Philosophical Studies (2005) 124: 313–330

FRANCESCO PAOLI

THE AMBIGUITY OF QUANTIFIERS

ABSTRACT. In the tradition of substructural logics, it has been claimed for a long time that conjunction and inclusive disjunction are ambiguous: we should, in fact, distinguish between 'lattice' connectives (also called additive or extensional) and 'group' connectives (also called multiplicative or intensional). We argue that an analogous ambiguity affects the quantifiers. Moreover, we show how such a perspective could yield solutions for two well-known logical puzzles: McGee's counterexample to modus ponens and the lottery paradox.

1. THE AMBIGUITY THESIS: FROM CONNECTIVES TO QUANTIFIERS

In the tradition of relevant and substructural logics, it has been argued at length (e.g. by Anderson and Belnap, 1975; or by Read, 1988) that the conjunction "and" and the *inclusive* disjunction "or" of ordinary English are ambiguous connectives. To summarize in a quick fashion such a thesis, let us focus on disjunction and consider Stalnaker's (1975) celebrated example:

(1) Either the butler or the gardener did it.

I can assert this sentence on either one of two different grounds. Assume that I know for sure that the butler is guilty of the crime at issue; this gives me the right to assert that the culprit was either the butler or the gardener, for the latter sentence follows logically from my conviction. Remark, however, that such an assertion can be true even if the gardener has a cast-iron alibi and is not even remotely a suspect: my sole ground for asserting the disjunction is my belief in the truth of one disjunct, the other one being, at least in this case, totally irrelevant to the truth value of the whole compound.

On the other side, suppose that I carried out some investigations and detected that there were only two individuals on the scene of crime – just the butler and the gardener. Even if I don't know, at present, who is the actual culprit, I can safely maintain that if it wasn't the butler, it was the gardener, and if it wasn't the gardener, it was the butler. This gives me the right to assert that the culprit was either the butler or the gardener – but the ground for such an assertion, this time, is completely different. Here, the disjuncts must be relevant to each other. On the other hand, I need not accept either the former or the latter: it is their mutual connection that produces the acceptance of the disjunction, not the previous acceptance of at least one of the disjuncts.

The preceding opposition has been termed in many different ways in the literature. This lack of terminological uniformity is hardly surprising given the fact that many vernacular traditions, each one with its own nomenclature, coexist side by side in the field of substructural logics. To name but a few examples, relevant logicians speak of *extensional vs intensional* connectives, while linear logicians employ the pair *additive/multiplicative*. However, since the former pair is too heavily loaded with past philosophical connotations and the latter is in our opinion quite misleading,[1] we prefer to follow Casari (1997) in his use of the terms *lattice-theoretical* and *group-theoretical*.[2]

Now, what is important is that our vague and informal remarks about "acceptance" and "assertion" are by no means the only way to characterize the distinction between lattice-theoretical and group-theoretical connectives, which can be made far more precise by specifying their *inferential role* in deductive arguments – e.g. in terms of introduction and elimination rules in the context of a natural deduction calculus. Thus, group-theoretical disjunction (for which we use the symbol $\oplus$) abides by disjunctive syllogism as an elimination rule and by a version of conditional proof as an introduction rule, while lattice-theoretical disjunction (for which we use the usual symbol $\vee$) obeys the proof-by-cases schema as an elimination rule and the principles of addition as introduction rules; as for conjunctions, the lattice-theoretical one (referred to by $\wedge$) has as elimination and introduction rules, respectively, simplification and a context-sensitive version of adjunction, while the group-theoretical one (referred to by $\otimes$) has, respectively, Schroeder-Heister's rule and a context-free version of adjunction (Table 1).

TABLE 1

Introduction and elimination rules for the substructural connectives

$\dfrac{\Gamma \Rightarrow A}{\Gamma \Rightarrow A \vee B} \quad \dfrac{\Gamma \Rightarrow B}{\Gamma \Rightarrow A \vee B}\ (\vee I)$	$\dfrac{\Gamma \Rightarrow A \vee B \quad A, \Delta \Rightarrow C \quad B, \Delta \Rightarrow C}{\Gamma, \Delta \Rightarrow C}\ (\vee E)$
$\dfrac{\Gamma, \neg A \Rightarrow B}{\Gamma \Rightarrow A \oplus B}\ (\oplus I)$	$\dfrac{\Gamma \Rightarrow \neg A \quad \Delta \Rightarrow A \oplus B}{\Gamma, \Delta \Rightarrow B}\ (\oplus E)$
$\dfrac{\Gamma \Rightarrow A \quad \Gamma \Rightarrow B}{\Gamma \Rightarrow A \wedge B}\ (\wedge I)$	$\dfrac{\Gamma \Rightarrow A \wedge B}{\Gamma \Rightarrow A} \quad \dfrac{\Gamma \Rightarrow A \wedge B}{\Gamma \Rightarrow B}\ (\wedge E)$
$\dfrac{\Gamma \Rightarrow A \quad \Delta \Rightarrow B}{\Gamma, \Delta \Rightarrow A \otimes B}\ (\otimes I)$	$\dfrac{\Gamma \Rightarrow A \otimes B \quad A, B, \Delta \Rightarrow C}{\Gamma, \Delta \Rightarrow C}\ (\otimes E)$

In the light of our preceding informal discussion, it is easy to become convinced that addition holds for $\vee$ but not for $\oplus$: if the butler did it, it does not follow that if it wasn't the butler, it was the gardener (who may be wholly unrelated to the misdeed). On the other hand, it is just as natural to say that disjunctive syllogism holds for $\oplus$ but not for $\vee$: if my ground for asserting that either the butler or the gardener did it is the belief that the butler did it, upon learning that the butler didn't do it I do not go on to infer that it was the gardener; rather, I may be willing to retract my previous disjunctive assertion. In such a perspective, Lewis's controversial "independent proof" of the *ex absurdo quodlibet* ($A \wedge \neg A \rightarrow B$) appears to be nothing more than a *fallacy of equivocation*: the disjunction which occurs therein is paralogistically taken to abide by both $\vee I$ and $\oplus E$, i.e. by the rules for *different* connectives.

So much for the thesis that the usual conjunction and disjunction *connectives* are ambiguous. Are the quantifiers affected by an analogous ambiguity? At least two reasons seem to suggest an affirmative answer. In the first place, it is commonly held that the universal and the existential quantifiers can be considered as sorts of infinitary counterparts, respectively, of conjunction and disjunction. If the latter constants are ambiguous, by analogy we might be inclined to suspect that this is the case for their infinitary mates as well. In the second place, if we apply the ordinary rules for quantifiers in peculiar contexts, we are in a position to develop paradoxical arguments which, although perhaps less spectacular, are akin in nature to the sentential paradoxes of material implication. For example, suppose we are informed about Oswald's plans to kill Kennedy in a solitary action, with no help by any accomplice. Later, we get conflicting information about whether he succeeded or not. From these data – by merely applying rules of quantification and identity – we are in a position to reach the unwarranted conclusion that *someone else* killed Kennedy:

(2)	Oswald killed Kennedy	assumption
(3)	Oswald did not kill Kennedy	assumption
(4)	Someone killed Kennedy	from (2)
(5)	Someone else killed Kennedy	from (3), (4)

In spite of these reasons, and although the ambiguity claim for connectives is widely agreed upon by researchers working in the areas of relevant and substructural logics, there is not a comparable agreement concerning the ambiguity of the universal and of the existential quantifiers. All the existing papers about quantification in relevant and substructural logics concern *lattice-theoretical* quantifiers;[3] group-theoretical quantifiers remain utterly mysterious. This state of affairs has its roots in a twofold difficulty: on the one side, specifying appropriate formal rules for such quantifiers seems a tricky issue; on the other, it is no easier to attach an intuitive meaning of whatever sort to them. There are good reasons to believe that the task of making sense of this distinction is the most important open problem in the area of substructural logics (Paoli, 2002).

2. LATTICE-THEORETICAL AND GROUP-THEORETICAL QUANTIFIERS

Indeed, we believe that a case for the ambiguity of quantifiers could be plausibly made. Consider once again the sentence (4) above. Let us try to figure out two possible contexts which could entitle a speaker to assert it.

Suppose first that you are an FBI officer in Dallas, soon after the president has been assailed. You can see that Kennedy has been shot and that he did not survive the attack, but you have no clue about the identity of the culprit. This certainly seems a sufficient ground to assert (4). On the other side, suppose that you are standing in the crowd, a short distance from Oswald, just before the crucial moment. You see Oswald drawing out his handgun, aiming at Kennedy, and shooting him. You can barely see the presidential car amidst the crowd, but you are pretty sure that the shot reached its target and had lethal effects. This also gives you a sufficient ground to assert (4). Two crucial questions arise: are the sentences respectively uttered, or believed, by the FBI officer in the former context (let us call it (4A)) and by the bystander in the latter (4B) just tokens of the same proposition, which is merely asserted for different reasons, or else are we in the presence of two distinct propositions? If the latter alternative is correct, is such an ambiguity lexical or structural?

At first sight, denying the ambiguity would seem a more plausible option: it avoids unnecessary proliferations of meanings and appears closer to common sense. Nonetheless, take the sentence

(6) If Oswald didn't kill Kennedy, someone else did

While (4A) seems to imply (6) – if Kennedy has been murdered, and if Oswald or any other suspect didn't do it, it means that someone else did – (6) does not appear to follow from (4B): the bystander's sole ground for asserting (4B), indeed, is the fact that he saw a particular person shooting. Should he learn that Oswald didn't kill Kennedy, he would not go on to infer the consequent of (6); rather, the warrant for his existential assertion would be undercut. He would have no special reason to believe that someone else in the crowd shot Kennedy: it would be more plausible for him to suppose that Oswald missed his target, or that the shot was not a deadly one.[4] Similarly to what happened in the example of the butler and the gardener, therefore, it is perfectly consistent to imagine a situation where (6) cannot be deduced from the assumption of (4B).[5] Since it is impossible that the same proposition both implies and fails to imply another one, it follows that (4) expresses different propositions according to the situation – i.e., it is ambiguous.

Now, even if we accept the ambiguity of (4), we are not yet entitled to claim that it originates from any of the lexical components in the sentence. It might be the case that, upon a more careful analysis of its logical form, we find out that some operator is involved which has wide scope in one case and narrow scope in the other: that is to say, we might be confronted with a *structural*, rather than *lexical*, ambiguity. For example, (4A) could be analyzed as $\Box\exists xK(x, k)$ (it is known that someone killed Kennedy), while (4B) could be rendered as $\exists x\Box K(x, k)$ (there is someone of whom it is known that he killed Kennedy). If we did so, the difference between (4A) and (4B) would be reduced to the famous distinction between *knowing who* and *knowing that* (Hintikka, 1962: the officer merely knows *that* Kennedy has been killed, while the bystander knows *who* killed Kennedy).

However, there is a problem with this approach. At least according to the traditional view of quantified epistemic logic, as well as according to common sense, the implication

$$\exists x\Box A(x) \rightarrow \Box\exists x A(x)$$

is valid. Since, as we have seen, (4A) implies (6), by transitivity of implication it would follow that also (4B) implies (6), which is not – as we argued above. The suggested analysis, as a consequence, must be faulty.

Thus, if (4) is ambiguous and such an ambiguity is lexical rather than structural, we must locate the ambiguity somewhere, and the only appropriate candidate seems the existential quantifier. We think that there are two different quantifiers at issue: (4A) contains a *group-theoretical* existential quantifier, for which we shall use the symbol Σ,[6] while (4B) contains a *lattice-theoretical* quantifier, which we shall refer to by means of the usual symbol $\exists$. In fact, the bystander's sole ground for asserting the existential is his belief in the truth of a single istance, the other ones being totally irrelevant to the truth value of the whole compound, while in the officer's sentence the instances must be relevant to one another, and it is the connection between such instances that produces the acceptance of the existential, not the previous acceptance of at least one instance. The analogies with the two modes of disjunction are striking; we therefore suggest to take Σ as the infinitary analogue of $\oplus$ and $\exists$ as the infinitary analogue of $\vee$.

To be sure, the transition from the sentential level to the first order level brings about an additional difficulty. Consider once again our FBI officer. He would be willing to assent to (4A) or to (6), but he certainly wouldn't to "If nobody else killed Kennedy, *I* did", since of course he would go as far as to rule out at least *himself* as a possible suspect. It seems clear, therefore, that the officer should not be included in the *domain* over which "someone" ranges in (4A). But how are we to specify, or circumscribe, such a domain? While it is natural to build these specifications into disjunctions, by explicitly naming each alternative, quantificational sentences can be expressed but in an abridged form. In ordinary discourse, therefore, the specification of the relevant domain will be often left implicit and will be mostly determined by contextual factors; in the context at issue, for example, "someone" can be taken to stand for "someone in the crowd", or some such description which leaves the officer out. Remark that a related problem arises in the field of adjectival semantics whenever the *comparison class* of a relative adjective (like "tall", "big") has to be identified.

If there are two modes of existential quantification, one does not see why *universal* quantification should not come in two different forms as well. After all, if sentences like (4) are ambiguous, it is plausible to surmise that also their *negations* be such, and negated existential sentences amount to universal sentences *via* the usual equivalences (to which substructural logics with involutive negations are committed). We thus assume that there is a group-theoretical universal quantifier, for which we use the symbol Π, and a lattice-theoretical universal quantifier, which we refer to by means of the symbol $\forall$. Then, what does it mean to deny (4)? Suppose first that "someone" is understood group-theoretically. Then (4) is, in a sense, about *all* the possible murderers of Kennedy: it means that if Oswald did not kill Kennedy, then if Alan didn't, then if Beatrice didn't, then . . . then Zoe did. To deny this, one has to provide, for each possible murderer, *independent* evidence that he/she didn't kill Kennedy. On the other hand, if "someone" is understood lattice-theoretically, (4) is not, in general, about all possible murderers; rather, it is about at *least one* individual x who is such that the belief that x killed Kennedy represents our ground for asserting (4). To deny this sentence, one has to provide evidence that x didn't kill Kennedy for *any* possible choice of such x. In game-theoretical terms: my opponent claims that someone killed Kennedy on the ground that some particular person did it; to confute her, I have to prove that x did not kill Kennedy for any particular x of her choice.

In the light of these remarks, we obtain additional clues about what is required to justify us in asserting sentences having the logical forms $\Pi x A(x)$, $\forall x A(x)$, $\Sigma x A(x)$, and $\exists x A(x)$. We are justified in asserting $\Pi x A(x)$ if, for each object in our domain of quantification, we have *separate* and *independent* evidence that it has the property expressed by A. We are justified in asserting $\forall x A(x)$ if, given an *arbitrary* object in our domain of quantification, we have evidence that it has the property expressed by A (unlike in the preceding case, thus, a *single* justification suffices to attribute A to all the objects in the domain). We are justified in asserting $\Sigma x A(x)$ if for any ordering of the objects in our domain of quantification we have evidence that, if the first object does not have A, then if the second does not have A, then ... then the last has A (and this makes sense at least for finite domains). Finally, we are justified in asserting $\exists x A(x)$ if for at least one object in our domain we have evidence that it has the property expressed by A.

What about the *rules* which should govern the inferential behaviour of our quantifiers? We shall not venture to suggest some. Suffice it to say that the standard introduction and elimination rules for the quantifiers in natural deduction calculi defuse the reference to infinity by the recourse to *variables*; this seems acceptable for the lattice-theoretical quantifiers, but not for the group-theoretical quantifiers, where the reference to infinity is a *real*, and not merely a *schematic*, one. It seems, then, that a proper treatment of such constants should inherently require the use of sequents whose antecedents contain infinitely many formulae, and of infinitary proof-trees. In any case, we are confident that the intuitive characterization of such quantifiers we tried to give in this paper could provide some hints which could turn out useful in the search for a rigorous systematization of the concept.

Now it is perhaps appropriate to add a *caveat*, lest the distinction between lattice and group quantifiers should be confused with other well-known dichotomies which superficially resemble it. It is tempting, for example, to borrow and extend to indefinite descriptions Donnellan's (1966) terminology concerning *definite* descriptions, in such a way as to identify the lattice-theoretical existential quantifier with a *referential* quantifier, and its group-theoretical mate with an *attributive* quantifier. As far as the Kennedy example is at issue, nothing seems to prevent such a move: in (4B), in fact, "someone" refers to the specific individual Oswald, whereas in (4A) it refers to whoever has the property of having killed Kennedy. We shall see in §3, however, that a more prudent attitude is recommendable, as it is possible to cast some doubts upon this identification and suggest that these dichotomies may significantly overlap but are not likely to coincide with each other.[7]

3. FIRST APPLICATION: MCGEE'S PARADOX

Let us now see how our ambiguity thesis can be put to good use in the solution of two well-known logical puzzles. Both paradoxes admit of variants where no quantified sentences occur (i.e. where existentials are replaced by disjunctions and universals are replaced by conjunctions), but since the ambiguity that affects connectives is strictly related to the ambiguity of quantifiers, as we have seen, our diagnosis would not change significantly – and, of course, the quantificational versions are more relevant to our primary concern here.

In a 1985 paper (McGee, 1985), Vann McGee points to a *prima facie* counterexample to modus ponens – not for the material conditional, but for the standard indicative conditional of English. He recounts the following story:

Opinion polls taken just before the 1980 election showed the Republican Ronald Reagan decisively ahead of the Democrat Jimmy Carter, with the other Republican in the race, John Anderson, a distant third. Those apprised of the polls believed, with good reason:
If a Republican wins the election, then if it's not Reagan who wins it will be Anderson.
A Republican will win the election.
Yet they did not have reason to believe:
If it's not Reagan who wins, it will be Anderson. (p. 472)

The puzzle poses an intriguing challenge. If modus ponens for the indicative conditional is not a reliable argument form, what mode of inference can be such? It is hard to find anything as undebatable as the inference from *A* and "If *A*, then *B*" to *B*. At first sight, however, McGee's counterexample seems incontrovertible. As always happens with any paradoxical argument, only three choices are open before us: either we claim that the argument is invalid, or we accept its conclusion, or else we reject one of the premisses. Let us examine these alternatives more closely.

The first option is embraced by McGee himself, who denies an absolute validity to modus ponens. As already remarked, this solution defies our most intimate convictions about logic, and can therefore be accepted only if no better explanation is forthcoming.[8]

The second route is taken by Sinnott-Armstrong et al. (1986), who adopt a typical Gricean escape: the sentence

(7) If it's not Reagan who wins, it will be Anderson

is true if we assume that Reagan is going to win, but has a low degree of assertability (it is "misleading"), since it is logically weaker than

(8) Reagan will win the election

which embodies the ground for asserting it. Apart from the independent objections that can be raised against the Gricean theory of implicature – on which we shall not dwell, since they are altogether well-known – we simply remark that such a route clashes with intuition: we feel, as most untutored speakers do, that there is a sense in which (7) is not only scarcely assertable, but plain *false* if there is a third competitor who has a better chance than Anderson of winning the race. Moreover, also a further claim by Sinnott-Armstrong et al. sounds unconvincing: they think that McGee has not refuted *real* modus ponens (if A and $A \rightarrow B$ are true, so is B), but at most an *epistemic* version of it (if A and $A \rightarrow B$ are *believed*, so is B). However, both humans and computers reason, more often than not, on the basis of assumptions; it is generally supposed that modus ponens cannot lead us astray in any circumstance, and not simply when such assumptions are hard facts. If I assume A and $A \rightarrow B$, I have every reason to think that I am in a position to conclude that B (under those assumptions, of course). If it is conceded that an argument like McGee's undercuts this belief, there is still enough to be worried about.

Finally, the third alternative is chosen by Katz (1999), according to whom

(9) If a Republican wins the election, then if it's not Reagan who wins it will be Anderson

is false. In fact, if we assume that Reagan will win the election, then the antecedent of (9) is true, while its consequent is false given the fact that Carter has better winning chances than Anderson. But it is generally agreed that a conditional with true antecedent and false consequent is false; therefore (9) is false. Like the preceding solution, also this one is at odds with our intuition: there is surely a sense in which (9) is true, for if a Republican different from Reagan should win the election, it would perforce be Anderson.

None of the explanations we considered, then, appears fully satisfactory. There is, however, an option we did not take into account. McGee's reasoning could be invalid, all right, but for a reason which is not the invalidity of modus ponens. Our suggestion is that the argument rests, exactly like Lewis' proof, on a fallacy of equivocation.[9] Consider the sentence:

(10) A Republican will win the election

This sentence occurs twice in the argument: once as the categorical premiss, and once as the antecedent of (9), the conditional premiss. The indefinite article "a", however, is used differently in these sentences: it is readily seen that in the categorical premiss it is used *lattice-theoretically,* while in the conditional premiss it is used *group-theoretically*. In the categorical premiss, in fact, my ground for believing that a Republican will win the election is my belief that Reagan is going to win, together with the knowledge of the fact that Reagan is, indeed, a Republican. The presence and the winning chances of other Republican candidates are immaterial under this respect. On the contrary, a possible ground for believing that a Republican will win the election, with "a" understood in its group-theoretical sense, could be given by a scenario where all the Republican candidates are well ahead of their Democratic competitors, although possibly none of them has a decisive edge over the others, so that if it's not the first Republican candidate who wins, then if it's not the second, then ... then it's the last one who wins (again, we remark that of course this must be the case for *any* ordering of the candidates). Notice moreover that (10) follows from (8) (or, to be precise, from "Reagan is a Republican and Reagan will win the election") by the rule $A(a) \vdash \exists x A(x)$, which is the quantificational counterpart of the addition rule – valid for lattice disjunction but not for group disjunction, as we saw at the outset. On the other hand, in (9), the consequent follows from the antecedent by means of what we identified above as an inferential pattern which is typical of the group-theoretical existential quantifier.

The premisses of McGee's argument, therefore, can be both true only if there is a fallacy of equivocation. In such a case, the argument has the form

$$A \to (B \to C), D \vdash B \to C$$

and so, though invalid, is not an instance of modus ponens. If the existential quantifier is interpreted lattice-theoretically in both premisses, (9) is false; if it is interpreted group-theoretically in both premisses, (10) is false. In both cases, the argument is unsound. Modus ponens seems to have been vindicated *ab omni naevo*.

We would like to conclude this section with a couple of remarks. First, McGee's paradox highlights in an illuminating way how an analysis of the lattice/group distinction for quantifiers in terms of an epistemic or doxastic operator ("it is known that ...", "it can be rationally believed that ...") having wide scope in one case and narrow scope in the other is bound to fail. Suppose in fact that (10) is analyzed as $\Box\exists x(R(x) \wedge W(x))$ in the conditional premiss and as $\exists x\Box(R(x) \wedge W(x))$ in the categorical premiss. McGee's argument would still go through as follows; let B be the formula $R(x) \wedge W(x)$:

$$\dfrac{\dfrac{\exists x\Box B \quad \exists x\Box B \to \Box\exists x B}{\Box\exists x B} \qquad \Box\exists x B \to \Box(\neg W(r) \to W(a))}{\Box(\neg W(r) \to W(a))}$$

Secondly, we still owe the reader an explanation of the reason why it is dubious to identify our distinction with Donnellan's one. In the original version of McGee's paradox, it seems pretty natural to say that the indefinite description "a Republican" is used referentially in (10) – because the speaker has some particular person in mind, namely Ronald Reagan – while it is used attributively in (9), for it points to whomever has the property of being a Republican. However, we can tweak the example so that the use of the description no longer looks referential, but the inference still goes through. Imagine a race with three Republicans – Reagan, Anderson, and Nixon – and a sole Democrat – Carter; and imagine that polls have shown Reagan and Nixon approximately tied for first place, with Carter a distant third and Anderson a still more distant fourth. Then one will believe both (10) and

(11) If a Republican wins the election, then if it's not Reagan or Nixon who wins, it will be Anderson

yet disbelieve

(12) If it's not Reagan or Nixon who wins, it will be Anderson.

Here there is no possibility that the speaker is using "a Republican" referentially, because he has no idea whether the winner will be Reagan or Nixon.

Does this variant of the paradox undermine our equivocation approach? We do not think so, because the ambiguity persists even here. For a start, remark that under the envisaged circumstances (10) is not at all about Anderson, while the antecedent of (11) is; they respectively mean

(13) Reagan will win the election $\oplus$ Nixon will win the election

(14) Reagan will win the election $\oplus$ Nixon will win the election $\oplus$ Anderson will win the election

Of course, (13) implies

(15) (Reagan will win the election $\oplus$ Nixon will win the election) $\vee$ Anderson will win the election

But this means that even reading (10) as (15) we cannot get rid of the ambiguity: (15) is accepted in virtue of the previous acceptance of one of the disjuncts, while the acceptance of (14) would rest upon the *connection* between the disjuncts. The former description is lattice-theoretical in nature, while the latter is group-theoretical.[10]

4. SECOND APPLICATION: THE LOTTERY PARADOX

McGee's paradox is not the only puzzle which can be given a plausible solution along the lines of our suggestion; also Kyburg's renowned *lottery paradox* can be attacked in a similar way. A version of the argument is the following (Nelkin, 2000):

Jim buys a ticket in a million-ticket lottery. He knows it is a fair lottery, but, given the odds, he believes he will lose. When the winning ticket is chosen, it is not his. Did he know his ticket would lose? It seems that he did not. After all, if he *knew* his ticket would lose, why would he have bought it? Further, if he knew his ticket would lose, then, given that his ticket is no different in its chances of winning from any other ticket, it seems that by parity of reasoning he should also know that every other ticket would lose. But of course, he doesn't know that; in fact, he knows that *not* every ticket will lose.

On the other hand, if Jim didn't know his ticket would lose, then can he know any empirical facts at all? If Jim does not know something that has an extremely high probability of being true (0.9999) and is in fact true then what can he know? (p. 372)

Let $W(x)$ stand for "ticket x will be the winning one". The paradox can be given the following natural deduction presentation:

$$\dfrac{\dfrac{\dfrac{\neg W(x)}{\forall x \neg W(x)} \quad \forall x \neg W(x) \rightarrow \neg \exists x W(x)}{\neg \exists x W(x)} \quad \exists x W(x)}{\perp}$$

Given an arbitrary ticket x, Jim knows that it won't win, and thus he is entitled to *assume* that it won't win. From this assumption it follows, by means of hardly disputable logical moves, that *no* ticket will win; but this contradicts the further assumption that *some* ticket will win after all.

If the paradox is formulated in terms of knowledge, the most popular way out is that of denying the premiss $\neg W(x)$ (Bonjour, 1985; DeRose, 1996): after all, Jim does not know for sure that, given an arbitrary ticket such as his, it will lose – he can only guess it, albeit with excellent chances of getting it right. However, there is a variant of the paradox which involves *rational belief* rather than knowledge; it can be plausibly denied that Jim knows that an arbitrary ticket will lose, yet not so much that it is rational for him to believe that it will. A straightforward rejection of $\neg W(x)$, therefore, precludes a uniform solution of both versions of the paradox (Nelkin, 2000).

In spite of this, several authors (e.g. Ryan, 1996; or Nelkin, 2000) have argued – although on different grounds – that it is not even rational for Jim to believe that $\neg W(x)$. Other writers (Kyburg, 1961; Foley, 1993) have focussed on the passage from $\neg W(x)$ to $\forall x \neg W(x)$

and claimed that, even though of every ticket it is rational for Jim to believe that it will lose, it is not rational for Jim to believe that every ticket will lose. If this line of reasoning were correct, the paradoxical argument would go through only if we assumed a sort of "doxastic Barcan formula" whose validity has often been cast into doubt in the literature (see e.g. Hintikka, 1962):

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\Box\neg W(x)}{\forall x\Box\neg W(x)} \quad \forall x\Box\neg W(x)\to\Box\forall x\neg W(x)}{\Box\forall x\neg W(x)} \quad \Box\forall x\neg W(x)\to\Box\neg\exists x W(x)}{\Box\neg\exists x W(x)} \quad \Box\neg\exists x W(x)\to\neg\Box\exists x W(x)}{\neg\Box\exists x W(x)} \quad \Box\exists x W(x)}{\perp}$$

(the other principles used in the proof are either straightforward or can be made plausible).

Since the former approach clashes with the intuitions of many people and the latter is subject to specific objections (see e.g. Nelkin, 2000), once again we find ourselves in the uncomfortable position of lacking a fully satisfactory solution to our paradox. On the other hand, if we adopt the thesis of the ambiguity of quantifiers, it is easy to see where the argument breaks down: all the quantifiers in the proof are lattice-theoretical, except for the one in $\exists xW(x)$, which is group-theoretical. In fact, $\forall x\neg W(x)$ cannot mean that *every* ticket will lose – since Jim knows that the lottery is fair, he would be wrong in making this assumption. After all, the mere fact that he bought his ticket shows that he deems such an assumption false. What is rational for him to believe is that *any* ticket will lose, i.e. that given an arbitrary ticket on sale, that ticket will lose. This quantifier is therefore lattice-theoretical. On the contrary, the existential quantifier contained in the other premiss of the argument must perforce be group-theoretical: when Jim agrees that some ticket will eventually win in a fair lottery, his ground for asserting this is not the belief that some specific ticket will win – if it were, it would be enough for him to buy *that* ticket and he would get the prize. But, alas, that is not the case: unfortunately, the justification of Jim's belief is simply his awareness that if ticket no. 1 loses, then if ticket no. 2 loses, then . . . then ticket no. 1,000,000 has to win.

Although this unhappy circumstance dramatically reduces Jim's chances of winning the lottery, it enhances our prospects of giving an acceptable solution to the paradox. In fact, this means that the real logical form of the argument is the following:

$$\dfrac{\dfrac{\dfrac{\neg W(x)}{\forall x \neg W(x)} \quad \forall x \neg W(x) \to \neg \exists x\, W(x)}{\neg \exists x\, W(x)} \qquad \Sigma x\, W(x)}{\perp}$$

So the last inference is wholly illegitimate: no contradiction is produced, since the first premiss is not the negation of the second.

ACKNOWLEDGEMENTS

A version of this paper was presented at the SILFS Conference in Rome (June 2002). We are indebted to the audience, and especially to Enzo Fano and V.M. Abrusci, for their precious comments. We also thank Graham Priest and Greg Restall for their remarks on earlier drafts of this paper, and an anonymous referee for raising some important objections and pointing out several inaccuracies.

NOTES

[1] After all, it is not very natural to call "addition" an *idempotent* connective.

[2] This usage is motivated by the algebraic semantics for substructural logics: see Paoli (2002).

[3] An exception is O'Hearn and Pym (1999), where – however – group-theoretical quantifiers are admittedly not treated as the infinitary counterparts of the corresponding conjunction and disjunction connectives. See also Montagna (200+) for a related distinction in fuzzy logics.

[4] To put it differently: (4A) can be plausibly taken as synonimous with "Oswald killed Kennedy $\oplus$ someone else killed Kennedy", which is in turn equivalent to (6). On the other hand, (4B) would only seem to entail "Oswald killed Kennedy $\vee$ someone else killed Kennedy", whence (6) does not follow.

[5] Notice that here Adams' (1970) well-known distinction between (6) and its "hadn't-would" counterpart is not at issue: our bystander would assent neither to (6) nor to its subjunctive variant.

[6] The reason is once again algebraic. Whereas it is natural to interpret lattice-theoretical existential quantifiers via arbitrary lattice joins, it is less clear how to interpret their group-theoretical counterparts. It seems plausible to suppose that they should be interpreted as limits of appropriate series (i.e. sequences of partial sums, where the sum is the groupoidal operation which interprets "$\oplus$") and this is why the symbol "Σ" seems appropriate.

7 An interesting passage by Bertrand Russell seems relevant to our distinction and, in general, to the preceding discussion. In his *Principles of Mathematics*, 2nd edn. (Russell, 1937), Russell offers an exemplary analysis of the abundance of expressions which – in English as well as in other natural languages – can be employed to refer to generality and particularity, and especially of the phrases *Every a*, *Any a*, *An a*, and *Some a*. Russell exemplifies this distinction by considering the case of a set consisting of finitely many elements, say $a_1, \ldots, a_n$. In his own words: "*Every a* denotes a_1 and denotes a_2 and ... and denotes a_n. *Any a* denotes a_1 or a_2 or ... or a_n, where *or* has the meaning that it is irrelevant which you take. *An a* denotes a_1 or a_2 or ... or a_n, where *or* has the meaning that no one in particular must be taken [...]. *Some a* denotes a_1 or denotes a_2 or ... or denotes a_n, where it is not irrelevant which is taken, but on the contrary some particular *a* must be taken" (p. 59). In another passage of the same work, he provides concrete natural language examples to show that his taxonomy is sound and applicable to English.

8 Other commentators partially agree with McGee. Piller (1996) suggests that at least two theories of conditionals (the Adams-Appiah theory and the implicature approach) can properly account for McGee's counterexample, showing that it creates no real problem to modus ponens. However, he raises independent objections against such theories, concluding thereby that McGee's negative attitude towards modus ponens should be – at least provisionally – upheld. Cp. also Aune (2002) for a cautious viewpoint on the issue. We finally mention the opinion of Gauker (1994), who believes that, although McGee has not refuted modus ponens, his argument can be successfully recast into a counterexample for modus *tollens*.

9 A similar proposal was advanced by Lowe (1987), although he located the ambiguity elsewhere in the argument.

10 We thank an anonymous referee for pointing to this variant of the paradox. The same referee suggested a further version. Imagine a speaker who knows that there are two Republicans in the race, Reagan and Anderson, and knows that one of them has pulled decisively ahead of the Democrat in the polls while the other has fallen dismally behind. Such a speaker will still believe the premisses of the original argument and disbelieve the conclusion, but cannot use "a Republican" to refer to the forerunner, because the speaker does not know who the forerunner is. In this case, we think that the categorical premiss can have two equally legitimate readings; on one of them, it has the same meaning as the antecedent of the conditional premiss ("either Reagan or Anderson will win, but it is not known which; if it's not Reagan it is Anderson, and if it's not Anderson it will be Reagan"). Here no equivocation is lurking, but the paradox is solved all the same because, on such a reading, the conclusion is *true*, being simply a reformulation of the categorical premiss.

REFERENCES

Adams, V.M. (1970): 'Subjunctive and Indicative Conditionals', *Foundations of Language* 6, 89–94.

Anderson, A.R. and Belnap, N.D. (1975): *Entailment*, vol. 1, Princeton, NJ: Princeton University Press.

Aune, B. (2002): 'Against Moderate Rationalism', *Journal of Philosophical Research* 27.

Bonjour, L. (1985): *The Structure of Empirical Knowledge*, Cambridge, MA: Harvard University Press.

Casari, E. (1997): 'Conjoining and Disjoining on Different Levels', in M.L. Dalla Chiara et al. (eds.), *Logic and Scientific Methods* (pp. 261–288), Dordrecht: Kluwer.

DeRose, K. (1996): 'Knowledge, Assertion, and Lotteries', *Australasian Journal of Philosophy* 74, 568–580.

Donnellan, K. (1966): 'Reference and Definite Descriptions', *Philosophical Review* 75, 281–304.

Foley, R. (1993): *Working Without a Net*, Oxford: Oxford University Press.

Gauker, C. (1994): *Thinking Out Loud*, Princeton, NJ: Princeton University Press.

Hintikka, J. (1962): *Knowledge and Belief*, Ithaca, NY: Cornell University Press.

Katz, B. (1999): 'On a Supposed Counterexample to Modus Ponens', *Journal of Philosophy* 96, 404–415.

Kyburg, H. (1961): *Probability and the Logic of Rational Belief*, Middletown: Wesleyan University Press.

Lowe, E.J. (1987): 'Not a Counterexample to Modus Ponens', *Analysis* 47, 44–47.

McGee, V. (1985): 'A Counterexample to Modus Ponens', *Journal of Philosophy* 82, 462–471.

Montagna, F. (200+): 'Storage Operators and Multiplicative Quantifiers in Many-Valued Logics', unpublished manuscript.

Nelkin, D.K. (2000): 'The Lottery Paradox, Knowledge, and Rationality', *Philosophical Review* 109, 373–409.

O'Hearn, P.W. and Pym, D.J. (1999): 'The Logic of Bunched Implications', *Bulletin of Symbolic Logic* 5, 215–244.

Paoli, F. (2002): *Substructural Logics: A Primer*, Dordrecht: Kluwer.

Piller, C. (1996): 'Vann McGee's Counterexample to Modus Ponens', *Philosophical Studies* 82, 27–54.

Read, S. (1988): *Relevant Logic*, Oxford: Blackwell.

Russell, B. (1937^2): *The Principles of Mathematics*, London: Allen & Unwin.

Ryan, S. (1996): 'The Epistemic Virtues of Consistency', *Synthese* 109, 121–141.

Sinnott-Armstrong, W., Moor, J. and Fogelin, R. (1986): 'A Defense of Modus Ponens', *Journal of Philosophy* 83, 296–304.

Stalnaker, R. (1975), 'Indicative Conditionals', *Philosophia* 5, 269–286.

Dipartimento di Scienze Pedagogiche e Filosofiche
Università di Cagliari
Via Is Mirrionis 1, 09123 Cagliari
Italy
E-mail: paoli@unica.it