

Published:

Louis Caruana, 'Is Science eliminating Ordinary Talk?', *Forum Philosophicum* 4 (1999), 25 – 39.

Is Science eliminating Ordinary Talk?

Discussions dealing with natural science, philosophy and common sense are bound to draw on long-standing debates dealing with realism, methodology of science, philosophy of mind, metaphysics, epistemology, theories of meaning, and other topics. Instead of presenting a broad overview of these main trends, which will necessarily be superficial, I will do a kind of case study. The aim is to present just one particular debate which is of relevance to current research. The presentation is meant to give a taste of how these various long-standing debates are brought to bear on a specific issue. In this way, the very practice of engaging in a particular area of philosophy of science will serve as a platform from where the major areas can be seen in actual operation. The paper has four sections: the nature of ordinary talk; the ontological implications of this; the recently proposed account of the mental; an evaluation.

I. Ordinary talk

Few things can be said to be more ordinary than talking about one's own beliefs and about those of others. One may say that these operations form part of a simple skill people have to understand one another in everyday life. In spite of its ordinariness, this ability has recently come under attack. It has been asserted that some very serious mistakes are made whenever we exercise this ability. During the many centuries that people have used this skill, such mistakes were inevitable. But now things have to change. We now have new scientific evidence that shows that making these mistakes should be avoided. Are we on the threshold of a major conceptual revolution?

To start with, we need to have a closer look at what this everyday skill involves. The ability of people to understand one another in everyday life has certainly a much longer history than natural science as we know it. People describe one another's behaviour in terms of key concepts, like the concept of 'belief'. Without even a trace of scientific education, if I see one of my guests deliberately avoiding the cat, I may explain that behaviour by saying to myself that the person has a belief about cats, namely the belief that they may easily become aggressive. Obviously, I may be mistaken. The person may have another belief. He may in fact have the belief that he is allergic to cat fur. But the point I want to focus on here is not that he is assumed to have *this* belief or *that* belief, but just that he is assumed to have a belief of something or other. We employ the term 'belief', and subsequently attribute a content to it, as part of our natural way of accounting for each other's behaviour. This skill is not one dealing only with description. In everyday life, we also use concepts

like that of 'belief' and 'desire' to make predictions about how people will behave. Hence, in the above example, if my other two cats come along to the sitting room and occupy all the available chairs, I can predict that my guest will be somewhat embarrassed when looking for a place where to sit.

The crucial question here is the following. Can this ability or skill that people have of describing each other's behaviour and even of making predictions about them, using key concepts like 'belief' and 'desire', be conceived of as a scientific theory? Given the mathematical structure associated with normal scientific theories, one may find it difficult at first sight to call this skill a scientific theory. Indeed, scientific theories are often viewed as intricate hypothetico-deductive systems consisting of covering laws which act as a small set of axioms or hypotheses from which everything else can be shown to follow as a deductive consequence. And this is quite far from what happens in everyday language about human behaviour. Nevertheless, one must recall that ideas about the nature of scientific theories have seen recent developments. It progressively became apparent that the hypothetico-deductive view of theories will inevitably leave out major scientific theories, like Darwin's theory of evolution. Recent work in fact has broadened the previous understanding in such a way that now a scientific theory is considered to be better described, in general, as an attempt at binding together in a systematic fashion the knowledge that one has of some particular aspect of the world of experience.¹ The main ingredient is not a deductive structure but a set of models which may not all have empirical meaning at a given time, but may gain such a meaning in so far as they can be applied fruitfully to specific areas of empirical inquiry. On this wider and more realistic understanding of theory, the ability people have of describing and predicting each other's behaviour could, according to some philosophers, qualify as a scientific theory.

The main claim is that behind this apparently spontaneous use of terms like 'beliefs' and 'desires' lies an implicit theory which we come to learn and live by as we grow up. This theory is not any different in kind from any other scientific theory. Two typical reasons for this claim are the following.² First, there is the argument from simplicity. If we call this human ability a theory, then we arrive at a simple and unifying organisation of most of the major topics in the philosophy of mind. This unification of our understanding is always taken as a criterion of genuine scientific progress. In this particular case here, this unification and simplification occurs as regards at least one difficult problem of traditional philosophy. This is the problem of other minds. In brief, it concerns the question whether, given the natural way we use mental terms like 'pain', and 'belief', we can ever make sense of attributing them to others. The main problem here arises because these mental terms seem to obtain their meaning exclusively from our relating them to our own experiences. But this mysterious capacity to jump from my mind to the mind of others ceases to be mysterious and problematic once we start considering our everyday linguistic skill as a scientific theory. My conviction that another individual is the subject of certain mental states is neither a deduction nor an induction on my part but a simple, straightforward explanatory hypothesis. When I say 'You are in pain', I am just presenting my scientific theory about you. It is as simple as that. Since this problem of other minds is thus solved by assuming that everyday discourse on mental behaviour

¹ Popper 1959, Kuhn 1970, Newton-Smith 1981.

² Churchland 1981.

is a scientific theory, then the assumption must be right. It is not a mysterious linguistic skill but a scientific theory like any other. The second typical argument concerns the logical structure of expressions in everyday discourse and the logical structure of statements constituting scientific theories. Close inspection of our daily use of mental attributes reveals that there seems to be no difference between these two logical structures, in spite of the spontaneity and carelessness associated with everyday utterances. For example, in science we say that whenever a given mass is acted upon by a force, a definite acceleration results. In everyday discourse we likewise say that whenever a given person is acted upon by a certain fear, a definite desire to avoid the source of fear results. The logical structure of relations between force and acceleration is the same as the logical structure of relations between fear and desire. Compare:

$$(x) (f) (m) [(x \text{ has a mass } m) \ \& \ (x \text{ suffers a force } f)] \supset (x \text{ accelerates at } f/m)]$$

and

$$(x) (p) [(x \text{ fears that } p) \supset (x \text{ desires that } \sim p)].$$

So this is another reason for holding that our everyday ability to describe and predict human behaviour does not have a special status but is in fact just another empirical theory. The term *folk-psychology* has been suggested with this point in mind: folk-psychology is the simple alternative to, or the precursor of, scientific-psychology.

II. Ontological implications

If one accepts these arguments, then the two kinds of psychology become just two discourses of the same empirical kind. This is a crucial claim. It shows that we are dealing with two theories being employed simultaneously to account for the same reality. And there is seldom peaceful co-existence of theories when they overlap. Folk-psychology is in fact inconsistent with scientific-psychology. The clash occurs at the deep level. To appreciate this, it must be recalled that, in general, scientific theories oblige their protagonists to adopt an implicit ontology. A good look at how statements within a theory are used will reveal what objects are being considered real and in what sense. It will reveal the ontological commitment of users of the theory.³

³ Cf. Quine 1981. The term 'ontology' is here being used not in its general sense of the science of being in general, but in a derivative sense. It here refers to the set of things whose existence is acknowledged by a particular theory or system of thought. In the course of history, there have been famous debates concerning the question whether existence is a predicate or not. In current analytic philosophy of science, one usually refers to the study by Quine 1961. Here one finds his claim that 'to be is to be the value of a bound variable'. For a given theory, identifying what kinds of things the scientist quantifies over reveals what that scientist takes there to be. If the theory in question is well-established, then its ontology merits a high degree of our credibility. On the contrary, if the theory is not well-established, its ontology merits a low degree of credibility. Interesting areas of debate arise when very well established scientific theories seem to require an ontology that goes against traditional metaphysical principles. It is in this sense that Quantum Mechanics is said to challenge traditional metaphysics, especially as regards the principle of causation 'every event must have a cause' and the principle of continuity '*Natura non facit saltus*'. There is another point worthy of note regarding Quine's idea of Ontological Commitment. He admits *only* those kinds of things that *need* to be referred to in stating what we know. This makes sense only in so far as we can have a clear idea of what we

Folk-psychology, considered as a scientific theory, must have its own ontological baggage which is carried around by its users. To get an idea of what this is, we have to investigate how we use simple notions like ‘belief’, ‘desire’, ‘fear’, and such like in our everyday description and prediction of human behaviour. Consider my previous example. I see my guest deliberately avoiding the cat. I attribute a certain belief to him, namely that he thinks cats may get unpleasantly aggressive without warning. In making this attribution, what deep-level assumptions about the mental am I committing myself to?

Some may perhaps be tempted to reply here that, in using folk-psychology in this way, I am implicitly assuming that there is a definite piece of my guest’s brain that corresponds to the belief I am talking about. Just as I can assume that he has a few coins in his pocket because of the jingle I hear when he moves around, so also I assume that there are a few special neurones in his brain because of the behaviour I am observing. I may even call these neurones: ‘the fear-of-cats neuronal structure’. This reply is a typical physicalist response. It is not the only possible kind of response. In fact, it represents one extreme position. Folk-psychology does not commit me to be a physicalist. A physicalist explanation in this case does seem to be consistent with folk-psychology. But going to the extreme of physicalism is adding more to what users of this theory are normally committed to.

The minimal ontological commitment in my using folk-psychology consists simply in the presupposition that my guest’s store of beliefs and memories are in fact a collection of functionally discrete states. In other words, when talking of beliefs, and even identifying them as individual beliefs about some particular thing or situation, I am implicitly assuming that the holder of these beliefs has an interconnected collection of these discrete states. Hence, if I claim that my guest has a belief that the cat is to be avoided because it may get aggressive, then the cognitive state my guest is in should, I take it, be described by saying that he has one definite and identifiable sentence-like structure corresponding to that one distinct belief, namely the belief that the cat is to be avoided because it may get aggressive. The important point here is that by attributing a belief to my guest, I am manifesting my deeper assumption that his cognitive faculty operates in a way that involves discrete parts. Hence, I may eventually observe other elements of my guest’s behaviour which I explain by saying that he has other beliefs about cats, and other beliefs about other things. These beliefs are held by him, in a sense, one next to the other. I may even arrive at a point when I start worrying that some of his beliefs may clash with others, and start wondering whether he is to be considered fully rational as regards his attitude towards animals. All such claims are the result of my applying the theory I have been brought up with, namely folk-psychology. And all of them are built upon the deep assumption that the mental is divisible into discrete parts, sometimes called modules. The minimum that folk-psychology seems to impose, therefore, as its own ontological baggage is a *modular* account of the mental according to which the mental is neatly divisible into bits each corresponding to what we call beliefs, desires, and so on.⁴

need. Quine, and many pragmatists, do not seem conscious of the difficulties here: should we limit ourselves to immediate needs or include also possible needs?

⁴ This modularity of the mental is a kind of propositional modularity. It should be distinguished from the idea that the mind may function in different modes. In this paper, the term module means a discrete part that may cause behaviour. Consequently, according to the modular account of the mental, a person

III. A new account of the mental

The result of our philosophical investigation up to now may seem rather trivial. A crisis, however, starts to emerge when we realise that this modular ontology of the mental clashes with modern accounts derived from scientific-psychology, and from very plausible neuro-scientific theories of the brain. These modern accounts of the mental depend on *connectionism*. To enable the reader to appreciate the full extent of the disagreement with folk-psychology, I will proceed by giving a quick, simplified overview of what the new theory is about, and then examine, as I did before, the implicit ontological assumptions that have to be made when applying this theory to explain the mental.

Connectionism is the name given to a certain research programme in computer science and especially in its application to cognitive science.⁵ When discussed in the context of computer science, it is sometimes referred to as Parallel Distributed Programming. A network here can be imagined as a number of points in space with connections between them. The points, or units of the network, are simple processors and are usually considered not just as a homogenous cloud but as grouped into a number of distinct populations or layers. The connections between units of one layer with units of the next layer are extremely numerous and, moreover, are not all of the same strength. Information passes from one layer to another from processor to processor, but the itinerary of information as it passes through the network is not linear. It is not, in other words, along a single line made up of connections from one node to another node to another node, and so on. On the contrary, it is spread out: it passes along various connection pathways involving many nodes at the same time. This aspect is what makes it *parallel* processing. Moreover, any individual connection or node participates not only in the processing or storage of one particular item of information but in the simultaneous processing and storage of many different items of information. This explains why the system is called parallel *distributed* processing.

The usefulness of such processing is undisputed. One of the interesting features of such networks is their ability to program themselves. They have autonomous procedures for tuning their weights to eventually perform some specific computation. Such a learning procedure usually depends on a training period during which the network is presented with sample input-output pairs. For example, a network may be taught how to distinguish between sonar echoes bouncing back from submarine rocks and similar echoes bouncing back from explosive mines.⁶ Scientists record fifty different rock echoes and fifty different mine echoes. These are digitised and used to train, or standardise, the network. The training is done with input-output pairs. One should realise at this point that, during training, the network uses nothing else but the inputs in order to distinguish different output classes: there is no other source of information that the network has access to. After this training procedure, the network could be used to distinguish between new mine and rock echoes outside the hundred recorded examples. The interesting thing is that such networks can work and

engaging in some cognitive activity, like answering a question, has some mental states that are activated and others not.

⁵ Smolensky 1988, Ramsey et al. 1991.

⁶ Churchland 1992, p. 346.

expand in their detecting ability on their own. They optimise their weights for a problem by just using the input-output pairs in the training data. It is true that the experimenter can always observe what the values of all the connections are, no matter how deep in the network they are. However, the experimenter is never obliged to set the connection-weights by hand. It is in this sense that the networks can be said to learn on their own.

When we apply this model to explain the functioning of the brain, we have a node corresponding to a brain cell with a particular specialisation. In the case of vision, the first layer of the system consists of a row of sensory nodes corresponding to the cells that are found on the retina. Each of these nodes projects a fibre towards a second layer of nodes which can conveniently be imagined as situated 'deeper down' into the system. In reality, these nodes of the second layer correspond to the small parts of the brain that are connected directly to the optic nerve and are called by neuro-physiologists the lateral geniculate nuclei. The connection of each retinal node all along the optic nerve is the normal threadlike part of that nerve cell, called an axon, which conducts impulses. As each axon arrives at the brain, it can make many thousands of connections as it fans out to reach each node of this second layer. Since these nodes of the second layer are not directly connected to outside stimulus, they are referred to as hidden units. They are connected in turn to units of a third layer, and so on, until the tree-like, neural organisation arrives at nodes that are not connected to other nodes but to organs of the body. These are called output nodes, in the sense that they can be the initiators of motor action elsewhere in the system, movement of muscles, activation of other parts of the brain, and so on. The complexity of real brains is staggering, not only because of the many layers, but also because of feedback mechanisms that are not included within this proposed network structure. The simplest model we can have, and one which we can reasonably understand, has just three layers of nodes and no feedback: first an input layer of nodes, each connected to numerous nodes of a second, hidden layer, and again, each of these second layer nodes connected to numerous output nodes at a third and final layer. This simple three-layer structure, consisting of input-nodes, hidden-nodes and output-nodes, is a simplified but faithful description of what exists in the brain. An input stimulus such as light produces an effect of a definite strength on a number of input nodes. This then conveys from each activated node a signal of proportional strength to the many hidden units. The connections do not all have the same size or weight. A given hidden or output unit sums up the effects incident from the many input units connected to it.

In what way does this account of the mental clash with folk-psychology? It clashes because of its different presupposed ontology. Accepting a connectionist model of the mental means accepting also an implied view of what beliefs are. We have already mentioned how accepting folk-psychology obliges us to adopt a modular ontology according to which mental activity is neatly divisible into bits each corresponding to what we call beliefs, desires, and so on. In contrast to this, accepting a connectionist account of the mental will not involve any possible or actual carving out of single, localised states which could correspond to a specific belief or desire. The ontology behind a connectionist account does not involve modularity. This is so because a parallel distribution network will always process and store information in the connection weights in a way that is *widely distributed* rather than *localised*. If my

guest believes that the cat may get aggressive, then, according to this model, there is no definite part of his cognitive system that is in operation while the rest is idle. The belief is stored in a *holistic* way. Moreover, there is a simple interpretation only of the external part of the network proposed by the connectionist account. In fact, as regards vision, what happens at the outer input nodes can be taken to correspond to external stimuli hitting the retina, stimuli with a definite shape and colour. But what happens at the next level of nodes, the hidden level, has no symbolic interpretation. One cannot say in what sense the original shape of the external object, or its colour, is retained within the system. The hidden level is inaccessible to our schematising.

So we have now arrived at an impasse. We have two theories to explain the same reality, namely mental activity. The two theories are not reconcilable at least on the important ontological issue of modularity or holism. How is one to proceed?⁷

Common sense dictates that, when such a clash of theories occurs, if one theory cannot be smoothly reduced to the other, then they cannot both be right. One should be eliminated. The history of science is full of examples of elimination of one theory by another. A well known example is that of the phlogiston theory of combustion. As historians tell us, during the seventeenth century, chemistry was dominated by the idea of this so-called 'fire principle'.⁸ When a body burns, the explanation was that it lost phlogiston. The overthrow of this theory gradually became inevitable as a result of the work of Lavoisier, especially of his publication of a scientific note in 1772. This paper showed that sulphur and phosphorus when burned increased in weight because they absorbed 'air'. What is so dramatic about this historical example is that the two competing theories were not just superficially different but fundamentally opposed: one made central claims that were the direct negation of those made by the other. Where the old theorist saw a loss of phlogiston, the new saw a gain of oxygen. So there was no question of harmonising them. One theory had to give way.

The same seems to be true of the case of the theories about mental activity. According to some present philosophers and cognitive scientists, we are in a similar position as the chemists of the 18th century. We have two competing theories which are not just different but make mutually contradictory claims: one assumes modularity of the mental, the other holism. Since it is very likely that connectionism is closer to the truth than folk-psychology, the latter has to go. We should eliminate the use of terms like 'belief' and 'desire'.

IV. Evaluation

⁷ The connectionist model of the brain may also include some degree of superficial modularity. We can consider the brain as an ensemble of a number of quasi-independent agglomerations, each of which is a parallel-distributed network corresponding to a particular function, such as vision, hearing, and so on. Furthermore, we could consider these groups to be weakly interacting, in the sense that just a few small weights connect one to the other. This coarse modularity does not affect the main argument in this paper because, even if we have it, any single belief will be stored holistically in a number of different agglomerations. The holism remains.

⁸ Holmes 1985.

Is this argument convincing? Let us take a careful look at each step it involves. There are four: (1) folk-psychology obliges us to accept an ontology involving the modularity of the mental; (2) connectionism obliges us to accept an ontology involving holism of the mental, which is here taken to be the exact opposite of modularity of the mental; (3) mental activity can be better explained by connectionism; and therefore (4) folk-psychology has to be eliminated to give way to connectionism. Each of the steps seems to be well founded. So what could be wrong? A vulnerable spot may have been overlooked at the first step. The very starting point of the entire argument should be examined more carefully: is folk-psychology an empirical theory? Can we legitimately hold that our day to day skill of describing and predicting each other's behaviour constitutes indeed a discourse of the same kind as the oxygen theory, or any other scientific theory?

For a better evaluation of the entire issue, perhaps a historical example other than the phlogiston case should be examined. Consider the solar system. The long debate concerning the two competing systems of geo-centrism and helio-centrism came to a definite end. The outcome is the same as in the case of phlogiston: one theory is eliminated and forgotten. And yet — in this case, not everything that has to do with geo-centrism has in fact been eliminated from our everyday linguistic practices. One would have expected that, after so many years since the great debates about this point, the elimination process would have been over by now. But, lo and behold, some aspects of geo-centrism linger on: we still say 'The sun is rising' and 'The sun is setting'. Shouldn't such expressions have been abandoned long ago, if everyone is convinced that helio-centrism is in, and geo-centrism is out? What lesson can be drawn from this kind of linguistic *immunity to theory-elimination*?

The main lesson has to do with the meaning of the terms we use in everyday life. It is an indisputable fact that expressions about sunrise and sunset are still being used even though the accepted scientific explanation of the movement we are referring to does not involve any movement of the sun. This everyday use should be taken neither as a careless or approximate way of talking, nor as an embarrassing and persistent collective mistake. It should be taken as how things actually are. *It* should be our starting point in our search for a theory of meaning. We will then restart our reflections with the insight that sentences involving sunrise and sunset cannot be considered simply as statements referring to the underlying mechanism of Nature. The main source of their meaning is not to be sought in the causal structure of the material world. The role played by propositions in everyday life goes far beyond the simple predictive strategy we associate with discourse which is strictly scientific. Think how strange it would be for someone to correct you by saying: 'No, don't talk about the sunrise: we should say the earth has rotated to where it was yesterday at this time.' Nor should the main source of the meaning of our everyday utterances be sought in the underlying mechanism of the utterer, that is in the hidden structure of his or her brain. When we cannot understand a word and ask for its meaning, we are given the various ways in which it is usually used: we are given an idea of how it is embedded in human customs and institutions.⁹ Think how silly it would be if, when asking for the meaning of an unfamiliar word, we are shown a detailed diagram of the brain and told which neuronal group is activated when that word is in use.

⁹ Wittgenstein 1953, §§ 337, 525.

The basic starting point, therefore, is that any utterance obtains its meaning not from the material structure of the world or of the brain of the utterer, but from the context of its use. The centre of my thought and meaning is not my singular brain, as if my brain secretes meaningful utterances which I stick onto facts around me. Neither is the centre of my thoughts and meaning the material world, as if I receive the meaning of my utterances from the mechanism in operation in the physical universe. Rather, the centre of my thought and meaning is our plural, institutional and cultural world.

Let us now apply this general principle to the specific case of the various kinds of discourse about human behaviour. One recalls that those in favour of the elimination of the use of terms like ‘belief’ and ‘desire’, who are normally called eliminativists, hold that our everyday discourse about each other’s behaviour is a theory, which they call folk-psychology and whose role is to describe hidden mechanisms.¹⁰ Hence, suppose we say

S1 Tom avoids touching cats because he believes they may easily become aggressive.

Eliminativists claim that by saying this we are in fact saying that Tom has a belief *somewhere lodged inside him which causes him* to engage in cat-avoiding behaviour. Eliminativists work always within the mechanistic framework. They cannot conceive what else we may be saying by saying S1. They see no other alternative. And they thus miss the most important characteristic of the meaning of everyday utterances, namely that the meaning of an utterance is the *function* of the utterance not in a mechanism but in the culture in which it is embedded, the function in the complex combination of human institutions, customs, practices, habits and so on.¹¹

Eliminativists therefore cannot produce convincing reasons to justify their crucial premise that folk-psychology is an empirical theory. In spite of some similarity in logical structure between statements of everyday discourse about behaviour and statements of science, the former are not of the same kind as the latter. Suppose we ascribe a content to the mental state of my guest Tom, for example by holding S1. By doing this we are not saying anything, in normal situations, about the mechanism in his brain. What we are saying has to do with how Tom behaves within the culture he is embedded in. Statement S1 is not about the neurones in his brain, but about his acting differently from others during the party when the cat came into the room, perhaps about his previous awful experience with the neighbour’s cat, perhaps about his mistrust of our usual underestimation of animal double-think, and so on. S1 is linked continuously with an enormous chain of social and cultural situations forming an organic whole without separable parts. Likewise when we hold that

S2 Jane thinks the sun has risen

we are not saying anything, in normal situations, about her convictions regarding geocentrism or helio-centrism. What we are saying has to do with Jane’s getting up, her having breakfast, her rushing off to work, her feeling fresh at the beginning of a new

¹⁰ Typical eliminativists are the P.M. Churchland, P.S. Churchland, S. Stich.

¹¹ McDonough 1989, 1991.

day, her feeling concerned that the dreaded interview has got a day closer, and so on, including, if it were possible, all the distant ramifications of meaning embedded in Jane's personal, social, and cultural life.

When holding that everyday discourse dealing with describing and predicting each other's behaviour is an empirical theory, in the sense that it is a kind of psychology about the internal workings of the brain, eliminativists are distorting their object of study. They distort it by trying to make it fit into an exclusively mechanistic world view. Admittedly, some aspects of our world, including our brain, can indeed be viewed in terms of push-pull relations on the model of a machine. But, the fact that this mechanistic world view is useful in some areas does not mean that it is useful in all areas. As we have seen, the meaning of an utterance is plausibly accounted for only if we bring in the complexity of the function the utterance has in the culture in which it is used. The mechanistic world view would have us believe that the meaning of an utterance is only related to some underlying mechanism: either the one outside us, or the one inside our brain. People's ability to describe and predict each other's behaviour, even though this ability involves terms like 'beliefs' and 'desires' which may sound scientific, is certainly not a scientific theory dealing with mechanisms. This ability is not a cheap kind of scientific psychology, which should be identified as such by the term folk-psychology. It is a skill which forms part of our life. If we reduce this complex source of linguistic behaviour to a mechanistic scientific theory we would be missing its essential nature. We would be making a mistake just as we would if we were to reduce culture to a rigid set of presuppositions. The meaning of the utterances of our everyday linguistic practice is not to be sought towards the centre of the human individual, the brain, but towards the centre of human practice. Seeking the source of our meaning within the mechanism of the brain of each individual, who would presumably possess internal units called beliefs, desires and so on, is to work in the wrong direction.

It is evident therefore that the vulnerability of the eliminativist program lies in its basic presupposition, namely its initial claim that folk-psychology is an empirical theory. One should note here, however, that, having identified this flaw in the argument of eliminativists is not to say that philosophers cannot learn from the advance of science. Indeed they can. In this paper we have an example of how a careful evaluation of the alleged competition between the ontology of folk-psychology and that of connectionism can be useful. It has redirected our reflections to the fundamental source of our language, namely the meaning-generating community as a complex whole.

So, what should be eliminated? Should we perhaps eliminate the now apparently ordinary use of the term folk-psychology? If by using the term we are consciously or unconsciously holding that our linguistic practices constitute an empirical theory about the internal mechanism in our heads, we should.¹²

¹² I am thus going further than those who retain the term 'folk-psychology' and deny it is a predictive device (Morton 1996). What we do with words in our everyday life is not a predictive device because it is neither predictive nor a device. Retaining the term 'folk-psychology' gives the wrong impression that one is dealing with a theory.

Louis Caruana SJ
Pontificia Università Gregoriana
Roma

Acknowledgements:

Thanks are due to Dr. Coryn Bailer-Jones, Cavendish Laboratory, Madingley Road, Cambridge, for helpful comments.

References

- Churchland, P.M., 'A Deeper Unity: Some Feyerabendian Themes', in: *Cognitive Models of Science*, ed. R. Giere, Minnesota Studies in the Philosophy of Science Volume XV, Minneapolis: University of Minnesota Press, 1992, pp. 341-365.
- , 'Eliminative Materialism and the Propositional Attitudes', *The Journal of Philosophy*, LXVIII, 2, 1981, pp. 67-90.
- Greenwood, J.D. (ed.), *The Future of Folk-Psychology: Intentionality and Cognitive Science*, Cambridge: Cambridge University Press, 1991.
- Holmes, Frederic L., *Lavoisier and the Chemistry of Life: an Exploration of Scientific Creativity*, Wisconsin: University of Wisconsin Press, 1985
- Kuhn, T., *The Structure of Scientific Revolutions*, Chicago: University of Chicago Press, 2nd ed., 1970.
- McDonough, R., 'Towards a Non-Mechanistic Theory of Meaning', *Mind*, vol. XCVIII, No. 389 (1989), pp. 1-22.
- , 'A Culturalist Account of Folk-Psychology' in: Greenwood (1991), pp. 263 - 288.
- Morton, Adam, 'Folk-psychology is not a Predictive Device', *Mind* vol. 105 (1996), pp. 119-137.
- Newton-Smith, W. H., *The Rationality of Science*, London: Routledge and Kegan Paul, 1981.
- Popper, Karl, *The Logic of Scientific Discovery*, London, 1959.
- Quine, W.V.O., 'Things and their place in theories', in: *Theories and Things*, Cambridge, Mass.: Harvard University Press, 1981, pp. 1-23.
- , 'On what there is', in: *From a Logical Point of View*, Cambridge, Mass., Harvard University Press, 1961.

Ramsey, W., Stich, S., Garon, J., 'Connectionism, Eliminativism, and the Future of Folk-psychology', in: Greenwood (1991), pp. 93-119.

Smolensky, P., 'On the proper treatment of Connectionism', *The Behavioural and Brain Sciences*, vol. 11 (1988), pp. 1-23.

Wittgenstein, L., *Philosophical Investigations*, trans. G. E. M. Anscombe, Oxford: Blackwell, 1953.