $\exists(\psi|\mathcal{G}_s) \neq \mathbf{0}$. We now make an additional observation $\exists(\varphi|\mathcal{G}_s)$. We may assume that $\exists(\psi|\mathcal{G}_s) \wedge \exists(\varphi|\mathcal{G}_s) \neq \mathbf{0}$. By monotonicity $\exists(\exists(\varphi|\mathcal{G}_t)|\mathcal{G}_s) \geq \exists(\varphi|\mathcal{G}_s)$. Now

[11] Default rules consisting of open formulas would not be allowed in Reiter's set up, but here they are.

suppose that $\exists(\psi|\mathcal{G}_s) \wedge \exists(\varphi|\mathcal{G}_t) = \mathbf{0}$. Then also $\exists(\exists(\psi|\mathcal{G}_s) \wedge \exists(\varphi|\mathcal{G}_t)|\mathcal{G}_s) = \mathbf{0}$, and by Frobenius, $\exists(\psi|\mathcal{G}_s) \wedge \exists(\exists(\varphi|\mathcal{G}_t)|\mathcal{G}_s) = \mathbf{0}$; whence $\exists(\psi|\mathcal{G}_s) \wedge \exists(\varphi|\mathcal{G}_s) = \mathbf{0}$, a contradiction. It is crucially important for this argument that the observation $\exists(\varphi|\mathcal{G}_s)$ is performed at level s. If the observation were less accurate, say $\exists(\varphi|\mathcal{G}_r)$ for $r < s$, then the assumption $\exists(\psi|\mathcal{G}_s) \wedge \exists(\varphi|\mathcal{G}_r) \neq \mathbf{0}$ does not contradict $\exists(\psi|\mathcal{G}_s) \wedge \exists(\varphi|\mathcal{G}_s) = \mathbf{0}$. In this sense the observation has to be maximally informative relative to the background knowledge. Furthermore, since the Frobenius property is used, we need conditioning on Boolean algebras (or at least bi-Heyting algebras).

Motivated by these considerations we shall take a default to be a rule of the form

$$\exists(\varphi(x)|\mathcal{G}_s)/\exists(\varphi(x)|\mathcal{G}_t).$$

The rule should be interpreted as: 'if I have observed X at stage s and have ascertained that X satisfies φ, where s represents the maximum available accuracy, then I may assume that observing X at stage t will still yield that X satisfies φ.'[12]

Now the following natural question arises: suppose we have thus nonmonotonically concluded that for all t, $\exists(\varphi(x)|\mathcal{G}_t)$, what does this tell me about the real world? Can I conclude from this, albeit nonmonotonically, that φ? This is clearly question 2 above. The answer is generally 'no', although for positive formulas we can do better. This will be seen when we finally return to where we began, with the connection between granularity and conditional quantification.

Let $\mathcal{M}$ be the 'real world', and let $\{\mathcal{M}_s \mid s \in T\}$ be a set of simple, for instance finite, coarse-grained approximations to $\mathcal{M}$. We may assume that the set $\{\mathcal{M}_s \mid s \in T\}$ has the structure of an *inverse system*:

Definition 23 *Let T be a set directed by a partial order $\geq$; i.e., $\geq$ is reflexive, transitive, anti-symmetric, and for $s, t \in T$ there is $r \in T$ such that $r \geq s, t$. An inverse system (indexed by T) is a structure $\langle \mathcal{M}_s, h_{st}\rangle_{s,t\in T}$ with*

1. *for each $s \in T$, $\mathcal{M}_s$ is a structure for the signature σ_s;*
2. *for any R in the union of the signatures there is $t \in T$ such that R is in σ_s if $s \geq t$.*
3. *for each $s, t \in T$ with $s \geq t$ there is a homomorphism $h_{st} : |\mathcal{M}_s| \longrightarrow |\mathcal{M}_t|$, satisfying for each R in $\sigma_t \cap \sigma_s$*

$$\{\langle h_{st}(d_1), \ldots, h_{st}(d_n)\rangle : \langle d_1, \ldots, d_n\rangle \in R_s\} \subseteq R_t;$$

, where $R_s(R_t)$ is the interpretation of R on $\mathcal{M}_s(\mathcal{M}_t)$[13]

[12] The difference between Reiter's interpretation and ours is that we take the consistency of the justification to be relative to a stage s, whereas in Reiter's case it refers to an extension of the default theory. Here, we shall forego a discussion of the possible notions of extensions applicable in this context; see [22].

[13] It is mainly for conceptual reasons that we allow the signatures of the models to vary, since we may wish to say that a predicate is not yet applicable at a certain stage.

4. h_{rr} *is the identity on* $\mathcal{M}_s$*, and for* $s \geq t \geq r$*,* $h_{sr} = h_{st} \circ h_{tr}$*.*

Requiring that T is directed is tantamount to assuming that the collection of coarse-grained approximations is consistent, in the sense that any two approximations $\mathcal{M}_s, \mathcal{M}_t$ can be 'fused' into an approximation $\mathcal{M}_r$ which projects homomorphically onto both $\mathcal{M}_s$ and $\mathcal{M}_t$. Another way of viewing the consistency requirement is that (under an additional topological restriction) an inverse system has an inverse limit $\mathcal{M}$ such that $\mathcal{M}$ is the smallest structure which projects homomorphically onto all the $\mathcal{M}_s$.

Definition 24 *Let* $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$ *be an inverse system. Its inverse limit*

$$\mathcal{M} := \lim_{\leftarrow T} \mathcal{M}_t$$

is defined as follows

1. *the domain* $|\mathcal{M}|$ *consist of the* threads *in the product* $\Pi_{t \in T} |\mathcal{M}_t|$*; i.e., functions* $\xi : T \longrightarrow \bigcup_{t \in T} |\mathcal{M}_t|$ *satisfying:* $\xi_t \in |\mathcal{M}_t|$*, and* $h_{st}(\xi_s) = \xi_t$ *for* $s \geq t$*.*

2. *the interpretation of the predicates is given by: for each* R *there exists* $t \in T$ *such that for all threads* $\xi^1, \ldots, \xi^n$

$$R(\xi^1, \ldots, \xi^n) \Longleftrightarrow \forall s \geq t : R_s(\xi^1_s, \ldots, \xi^n_s)$$

The inverse limit $\mathcal{M}$ is a submodel of the direct product $\Pi_{t \in T} \mathcal{M}_t$; however, the domain of this submodel might be empty. Under the additional assumption that the $\mathcal{M}_s$ are finite this cannot be so. The proof rests on the fact that the discrete topology on the $\mathcal{M}_s$ is compact and makes the bonding mappings continuous.

Theorem 11 *Suppose* $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$ *is an inverse system of finite models. Then* $|\mathcal{M}|$ *is non-empty.*

Now let $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$ be an inverse system with inverse limit $\mathcal{M}$. We now associate a family of conditional quantifiers to the $\mathcal{M}_s$. Let $\mathcal{B}_s$ denote the algebra of first-order definable sets of assignments on $\mathcal{M}_s$, and let $\pi_s : \mathcal{F}_{\mathcal{M}} \longrightarrow \mathcal{F}_{\mathcal{M}_s}$ denote the projection from the assignment space of $\mathcal{M}$ to that of $\mathcal{M}_s$. $\pi^{-1}(\mathcal{B}_s)$ is an algebra on $\mathcal{F}_{\mathcal{M}}$, but not yet quite the one we want, since it does not take account of the information in models $\mathcal{M}_t$ where $t \leq s$. Therefore we define $\mathcal{G}_s = \bigcup_{t \leq s} \pi_t^{-1}(\mathcal{B}_t)$, an algebra which does take account of what has gone on 'before'. For simplicity assume that the quantifier $\exists(\cdot|\mathcal{G}_s)$ is generated by the equivalence relation determined by $\mathcal{G}_s$. If $s \geq t$, then it is clear that the family $\{\exists(\cdot|\mathcal{G}_s) \mid s \in T\}$ is a martingale, since $s \geq t$ implies by construction $\mathcal{G}_t \subseteq \mathcal{G}_s$. As observed above, if T is directed, $\lim_{s \in T} \exists(\varphi|\mathcal{G}_s)(f) = 1$ iff $\exists s_0 \forall s \geq s_0 \exists(\varphi|\mathcal{G}_s)(f) = 1$ iff $\bigwedge_{s \in T} \exists(\varphi|\mathcal{G}_s) = 1$.

We are now interested in cases where we have a strong form of convergence: $\varphi = \lim_{s \in T} \exists(\varphi|\mathcal{G}_s)(f) = \bigwedge_{s \in T} \exists(\varphi|\mathcal{G}_s)$, where the $\mathcal{G}_s$ are Boolean algebras.

Lemma 24 *Suppose $\mathcal{M}$ is an inverse limit of an inverse system whose models $\mathcal{M}_s$ are finite. Then the associated martingale $\{\exists(\cdot|\mathcal{G}_s) \mid s \in T\}$ converges on positive formulas.*

PROOFSKETCH. Using compactness and the fact that positive formulas are preserved by homomorphism, one shows that for positive φ, if for all s, $\mathcal{M}_s \models \varphi[\pi_s(f)]$, then $\mathcal{M} \models \varphi[f]$. For positive formulas we have the relationship $\exists(\varphi|\mathcal{G}_s) \subseteq \pi_s^{-1}\{g \in \mathcal{F}_{\mathcal{M}_s} \mid \mathcal{M}_s \models \varphi[g]\}$. It then follows that $\bigwedge_{s\in T} \exists(\varphi|\mathcal{G}_s) \leq \bigwedge_{s\in T} \pi_s^{-1}\{g \in \mathcal{F}_{\mathcal{M}_s} \mid \mathcal{M}_s \models \varphi[g]\} \leq \varphi$. The converse direction follows from $\varphi \leq \exists(\varphi|\mathcal{G}_s)$. □

The proof uses essentially that the $\mathcal{G}_s$ generate a compact totally disconnected topology on $\mathcal{M}$. Without the compactness assumption, we only have that convergence holds for universal formulas. More importantly, however, convergence may fail even for negations of predicates, cf. [21]. This can be verified for example on $\langle 2^\omega, A\rangle$, written as the limit of the inverse sequence of models $\langle\{0,1\}^n, A\rangle$, such that the interpretation of A on 2^ω is a single infinite branch. We will now pause to connect the above results once more to perception.

9.1 Digression: illusory conjunctions continued

Treisman and Gelade [17, p. 113] remark that there are two ways in which search tasks can be difficult. (1) Targets and distractors may be difficult to discriminate, and so require serial fixations with foveal vision; this may happen both for single features and for conjunctions. (2) In case there are two or more items present, identification of conjunctions requires focussed attention which serially scans each item. The second form of search was discussed in section 7. Here we briefly treat the first form.

In [21], [19] and [20] it is argued at length, following Marr [12], that a formal model of the higher stages of visual perception is best clad in the form of an inverse system of first order structures, so that $s \geq t$ means that $\mathcal{M}_s$ is a more refined description than is $\mathcal{M}_t$; equivalently, that $\exists(\cdot|\mathcal{G}_s)$ is a finer-grained filter than is $\exists(\cdot|\mathcal{G}_t)$. The visual analogue of moving from s to t would thus be changing foveal acuity, which allows one to see sharper. The following result then gives a logical analogue of Treisman and Gelade's first search procedure in the case of a conjunction.

Lemma 25 *In the situation outlined above, assuming compactness, for all $s \in T$ there exists $t \geq s$ such that for positive φ, ψ, $\exists(\varphi|\mathcal{G}_t) \wedge \exists(\psi|\mathcal{G}t) \leq \exists(\varphi \wedge \psi|\mathcal{G}_s)$.*

PROOF. We have $\varphi \wedge \psi = \bigwedge_t \exists(\varphi|\mathcal{G}_t) \wedge \exists(\psi|\mathcal{G}t) \leq \exists(\varphi \wedge \psi|\mathcal{G}_s) \leq C$, for any $C \in \mathcal{G}_s$ such that $\varphi \wedge \psi \leq C$. Since the $\exists(\varphi|\mathcal{G}_t) \wedge \exists(\psi|\mathcal{G}t)$ are closed and C is clopen, the result follows by compactness. □

The result says that parallel detection of features, if sufficiently accurate, may give a reliable indication of the presence of a true conjunction, thus providing a second route to the elimination of illusory conjunctions. □

We now return to the logical questions raised by martingale convergence. One might conjecture that on inverse systems of compact models, martingale convergence holds at most for positive formulas. This is not true, since convergence holds as well for the following class of formulas

Definition 25 *A formula $\varphi(x)$ is* locally positive *if there exist sentences $\psi_1, \ldots, \psi_n$ satisfying $\models \psi_1 \vee \ldots \vee \psi_n$ and for $i \neq j$, $\models \neg(\psi_i \wedge \psi_j)$, and positive formulas $\theta_1, \ldots, \theta_n$ such $\models \forall x(\varphi(x) \leftrightarrow \bigwedge_{i \leq n}(\psi_i \wedge \theta_i))$. The notion* locally universal *is defined analogously.*

The formula $(\psi \wedge A(x)) \vee (\neg\psi \wedge B(x))$, where ψ is the formula 'there exist at least 2 elements', is an example of a formula which is locally positive, but not positive.

Lemma 26 *φ is locally positive (universal) if there exist finitely many non-equivalent positive (universal) θ_i such that $\models \bigvee_i \forall x(\varphi(x) \leftrightarrow \theta_i)$.*

By a compactness argument one then shows that a formula φ is locally positive (universal) iff for each model $\mathcal{M}$ there exists a positive formula θ such that φ is equivalent to θ on $\mathcal{M}$.

Lemma 27 *Under the same assumptions as above, the martingale $\{\exists(\cdot|\mathcal{G}_s) \mid s \in T\}$ converges on locally positive formulas. Without compactness we have convergence on locally universal formulas.*

Proof. This is immediate from lemma 26. □

We therefore conjecture, for the general case without compactness assumptions

Conjecture 1 *If all martingales derived from an inverse system converge on φ, then φ is equivalent to a locally universal formula.*

References

[1] N. Alechina and M. van Lambalgen. Generalized quantification as substructural logic. *Journal of Symbolic Logic*, 4(3):177–189, 1996.

[2] J. Barwise. Scenes and other situations. *Journal of Philosophy*, LXXVIII:369–397, 1981.

[3] G. Birkhoff. Lattices in applied mathematics. In *Lattice theory*, pages 155–184. American Mathematical Society, Providence RI, 1961. Proceedings of Symposia in Pure Mathematics Vol. II.

[4] D.P. Ellerman and G.-C. Rota. A measure theoretic approach to logical quantification. *Rend. Sem. Mat. Univ. Padova*, 59:227–246, 1978.

[5] S. Graf. On the existence of strong liftings in second countable topological spaces. *Pacific Journal of Mathematics*, 58(2):419–426, 1975.

[6] P. R. Halmos. Algebraic logic, I: Monadic boolean algebras. *Compositio Mathematica*, 12:217–249, 1952.

[7] A. Herskovits. Language, spatial cognition and vision. In O. Stock, editor, *Spatial and Temporal Reasoning*, pages 155–202. D. Reidel Publishing Company, Dordrecht, 1997.

[8] J.R. Hobbs. Granularity. In *Proceedings IJCAI85*, pages 1–4. AAAI, 1985.

[9] A.S. Kechris. *Classical descriptive set theory.* Graduate texts in mathematics. Springer-Verlag, New York, 1995.

[10] H.J. Keisler. Probability logics. In J. Barwise and S. Feferman, editors, *Model-theoretic logics*, chapter 16, pages 579–596. Springer-Verlag, New York, 1985.

[11] F.W. Lawvere. Intrinsic co-Heyting boundaries and the Leibniz rule in certain toposes. In A. Carboni, M.C. Pedicchio, and G. Rosolini, editors, *Category Theory*, pages 267–297. Springer-Verlag, 1991. Lecture Notes in Mathematics 1488.

[12] D. Marr. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information.* W.H. Freeman, San Fransisco, 1982.

[13] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[14] C.I. Steinhorn. Borel structures and measure and category logics. In J. Barwise and S. Feferman, editors, *Model-theoretic logics*, chapter 16, pages 579–596. Springer-Verlag, New York, 1985.

[15] C.I. Steinhorn. Borel structures for first order and extended logics. In A. Sčedrov L.A. Harrington, M.D. Morley and S.G. Simpson, editors, *Harvey Friedman's research on the foundations of mathematics*, pages 161–178. North-Holland, Amsterdam, 1985.

[16] A.M. Treisman. Features and objects. *Quarterly Journal of Experimental Psychology*, 40A:201–237, 1988.

[17] A.M. Treisman and G. Gelade. A feature-integration theory of attention. *Clarendon Press*, 12:97–136, 1980.

[18] A.M. Treisman and H. Schmidt. Illusory conjunctions in the perception of objects. *Clarendon Press*, 14:107–141, 1982.

[19] J. van der Does and M. van Lambalgen. A logic of vision. Research report LP–97–07, ILLC, University of Amsterdam, 1997. To appear in *Linguistics and Philosophy.*

[20] M. van Lambalgen. Some logical constructions suggested by vision. In P. Blackburn and J. Seligman, editors, *Proceedings of ITALLC98*, Lecture Notes. CSLI Publications, Stanford, California, t.a.

[21] M. van Lambalgen and J. van der Does. A logic of vision: Preliminaries. In G.R. Renardel De Lavalette and C. Condaravdi, editors, *Proceedings of the Stanford LLC-workshop (prov. title)*, Lecture Notes. CSLI Publications, Stanford, California, t.a.

[22] F. Voorbraak. A nonmonotonic observation logic. Research report LP–97–09, ILLC, University of Amsterdam, 1997.

[23] D. Williams. *Probability with martingales.* Cambridge University Press, Cambridge, 1991.

[24] F. B. Wright. Convergence of quantifiers and martingales. *Pacific Journal of Mathematics*, 6:296–307, 1962.

[25] F. B. Wright. Boolean averages. *Canadian Journal of Mathematics*, 15:440–455, 1963.