The Wysiwyg Universe

by

Simon Free

http://wysiwyg.stanford.edu/

Document Version 1.1, April 2002

Introduction

WYS.I.WYG /'wi-zE-"wig, adjective: ***W**hat **Y**ou**S**ee**I**s**W**hat**Y**ou**G**et*

Much has been written in the fields of philosophy, psychology and physics regarding the role of the observer in understanding the nature of reality. Physicists are zeroing in on what they hope to be a "Theory of Everything", a set of mathematical equations which describe all that there is to know about the basic building blocks of the Universe. At the same time, psychologists and neurobiologists are closer than ever to a complete map of how the brain allows us to see, think, feel, and even dream.

Yet despite these advances in our scientific understanding, there remain some unsettling gaps between what we *say* about the world and what we actually *experience*. For example, no description of water, no matter how poetically vivid or mathematically correct, will quench your thirst. Similarly, the abstract beauty of Maxwell's electromagnetic field equations seems pale by comparison to the myriad hues of a sunset. Even the most elegant theories of the origins of the Universe seem somehow devoid of substance and leave us uncomfortably unsatisfied. In the words of physicist Stephen Hawking: "What is it that breathes fire into the equations and makes a universe for them to describe?"

It will be argued here that these explanatory gaps arise from a simple ontological error commonly committed when discussing what we take to be reality. At the root of this confusion is the assumption of a world *beyond* what we experience. This assumption generally begins as a *model* of a set of experiences which is formulated either verbally or in the language of mathematics. If a model is successful, it allows us to make predictions about observations within a certain domain of experience. However, the more successful a model becomes, the greater is the tendency to speak of the model as though it describes the way the world really *is*. But this confuses the map for the territory. Few would argue that a map of the Rocky Mountains actually *is* the mountain range, no matter how detailed or elaborate the map. Similarly, no neurological description of the taste buds in your tongue and their connections to your brain will enable you to experience a flavor you've never tasted before.

It will be shown below that by adopting two simple postulates regarding the relationship between experience and reality, we can eliminate the confusion which tends to arise between our models of reality and reality itself. But before we take the plunge, let us examine two particular cases where the apparent gap between model and experience seems just too great to cross.

The Quandary about Qualia

There is a long standing debate in philosophy and psychology on the nature of *qualia* which are defined as those elementary forms of experience such as color, taste, sound etc., with which each of us is intimately familiar. Despite great strides in the understanding of brain function over the past century, little progress has been made toward the elucidation of how the firing of neurons in the brain can give rise to these subjective qualitative states.

For example, a pin prick to the finger causes nerve endings beneath the skin to send impulses to pain centers in the spinal cord and brain. These centers are connected to reflex motor actions which cause the withdrawal of the finger away from the pin. However, we not only react to the pin prick in this neuromechanical fashion; we also *feel* pain. How this subjective feeling of pain arises from the firing of neurons in the brain is a mystery. To put it another way, we have no idea why neurons firing in one part of the brain give rise to pain, while neurons firing in a different part of the brain give rise to sounds, and still others produce the experience of colors.

It appears that qualia are made of a different "stuff" than the cells in the brain or the molecules which make up the cells. In fact, many have argued that there can never be a physical explanation of how we come to have experience. Others deny that we really have experiences at all! This apparent ontological "gap" between the physical world and the world of experience is what philosophers call the *hard problem* of consciousness or the *mind-body problem*.

Quantum Quirkiness

The second domain in which experience does not seem to follow directly from the physical world is the realm of the very small, such as the level of atomic and subatomic processes. At the more macroscopic levels of experience, the process of measuring the property of an object, such as the position of a rock, generally yields a single value for any given state of the object. However, in the world of electrons and other subatomic entities, measuring the property of an object can yield more than one result, even though the initial state of the object is the same from one measurement to the next. For example, an electron has a property called "spin" which can take on only one of two values, "up" or "down". (Think of spin as a little arrow attached to the electron, where the arrow can point either up or down but not in between.) If a suitable measurement of spin is made on a large number of identically prepared electrons, half the time we get the value "up" and the other half the time we get the value "down". For any given electron, there seems to be no way to predict which value will be measured.

The branch of physics called *quantum mechanics* deals with this uncertainty by describing objects such as electrons using a formula called the *wave function*. The wave function of a quantum mechanical system prescribes both the set of possible values one can obtain during a measurement of a given property, as well as the probabilities of obtaining each result. For example, when the wave function is computed for an electron passing through a magnetic field, the formula admits just two possibilities: spin up and spin down. It also prescribes probability values of 0.5 for obtaining the value "up" and 0.5 for obtaining the value "down". So when measuring the spin of an electron, quantum mechanics can only tell us that the probability of measuring spin "up" is 0.5 and the probability of measuring spin "down" is 0.5. We will definitely obtain one or the other value, but *which* value we obtain is apparently determined by the toss of a coin.

Needless to say, many physicists are not happy with this apparent irreducible randomness in the subatomic realm. Consequently, a number of theories and conjectures have arisen over the years in an attempt to understand what actually *determines* the value of a quantum measurement. In the case of electron spin for example, some physicists speculate that, prior to a measurement, the wave function of the electron actually describes a state of being *both* spin up *and* and spin down--a 50-50 superposition of opposite spin states. Since we know that we only ever observe *either* spin up *or* spin down, this superpositional state is not a physical one. To turn this non-physical state of superposition into a real value of the spin, the process of measuring the spin is said to "collapse" the wave function into one of its particular values. This seems to imply that the process of observation actually helps to create the reality which we ultimately observe.

Imagine how strange this would seem at a macroscopic level. Suppose that a rose can appear either red or yellow. On half of the occasions we look at the rose it appears red, and the other half of the time it appears yellow. As much as we study the other properties of the rose, there seems to be no way of telling before hand which color we will see--only that there is a 50-50 chance it will be red or yellow. So what color does the rose have when we are *not* looking at it? Is it *both* red and yellow? Is it orange? Does it even have a color at all? Or is the question even meaningful? And what actually happens when we do look at the rose? We only ever see one color at a time so how does the unobserved and ambiguous color state of the rose (whatever this means) actually become the color we observe? Does the very process of our looking at the rose actually *cause* one color to be selected over the other?

The Wysiwyg Hypothesis

To clarify the respective roles of experience (qualia), models (science) and reality, we will begin with two simple postulates.

First Postulate: **Experience is the only reality for an observer.**

Second Postulate: **There is no reality other than experience.**

In general, postulates are assumed to be irreducible assumptions of a theory. However, one normally chooses these assumptions because they seem reasonable or necessary, at least on the surface. In this case, the two postulates may or may not seem self-evident. We therefore begin with arguments for their reasonableness before proceeding to the issues of qualia, quantum mechanics and a few other puzzling issues in science and mathematics.

Arguments for the First Postulate

Since the issue is experience, and you are the only one who can examine your own experience, ask yourself the question: "What do I know of reality besides what I experience?" Suppose you answer: "I know that I have a brain in my head which allows me to have experience." Of course, you don't directly experience your brain--you can't even see it. So the "knowledge" that you have a brain in your head is an assumption based on what you have learned from the cultural database of knowledge that our society has collected to this point in time. Do you think this knowledge in any way affects the way you experience, say, color? Most non-scientists are probably completely unaware of the detailed neurobiology behind color vision, yet they see colors just fine. So the model we have of how brain activity enables us to perceive color is not the same as the experience of color itself. Hopefully this much is clear.

Next you might argue: "While our model of brain function does not determine how I see colors, surely the function of the *physical* brain does." This is where the "map for the territory" error creeps in. To demonstrate how this comes about, imagine that you could peer inside your own head and actually see your brain in action. (In fact, we can do something just like this using the technology of functional MRI.) By viewing the workings of your own brain, what additional knowledge do you now have that relates to your experience of color? Let us focus on the visual subsystems which neuroscientists know to be involved in color vision. At first, you may not identify the various neural components in front of you. They appear to you as various shapes and sizes--in other words, as visual *experiences*. If you have read about the brain mechanisms behind color vision or you are a researcher in the field, you will have associations between these visual experiences and the names and functions of each component. In the end you could put together a verbal explanation (a model) of how you see color: "Light passes through the lens of my eye stimulating the rods and cones in the retina. The rods and cones activate my optic nerve which sends impulses to my visual cortex. Here I can see specialized neurons that have evolved to detect the ratios of wavelengths in the original light. For example, whenever I experience the color red, I see that cortical cells in region A are active. Whenever I experience the color blue, I see that cells in region B are active."

In the end you have correlated one experience with another: your subjective experience of color on the one hand, and your subjective experience of your brain's activity (as observed in our imaginary scenario) on the other hand. This is actually very useful knowledge. For example, it allows you to predict that if you were to remove or otherwise damage the cortical cells in region A, you would probably have difficulty experiencing the color red. However, what you have *not* done is observe how a physical process (brain activity) *gives rise* to your experience. What we call the "physical brain" is actually a collection of experiences just like the experiences we are trying to explain. It makes no difference how detailed a description or model we make of the brain--right down to the quantum mechanical fluctuations in the spacetime continuum if you like: each component of the description will arise from its own collection of experiences and, as with the map of the mountain range, even the most articulated model of the brain is not the brain itself. The distinction between the physical world and the experiential world is therefore seen to be artificial; hence our first postulate: *Experience is the only reality for an observer.*

Arguments for the Second Postulate

It is one thing to say that experience is the only reality for a *given* observer (yourself for example). The second postulate takes this a step further: *There is no reality other than experience*. For thousands of years, thinkers in all cultures have wondered what is the "true nature" of reality behind our experience. Tribal societies imagine animated spirits behind objects and events. Religious leaders believe in the existence of an overarching theological plan. Philosophers have argued for a world of ideal platonic forms behind the objects we perceive with our senses. And scientists postulate a physical or mathematical foundation to reality. But for thousands of years, no one has come up with an answer that satisfies everyone. Perhaps there is a very simple reason for this: *Experience is all there is*.

It should be clear from the arguments for the first postulate that ideas about a "reality behind experience" are really just models in disguise. Such ideas rely on constructions using verbal language or mathematics which postulate unobservable entities that give rise to our experience. But this is just what a model is, and the hypothesized unobservable entities are just place holders for correlations among experiences. Furthermore, models are changing all the time as observers discover new phenomena to be explained. For example, the experience of gravity has been conceptualized and modelled very differently from the time of Aristotle's physics, through Newton's dynamics, and on to Einstein's theory of General Relativity. But rocks fall to the ground all the same. In any case, a model cannot be *identified* with something outside our experience. There is no way out of the experiential loop. Hence our second postulate: *There is no reality other than experience.*

Applications of the Wysiwyg Hypothesis

The arguments given above are meant to show that the postulates are not inherently unreasonable. Their usefulness will become apparent when they are employed to solve a number of long standing puzzles in philosophy and science. We now proceed to the application of the postulates to the problem of qualia, the collapse of the wave function in quantum mechanics, and to a number of other troublesome issues, including the uncanny fit between mathematics and experience.

Qualia, Qualia, All the Way Down

Let us return to the example given above wherein you were able to view the activity of your own brain. Let E be an experience, such as the experience of the color "red". Suppose that whenever you experience E, you observe activity in a collection of cells, C, in your brain. Note that C is itself an experience: the experience you have when observing your brain cells and their activity. We denote by the notation, E ~ C, the correlation you have established between E and C. By the first postulate, E ~ C is the only reality you know: any additional facts you might acquire about the collection of cells, C, would simply add to the collection of experiences which we are calling C. By the second postulate, there is no other reality besides E ~ C. In other words, there is no hidden "physical reality" behind C nor is there some "ethereal reality" behind E. Since E ~ C is simply a correlation between like entities (experiences) there is no ontological gap which needs to be bridged between the "physical" world of neurons and molecules and the "experiential" world of mind or consciousness. In the end, there is only experience, all the way down.

Exorcising Quantum Ghosts

Recall from the introduction that measurements made on identically prepared quantum mechanical systems can lead to different results. Which result arises during the measurement appears to be random, but each result can be assigned a probability of occurrence by computing the wave function of the system before the time of the measurement. The philosophical issue which concerns physicists is that it seems to be impossible to assign a definite value of the measured property to the object under study before someone actually observes it.

Let us apply our postulates to the specific case of electron spin. What do we experience during the experiment? We experience a series of spin values, one for each electron which passes through our apparatus. If we count up the number of spin up electrons versus the number of spin down electrons, we find that they appear in roughly equal numbers. That's it. That is all we experience. According to the first postulate, reality for the experimenter is simply this: when you measure the spin of a batch of identically prepared electrons, about half of them have spin up and the other half have spin down. At the present time, there does not seem to be a way to determine which value (up or down) one will obtain for a given electron at the time of measurement.

So let E be the experience of measuring the electron spins. And let C be the collection of recordings of those spins--for example, writing a tally of "ups" versus "downs" on a piece of paper. Note that C is also an experience--the experience of recording the spins. After all, if we did not record our results (even in memory), there would be no way to do the analysis. We can then establish a correlation, E ~ C between the two experiences which says that the measurement of electron spins tends to generate a record of equal numbers of spin ups and spin downs.

Physicists can take this correlation one step further. Quantum mechanics prescribes a precise method for actually computing the wave function of a system from its energy. The energy of the system is in turn computable from other well known properties of the system. We can therefore compute the wave function which represents the spin experiment ahead of time and thereby *predict* what the results of the measurements will be. The wave function therefore provides a *model* for how the system will behave. Since the wave function predicts the probability of obtaining spin up or spin down in the experiment, we can measure the spin of a large number of electrons and see if the relative numbers of spin ups and spin downs reflect the predicted probabilities (they do). In essence, the wave function takes the place of the correlation symbol, ~, between our experience of measuring spins (E) and our experience of recording the results (C). The only difference is that we now have a formula which enables us to predict the results ahead of time.

What happens next is where all the confusion arises. Rather than recognizing the wave function as a succinct representation of the correlation between measurements and results, some physicists would like to endow the wave function with a reality beyond its role as a mathematical formula. The argument goes something like this: since the wave function appears to represent the state of the system before a measurement is made, then the system must actually *be* in that state before the measurement is made. Besides confusing the map for the territory, this is an especially puzzling view to take of the wave function, since the very role of the wave function is to represent the system *before* it has been observed. The ontological status of the unobservable wave function is therefore, in principle, untestable.

If we apply our second postulate, *there is no reality other than experience*, to the problem of quantum mechanical measurement, we run into no trouble at all. In the case of the electron spin experiment, our experience includes the measurement of the spin, the recording of the results, and the wave function (formula) which correlates one with the other. We recognize the wave function for what it is, a recipe for generating probability predictions. There is no ghostly wave function in the experiment itself and therefore no need to collapse it when making a measurement. The observer therefore plays the same role as observers always do--they have experiences.

Mathematics and Experience

(A more rigorous treatment of this section for the mathematically minded will appear in a forthcoming Appendix. The version presented here is intended for a more general audience.)

One of the more interesting puzzles in the history of physics and mathematics is why mathematics applies to the world at all. Einstein once said, "The most incomprehensible thing about the world is that it is comprehensible." Other physicists and mathematicians have asked: Why should a symbolic system as abstract as mathematics be so good at describing physical reality?

Our two postulates about experience and reality allow us to solve this mystery. First of all, it helps to know that all branches of mathematics are actually variations on a single theme. This common starting point is the definition of a *set* of objects, together with a *structure* defined among the elements of the set. A structure is defined simply as a collection of pairings between elements of the set. In fact, a structure on a set is another set, a set where the elements are certain pairings of elements from the original set.

For example, consider the set whose elements are the words "apple", "banana", "red", "yellow". We write this set using the notation {apple, banana, red, yellow}. An example of a structure on this set would be the pairings {{apple,red},{banana,yellow}}. Notice how the pairings of elements from the original set become new elements in the structure set. You can immediately begin to see how such structures can define a model. For example, the structure given above defines pairings between "apple" and "red", and "banana" and "yellow". If we were interested in the colors of various fruits, this model of fruit colors would predict the right colors for apples and bananas. A different model, say {{apple,yellow},{banana,red}} would make the wrong predictions.

If this sounds familiar, it should. As we have seen, a model of a domain of experience is a correlation, E ~ C, between the set of experiences to be explained (E) and the set of recordings (C) made during these experiences. By our first postulate, recordings of results are just other experiences (like writing down numbers on a piece of paper). We can therefore combine the two sets of experiences, E and C, into a single set of experiences which we could call U. The combined set U would look something like this: {e_1, e_2, e_3 ..., c_1, c_2, c_3 ...} where e_1, e_2, etc., are experiences in the set E, and c_1, c_2, ... etc., are recordings (also experiences) in the set C. By our second postulate, these experiences are all we have. So the best we can do to model our experiences is to define a collection of associations between elements of this set. Such a model would look something like this: {{e_1,c_1}, {e_2,c_2}, {e_3,c_3}, ...} which implies that experience e_1 predicts or correlates with recording c_1 (and vice versa), e_2 predicts or correlates with c_2, and so on. But such a collection of parings is none other than a mathematical structure as defined above. We have therefore shown that *any* model of experience is by definition a mathematical structure. Conversely, any mathematical structure can be viewed as a model, although the vast majority of such structures will be very *poor* models when it comes to predicting results.

Physics 101

It might seem at first glance that this definition of a model bears little resemblance to the formulas commonly found in physics, such as Newton's law of motion, F=ma, or the wave function in quantum mechanics. However, we can easily show that these formulaic models are just special cases of the more general mathematical structures defined above.

Take the example of objects falling to the earth. The great Italian physicist Galileo Galilei recorded the time it took a ball to roll down a long inclined plane by using his pulse as watch (and later, a water clock). We see immediately that Galileo's experiment involved correlating one experience with another; namely, the experience of viewing the rolling ball at different points along the inclined plane, and the experience of feeling his pulse. Let E be the set of experienced pulse beats and let C be the set of positional recordings of the rolling ball. The set E can be written as {p_1, p_2, p_3, ..., p_N} where p_k corresponds to k^{th} pulse beat since the release of the ball at the top of the inclined plane. The set C can be written as {d_1, d_2, d_3, ..., d_N} where d_k is the distance the ball has rolled down the plane by the k^{th} pulse beat. The simplest model of the results would be the structure {{p_1,d_1}, {p_2,d_2}, ..., {p_N,d_N}} which says that the object will be found at distance d_1 at pulse p_1, at d_2 at pulse p_2, etc.

Of course, this model would only work for one angle of the inclined plane. Fortunately, both pulse beats and distances can be given *numerical* representations. If we assume that the pulse beats are equally spaced apart, then we can represent the pulse beats with integers and write this new set as T = {1, 2, 3, ..., N} and call it "time". Next, regardless of the units used to measure distance, we can write their values relative to the distance traveled during the first pulse beat. What Galileo discovered was that, independent of the angle of the plane, the distance measurements took on the relative values D = {1, 4, 9, ..., N^2}. Our new model therefore relates time and distance expressed numerically and looks like this: {{1,1}, {2,4}, {3,9}, ... {N,N^2}}. Even with Galileo's crude timing measurements, it was clear that these numbers indicated that the distance traveled by the ball (the second number in each pair) was related to the *square* of the time since the ball was released (the first number in each pair). What's more, this squaring relationship did not depend on the angle of the inclined plane or the mass of the ball used to do the experiment. These results implied two important properties of falling objects: that objects do not fall to the ground with a constant speed but are *accelerating*; and that this acceleration does not depend on the mass of the object. This latter result actually invalidated the model of falling objects proposed by Aristotle who thought that heavier objects fall faster than lighter ones. Only by slowing down the rate of motion by using the inclined plane was Galileo able to collect the experiences required to disprove Aristotle's claim.

Given the numerical form of Galileo's model, we can write the results in more succinct form as $D = kT^2$ where k is a constant of proportionality related to the angle of the inclined plane. It is this form of the model which we generally recognize as a mathematical or physical "law". But note that Galileo himself actually started with a formulation of his results exactly like the set notation used above. It was the *numerical* representation of the experiment that made the results amenable to a simple formulaic analysis. And this numerical representation was possible only because the experiences under study were of such a nature that numbers make a good representation of the results. How else would you choose to represent distance? Or what could be more natural than counting your pulse to measure time? The fact that certain aspects of our experience such as distance, time, mass, force, temperature, brightness, etc., can be easily represented numerically is the reason we can turn our observations of correlations among such properties into numerical formulas.

Genetic Gymnastics

Observations in other domains of experience, such as how the genes in biological cells direct the construction of proteins, do not lend themselves so naturally to numerical representation. Nonetheless, we can still generate a model of protein synthesis in terms of the genetic code which allows us to fully understand how genetic variation can produce changes in the structure and function of biological organisms. Furthermore, this model qualifies as an example of a mathematical structure using our more general definition.

Biologists have known since the 1950's that all proteins are made up of combinations of 20 different amino acids. They also know that genes are composed of combinations of just 4 nucleotides. Somehow genes direct the construction of proteins but how? The set of all possible amino acid sequences and the set of all possible nucleotide sequences are both potentially *very* large. For example, even smaller proteins consisting of only 100 amino acids could come in potentially 20^{100} variations. This number is so huge (larger than the total number atoms in the universe) that it defies imagination. However, it can still be treated as a set. Similarly, there are approximately 100,000 genes in human cells which code for proteins and each of these genes contains anywhere from 1,000-2,000 nucleotides. So a gene can come in roughly 4^{1000} variations which is again an unthinkably large number. The challenge before the biologists was to figure out which sequences of nucleotides in genes coded for which sequences of amino acids in proteins. And you thought Galileo had it tough!

The story of how this coding problem was finally solved is one of the most intriguing pieces of detective work in the history of science. Theorists from physics, information processing, mathematics and biology all had a crack at it. The key was the realization that a genetic "alphabet" of four letters (nucleotides) could easily code for 20 amino acids if these letters formed at least three letter words, one word for each amino acid. For example, with four letters to choose from, one letter words give us only 4 possibilities. Two letter words yield 4x4 = 16 possibilities--still not enough to code for 20 amino acids. Three letter words come in 4x4x4 = 64 variations which would be more than enough to account for the 20 amino acids. Much theoretical speculation was aimed at understanding how these 64 genetic codes could be reduced to just 20 unique amino acid codes.

In the end, the solution had to be determined empirically, that is, by actually looking inside the cells to see how segments of a gene were transcribed into pieces of a protein--just as Galileo had to record the positions of the rolling ball at each beat of his pulse. The result was a mapping between three letter genetic sequences and amino acids which bore no resemblance to the theoretical models. As we might expect using 20-20 hindsight, Nature just threw things together one code at a time over the course of evolution as necessity demanded invention.

The net result is a built-in redundancy whereby multiple genetic triplets code for the same amino acid. In this way, genes have a tolerance for genetic errors which the theorists hadn't considered.

The empirical data laid out a mapping between three letter genetic codes and amino acids. Our mathematical problem can therefore be recast in terms of the set of amino acids, A, and the set of three letter genetic codes, C. If $A = \{a_1, a_2, a_3, ..., a_{20}\}$ and $C = \{c_1, c_2, c_3, ..., c_{64}\}$ where there are 20 elements in the set A and 64 elements in the set C, then the model of the empirical results takes the form: $\{\{a_1,c_1\}, \{a_1, c_{27}\}, \{a_2, c_9\}, \{a_2, c_{58}\}, ..., \{a_{20}, c_{16}\}, \{a_{20}, c_{32}\}\}$ where we have indicated that a given amino acid may be coded by more than one genetic triplet.

This is our mathematical model of the genetic code. We have defined a structure of pairings between genetic triplets and amino acids. One might wonder if we could take this a step further, as Galileo did for time and distance, and come up with a formula that would allow us to *compute* the value of an amino acid from the genetic code(s) that stand for it. This would require that we had some meaningful way of assigning numerical values to both the set of amino acids and the set of genetic codes. Of course, we could arbitrarily label the amino acids {1, 2, 3, ..., 20} and the codes {1, 2, 3, ..., 64}, but these values do not bear any relationship to the properties of amino acids or nucleotides. In other words, what would stop us from labeling the same amino acids in reverse order; i.e., {20, 19, 18, ..., 1}?

So while it does not make sense to look for a formulaic version of our genetic model, it does make sense to figure out how genetic triplets are actually translated into amino acids using the processes available to a cell. The series of steps involved in this process is now well understood and is called *protein synthesis*. Segments of genes containing many coding triplets are lined up with another type of genetic molecule called messenger RNA. Messenger RNA then takes its mirror copy of the genetic triplets into the cytoplasm of the cell where some very mechanical processes in structures called ribosomes line up free floating amino acids with the appropriate genetic triplet along the length of the RNA molecule. When the assembly process is complete, the string of amino acids peels away as a freshly made protein molecule.

Summary

Hopefully it is now clearer why the world appears to obey mathematical laws. For when we say that the world "obeys" mathematical laws, we simply mean that there are correlations (pairings) among elements of our experience which enable us to make predictions from one experience to another. For some domains such as physics, the correlations between experiences can actually be given functional form, like the wave function in quantum mechanics or F=ma in Newtonian dynamics. These formulaic models are possible because the observations of interest can be given a numerical representation. Since formulas and equations are what most of us think of as mathematics, we marvel at how such abstract looking constructs seem to fit so nicely with our experiences.

But from what we have shown above, even the most complex functional equations can be defined as examples of sets with structure. Conversely, models found in other sciences, such as the genetic code in biology, are no less mathematical in their essence than those found in physics. In these cases, the relevant objects under study do not lend themselves to numerical representation, but the structural mapping of one set of observations to another still forms a mathematical model.

Unfortunately, the undeniable success of numerically precise models in physics has generated an implication that these models are somehow more "fundamental" with regard to an ultimate understanding of reality than the models used in, say, biology or psychology. However, it is important to realize that even the most successful grand unified theory in physics will be essentially mute when it comes to many of the more macroscopic aspects of our experience. For example, the structure of the Empire State building is entirely compatible with the laws of physics, but those same laws cannot predict the existence of the Empire State building. In the same way, the model of the genetic code described above cannot be reduced to the laws of physics: the genetic code is a result of billions of years of haphazard biological evolution, and while these evolutionary processes are consistent with the laws of physics, they are not necessitated by them.

The Wysiwyg hypothesis allows for "fundamental laws" at all levels of experience. There is no *a priori* reason to expect that one domain will be in some sense more fundamental than all the rest--after all, all such domains are ontologically equivalent insofar as they are all elements of our experience. For example, attempts to force an explanation of qualia in terms of Planck-scale quantum fluctuations in the spacetime continuum seem somewhat misguided. From our experience of brain function, the relevant level of description is that of neurons, their interconnections and their chemical messengers. And it is at this level that we are most likely to find the relevant correlations between qualia and brain mechanisms.

Experimental Support for Wysiwyg

The Wysiwyg postulates seem to clarify many otherwise confusing issues in philosophy and science. However, is there some empirical set of observations which give the postulates a firmer ground on which to stand? Surprisingly, the answer is a qualified "yes".

The Wysiwyg postulates claim that there is no "objective" reality other than experience--nothing behind the scenes pulling the strings or holding everything together. A similar implication can be drawn from the fallout of a paradox first described in 1935 by Einstein and two co-authors, Boris Podolsky and Nathan Rosen. It is therefore referred to as the Einstein-Podolsky-Rosen (EPR) paradox.

The EPR paradox was meant to show that the quantum mechanical description of the physical world was not complete. The apparently irreducible randomness at the quantum level of matter greatly annoyed Einstein who felt that "God would not play dice with the Universe." He therefore speculated that there must be "hidden variables" lurking behind quantum measurements which, although they had not yet been observed, would allow us to explain the random behavior of the results once the variables had been identified.

The basic setup of the EPR paradox is fairly easy to describe in general terms. Suppose we have two electrons A and B which are initially prepared such that if one of them has spin "up" then the other *must* have spin "down" and vice versa. Such spins are said to be "correlated" and can actually be produced in the lab. Now imagine that, before we actually measure the spin of either electron, we allow them to fly apart so that they are now separated by a very large distance. If we then measure the spin of electron A and find that it is "up", quantum mechanics must predict that a measurement of the spin of electron B will yield a value of "down". The question becomes: how does electron B "know" that our measurement on electron A resulted in a value of spin up? We can move the two electrons as far apart as we like so that even a light signal could not communicate the result of the measurement on A fast enough for B to be informed of the result. Einstein called this hypothetical result "spooky action at a distance" and concluded that the only way quantum mechanics could handle such a situation was to admit that additional hidden variables were at work which determined the spins of both electrons before they were separated. In other words, what you see is only part of what you get.

Due to the technical sophistication required to actually carry out the correlated spin experiment, a resolution of the EPR paradox had to wait another 30 years. In the meantime, an Irish mathematical physicist named John S. Bell took a keen interest in the EPR argument. Bell wondered, what if we accept the EPR hypothesis and assume that there *are* hidden variables behind the spin experiment? Bell was then able to prove a theorem which showed that for *any* such hidden variable theory, he could derive an equation (actually, an inequality) which put severe constraints on the types of experimental outcomes such a theory could predict. The inequality derived from *Bell's Theorem* for the case of correlated electron spins clearly showed that the results predicted by quantum mechanics could *not* be accounted for by any hidden variable theory. What's more, Bell's work led to the design of actual experiments which could test Bell's inequality in the lab and hence determine whether or not quantum mechanics was correct, despite the EPR paradox.

A series of experiments followed, the most famous being that of Alan Aspect and co-workers. These experiments are very similar in principle to the original EPR thought experiment. The results are fairly unequivocal: quantum mechanics is right and Einstein was wrong.

Bell's Theorem and the subsequent experimental vindication of quantum mechanics by Aspect and others has been widely regarded as "the most profound discovery in science." To appreciate the full impact of Bell's work, the EPR paradox can be recast so that it pits quantum mechanics against three fundamental assumptions about reality:

1. Deductive logic is valid.
2. No signal can travel faster than light.
3. There is an objective world independent of our experience.

Since the quantum mechanical predictions for the results of the Aspect experiment were confirmed, one or more of these assumptions must be wrong. There is little reason to date to reject either deductive logic or the wealth of experimental evidence behind the maximum speed limit imposed by light. This leaves assumption 3 which, if rejected, becomes the equivalent of our second Wysiwyg postulate. In this way, one can see that the Wysiwyg hypothesis is not so far fetched as one might seem--in fact, it appears to find direct experimental confirmation in the physicists' lab.

Wysiwyg versus Mysticism

The Wysiwyg hypothesis puts all the emphasis on experience. Another pattern of thinking which tends to emphasize experience is "mysticism". No doubt there will be a temptation to equate the Wysiwyg view with the picture of the world described by various mystical schools. However, this would be a mistake.

First of all, unlike science, there is no one set of procedures or generally accepted truths about the world which we can lump into a single category called "mysticism". Scientists, by and large, speak the same language the world over. The results found in one lab need to be reproducible in anyone else's lab before the results enter the general database of scientific knowledge. New theories in science must be testable, at least in principle, and any data collected which contradicts the predictions of a theory necessarily invalidates it. It is very difficult to make a living in this day and age pedaling false science--someone will call your bluff in very short order.

Mysticism generally begins with a simple process of reflection called *meditation*. The basic idea is to allow experience simply to "happen", without preconceptions and without intellectual analysis. In of itself, this may be the only way to truly appreciate the essence of experience as a foundation for all that is real. But from this simple beginning, whole schools of mystical thought have sprung forth over the millennia. Countless volumes are written on the insights gained through meditation regarding the true nature of reality and our place in the cosmic order of things. Alas, as testable models of the world, most of these descriptions fall far short of even basic commonsense checks which anyone can perform in their daily lives. Needless to say, without testability, many such ideas are ripe for suggestibility, preconception and autosuggestion. Furthermore, many such descriptions date back thousands of years and seem starkly out of place in today's world. It would be as if we continued to follow Aristotle's model of physics simply because he was considered to be an authority two thousand years ago!

As we have seen in previous sections, the Wysiwyg Universe is a Scientific Universe. There seems to be no end to the correlations among the rich and varied experiences each and every one of us enjoys on a daily basis. And these correlations tend to be agreed upon by anyone who cares to look. Should we happen to find new experiences during the course of meditation, these will simply add to our repertoire of experiences. If there are correlations between these experiences and others, time will reveal their structure and they will become generally known amongst a greater number of people.

Perhaps the closest parallel one can find between the Wysiwyg hypothesis and mysticism can be summed up by the famous anonymous Zen saying:

Before satori (enlightenment), the mountain is a mountain
During satori, the mountain is no longer a mountain
After satori, the mountain is a mountain again

Questions and Answers

Wysiwyg and the Brain

When describing the relationship between experience and the brain, we seem to confront a classic chicken-and-egg problem: What comes first, experience or the brain?

When studying brain activity and its relation to experience, there is a tendency to draw a rather arbitrary balloon around the head and say: everything inside the skull we will call "internal" and everything outside the skull we will call "external". Most of us go along with this since, after all, brains are very much hidden away inside their owner's heads, unlike, for example, hands or faces. Like most internal organs, we can't see our brain, or hear it, touch it, smell it or--ahem--taste it. So when we are confronted with the problem of, for example, how a brain hears the singing of a bird, we are tempted to say: the source of the sound is in the "external world" and there we find patterns of vibrating air molecules. But the brain operates in the "internal world" and here we find only the activity of neurons. So where exactly is the sound? Is it in the external world or the internal world? And if it is in the internal world, how can the firing of neurons actually *be* sound? Of course, this is where the whole mind-body problem sneaks in.

While this division of the world into what's inside the head versus what's outside the head makes sense from a purely engineering point of view--for example, how could we make a computer recognize the sound of a bird--it has no ontological basis: the brain is just as much an object in the world as are birds and people and air molecules. When we stop to think about it, no one would argue this point, but our long entrenched speaking habits often lead us down an unnecessary dualistic path.

The root of the confusion arises from our switching back and forth between a first-person perspective and a third-person perspective. For example, if you are listening to a bird singing on a branch before you, a neuropsychologist might say that the singing does not exist in the "external" world, only in your brain (the "internal world"). What exists in the external world are just patterns of vibrating air molecules which emanate outward from the bird. These vibrations then impact upon the membranes of your inner ear, which in turn stimulate your auditory nerves, and ultimately the parts of your brain responsible for sound. For this reason, your brain is said to *construct* the experience of sound from the raw sense data of vibrating air. Of course, the psychologist doesn't actually have access to your experience, so the best he can do is say that your brain constructs a neural representation of a bird singing. This is the third-person perspective--the neurospsychologist is the third individual viewing the interaction between two other individuals, you, or rather your brain, and the bird.

The first-person perspective is simply what you experience. From this perspective, the experience of the bird song really is *over there* coming from the bird sitting on the branch--not somewhere in your brain. In fact, you have no sensation of your brain activity at all. Indeed, if you had been raised by wolves rather than in a human society, you probably wouldn't even know that you had a brain inside your head. But you would presumably still have the experience of the bird singing over there on the branch. No matter how much physics or neurobiology someone throws at you, from your point of view, the bird is experienced as over there on the branch, not inside your head.

Let's now bring these two perspectives together using a simple thought experiment. There you are, sitting in the garden, experiencing a bird singing on a branch. To better understand the relation between the sound of the bird and the activity of your brain, you attach a functional MRI device to your head. The device has a large color display which shows in fine detail the activity of the various parts of your brain. In this way you have taken your brain outside of your head without having to cut open your skull. As a result, your brain can now be seen for what it is: another object of your experience just like the bird.

As the bird continues to sing, the MRI display reveals that various parts of your auditory cortex light up, along with regions of your association and visual cortex, and parts of the language areas in your left hemisphere (if you are right handed). There is even significant activity in your medial forebrain bundle which reflects the fact that you are enjoying the sound!

You now have both first-person and third-person perspectives on your situation. You have the experience of the bird singing, and you have all the data you need to understand how the brain constructs the experience from the pattern of vibrating air striking your ear. Or do you? The MRI display appears to you as a visual experience which is very nicely correlated with the auditory experience of the singing bird. In fact, the correlation is of the very same nature as you might notice if you viewed a closeup of the motion of the bird's beak while it made the sound. In both cases you see a correlation between the sound experience and a visual experience. And that's it. Just correlations among experiences--mathematical structures which allow us to predict one experience from another. No ontological distinction between "internal" and "external" worlds. No need to wonder how the "physical" brain and environment give rise to our "mental" experience.

In the Wysiwyg view, there is only experience. As an object in the world, the brain is a collection of correlated experiences--the experiences we have when studying the organ inside peoples' head. This structured set of experiences which we call the brain can in turn be correlated with other experiences we have such as the singing of birds, the color of the sky, sensations of hot and cold, hunger and thirst, etc. In this way we can see that experience is not a construction of the brain; rather, the brain is a structure in our experience.

Wysiwyg and Occam's Razor

Doesn't the Wysiwyg hypothesis violate Occam's Razor?

When first presented with the Wysiwyg hypothesis, one might object to the second postulate as follows: "I can accept that experience is the only reality for an observer. But how do we account for all the myriad correlations among our experiences without reference to an 'external world?' In other words, how can experience be all that there is?"

First of all, we must be clear about what is meant by the phrase "external world". As we saw in the previous question on *Wysiwyg and the Brain*, we sometimes speak as if "external world" simply means everything "outside the brain". But the brain is clearly in the "external world" as much as any other object and certainly cannot be given any special ontological status *a priori*. So by "external world" we must mean "outside experience". Not just outside current experience, but outside all experience, even in principle. (Don't forget that observations we make using telescopes or microscopes or particle accelerators, etc., are all still part of our experience.) The Wysiwyg hypothesis postulates that there is no reality other than experience. The objection raised above can now be rephrased as: "How do we account for all the myriad correlations among our experiences without reference to anything outside experience?"

When evaluating competing models of some domain of our experience, scientists often use a principle called *Occam's Razor*. William of Occam was a medieval philosopher who first put forth the idea that, given a choice between two explanations of some phenomenon, we should choose the simplest--the explanation which requires the fewest assumptions. For example, if you assume that the Sun and the rest of the solar system revolve about the Earth, as did the ancient Greek astronomer Ptolemy, then you need to also assume that the other planets follow very peculiar orbits in order to account for the observed motions through the night sky. On the other hand, if you assume that the Earth and the other planets revolve about the Sun, as Copernicus argued, then no further assumptions are required to account for the observed motions. Both the Ptolemic model and the Copernican model can explain the observed motions of the planets; however, the Copernican model makes far fewer assumptions and is therefore deemed preferrable.

Readers familiar with Occam's Razor might argue that the second Wysiwyg postulate requires an infinite number of assumptions; namely, one assumption for every observed correlation among experiences. Wouldn't it be simpler to postulate the existence of variables or entities beyond experience to account for these correlations? If the number of these variables or entities is fewer than the number of correlated experiences, then this non-Wysiwyg view might be considered simpler and therefore preferrable to the assumption that there is no reality other than experience.

There are a number of problems with this argument, some of which have already been spelled out in the earlier parts of this document. First of all, Wysiwyg makes no assumptions regarding correlations among experiences; they are simply observed. If we wish to *model* these correlations for a *particular* domain of experience, then we have to choose among the large number of all possible models (mathematical structures) which include the observed correlations as a subset. And it is at this stage that Occam's Razor can play a role in choosing among different models. Occam's Razor was not intended to be applied to the set of *all* experiences--or to *any* experience for that matter--only to *models* of a subset of our experiences. And even if we could apply the principle to the Wysiwyg postulates, which sounds simpler: that there is no reality other than experience; or, that in addition to everything we experience (including what we observe through microscopes and telescopes and infrared glasses etc.), there is a whole "world beyond experience" which we cannot observe but which somehow holds it all together? In this respect, Occam's Razor seems to actually favor Wysiwyg.

But there is an even more compelling reason to let go of the idea of a "world outside experience": it is a mistake to assume that by simply referring to a world outside experience, we have then accounted for the observed correlations among experiences. For one could just as well ask: how then do we account for the structure in this other world beyond experience? We can best illustrate this point with a specific example.

If we regularly experience the sight of lightning followed by the sound of thunder, we could sum up our experience with the phrase (simple verbal model) "thunder follows lightning". Suppose we then ask the question: "But *why* does thunder always follow lightning?" If we are simply making things up as we go along we might answer: "Thunder is the sound of Thor grumbling whenever he is blinded by a flash of lightning." Of course, then we have to explain who Thor is, why he doesn't simply cover his eyes during a thunder storm, and why he never seems to appear in public. But if we are doing science, the first thing we do is broaden our set of experiences by looking for intermediate phenomena (experiences) between the sight of lightning and the sound of thunder. So we set up some lightning rods and some recording devices and wait for the next storm.

When the data are collected, we notice the following chain of experiences: electrical surge, bright light, high temperatures, low pressure near the electrical surge, high pressure radiating out from the region of low pressure, loud sound. We put these together verbally and say, "An electrical surge through the air leads to a flash of light and a hot core of air around the surge. This hot air expands (low pressure), compressing the surrounding air (high pressure). When the high pressure air hits our ears we hear a loud sound." Notice that at no point in this explanation have we appealed to entities or processes beyond our experience. But by expanding our collection of observations, we are able to produce a more detailed model of why thunder follows lightning.

Of course, one can play the "why" game to any level of detail you like: Why is there an electrical surge? Why does hot air expand? Why is there an atmosphere at all? Why does the Earth exist? And so on. In the end, you will always end up with a set of correlated experiences. What's more, the resulting explanation in terms of these correlations is just what we call a scientific model (physics in this example). If you are studying a reductive phenomenon like lightning and thunder, then you could conceivably collect a (nearly) "complete set" of predictive pairings which accounts for (almost) every "why" question in the explanatory chain. Some of the final questions might go something like: Why are there electrical charges? Why are there electrons? Why is there matter at all? Why was there a Big Bang? Why did Nothing suddenly become Everything? And then you are stumped, because the correlation {Nothing, Everything} is a pairing that makes no sense from a scientific point of view and cannot be further reduced.

This is all just a long winded way of saying that the set of experiences fall into predictive pairings because that is the way the world is. If this were not the case, we wouldn't be having this discussion. As far as we can tell, science never answers the ultimate "why" questions. Why is the experience of "red" red rather than blue? Why does space have three dimensions instead of twenty? Why did the Universe spring from Nothing? Why does anything exist at all? The best science can do is to model the predictive pairings found among our experiences so that we can in fact use one to predict another. If you want to know the nature of your experiences "directly", then you should study your experiences themselves rather than a model of their correlations. The process of studying your experiences directly is called *meditation*.

Wysiwyg and Idealism

`Isn't the Wysiwyg hypothesis the same as Idealism?`

The philosophical school called *Idealism* dates back to Bishop George Berkeley, an Irish philosopher and Anglican clergyman living in the 1700's. While the Wysiwyg hypothesis may seem similar to Idealism on the surface, the two approaches are actually quite distinct. In particular, Berkeley's Idealism is really a kind of *theological epistemology*, as many thinkers at the time were trying to reconcile the new empiricism of John Locke with their own religious convictions.

First of all, the Wysiwyg hypothesis is contained entirely within the two postulates which are repeated here as a reminder:

- Experience is the only reality for an observer.
- There is no reality other than experience.

Berkeley used the word "idea" where we use "experience" or "sensation"; hence, the term "idealism". The modern connotations of the word "idea" suggests something beyond basic experience, such as "concept" or "thought." However, at the time of Berkeley's writing, "idea" denoted something much simpler and we can safely substitute "experience" in its place.

Berkeley assumed that each of us has a "mind" and argued that "matter" cannot be conceived to exist independent of the mind. In essence, Berkeley sets himself up for a kind of implicit dualism: mind versus matter. However, according to Berkeley, matter can only exist if it is perceived by the mind. Furthermore, to account for the apparent ability of things to exist even when one is not observing them (which apparently contradicts his original premise), Berkeley held that all things must be constantly perceived by God. (Remember, Berkeley was a member of the clergy as well as philosopher.)

By contrast, Wysiwyg simply says that each of us has experiences and that there is no reality other than experience. "Matter" and "mind" are labels we have for certain aspects of our experience. The notion of objects outside our experience is part of our *model* of the world. It permits us to summarize our experience and make predictions about other experiences. However, there is no reality outside experience.

To the non-philosopher (which includes the author), Idealism and Wysiwyg will probably seem like birds of a different feather. But just in case the distinction is still unclear, we can summarize the differences as follows:

Idealism	Wysiwyg
Reality consists of mind and matter but...	Reality is experience. "Matter" and "mind" are labels for certain aspects of our experience.
...Matter does not exist outside of the mind.	There is no reality other than experience. "Inside" and "outside" are also labels for certain aspects of our experience.
Objects exists even in the absence of observers because they are continually perceived by God.	The notion of objects outside our experience is part of our model of the world. Models of reality should not be mistaken for reality itself.

Ever since Berkeley put forth is ideas, philosophers have been somewhat annoyed to find that Idealism is very hard if not impossible to refute. What bothers them most is that an "external world" clearly seems to exist "outside of the mind". Dr. Samuel Johnson was a prominent physician and contemporary of Berkeley. In James Boswell's biography of Johnson, Boswell recounts:

After we came out of the church, we stood talking for some time together of Bishop Berkeley's ingenious sophistry to prove the nonexistence of matter, and that every thing in the universe is merely ideal. I observed, that though we are satisfied his doctrine is not true, it is impossible to refute it. I never shall forget the alacrity with which Johnson answered, striking his foot with mighty force against a large stone, till he rebounded from it -- "I refute it thus."
A refutation of Idealism, maybe. But what could be a more poignant confirmation of Wysiwyg than a stubbed toe!

Avoiding Solipsism

The Wysiwyg postulates seem to imply that every individual knows only what they experience on their own. What about the knowledge shared with other individuals?

Shared knowledge is learned by having experiences--listening, reading, viewing film or video, observing others, etc. But what we acquire in this manner is a shared *model* of the world. For example, an individual who has never experienced snow directly might utter the phrase "snow is cold". This knowledge might result from reading books, watching television, or listening to someone who has experienced snow. But of course, the knowledge that "snow is cold" is not the same as experiencing snow itself. And as far as we can tell, there is no way that one individual can have someone else's experience. So in the end, if you really want to know something about the world for yourself, you need to experience it yourself. All the knowledge in the world about chocolate cake won't taste as good as the real thing. But it will come in very handy if you want to bake the cake in the first place.

He Said, She Said

How do we explain the ability of two or more individuals to agree on the identification of an object without reference to a world beyond experience?

Suppose you and I meet for the first time. We might not even speak the same language. Assume that I have an experience, *B*, to which I assign the verbal label, *Blue*. At the same time, suppose you are having the experience, *b*, and that my verbal utterance generates in you the experience, *blah*. Then for me, *B* ~ *Blue*, and for you *b* ~ *blah*. If at some later time I utter *Blue*, you will experience *blah*. If at the time you are also experiencing *b*, then by your correlation, *b* ~ *blah*, you will agree with my utterance. If you are *not* experiencing *b*, then you will disagree.

So by using a verbal label when I have the experience *B*, I can induce in you a corresponding label for your experience *b*. Note that as far as we know, there is no way either of us can directly compare our experiences, *B* and *b*. It is only through the intermediate experience of language that we can agree on the correlations of our other experiences. Of course "language" here simply means any method by which we can signal each other. One need only try to understand two people talking in an unfamiliar foreign language to appreciate the arbitrariness of assigning labels to experience. Yet once the labels are habitually assigned, two individuals can (generally) convey the sense of their separate experiences to each other.

Isn't Wysiwyg just another model?

Actually, Wysiwyg is just WYSIWYG! :-) Just remember that the finger pointing at the Moon is not the Moon.

Appendix

Coming Soon!

Acknowledgments

Many thanks to Jeffrey A. Gray whose penetrating discussions on the nature of consciousness got the whole ball rolling for this document.

References

Harrison, David M. (1999), *Bell's Theorem*

Hayes, Brian (1998), *The Invention of the Genetic Code*

Whitaker, Andrew (1998), *John Bell and the Most Profound Discovery of Science*

Document Version 1.1, April 2002
