Intellectica, 1992 1/2, 13-14, pp. 79-103

Andrea BONOMI

Persistent Truths

I. A Cartesian Background

In his *Fourth Meditation* and in some letters Descartes presents a theory of judgement in which the reference to the limitations of our knowledge plays an essential role. The picture he suggests is the following.

Each of us is provided with a restricted perspective from which he can get only partial information about the world. We can imagine expanding this perspective to a wider one (so that more things are known), then to another one, and so on. But each of these states of information is in essence limited; it cannot coincide with God's point of view (where everything is either true or false in a definite way), because at each step only a *finite* amount of information is added. Moreover, at any step it is possible that some judgements assumed as true turn out to be false, in view of the expansion of knowledge. So the question is: how to "restrain" the faculty of judging within the limits of evident or certain ideas, if there is no guarantee about the persistence of judgements, no obvious way of telling what is really certain (and therefore true) from what is only assumed to be such? Now, my opinion is that the *Cogito* argument is supposed to give a sample of evidence which presents, according to Descartes, a guarantee of *persistence*[1].

But this is not the point on which I am going to concentrate. For the present purpose it is enough to emphasize two features of the Cartesian theory of judgement: (i) the opposition between the "completeness" of God's perspective (where, so to speak, a principle of epistemic bivalence holds: for every statement A, either A is known to be true or A is known to be false) and the "narrowness" of our segments of information; (ii) the requirement that certainty entails persistence through increasing segments. In the next section I shall try to show that neglecting this requirement is one of the causes which give rise to Gettier's problem.

II. Gettier's problem: a diagnosis

The necessary and sufficient condition for the truth of a sentence of the form:

(1) *s* knows that A

is often given in terms of the following clauses:

(2) (i) A is true
(ii) *s* believes that A
(iii) *s* is justified in believing that A.

Now, Gettier's problem shows that what (2) expresses is only the necessary condition for the truth of (1): (2) is not in itself strong enough to also supply the sufficient condition, *unless the concept of*

---

1 An analysis of the *Cogito* argument from this point of view is presented in Bonomi (1990).

justification is properly specified. Here is one of the two counterexamples discussed by Gettier[2].

Suppose that s has some good reasons to be convinced of the truth of:

(3) *t* is the so-and-so
(3') *t* is P.

Therefore s is apparently justified in believing that:

(4) The so-and-so is P.

But suppose also that actually *t'* — *not* t — is the so and so and that *t'* is P, although *s* is completely in the dark about this[3]. Then we have: (4) is true (even if (3) — from which *s* inferred (4) — is false); *s* believes that (4) is true; and, finally, *s* seems to be justified in believing that (4) is true. So, according to (2), it would be true that s *knows* that the so and so is P, which is absurd, for s has no reliable information about it: what he relies on is only a mistaken basis of inference.

We must then wonder why, without suitable qualifications, (2) is inadequate. And the answer I suggest is that the concept of justification is too vague. How can we make it more definite? Well, let us look at the very structure of (2). If for the time being we disregard (iii) — which refers to the notion to be clarified —, the remaining points indicate two possible directions of inquiry. As a matter of fact, (i) has to do with the way in which *s*'s knowledge is supported by relevant *facts*. And plain truth is the proposed requirement. On the other hand, (ii) appeals to *the subject's* beliefs. The problem is that (2) says nothing interesting from either point of view: truth (in (i)) is a very rough requirement (since our beliefs can be true by mere accident). And the

---

2 See Gettier (1963).
3 As a matter of fact *s* might even not to know that *t'* exists.

same holds of the requirement in (ii), for no reference is made to the inferential structure of s's system of beliefs. Due to this coarseness on both sides, such a definition as (2) can give rise to different kinds of predicaments. And my idea is that behind the so-called Gettier's problem, in its countless versions, there actually lurk *different* (although correlate) problems. And, perhaps, different kinds of solution must be searched for.

In what follows I shall deal with the question of "external" adequacy raised by (i). The reason of this choice is that in Gettier's original puzzle, that I have just summarized, the deduction made by s is perfectly sound. No failure can be detected in the inferential structure of his / her system of beliefs. So the problem lies completely on the side of (i), for it has to do with the information the subject has about the current state of the world.

I shall try to solve this problem by referring to a *more "fine-grained" notion of truth* which accounts for the pieces of evidence available to *s*. In this way a solution of Gettier's puzzle can be found independently of any consideration about the inferential requirements which s's system of beliefs has to meet. (Yet a reference to these requirements is needed to cope with more sophisticated versions of the puzzle[4].)

The concept of local truth I am going to introduce is not the classical one (as in 2.i), but it is strong enough to guarantee *persistence* through increasing segments of information, which was mentioned in the previous section as a possible criterion of justifiability. That concept will be formalized in the next section. For the time being let us consider again Gettier's example. Suppose then that the point of view from which s sees the world is just a little bit richer and that the *only* additional information it incorporates is the truth of

---

4 In formal terms this means that a proper system of axioms should be introduced not only for the truth-operator 'T' (see sect. V), but for the operator 'B' as well.

(5) *t* and the so-and-so are different individuals.

It is plausible to argue that in such circumstances s would not believe that the so-and-so is P, for he could not believe that t and the so-and-so are the same individual, which was an essential premise of the inference at issue. So, what happens is that, in this case, believing that the so and so is P is not preserved through growing segments of information. In other words, the condition which is not satisfied here is the persistence, in any better state of knowledge, of the reasons supporting *s*'s judgement. And that this condition is an essential part of our intuitive characterization of knowing is wittnessed by the oddness of:

(6)? I know that A, but A might be false.

What is not acceptable in (6) is the idea that, with more information available, we can give up something we *know*. This is what is missing in (2), where the truth-condition for '*s* knows that A' does not entail the requirement of persistence. But it is time to see what kind of formal treatment can suit the main ideas I have been expounding.

III. Segments of information

Since the problems we are dealing with are independent of the iteration of such epistemic operators as 'B' and 'J', to simplify things we shall assume that in our formal language L^J these operators apply only to formulas in which they do not occur[5], that is formulas of a standard first-order language L.

So let L be any first-order language with a denumerable set C of individual constants. '¬', '∧' are its primitive connectives, and '(∀x)'

[5] P. Casalegno has extended this semantic framework to the case of iterated operators.

(where 'x' stands for an individual variable) its universal quantifier. The notion of well-formed formula (wff) of L is defined inductively by the usual rules.

Now, assuming that T is the set of the terms of L (individual variables and constants, in our case), let L^J be an extension of L which is characterized by the following additional rule:

If $s \in T$ and A is a wff of L, then B(s,A), J(s,A) are wffs of L^J.

Associated to such a language as L is the usual notion of (first-order) L-model, i.e. a pair $M=\langle D,I \rangle$, where D is a domain of entities and I a function such that $I(t) \in D$ if t is an individual constant and $D^n \supseteq I(P^n)$ if P^n is a predicate letter (for every $n \geq 1$). If $M=\langle D,I \rangle$ is such a model, a submodel of M is any pair $\mu=\langle D,I_\mu \rangle$, such that, for every individual constant t, $I_\mu(t)=I(t)$ and, for every predicate letter P^n, $I_\mu(P^n)=\langle X,Y \rangle$, where $I(P^n) \supseteq X$, $D^n \backslash I(P^n) \supseteq Y$ (that is, Y is a subset of the complement of $I(P^n)$: therefore $X \cap Y = \emptyset$ and $D^n \supseteq X \cup Y$; if $D^n \neq X \cup Y$ for some predicate P^n such that $I_\mu(P^n)=\langle X,Y \rangle$, μ will be said to be a *proper* submodel of M). From now on X and Y will be designated respectively by '$I_\mu{}^+(P^n)$' (the "extension" of P^n in μ) and '$I_\mu{}^-(P^n)$' (the "counterextension" of P^n in μ). Let $\mathbf{M}$ be the set of the submodels of M and $\leq$ the relation on $\mathbf{M}$ such that $\mu \leq \nu$ iff, for any P^n, $I_\nu{}^+(P^n) \supseteq I_\mu{}^+(P^n)$ and $I_\nu{}^-(P^n) \supseteq I_\mu{}^-(P^n)$. Notice that $\leq$ is a partial order with a least element (which is the submodel where both the extension and the counterextension of every predicate letter are empty) and a greatest one (where, for every predicate letter, its extension and its counterextension cover the entire domain; this submodel is M itself). In fact, $\langle \mathbf{M},\leq \rangle$ is a lattice: for every $\mu,\nu \in \mathbf{M}$, their join is the submodel σ such that, for every predicate P^n, the extension of P^n in σ is the union of the extensions of P^n, respectively, in μ and ν, whilst the counterextension of P^n in σ is the union of the counterextensions. (To get the dual definition of the meet of μ and ν, replace union by intersection.) This lattice is determined by a Boolean algebra <**M**,0,1,*,∪,∩>, where:

0 and 1 are respectively the least element and the greatest element described above;

μ^* (the complement of μ) is the submodel such that, for every P^n, the extension of P^n in μ^* is $I(P^n)\backslash I_\mu^+(P^n)$, and its counterextension is $(D^n\backslash I(P^n))\backslash I_\mu^-(P^n)$;

$\mu \cup v$ is the submodel described as the join of μ and v;

$\mu \cap v$ is the submodel described as the meet of μ and v.

Given the extended language L^J, a model M^J for L^J is a structure <M,**P**,[ ] ,Φ,F>, where:

(i) M = <D,I> is a standard (first-order) L-model (**M** is the set of its submodels, in the sense defined above);

(ii) **P** is the set of propositions[6];

(iii) [ ] is a function from the set of wffs of L to the set **P** of propositions. So if A is a wff of L, [A] is the proposition associated with it;

(iv) Φ is a function from D to **M** That is, Φ assigns to each[7] individual u a submodel μ in **M**. When $\mu=\Phi(u)$, μ is called a *local structure for u*; μ represents the segment of information (about the complete model M) available to u.

---

[6] For our purposes it is not essential to determine what propositions are. In a suitable property theory they might be 0-place relations (in intension). Roughly speaking, if A is a wff and π a sequence of (occurrences of) variables, '$[A]_\pi$' is a term which denotes a property if π is a sequence of a single (occurrence of) variable, a binary relation if π is sequence of two (occurrences of) variables, etc. So [A] (with π empty) would denote a 0-place relation. This notation will be used to designate propositions (in the semantic metalanguage).

[7] It would be more natural to restrict the domain of Φ (and of F also) to the set of individuals that can have intentional attitudes. But this modification is not relevant here.

(v) F is a function from D to the power set of **P**. In other terms, F assigns a set of propositions to each individual u in D.

As to the intuitive justifications of such a framework, we can conceive the standard first-order model M as a *complete* picture of a certain state[8] of the world. As specified in (iv), a *partial* submodel μ in **M** can be assimilated to a segment of information about this state of the world M. Notice that, since all the submodels have the same domain D (i.e. the domain of the complete model M), partiality, here, is treated in terms of lack of local information about individuals, rather than in terms of their possible "inexistence" in local domains[9]. Every epistemic subject s is linked, through Φ, to a local structure Φ(s)=μ, where μ (a submodel of M) is the segment of information available to s. Finally, F specifies the set of beliefs[10] that s has in the given state of the world.

IV. A minimal system

In order to reconstruct Gettier's paradox from a formal point of view, I shall start from a very weak system, where the truth of 's knows that A' is simply defined in terms of the truth of A (as in (2.i)) and the truth of 's believes that A' (as 2.ii)). The paradox will be

---

[8] At some given instant. For the sake of simplicity, time is left aside in the present framework. Local structures must be conceived as representing increasing information about a *single* state of the world.

[9] But in a different approach, which is of course possible, the domains of the submodels can be proper subsets of the "absolute" domain D.

[10] No restriction is imposed here on F with respect to μ. But, for example, it would be quite natural to require, by suitable axioms, that, given a local structure Φ(s)=μ, if A expresses a simple state of affairs which obtains in μ, then [A]∈F(s) (so that it turns out to be true that s believes that A). Or to impose that F(s) be closed, e.g., under conjunction (so that if [A]∈F(s) and [B]∈F(s), then [A∧B]∈F(s). The point is that, in the present paper, my purpose is not to propose a proper axiomatization for the belief-sentences, but only to analyze the relations between beliefs and guaranteed justifiability. See sect. IV and V.

removed by replacing (2.i) with a stronger condition, which is incorporated into the formal system of sect. 5.

Coming now to the formal definitions, let M^J=<M,**P**,[ ],Φ,F> be a model for the language L^J (where M=<D,I> is a standard first-order L-model), and μ any submodel in **M**. Furthermore, let *a* be an assignment of values in D to the individual variables (whilst a^u_x is the assignment such that $a^u_x(x)=u$ and $a^u_x(y)=a(y)$ if y is distinct from x, where u∈D). Finally let E be a mapping of the terms into the domain D such that E(t)=*a*(t) if t is a variable and E(t)=I(t) if t is a constant. We shall define now what it means for a wff A (of L) to be validated (respectively, invalidated) in μ by *a* - in symbols: μ,*a* ||— A (resp. μ,*a* —|| A) -, and what it means for a wff A (of the extended language L^J) to be satisfied in M^J by *a* - in symbols: M^J,*a* ||— A.

$\mu,a \Vdash P^n(t_1,\ldots,t_n)$ iff $<E(t_1),\ldots,E(t_n)> \in I_\mu^+(P^n)$

$\mu,a \dashv\vdash P^n(t_1,\ldots,t_n)$ iff $<E(t_1),\ldots,E(t_n)> \in I_\mu^-(P^n)$

μ,*a* ||— ¬A iff μ,*a* —|| A

μ,*a* —|| ¬A iff μ,*a* ||— A

μ,*a* ||— A ∧ B iff μ,*a* ||— A and μ,*a* ||— B

μ,*a* —|| A ∧ B iff μ,*a* —|| A or μ,*a* —|| B

μ,*a* ||— (∀x)A iff μ,a^u_x ||— A for every u∈D

μ,*a* —|| (∀x)A iff μ,a^u_x —|| A for some u∈D

M^J,*a* ||— $P^n(t_1,\ldots,t_n)$ iff $<E(t_1),\ldots,E(t_n)> \in I(P^n)$

M^J,*a* ||— ¬A if not M^J,*a* ||— A

$M^J, a \Vdash A \wedge B$ iff $M^J, a \Vdash A$ and $M^J, a \Vdash B$

$M^J, a \Vdash (\forall x)A$ iff $M^J, a^u_x \Vdash A$ for every $u \in D$

$M^J, a \Vdash B(s,A)$ iff $[A] \in F(E(s))$ (i.e., iff [A] belongs to the set of propositions F assigns to the individual denotated by 's')

$M^J, a \Vdash J(s,A)$ iff $M^J, a \Vdash A$ and $M^J, a \Vdash B(s,A)$.

(Notice that the last clause — which gives the truth-condition concerning the operator 'J' — makes no reference to local structures[11].)

A wff A of L is validated (resp. invalidated) by the submodel μ - in symbols: $\mu \Vdash A$ (resp. $\mu \dashv\vdash A$) - iff $\mu, a \Vdash A$ (resp. $\mu, a \dashv\vdash A$) for every assignment a. Moreover a wff A of the extended language L^J is true (resp. false) in the model M^J - in symbols: $M^J \Vdash A$ (resp. $M^J \dashv\vdash A$) - iff $M^J, a \Vdash A$ for every (resp. no) assignment a. Finally, A is valid (in symbols: $\Vdash A$) iff $M^J \Vdash A$ for every model M^J. It is easy to prove that, since M^J can be seen as a "total" evaluation function, for every sentence A of L^J either $M^J \Vdash A$ or $M^J \dashv\vdash A$ (whilst there are local structures μ and sentences A such that neither $\mu \Vdash A$ nor $\mu \dashv\vdash A$).

Assuming some standard axiomatization of the predicate calculus (PC), let us now consider a very weak theory T^J whose axioms are all

---

[11] This means that the semantics illustrated in the present section makes an *empty* use of the set of submodels. (So, by omitting Φ, a model M^J might be defined as a structure <M,**P**,[ ],F>.)

Yet submodels are introduced at this early stage in order to present the semantics of sect. V as an extension of the structure defined here. Thus, suitable comparisons will be made possible thanks to this unified framework.

the wffs of L^J which are either axioms of PC or of one of the following forms:

(A1) $J(x,A) \rightarrow A$
(A2) $J(x,A) \rightarrow B(x,A)$
(A3) $A \wedge B(x,A) \rightarrow J(x,A)$[12].

The soundness of T^J can be proved easily. The axioms and rules of PC will be ignored since the models for L^J behave classically with respect to PC. So only (A1)-(A3) will be considered. On the other hand, that (A1)-(A3) are valid follows immediately from the truth-conditions for 'J(s,A)'.

The completeness proof is omitted.

V. Removing the inadequacy: certainty as local truth

In the simple example of sect. II the situation is roughly the following:

(i) *s* is convinced of the truth of sentences A and B.
(ii) C is an immediate consequence of A and B.
(iii) Because of (i) and (ii), *s* believes that C.
(iv) A is false.
(v) C is true.

In such circumstances, we said, it would be questionable to evaluate the sentence

(*) *s* knows that C

---

12 In view of further developments, the different principles are kept separate. But a single axiom (expressing this minimal view) would of course be sufficient here: $J(x,A) \leftrightarrow A \wedge B(x,A)$.

as true. Now, this is exactly what happens in the semantics I have just presented. According to the relevant definition, (iii) and (v) are a *sufficient* condition for the truth of (*). So our problem is to present a semantics (in the same framework as before) where the notion of (justified) knowledge is more restrictive. As a result, in the circumstances (i)-(v), (*) would be falsified in a principled way.

Our analysis will rest on the following remark: in the semantics of sect. IV the submodels in **M** (or local structures, as we called them) are an unexploited resource. They are not mentioned in the truth-condition concerning the epistemic operator 'J'. How they can play a significant role is shown by a new reflection on Gettier's example in terms of the present theoretical framework.

Statements (i)-(v) depict a situation where the subject *s* has some (possibily erroneous) beliefs and a set of certainties[13] about the world. In terms of our semantics, this set is simulated by μ (a partial submodel of the complete model M). Now, coming back to Gettier's example, C cannot be counted among the propositions validated by μ, i.e. the propositions whose truth is ascertainable on the basis of the data available to *s*. It will be remembered that the reason why s feels committed to the truth of C (i.e. the sentence 'The so-and-so is P') is that s assumes the truth of A, and C is an immediate consequence of A (and other assumptions). But A, in its turn, cannot be validated by μ. So we have:

(a) $M \Vdash C$; (b) not $\mu \Vdash C$.

(a) says that C is true, whilst (b) specifies that C is *not* validated by s's segment of information: more exactly, that it is not true in μ (which

[13] In the present context, 'to be certain' is treated as a "factive" predicate That is, if I have the certainty that P, P is true. In the same way, 'information' means truthful information.

does not mean, of course, that C is false in μ). This is the critical point disregarded by the truth-condition for 'J' given in sect. IV.

The argument presented at the outset is an instructive one from this point of view. Let us resume it briefly in terms of the formal notions at issue. Suppose that ν is like μ (s's segment of information), but with the additional information that A is false (i.e. $\mu<\nu$ and $\nu \dashv\!\vdash$ A). Now, if s were in the segment, he/she would question the truth of C (for A, whose truth was believed to entail the truth of C, turns out to be false) and the persistence of s's epistemic attitude towards C would be lost. But in the sense I have in mind, *s* is *justified* in saying 'I know that C', given his / her segment of information μ, only if no larger segment ν would lead s to change her / his view: i.e. only if μ is *already* rich enough to validate C.

As to the truth-condition we have to modify in our semantics, the requirement which is to be specified is that C be validated in the partial submodel μ (i.e. the subject's segment of information). In other terms, in the recursive definition of sect. IV, the old clause concerning the operator 'J' must be replaced by the following:

(K) $M^J, a \Vdash$ J(s,A) iff $\mu, a \Vdash$ A and $M^J, a \Vdash$ B(s,A), where $\mu=\Phi(E(s))$ (i.e., μ is the submodel in $\mathbf{M}$ that Φ assigns to the individual denotated by *s*).

It is easy to see that, with this new definition, (A3) is no longer valid. (To falsify it, think of a situation where s believes that A, A is true, but A is not true in s's segment of information μ. As expected, the situation described by Gettier is the case in point!)

In our example, since C is not validated by the relevant segment of information, (*) turns out to be false in the circumstances described by (i)-(v), which is the intuitive outcome we pursued. So, the crucial notion is: *being validated in the subject's segment of information* or, in terms of our semantics, being true in the relevant local structure.This notion of local truth can be axiomatized in different ways[14]. I intend to choose one that involves a new operator. The extended language which is to be used is the same as L^J except for this new formation rule:

If s is a term and A is a wff of L, then T(s,A) is a wff.

That is, like 'J' and 'B', 'T' applies to a term *s* and to a formula A of the given first-order language (i.e. a formula without[15] occurrences of 'J', 'B' and 'T' itself) to generate the formula 'T(s,A)'. The intuitive meaning of such a formula is: A is true in the submodel assigned to *s*, A is validated by the subject's local structure. Formally, given a model $M^J = \langle M, \mathbf{P}, [\ ], \Phi, F \rangle$, the relevant truth-condition that is to be added is the following:

(CT) $M^J, a \Vdash T(s,A)$ iff $\mu, a \Vdash A$, where $\mu = \Phi(E(s))$ (i.e. μ is the submodel in $\mathbf{M}$ that Φ assigns to the individual denotated by *s*).

In this extended language, the truth-condition (K) concerning the operator J can also be expressed in the following terms:

(K') $M^J, a \Vdash J(s,A)$ iff $M^J, a \Vdash B(s,A)$ and $M^J, a \Vdash T(s,A)$.

Here is a complete set of axioms for this modified semantics. The first one accounts for the new truth-condition (K'):

(B1) $J(x,A) \leftrightarrow B(x,A) \wedge T(x,A)$.

A second axiom is connected with the monotonicity of the partial semantics determined by the local structures in $\mathbf{M}$: if a sentence of L

14 P. Casalegno has given a complete set of axioms for this notion in the language L^J itself, i.e. without new operators.

15 The reason of this simplification is, once more, that the iteration of the epistemic operators is not considered here, although an extension of the present framework in this direction is available.

turns out to be true in a submodel (assigned to any individual *s*), it will be true in every larger submodel, in particular in M itself (the "complete" given model). This axiom is:

(B2) $T(x,A) \rightarrow A$.

The remaining axioms express necessary and sufficient conditions for evaluating formulas in local structures (using negation, conjunction and universal quantification as primitive):

(B3) $T(x,A \wedge B) \leftrightarrow T(x,A) \wedge T(x,B)$

(B4) $T(x,\neg(A \wedge B)) \leftrightarrow \neg(\neg T(x,\neg A) \wedge \neg T(x,\neg B))$
(or, equivalently: $T(x,A \vee B) \leftrightarrow T(x,A) \vee T(x,B)$)
(B5) $T(x,\neg\neg A) \leftrightarrow T(x,A)$
(B6) $T(x,(\forall y)A) \leftrightarrow (\forall y)T(x,A)$
(B7) $T(x,\neg(\forall y)A) \leftrightarrow \neg(\forall y)\neg T(x,\neg A)$
(or, equivalently: $T(x,(\exists y)A) \leftrightarrow (\exists y)T(x,A)$).

Once again, the completeness proof is omitted.

A proper analysis of the notion of certainty as local truth would involve other related concepts: in particular, evidence and necessity. Their connections — in the formal framework adopted here — will be the topic of a separate paper. Here are some informal hints.

VI. Evidence, certainty and necessity: a preliminary view

What about this new characterization of (justified) knowledge? The new truth-condition (K) expresses the restriction mirrored in (B1): knowing that A entails, for a subject s, that A is validated by the subject's segment of information. Independently of other conditions (the condition that s believes that A, in particular), this restriction neutralizes an aspect of the problem of logical omniscience. As a matter of fact, suppose that A is a tautology formed by atomic sentences which are not validated by the relevant local structure. Therefore, in virtue of the new clause, the assertion that s knows that A is falsified, for A is not validated by that structure. The "absolute" knowledge of logical truths is not presupposed.

But the question of logical omniscience is still lurking. Take for example this situation, suggested by S. Zucchi (personal communication):

(I) p and q are simple atomic sentences whose plain truth is completely evident to *s*.

(II) for *erroneous* reasons s believes that not(not p or (not q and p)). [For example, *s* is convinced that this complex sentence is a tautology.]

According to the new truth-condition (K), the sentence

(*) *s* knows that not(not p or (not q and p))

is true, for the embedded sentence is: (i) true; (ii) believed by s; (iii) true in the subject's segment of information μ (because p and q are true in μ).

There can be reasons for saying that a treatment which evaluates (*) as true in the circumstances described by (I)-(II) is inadequate. It depends on what you have in mind about the notion of (justified) knowledge. Unfortunately this is not a clear-cut concept in *intuitive* terms. Even the weak (traditional) definition formalized in sect. IV can be appropriate with respect to some pre-theoretical uses of the word.

As to the problem of inadequacy raised by Zucchi's example, what I want to stress is the *different* nature of this alleged inadequacy with respect to the limitations of the traditional definition. In Gettier's example, this definition is responsible for the following predicament: we are allowed to ascribe to s the property of knowing that C (the sentence 'The so and so is P') *although* C is nothing s *could verify* on the basis of the available information. (See (i)-(v) in sect. V, where the only constituent of C — i.e. C itself — is *not* validated by s's local structure.) But this is not the case of our last example, where p and q — the constituents in the statement believed by *s* — are validated by the relevant submodel. (See (I) and (II)). In other terms, what *s* believes is supported here by the given basis of information, whilst in Gettier's example there is no such possibility of verification. The relevant *data* are available to the subject in one case, but not in the other one.

Therefore, even if more restrictive than the traditional definition, the characterization of (justified) knowledge given in (K) is liberal enough to evaluate (*) as true in the circumstances described by (I) and (II)[16]. We have just seen in what sense this outcome can be plausible. But, in another sense, (K) is still restrictive. Take, once more, Descartes' analysis of judgements. In the *Regulae* he makes a distinction between *evidence* and *certainty*. What is evident must be grasped by a single direct act of the mind (intuition). It is something *given* and simple. But certainty can also be *acquired* by means of complicated deductive chains: 'Many things are known with certainty, though not by themselves evident, but deduced from true and known principles by the continuous and uninterrupted action of a mind [...]. Hence we distinguish this mental intuition from deduction by the fact that into the conception of the latter enters a certain movement or succession, into that of the former there does not.'[17] This is why a complex statement, for example, can be certain without being evident. So, Descartes' distinction reminds us that different levels of justification can be individuated:

---

[16] To prevent this outcome some requirements on the inferential structure of the subject's set of beliefs must be introduced. This can be done by giving proper axioms for the operator 'B'.

[17] Rule III (Descartes, 1985: 8).

evidence — the status of some (privileged) *elementary* statements, whose truth is grasped by a *simple* mental act;

certainty — the status of:
(i) evident statements; or
(ii) (complex) statements whose ultimate components are verified, but whose truth is grasped by some mental *operation*;

necessity — the status of:
(i) statements that are certain; or
(ii) statements whose truth is grasped (by purely logical arguments) *independently* of the verifiability of their components.

In our framework, evident statements (for a subject *c*) can be described as (selected[18]) atomic statements *validated* by μ, i.e. by *c*'s segment of information. On the other hand, certainty can be ascribed to statements of any complexity, provided that they are validated by μ. Because of the monotonicity of the partial semantics determined by our submodels, the truth of a statement in a given submodel is *persistent* through growing submodels. So certainty, in the particular sense selected here, is the notion the present paper focuses on, and is formally expressed by the operator 'T' (local truth).

Finally, how to characterize necessity is a question I intend to face in connection with the problem of finding weaker criteria of justification.

The discussion of Zucchi's example showed that, roughly speaking, the criterion formalized here (after the introduction of the operator 'T') is certainty, in the special sense defined above. In particular, we have just seen that it is liberal enough not to filter out the example (*). But it can be considered restrictive from another point of view.

---

18 To say *what* kind of statements can play this role is part of the philosophical metatheory. For example, according to Descartes the statement 'I exist' in the *Cogito* argument is meant to represent a sample of evidence.

Suppose I have in front of me a pack of cards. I take out a covered card. In this situation, it might be argued, it would be plausible to say I *know* that the card is either black or red, although I have evidence for neither of these alternatives. But in our system of sect. V this is not allowed. For, let a be the intended card, B the property of being black and R the property of being red. Now, the truth of 'J(c,Ba∨Ra)' entails the truth of 'T(c, Ba ∨ Ra)', which in its turn entails either the truth of 'Ba' or the truth of 'Ra' in *c*'s segment of information. But neither holds. So the statement that (in the given circumstances) *c knows* that the card is either black or red is falsified.

Shall we conclude that the criterion of justification formalized by 'T' (what we called certainty) is implausible? Not necessarily. After all there is *no* clear *pre-theoretical* idea of justification which can be referred to. Beyond a certain point, our intuition is helpless. A theory is needed. Now, axioms (B1) - (B7) capture (in part) a situation where the relevant criterion is *certainty*, in the above sense. But, in *this* sense, the certainty of a statement — as local truth in the segment of information μ — entails the validation by μ of (some of) its components. So certainty, as defined here, entails the "downward"[19] closure of the set of the statements that it characterizes. And, *within* this theoretical characterization, it is quite plausible to say that to know, for example, that A or not-A we must have enough information to verify A or to verify not-A.

But is it possible, in our semantical framework, to characterize a more liberal criterion of justification, so that necessity (in the epistemic sense captured by the third notion in the above classification) — instead of certainty — is referred to? To answer this question[20], suppose that,

---

[19] For example, downward closure of local truth under disjunction and conjunction is given respectively by axioms (B3) and (B4) — in the left to right sense of the biconditional.

[20] The details of this solution will be studied in a separate paper. What I present here is only a rough sketch.

given a submodel $\mu=<D,I_\mu>$ in **M**, a *local completion* of μ is defined as a structure $\sigma=<D,I_\sigma>$ where, for any P^n, $I_\sigma^+(P^n)\supseteq I_\mu^+(P^n)$, $I_\sigma^-(P^n)\supseteq I_\mu^-(P^n)$ and $I_\sigma^+(P^n)\cup I_\sigma^-(P^n)=D^n$. That is, σ is a "total" structure which extends μ. We are now in a position to define the necessity of a statement A (with respect to μ) as the truth of A in all the local completions of μ. In this sense, a statement can be "necessary" (with respect to the segment of information μ) although its components are not validated by μ. This holds, for example, of the statement 'A or not A', which is true in all the completions of μ although it is not validated by μ (because A is neither true nor false in μ). So, if in our axiom (B1) the formula 'T(x,A)' (which expresses the condition that A is "certain" — i.e. validated — in the relevant segment of information μ) is replaced by the formula 'N(x,A)' (which expresses the "necessity" of A with respect to μ), the truth of 'J(x,A$\vee$B)' does not entail the validation by μ of either of the statements A and B any more. And I am justified in *knowing* that the covered card is black or red although neither alternative is verified by my segment of information.

Thus, the notions of evidence, certainty and necessity (as defined above) individuate three possible criteria of justification, according to a decreasing order of narrowness. If the strongest condition (i.e. evidence) is assumed, all those beliefs — like in Zucchi's example — which do not coincide with (selected) elementary states of affairs are filtered out. This is not true of certainty. On the other hand, this second criterion can still filter out beliefs that are not founded on *verified* information, as in the example of the covered card. Finally, this restriction is removed if the third criterion — necessity — is referred to, although it is still strong enough to block Gettier's (original) counterexample.

Appendix. A remark on quantifiers and monotonicity

Let us consider the extensional part of L^J — i. e. L — and the semantics for L which was built in relation to the set $\mathbf{M}$ of the submodels of a given standard model M. This semantics is monotonic. As a matter of fact, if $\mu \leq \mu'$ and if a sentence A of L — i.e. a sentence where there is no occurrence of intensional operators — is true (false) in μ, then A is still true (false) in μ'.

A probable objection is the following. If $\mu, \mu' \in \mathbf{M}$ and $\mu \leq \mu'$, this means that μ is, in its turn, a submodel of μ'. In other terms, μ is "smaller" than μ'. Therefore, what is expected is that a universal sentence which is true in μ can turn out to be false in μ'. From an intuitive point of view, such a sentence as 'All the books are on the desk' can be true in that particular segment of reality which is my office, but false in a larger one (the Department where I work, for example). Natural languages, in this sense, are not monotonic. So — it might be argued — the monotonicity of your semantics is guaranteed by an unnatural device: according to your definitions, for a sentence like '$(\forall x)Px$' to be true in a submodel μ, the open formula 'Px' must be satisfied by *all the individuals in D*, which is the domain of the "absolute" model M. This means that what is referred to is not the restricted class of objects which is relevant in μ, but the whole universe. In this way monotonicity is insured — since there are no unpleasant surprises when passing from μ to the larger structure μ' —, but the picture you present is counterintuitive. As a matter of fact the truth of such a sentence as 'All the books are on the desk' entails that the relevant domain is the set of things in my office, not the whole universe of things.

That is true. For the sake of simplicity, the treatment of the quantifiers I have presented here is narrow. But this restriction can easily be removed. The point is that, in a generalized version of this semantics, the submodels in $\mathbf{M}$ can be conceived of as contexts which determine what is the relevant universe for evaluating quantified sentences. So, suppose that statements, rather than sentences, are evaluated — where, according to an idea suggested by Bar-Hillel, a statement is an ordered pair $\langle A,\sigma\rangle$ formed by a sentence A and a context σ (a submodel in **M**, in our framework). And if the truth-conditions are designed to deal with statements, in the case of universal quantification we can have something like this:

$\mu,a \Vdash \langle(\forall x)A,\sigma\rangle$ iff $\mu,a^{u}_{x} \Vdash \langle A,\sigma\rangle$ for every $u \in D^*$, where D^* is the (partial) universe of the submodel σ.

Intuitively speaking, the submodel μ is used to evaluate the sentence, whilst σ is referred to as a proper context in order to fix the relevant universe of quantification. So, our semantics can precisely account for the phenomena mentioned in the objection at issue. But this problem is the topic of a separate paper, where indexical terms, quantifiers and some uses of definite descriptions are treated from the point of view of partial submodels. For the time being, I shall content myself with concentrating on this question: is it really true that examples like the one produced by our objector call for a non-monotonic treatment of quantifiers?

Imagine this situation. I live in the countryside. My family includes five nice dogs. Unfortunately, some of them are extremely lively and, when they run off, the cats in the neighborhood are in serious danger. All of a sudden I hear some dog bark in the distance. I look anxiously my wife, who has just inspected the garden, but she reassures me. 'All the dogs are sleeping' she says.

It is not true, of course, that all the dogs, in the absolute sense, are sleeping. As a matter of fact, there are some barking very loudly. So what shall I conclude? That my wife is temporarily deaf? Or that she told me a deliberate lie? In the circumstances I described above, it is obvious that neither conclusion is appropriate. My interpretation of her statement is that *all the dogs we have in mind* — i.e. the set of dogs selected with respect to a particular situation — are sleeping. Under this natural interpretation, my wife's statement is true. But notice: that all *these* dogs are sleeping is still true if a larger segment of reality is taken into account.

What is interesting, in our trivial example, is that when I accept my wife's statement as *true*, the situation I am referring to is large enough to involve dogs that are *not* sleeping. (Am I not hearing some dog bark?) But this characteristic does not lead me to consider the statement false. The point is that, besides this large segment of the world with respect to which the statement is evaluated, there is another one — used as a context — that I have in mind when I must select the relevant universe of quantification. And this is a smaller segment. As a result, it turns out to be true, with respect to the more inclusive situation, that *all* the dogs individuated by the smaller situation are sleeping. More in general: if a class X of objects, individuated by a given segment σ, has the property P in μ, then X keeps on having the property P in μ', where $\mu \leq \mu'$. If the submodels in **M** are treated not only as structures which provide a (partial) evaluation of the sentences, but also as contexts of use (which, combined with sentences, form *statements*), it is easy to prove that, if $\mu \leq \mu'$ and the statement $<(\forall x)Px,\sigma>$ is true[21] in μ, then this statement is also true in μ'[22]. In other terms, in this case universal quantification does not preclude monotonicity. And this is intuitive, since — in our example — if all the dogs individuated by a

[21] μ, the submodel used to evaluate the sentence, can coincide with the context σ itself, for contexts are in general submodels. (See note 25). So, in our example, if σ is the smaller segment of reality, it is true in σ that all the dogs individuated by σ are sleeping. And that *these* dogs (whose individuation still depends on σ) are sleeping is true also in the larger segment, where other dogs are not sleeping.

[22] The idea is that the submodels in **M** play a double role: as contexts, they contribute to fix the referents of the terms, the domain of the quantifiers, etc.; as model-theoretic structures they determine the proper "local" truth-conditions. As I have already said, this approach is to be developed in a separate paper.

given segment of reality have the property of sleeping, that those dogs are sleeping is still true when a larger segment of reality (where, possibly, other dogs come on the scene) is taken into account.

As a consequence, we need not give up monotonicity. What we need is a treatment of the quantifiers which allows our semantics to express this reference to a context. And the system of submodels in **M** makes this treatment possible. In some sense, the truth-conditions presented in sect. IV are restricted to the borderline case: i.e. when the context is the absolute model M itself, so that the whole domain D is selected. As I said, it is possible — and interesting — to remove this restriction. But what is important to emphazise, here, is that this move does not necessarily entail the rejection of monotonicity.

So partial domains (of the submodels) and monotonicity are not incompatible. As a consequence, the semantics presented in sect. IV is susceptible to interesting developments without undermining our assumptions about monotonicity. As you will remember, in that semantics the incompleteness or partiality of the submodels does *not* have to do with *ontology*[23] (for we have an invariant domain), but, so to speak, with information or knowledge. In a submodel an individual can be totally indeterminate with respect to properties and relations (if there is no information about it), and this epistemological sense of "inexistence" is perhaps more interesting than others. But now this view can also be expressed in "ontological" terms, without losing monotonicity. Take a semantics which allows the local domains to be (proper) subsets of the domain D in the "complete" model M, e.g. by defining "local" existence (with respect to a subject s) as follows:

(DET) $E(s,t)$ iff $(\exists P)(T(s,P(t)) \vee T(s, \neg P(t)))$.

[23] There are no non-denoting terms (as for example in free logic) with respect to the local domains.

This means that t "exists" in the segment of information μ assigned to s iff there is some property P such that either P or its negation can be definitely ascribed to t in μ. Incidentally, (DET) fits the conceptual framework illustrated in sect. I, since it is very close to the so-called Thomas' principle assumed by Descartes[24].

The semantical counterpart of this definition is the following:

(DEG) $\mu \Vdash E!(x)$ if and only if for some predicate P, either $\mu \Vdash P(x)$ or $\mu \Vdash \neg P(x)$.

That is, the "local" existence of x in a submodel is expressed in terms of x's being either in the extension or in the counterextension of some predicate in μ. And the expansion of information through growing submodels is mirrored by a parallel expansion of the partial domains.

Andrea BONOMI
Dipartimento di Filosofia
Università di Milano
via Festa del Perdono, 7
20122 MILANO (Italia)

References

BONOMI, A. (1990). *The Cogito Revisited.* Mimeo: University of Milan.

DESCARTES, R. (1985). *The Philosophical Works.* (Ed. by E.S. Haldane: and G.R.T. Ross). Cambridge: Cambridge University Press.

GETTIER, E. (1963). Is Justified True Belief Knowledge ? *Analysis*, 23, 121-123.

---

[24] 'According to the laws of true Logic, the question "does a thing exist?" must never be asked unless we already understand *what* the thing is.' See *Reply to Objections I* (Descartes, 1985: 13).