institute for logic, language and computation

MICHIEL VAN LAMBALGEN, JAAP VAN DER DOES

A Logic of Vision

LP-96-14, *received*: November 1996

ILLC Research Report and Technical Notes Series
Series editor: Dick de Jongh

Logic, Philosophy and Linguistics (LP) Series, ISSN: 0928-3307

Institute for Logic, Language and Computation (ILLC)
University of Amsterdam
Plantage Muidergracht 24
NL-1018 TV Amsterdam
The Netherlands
e-mail: illc@fwi.uva.nl

A Logic of Vision

Michiel van Lambalgen Jaap van der Does*
{michiell, jvddoes}@fwi.uva.nl
WINS : University of Amsterdam
Plantage Muidergracht 24
NL–1018 TV Amsterdam
The Netherlands

November 11, 1996

> *Our mistake lies in supposing that things present themselves as they really are, names as they are written, people as photography and psychology give an unalterable picture of them. But in reality this is not at all what we ordinarily perceive. We see, we hear, we conceive the world in a lopsided fashion.*
>
> Marcel Proust, *La Fugitive*

Abstract

This article is on logical aspects of uncertain, unstable perceptual information. In a review of logics of direct perception reports we stress that their semantics is often based on perfect, irretractable 'points'. We argue that this way of modelling has to be replaced by a more principled one which takes the retraction of perceptual information seriously. To do so, our logic tries to stay close to the psychological models of perception developed in Marr 1982 and subsequent work. In such models perception is a multi-layered process. The different layers have filters of different gradation, which makes perception at each of them approximate. Indeed, our main tasks will be to formalise the layers and the ways in which they may refine each other, and to develop a logic in which description varies with such measures of refinement. Within such a framework, perception can be analyzed as consisting of a non-veridical, approximative core, which becomes veridical by our expectation that what is perceived will remain the case. We show in detail that this non-monotonic view on perception can be obtained by combining a particular kind of inverse limit for first order models with the notion of conditional quantification in van Lambalgen 1996.

*This research is part of the PIONIER-project 'Reasoning with Uncertainty' sponsored by the Netherlands Organization for Scientific Research (NWO) under grant PGS 22–262. See URL `http://www.wins.uva.nl/research/pion/index.html` for further information on the project.

1 Introduction

It often happens in life that we perceive a person (or an object) only partially, but are nonetheless able to reason about this person and to incorporate him or her in our schemes. Such a situation has been eloquently described by Marcel Proust in *La Fugitive*

> Had I been obliged to draw from memory a portrait of Mlle d'Eporcheville, to furnish a description of her, or even to recognise her in the street, I should have found it impossible. I had glimpsed her in profile, on the move, and she had struck me as being simple, pretty, tall and fair; I could not have said more. But all the reflexes of desire, of anxiety, of the mortal blow struck by the fear of not seeing her if my father took me away, all these things, associated with an image of which on the whole I knew nothing, and as to which it was enough that I knew it to be agreeable, already constituted a state of love. (Penguin Classics ed., p. 577-8)

At the other extreme we may have a situation in which we perceive an object in its entirety, but nonetheless misidentify it. Again Proust furnishes an example

> I opened the Figaro. What a bore! The main article had the same title as the article which I had sent to the paper and which had not appeared. But not merely the same title... why, here were several words that were absolutely identical. This was really too bad. I must write and complain. But it wasn't merely a few words, it was the whole thing, and there was my signature... It was my article that had appeared at last! But my brain which, even at that period, had begin to show signs of signs and to tire easily, continued for a moment longer to reason as though it had not understood that this was my article, like an old man who is obliged to complete a movement that he has begun even if it has become unnecessary, even if an unforeseen obstacle, in the face of which he ought at once to draw back, makes it dangerous. (*ibidem*, p. 579)

In this article, we aim to give a semantics for perception reports which is sufficiently flexible to accomodate these examples, as well as the more common cases in which we correctly infer the identity of objects on the basis of partial information. Our methodology is akin to that of Barwise's work on naked infinitive perception reports. He maintains that our intuitions about the semantics of perception are fairly clear, 'but to square with these intuitions, the model-theoretic machinery [of possible worlds semantics] became quite bizarre' (Barwise 1981). For this reason, he rather prefers a semantics which 'fits nicely with a realist philosophical-psychological account of perception to be found in the writings of Dretske, R.J. Hirst, and J.J. Gibson' (*ibidem*). In the same vein, we strive to develop a logical model of vision along the lines of Marr 1982, and show how it complies with the semantical data. This in contrast to a more instrumentalist approach, which would allow us to explain the semantic facts in purely abstract terms (van der Does 1991 is an example).

There is another methodological aspect of our approach worth mentioning. The dynamic turn in logical semantics for natural language was an eye-opener for all people working in the field. Yet, other areas of research, such as probability theory, have been 'dynamic' for a long time. The semantics of perception reports developed in this paper is based on conditional quantification, which resulted from a logical analysis of conditional expectation. As a consequence, the dynamics of this probabilistic notion is imported naturally into the semantics. Indeed, we shall make heavy use of it since retraction of information is crucial to perception. But before we show in section 4 how this can be, let us first review some basic facts and received opinions concerning the perception of objects and scenes.

2 Direct perception

This article is mainly concerned with description of direct perception. Linguistically, this can take several forms; e.g., the 'simple' perception in (1), the naked infinitive form in (2), the gerundive form in (3), and the complementised form in (4).

(1) Jack saw Sharon.

(2) Jack saw Sharon wash her face.

(3) Jack saw Sharon washing her face.

(4) Jack saw that Sharon washed her face.

There are subtle semantic differences between the perception reports, which have to do with the kinds of object perceived. Sentence (1) describes the perception of objects, (2) and (3) that of scenes, while (4) gives the content of what is seen. Also, (1-3) report on what is seen directly, whereas (4) may state a conclusion inferred from what is actually perceived. The difference between (2) and (3) is aspectual; (2) concerns a finished action, (3) an ongoing one.

These differences also manifests themselves in the restrictions placed on the semantic ingredients of perception reports: (i) the part of objective reality perceived, (ii) the perceiver's field of vision, and (iii) the semantic content of the perceptual description. Direct perception reports require an immediate link between fields of vision and semantical content. Since visual fields normally represent part of objective reality, this link often gives a close connection between reality and semantic content. But of course, such a connection is not necessary. This is shown by misperceptions and hallucinations.

(5) Macbeth saw a dagger move in front of his eyes.

In case (5) is true, Macbeth's immediate perceptions should truly involve a dagger moving in front of his eyes, even though this is fysically impossible. In what follows we concentrate on the perception of objects, as in (1), and on naked infinitive perception reports, as in (2).

2.1 Perceiving objects

Simple expression does not mean simplicity of semantics. Although 'Snow is white' is well-suited to explain one's theory of truth, the semantics of a mass term as 'snow' is still in many ways open. Analogously, sentences like (6a) are used in textbooks to show that transitive verbs can be interpreted as two-place relations; here 'to see' as 'S'.

(6) a. Jack saw Sharon.
b. $S(j, s)$.

But to do so one has to idealise from considerable semantic facts. The introduction has already made clear that Jack need not see much of Sharon to make (6) true. To perceive an object it is often sufficient to see one of its (prototypical) parts. In fact, the truth conditions are even weaker, in that parts of representations may go proxy for the objects themselves. Most people will only have seen enlivened celluloid or electro-magnetic representations of famous Sharon, rather than Sharon herself. The important point is that the weak conditions of truth make simple descriptions of perception highly fallible. An amusing example is furnished by the following children's story.

(7) As always, poor Jack was short of money. But today Fortune was at his side. Looking from a window of his parent's penthouse just above the 17th floor, his sharp eyes saw a dime lying on Main Street. He rushed downstairs, checking every now and then whether some lucky bastard would find it before him. A miracle happened... At the 11th floor, it turned out to be a quarter, at the 5th floor even a dollar! How great his disappointment, when out of breath at ground floor he noticed to have chased a trash can.

The story nicely illustrates another crucial aspect of perceiving objects. Perceptions and representations not only concern parts of objects, they also come with a certain granularity. In fact, different parts of a single field of vision may have varying grades of precision:

(8) While focussing on Sharon, Jack saw Maria vaguely out of the corner of his eye.

Change of granularity could induce change of truth value. We sometimes have to retract our descriptions if we come to see more of an object, or if the accuracy of the relevant representations and perceptions alter. In the next section we shall see that similar adjustments are typical of more complex perception reports. But even at this point we observe that a psychologically realistic semantics should contain two ingredients. Firstly, it should use models which refine each other in several ways (with infinite precision in the limit). Secondly, the semantics should formalise retractability of descriptions relative to a certain degree of precision. These ingredients are developed in section 3, and applied in section 4.

2.2 Naked infinitive perception reports

Ideally, naked infinitive reports (whence: 'NI reports') are used to report on the perception of reality. Indeed, in most semantics an NI report such as 'Jack saw Sharon walk' is true iff Jack stands in the 'see'-relation with a part of the world in which Sharon walks.[1] However, this approach makes perception highly factual, which can only be sustained under the perfect circumstance where our perceptual field truly represents reality. Normally the connection between reality and our fields of vision is rather uncertain. This looseness between reality on the one hand, and vision and semantics on the other makes NI reports retractable over time; cf. (9) which is perfectly in order.

(9) From a distance, Jack saw a dime fall on main street, but on coming closer he saw it was a dollar.

The intimate relationship between semantic content and reality is often used to explain the logical transparancy of NI reports. As soon as we take the uncertainty in this relation seriously, as in the present semantics, most of the transparancy is lost. To see this, let us revisit some well-known logical principles which have been considered to hold for NI reports.

2.2.1 Partial perception

One of the most basic non-inferences concerns the interplay between perception of objects and NI reports; (11) does not follow from (10).

(10) Jack saw Sharon, and Sharon winked.

(11) Jack saw Sharon wink.

There is a twofold explanation of this fact: either Sharon's action is not within Jack's visual field, or it is too subtle to be discerned by him. The two possibilities combined identify a range of vision with a coarsened part of reality.

2.2.2 Veridicality

Veridicality is the principle which allows us to conclude (13) from (12).

(12) Jack saw Sharon was her face.

(13) Sharon washed her face.

[1] See Barwise 1981, Higginbotham 1983, Kamp 1984, Asher and Bonevac 1985, Landman 1986, Asher and Bonevac 1987, Muskens 1989, van der Does 1991, Hendriks 1993, Koons 1996, among other people; also for a discussion of the logical principles.

Since (13) describes 'the world' rather than Jack's field of vision, veridicality does not comply with the acceptability of (9). This is also clear from (5) and (14) in which the NI complements highly dependent on the viewer's perspective.

(14) Jack saw Sharon walk left of Maria and Jill saw Maria walk left of Sharon.

Of course it does not follow from (14) that Sharon walked left and right of Maria, as veridicality would have it. Negative NI complements also block veridicality.

(15) Jack saw no girl wink.

(16) No girl winked.

In general, we shall argue that the principle of veridicality can only be sustained in idealised situations, which makes it a default rule rather than a logical inference.

2.2.3 Boolean connectives

Along similar lines, the validity of some of the other logical principles can be questioned.

Conjunction. On the whole, there is consensus concerning the equivalence of (17) and (18).

(17) Jack saw Sharon wink and Mary smile.

(18) Jack saw Sharon wink and Jack saw Mary smile.

It is indeed hard to find counterexamples to the transparancy of 'to see' for conjunctive NI complements. Perhaps the following is one, since (20) does not appear to be a consequence of (19)

(19) Jack saw an ant walk nearby and Jack saw a beetle walk at a distance.

(20) Jack saw an ant walk nearby and a beetle walk at a distance.

The invalidity should be due to our impossibility to focus on scenes at different distances at the same time. The conjunction in the premiss allows a short lapse of time to make both conjuncts true, but not so for the conjunction in the NI complement.

Disjunction. In case of the logical connective 'or' the question is whether (21) is equivalent to (22).

(21) Jack saw Sharon smile or stare.

(22) Jack saw Sharon smile or Jack saw stare.

There is a natural tendency to interpret the disjunction in (21) exclusively, which would block the inference from (21) to (22). But Grice has argued convincingly that this effect is pragmatic, not part of the semantics. And for inclusive disjunction (21) and (22) are equivalent. See also the discussion in example 6 below.

Negation. What about negative NI complements? In the literature, one normally takes (23) to imply (24) but not conversely.

(23) Jack saw Sharon not cry.

(24) Jack didn't see Sharon cry.

On the assumption that 'to see' denotes a relation between an object and a factual scene, we only have a non-trivial implication if there are negative facts of some sort. However, in some of us lives a Mr. X, who in a famous discussion on logical atomism with Bertrand Russell doubted the existence of such facts (Russell 1988, pp. 215–16). Mr. X would maintain that the logical form of the premiss has no negation within the scope of the perception operator. Besides, there are those who read the premiss as 'Jack saw Sharon refrain from crying', with 'to refrain from crying' the antonym of 'to cry' (Higginbotham 1983). Formally, this corresponds to introducing disjoint positive and negative extensions of a relation in the manner of Feferman 1984. But then we essentially stay within the realm of positive information, and keep Mr. X satisfied.

Denials in dialogues also ask for a special treatment of negative information in perception reports. Consider (25).

(25) 'Did you see that hawk there?'
'I saw *something*, but it was not a hawk.'

It makes perfect sense to retract the last sentence by saying 'No, you're right, it *is* a hawk'.

2.2.4 Quantifiers

Intuitions similar to those concerning veridicality and the connectives also influence our judgements on quantificational behaviour. Apart from scope phenomena *pur sang*, there is the question to what extent NI reports are epistemically neutral; i.e, a perceptual field does not alter the interpretation of quantifiers within the scope of 'to see'. If this is not so, someone's visual field will determine the extent in which quantifiers may be imported into or exported out of the scope of a perception verb.

It seems we have to following situation. If NI reports are epistemically neutral, quantifiers may be moved freely into and out of the scope of 'to see'. Then, (26) is equivalent with (27), (28) with (29), and (30) with (31).

(26) Jack saw a girl swim.

(27) A girl is such that Jack saw her swim.

(28) Jack saw no girl swim.

(29) No girl is such that Jack saw her swim.

(30) Jack saw every girl leave his party.

(31) Every girl is such that Jack saw her leave his party.

But if NI reports are not neutral in this sense, all the pairs are independent of each other. Jack may be too confused to perceive a 'real' girl as a girl, and what he perceives to be a girl need not be one.

This ends our discussion of the interaction between connectives, quantifiers, and the perception verb. We now give an informal sketch of the semantics developed in the remainder of our paper.

2.3 Sketch of the formal semantics

Logics for perception reports are often based on the assumption that the perceived objects are 'points' which cannot be refined any further; all partiality comes from their properties. However, in the previous sections we have seen many cases where it is more natural to assume that these objects are underdetermined. We therefore favour an approach where the infinitely precise points arise in the limit of increasingly refined stages. In the formalisation each stage is a first order model; and the properties of an object may vary with the stages, although we normally expect they do not. Despite such instability, it is possible to define an inverse limit, which we take to represent reality.

Reality is but one side of perception's coin, we also need linguistic means to describe it. Given the evidence that the dynamics of retraction is crucial to the semantics of perception reports, this cannot be just a 'static' logic. Instead, we shall use so-called conditional quantification (van Lambalgen 1996) which is rather different from the generalisation of quantifiers in Mostowski 1957 and Lindström 1966. Conditional quantification singles out the logical core of conditional expectation in probability theory. It offers a natural way to relativise quantification to varying measures of accuracy, and is hence well-suited for our purposes. To be more precise, we model a range of vision as a stage in the above limit construction, and show that quantification local to a visual field can be mimicked in the inverse limit. We prove moreover that background laws of perception hold in the inverse limit, if they hold at each stage. Within this framework, perception can be analyzed as consisting of a non-veridical, approximative core, which becomes veridical by our expectation that what is perceived will remain the case.

As we have said before, David Marr's theory of vision is the heuristic backdrop against which the formal theory is developed (Marr 1982). The next section starts with the essentials of his theory.

3 Vision, and a blurred view on logic

We have seen in section 2.2.2 that the principle of veridicality is an idealisation which does not allow for the retraction of perception reports. We believe that retraction is a very real phenomenon, and that any semantics for perception reports should account for this. Moreover, a semantics should also allow for partially perceived objects. This could possibly be achieved by introducing partial objects in the domain, but we favour a principled solution in which, roughly speaking, partial perception is the rule not the exception.

The semantics for perceptual expressions introduced here is characterised by the following features:

1) it is completely model theoretic in nature;
2) it tries to stay close to psychological models of perception;
3) it takes veridicality to be a defeasible principle which allows for the possibility to retract a perception report.

The reader might think there is a certain tension between 1) and 2), since typically the psychological models involve mathematical constructs such as Gaussians, Laplace operators etc., which one would not like to have in one's semantics. Indeed, it is incumbent on us to show that these psychological theories contain a model theoretic 'core' that is relevant to a semantics of perception. We believe that the two central notions here are 'inverse limit' and 'conditional quantification'; whether these indeed capture the semantically significant part of psychological modelling we must leave for the reader to judge. In any case, whatever the fate of this proposal, we are in agreement with Marr when he writes, criticising Gibson's 'realistic' approach

> The underlying point is that visual information processing is actually very complicated, and Gibson was not the only thinker who was misled by the apparent simplicity of the act of seeing. The whole tradition of philosophical inquiry seems not to have taken seriously enough the complexity of the information processing involved (Marr 1982, p. 30).

We maintain that, in order to explain the logic of perception, some of these complications have to be imported in the model theoretic machinery.

Inevitably, this section will be rather technical. We have tried to organise the presentation in such a manner that the main thrust of the argument can be followed also by those not willing to delve into the technicalities. Section 3.1 gives a rapid introduction to David Marr's theory of vision in so far as it is relevant to our concerns. In section 3.2 we extract from his work two model theoretic notions, that of an inverse system of models, and the inverse limit thereof. A model theoretic correlate of a third notion, that of a filter (in the sense of stochastic control, not in the familiar logical sense), is studied in section 3.3. The three notions are linked in section 3.4.

3.1 David Marr on vision

We start by explaining the psychological motivation underlying the model theory. Here, we base ourselves on an abstract account of Marr's theory of vision (1982). Of course, basing one's semantics on an empirical theory brings with it the danger that the empirical theory is wrong; indeed, it has been claimed that Marr's views are 'almost completely wrong' (Mayhew, as quoted in Boden 1988, p. 74). Nevertheless, we hope to convince the reader that Marr's theory is extremely suggestive from a model theoretic point of view. In particular Marr's idea of a hierarchy of three dimensional models has a good model theoretic correlate; 'good' in the sense that the associated preservation and non-preservation theorems may shed some light on the logic of perception. The model theory is sufficiently abstract to be compatible with other approaches based on the idea of a hierarchy of perceptual models, such as P.K. Allen's (1987).

Marr's fundamental idea is that vision is in many ways approximate. Filtering takes place at many of the earlier levels of visual processing, leading up to the so called primal sketch; and, at the other end, the perception of 3-D objects and scenes takes place by means of a hierarchy of ever more refined, but never completely accurate models.

Here we shall concentrate on the last stage of the visual process, although the proposed mathematical model describes the earlier stages as well.

Seeing a 3-D object involves two processes: constructing an image from visual data, and matching the image to a catalogue of 3-D models, where the matching is based on some salient features derived from the image. At this point we can do no better than reproduce the following illustration from Marr 1982, p. 306.

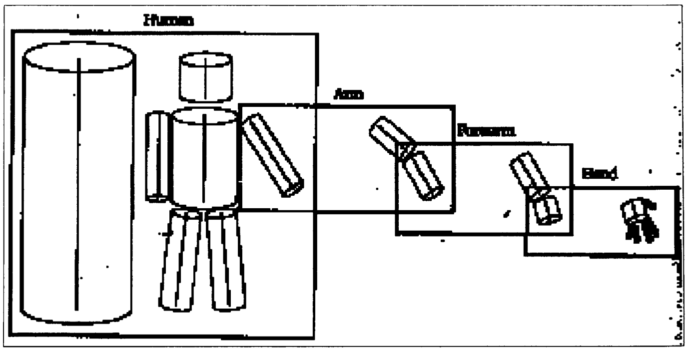

Refinement of an arm

What is depicted is an increasingly detailed series of models of an arm. Obviously, this series can be extended further by detailing the shape of the fingers, by replacing the cylinders by less rigid shapes etc. Marr's point is that we recognise an object in the real world in terms of these 3-D models, and that we may often use a rather rough approximation to correctly identify the object.

> Modularity [...] allows the representation to be used more flexibly in response to the needs of the moment. For example, it is easy to construct a 3-D model description of just the arm of a human shape that could later be included in a new 3-D model description of the whole human shape. Conversely, a rough but usable description of the human shape need not include an elaborate arm description. Finally, this form of modular organisation allows one to trade off scope against detail. This simplifies the computational processes that derive and use the representation, because even though a complete 3-D model may be very elaborate, only one 3-D model has to be dealt with at any time, and individual 3-D models have a limited and manageable complexity (Marr 1982, p. 307).

There exists an interplay between the clues derived from an image and the matching process (cf. Marr 1982, p. 321): after a 3-D model has been selected (guided by the image), it can be used to search for additional clues in the image; in turn, these can be used (when necessary) to match the image to a more detailed 3-D model. However, it may turn out to be impossible to find a more detailed 3-D model *of the kind we expected*. Indeed, like all computationally efficient heuristics, the use of such approximate models brings with it the possibility of error: what is identified as a real arm with respect to a given approximation may turn out to be something else (e.g., a wooden arm) when 'looking closer', i.e. with respect to a more refined approximation. (This point is not much emphasised in Marr 1982 though.) In any case a theory such as Marr's is well-suited to account for partial perception of an object: this is simply the matching of an object to a 3-D model without an expectation as to the direction in which the model can be refined. These observations suggest a formal semantics for visual reports in terms of approximate models and a stability condition. For instance, informally still, the expression

(32) I see an arm

can be taken to mean the conjuction of (i) and (ii).

i) with the present approximation the object that I focus on is identified as an arm;
ii) I expect this to be the case for every more refined approximation.

That is, the arm reported on in (32) is viewed as a (possibly infinite) series of ever more accurate approximations; recognizing something as an arm means finding a matching 3-D model somewhere in this series. The stability condition says that we could also find less approximate models in this series, if we would care to submit the image to more elaborate processing. By contrast, if we say: 'What I see looks like an arm' we imply only condition (i), not (ii).

3.2 Inverse limits

Consider again Marr's suggestive example of 3-D models of an arm. Viewed abstractly, what we see is a series of first order models, composed of objects and relations between them, together with a mapping specifying how an object occurring at one level is decomposed at the next. This situation can be represented by means of an *inverse system* of first order models. The basic ingredient is the following

Definition 1 Suppose $\mathcal{M}$, $\mathcal{M}'$ are first order models, with domains D, D'. Let $D^{<\omega}$ denote the set of finite sequences of elements of D. A *continuous onto mapping* $h : \mathcal{M} \longrightarrow \mathcal{M}'$ is a surjective function $h : D^{<\omega} \longrightarrow D'^{<\omega}$ such that for each formula $\varphi(x_1 \ldots x_n)$ of $\mathcal{M}'$,

$$\{w \in D^{<\omega} : \text{length}(h(w)) = n \text{ and } \mathcal{M}' \models \varphi(h(w))\}$$

is definable on $\mathcal{M}$ by means of a first order formula (without parameters). In this situation we write $\mathcal{M} \longrightarrow_h \mathcal{M}'$.

Intuitively the model $\mathcal{M}$ is a refinement which decomposes the objects in $\mathcal{M}'$; $h^{-1}(a)$ is the set of components of a in the refinement $\mathcal{M}$. The (technically convenient) surjectivity requirement says that all objects of $\mathcal{M}'$ are decomposed. It will become clear later why we call these mappings continuous.

In the example, we could let $\mathcal{M}'$ be a model $\langle D'; A, \ldots \rangle$, with a unary predicate A for 'arm'. $\mathcal{M}$ could be a model $\langle D; U, F, J, \ldots \rangle$, where U and F are unary (for 'upperarm' and 'forearm') and binary J (for 'joined')

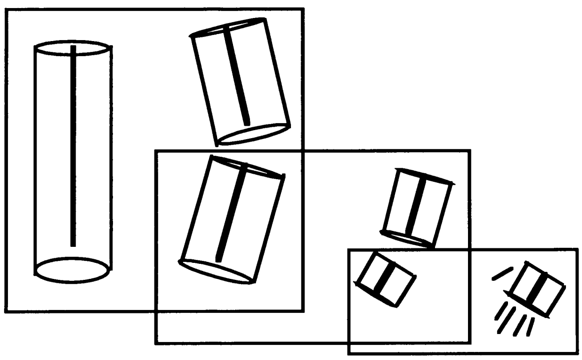

Arm – Forearm – Hand[2]

Clearly $\mathcal{M}$ represents the same situation in more detail than $\mathcal{M}'$. The continuous onto mapping h should be such that for all $d, e \in D$:

$$\text{if } \mathcal{M} \models U(d) \wedge F(e) \wedge J(d,e), \text{ then } \mathcal{M}' \models A(h(\langle d, e \rangle))$$

[2]After Marr 1982, p. 306.

But not every element in A may be decomposed in this manner. In the latter case, we need other relations to decompose the remaining objects to satisfy the surjectivity requirement.

It is convenient to present the gluing of one model onto another (cf. the picture) in the following manner: the arms of $\mathcal{M}'$ which are decomposed in $\mathcal{M}$ in a way compatible with the true notion of arm, are present in $\mathcal{M}$ *as objects*, not just as *sets* of objects satisfying certain relations. This leads to the formal notion of a model $\mathcal{M}$ being a refinement of a model $\mathcal{M}'$ with respect to a predicate A

Definition 2 $\mathcal{M} \longrightarrow_h \mathcal{M}'$ *is a refining pair with respect to the predicate* A if the folowing properties hold:

1. $\mathcal{M} \models A(a)$ implies $\mathcal{M}' \models A(h(a))$;
2. $\mathcal{M} \models A(a) \wedge A(b) \wedge a \neq b$ implies $h(a) \neq h(b)$;
3. $\mathcal{M}' \models \exists x A(x)$ implies $\mathcal{M} \models \exists x A(x)$.

$\mathcal{M} \longrightarrow_h \mathcal{M}'$ is a *proper refining pair* with respect to the predicate A if $h(\{d : \mathcal{M} \models A(d)\})$ is not definable on $\mathcal{M}'$ (hence it is in particular a proper subset of $\{e : \mathcal{M}' \models A(e)\}$). The definition has an obvious extension to simultaneous refinements of several predicates.

If $\mathcal{M} \longrightarrow_h \mathcal{M}'$ is a refining pair with respect to A, we may glue $\mathcal{M}$ onto $\mathcal{M}'$ by identifying the A-elements of the two models which are correlated by h; this is the content of Marr's picture. If $\mathcal{M} \longrightarrow_h \mathcal{M}'$ is a proper refining pair with respect to A, the difference between $\mathcal{M}$ and $\mathcal{M}'$ may show up in the fact that a statement $\exists x(A(x) \wedge \varphi)$ is true on $\mathcal{M}'$, whereas it becomes false on $\mathcal{M}$. However, we shall see below, in section 3.4 that the rough approximation of 'arm' given by the model $\mathcal{M}'$ may be reproduced inside $\mathcal{M}$. We now introduce arbitrary sequences of refinements.

Definition 3 Let $\langle T, \leq \rangle$ be a directed set, i.e. $\leq$ is transitive and reflexive and for each $t, t' \in T$ there exists $s \in T$ such that $t, t' \leq s$. An *inverse system of models indexed by* T is a tuple

$$\langle \{\mathcal{M}_t : t \in T\}, \{h_{st} : s, t \in T, t \leq s\} \rangle$$

such that for each $s, t \in T$ with $t \leq s$, $\mathcal{M}_s \longrightarrow_{h_{st}} \mathcal{M}_t$.[3]

Observe that the set of 3-D models (in the sense of Marr) will have the structure of an inverse system of models (in our sense): the condition of directedness says that two refinements will themselves have a common refinement. In line with our intuitive motivation, we shall mostly be interested in inverse systems which refine a distinguished predicate A in the following sense

[3]The customary notion of inverse limit in model theory (cf. Chang and Keisler 1990, p. 322; 1977, p. 243) is slightly different because it is based on the algebraic analogue. In our context the topological analogue is the more useful one.

Definition 4 The inverse system $\langle\{\mathcal{M}_t : t \in T\}, \{h_{st} : s,t \in T, t \leq s\}\rangle$ *refines* A if there exists some maximal chain $C \subseteq T$ with for all $s \leq t$: $\mathcal{M}_s \longrightarrow_{h_{st}} \mathcal{M}_t$ is a proper refining pair with respect to A.

Although our perceptual models of reality are always approximate, reality itself is assumed to be precise. Consider once again the example of the arm. The situation was the following: there were models $\mathcal{M}' := \langle D'; A, \ldots\rangle$, with a unary predicate A for 'arm' and $\mathcal{M} = \langle D; U, F, J, A, \ldots\rangle$, where U and F are unary (for 'upperarm' and 'forearm') and binary J (for 'joined'); moreover there was a mapping $h : \mathcal{M} \longrightarrow \mathcal{M}'$ such that

$$\text{if } U(d) \wedge F(e) \wedge J(d,e) \text{ then: } A(h(\langle d,e\rangle))$$

for all $d, e \in D$. We now want to think of h as a mapping η on assignments $D^{\text{Var}} \longrightarrow D'^{\text{Var}}$, in the following manner. Rewrite $(*)$ as, for all $d, e \in D$

$$\text{if } \mathcal{M} \models U(x) \wedge F(y) \wedge J(x,y)\ [f] \text{ then: } \mathcal{M}' \models A(z)\ [\eta(f)]$$

with the understanding that $f(x) = d$, $f(y) = e$ and $\eta(f)$ should be such that $\eta(f)(z) = h(\langle d,e\rangle)$. To ensure that the values of x and y are components of the value of z, we need a bijection $V : Var^{<\omega} \longrightarrow Var$, satisfying $V(\langle x,y\rangle) = z$. Given h and V, we may define η by $\eta(f)(z) = h(f|V^{-1}(z))$. By putting a topology on the set of assignments we arrive at a general definition of the η of interest.

Definition 5 A set of assignments on $\mathcal{M}$ (i.e., a subset of D^{Var}) is *clopen* if it is of the form $\{f : \mathcal{M} \models \varphi\ [f]\}$ for a first order formula φ. A function $\eta : D^{\text{Var}} \longrightarrow D'^{\text{Var}}$ is *continuous* if for a clopen $C \subseteq D'^{\text{Var}}$, $\eta^{-1}(C)$ is clopen.[4]

Definition 6 An *inverse system* is a tuple

$$\langle\{\mathcal{M}_t : t \in T\}, \{\eta_{st} : s,t \in T, t \leq s\}\rangle$$

such that T is directed and for $s,t \in T$ with $t \leq s$: $\eta_{st} : \mathcal{M}_s \longrightarrow \mathcal{M}_t$ is continuous and surjective.

Definition 7 (Adapted from Engelking 1989, p. 98) Let

$$\langle\{\mathcal{M}_t : t \in T\}, \{\eta_{st} : s,t \in T, t \leq s\}\rangle$$

be an inverse system. A *thread* is a sequence $(\xi_s)_{s\in T}$ such that each ξ_s is an assignment on $\mathcal{M}_s$, and for all $s,t \in T$ with $t \leq s$: $h_{st}(\xi_s) = \xi_t$. The *inverse limit* $\langle\mathcal{M}, \mathcal{F}\rangle$ of $\langle\{\mathcal{M}_t : t \in T\}, \{\eta_{st} : s,t \in T, t \leq s\}\rangle$ is constructed as follows. Given a thread $\xi = (\xi_s)_{s\in T}$, define an assignment f_ξ by $f_\xi(x) = (\xi_s(x))_{s\in T}$. The domain D of $\mathcal{M}$ will be

$$\{(\xi_s(x))_{s\in T} : x \in \text{Var}, (\xi_s)_{s\in T} \text{ a thread}\}$$

[4]'Clopen': means closed and open. We think of the definable sets as a basis for the topology. Since the complement of a definable set is again definable, the basis consists of clopen sets. If h is continuous in the usual sense, $h^{-1}(O)$ is open for open O and $h^{-1}(F)$ is closed for closed F, hence our definition.

but we take only a subset of D^{Var}, namely $\mathcal{F} := \{f_\xi : \xi \text{ a thread}\}$ as the set of admissible assignments. If $R(x_1, \ldots, x_k)$ is a relation, we define its interpretation (as a clopen set of admissible assignments) by[5]

$$f_\xi \in R(x_1, \ldots, x_k) \text{ iff: } \xi = (\xi_s)_{s \in T} \ \& \ \forall s \in T(\xi_s \in R(x_1, \ldots, x_k) \text{ on } \mathcal{M}_s)$$

Note: the inverse limit is specified as a model $\mathcal{M}$ *plus* a set of admissible assignments $\mathcal{F}$.

Theorem 1 *If* $\langle \{\mathcal{M}_t : t \in T\}, \{\eta_{st} : s, t \in T, t \leq s\} \rangle$ *is an inverse system of* ω_1*-saturated models, its inverse limit exists (i.e. its domain is nonempty).*

PROOFSKETCH. The result follows from theorem 3.2.13 in Engelking 1989. The assumption of ω_1-saturation (cf. Chang and Keisler 1990, ch. 5; 1977, 214) is necessary to push the requisite topological argument through: it ensures that the spaces of assignments can be made compact Hausdorff. □

Since any infinite model has an elementary extension to an ω_1-saturated model, the assumption appears to be logically harmless. One might try to restrict oneself to finite $\mathcal{A}_s$, even though the inverse limit will be infinite; however, as excercise 2.5 in Engelking 1989, p. 104, shows, the domain of the inverse limit of an inverse system of finite structures can be empty.

At this point it is appropriate to get rid of a technical nuisance. If we construct the inverse limit, not only the domain should be nonempty, but of course also some predicates should receive nonempty interpretation. By definition of the inverse limit, the interpretation of a predicate A is empty if it is empty on an $\mathcal{M}_s$. But this is precisely what happens if one takes Marr's picture seriously: there is a stage in which arms have not yet made their appearance. In order to avoid this problem, there appear to be two options. The elegant solution is to define a *filtered* inverse limit, i.e. to define the inverse limit as a submodel of a reduced product. In this case, a predicate has nonempty interpretation on the inverse limit if it is nonempty from a certain stage onward. Unfortunately, this creates complications for the treatment of conditional quantifiers below, so we opt for a less elegant alternative: at those stages at which arms have not appeared yet, we interpret A is the full domain, and *mutatis mutandis* for n-ary relations.

So far, the models $\mathcal{M}_s$ were viewed as approximations to reality, i.e. the inverse limit. In the next section we study the opposite direction: how to blur reality to get an approximation.

[5]This may look formidable, but it is a standard construction of algebraic logic: the inverse limit is a subdirect product of a collection of cylindric set algebras.

3.3 Conditional quantification

Consider clause (i) of the proposed semantics for 'I see an arm': with the present approximation the object that I focus on is identified as an arm. The logical form of this statement is some kind of existential quantification, namely 'there exists an arm in my (filtered) perceptual field', the filter corresponding to the degree of approximation. (The meaning of the word 'filter' here has nothing to do with its logical meaning; one should think rather of, say, a UV filter in photography, or the Dolby filter on a cassette deck.) However, this notion of quantification is *nonstable*: with respect to a finer filter, the existential statement may become false. The technical problem is, how to incorporate the filter into the existential quantifier. We first discuss one way of doing this, and then (in section 3.4) we return to the question of formally representing Marr's hierarchy of 3-D models.

The novel notion of a 'conditional quantifier' (cf. Van Lambalgen 1996) originated from an attempt to extract the logical content of conditional expectation in probability theory.[6] Since this heuristic motivation is directly relevant to our theme, vision, we reproduce it here, apologizing to those readers who are not familiar with probability. In any case we shall first give a simple example which hopefully succeeds in illustrating the main idea.

Example 1 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 $\epsilon, \delta \ll 1$; Y takes value $\frac{1}{2}$ on B_0 and value $-\frac{1}{2}$ on B_1. 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 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. If we can measure only $X + Y$, not X, this implies that we cannot determine the expectation $\mathbf{E}(X)$ of X, which would equal $\frac{1}{2}$, but only its expectation $\mathbf{E}(X|\mathcal{B})$ with respect to $\mathcal{B}$, which equals $\frac{1}{2} + \frac{1}{2}\epsilon - \frac{1}{2}\delta \neq \frac{1}{2}$. What we are after is a logical analogue of the expression '$\mathbf{E}(X|\mathcal{B})$'.

Example 2 A more realistic example is furnished by a typical instance of filtering. As in Marr (1982, pp. 54–61), we may think of a filter as a blurring operation (in his case applied to an image), usually by means of a Gaussian. Formally, we have a random variable X on a sample space Ω (measurable with respect to an algebra $\mathcal{B}$) which we want to measure; however, due to, for example, noise, we can only observe $X + c.\xi$, where ξ is a Gaussian with mean 0 and standard deviation 1. $X + c.\xi$ is measurable

[6] There does exist a (largely forgotten) body of work studying quantifiers over Boolean algebras in a probabilistic context; cf. Wright 1963 and references given therein. What is new in van Lambalgen 1996 is the addition of these quantifiers to first order logic. This requires techniques beyond Stone duality.

with respect to an algebra $\mathcal{G}$, which need not be the same as $\mathcal{B}$ (when both $\mathcal{B}$ and $\mathcal{G}$ are taken to be minimal). Hence we can only determine properties of X as filtered through $\mathcal{G}$; this is represented by the conditional expectation $\mathbf{E}(X|\mathcal{G})$. $\mathbf{E}(X|\mathcal{G})(\omega)$ is a $\mathcal{G}$-measurable random variable on Ω, representing the questions that can be asked about X from the point of view of $\mathcal{G}$. Suppose an experiment has been performed. The only information available to us regarding which sample point ω has been chosen is the value of $Z(\omega)$, for each $\mathcal{G}$-measurable random variable Z. This determines a measurable set $A \in \mathcal{G}$. Then the expected value of X given this information, $\int_A X dP$, equals $\int_A \mathbf{E}(X|\mathcal{G}) dP$; from the point of view of $\mathcal{G}$, no other questions about X can be answered. The Radon-Nikodym theorem is used to show that a function $\mathbf{E}(X|\mathcal{G})(\omega)$ with these properties exists. A statement of the form '$\mathbf{E}(\varphi|\mathcal{G}) > 0$ a.s.' denotes a strong form of existentially quantifying a formula φ (represented by an indicator function) relative to the available information, codified in $\mathcal{G}$.

We shall take conditional expectation as our heuristic model; the form of quantification we are interested in will be called *conditional quantification*. Sometimes we shall also speak of *filtered quantifiers*, especially when it is helpful to think of the conditioning algebras as filters (in the sense of blurring operations).

Consider first a rather roundabout way of formalising existential quantification on a model $\mathcal{M} = \langle D, \ldots \rangle$. Let $\mathcal{B}$ be the algebra of first order definable subsets of D^{Var} (D^{Var} is the set of assignments on $\mathcal{M}$); $\mathcal{G}_x$ denotes the subalgebra of $\mathcal{B}$ generated by formulas with free variables in $\text{Var} - \{x\}$. Then $\exists x$ is the unique surjective mapping $\mathcal{B} \longrightarrow \mathcal{G}_x$ satisfying

1) $\exists x \mathbf{0} = \mathbf{0}$, $\exists x \mathbf{1} = \mathbf{1}$;
2) $\varphi \leq \psi$ implies $\exists x \varphi \leq \exists x \psi$;
3) $\varphi \leq \exists x \varphi$; and
4) $\exists x(\varphi \wedge \exists x \psi) = \exists x \varphi \wedge \exists x \psi$.

Intuitively speaking, $\mathcal{G}_x$ represents the situation that we have precise information about variables other than x, and no information about x. Indeed, we may view $\mathcal{G}_x$ as the set of questions about variables (of the type 'does y satisfy ψ?') that can be asked and answered in this situation. We assume that all questions can be formulated in a first order language; since $\mathcal{G}_x$ contains all formulas in free variables other than x, it represents maximal information about those variables. Note that this way of analysing the existential quantifier is very different from the one which led to Mostowski-Lindström generalised quantifiers: there, $\exists$ is interpreted as the set of nonempty subsets of the domain, and the truth condition reads

$$\mathcal{M} \models \exists x \varphi(x) \text{ iff } \{d \in D : \mathcal{M} \models \varphi(d/x)\} \in \exists$$

The difference is that Mostowski-Lindström generalised quantification follows the Fregean tradition in analysing variables away (variables are just place holders), whereas here quantifcation is defined with respect to a given variable. This is because we think of a variable as an observable quantity, much like the random variables of probability theory. The resulting very non-Fregean treatment of variables is explained in greater detail in van Lambalgen 1996.

Suppose we have a formula $\varphi(x, y)$ and information $\psi(y)$ about y. The question 'does x satisfy $\varphi(x, y)$?' cannot be answered on the basis of $\mathcal{G}_x$ alone. The most we can say is either that it is consistent with $\psi(y)$ that $\varphi(x, y)$, i.e., $\exists x(\varphi(x, y) \wedge \psi(y))$, or that any x satisfies $\varphi(x, y)$, i.e., $\forall x(\varphi(x, y) \wedge \psi(y))$, etc. Viewed in this manner, the logical existential quantifier becomes analogous to the probabilistic conditional expectation considered above. This then leads to a consideration of arbitrary subalgebras $\mathcal{G}$ of $\mathcal{B}$, representing the available information about variables. For instance, if $\mathcal{G}$ is a proper subalgebra of $\mathcal{G}_x$, this means that for some variable y, not all possible questions about y can actually be asked (e.g., for lack of a sufficiently accurate measurement device). Below we shall give several examples illustrating this possibility.

The resulting new notion of generalised quantification is called *conditional quantification*; we write $\exists(\bullet|\mathcal{G})$ for the existential quantifier conditional on the algebra $\mathcal{G}$. Generalizing both from $\exists x$ and conditional expectation, we obtain the following set of properties that $\exists(\bullet|\mathcal{G})$ should satisfy.

1) $\exists(\varphi|\mathcal{G}) \in \mathcal{G}$. This says that the truth value of $\exists(\varphi|\mathcal{G})$ is determined solely on the basis of the information codifiable in $\mathcal{G}$; and implicitly also that $\exists(\varphi|\mathcal{G})$ is a function on the assignment space. This is analogous to requiring that $\mathbf{E}(\bullet|\mathcal{G})$ is $\mathcal{G}$-measurable. Note that we also implicitly require that $\mathcal{B}$ is closed under $\exists(\bullet|\mathcal{G})$ (this requirement is the most difficult to satisfy). In the following statements, $=$ and $\leq$ refer to the ordering in $\mathcal{B}$ (of course, $\leq$ equals $\subseteq$).
2) $\exists(\mathbf{0}|\mathcal{G}) = \mathbf{0}, \exists(\mathbf{1}|\mathcal{G}) = \mathbf{1}$;
3) $\varphi \leq \psi$ implies $\exists(\varphi|\mathcal{G}) \leq \exists(\psi|\mathcal{G})$ (monotonicity);
4) $\varphi \leq \exists(\varphi|\mathcal{G})$ ($\exists(\bullet|\mathcal{G})$ is increasing);
5) $\exists(\varphi \vee \psi|\mathcal{G}) = \exists(\varphi|\mathcal{G}) \vee \exists(\psi|\mathcal{G})$ (additivity)
6) $\exists(\varphi \wedge \psi|\mathcal{G}) = \exists(\varphi|\mathcal{G}) \wedge \psi$, where $\psi \in \mathcal{G}$ ('taking out what is known').

Property 6) is the analogue of the following property of conditional expectations: if Z is $\mathcal{G}$-measurable, then $\mathbf{E}(XZ|\mathcal{G}) = Z\mathbf{E}(X|\mathcal{G})$ a.s.. Since $\mathbf{E}(Y|\mathcal{G})$ is itself $\mathcal{G}$-measurable, we have the fundamental identity

$$\mathbf{E}(X\mathbf{E}(Y|\mathcal{G})) = \mathbf{E}(X|\mathcal{G})\mathbf{E}(Y|\mathcal{G}) \text{ a.s.}$$

In logic this corresponds to the so-called Frobenius property

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

Note that 2) and 6) imply that $\exists(\bullet|\mathcal{G})$ is the identity on $\mathcal{G}$.

All properties mentioned so far are satisfied by a quantifier satisfying (*) in the following definition.

Definition 8 Let L be a first order language, $\mathcal{M}$ a model for L, $\mathcal{F}$ the set of assignments on $\mathcal{M}$, $\mathcal{B}$ the algebra of definable subsets of $\mathcal{F}$ generated by $\mathcal{M}$, $\mathcal{G} \subseteq \mathcal{B}$. $\exists(\bullet|\mathcal{G})$, *an existential quantifier conditional on* $\mathcal{G}$, is any mapping $\mathcal{B} \longrightarrow \mathcal{G}$ satisfying

(*) For all $\psi \in \mathcal{G}$: $\varphi \leq \psi$ if and only if $\exists(\varphi|\mathcal{G}) \leq \psi$.

Condition (*) is known as a Galois correspondence. The reader may wish to check that in the particular case $\mathcal{G} = \mathcal{G}_x$, (*) is a way of stating the familiar left and right introduction rules for $\exists x$. The Galois correspondence suggests that we define $\exists(\varphi|\mathcal{G})$ as $\bigwedge\{\psi \in \mathcal{G} : \varphi \leq \psi\}$, but $\mathcal{G}$ need not contain this infinite infimum (in the case of $\mathcal{G}_x$ it does, though). We shall see later what to do about this, but while reading through the following set of examples it is helpful to keep the interpretation just given in mind.

Example 3 *'Blurring of individuals' In nuce*, the following example describes our approach to a semantics of perception. A statement like 'y sees x' is rendered formally as: '$\exists x[\varphi(x) \wedge S(y,x)]$', where φ defines a unique x. Symbol S gives the denotation of 'to see', and it delimits the set of objects $\{d : Sad\}$ seen by a. The quantifier $\exists x$ ranges over completely accurate objects; in formal terms, the elements of the inverse limits constructed in the previous section. To accomodate actual perception, which always has finite precision, we replace the quantifier $\exists x$ by a filtered quantifier $\exists(\bullet|\mathcal{G})$, where $\mathcal{G}$ represents the degree of blurring. Here we present only a simple case.

Let $\mathcal{M} = \langle\{1,2,3,4,5\}, S, W, U\rangle$, with '$S(a,b)$' for: a sees b, 'W' for: West, and 'U' for: up. In particular, $S = \{\langle 2,1\rangle, \langle 2,3\rangle, \langle 4,5\rangle\}$, $W = \{1,5\}$, $U = \{1,2,3\}$.

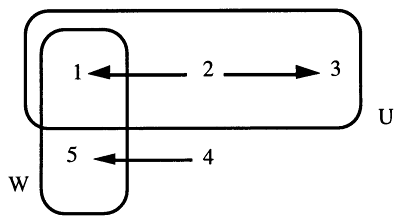

Up west

Put $\varphi(x,y) = S(y,x) \wedge W(x) \wedge U(y)$, then for f such that $f(y) = 2$: $\mathcal{M} \models \exists! x\varphi(x,y)$ $[f]$ ('y occupies the 'up' position, and looking to the West y sees exactly one object'); in fact

$$\begin{aligned} &\{f : \mathcal{M} \models \exists x\varphi(x,y)\ [f]\} \\ = \quad &\{f : \mathcal{M} \models \exists(\varphi(x,y)|\mathcal{G}_x)\ [f]\} \end{aligned}$$

$$= \{f : f(y) = 2\}$$

The algebra $\mathcal{G}_x$ represents the case that the viewer y has complete information, both about its own position ('up') and the direction in which it is looking ('west'). Let us now vary this situation, for instance by depriving the viewer of the information that he is in the 'up' position. The algebra $\mathcal{G}$ on the set of assignments, corresponding to this state of affairs, is determined by the formula algebra generated by the set $\{S(y,x), W(x)\}$. We now have

$$\begin{aligned} &\{f : \mathcal{M} \models \exists(\varphi(x,y)|\mathcal{G})\ [f]\} \\ &= \bigwedge\{\psi \in \mathcal{G} : \varphi \leq \psi\} \\ &= \{f : \mathcal{M} \models S(y,x) \wedge W(x)\ [f]\} \\ &= \{f : (f(y) = 2 \wedge f(x) = 1) \vee (f(y) = 4 \wedge f(x) = 5)\} \end{aligned}$$

Similarly, if y is so disoriented that he does not know whether he is looking East or West, we may describe his predicament by the algebra $\mathcal{H}$ generated by $\{S(y,x), U(y)\}$. In this case

$$\begin{aligned} &\{f : \mathcal{M} \models \exists(\varphi(x,y)|\mathcal{H})\ [f]\} \\ &= \bigwedge\{\psi \in \mathcal{H} : \varphi \leq \psi\} \\ &= \{f : \mathcal{M} \models S(y,x) \wedge U(x)\ [f]\} \\ &= \{f : f(y) = 2 \vee (f(x) = 1 \wedge f(x) = 3)\} \end{aligned}$$

We trust the reader can subject y to a still more savage experiment.

In order to highlight some formal features of this example, a technical point: we adopt the following convention concerning *free variables*

$$\mathrm{FV}(\exists(\varphi|\mathcal{G})) = \mathrm{FV}(\varphi) \cap \bigcup\{\mathrm{FV}(\psi)|\psi \in \mathcal{G}\}$$

Notice that the number of free variables need not be reduced by conditional quantification! Indeed, in the formula $\exists(\varphi(x,y)|\mathcal{G})$ x is neither fully bound nor fully free; whereas ordinary existential quantification $\exists x$ has the effect of abolishing all restrictions on x, a filtered quantifier need only liberalise restrictions on x to some extent, without fully abolishing them.

A second feature is: $\exists(\varphi(x,y)|\mathcal{G}) \in \mathcal{G}$. Unfortunately, this is lost in more intricate cases, the subject of our next example.

Example 4 *'Truck through door?'* We now replace the dichotomy up-down by a continuum of possibilities. A truck has to pass through a narrow door in a wall; a person behind the truck checks whether this is possible. In this case the viewer must choose his position rather accurately: too far up or down means that he will see only one side of the truck, if he comes too close the truck will occlude the door, but if he goes too far to the left he cannot accurately estimate the distance between truck and doorposts. See the picture.

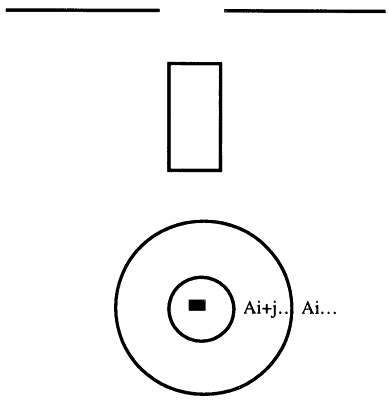

Truck through door

Let $O(x)$ be the predicate 'x is a (sufficiently large) opening', $S(y,x)$ the relation 'y sees x' (where the viewer y is identified by the coordinates of his position in the plane) and let $A_n(y)$ be a countable collection of predicates denoting open sets in the plane. We assume the following logical relations between these predicates

1) $\forall n \forall x \forall y (O(x) \wedge S(y,x) \rightarrow A_n(y))$
2) $\forall n (A_{n+1}(y) \rightarrow A_n(y))$
3) $\forall n \exists y (\neg A_{n+1}(y) \wedge A_n(y))$
4) $\forall n \exists y (A_n(y) \wedge \forall x (O(x) \rightarrow \neg S(y,x)))$.

Condition 4) expresses that it is hard to find the exact position from which an opening can be accurately observed; each open set A_n contains positions from which no opening is visible.

Suppose $\mathcal{G}$ is the Boolean algebra generated by

$$\{O(x)\} \cup \{S(y,x)\} \cup \{A_n(y) : n \in \omega\}$$

This algebra represents the situation that the viewer has no precise information about his location y; the only available information is in the form of the open sets $A_n(y)$. This has a twofold consequence

a) $\exists(O(x) \wedge S(y,x)|\mathcal{G}) \notin \mathcal{G}$,
b) $\exists(O(x) \wedge S(y,x)|\mathcal{G})$ is not first order definable.

For a), since $O(x) \wedge S(y,x) \rightarrow A_n(y)$, we must have

$$\exists(O(x) \wedge S(y,x)|\mathcal{G}) \rightarrow A_n(y)$$

If $\exists(O(x) \wedge S(y,x)|\mathcal{G}) \in \mathcal{G}$, then $\exists(O(x) \wedge S(y,x)|\mathcal{G})$ must be a Boolean combination of the A_n, say ψ. Write ψ in distributive normal form, then since $\psi \rightarrow A_n(y)$, negative occurrences of the A_n cancel out. It follows that ψ is equal to some $A_k(y)$. However, this conflicts with

$$\exists y(A_n(y) \wedge \neg A_{n+1}(y))$$

For b), let $\mathcal{M}$ be an ω-saturated model of the theory 1–4). On $\mathcal{M}$, $\exists(O(x) \wedge S(y,x)|\mathcal{G})$ must equal $\bigwedge A_n$. We show that $\bigwedge A_n$ is not first order definable. Let χ be a first order formula defining $\exists(O(x) \wedge S(y,x)|\mathcal{G})$, then for all n: $\chi \to A_n(y)$. The set

$$\{\neg\chi\} \cup \{A_n(y) : n \in \omega\}$$

must be finitely satisfiable, otherwise for some k, $A_k(y) \to \chi$, whence $\chi \leftrightarrow A_k(y)$, in contradiction with $\exists y(A_n(y) \wedge \neg A_{n+1}(y))$. Hence on $\mathcal{M}$, $\bigwedge A_n = \exists(O(x) \wedge S(y,x)|\mathcal{G})$ is strictly larger than χ. In particular $\exists(O(x) \wedge S(y,x)|\mathcal{G})$ is not equivalent to $\exists x(O(x) \wedge S(y,x))$.

Since we require $\exists(\varphi|\mathcal{G}) \in \mathcal{G}$, the moral of the example is that the set-up chosen hitherto is too narrow. But before we remedy the situation, we give a fifth example concerning a logical analogue of probabilistic conditionalisation.

Example 5 So far we conditioned quantifiers on restricted information; in probabilistic terms, we consider a subalgebra of the full algebra, but retain the measure. Sometimes, however, the measure may change by *conditionalisation*: if P is the original ('a priori') measure, and we know for sure that B has happened (where $P(B) > 0$), then a new ('a posteriori') measure P' is determined by

$$P'(A) = P(A|B) = \frac{P(A \cap B)}{P(B)}$$

Logically this corresponds to taking the universe of the conditioning algebra to be a subset of the full set of assignments. For instance, if $\mathcal{M}$ is a model with domain D, we may move to a submodel $\mathcal{M}'$ of $\mathcal{M}$ with domain $D' \subseteq D$, with corresponding set of assignments $\mathcal{F}' = D'^{\mathrm{Var}}$. Let $\mathcal{H}$ be an algebra on $\mathcal{F}$. We briefly study the properties of $\exists(\bullet|\mathcal{H})$. Its defining conditions are

i) $\exists(\psi|\mathcal{H}) \in \mathcal{H}$
ii) $\mathcal{F} \cap \psi \subseteq \theta$ iff $\exists(\psi|\mathcal{H}) \subseteq \theta$, for all $\theta \in \mathcal{H}$.

For instance, we get (a), but also (b)

a) $\exists(\mathbf{0}_{\mathcal{H}}|\mathcal{H}) = \mathbf{0}_{\mathcal{H}} = \exists(\mathbf{0}|\mathcal{H}) = \mathbf{0}$
b) $\exists(\mathbf{1}_{\mathcal{H}}|\mathcal{H}) = \mathbf{1}_{\mathcal{H}} \neq \exists(\mathbf{1}|\mathcal{H}) = \mathbf{1}$

All other properties of $\exists(\bullet|\mathcal{H})$ are relativised to $\mathcal{F}$ as well; e.g., we get, for all ψ , $\mathcal{F} \cap \psi \subseteq \exists(\psi|\mathcal{H})$, whereas $\psi \subseteq \exists(\psi|\mathcal{H})$ now fails. It is especially the combination of the two processes, filtering and conditionalisation, that is useful for a logic of vision. Conditionalisation corresponds to the submodel of reality determined by the perceptual field, while filtering represents the degree of accuracy with which we we can perceive that submodel.

What follows is a technical aside, not directly relevant to the line of argument of the paper. In example 4 ('truck through door') we have seen that since conditional quantifiers need not be first order definable, they cannot be functions $\mathcal{B} \longrightarrow \mathcal{B}$; a fortiori, subalgebras $\mathcal{G} \subseteq \mathcal{B}$ can be too small to be consistent with $\exists(\bullet|\mathcal{G}) \in \mathcal{G}$. In order to prove existence of conditional quantifiers we have to enlarge both $\mathcal{B}$ and $\mathcal{G}$ to the σ-algebras generated by them.

Definition 9 A σ-algebra on X is a set of subsets of X closed under countable unions and complements, and containing X. If $\mathcal{B}$ is a Boolean algebra, the σ-algebra generated by $\mathcal{B}$ is the smallest σ-algebra containing $\mathcal{B}$.

However, after enlarging the algebra we face a fresh problem. A σ-algebra is uncountable, hence the infimum necessary to define $\exists(\bullet|\mathcal{G})$ will in general be an infimum over uncountably many sets, so will not be in $\mathcal{G}$. In Van Lambalgen 1996 it is shown how to overcome this problem. The proof is again too complicated to be included here. The trick is to leave $\exists(\bullet|\mathcal{G})$ undefined on a 'small' set of assignments, where the relevant notion of smallness is given by the following definition from topology adapted to the present context

Definition 10 A set N of assignments on $\mathcal{M}$ is *meagre* if there exists a set of first order formulas $\{\varphi_n : n \in \omega\}$ such that

i) $N = \{f : \text{for all } n, \mathcal{M} \models \varphi_n\ [f]\}$
ii) there is no formula χ, such that for all n: $\mathcal{M} \models \forall\vec{x}(\chi \rightarrow \varphi_n)$

Here, $\forall\vec{x}$ denotes universal closure.

Definition 11 Let L be a first order language, $\mathcal{M}$ a model for L, $\mathcal{B}$ the Borel σ-algebra generated by the L-definable sets of assignments, $\mathcal{G} \subseteq \mathcal{B}$ a sub σ-algebra. $\exists(\bullet|\mathcal{G})$, *an existential quantifier conditional on* $\mathcal{G}$, is any mapping $\mathcal{B} \longrightarrow \mathcal{G}$ satisfying: there exists a meagre set N such that outside N: for all $\psi \in \mathcal{G}$, $\varphi \leq \psi$ if and only if $\exists(\varphi|\mathcal{G}) \leq \psi$.

Theorem 2 *Any first order model $\mathcal{M}'$ has an elementary extension to a model $\mathcal{M}$ such that for any countably generated sub σ-algebra $\mathcal{G}$ of $\mathcal{B}$ (the σ-algebra generated by the first order definable sets of assignments on $\mathcal{M}$), a version of conditional quantification relative to $\mathcal{G}$ exists.*

While we have included these results here to convince the reader that conditional quantifiers exist, they are not indispensible for an understanding (as opposed to a rigorous treatment) of the proposed semantics of perception reports. The reader may continue to think of conditional quantifiers as given by definition 8, and as exemplified by the first example, although in the next subsection we shall sometimes implicitly use the fact that conditional quantification can be applied to Borel sets.

One thing is crucially important: a conditional quantifier is determined by its conditioning algebra, considered as algebra of sets of assignments; hence the variables which occur in the formulas defining the sets of assignments are an essential (non-Fregean) ingredient in this notion of quantification. The conditioning algebras *cannot* be thought of as determined by subsets of (powers of) the domain (eliminating explicit reference to variables). The reason is that the conditioning algebras may contain formulas of arbitrary arity. This makes conditional quantification different from Mostowski-Lindström quantification, where each quantifier can only take as arguments tuples of relations of which the arities are fixed. As a consequence, if we aim at a formalism in which both types of quantifiers can occur simultaneously, the definition of Mostowski-Lindström quantification has to be generalised to allow relations between sets of assignments.

We now return to the inverse limits of section 3.2, and show how conditional quantifiers can be interpreted in them.

3.4 Approximate models formalised

In section 3.2 we argued that reality, as relevant to the logic of perception reports, should be constructed as an inverse limit of approximations. In the previous section we introduced conditional, or filtered, quantifiers, which allow one to blur reality. Putting these two ideas together, an obvious question arises: if $\mathcal{M}$ is the inverse limit of the inverse system $(\mathcal{M}_s)_{s \in T}$, can we recapture the approximations $\mathcal{M}_s$ by applying a suitable filter to $\mathcal{M}$? If so, the $\mathcal{M}_s$ would have a dual role: as refinements or approximations (when viewed in relation to $\mathcal{M}_t$ for $t \leq s$), or as a blurring of the real world, as represented by $\mathcal{M}$. The answer is 'yes', and to prove this we shall from now on essentially use the fact that conditional quantifiers 'live' on algebras of sets of assignments.

We ask the reader to once again took a look at definition 7 of inverse limit. There is a way of approaching this definition, which relates it to the non-Fregean view of variables briefly discussed in the previous subsection. A common situation in probability theory is that one wants to refine a sample space Ω' into a sample space Ω. This means that points of Ω' correspond to sets of points of Ω; intuitively, elementary events of Ω' are further subdivided in Ω. Let $\mathcal{G}$ be a σ-algebra on Ω and $\mathcal{G}'$ a σ-algebra on Ω'. The refinement is then given by a surjective map $\eta : \Omega \longrightarrow \Omega'$ such that $\eta^{-1}(\mathcal{G}')$ is a subalgebra of $\mathcal{G}$. If X is a random variable on Ω, it can be 'blurred' to a random variable on Ω' by taking the conditional expectation $\mathbf{E}(X|\eta^{-1}(\mathcal{G}'))$. The analogy between this situation and ours is the following.

We think of both variables and formulas as random variables on a space, namely the set of assignments. Let $\mathcal{M}$ be a model with set of assignments D^{Var}. If x is a variable, it corresponds to a function $X : D^{\text{Var}} \longrightarrow D$ by

$$X(f) = f(x)$$

for all $f \in D^{\text{Var}}$. Similarly a formula $\varphi(x_1, \ldots, x_n)$ corresponds to a function F which maps the 'random variables' $X_1, \ldots, X_n$ to the 'random variable' $F(X_1, \ldots, X_n) : D^{\text{Var}} \longrightarrow \{0, 1\}$ by

$$F(X_1, \ldots, X_n)(f) = 1 \text{ iff } \mathcal{M} \models \varphi(x_1, \ldots, x_n)\ [f]$$

The meaning of this move is that variables acquire individuality, instead of being 'explained away' as Quine would have it. This is suggestive when we seek to apply probabilistic concepts to logic. For instance, as in the probabilistic case, we now want to compare a formula on D^{Var} with its 'blurred' version on D'^{Var}; given the analogy between formulas and random variables, conditional quantification seems suited to play the role of conditional expectation here. So we are given an inverse system $\langle \{\mathcal{M}_t : t \in T\}, \{\eta_{st} : s, t \in T, t \leq s\} \rangle$; we shall try to relate truth on model $\mathcal{M}_s$ with truth, 'filtered' by a suitable algebra, on the inverse limit. We first introduce the various algebras that we shall need.

Lemma 3 *Let $\langle \{\mathcal{M}_t : t \in T\}, \{\eta_{st} : s, t \in T, t \leq s\} \rangle$ be an inverse system. For any $s \in T$, if $\mathcal{M}_s$ is a model with set of assignments $\mathcal{F}_s$, let $\mathcal{G}_s$ be the Boolean algebra of clopen subsets of $\mathcal{F}_s$. If $t \leq s$, then since the mapping η_{st} is continuous, $\eta_{st}^{-1}(\mathcal{G}_t)$ is a subalgebra of $\mathcal{G}_s$.*

Definition 12 Let $\langle \mathcal{M}, \mathcal{F} \rangle$ be an inverse limit. We define a projection $\pi_s : \mathcal{F} \longrightarrow \mathcal{F}_s$ by $\pi_s(f_\xi) = \xi_s$, where $\xi = (\xi_t)_{t \in T}$ is a thread.

That the projection mappings π_s are surjective follows from corollary 3.2.15 in Engelking 1989. Also, $\pi_s = \eta_{ts} \circ \pi_t$, for $s \leq t$.[7]

The π_s are continuous with respect to the Tychonoff topology on $\Pi_{t \in T} \mathcal{F}_t$; in this topology, $\mathcal{F}$ is a closed subset of $\Pi_{t \in T} \mathcal{F}_t$. Note, however, that there is a second natural topology on $\mathcal{F}$, determined by the formulas. Recall that, if $R(x_1, \ldots, x_n)$ is a relation, we define its interpretation (as a clopen set of admissible assignments) by

$$f_\xi \in R(x_1, \ldots, x_k) \text{ iff: } \xi = (\xi_s)_{s \in T} \ \&\ \forall s \in T(\xi_s \in R(x_1, \ldots, x_k) \text{ on } \mathcal{M}_s)$$

This shows that, from the point of view of the Tychonoff topology, R is of the form $\bigcap_{s \in T} O_s$. If T is countable, sets open in the second topology are Borel, in fact G_δ, with respect to the first topology. Henceforth we work under the assumption that T is countable. In this case the π_s are Borel measurable with respect to the formula topology; this is useful when we represent the approximating models $\mathcal{M}_s$ on the inverse limit.

Indeed, let $\langle \mathcal{M}, \mathcal{F} \rangle$ be the limit of the inverse system, with $\mathcal{B}$ the corresponding σ-algebra generated by the restriction to $\mathcal{F}$ of the product topol-

[7] **commutative diagram to be drawn.**

ogy on $\Pi_{t\in T}\mathcal{F}_t$. In this case the first order definable subsets of $\mathcal{F}$ determine a subalgebra of $\mathcal{B}$. By means of the projections π_s, each algebra $\mathcal{G}_s$ of clopen subsets of $\mathcal{F}_s$ can be identified with a subalgebra $\pi_s^{-1}(\mathcal{G}_s)$ of $\mathcal{B}$. We abbreviate $\pi_s^{-1}(\mathcal{G}_s)$ by $\mathcal{B}_s$. The next results concern the relation between truth on $\mathcal{M}_s$, and truth on $\mathcal{M}$, conditioned on $\mathcal{B}_s$.

Definition 13 A formula is *positive primitive* if it is equivalent to a formula in which only $\vee, \wedge, \exists$ occur. A formula is *positive* if it is equivalent to a formula in which only $\vee, \wedge, \exists, \forall$ occur.

The importance of positive primitive formulas (also called *geometric* formulas in the constructive tradition) was emphasised by Mönnich in his work in progress on the semantics of perception reports (personal communication).

Theorem 4 *For any formula ψ, let ψ_s denote $\{f \in \mathcal{F}_s : \mathcal{M}_s \models \psi\ [f]\}$. If ψ is positive primitive, we have $\psi \subseteq \pi_s^{-1}(\psi_s)$ (where ψ, $\pi_s^{-1}(\psi_s)$ are considered as elements of $\mathcal{B}$).*

PROOF. If ψ is atomic, this is the truth definition of ψ on $\mathcal{M}$. Cases $\vee$, $\wedge$ are trivial. Suppose $f_\xi \in \exists x\varphi$, then there exists a thread ξ' such that $f_{\xi'} =_x f_\xi$ and $f_{\xi'} \in \varphi$. But $f_{\xi'} =_x f_\xi$ implies $\xi'_s =_\xi \xi_s$, whence $\xi_s \in (\exists x\varphi)_s$. □

Corollary 5 *If ψ is positive primitive, $\exists(\psi|\mathcal{B}_s) \subseteq \pi_s^{-1}(\psi_s)$.*

PROOF. Since $\pi_s^{-1}(\psi_s) \in \mathcal{B}_s$ and $\psi \subseteq \pi_s^{-1}(\psi_s)$, the statement follows from the Galois correspondence. Note that we essentially use the fact that $\exists(\bullet|\mathcal{B}_s)$ can be applied to Borel sets. □

This result is best possible in the sense we cannot even expect it to hold for negated atomic formulas: $\neg\alpha$ may be false on $\mathcal{M}$ because it fails on a coordinate t different from s.

Strong additional conditions on the η_{st} render the theorem true for positive formulas as well, but in the general case $\forall$ is troublesome. The problem is, roughly speaking, that the inverse limit $\langle\mathcal{M}, \mathcal{F}\rangle$ comes with a set of admissible assignments $\mathcal{F}$, which is in general a proper subset of the set of all assignments on its domain. As a consequence, it is 'too easy' for a universal formula to become true on $\mathcal{M}$. We do not know whether this is an artefact of our formulation, or a reflection of an essential feature of the logic of perception.

But the corollary can be strengthened in the sense that in some cases we may have equality for positive primitive formulas. For convenience, we restate the definition of a proper refining pair in the present format.

Definition 14 $\mathcal{M} \longrightarrow_{\pi_s} \mathcal{M}_s$ is a proper refining pair with respect to the positive primitive formula ψ if there does not exist a formula τ such that $\psi \subseteq \pi_s^{-1}(\tau_s)$ and τ_s is properly contained in ψ_s.

Corollary 6 *If $\mathcal{M} \longrightarrow_{\pi_s} \mathcal{M}_s$ is a proper refining pair with respect to* *the positive primitive formula* ψ, *then either* $\exists(\psi|\mathcal{B}_s) = \mathbf{0}$ *or* $\exists(\psi|\mathcal{B}_s) = \pi_s^{-1}(\psi_s)$.

PROOFSKETCH. Assume $\exists(\psi|\mathcal{B}_s) \neq \mathbf{0}$, and that there is some $C \in \mathcal{B}_s$ such that $\exists(\psi|\mathcal{B}_s) \subseteq C \subset \pi_s^{-1}(\psi_s)$. We only do the case when there is clopen $C' \subseteq \mathcal{F}_s$ such that $C = \pi_s^{-1}(C')$. We then have $C' = \pi_s(\pi_s^{-1}(C')) \subset \psi_s$, and $\psi \subseteq C$, whence $\pi_s\psi \subseteq \pi_s C = C' \subset \psi_s$. This contradicts the fact that $\mathcal{M} \longrightarrow_{\pi_s} \mathcal{M}_s$ is a proper refining pair with respect to ψ. □

We add a few clarificatory remarks. The purpose of this result and the next is to relate two processes: one going from the approximations $\mathcal{M}_s$ to the inverse limit $\mathcal{M}$, and one going in the opposite direction, from $\mathcal{M}$ to its blurred versions. If A is a predicate, one may view the formula $\exists(A|\mathcal{B}_s)$ as the best estimate of A, given the 'filter' $\mathcal{B}_s$. This estimate should be contained in $\pi_s^{-1}(A_s)$, which is what $\mathcal{M}_s$ proposes as an estimate of A. However, one doesn't necessarily have equality, since $\mathcal{M}_s$ itself might actually have a better estimate of A than A_s, namely some formula θ implying A, but not equivalent to it! In the running example, where A stands for 'arm', this would be the situation for a model $\mathcal{M}_s$ where

i) A is already decomposed, e.g. in upper arm and forearm, and
ii) A_s properly contains the set of objects decomposed in upper arm and forearm.

If this is not the case, i.e. if $\mathcal{M} \longrightarrow_{\pi_s} \mathcal{M}_s$ is a proper refining pair with respect to A, then we do have $\exists(A|\mathcal{B}_s) = \pi_s^{-1}(A_s)$.

In this context it is of interest to observe that, if the inverse system is refining with respect to A, but not properly refining, i.e. if there does exist θ such that $\mathcal{M} \models \forall\vec{x}(A \rightarrow \theta)$ and $\mathcal{M} \models \forall\vec{x}(\theta \rightarrow A)$, but $\mathcal{M}_s \models \forall\vec{x}(\theta \rightarrow A)$, then there will be a $t \geq s$ such that $\mathcal{M}_t \models \forall\vec{x}(\theta \rightarrow A)$;[8] this uses the theorem on basic Horn formulas below. Hence the assumption that $\mathcal{M} \longrightarrow_{\pi_s} \mathcal{M}_s$ is a proper refining pair with respect to A, means that $\mathcal{M}_s$ represents our best guess concerning A; we couldn't have done better at stage s. Viewed in this light, the following result, which relates truth in $\mathcal{M}_s$ with 'filtered truth' on the inverse limit, should come as no surprise. The proof uses ω_1-saturation of the models $\mathcal{M}_s$.

Theorem 7 *Let* φ *be a formula of the form* $\exists x\theta$, *where* θ *is a conjunction of predicates, and* $\mathcal{M} \longrightarrow_{\pi_s} \mathcal{M}_s$ *a proper refining pair with respect to* θ *('*θ *is an atom in* $\mathcal{B}_s$*'). If* x *is a variable, let* $\mathcal{G}_{s,x}$ *be the subalgebra of* $\mathcal{G}_s$ *determined by formulas not containing* x *free, and let* $\mathcal{B}_{s,x}$ *be its inverse image under* π_s. *Then we have either* $\exists(\psi|\mathcal{B}_{s,x}) = \mathbf{0}$ *or*

$$\mathcal{M}_s \models \exists x\theta\ [g] \text{ iff: for all } f_\xi \text{ such that } \xi_s = g, \mathcal{M} \models \exists(\theta|\mathcal{B}_{s,x})\ [f_\xi]$$

[8]As before, $\forall\vec{x}$ stands for universal closure.

PROOF. Suppose $\mathcal{M}_s \models \exists x\theta\ [g]$, and f_ξ satisfies $\xi_s = g$. We have to show that $f_\xi \in \exists(\theta|\mathcal{B}_{s,x})$. There exists $g' =_x g$ such that $\mathcal{M}_s \models \theta\ [g']$. Corollary 3.2.15 in Engelking 1989, p. 142, assures that there is a thread $(\xi'_t)_{t\in T}$ such that $\xi'_s = g'$ (this essentially uses the fact that the assignments spaces $\mathcal{F}_s$ are compact, i.e. the ω_1-saturation of the models $\mathcal{M}_s$). Since θ is an atom, we have $f'_\xi \in \exists(\theta|\mathcal{B}_s)$, hence $f'_\xi \in \exists(\theta|\mathcal{B}_{s,x})$. Now $\exists(\theta|\mathcal{B}_{s,x}) \in \mathcal{B}_{s,x}$, hence it is of the form $\pi_s^{-1}(C)$, where $C \in \mathcal{G}_{s,x}$. Then we have $f'_\xi \in \exists(\theta|\mathcal{B}_{s,x})$ implies $g' = \pi_s(f_{\xi'}) \in C$, but since C cannot distinguish between assignments with different x-coordinates, it must follow that $g = \pi_s(f_\xi) \in C$, i.e. $f_\xi \in \exists(\theta|\mathcal{B}_{s,x})$.

For the other direction, assume the thread $(\xi_t)_{t\in T}$ is such that $\xi_s = g$ and $\mathcal{M} \models \exists(\theta|\mathcal{B}_{s,x})\ [f_\xi]$. By theorem 4, $\exists x\theta \subseteq \pi_s^{-1}((\exists x\theta)_s)$, hence also $\theta \subseteq \pi_s^{-1}((\exists x\theta)_s)$. Since $\pi_s^{-1}((\exists x\theta)_s) \in \mathcal{B}_{s,x}$, the Galois property shows that $\exists(\theta|\mathcal{B}_{s,x}) \subseteq \pi_s^{-1}((\exists x\theta)_s)$. □

A few applications of the Galois conditions show that $\exists(\theta|\mathcal{B}_{s,x}) = \exists x\exists(\theta|\mathcal{B}_s)$, which we shall use in section 4 to study the in- and export of quantifiers.

We regard the inverse limit $\mathcal{M}$ and the conditional quantifiers as the fundamental concepts here. We assume that perception always gives us the best estimate allowed by a filter; this estimate is represented formally by a condionally quantified statement. The theorem then tells us under what circumstances truth on $\mathcal{M}_s$ and truth filtered by $\mathcal{M}_s$ coincide. This will be put to use in the next section, on the semantics of perception reports.

The result can easily be extended to formulas of the form $\exists\vec{x}\theta$. Observe that every positive primitive formula is a disjunction of such formulas. Hence if a positive primitive formula φ can be written in such a way that for each of the disjuncts $\exists\vec{x}\theta$, $\mathcal{M} \longrightarrow_{\pi_s} \mathcal{M}_s$ is a proper refining pair with respect to θ, the theorem holds for φ as well.

Thus far we have emphasised the changes that may occur when we move from a given approximation to a more refined one. However, one would also expect some general features of the world to be preserved, say concerning possible spatial arrangements of shapes. For a discussion of the importance of this topic, cf. Marr 1982, Ch. 5. Logically speaking, these constraints can often be expressed by means of basic Horn formulas.

Definition 15 A formula is basic Horn if it is of the form $\forall\vec{x}(\varphi \to \psi)$, where φ is positive primitive and ψ is atomic or $\bot$. A theory is basic Horn if it consists of special Horn sentences.

Proposition 8 *Let* $\langle\{\mathcal{M}_t : t \in T\}, \{\eta_{st} : s, t \in T, t \leq s\}\rangle$ *be an inverse system, with inverse limit* $\mathcal{M}$. *If* Γ *is a basic Horn theory true on all* $\mathcal{M}_s$, *then* Γ *is true on* $\mathcal{M}$.

PROOF. Choose $\forall\vec{x}(\varphi \to \psi)$ in Γ, and suppose $\mathcal{M} \models \varphi\ [f_\xi]$. By theorem 4, $\varphi \subseteq \pi_s^{-1}(\varphi_s)$, whence $\mathcal{M}_s \models \varphi\ [\xi_s]$. If ψ is atomic, $\mathcal{M}_s \models \psi\ [\xi_s]$. Since this holds for all $s \in T$, we have $\mathcal{M} \models \psi\ [f_\xi]$ by definition of the inverse limit. If ψ is $\perp$, it follows that $\mathcal{M} \models \varphi\ [f_\xi]$, and we are done. □

This finishes our introduction of the basic concepts. Next we show how they can be used to define a semantics of direct perception.

4 Perception, or: non-monotonicity, the hard way

Now that the formal apparatus is in place, we shall develop a logical semantics for direct perception reports in two steps. First, section 4.1 shows how fields of vision are part of reality (as an inverse limit). This is used to interpret direct perception reports of form '$S(a, b)$' for: a sees b, and of form '$\text{SEE}(a, \varphi)$' for: a sees φ. We discuss some of the inferences of section 2 in terms of this semantics. However, it only gives perception relative to an approximation, which is non-veridical. The important veridical part expects that what is perceived will remain the case for more refined approximations. Such expectations are defeasible, and should therefore be given in terms of non-monotonic rules. This is what we shall do in section 4.2.

4.1 Non-veridical perception

The semantics for non-veridical perception is based on the idea that a field of vision approximates part of reality. Since reality is already identified with an inverse limit of first order models, the only thing left for a semantics of non-veridical NI reports is to specify the rôle of approximation and partiality. Given the discussion in section 3, it is clear that an approximation should be one of the stages of an inverse system, and in such a stage s partiality consists of taking a submodel $\mathcal{M}'$ of $\mathcal{M}_s$.

Section 2 has an example which indicates that someone's range of vision may come with different granularities (cf. (19)). This would mean that a visual field consists of several approximations, not just one. In what follows, however, we assume that each viewer comes with a unique range of vision.

Definition 16 Let S be the two place relation corresponding to 'x sees y'. A *perceptual model* is a tuple $\langle \mathcal{M}, S, \mathcal{F}, \Sigma \rangle$ where

i) $\langle \mathcal{M}, S; \mathcal{F} \rangle$ is the inverse limit of a system $(\langle \mathcal{M}_s, S_s \rangle)_{s \in T}$ (T countable and directed, $\mathcal{F}$ the set of admissible assignments); and
ii) Σ is a function from $\text{dom}(S)$ to $\{\mathcal{M} : \mathcal{M} \subseteq \mathcal{M}_s,\ s \in T\}$.

We often write $\mathcal{M}_a$ for the submodel that Σ assigns to $a \in \text{dom}(S)$ in a certain $\mathcal{M}_s$.

In a perceptual model, the symbol S has a double rôle to play: it gives the denotation of the transitive verb 'to see', and it delimits the set of objects $\{d : Sad\}$ seen by a. The function Σ is used to interpret the counterparts $\text{SEE}(a, \varphi)$ of NI reports 'a sees φ', where a is a term and φ a formula. For the moment we disallow iterations of 'SEE' (as in van der Does 1991).[9]

[9]Presumably, iterations of SEE can be obtained by setting: $\text{SEE}(a, \text{SEE}(b, \psi)) :=$

Definition 17 Let $\mathcal{M}$ be a perceptual model and $\mathcal{F}$ its set of admissible assignments. The *interpretation of formulas* $\text{SEE}(a,\varphi)$ in $\mathcal{M}$ is defined by

$$\begin{array}{rcl} \text{SEE}(a,\psi) & := & \exists(\psi|\mathcal{A}_a) \\ \text{SEE}(a,\neg\chi) & := & \mathcal{F} - \text{SEE}(a,\chi) \\ \text{SEE}(a,\exists x\chi) & := & \exists(S(a,x)\wedge\chi|\mathcal{A}_{a,x}) \end{array}$$

Here ψ is atomic, conjunctive or disjunctive. $\mathcal{A}_a$ is the Boolean algebra of clopen subsets of the set of assignments for $\Sigma(a)$ (thought of as a subalgebra of the σ-algebra $\mathcal{B}$ for $\langle\mathcal{M},S\rangle$, as before). $\mathcal{A}_{a,x}$ is the subalgebra of $\mathcal{A}_a$ generated by the formulas in which x does not occur free.

Let us consider how the semantics fares with respect to the non-veridical principles discussed in section 2.

4.1.1 Retract of perception

Recall the children's story on the apparent magic of dimes (i.e. (7), section 2.1). In this story, we describe one and the same object x as it goes from one stage to another. Due to the change of granularity in Jack's perceptions at these stages, its properties vary. Formally this corresponds to conditionally quantifying over threads ξ in the inverse limit, such that at consecutive stages s_i $\pi_{s_i}(\xi)(x)$ has the property of being a dime, a quarter, a dollar, a trash can...

4.1.2 Partial perception

In a perceptual model $\mathcal{M}$ the function Σ assigns to a viewer a a submodel of an approximation $\mathcal{M}_s$. Consequently, (10) does *not* entail (11).

(10) Jack saw Sharon, and Sharon winked.

(11) Jack saw Sharon wink.

Let s be a constant denoting the thread Sharon, and let $W(x)$ denote the predicate 'x winks'. 'Jack saw Sharon' can now be modelled as $\exists(x=s|\mathcal{H})$ and 'Jack saw Sharon wink' as $\exists(x=s\wedge W(x)|\mathcal{H})$ (conditional quantifiers only work on real formulas). Then one may have

$$\exists(x=s\wedge W(x)|\mathcal{H}) < \exists(x=s|\mathcal{H})$$

even though $W(s)$ (cf. example 5).

4.1.3 Boolean connectives

It is easily seen that non-veridical perception validates the following inferences

i) $\text{SEE}(a,\varphi\wedge\psi) \leq \text{SEE}(a,\varphi)\wedge\text{SEE}(a,\psi)$

ii) $\text{SEE}(a,\varphi\vee\psi) = \text{SEE}(a,\varphi)\vee\text{SEE}(a,\psi)$

iii) $\text{SEE}(a,\neg\varphi) = \neg\text{SEE}(a,\varphi)$

$\exists(\text{SEE}(b,\psi)|\mathcal{A}_a)$. This would require a study of principles for $\exists(\exists(\psi|\mathcal{A}_b)|\mathcal{A}_a)$. E.g., $\text{SEE}(a,\text{SEE}(a,\varphi)) = \text{SEE}(a,\varphi)$ by 'taking out what is known' and $\text{SEE}(a,\varphi)\in\mathcal{A}_a$. This should be compared with Kamp 1984.

Principle (i) follows from the monotonicity of $\exists(\bullet|\mathcal{A}_a)$. It's converse only holds for veridical perception, as we shall see shortly.

Principle (ii) is just the additivity of $\exists(\bullet|\mathcal{A}_a)$, but sometimes a more subtle analysis is available as well. Although a predicate A and its negation $\neg A$ determine disjoint sets of assignments, the filtered predicates $\exists(A|\mathcal{H})$ and $\exists(\neg A|\mathcal{H})$ overlap if $A \notin \mathcal{H}$. We interpret a statement $\text{SEE}(j, A(x))$—'Jack sees that x is A'—by means of a formula $\exists(A(x)|\mathcal{H})$ with $\mathcal{H}$ given by Jack's approximation. The world assigns value $f(x)$ to x, and we take $\mathcal{M} \models \exists(A(x)|\mathcal{H})\ [f]$ to mean that 'I see that $f(x)$ is A (with filter $\mathcal{G}$)'. Now we can see more precisely how the validity of the inference from (33) to (34) depends on the kind of disjunction.

(33) Jack saw Sharon smile or not smile.

(34) Jack saw Sharon smile or Jack saw Sharon not smile.

The inference appears to be invalid if 'or' is interpreted exclusively, because we may be uncertain whether Sharon is smiling or not. If 'or' is interpreted inclusively and 'smile' is interpreted as a filtered predicate, the latter feature is reproduced exactly.

Notice that Mr. X, who made his appearence in section 2.2.3, would be statisfied with principle (iii); it is based on a treatment of negation where precepts are essentially positive. An alternative would be to interpret negative formulas in the manner of Feferman 1984 by means of positive and negative extensions. According to lemma 8 this would work because requirements of non-overlap are Horn; but we have to leave the details to the reader. A non-monotonic treatment of negation will be given in section 4.2.1.

4.1.4 Quantifiers

Recall the following pair of sentences, illustrating in- and exportation of $\exists$

(26) Jack saw a girl swim.

(27) A girl is such that Jack saw her swim.

We remarked that neither of these sentences seems to imply the other, if one drops veridicality. On the other hand, the usual translation into first order logic, or even generalised quantifier logic, does give equivalence (cf. van der Does 1991). We shall now translate these sentences in the framework of conditional quantification; it is then easily seen under what conditions one or both of the two implications fail.

So far we did not have occasion to apply ordinary quantifiers to formulas involving conditional quantifiers, as we need to do now.

Let V be the set of assignments corresponding to Jack's restricted perceptual field, and let $\mathcal{H}$ be an algebra on V, given by Σ, filtering his field of vision. Whenever this is convenient, we shall assume that V is determined by the formula $D(x) \equiv S(j, x)$ ('D' for 'domain'). Sentence (26) then becomes '$\text{SEE}(j, \exists x(G(x) \wedge S(x)))$'. Disregarding $D(x)$ for a moment, this has interpretation (26'a), which was shown to be equivalent to (26'b), just after theorem 7

(26') a. $\exists(G(x) \wedge S(x)|\mathcal{H}_x)$
b. $\exists x\exists(G(x) \wedge S(x)|\mathcal{H})$

Further, (27) is translated into '$\exists x(G(x) \wedge \text{SEE}(j, S(x)))$' which means

(27') $\exists x(G(x) \wedge \exists(S(x)|\mathcal{H}))$

In the direction from (26') to (27'), we distinguish several cases.

Case (a) $G(x) \in \mathcal{H}$: 'Jack recognises a girl when he sees one, and he can see all girls'. By 'taking out what is known' we have

$$\exists(G(x) \wedge S(x)|\mathcal{H}) = G(x) \wedge \exists(S(x)|\mathcal{H})$$

whence also

$$\exists x\exists(G(x) \wedge S(x)|\mathcal{H}) = \exists x(G(x) \wedge \exists(S(x)|\mathcal{H}))$$

Case (b) $G(x) \wedge D(x) \in \mathcal{H}$, where D may be nontrivial: 'Jack recognises a girl when he sees one, but he might not see all girls'. Let us first try to derive (26') from (27').
We have

$$\exists(G(x) \wedge D(x)|\mathcal{H}) = G(x) \wedge D(x)$$

Therefore

$$\begin{aligned} & D(x) \wedge G(x) \wedge \exists(S(x)|\mathcal{H}) \\ & = \exists(G(x) \wedge D(x)|\mathcal{H}) \wedge \exists(S(x)|\mathcal{H}) \\ & = \exists(\exists(G(x) \wedge D(x)|\mathcal{H}) \wedge S(x)|\mathcal{H})) \\ & = \exists(G(x) \wedge D(x) \wedge S(x)|\mathcal{H}) \\ & \subseteq \exists(G(x) \wedge D(x)|\mathcal{H}) \end{aligned}$$

Now notice that $\exists(S(x)|\mathcal{H})$ should be a subset of $D(x)$, since $\mathcal{H}$ represents Jack's perceptual field. Then actually

$$G(x) \wedge \exists(S(x)|\mathcal{H}) \subseteq \exists(G(x) \wedge D(x)|\mathcal{H})$$

whence $\exists x(G(x) \wedge \exists(S(x)|\mathcal{H})) \models \exists x\exists(G(x) \wedge D(x)|\mathcal{H})$.

This result may seem surprising; couldn't it be the case that there is a girl whom Jack perceives as swimming, without actually being aware that it is *she* who swims? No, because the anaphor 'her' in the scope of 'see' in (27) is taken to imply that the girl is imported in Jack's visual field; since he correctly identifies girls, (26) follows. Of course, we have just restated the above proof in plain English.

Case (c) $G(x) \notin \mathcal{H}$, V is the full set of assignments: 'Jack can see all girls, but he cannot identify them correctly'. In this case the derivation of (26') from (27') fails by an argument similar to example 4 in section 3.3. Let $\mathcal{H}$ be such that

$$G(x) \wedge S(x) = \mathbf{0} \in \mathcal{H}$$

but for every $A \in \mathcal{H}$ such that $\forall x(S(x) \to A(x))$:

$$\exists x(A(x) \wedge G(x))$$

(This is only possible when $G(x) \notin \mathcal{H}!$). Then $\neg\exists x\exists(G(x) \wedge S(x)|\mathcal{H})$, but $\exists x(G(x) \wedge \exists(S(x)|\mathcal{H}))$. Notice that in this situation, (27) could be formulated more accurately as 'A girl is such that Jack sees *it* swim'.

We now turn to the other direction, from (26') to (27'). Here the analysis has to be considerably more subtle. Consider the sentence $\exists x(G(x) \wedge \exists(S(x)|\mathcal{H}))$. If it is false, then $G(x) \wedge \exists(S(x)|\mathcal{H}) = \mathbf{0}$. However, $\mathbf{0}$ is an element of any Boolean algebra; hence $\mathbf{0}$ covers $G(x) \wedge S(x)$ and we have $\exists(G(x) \wedge S(x)|\mathcal{H}) = \mathbf{0}$. So, for instance, if there are no swimming girls in Jack's perceptual field, then Jack would know this. This is an irksome consequence of the fact that $\exists(\varphi|\mathcal{G}) \in \mathcal{G}$ and that $\mathbf{0}$ is allowed as a possible upper estimate for φ, an idealisation that is useful otherwise. Nevertheless, the failure of the implication from (26') to (27') can easily be established when we revert to the approximating models $\mathcal{M}_s$. We now model Jack's suboptimal perception of swimming girls by

$$\mathcal{M}_s \models \exists x(G(x) \wedge S(x))$$

hence there exists g such that $\mathcal{M}_s \models \exists x(G(x) \wedge S(x))$ $[g]$. By surjectivity of the projections π_s we may assume that there exists a thread $\xi = (\xi_t)_{t\in T}$ such that $g = \xi_s$. If Jack did the best he could as regards perceiving swimming, we have $f_\xi \in \exists(S(x)|\mathcal{B}_s)$, but we may easily have $\mathcal{M} \not\models G(x)$ $[f_\xi]$. Only in the case when $G(x) \in \mathcal{B}_s$ does it follow that $A \models G(x)$ $[f_\xi]$.

If we now look back at the previous argument, which seemed to show that (26') does imply (27'), we can see clearer why it clashes with intuition: the argument is by contradiction, hence constructively not acceptable. We submit that the intuitive reading of the implication 'if (26) then (27)' is constructive; (26) should provide evidence for (27), so given x satisfying $\exists(G(x) \wedge S(x)|\mathcal{H})$, transform it into an x satisfying $G(x) \wedge \exists(S(x)|\mathcal{H})$. This is manifestly impossible, so the implication should fail. It is possible to recast this argument for the non-validity of 'if (26') then (27')' in the framework of conditional quantifiers, provided one adopts a forcing definition of truth, but we shall leave the matter here.

This section has shown that non-veridical perception reports already uses much of the intricacies of the logic of vision, but the situation becomes even more interesting in case of veridical perception.

4.2 Veridical perception

Recall that we took the expression 'I see an arm' to mean the conjunction of (i) and (ii).

i) 'with the present approximation the object that I focus on is identified as an arm',
ii) 'I expect this to be the case for every more refined approximation'.

The second condition is evidently non-monotonic: more precise information may contradict the expectation expressed in (ii).

The expectation that a percept is stable under an increasing series of ever finer approximations can be captured by means of a non-monotonic rule governing the conditional quantifiers.

Definition 18 Let $\langle T, \leq \rangle$ be a directed set. $C \subseteq T$ is *cofinal* in T if for every $t \in T$, there is $s \in C$ such that $t \leq s$. A set of algebras $\{\mathcal{H}_s : s \in T\}$ indexed by T is called a *martingale* if $s \leq t$ implies $\mathcal{H}_s \subseteq \mathcal{H}_t$. The martingale property can be stated equivalently in terms of conditional quantifiers as follows: $s \leq t$ implies for all φ, $\exists(\varphi|\mathcal{H}_t) \subseteq \exists(\varphi|\mathcal{H}_s)$.

We model reality as the inverse limit of the inverse system $\langle \{\mathcal{M}_t : t \in T\}, \{\eta_{st} : s, t \in T,\ t \leq s\} \rangle$. The collection of algebras $\{\mathcal{B}_t : t \in T\}$ then forms a martingale; this follows from the commutativity of the diagram in definition 12.

In section 3.4 we argued that positive primitive information φ, when evaluated on a degree of approximation $\mathcal{M}_s$, can also be represented as $\exists(\varphi|\mathcal{B}_s)$ on the inverse limit $\mathcal{M}$. For positive primitive φ we may then introduce a rule veridicality$_1$ formalising the expectation (ii):[10]

$$\frac{\exists(\varphi|\mathcal{B}_s)}{\exists \text{ cofinal } C \subseteq T(s \in C \ \&\ \forall t \in C \exists(\varphi|\mathcal{B}_t))}\ \mathrm{V}_1$$

We take a cofinal subset of T containing s rather than the upper cone of s in T because the series of ever finer approximations may contain gaps. Also, the restriction to positive formulas is natural, since veridicality fails for negative ones (section 2.2.2).

The rule is called 'veridicality' because we believe that the principle of veridicality discussed in section 2 is not a logical rule, but rather a default assumption, valid perhaps in most circumstances, but an assumption nonetheless.

One may object to this formulation that the principle of veridicality should sanction inferences of the form

$$\frac{\text{I see an arm}}{\text{There is an arm}}$$

without mentioning such things as approximations. Formally, this would correspond to the principle veridicality$_2$: 'from SEE(s, φ) infer φ', or

$$\frac{\exists(\varphi|\mathcal{B}_s)}{\varphi}\ \mathrm{V}_2$$

followed by $\exists$-introduction. The rule says 'if it is consistent to assume φ with respect to approximation s, then assume φ'. This is indeed the standard format for a non-monotonic rule; it is similar to, say, autoepis-

[10] This rule cannot be expressed in terms of 'SEE' because it involves quantification over a set of approximations.

temic logic (Moore 1985), in that consistency is represented by an object-language operator, but it differs from other approaches in that consistency always comes with a degree. Thus there is a natural internal notion of the 'strength' of premisses and conclusions; this should be compared to the external assignment of strength in Nute 1994.

From a logical standpoint the second rule seems more acceptable than veridicality$_1$, since there is no second order quantifier over cofinal subsets of the directed set.

The two rules are equivalent, however. This is the content of the 'martingale convergence theorem'. We give a formulation stripped of technical details; for the full formulation and the proof the reader is referred to van Lambalgen 1996.

Theorem 9 *Suppose $\{\mathcal{H}_s : s \in T\}$ is a martingale. Let $\mathcal{H}$ be the σ-algebra generated by $\{\mathcal{H}_s : s \in T\}$. Then for all formulas φ*

$$\lim_{s \in T} \exists(\varphi|\mathcal{H}_s) = \exists(\varphi|\mathcal{H})$$

That is, if f is an assignment such that for all $s \in T$, $f \in \exists(\varphi|\mathcal{H}_s)$, then $f \in \exists(\varphi|\mathcal{H})$; if there exists s_0 such that for all $s \in T$, $s \geq s_0$ implies $f \in \exists(\varphi|\mathcal{H}_s)$, then $f \in \exists(\varphi|\mathcal{H})$.

Now consider the martingale $\{\mathcal{B}_s : s \in T\}$. The least σ-algebra containing all the $\mathcal{B}_s$ is $\mathcal{B}$, the σ-algebra determined by the topology on the inverse limit $\mathcal{M}$; this follows from proposition 2.5.5 in Engelking 1989, p. 99. This remains true if s takes values in a cofinal subset $C \subseteq T$. An application of the martingale convergence theorem shows that

$$\lim_{t \in C} \exists(\varphi|\mathcal{B}_t) = \exists(\varphi|\mathcal{B})$$

However, since $\varphi \in \mathcal{B}$, we have

$$\lim_{t \in C} \exists(\varphi|\mathcal{B}_t) = \varphi$$

Hence the two rules are equivalent.

This is useful, because we can now understand why 'see', with meaning given by conditions (i) and (ii) above, is often assumed to distribute over conjunction, at least when applied to positive formulas. Indeed, the rule 'from $\text{SEE}(s,\varphi)$ and $\text{SEE}(s,\psi)$ conclude $\text{SEE}(s,\varphi \wedge \psi)$' is valid, non-monotonically

$$\cfrac{\cfrac{\exists(\varphi|B_s) \wedge \exists(\psi|B_s)}{\varphi \wedge \psi}}{\exists(\varphi \wedge \psi|B_s)}$$

4.2.1 Negation

In the previous section we took care to formulate our results for positive formulas; this is because the occurrence of negation in perception reports seems to necessitate a special treatment. Recall the piece of dialogue in (25), from section 2.

(25) 'Did you see that hawk there?'
'I saw *something*, but it was not a hawk.'

Again, it makes perfect sense to retract the last sentence by saying 'No, you're right, it *is* a hawk'. *Prima facie*, this exchange causes a problem for the analysis presented here, because by definition of the inverse limit, if $\mathcal{M}_s \models \neg H(b)$, there is no refinement ξ of b in $\mathcal{M}$ such that $\mathcal{M} \models H(\xi)$.

A solution of this difficulty can be found by analysing the sentence 'I saw *something*, but it was not a hawk' as a conclusion arrived at by means of a non-monotonic argument. In a model $\mathcal{M}_s$ we have a predicate H, which is decomposed in various ways in models $\mathcal{M}_{t_1}, \ldots, \mathcal{M}_{t_k}$, for $s \leq t_1, \ldots, t_k$. By themselves, these models can be thought of as specifying features of hawks, such as colour, form-in-motion, or sound; as viewed from the inverse limit $\mathcal{M}$, which represents reality, they act as filters which to a greater or lesser extent blur that reality. Non-monotonically, we can now argue as follows: 'If this were a hawk, then it would have such-and-such colour, it would shriek like this, etc. It shows none of these features, so it cannot be a hawk.' However, this conclusion might well be mistaken, e.g., because we did not take into account all possible refinements $\mathcal{M}_{t_i}$ (we might have overlooked the refinement corresponding to a juvenile coat),[11] or because the distance between us and the bird was such that we could not actually apply the filters $\mathcal{M}_{t_1}, \ldots, \mathcal{M}_{t_k}$ to the bird, but rather applied much coarser filters.

5 Conclusion

In this article we have developed a semantics for direct perception reports based on a logical variant of Marr's theory of vision (Marr 1982). The main idea was to develop a non-monotonic logic for veridical perception, on the basis of approximate, non-veridical perception. To formalise this idea, we modelled reality as an inverse limit of first order models which refine each other in different ways. Reality is taken to be infinitely precise. We have shown in detail that the informal findings on the logic of perception which we took as our point of departure, have formal counterparts in the model developed. Besides, the study of principles for negation, veridicality, and the scope behaviour of quantifiers, has led to unexpected new insights.

[11]A good example of this kind of situation is furnished by a type of buzzard, *buteo rufinus*, whose colour shows two phases: a common light phase and a rare chocolate brown, almost black, phase. Not knowing that the latter phase exists may easily lead to misidentification.

References

Allen, P.K. 1987. *Robotic Object Recognition through Vision and Touch.* Dordrecht: Kluwer Academic Publishers.

Asher, Nicholas M., and Daniel Bonevac. 1985. How Extensional is Extensional Perception? *Linguistics and Philosophy* 8:203–228.

Asher, Nicholas M., and Daniel Bonevac. 1987. Determiners and Resource Situations. *Linguistics and Philosophy* 10:567–596.

Barwise, Jon. 1981. Scenes and other Situations. *Journal of Philosophy* LXXVIII:369–397. Also in Barwise 1989.

Barwise, Jon. 1989. *The Situation in Logic.* Lecture Notes, Vol. 17. Stanford, California: CSLI Publications. Distributed by Cambridge University Press.

Boden, Margaret A. 1988. *Computer Models of Mind.* Problems in Behavioural Sciences, Vol. 4. Cambridge: Cambridge University Press.

Chang, C.C., and H.J. Keisler. 1990. *Model Theory.* Studies in Logic and the Foundations of Mathematics, Vol. 73. Amsterdam: North-Holland. Third edition. First edition 1973, second edition 1977.

Engelking, Ryszard. 1989. *General Topology.* Sigma series in Pure Mathematics, Vol. 6. Warzawa: Heldermann. Revised and completed edition.

Feferman, Solomon. 1984. Towards Useful Type-Free Theories I. In *Truth and the Liar Paradox*, ed. R.L. Martin. 237–287. Oxford: Oxford University Press.

Hendriks, Herman. 1993. *Studied Flexibility. Categories and Types in Syntax and Semantics.* Doctoral dissertation, ILLC, University of Amsterdam., Amsterdam. ILLC Dissertation Series 1993–5.

Higginbotham, James. 1983. The Logic of Perceptual Reports: An Extensional Alternative to Situation Semantics. *Journal of Philosophy* LXXX:100–127.

Kamp, Hans. 1984. A Scenic Tour through the Land of Naked Infinitives. Unpublished manuscript, Univeristy of Texas, Austin.

Koons, Robert C. 1996. Information, Representation and the Possibility of Error. In *Logic, Language, and Computation*, ed. Jerry Seligman and Dag Westerståhl. Chap. 22, 333–345. Lecture Notes. Stanford, California: CSLI Publications.

Landman, Fred. 1986. *Towards a Theory of Information.* Doctoral dissertation, Department og Philosophy, University of Amsterdam, Amsterdam. Also published with Foris.

Lindström, Per. 1966. First-order Predicate Logic with Generalized Quantifiers. *Theoria* 32:186–195.

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

Moore, R.C. 1985. Semantical considerations on nonmonotonic logics. *Artificial Intelligence* 25:75–94.

Mostowski, Andrzej. 1957. On a Generalization of Quantifiers. *Fundamenta Mathematicae* 44:12–36.

Muskens, Reinhard. 1989. *Meaning and Partiality.* Doctoral dissertation, University of Amsterdam, Amsterdam.

Nute, D. 1994. Defeasible Logic. In *Handbook of Logic in Artificial Intelligence and Logic Programming (vol. 3)*, ed. Dov M. Gabbay, C.J. Hogger, and J.A. Robinson. Chap. 10, 353–395. Oxford: Clarendon Press. Volume Co-ordinator: D. Nute.

Russell, Bertrand. 1988. The Philosophy of Logical Atomism. In *Logic and Knowledge*, ed. Robert Charles Marsh. 177–281. London: Unwin Hyman Ltd. First published in 1956.

van der Does, Jaap. 1991. A Generalized Quantifier Logic for Naked Infinitives. *Linguistics and Philosophy* 14:241–294.

van Lambalgen, Michiel. 1996. Conditional Quantification, or Poor Man's Probability. Technical report. Department of Mathematics, Computer Science, Physics, and Astronomy (WINS), University of Amsterdam. In preparation.

Wright, Fred B. 1963. Boolean Averages. *Canadian Journal of Mathematics* 15:440–455.

Contents

institute *for* logic, language *and* computation

ILLC Research Reports and Technical Notes

Coding for Reports and Dissertations: *Series-Year-Number*, with LP = Logic, Philosophy and Linguistics; ML = Mathematical Logic and Foundations; CT = Computation and Complexity Theory; X = Technical Notes; DS = Dissertations. All previous ILLC-publications are available from the ILLC bureau. For prepublications before 1994, contact the bureau.

CT-95-01 Marianne Kalsbeek, Yuejun Jiang *A Vademecum of Ambivalent Logic*

CT-95-02 Leen Torenvliet, Marten Trautwein *A Note on the Complexity of Restricted Attribute-Value Grammars*

CT-95-03 Krzysztof Apt, Ingrid Luitjes *Verification of Logic Programs with Delay Declarations*

CT-95-04 Paul Vitányi *Randomness*

CT-95-05 Joeri Engelfriet *Minimal Temporal Epistemic Logic*

CT-95-06 Krzysztof Apt, Rachel Ben-Eliyahu *Meta-variables in Logic Programming, or the Praise of Ambivalent Syntax*

CT-95-07 Frans Voorbraak *Combining Unreliable Pieces of Evidence*

CT-96-01 Peter van Emde Boas *The Convenience of Tilings*

LP-95-01 Marten Trautwein *Assessing Complexity Results in Feature Theories*

LP-95-02 S.T. Baban, S. Husein *Programmable Grammar of the Kurdish Language*

LP-95-03 Kazimierz Świrydowicz *There exist exactly two Maximal Strictly Relevant Extensions of the Relevant Logic R^**

LP-95-04 Jaap van der Does, Henk Verkuyl *Quantification and Predication*

LP-95-05 Nataša Rakić *Past, Present, Future and Special Relativity*

LP-95-06 David Beaver *An Infinite Number of Monkeys*

LP-95-07 Paul Dekker *The Values of Variables in Dynamic Semantics*

LP-95-08 Jaap van der Does, Jan van Eijck *Basic Quantifier Theory*

LP-95-09 Jeroen Groenendijk, Marin Stokhof, Frank Veltman *Coreference and Modality*

LP-95-10 Jeroen Groenendijk, Martin Stokhof, Frank Veltman *Coreference and Contextually Restricted Quantification*

LP-96-01 Renate Bartsch *Understanding Understanding*

LP-96-02 David Beaver *Presupposition*

LP-96-03 Theo M.V. Janssen *Compositionality*

LP-96-04 Reinhard Muskens, Johan van Benthem, Albert Visser *Dynamics*

LP-96-05 Dick de Jongh, Makoto Kanazawa *Angluin's Theorem for Indexed Families of R.E. Sets and Applications*

LP-96-06 François Lepage, Serge Lapierre *The Functional Completeness of 4-value Monotonic Protothetics*

LP-96-07 Frans Voorbraak *Probabilistic Belief Expansion and Conditioning*

LP-96-08 John Case *The Power of Vacillation in Language Learning*

LP-96-09 Jaap van der Does, Willem Groeneveld and Frank Veltman *An Update on* Might

LP-96-10 Jelle Gerbrandy and Willem Groeneveld *Reasoning about Information Change*

LP-96-11 Renate Bartsch *Propositional Attitudes in Dynamic Conceptual Semantics*

LP-96-12 Paul Dekker *Reference and Representation*

LP-96-13 Rens Bod *Data-Oriented Language Processing: An Overview*

LP-96-14 Michiel van Lambalgen, Jaap van der Does *A Logic of Vision*

ML-95-01 Michiel van Lambalgen *Randomness and Infinity*

ML-95-02 Johan van Benthem, Giovanna D'Agostino, Angelo Montanari, Alberto Policriti *Modal Deduction in Second-Order Logic and Set Theory*

ML-95-03 Vladimir Kanovei, Michiel van Lambalgen *On a Spector Ultrapower of the Solovay Model*

ML-95-04 Hajnal Andréka, Johan van Benthem, István Németi *Back and Forth between Modal Logic and Classical Logic*

ML-95-05 Natasha Alechina, Michiel van Lambalgen *Generalized Quantification as Substructural Logic*

ML-95-06 Dick de Jongh, Albert Visser *Embeddings of Heyting Algebras (revised version of ML-93-14)*

ML-95-07 Johan van Benthem *Modal Foundations of Predicate Logic*

ML-95-08 Eric Rosen *Modal Logic over Finite Structures*

ML-95-09 Hiroakira Ono *Decidability and Finite Model Property of Substructural Logics*

ML-95-10 Alexei P. Kopylov *The Undecidability of Second Order Linear Affine Logic*

ML-96-01 Domenico Zambella *Algebraic Methods and Bounded Formulas*

ML-96-02 Domenico Zambella *On Forcing in Bounded Arithmetic*

ML-96-03 Hajnal Andréka, Johan van Benthem & István Németi *Modal Languages and Bounded Fragments of Predicate Logic*

ML-96-04 Kees Doets *Proper Classes*

ML-96-05 Søren Riis *Count(q) versus the Pigeon-Hole Principle*

ML-96-06 Angelo Montanari, Alberto Policriti *A Decidable Theory of Finitely-Layered Metric Temporal Structures*

ML-96-07 Angelo Montanari, Adriano Peron, Alberto Policriti *Decidable Theories of ω-Layered Metric Temporal Structures*

ML-96-08 Johan van Benthem, Angelo Montanari, Giovanna D'Agostino, Alberto Policriti *Modal Deduction in Second-Order Logic and Set Theory - II*

ML-96-09 Angelo Montanari, Maarten de Rijke *Decidability in Metric Temporal Logic*

ML-96-10 Vladimir Kanovei *On a Dichotomy related to Colourings of Definable Graphs in Generic Models*

ML-96-11 Domenico Zambella *Forcing in Finite Structures (revised version of ML-96-02)*

ML-96-12 Jon Barwise and Johan van Benthem *Interpolation, Preservation, and Pebble Games*

ML-96-13 Lex Hendriks *Intuitionistic Propositional Logic with only Equivalence has no Interpolation*

X-95-01 Sophie Fischer, Leen Torenvliet *The Malleability of* TSP_{2Opt}

X-96-01 Ingmar Visser *Mind Rules: a Philosophical Essay on Psychological Rules and the Rules of Psychology*

DS-95-01 Jacob Brunekreef *On Modular Algebraic Protocol Specification*

DS-95-02 Andreja Prijatelj *Investigating Bounded Contraction*

DS-95-03 Maarten Marx *Algebraic Relativization and Arrow Logic*

DS-95-04 Dejuan Wang *Study on the Formal Semantics of Pictures*

DS-95-05 Frank Tip *Generation of Program Analysis Tools*

DS-95-06 Jos van Wamel *Verification Techniques for Elementary Data Types and Retransmission Protocols*

DS-95-07 Sandro Etalle *Transformation and Analysis of (Constraint) Logic Programs*

DS-95-08 Natasha Kurtonina *Frames and Labels. A Modal Analysis of Categorial Inference*

DS-95-09 G.J. Veltink *Tools for PSF*

DS-95-10 Giovanna Cepparello *Studies in Dynamic Logic*

DS-95-11 W.P.M. Meyer Viol *Instantial Logic. An Investigation into Reasoning with Instances*

DS-95-12 Szabolcs Mikulás *Taming Logics*

DS-95-13 Marianne Kalsbeek *Meta-Logics for Logic Programming*

DS-95-14 Rens Bod *Enriching Linguistics with Statistics: Performance Models of Natural Language*

DS-95-15 Marten Trautwein *Computational Pitfalls in Tractable Grammatical Formalisms*

DS-95-16 Sophie Fischer *The Solution Sets of Local Search Problems*

DS-95-17 Michiel Leezenberg *Contexts of Metaphor*

DS-95-18 Willem Groeneveld *Logical Investigations into Dynamic Semantics*

DS-95-19 Erik Aarts *Investigations in Logic, Language and Computation*

DS-95-20 Natasha Alechina *Modal Quantifiers*

DS-96-01 Lex Hendriks *Computations in Propositional Logic*

DS-96-02 Angelo Montanari *Metric and Layered Temporal Logic for Time Granularity*

DS-96-03 Martin H. van den Berg *Some Aspects of the Internal Structure of Discourse: the Dynamics of Nominal Anaphora*

DS-96-04 Jeroen Bruggeman *Formalizing Organizational Ecology*