

On the Evolution of Determinate Information

Conrad Dale Johnson
June 26, 2013

I Where Do Measurement-Contexts Come From?

When Wheeler coined the phrase "It from Bit", his idea was not that the world consists of logical bits, or that physical dynamics works through digital computation. He was thinking about the fact that the most elementary entities we know of, in quantum mechanics (QM), don't have intrinsically definite states and properties. Rather, information about quantum systems is determinate exactly to the extent that it's determined by measurements. He wrote –

'It from bit' symbolizes the idea that... reality arises in the last analysis from the posing of yes-no questions and the registering of equipment-evoked responses; in short, that all things physical are information-theoretic in origin and that this is a participatory universe.[1]

While a digital basis for physics seems attractive to many, few have pursued this strange notion that the world constantly creates all its information through a process of physical question-and-answer.[2] This isn't surprising, since the role of measurement in quantum physics is so mysterious. Despite decades of analysis, we still have no definite idea of what constitutes a measurement, or why determinate information should depend on observer participation. On the other hand, nothing seems clearer or simpler conceptually than binary data and computational logic.

This essay argues for Wheeler's original idea – not just that the physical world is made of information, but that information depends on measurements. Not only those made by observers in a lab, of course. Wheeler imagined the universe itself as an information-defining system, in which we and our measuring devices participate on the same basis as everything else.[3] My goal here is to show that this notion can make sense if we think of measurement-processes in evolutionary terms.

Like Wheeler, I think what QM shows us is that information is not "just given" anywhere – neither in the intrinsic nature of elementary entities, nor in a Platonic realm of pure logic. There's no such thing as information without a context that actually defines it. There get to be definite answers in the world only to the extent that there are meaningful questions – that is, physical situations in which more than one outcome is possible, and where it makes a difference which outcome actually occurs.

When we set up a measurement, we're setting a trap for a certain type of information. Normally we assume that this data is given in advance, and observing it just replicates the data in some other medium. And this assumption works, except in quantum physics. There it turns out that setting a particular kind of trap, asking a specific question, is required for an answer to exist.

So we're supposing here that these information-traps are actually being set everywhere – that they are in fact what the observable universe is made of. We're going to see if we can understand how

these physical “questions” work and where they come from, by imagining a process that evolves through natural selection – a process in which measurement-contexts constantly arise from the results of previous measurements, made in other such contexts.

II Measurements Depend on Prior Measurements

Since the nature of measurement in quantum theory is so problematic, let’s begin with what we know about it in daily life and experimental physics. Observing, after all, is something we all do all the time, and physicists do it better than anyone. They know all about what it takes to measure physical parameters with near-maximal precision. But our deep familiarity with observing and measuring things hasn’t made this easy to conceptualize, for two reasons.

On the one hand, nothing seems simpler for us than seeing, hearing and touching things. Awareness of our environment is so basic to our existence that it feels effortless. Even though we know that many different complex neural processing-systems support this awareness, we still tend to think of observing as if it were simple and direct – as in the ancient notion that the things we see somehow imprint themselves on our minds, like a seal pressed into soft wax. And we conceive measurement too as a simple matter of comparing one thing with another, as when we measure a piece of wood with a ruler, or weigh something on a balance-scale. We rarely consider all the complex physics that’s actually required, at the atomic level, to provide us with things like rulers and balance-scales.

So part of our problem with understanding measurement is that we tend to think it should be simple. The other part is that when we do consider what it takes to measure anything, we run into a forbidding tangle of interdependencies. The key point is that setting up an arrangement to measure any parameter typically involves measurements of several other parameters. In experimental practice, there’s always a lot of predetermined information built into the structure of the measurement-context. To measure a particle’s momentum by observing its path through a magnetic field, for example, we need to know the particle’s mass as well as the field-vector at each point.

Essentially all measurements come down to determining distances in space and/or time, which seem simple and basic. But we only have things we can use as clocks and measuring-rods, in our universe, because of the precisely uniform structure of atomic electron-shells, with their ability to link atoms together into quasi-rigid molecules. And this structure again involves a complex interdependency of many different laws and principles. Yet without this finely-tuned atomic structure there would be no standard distances or frequencies, and no way to set up an arrangement to measure anything.

Because our universe works the way it does, we can take for granted that it provides us with ways of observing all its parameters – that the right kind of background-information will always be available. But in principle, *any kind of measurement is possible only because other kinds of measurements have already been made*. There can be no such thing as a single measurement-event; a context of other measurements is always needed.

There’s no mystery about how physical parameters depend on each other. Every equation in physics gives us a recipe for a measurement: if $F=ma$, that means we can measure a force by determining a relevant mass and acceleration. But it’s not clear how to map out this kind of inter-referential structure, or which parameters should be considered more basic than others, since there are always

several ways to measure any one of them. And anyway, we're not used to thinking of this context-structure as having any fundamental significance.

Traditionally we imagine the world as inherently factual, and the goal of physics is only to show that the facts have a simple and logical underlying structure. If we have two different forces, say electromagnetic and gravitational, we don't ask how they help to define each other, or how each of them contributes something different to the informational environment. Instead we want to show that these forces are ultimately the same, a single unified structure of given fact. So the referential complexity inherent in measurement seems like the opposite of what we're looking for as a foundation for physics.

Yet we have in biology an example of a fundamental information-generating process that's also inherently complex, involving many diverse and interdependent functions. This is the process of self-replication that underlies Darwinian evolution. Even the first self-replicating systems were probably not at all simple – I imagine an accidental collection of organic molecules that happened to be able to serve to catalyze each other's synthesis. And as life evolved into more and more complex forms, reproduction has always been by far the most complex and difficult process any living organism undergoes. What makes this process powerful enough to generate life isn't that it's in any way simple or logical, but that it produces variants that are subject to random selection.

I'm proposing here that the process we call measurement – including the communication of the results as the basis for setting up further measurements – can also evolve through accidental selection. This works very differently from self-replication,[4] as we'll see. And it's a much more difficult process for us to envision. Reproduction is at least conceptually simple, though hard to carry out in the physical world. But we don't yet have anything like an adequate concept of measurement.

Here I'm trying to develop such a concept, one that makes sense of what we know both in QM and in ordinary experience. Its basis is just the fact that measurements need a context of other kinds of measurements. I'll try to show that this already implies the kind of system of interactions that can evolve through natural selection.

III Measurement in the Quantum Domain

We're imagining the world not only as a set of facts, but also as a web of many kinds of interaction-contexts that trap those facts and make them observable, defining each in terms of other relevant facts, defined in other contexts. To make this idea less abstract, let's look at the role of measurement in QM.

In the quantum domain we're at the edge of what's physically observable – given that interaction occurs in discrete, momentary, one-on-one connections involving the minimal quantum of action. However such a web is constructed, at bottom it's very coarsely-woven. Whether or not there's an adequate context to define a certain fact can depend on whether a single interaction takes place.

In contrast, at the macroscopic scale the web is extremely dense and finely-woven; robustly redundant measurement-contexts are everywhere. Everything is constantly observed in many ways by countless other things from various viewpoints. For this system to work, these measurements have to

support and agree with each other; they have to define and maintain certain common structural principles that define (and are defined by) those agreements. And the network so dense, the degree of self-consistency it achieves is so precise, that we never need to pay attention any to it. It's as though the facts were just given and inherently well-defined in each thing by itself.

As for the quantum level – let's consider an example. Photons convey two types of data, in their energy-momentum and spin-polarization. For either type to be measured, the photon has to get absorbed by a charged particle, say an atomic electron, which then changes its own energy-level and spin-state, which also changes the configuration of the atom's electron-shells. That in turn changes the way the atom interacts with other atoms, which can change the shape of a molecule in which the atom participates, and so on.

Each of these interactions is different, and none of them just replicates data from one system to another. We may say that a certain amount of momentum gets transferred from the photon to the electron, but this is an abstraction – physically what happens is that the electron's state changes in a way that's correlated with an opposite change in the state of the particle that emitted the photon. Only on paper does momentum exist as a piece of static data, like a sequence of bits. In the physical world, the mere replication of given data hardly occurs, unless through human intervention. That's why life, based on complex processes that can replicate information, is very rare in the universe.

Now any measurement involves a chain of events like this, where a change in one system makes a difference to another kind of system, in a different context. All these interactions are well-understood theoretically. But none of them yet constitutes a measurement; none by itself defines any information. At this level we don't yet have any physical questions; the contexts in which these events occur can only correlate information between systems, not trap it into a definite answer. So QM represents these systems as superpositions of all their possible states. When the photon gets absorbed, its possibilities get correlated with those of the electron in a combined superposition, and the same occurs with the atom, molecule and so on.

The "measurement problem" in QM comes down to this question – at what point then does a measurement occur? When and how do superpositions "collapse" to give a factual result? Since there seems to be no issue with measuring things in classical physics, this seems to be an issue unique to QM. But I think the reason it remains unresolved, after so many decades, is that we have no idea what a measurement is, even in classical physics. All we know is that they happen – that we do have access to well-defined information about the things around us. But the linearity of the quantum equations tell us no such thing can happen. When things interact, their superpositions never collapse; they only get entangled in larger superpositions.[5]

Objectively, I think this is right. There's no point at which a system physically changes from a superposition to a definite state. But let's pull back from looking at individual events, and gradually widen our viewpoint, taking account of more and more of the web of interaction surrounding these atoms and molecules. At some point, we know, we'll reach a higher level of structure that does provide measurement-contexts, where many kinds of interaction work together to set traps for specific information, that helps specify other information.

Even at that point, though, there's no reason to think that any objective physical collapse occurs. If we could stand outside the universe and "see" the web objectively, not as participants in it,

presumably we'd still find only superpositions of all possible correlated interactions. But from the viewpoint of local systems inside the web, there's a tiny subset of all these possible interactions that happens to be able to define a mutually-supporting set of coherent facts. For systems in this sub-network, interactions that happen to fit its self-defining structure do define and communicate specific information, that contributes to contexts defining other information. Interactions that don't happen to fit this structure aren't physically eliminated; they're just irrelevant to the ongoing process that makes things within this network observable to each other.

Let's compare this scenario to the parallel situation in biology. If we look at what happens between individual molecules in a living cell, nothing we see is distinctively alive; we see random interactions just like those that happen in a test tube. But as we pull back and take in more and more of the scene, we start to see complex cycles of interaction that regenerate complex structures made of many molecules, and so forth – a distinctly biological organization that's much harder to imitate *in vitro*. But it's not until we get to the level of the organism as a whole that we reach something truly alive in itself. Only at this level do we see what makes all of this finely-tuned complexity possible, since it's only the organism as a whole that can successfully reproduce.

Likewise what it means to be determinate or observable in physics can't be defined at the level of individual interactions. Only at the macroscopic level of classical physics can we count on the full functionality of the web of measurement-contexts, giving us an objectively determinate and precisely observable world. In living organisms, there's no specific level of complexity at which molecular processes are suddenly complex enough to be alive. Likewise there's no specific level of complexity at which superpositions collapse into definite states. In both cases we're talking about a primary functionality that's not reducible to lower-level processes.

IV How Measurement Evolves

There are important analogies between biological evolution and the process we're trying to imagine here, but there are also basic differences. Biological organisms split in two, or else build copies of themselves. Measurements don't replicate themselves or their own contexts; instead they select certain information and pass it on to set up contexts for other kinds of measurements. In both cases, though, many different interdependent processes evolve, all ultimately serving just to keep the evolutionary process going.

In biology, most information needed to maintain the reproductive process eventually got encoded in the static linear sequences of DNA molecules. In physics something similar developed, though there was no static storage-mechanism available within quantum systems. Instead, the possibility of successful measurement is maintained by constantly reproducing a complex structure of static meta-information in the interactive environment itself – this structure being what we know as the laws and principles of physics.

No kind of measurement would be possible, of course, if interaction were chaotic or random – as it is in the quantum vacuum, or as it was in the early stages of our universe's history. In the present-day universe, a great diversity of principles support the many different ways of measuring things. We have spacetime geometry, principles of dynamics that apply to all interactions, specific laws for the various forces; there are quantum rules supporting the stability of atoms and nuclei, and a menagerie

of partly-understood regularities of particle physics. All these principles are empirically verified – so they're all regularities defined by the web of measurements. And most of these regularities, possibly even all of them, are needed in some way to make one or another kind of measurement possible.

Now the one thing we know for sure about our universe, from empirical observation, is that at least some things about it are observable. And for any facts to be observable, an appropriate context of other facts is needed to define them, and those facts must also be observable in some context. Facts have to be consistent with facts defined from other viewpoints, so information from one context can be communicated to set up other contexts. And the observed facts themselves have to be able to define all this structure, without depending on any underlying unobservable reality. Whether or not things obey any absolute laws, whether or not there's any absolute spacetime structure behind the scenes, etc. – just doesn't matter. Only what's observable can contribute to the context-structure of empirical information.

So the mere fact that anything is measurable shows that there exists here a self-defining information-structure of a very special kind – one that can define all its own facts, parameters and principles in terms of each other. This sort of system is as remarkable in its own way as a living organism. And the question is whether, like living organisms, we can understand this kind of system as evolving from more primitive kinds of interaction-webs that were also capable of defining their own information.

If we look again at a living cell close-up, at the molecular level, and then gradually back out to take in higher and higher levels of structure, we see all kinds of remnants left over from earlier stages in the evolution of life. Many of the catalytic molecular cycles we see were already going on in the cells of pre-Cambrian life-forms – as evidenced by the fact they're shared in common by all living organisms today. When we get to the level of individual cells (and some of their organelles, like mitochondria), we're looking at entities that were once independently self-replicating organisms themselves, just like single-celled bacteria. They still contain their own DNA and replicate themselves, though only in the context provided by the organism as a whole.

Do we find something like this in the physical environment as well – remnants of more primitive self-defining systems? Take electromagnetism, a unique structure in which moving electric charges create magnetic fields, and changing magnetic fields create electric fields. This structure depends only on the discrete binary $+/-$ charge of particles, not on their real-valued masses or other properties. Further, most of the electromagnetic field-structure can be defined without reference to the metric structure of spacetime. So it seems quite possible that this is a remnant of a more primitive kind of self-defining system, from which our current system evolved.

One of the most remarkable things about the principles and parameters of physics is how many of them there are, and how different are the levels of complexity at which they operate. We've made tremendous progress toward the goal of unification, yet the Standard Model (plus gravity) still gives us a bewildering diversity of elementary particles and forces. The only way such basic differences get explained, currently, is through the *ad hoc* assumption of "spontaneous symmetry-breaking." An evolutionary theory, on the other hand, would look for functional relationships between the different structural elements, giving each a unique role in maintaining the common informational environment. And it would interpret their commonalities – e.g. the inverse-square law of both electromagnetism and gravity – as evidence that more complex structures evolved from more primitive ones.

The basic mechanism of this evolutionary process would be just as Wheeler envisioned – the random selection that occurs in the course of measurements. At any stage in evolution, the environment is able to define and communicate certain kinds of facts, creating contexts that set traps for new facts of the same kinds. It would define certain constraints on random selection, to prohibit some results and make some more probable than others. The constraints would just be the observable regularities of structure that are needed to let this fact-structure redefine itself again and again in new contexts. Interactions that don't happen to "obey" these rules are irrelevant, merely "virtual." Since they don't participate fully in the current information-defining structure, they can't make any definite difference to it.

Such a system would continually generate new facts, and reproduce the conditions under which new facts can be generated, by reproducing the static order of universal principles as well as a common body of historical fact. The *de facto* "purpose" of this process is just to keep itself going – yet the underlying randomness could always result in new kinds of situations. So it's conceivable that new kinds of structure could emerge to define and integrate new kinds of information.

Let's imagine a certain point in this process, where the current informational environment traps certain information. Part of this information is required by current laws; part of it is just randomly selected and passed on, as historical fact. Suppose it becomes possible to define some new kind of regularity that would further constrain this randomness, while preserving consistency with the current order. The new rule would still have to be able to define itself. That is, the sub-network of interactions that happened to "obey" this new rule would need to define new kinds of fact and create new contexts to trap these facts. And again, for systems that happen to participate in this more restrictive sub-network, most interaction in the previous environment would now be only "virtual". They would still be there, and would still play a role in defining the earlier layers of the interaction-web, but they wouldn't make any directly observable contribution in the new environment.

Overall, then, we're imagining a repetitive process in which each measurement-context defines a set of possible outcomes, based on inputs from other measurements. To the extent that a selection from this set can determine new information that contributes to set-up of another such context, then a measurement has taken place – so long as the result of that measurement also becomes meaningful in the set-up of another context, and so on. Essentially we're defining "measurement" recursively, in terms of itself – a measurement takes place insofar as some other measurement results from it.

But this is exactly how "reproduction" is defined in biology. Successful reproduction doesn't mean that an organism creates an exact copy of itself (since errors in replication are also needed, to produce variation subject to selection). The evolutionary criterion is that an organism reproduces successfully if and only if at least one of its offspring also succeeds in reproducing itself. This is more than a theoretical definition, since every ancestor of every organism on earth has actually met this requirement.

In physics, each measurement narrows down the range of possibilities that are passed on as a basis for future measurements, creating more specific situations in which more determinate information can be defined. At the bottom of all this is a quantum vacuum, where anything is possible, and nothing is determinable. But long-term, the observable universe will gradually become more determinate, as more ways of determining things arise. Eventually the possibilities for what can happen next, in each situation, get narrowed down as much as they can be, given the indeterminacy of the foundations.

So in our universe today, at the macroscopic level, everything that can be observed involves a mutual selection among trillions of events, all satisfying dozens of layers of structural constraints, making a world that looks as much as possible like a classically deterministic cosmos.

V What Physical Theory Can Explain

Whether this roughly-sketched fantasy really makes sense of course remains to be seen. But I think this approach has the potential to explain much about the physical world that can't be explained by any theory in which deterministic laws and principles underlie the structure of observed phenomena. While classical physics more than proves the value of deterministic principles in physics, I don't think there's any way they can account for the fact that there are observable phenomena in the world.

Beyond that, I've already mentioned the problem of accounting for the great diversity of physical structure in any "unified" theory, since the very notion of unification tends to discount differences as insignificant or secondary. It's not clear how this kind of theory could ever give meaning to the fine-tuning of physical parameters, in our universe, or explain the functionality of atomic structure. On the other hand, an evolutionary theory would be expected to account for the many ways in which atoms support the possibility of measurement.

I'd like to conclude by pointing out another issue that has to be ignored, in deterministic theories. What I have in mind is that nothing that happens in the physical world can be computed as precisely as it can be measured. To take a well-known example, there's no equation for the dynamics of just three point-like masses, interacting gravitationally. Even such a drastically simplified situation can only be approximated mathematically, even in Newtonian physics – as for General Relativity, I believe it has difficulty defining the dynamics of only two mass-points. To me at least, it seems clear that the physical world is incomparably more powerful as an information-processing system than any kind of mathematics or computational logic.

Nonetheless, physicists have taken it for granted for centuries that the world is somehow able to operate on deterministic principles, even though it everywhere provides instant and exact solutions to problems that are mathematically intractable. This used to be attributed to the infinite power of the Mind of God; today it's essentially just a scientifically acceptable form of magic. Of course, physicists don't necessarily imagine the world as computing itself in real time. More often they think of it as a complex mathematical pattern laid out over all of spacetime at once. In this pattern nothing has to calculate itself; it all just is.

Yet if the world were built on mathematical foundations, why should it be structured in a way that makes computation and mathematical analysis so hard? It's not difficult to invent equations or algorithms that give rise to extremely complex systems. But in our universe, equations govern one-on-one interactions between individual entities, and what happens in any situation depends on a non-linear combination of many different kinds of interactions going on at once. So whether or not we believe in mathematical miracles, this universe hardly looks like a system based on deterministic computation.

At the macroscopic scale, though, things do "obey" deterministic laws, to the highest precision we can measure. So if the equations aren't running this show, then why are they there? What role do

laws of physics play? And if at a deeper level things aren't fully determinate, if they turn out to obey laws only on average, then why does the world we observe end up looking so precisely factual and deterministic?

Such questions can hardly be asked in the traditional theoretical context. But I've tried to show that they could have sensible answers in an evolutionary theory.

Notes and References

[1] John A. Wheeler (1990), "Information, Physics, Quantum: The Search for Links", in *Complexity, Entropy, and the Physics of Information* edited by John A. Wheeler and Wojciech H. Zurek (Addison-Wesley Press)

[2] Carlo Rovelli is an important exception, with his projected derivation of QM from the structure of yes/no questions posed in measurements (though he does not attempt to analyze the structure of measurement-contexts). (1996-7, 2008), "Relational Quantum Mechanics", <http://arxiv.org/pdf/quant-ph/9609002v2.pdf>.

[3] John A. Wheeler (1983): "Law Without Law" pp. 182–213, in *Quantum Theory and Measurement* edited by Wheeler and Zurek (Princeton University Press)

[4] An evolutionary theory based on the self-replication of universes in black holes was proposed by Lee Smolin (1997), *The Life of the Cosmos* (Oxford University Press). Apart from the speculative nature of the premise, it's not clear how much about our universe could be explained by a process of selection with the special goal of reproducing black holes.

Closer to my approach is that of Wojciech Zurek, who has extended his theory of decoherence to suggest that information about quantum systems gets selectively replicated in the interactive environment – W. H. Zurek (2003), "Quantum Darwinism and Envariance", <http://arxiv.org/pdf/quant-ph/0308163v1.pdf>, (2009) "Quantum Darwinism", <http://arxiv.org/pdf/0903.5082.pdf>. This process may play an important role, but Zurek bypasses the key issue of the question-structure of measurement-contexts, and merely assumes that quantum systems "imprint" their information on the environment through entanglement.

[5] As Zurek generally points out, his decoherence theory doesn't address the question of "collapse". It shows that quantum superpositions subject to interaction with a chaotic environment may become less entangled with other systems and more like a set of possibilities for classical measurements. And he suggests that Quantum Darwinism goes further toward resolving the measurement problem. But many physicists feel this problem has already been solved by one interpretation or another. Its resolution will remain doubtful so long as we have no clear idea of what measurement means.