EPILOGUE

After a marathon colloquium on observer mechanics we were approached by Matthew and Ida. Rumor had it that their relationship was stormy: they were quite contentious, each frequently found the other's point of view entirely unintelligible, and neither hesitated to say so. We expected the worst. However the ensuing conversation proved, to our relief, to be relatively free of hostilities and at points even edifying. Matthew had recently enjoyed the upper hand in his arguments with Ida, and he led off.

Matthew: Quite an interesting, albeit long, colloquium. Ida and I agree on little, but we both agree that you've left a lot of questions unanswered. Can we talk?

One of us: Most certainly. What's on your minds?

Ida: Lots. But for starters I'll be blunt: Is a fork an observer? When you say things like "the objects of perception are other observers" it sure sounds like you're saying something of the sort.

O: Not at all. A fork is a conclusion, not an observer.

M: You did say that the objects of perception are other observers, didn't you?

O: To a first approximation, yes.

M: Well if a fork isn't an object of perception, I don't know what is. And from the statements "forks are objects of perception" and "objects of perception are other observers" it surely follows that forks are observers.

O: It certainly does. But we don't buy the first statement. Forks aren't objects of perception under our definition of that term.

I: Could you remind us of your definition?

O: Surely. The objects of perception for an observer or a participator are those entities with which it interacts in an act of perception.

M: Then you deny that when I look at a fork I am interacting with the fork?

O: That's right. But what about you? Would you want to assert that when you look at a fork, the entity you're really interacting with is the fork itself?

M: Not really. I guess I'd say that the entities I'm interacting with are the fun-

damental constituents of the fork—it's quarks and leptons and whatnot. But I don't think this'll buy you much. If my true objects of perception, the things I'm really interacting with when I perceive, are quarks, then it seems that you're committed to saying that quarks are observers, aren't you?

O: Not at all. What goes for forks goes for quarks. Quarks and leptons aren't what we're interacting with in perception any more than forks are. And since we don't think quarks are objects of perception we're not committed to saying that quarks are observers. In fact, we think they're not.

I: That sounds fine to me. But that's because ...

M: That's because you don't keep a respectable ontology.

I: Do you want to get into it now?

M: Sorry. No. Let's continue to discuss with them.

I: Fine.

M: How can you say that elementary particles aren't the objects of perception, given all that we know about the physics of light, the optics of the eye, and the physiology of the visual system? There's a known causal path beginning with distal elementary particles, continuing with emitted photons, followed by absorption of the photons by rods and cones, and concluding, after some complicated neural processing, with perception. It would seem that you're contradicting some well-established scientific facts.

O: We see no contradiction. An observer, given some premise s, perceives that interpretation or those interpretations which are given nonzero weight by its conclusion measure $\eta(s, \cdot)$. These interpretations are encoded in a systematic representational scheme that we call X. If some physical objects or physical properties are among the symbols employed by this scheme, then these may be perceived. But the observer isn't interacting with its own symbols, it's interacting with other observers. The conclusions an observer reaches are tied to statistical properties of the dynamics of this interaction. In short, observers interact with observers; physical properties are among the symbols employed by observers to represent aspects of this interaction. The scientific story you just told is fine as far as it goes. But, so to speak, "behind" the physical symbols is the dynamics of observers those symbols represent.

M: Your ontology is no more respectable than Ida's.

I: I warned you.

O: Is there something amiss in our definitions? We've endeavored to define observers, participators, reflexive frameworks and all of our participator dynamics in a manner free of formal errors. But perhaps we've failed somewhere.

M: I don't know if you have. I've not had a chance to examine it carefully.

O: But then what's wrong with our ontology? Our commitments are restricted to those of logic and set theory. Set theory has problems, certainly, but we are no more in trouble on this count than is contemporary physics.

M: The problem is that your account isn't naturalized, nor does it seem amenable to naturalization. Look, physicists are going about finding and listing the fundamental categories and properties of the world. The list needs work, no doubt, but when they're all through it'll contain things like charge, spin, mass, and charm, but not things like observers and participators. But you're making observers and such fundamental categories in your theory. That is *prima facie* implausible.

O: If we were taking the terms "observer" and "participator" as ill-defined or undefined primitives we might agree with you. Such primitives would be poor foundations for a theory of anything. The appropriate move would be to try to naturalize them at best, and completely abandon them at worst. But we take observer and participator as technical terms with rigorous definitions. And we don't share your ontological bias. Rather than naturalize these technical terms our project is to "perceptualize" the technical terms of physical theories. That is, we want to show rigorously how the categories and properties employed in physics could arise (1) from statistical properties of certain dynamics of observers and (2) as aspects of the representational schemes employed by certain observers to describe these statistical properties.

M: Good luck. Go "perceptualize" quantum field theory and then let's have dinner.

O: We might need a rain check. Things take time, but we have some interesting leads. It happens that many physical properties, like spin and mass, turn out to be properties of the representations of algebraic groups. We mentioned, you'll recall, groups, Hilbert bundles, systems of imprimitivity and the like in our colloquium. Well our bet is that the groups that crop up in physics are intimately related to the groups that crop up in our reflexive observer frameworks and to the symmetry groups of the transition probabilities of observer dynamical systems. And it appears that the Hilbert spaces so ubiquitous in quantum theory might arise from

the linearization of specialized reflexive frameworks. If so, then the notion of a physical state—namely a measure on the logic of subspaces of an appropriate Hilbert space—might be grounded in observer theory. And then quantum measurement theorists and perceptual theorists might have something substantial to talk about.

I: Sounds like an interesting direction to me. I was going to ask whether, although you deny that forks are observers, you would at any rate assert that forks are composed of observers. But what you just said suggests otherwise.

O: That's right. Physical objects are symbols employed by observers to represent aspects of their interactions with other observers. Physical objects aren't conglomerates of observers. Forks are no more composed of observers than a newspaper is composed of the people and events it describes. We don't endorse panpsychism.

M: Do you deny the existence of quarks or forks?

O: Not at all. Nor would we deny the existence of the symbols used by a Turing machine in its computations. Both forks and quarks are symbols employed by observers. Our view on forks and quarks shares something in common with the "internal realism" of Putnam. We agree with Putnam's rejection, on the one hand, of the metaphysical Realist view, say as put forward by Sellars, that denies the real existence of "middle-sized" physical objects such as ice cubes, and that grants existence only to the particles of physics and their occurrent properties. Putnam rejects it as embodying untenable dichotomies, for example a dichotomy between properties physical objects have "in themselves" versus properties projected on them by the mind. We also agree with Putnam's rejection, on the other hand, of complete relativism— relativism that goes beyond the acknowledgement of different "versions" of the world to the claim that truth is just consensus or some such. And we agree with Putnam that a quark, no less than a fork, is a version-relative notion; and that this impugns the ontological status of neither. Where we differ with Putnam, of course, is our specific proposal that quarks and forks are symbols employed by observers to represent properties of the dynamics of *participators*.

I: If matter isn't composed of observers, what about the converse? Would you say that observers are composed of matter?

O: No. According to our definition, an observer is composed of six parts: X, Y, E, S, π, and η.

I: Certainly. But although I have no trouble with this, I imagine Matthew

would not be comfortable without some assurance of token physicalism regarding observers. I mean, he'd say it might be fine to have an abstract definition of observers, but any particular instance must somehow be physically instantiated. Maybe you endorse some kind of property dualism: matter has physical properties and observer-theoretic properties.

O: We do have a notion of the instantiation of an observer or a participator, but it doesn't amount to token physicalism or property dualism.

M: This doesn't sound good. What was your notion of instantiation again?

O: Recall that each observer is an inferential system. It gets certain premises and reaches certain conclusions. If O is some observer, where does O get its premises? Well from other observers. There is some collection of observers, say T, whose conclusions or their deductive consequences are the premises for O. We called these observers "transducers" for O. They are the first level of instantiation of O. Of course each observer in the collection T has, in turn, its own transducers. And so on ad infinitum, presumably. You can picture the instantiation of O something like an infinite cone with its tip at O and getting wider and wider as one goes through successive levels of transducers.

M: Fine and dandy. But if there's no matter in the instantiation, how can you see one? I for one have never encountered a perambulating cone.

O: In an interesting way, matter is involved in our notion of instantiation.

M: Pulling in your horns pretty quickly, sounds to me.

O: Not really. The idea is that there are many levels—infinitely many levels—of observer dynamics taking place in the instantiation cone of O. The way O represents the statistical properties of the dynamics of observers in its instantiation depends on how far down in the cone that dynamics is taking place. For dynamics near the top of the cone, close to O itself, the representation tends to be more "psychological" whereas as one goes down the cone the representation becomes more "neurobiological" and then more "physical" and then ..., well there's no bottom that we know of.

I: Then you would deny a principled distinction between mind and body?

O: Yes. "Mind" and "body" are convenient terms to distinguish between levels of the instantiation cone for a given observer. Higher levels, or rather an observer's representation of the dynamics at these higher levels, are "mental." Somewhat lower levels are its "body." Even lower levels, unrepresented and as yet unexplored, presumably also exist. And since what

is mind and what is body are relativized to the observer, a dynamics which is mental relative to one observer may be physical relative to another, and vice versa.

M: This sounds worse all the time. Physical properties and physical entities must be the ontological bedrock. Any theory of mind must be built upon these, or at least be compatible with these.

I: Actually, I like this observer story. I've always felt that the physical could be reduced to the mental, or eliminated entirely.

O: Let's be careful here. We don't really claim to be reducing the physical to the mental. We're saying that both the physical and the mental are derivative upon something more fundamental, namely an infinite "hierarchy" of levels of observers/participators in dynamical interaction. And we certainly don't claim to be eliminating the physical. If anything, we're proliferating the physical. Since what's physical depends on which level of the observer hierarchy you're talking about, there's no such thing as *the* physical world, but rather there are infinitely many "physical worlds."

I: This is starting to sound rather like the monadology of Leibniz, what with hierarchies of perceivers and all.

O: There is some resemblance, but there are important differences. First, monads are rather loosely defined. Certainly not well enough to attempt to build a quantitative science on them. Observers and participators, on the other hand, have precise mathematical definitions. Second, whereas Leibniz postulates a preestablished harmony between the activities of non-interacting monads, observer theory postulates a stochastic dynamics of interacting participators, with markovian kernels that have been characterized precisely. Third, as one goes down the levels of monads, one encounters successively impoverished modes of perception. Whereas it is completely compatible with the observer-theoretic hierarchy that different levels are, in some sense, equally rich—just different. And finally, we don't yet know if our hierarchy reaches to the City of God.

I: And I take it, given your account of mind and body that, *pace* Berkeley, you would not want to say that physical objects, elementary particles, and so on are existentially dependent on one's mind?

O: There are a couple of reasons why we wouldn't say that. First, there is no mathematically precise definition of mind that is generally accepted, functionalist accounts not withstanding. So such a statement would not, for now, be at a level of precision required for dialogue with, say, a quantum measurement theorist. Second, whatever notion of mind eventually does

emerge, we anticipate that it will be derivative upon the notion of hierarchies of dynamical systems of observers. So that the more fundamental issue is the existential dependence of physical objects on observers and participators. And third, here the relationship isn't one of simple existential dependence. Physical objects and properties are, we have claimed, among the symbols or representations employed by observers and participators, so that they are in some sense "parts" of these observers/participators. But that's not the whole story. Each participator is not alone. There are dynamical interactions among participators, and although a given participator contributes to its own dynamics and, indirectly, to the dynamics in its instantiation, there is much more to these dynamics than just the contribution of this participator. Something "independent" of or only partially dependent on this participator is going on, as is clear from that fact that the other participators are governed by their own action kernels. But since, on our story, it is statistical properties of this dynamics that are represented by the given participator's symbols, and since this dynamics is partially "independent" of that participator, it follows that the tokening of particular symbols by the participator depends in part on the participator and in part on its "environment." The statement that such symbols are existentially dependent upon the participator is just too simple. It only catches one part of the whole elephant.

I: This sounds a bit Kantian, the idea of having a supersensible realm which is behind the realm of experience and which, in some fashion, drives that realm of experience.

O: Perhaps a bit. But the differences are crucial. For Kant the supersensible thing-in-itself is unknowable and not a potential subject of scientific enquiry. For us the supersensible realm of participators in dynamical interactions is, although perhaps not directly knowable, still a subject of scientific enquiry. There is nothing unusual about exploring the unobservable through science. That's how we know about quarks and thermonuclear processes in the sun. Similarly for participators and observers. We can't see them directly, but we can legitimately construct theories of participators and their dynamics, and then look for empirical consequences that can be checked.

I: I take it then that you don't embrace phenomenalism?

O: Right. We don't take "elementary sensations" such as colors, sounds, spaces, and times as the constituents out of which physical objects are constructed or as the foundation, incorrigible or otherwise, upon which all else is built. Quite to the contrary, we take unobservable entities—observers

and participators—as the (far from incorrigible) foundation. Sensations no less than physical objects are, for us, the corrigible conclusions of the nondemonstrative inferences of observers. This view of sensations is, by the way, one reason why we took time to point out the problems in defining transduction. By rejecting the widely held notion that one can point to a distinguished single stage of transduction in, say, vision, we were rejecting even the more recent (and nonphenomenalist) suggestions of a demure foundational status for sensations. Instead, we relativize the definition of transduction to the observer, so that what is "directly detected" relative to one observer is, relative to another, the conclusion of a nondemonstrative inference. If there are epistemological foundations, they are not to be found in sensations.

I: Can perception, then, yield knowledge? Say knowledge in the traditional sense of justified true belief?

O: For a specific premise s, the conclusion $\eta(s,\cdot)$ of a participator is true if the probabilities it assigns to the interpretations in $\pi^{-1}(s) \cap E$ match the actual probabilities generated by the dynamics in which it participates. The conclusion is justified if η is a regular conditional probability distribution (rcpd) with respect to π of the stationary measure of this dynamics. Justified true conclusions are, perhaps, knowledge.

I: It would seem, then, that your participators could have perceptual knowledge without being certain that they have it?

O: That's right. In fact it seems they can't be certain. A participator cannot determine if it's interpretation kernel η is the appropriate rcpd.

I: How then can it happen that η turns out to be the appropriate rcpd? This seems a rather unlikely coincidence.

O: For this to happen there must be an appropriate relationship between the action kernels and the interpretation kernels of the various interacting participators. The action kernels, you see, determine the transition probability of the participator dynamics, and it's the stationary measures of this stochastic dynamics for which the interpretation kernels must be rcpd's. Whether there are particular strategies that participators could adopt (for example, special kinds of action kernels) to enhance their chances of true conclusions is a topic for further research. But it is clear that a participator P "wants" not only its own conclusions to be true, but also the conclusions of participators at the various levels of its instantiation to be true as well—for these conclusions eventually become P's premises. This means that at the various levels of dynamics in the instantiation cone of P the conclusions of the instantiating participators should also be the appro-

priate rcpd's of the stationary measures of their respective levels. At the "biological" levels of the instantiation cone of P this matching of rcpd's to measures might appear, relative to P, as an appropriate "adaptation" of the instantiation of P. Thus on this particular point we apparently agree with evolutionary epistemologists such as Popper and Campbell: there is a continuity, a formal similarity, between the processes which eventuate in knowledge and those which eventuate in biological adaptation.

I: But you don't seem to buy their representationalism.

O: Not to the extent that they take the external world, the world that is to be in some fashion represented, as a physical world—a world of forks and quarks. For us the represented world, the "World 1" in Popper's terminology, is the world of observers and participators in hierarchies of dynamics. The world of forks and quarks is a world of representations, not the world to be represented.

M: You seem to put a lot of weight on the dynamics of participators, and I'm not sure I have an intuitive grasp of this dynamics. Can you help with an example?

O: We can try. But remember, there are three kinds of dynamics of relevance to a given participator P. First, there's the dynamics on the reflexive framework of P itself. Here P is interacting with other participators that are on the same framework with it. Second, there are the various dynamics going on in P's instantiation. The asymptotic properties of these dynamics eventuate as premises for P. And third, there are the various dynamics "above" P. For P is itself involved in the instantiations of higher observers and participators.

I: This just follows from the hierarchy of participator dynamics you mentioned before, right?

O: Exactly. Now let's look at an analogue of the second type of dynamics, the dynamics of P's instantiation leading to a premise for P's own inferences. We say "analogue" because this example isn't really a participator dynamics at all, but an example drawn from the physical realm to help intuitions. We discussed a real example, you'll recall, when we talked about instantaneous rotation observers and the incremental rigidity scheme. So if the analogue doesn't help, forget it and think about the real example.

M: Enough caveats. Do proceed.

O: Consider your perception of pressure and shape when you press your fingertip on a corner of a table. The physiologists tell us that your sensory experience is dependent upon various cutaneous mechanoreceptors, such

as Pacinian corpuscles and Merkel cells. The physicists tell us that at an even more microscopic level your sensory experience is dependent upon the dynamics of many atomic and subatomic particles. Both the table and your finger are composed of such particles and, before touching your finger to the table, the dynamical systems of particles for the finger and for the table have each their own kind of stability—as evidenced by the fact that each retains its own shape over time. But now when you place your finger on the table corner you are bringing these two dynamical systems into contact and letting them interact, with the consequence that a new stability of the table/finger system of particles is reached. Of course the new stability requires more give and take on the part of the finger system than on the part of the table system, with the result that the original stability of the finger system becomes quite perturbed and gives way to a new, quite different, stability. This is evidenced by the new dented shape of the finger. It is this change in the stability of the dynamics for the finger system that is picked up by the mechanoreceptors and eventually experienced as pressure and shape. So here we get a glimpse of how stabilities of dynamical systems at a "lower" level can be premises for perceptual inferences at a "higher" one.

I: We aren't to infer from this example, however, that dynamical systems of physical particles are the same thing as dynamical systems of participators?

O: Right. Physical particles are not participators.

M: What leads you to suggest that the objects of perception are other observers or participators?

O: Chronologically the definition of observer came first. As early as 1980 we tried to write down a formal structure common to the theories of (e.g.) structure from motion, stereo, and shape from shading that had, at that time, recently been developed. This structure was refined over a period of seven or eight years, as we continued to study new theories of specific perceptual capacities and to develop our own theories of structure from motion and shape recognition. When we finally had in hand a formal definition of observer that we were reasonably comfortable with, the question naturally arose whether the same could be done for the objects of observation. It seemed we had a fundamental choice of strategies. We could either propose that the objects of perception have a formal description that is fundamentally distinct from that of observers, or we could propose that they have the same, or related, formal descriptions. Proposing a fundamental distinction seemed problematic: it would introduce a dualism, it

would require a new formalism for the objects of perception together with a justification for this formalism, and it would require a new formalism to interrelate observers with these objects. So too did the alternative—proposing that objects of perception are observers. It would require us to figure out a formalism to relate observers with other observers so that mutual observation became possible. And it required us to face the obvious objections of the "Is a fork an observer?" variety. When we discovered the possibility of constructing reflexive observer frameworks, and thereby the possibility of mutual observation between observers, we decided to further pursue the nondualistic alternative. The vindication or rejection of this approach must await the further development of the resulting theory and the testing of its empirical consequences.

M: Your nondualistic approach seems to have the unsavory consequence that physical events do not cause other physical events. Is that a correct reading?

O: So it would seem. Physical events are employed by observers to represent properties of the dynamics of participators. Any notion of cause must derive from this dynamics, not from the symbols used to represent this dynamics.

I: Perceptual learning was conspicuous by its absence from your colloquium. Does observer mechanics have anything to say about it?

O: Actually it was absent only in name. Our entire development of participator dynamics can be viewed as dealing with perceptual learning. As a result of observations a participator updates its perspective π and its class of possible perceptual conclusions η, all under the dictates of its action kernel.

I: But the participator dynamics you develop is markovian—one can make the best possible prediction about the future behavior of the dynamics based only on knowledge of its current state. Isn't this a rather special case, not really suited to be a general theory of perceptual learning? What about the possibility of learning that depends on a past history, not merely on the current state?

O: It is true that the participator dynamics is markovian, and that this means that the present state is the best possible predictor of future behavior. However, the dynamics of various subsets of the participator system are typically quite nonmarkovian. And the markovian formalism is far more general than at first it might seem. For if, in one formulation of the notion of a dynamical state, it happens that predictions of the future are best conditioned not just on the current state, but on a finite history of

states, then it is always possible to reconstruct the notion of state and the description of the dynamics such that it is markovian. Thus by formulating participator dynamics as markovian we include learning that depends on finite past histories of any length.

I: Also conspicuous by its absence was introspection. If you intend to extend observer theory beyond perception to include cognition as well, as it appears you do in your attempts to define the notion "cognitive" in terms of observer theory, then you can't ignore introspection.

O: Most certainly. We don't see a principled distinction between cognition and perception. The considerations that have led Fodor, for instance, to suggest that there is a distinction—namely, the relative isotropy and unencapsulation of cognitive or "central" processors, on the one hand, and the domain specificity and encapsulation of perceptual "input analyzers," on the other—can all be satisfied by the relativized notion of "cognitive" and "transductive" that we introduced in discussing the issue of theory neutrality of observation. So we do face the task of interpreting such activities as introspection in terms of observer theory. We have no detailed account of introspection at present. But perhaps introspection on a given reflexive framework is performed by participators on another framework which take as their premises finite sequences of dynamical states (or of conclusions) of the given framework. On this provisional construal of introspection, it is a finitary precursor of specialization.

M: Is observer theory falsifiable?

O: Certainly. Take, for instance, our "observer thesis"—that every single perceptual capacity has a natural description that is an instance of the definition of observer. This thesis can be disconfirmed by counterexample. And were a counterexample to appear, we would have to rework or abandon the theory.

M: If you say physical entities are symbols employed by observers to represent aspects of participator dynamics, then what do you say about spacetime itself?

O: Roughly, our ideas are like this. There is not one global time, but different times at different levels of the hierarchy of participator dynamics. As one goes down the instantiation cone, for instance, of a participator one finds that at each successively lower level the time scale "speeds up." This is because typically it is properties of the *asymptotic* behavior at the lower level that serve as premises for observers and participators at the next level up. At a given level of the hierarchy, moreover, there is not just one time either. Rather each participator has its own "proper time" which increases

with every "channeling" in which it participates. In short, the unit of time for a participator is a discrete act of observation. A single unit of time for one participator will correspond to many, perhaps infinitely many, units of time for each participator in its instantiation. Now the rate at which a participator channels with others depends on the "τ-distribution" and the group difference-in-perspective of the other participators with which it interacts. The τ-distribution, then, governs the way in which time and "distance" trade off, and it is the key to an observer-theoretic account of the relativity of physical spacetime. Much is yet to be worked out, but the big picture is again that spacetime is part of the scheme employed by participators to represent properties of their interactions with other participators.

I: What other areas for further development do you see in observer theory?

O: Many. Here's an abbreviated shopping list. We have already mentioned introspection, and the project of perceptualizing physical theory (but not, of course, in any phenomenalist or idealist sense of "perceptualizing"). It would be interesting to develop the notion that physical properties are properties of the dynamics of interacting participators, and that therefore these properties hold only so long as the dynamics continues. Pursuing this may lead to an observer-theoretic understanding of the Aspect experiments. We need to develop further the theory of specialization and to understand the flow of information up and down between the different levels of participator dynamics. Information flows up as the conclusions from below become premises above. And information flows down as the participators above move under the guidance of their action kernels and "carry with them" their instantiations as they move. But more theory and examples, in addition to the incremental scheme we already discussed, are clearly required here. We need to understand strategies by which the interpretation kernels of participators can become rcpd's of the appropriate stationary measures. Here perhaps some benefit will derive from study of formal models of natural selection. Perhaps also we should explore more "cooperative" models, models in which the participators don't compete but choose action kernels which maximize the likelihood of true perception for all participators in their dynamics. We must develop more explicitly the epistemological and ontological implications of observer theory. And ..., well there's much left to do. This is just a start.