

Adriano P. PALMA*

The wages of sin

La rançon du péché

Résumé : Après une brève description de ce que le fonctionnalisme visait à accomplir, je déploie un certain nombre de considérations qui suggèrent que si l'on entend faire une place plus importante à l'esprit et notamment une place plus importante sur le plan causal, il ne nous est pas d'un grand secours. Un bref final traite d'une proposition hyperminimaliste récemment développée par Fodor.

Starting around 1960, I developed a view in the philosophy of mind based on the analogy between a mind and a digital computer. I gave my view the name "functionalism," and under this name it has become the dominant view — some say the orthodoxy — in contemporary philosophy of mind. Hilary Putnam¹

STRANGER: We proposed as a sufficient mark of real things the presence in a thing of the power of being acted upon or of acting in relation to however insignificant a thing.

THEAETETUS: Yes Plato, *The Sophist*, 248 c-d

As every schoolboy knows, metaphysics is one thing, ontology is another, and epistemology is a third one. To conflate the three of them or any of them pairwise is a mistake. Or so I was taught. Given the quirky character of my education a few more words of terminological clarification will not be amiss. Anybody who finds the distinction trivial is

* Adriano P. Palma, Institute of Philosophy, National Chung Cheng University, 160, SAN-HSING, MING-HSIUNG, CHIA-YI 621, TAIWAN R.O.C. & CREA\Ecole Polytechnique, Paris. EMAIL: palmaa@phil.indiana.edu.

¹ Hilary Putnam, "Why Functionalism Didn't Work", in *Words and Life*; Cambridge, Mass. 1994; p. 441.

strongly invited to skip this paragraph. Ontology has to do with what there is in the widest possible sense. There are ontologies comprising only particulars, some only concrete particulars, some having both particulars and universals. Some will have room for abstract entities outside space and time and some will want to stick to the idea that what fundamentally is there is temporal and spatial, and so forth. I leave it as an easy exercise to identify different views and currents. Each one will set up a specific strategy, most often a specific argumentative strategy to show why the ontology presented ought to be that one. To use an example as illustration almost all nominalistic ontologies will have some special theory accounting for numbers or possibilities, since it appears intuitively clear that once one claims that concrete particulars in space-time are the basic furniture of the universe, numbers are not or, at least, don't look like concrete particulars at all. I am not taking sides here, I am just pointing out a distinction which will turn out to be useful later on. Metaphysics on the other hand I view as a series of doctrines which purport to give us ways of identifying, differentiating, etc. what is there. This sounds like too delicate a distinction but I think it does make a difference. One may very well have an ontology in which both Rolex watches and fake Rolex watches are readily available (this is most certainly the case in any point east of the Caucasus) and yet have no criteria of identity that will tell which is which even in principle. One may then in a case like that have a metaphysics of the origin: applied in the case at hand it delivers the doctrine that only Rolex watches coming out of the correctly placed Swiss factory are Rolexes (note that this is a criterion that may not have anything to do with the intrinsic property of the token watch one buys: it may keep time and look from any standpoint exactly like the original, it isn't one because it does not originate from the right source. Compare and contrast: if one has an ontology which does allow unreduced social objects, such as credit cards, one is by the same token very much better off having some criterion to take care of forgery (again a metaphysics of the origin can do, but a metaphysics of social acceptance may work as well for certain purposes at least, not as a legal defense, I gather.)

Epistemology is then, in a boringly traditional fashion, the set of theories that deals with how we come to know any of this. Epistemologies are then theories of justification in essence. Knowledge is, at least, justified true belief and how we come to acquire any justified true belief isn't in most cases very clear at all. Note that in lots of cases we do have views that collapse one level with another. W.V.O. Quine's great motto "to be is to be the value of bound variable" (within something like a first order formal theory) is a way of linking strictly the ontology with a metaphysics, in his case a metaphysics of nominalistic flavor *cum* a specific doctrine of

the metaphysical import of the fact that a theory can be rendered in a canonical form.

Indirectly, one can see the wisdom of the old ways by observing that any theory claiming two of those to be at bottom identical is not trivial and takes quite a bit of argument to establish. The above mentioned Quine is a case in point. Turning to a different period, Peirce held truth to be just acceptability under epistemically ideal conditions. To prove his case took him a massive amount of work and its results are far from obvious. On the face of it, it looks as though a plain reality principle, not to mention Popperian scruples, suggests that even our best present theory can be wrong and either say something false or entail something false. Ditto for the best conceivable theory.

I won't pass judgment on functionalism's sins immediately. I just remark for the nonce that the persuasive power of the doctrine depends on its silence about the ontology of the mental².

IT TAKES TWO TO TANGO

The time in which Putnam framed his seminal functionalist approach was very much preoccupied with the post-Rylean questions of mind-body identity. The driving force of his functionalism comes from an empirical consideration in itself very compelling. It appears very likely that I can't be the one and only entity with mental states. It seems overwhelmingly likely that different creatures are physically vastly diverse. Here philosophers have been particularly fond of using far-fetched examples involving Martians, Venusians, etc. The intuition's grip seems to me much stronger if one considers a simple case. An adult and a baby human have two very different brains. Under the assumption that both can feel pain or pleasure, what corresponds or realizes the mental states <being-in-pain> and

² Writes Ned J. Block; "Theories of the mind prior to functionalism have been concerned both with what there is and with what gives each type of mental state its own identity, for example what pains have in common in virtue of which they are pain... [functionalism] answers the metaphysical question without answering the ontological question." See Ned J. Block, "Functionalism (2)", in Samuel Guttenplan (ed.) *A Companion to the Philosophy of Mind*; Oxford, 1994; pp. 325-326. The line I am suggesting here claims that while functionalism makes a metaphysical claim it is compatible with any ontology. More precisely: a class of functionally defined states is PAIN because of the relations between inputs and outputs of anything which can be the realization of a Turing machine, or a finite automaton, etc. Note the "anything" there. If angels can glow purple whenever the Virgin's name is taken in vain, there is a proper functionally definable state according to which we can attribute 'pain' to the angels.

<feeling-pleasure> in their respective nervous systems has to be very different. Since we do not have any clear *a priori* reason to be chauvinistic about the attribution of mental states ("could a computer be in pain?" sounds like an empirical question and not one that can be solved by analysis of the concept of pain) we reach the conclusion that mental states are multiply realizable. Mental states can be "implemented", jargon has it, in all sorts of different natural or artificial minds. The minds in question have to possess the properties needed to be interpretable as probabilistic automata³. The relation between the mental and the neural/physical is to be seen on the model of the relation between hardware and programs in computers. The view was a direct response to the identity theorists. If one takes the identity view in its simplest way it turns out that it can never be the case of both a computer and myself being in pain, given the differences of the physical properties that would in the two cases underlie the mental state. Indeed the multiple realizability argument is persuasive against the identity theory⁴.

The plot gets somewhat thicker when a rather different concern is our focus. To approximate my point, let us consider the idea that mentation consists in, perhaps among other things, having certain mental contents. It seems to me that the chief reason we have a mentalistic language or conceptual scheme is to interpret our acting. "interpret" here has to be carefully handled since it is ambiguous, at least in the present philosophical climate.

One reading has it that our interpretations are ex-post rationalizations. Attributing mentality to entities is a matter of stance. As a stance the attributing is subject to all the fine differences in degrees one can think of. To the extent that one finds it useful one can attribute mental contents ("I want more light," "More light comes from that direction") to plants. I won't discuss directly this way of describing our attributing mental states. It is egregiously defended and espoused by Daniel Dennett. I am quite persuaded that it is one coherent way of being functionalist. I also happen

³ See Hilary Putnam, "The mental life of some machines" in Hector-Neri Castañeda (ed.), *Intentionality, Minds and Perception* ; Detroit 1967.

⁴ The multiple realizability does have troubles of its own. Ned J. Block famously presented a battery of considerations charging functionalist identification of mental states with computational/functional states with too much chauvinism and too much liberalism. Skipping the metaphors the charges are that functionalism both attributes mentality to entities such that the attribution is nearly absurd at the intuitive level (the People's Republic of China) and it denies mentation for many systems that just have a different psychology. (See Ned J. Block "Troubles with functionalism" in David M. Rosenthal (ed.), *The Nature of Mind*, Oxford-New York, 1991, pp. 211-228.)

to believe it encounters different kinds of difficulties and that it is probably false. This won't be important here since I am going to assume a form of realism about the mental.

A second and very different reading has it that mentation is not *per se* a matter of interpretation. There is indeed some fact of the matter, whether we do get it right or not, as to what has a certain mental state. Furthermore the kind of explanation we adopt by attributing mental states to entities is in essence causal. It is to me plain that the second interpretation may be called realism. Realism about mental states views attributions as subject to all possible difficulties of epistemically hard cases, none the less the ontology is, so to speak, there. The essential ingredient of realism about mentation is the idea that mental contents are causally efficacious. This may sound like a distinction without a difference or a form of philosophical sophistry.

The difference here is between theories taking mental states to be causally irrelevant and those which don't. Most notably Wittgenstein and many of his followers view mental states as a peculiar kind of ex-post rationalized explanation of instances of behavior, not as causes of behavior in the first place. A lot hangs on the force attributed to the 'because' in a sentence like "Wittgenstein went to Norway because he liked fjords." The realist about mentality takes 'because' to be the marker of a causal explanation, the intentional stance supporter takes it to be a result of application of charity principles and the like.

The tension was present already in Wittgenstein's dictum to the effect that a reason can never be a cause. We face a reproduction of the dualism we were avoiding. The traditional dualism of substances is replaced by a dualism of explanatory levels. For those who stick to some form of physicalism there appears to be an axiom that can't be given up. Lepore & Loewer express it thus: "the physical is causally closed."⁵ Being closed entails that any physical event has a causal genesis composed entirely of physical events. Hence the dilemma: if an instance of behavior on the part of an entity with a mind is a physical event (something many don't dispute), then mental content can't have any causal relevance to it. The famous alternative is Davidson's view that the mental contents (call them reasons for perspicuity) are causes but not in virtue of their being mental. They are causes in virtue of their being physical. One needs to see mental contents as both physical and mental. This is the metaphysical kernel of Davidson's anomalous monism. The plausibility of anomalous monism

⁵ Ernest Lepore and Barry Loewer, "Mind Matters" in Rosenthal, *op. cit.*, p. 261.

seems to me to rest on the shaky rock of the possibility of epiphenomenalism as a viable philosophy of mind.

I think a reasonable *desideratum* for a realist theory of the mind consists in the requirement that minds are causally relevant. Why a *desideratum*? because there is one alternative. One could consistently hold the view that mental contents are real and epiphenomenal. A word of explanation is not amiss on what epiphenomenalism is. The best and most straightforward instance of epiphenomenon I owe to an example by Jaegwon Kim. A moving object projects a shadow (under the proper conditions of illumination.) The shadow is a phenomenon and it exhibits all sorts of interesting regularities. However it does nothing with respect to the motion of the object. It does nothing in the sense of being causally irrelevant to the motion. One sense of epiphenomenal is the strong one. One could have the view that because of unknown correlations that are purely accidental every time event A causes event B to take place, event K takes place too⁶. In this very strong form epiphenomenalism could not care less about causation and causal links. All we observe is the presence of a sort of reliable coincidence, a Humean causation if one wishes to call it so. However its very strength is also its weakness if one looks after the types of correlation that afford some generalizable, or lawlike, expression.

The weak sense of epiphenomenalism is more of a worry here. Take event A to be the Landers earthquake of June 1992. Take event B to be the lifting of the ridge near Emerson fault in the Mojave desert. A caused also the motion of a large number of needles in seismographs across the world. Call the movement of the needles K. Now relative to the lifting of the ridge K is epiphenomenal in the weak sense. A is the causally relevant factor for B and for K. A is a cause of B in virtue of some of its properties and not of all of them. It is a (causal) property of A to be able to move the needles of seismographs. It is in turn the moving of needles that can get semantically valued, say (being a symbol of) "being an earthquake of degree 9." Now, relative to the lifting of the ridge, the semantic property, or the content of K, is epiphenomenal in the weak sense. There is a long causal and semantical story that can be told with measurements thrown in. However the content of K, the semantic content of K is something hovering over A and B. This very property is causally inert in lifting the ridge in the desert.

⁶ I have in mind the logical possibility that each time I light a cigar in Hong-Kong a rock moves on the Moon. Why this would happen is perfectly unclear and it is not at all clear that anybody would have any interest in looking at it.

Epiphenomenalism in the weak sense seems to be the proper view for mental states as well if one accepts the functionalist-computational view. Mental states as functional states are purely abstract. If the physical is causally closed it has to be the concrete implementation of mental states that is causally active. I submit that a functionalist view is committed to weak epiphenomenalism.

To see why suppose the belief that grass-is-green is realized by my present neural configuration, with all the appropriate functional inputs and outputs. If such a belief is causally relevant in my actively seeking to buy green shirts because I like the grass' color it can be such in virtue of a set of connections between neural states. <that-grass-is-green> plays no role in my going out to buy the green stuff since, by the prior definition of mental state and by the closure of the physical, an abstract state would not interact with my neural states.

A LUCKY HARMONY

Epiphenomenalism spells trouble only for the realist. The instrumentalist view can quite well live with a form of epiphenomenalism. As long as there is some reasonable degree of parallelism between attribution of states and the behavior of the system states are attributed to, the instrumentalist should not be bothered. The realist about minds on the other hand may be bothered by the futility of the mental. After all, short of a pre-established harmony, that my belief is about green grass has no causal relevance whatsoever. The computational state realized by the automaton I am does all the causal work relative to behavior: its content or its semantical property is otiose⁷.

One alternative notoriously exploited by Jerry Fodor is to see mental contents as dual entities with a narrow and a wide component. Narrow

⁷ Reaching it from a different route, John Heil makes a similar point in "Minds and Bodies."; in Warner R., & Szubka T., *The Mind-Body Problem*; Oxford 1994; pp. 156-164. Heil uses two ways to approach epiphenomenalism. One is supported by externalist theories of reference. If externalism is correct then the contents of a mental state referring to grass depend on its relational properties. Heil notes: "the causal powers of an object depend exclusively on that object nonrelational, here-and-now characteristics." (p. 158). Therefore mental contents are causally irrelevant. The second possibility is that contents are genetically built. The analogy Heil uses is chess pieces. Being a rook for a piece of wood depends on having played a role in chess games. Being a piece of the Chinese Great Wall is causally otiose with regards to a brick that breaks my glass. The causal powers of the brick depend on weight, shape, etc., and not at all on its historical origin. If contents are genetically properties of neural states in this sense, again mental contents are epiphenomenal relative to behavior.

contents are essentially individualistic, internalistic entities supervening on nonintentional monadic properties of the body (the brain most likely.) They are individualistic because they play the role of pushing individual behaviors. Their principles of individuation are not sensible to the distinction one may have externalistically. Narrowly conceived the desire <I want to drink water> would be not distinct between the two possible contents, as in the twin thought experiment cases in which I and my counterpart face environments differing with regards to the chemical composition of the liquid we both call "water" and which is phenomally not distinguished by our perceptual states. Narrow contents are internalistic since to adopt the usual line one could very well have the same narrow content individuated by "I see water in the glass" in the case in which the glass is full of water and in the case of the glass full of vodka. If this narrow content is what pushes behavior one will quickly be painfully aware or pleasantly surprised by the sharp distinction between narrow contents driving one's actions and the wide contents having to do with how the environment is hospitable to one's own narrowly conceived mental states⁸. This solution seems to, for the nonce at any rate, unstable.

⁸ Narrow contents are essentially individualistic, supervening on nonintentional and monadic properties of the body. They should be relevant to psychological explanation. Roughly they are a cognitive counterpart of Kaplan and Perry's characters. The story gets very quickly complicated here. For one thing the whole narrow vs. wide content machinery has had a hard time facing the assault by the externalist views, most notably of Hilary Putnam and Tyler Burge. Jerry Fodor himself in more recent times tends to see Burge's point as stronger and indeed compatible with an intentional psychology. I presented some problems for narrow contents in "Hopes and Doubts" in *European Review of Philosophy* ; vol. 1; CSLI Publications, Stanford, Calif., 1994. Those have been expanded upon by Elisabeth Pacherie in *Naturaliser l'intentionnalité*; Paris 1993, ch. VIII. Pacherie holds that narrow content isn't content anymore. Block in "What Narrow Content is Not" (See Barry Loewer and Georges Rey, eds.; *Meaning In Mind: Fodor and his Critics*; Oxford 1991) shows how on one construal narrow content collapses into syntax and on another plausible construal is purely holistic. I won't enter into the debate there. For one thing the discussion on narrow contents seem to be quicksand. For another Fodor is quite aware of his holding a sort of dual explanation: "... it is very important to keep clear on the difference between the following two questions. "What's the story about the properties of mental states in virtue of which they are subsumed by psychological laws?" and "What's the story about the properties of mental states in virtue of which they are engaged by mental processes?" If the computational theory of the mind is true, then the properties of mental states in virtue of which they are engaged in mental processes are intrinsic/syntactic. If informational theories of content are right, then the properties of mental states in virtue of which they are subsumed by psychological laws are extrinsic/relational." see Jerry Fodor, in *Meaning in Mind*, *op. cit.* p. 298.

The realist about the mental if wedded to a functional-computational view of mental states is stuck with what Block calls the 'Paradox of the Causal Efficacy of Content'. The paradox comes from being committed to three claims. Each one of them appears to be true, and yet they are jointly unsatisfiable:

1. The intentional content of a thought is causally relevant to its behavioral effects.
2. Intentional content reduces to meanings of internal representations.
3. Internal processors are sensitive only to the syntax of internal representations⁹

Syntax here has to be taken as tokens which are completely devoid of any meaning. Even if we were to find something like the shape <0>, perhaps by neurons that neatly sit in a circle, that would not mean zero. The processor that computes conjunctions will only be able to spit out a 0 if and only if it receives two 0's as input. It couldn't care less if 0 means zero, or one or anything else. Semantics, no matter which view one will prefer, will have to deal with something relational in ways in which syntax never will.

Too many laws and too many properties

One of the remedies called in to defuse the apparent fishiness of this form of parallelism has been the appeal to laws. Since I have nothing new to add to points that have been made elsewhere I simply rely on the conclusion drawn by Block¹⁰. Many different laws have been called into service. The solution turns out to leave the epiphenomenalism intact.

When I first presented some of the reasons to be suspicious of functionalism¹¹, it was objected to me that dealing with causal claims we

⁹ Ned J. Block, "Can the Mind Change the World?"; in MacDonald C. & MacDonald G. (eds.) *Philosophy of Psychology*; Oxford, 1995; p. 30.

¹⁰ In "Can the Mind Change the World?", *op. cit.*; pp. 37-39. I voice just a mild skepticism about the strength of our grasp of what is involved in the distinction between *ceteris paribus* and strict laws. Philosophers of science, most notably Nancy Cartwright, suggest that all laws may in some sense need contextual parameters to be fixed. Philosophers of language, most notably Stephen Schiffer, deny to nonstrict laws even the status of laws.

¹¹ At the 1994 meeting of the Italian Society of Analytical Philosophy. Thanks are hereby extended to all of those who took the pain to listen and answer. In scattered order Roberto Casati, Roberta De Monticelli, Mauro Dorato, Simone Gozzano, Francesco Orilia. The second incarnation of this piece has been read to all members, students and faculty, of the institute of philosophy at my university. To all of them my

always have the theoretical option to take properties to be causally relevant. The idea has its merits¹². Let us explore then the possibility that in the case of mental causation of behavior it is properties of states that are causally relevant. If a mental state is individuated functionally one may think of its properties as both mental and physical. Now which properties are causally relevant in the sense a realist would want them to be? As long as one accepts Block's premise 3 and the "closure" of the physical it would appear that the physical/syntactic properties of the mental state are doing all the causal work. One of the reasons many nowadays greet substance dualism with an incredulous stare is its interactionist mystery. It is never clear how a substance within space-time would be causally related to one without space-time. The duality of properties does rebut the incredulous stare by firmly placing mental and physical properties very much within the same framework. What it does not do is to assuage the qualms of the mental realist (in the sense sketched above.) For if the syntactic properties are doing all the causal work (in terms of causing behavior) the mere fact that there are mental properties that got attached to them is not enough to show that the mental (or intentional properties) are efficacious¹³. It has to be assumed that syntactic properties can have a

thanks too. A special word for Elisabeth Pacherie: though an exceptionally good researcher herself she played the role of schoolgirl making me realize how many times I am too quick, too allusive, assuming too much, and at the end of the day not clear enough for all schoolboys and schoolgirls. To her cruel sense of clarity and insight the shape of this paper owes much. To her wonderfully non cruel person I owe much that would not be easy to spell out. The paper is dedicated to her, humbly asking for patience when I am not up to her standards.

¹² I am not sure there is a clear sense in which a property in and by itself is causally efficacious. It would seem to me it is the instantiation of a property which is efficient. The point may be more notational than substantial though. In the paper nothing depends on what "carriers" of causality one's ontology will select at the end of the day.

¹³ The problem here seems to be left unsolved even by those who have a direct explanation of intentionality. John Searle thinks there is an intrinsic intentionality of the brain which has a sort of "dual aspect." See, e.g. the following passage: "My present state of consciousness with all its subjectivity, intentionality, causal powers, phenomenology, 'qualia', and the rest of it, is just a feature or set of features of my brain. And it is caused by the behavior of lower level elements of my brain in the same unmysterious sense of 'cause' that the solidity of this table with all its features is caused by the behavior at the lower level of molecules and at the same time the solidity is just a higher level feature or set of features of the entire system that is composed of those molecules." (see John R. Searle, in *John Searle and His Critics*; Oxford, 1991; p. 141.) Ned Block notes that Dretske has a similar problem vis-a-vis epiphenomenalism. It appears to me that all biological explanations of the fixation of

causal role via their computational implementations, i.e. via the brain states and processes materializing them. An open avenue can be to claim that while syntactic properties are the causally relevant ones (with regards to the brain that realizes them) the properties responsible for mental content are **second-order** properties defined over first-order ones. Arguably mental states as semantically valuable states can be seen as second order properties. To use a medical example: being anaesthetic is a second order property many different substances possess. Pentotal is anaesthetic in virtue of the (chemical) properties of thiopental sodium. Curare works as or is anaesthetic via the properties of alkaloids. Though my chemistry does not go beyond high-school level one could decide that the proper level of causal explanation is chemical¹⁴. Alkaloids as well as sodium have causal efficacy with regards to my being anaesthetized. The properties may in fact be rather different (in this case they just are: the kind of effects are different although both anaesthetic.)

Now, if this can be taken as a model for functionally conceived mental states, it seems to me that the realist about mental states would view functionalism as fouled by weak epiphenomenalism. The connection between the property, chemical or physical more generally, that seems to do all the causal work and the second order property (which ought to coincide with the mental state as a contentful state) can be as nomological as one likes. In fact one can describe a variety of laws in one sense or another that permit us to group together all anaesthetic substances as well as all states that can be the belief that grass grows. All those laws will be instances in which we have nomic correlations. Unless we are able to show that second order properties are causally efficacious the correlations themselves will be exactly what weak epiphenomenalism is describing¹⁵.

intentionality face a similar challenge. The issue though would have to be explored in detail.

¹⁴ I am less impressed than Block by the worry that there may not be a fundamental level of everything. At each level of explanation one can decide what the right causal structure should be. For brains that could very well be chemical, unless there are more serious reasons than what we have so far to call up quantum-theoretical phenomena.

¹⁵ Block in "Can the Mind Change the World?", *op. cit.*) presents a very limited case for the possibility of envisaging a sort of generalized placebo effect. Being told that substance alpha has the power to put me to sleep, I go to sleep if given substance alpha which has none of the first-order properties directly connected with sleep. I am not sure at all that we could generalize such a phenomenon to the causal efficacy of mental states. One of the reasons is mentioned by Block himself: it takes an intelligent being to be affected by placebo effects (no placebo for plants, or so it seems. The

Summary, Coda, and some very provisional conclusions

It seems to me that the orthodoxy of functionalism, as perhaps many orthodoxies, solves many problems and opens up a bunch of new ones, or reopens old ones. In the scorecard for what was really answered is the issue that multiple realizability brings about for any type-type identity theory for mental states. At the same time it makes an excellent use of the best technology available to model at least some of the mental functions. I see this as a real advantage.

What is a bit stickier is the sort of worry philosophers may have. In the philosophy of mind the *desiderata* are debatable. All I said here depends on the crucial assumption that mental contents have to be causally efficacious and not just relevant in explanatory terms. Mental contents can be explanatory at some level or other and remain epiphenomenal at the causally explanatory level. This I take to be a view that entails their fundamental uselessness in terms of the influence the mental has on our actions and all forms of behavior. This is a debatable point since I am excluding possibilities that are perfectly coherent (and at time they seem even more plausible given the very strangeness of the mind-body problem.) For instance the very possibility that mind are syntactic engines and semantic aspects are purely illusory. It would take another series of papers to see why it would be desirable to have real semantics, 'real' in the Platonic sense foreshadowed in the second quotation on top of this paper.

For the nonce all I want to point out is the following: functionalist readings of mental states are the logical ancestors of epiphenomenalism. Functionalism breeds epiphenomenalism on the assumption that there is no causal trade between the purely abstract and the very mundane implementation levels.

Detours via even more complex notions of laws and property theories seem to bring about more problems rather than less suspicious looking results. The argument for all of this is, in its skeleton form, trivial. The abstract realm of functional characterizations is not per se causally efficacious. Whatever is mental in a mental state seems to have something to do with the semantics of its contents. Those are, at least, not fully

second one is that if something akin to placebo effects is the right model of neural communication we would have a completely absurd picture of the brain. If so the entire computational model has to be dragged back to its Turingian origins.

On these points, on quantum mechanics, and many others, thanks are hereby expressed to A.L.F. for unstintingly disposing of her time to listen to my nightly rantings.

captured by the functional characterizations (be that because of externalism in the philosophy of language, or the specific theory of reference one adopts, or because of Putnam's worries about the ineliminably normative nature of meanings, etc.) What seems to be doing all the causal work is the concrete implementation of the mental state, the neural, or neural together with even more physical stuff, state underlying it. Hence it looks as though the mental *qua* mental is idle. Hence epiphenomenalism.

Is epiphenomenalism really bad? That depends on the kind of theory of mind one prefers. For my part as a sort of penultimate barricade I would tend to shy away from epiphenomenalism insofar as possible.

Nearly anything should be tried before giving in. "everything" including perhaps abandoning the ontological neutrality of the functional characterization of mental states.

The wages of the "sin" of not making any commitment in terms of the ontology of mental states are then paid with the currency of epiphenomenalism.

As for the redemption: a moderate dose of reimmersion in physiofunctionalism in the short run will certainly do *in lieu* of penance.

The suggestion may seem retrograde and backwards, and I won't even try to defend it.

I state it in every simple terms: what has to give is the idea that ontological neutrality is the best way. We may be forced to restart looking at local forms of reductions that would entail that nothing is common between my pains and my pleasures, my thoughts and their contents thereof and those of my Martian cousins. Local reductions in this sense would be species-specific and the multiple realizability would go by the board.

This it too quick and I am sure I bored any reader more than required.

Recondita armonia: Coda (only for those not bored enough)

The mind is its own place, and in itself
 Can make a heav'n of hell, a hell of heav'n.
 John Milton, *Paradise Lost*, Bk I, 254-255

Jerry Fodor in his Jean Nicod Lectures¹⁶ faced the difficulties and presented a line of reasoning he summarizes as follows: ... I'm not going to argue that psychological laws *should* be broadly construed, and I'm not going to argue against narrow content. A fortiori, I'm not going to argue that the notion of narrow content is incoherent or otherwise infirm. What I *am* going to argue is this: the considerations that have been supposed to show that an externalist construal of content won't meet the purposes of psychological explanation are, on balance, unconvincing. So maybe narrow content is *superfluous*." [Fodor 1994, 28]

The strategy Fodor adopts is to defuse the impact of what have been taken to be the main reasons in favor of the adoption of some sort of dual account of mental content (narrow/wide, external/internal, etc.) The troublemaking cases are neatly placed in the categories of Twins, Experts, and Fregean. The Twins and the Experts are cases alleged to present unsurmountable difficulties for internalism, the Fregean cases are all the clones of the Morning Star which is the Evening Star and, all the same, the Babylonian rational agent kept thinking that 'The Morning Star is bright' is *not* a notational variant of 'The Evening Star is bright', due to lack of astronomical sophistication or something like that.

I myself find the Twin cases unconvincing at the end of the day and not only for Fodor's reasons. Twin cases are thought experiments in which Twin environments or Twin agents are placed but one, usually rather small element is changed. One could conceive of some possible world in which everything is like here, but water is not composed by the same chemical elements, even remaining the main way of quenching one's thirst. They work as intuition pumps and I don't find much force at all in the intuitions that are constantly put to use¹⁷. Much is there to appreciate in Fodor's treatment of the Expert cases. These are deferential uses of concepts. There are demonstrably many cases in which using a concept we defer to

¹⁶ Now published in *The Elm and the Expert*, Cambridge, Mass: MIT press 1994. Numbers in square brackets [Fodor 1994, #] are all from the published text.

¹⁷ Similar doubts on the intelligibility of thought experiments carried out in the philosophy of mind have been expressed by Kathleen V. Wilkes in her *Real People*, Oxford 1988 (see in particular Ch. 1 and 2.) The gist of her points is that quite simply we have a very slippery grasp, if any at all, of what is possible under the purported environmental conditions described by the fictional setting. To use one example: it may very well be possible, or so it seems to me at least under several notions of possibility, that nothing at all can have the same phenomenal properties as water and at the same time fail to be H₂O. Unless of course one starts varying more and more of the background involved, getting a weaker and weaker grip on the content of our "linguistic" intuitions.

an expert opinion on its extension, or even on the precise linguistic usage. Many of us don't really know much about glasses spin except that it is an interesting subject very much studied in Denmark and in Rome (this is the extent of my precise understanding of it, for instance.) We defer to experts to get the right extension conditions. Fodor has the excellent idea of seeing this as a case of exploitation, rather than as a case of linguistic division of labor. We exploit experts much as we could exploit mechanical arms in a radioactive environment. The Experts are more pervasive than one can think: even fairly everyday concept such as chair or sofa need an expert to defer to [see Fodor 1994, 36]: the hiring of an interior decorator to get one's sofa properly color coordinated is akin to asking what 'sofa' means to the decorator, the expert in the right conceptual field.

The Fregean cases are a bit more difficult and I would venture to say that even Fodor leaves much to be desired in his treatment. The key idea exploited in *The Elm and the Expert* is that while psychological laws are laws, they are such only with the *everything else being equal* qualifier. The move is to claim that Fregean cases (Oedipus who wants to marry Jocasta-*qua*-"Jocasta" and does not want to marry Jocasta-*qua*-"my mother", and all the related affairs in which poor Jocasta is his mother after all) are indeed empirically given as exceptions. If they proliferate psychology as a system of intentional laws is impossible: "... *ceteris paribus laws tolerate exceptions*, so long as the exceptions are unsystematic." [Fodor 1994, 39]

Fregean cases aren't systematic, but unlike water-as-H₂O and water-as-XZ in parallel universes, they do happen a lot [Fodor 1994, 49]. And this alone should give us pause. To hang the theory on an empirical claim about the frequency of Fregean cases is risky indeed and maybe not necessary.

Now I don't think I am able to consider all the possible variations Fodor looks upon, the main one being his refutation of Quine unscrutability of reference.

I prefer to finish with a remark more directly connected with the main body of what interests me here, namely that one line of thought in philosophy engenders epiphenomenalism. I remain ultimately neutral, and not penultimately though, on whether epiphenomenalism itself should be embraced as a welcome consequence or whether we should struggle when dead from the waist down to see it explained out.

The solution presented by Fodor is to the effect that the narrowness of content (if any) is captured by the semantics of the piece of mentalese sentence that encodes a proposition. These sentences are indeed as

finegrained as needed to be descriptive of all modes of presentation we can come up with. It looks as though it is the computational role that does all the explanatory work (in terms of explanations of behavior.) Again narrow content is shown to be, though quite possibly not non-existent, superfluous. Fodor claims this to be the line of thought pursued by Stich [Fodor 1994, 50-51] and defends his contention via a renewed form of harmony. "One can, I think, imagine a world where everything is delicately balanced in the following way: Content is broad, the metaphysics of content is externalist (e.g. causal/informational), and modes of presentation are sentences of Mentalese. Modes of presentation with similar causal histories (or nomic affiliations; anyhow with similar broad contents) overlap enough in their syntax to sustain robust psychological generalizations. But not enough to make the minds that these generalizations subsume homogeneous under *syntactic* description." [Fodor 1994, 52-53]

The picture painted is Leibnizian and the delicate balance depends on the following contingent fact (or hoped for fact): the causal, historical path of production of broad contents is such that its mentalese encodings are sufficiently similar to warrant our subsuming intentional explanations that are robust. Fodor notes that "maybe our world *is* like that, whether it is, *is* strictly an empirical issue." [Fodor 1994, 53]

While I think it is perfectly correct to see this as an empirical issue, I find no reason to believe that differently assembled lines of broad contents, even when they churn out the same product contentwise, should have produced modes of presentation similar in the required sense.

Or perhaps I just see much less harmony in *this* world, and that's where I am.

Adriano P. PALMA
INSTITUTE OF PHILOSOPHY
NATIONAL CHUNG CHENG UNIVERSITY
160, SAN-HSING, MING-HSIUNG
CHIA-YI 621, TAIWAN R.O.C.
& CREA\Ecole Polytechnique, Paris

Bibliography

- Castañeda, Hector-Neri (1967) *Intentionality, Minds, and Perception*, Detroit.
Fodor, Jerry (1994) *The Elm and the Expert*, Paris and Cambridge, Mass.
Guttenplan, Samuel A (1994) *Companion to the Philosophy of Mind*, Oxford.
Loewer, Barry & Rey, Georges (1991) *Meaning in Mind: Fodor and his Critics*, Oxford.

- MacDonald, C. & MacDonald, G. (eds) (1995) *Philosophy of Psychology*, Oxford.
- Pacherie, Elisabeth (1993) *Naturaliser l'intentionnalité*, Paris.
- Putnam, Hilary (1994) *Words and Life*, Cambridge, Mass.
- Rosenthal, David M.(1991) *The Nature Of Mind*, Oxford-New York.
- Warner R; & Szubka T.(1994) *The Mind-Body Problem*, Oxford.
- Wilkes, Kathleen V. (1988) *Real People*, Oxford.