

On the Ontology of Information

It has reached the level of cliché to claim that we live in an information-centered society. Mass media, especially in the form of the Internet, provide us with historically unprecedented access to all kinds of information about civic affairs, geopolitics, consumer products, commercial enterprises, and other human affairs from the mundane to the sublime. The quantity and rate of information delivery have wrought and are yet working profound changes in social structure, mental habits, law, politics, and commerce in the developed and developing worlds. The considerable cultural force of what is loosely and popularly called “information” suggests the benefit of a careful ontological assay of the thing.

An evaluation of the ontic status and composition of information rightly begins with the theorist who is considered to be the founder of information theory: Claude Shannon. We will review his account of information, as well as its more philosophically oriented recapitulation in Fred Dretske. We will evaluate whether information must be considered in connection with its semantics, then whether it is possible not just to transmit information, but also to store it over time. Finally, after briefly observing information’s encounter with the classical one-in-many problem, we will be prepared to evaluate what kind of a thing information is and how it might persist over time.

Shannon's Account

The first mathematically rigorous treatment of information is generally considered to be Claude Shannon's "A Mathematical Theory of Communication" published in *The Bell Systems Technical Journal* in 1948. In this seminal essay, Shannon conceives of a communication system as one in which a source sends information through a channel, where the signal is corrupted by noise, then is recovered at the destination.¹ The information at the source consists of a series of symbols chosen from a finite set of possibilities, the channel is some medium able to transfer the symbols with an arbitrary (and potentially imperfect) degree of faithfulness, and the destination is simply a place where the symbols are recovered, where some processes may be in place to remediate the effects of the noise in the channel.

Shannon's essay made at least two contributions to the state of the art, which have endured as central principles of communications engineering in the half-century since the essay's publication. The first is the formula which has become known as Shannon's Law. It describes a nomic relation which holds between three quantities: the maximum rate of accurate information transfer, the speed at which symbols can be sent through a channel, and the degree to which noise can be expected to corrupt symbols in transmission.² This law continues to govern the design of communication systems today, and has achieved the status of a physical law beyond any interesting dispute.

¹ Claude Shannon, "A Mathematical Theory of Communication," *The Bell Systems Technical Journal*, vol. 27 (July, October 1948), 2.

² *Ibid.*, 43.

The information-theoretic insight which is prior to Shannon's Law (and which is of greater metaphysical import) is the analogy between information and *entropy* in physical systems.³ Physical entropy is roughly the degree of disorderliness in a physical system, or the amount of energy in the system that is not available to do work. The higher the entropy of a system, the less work it is able to do, or put more loosely, the less organized the system is.⁴ For example, two tanks of water—one hot, one cold—might be employed as the thermal reservoirs in a heat engine by virtue of their temperature difference.⁵ The orderliness of the separation of hot from cold can be employed to do useful mechanical work. If the tanks are mixed together to form a less orderly, more mixed-up tank of water of a uniform temperature, the energy in the water can no longer be employed to do work. Under these circumstances, the entropy of the system is said to *increase*.⁶

Similarly, information entropy measures the amount of “disorder” of an information source. The information generated by the roll of a six-sided die is said to have higher entropy than that generated by the flip of a coin, since there are more possible outcomes involved in the roll of the die, or alternatively, the die roll is more “mixed-up” than the coin toss.⁷ Put differently, information entropy is the degree to which a symbol passing through a communication system informs the destination about states of affairs at the source. This important concept has found further elucidation in

³ Ibid., 11.

⁴ Francis Sears, Mark Zemansky, and Hugh Young, *University Physics*, 7th ed. (Reading, MA: Addison-Wesley, 1987), 440.

⁵ Ibid., 444.

⁶ Ibid., 443.

⁷ Shannon, 13.

Fred Dretske's 1981 work, *Knowledge and the Flow of Information*, which we will consider next.

Dretske's Account

Dretske provides a helpful example to explain the concept of the *amount* of information in quantitative terms, following Shannon's earlier example.⁸ We will adapt that example to our purposes here. Suppose there is an electronic commerce web site, *ephilosophers.com*, which allows users to shop for philosophy from the privacy of their homes. Suppose *ephilosophers.com* is offering a promotion in which a small number of users are selected to receive a complimentary philosophical treatise delivered to them free of charge. Terry, a frequent *ephilosophers.com* customer, is selected as a winner of the contest, and is told he will soon be receiving one of eight possible treatises chosen randomly from among "On What There Is," "On Sense And Reference," "On The Various Kinds of Distinctions," a passage from Aquinas' "Structures of Things In General," and four other writings of common interest to the philosophical community.

In this simple example, there are eight possibilities at the source of the information transmission. It is not terribly important to our purposes whether the choices are equally likely, but for simplicity's sake we will assume they are.⁹ The mere existence of un-actualized possibilities at the source does not in itself constitute information. Instead, information is realized when the eight possibilities are reduced to one. When a process

⁸ Fred Dretske, *Knowledge and the Flow of Information* (Cambridge, Massachusetts: MIT Press, 1981), 4.

⁹ It could have been the case that the probabilities were not equal. "On Sense and Reference" may have constituted 50% of the prizes, "On What There Is" may have had 20%, and the remaining six essays may have taken 5% each. This would complicate the equations presented, but it would not materially change the issues at stake.

inside philosophers.com actualizes the choice of one of the treatises, then mediates the choice to Terry, a transmission of information can be said to have occurred. Dretske gives us an equation for calculating the information content in bits of this state of affairs:

$$(1) \quad I(s) = \log_2 n$$

where n is the number of uniformly likely choices.¹⁰ The logarithm base two, which yields results in units of *bits*, was a convention in force at the time of Shannon's original writing.¹¹ In the case of our eight possible prizes, $I(s)$ yields a result of 3 bits. We may call this either the *average information* generated at the source or, in Shannon's terms, the *entropy* of the source. It doesn't matter how many possibilities there were to begin with, or even whether Terry was told precisely which essay he would be receiving. If, for instance, Terry's email told him to expect either Quine or Suarez, then he still would have been the recipient of some information: in this case, $\log_2 4 = 2$ bits. The key insight here is, as Dretske says, that "the information associated with a state of affairs has to do, simply, with the extent to which that state of affairs constitutes a reduction in the number of possibilities."¹²

Sidestepping Semantics

It is worth a brief excursus to disentangle information from its frequent companion, *meaning*. There are at least three reasons for denying that meaning and information are identical. First is the observation that not every meaningful sentence contains

¹⁰ Dretske, 7.

¹¹ Shannon, 1.

¹² Dretske, 8.

information. The sentence “Quine is the author of ‘On What There Is’” is meaningful (that is, it refers to an author and an essay both well-known in certain circles, it predicates of the author that he has the property of having written the essay, it is veridical, etc.) but does not contain information, because it does not constitute a reduction of possibilities in any real state of affairs. In reality, it is not the case that either Quine or Chisholm might have been the author of “On What There Is,” and that upon reading this sentence we can learn which man did the deed. It is a fact that Quine was, and it cannot really be the case that he was not.

Second, accordingly to Dretske, a single statement can be shown to bear information content that is different from its meaning. The statement “Terry is teaching class” may be said to mean (if the reader will indulge a brief episode of hermeneutical intentionalism) that Terry is teaching class. It cannot, under any credible hermeneutical framework, be said to mean that Terry is either teaching class or getting his hair cut. However, the statement “Terry is teaching class” does in fact *carry the information* that Terry is either teaching class or getting his hair cut.¹³ This is a consequence of Dretske’s concept of “analytical nesting,” in which the information that s is F carries the information that t is G if t being G is an analytical consequence of s being F .¹⁴ Analytically speaking, it must be the case that Terry is either teaching class or getting his hair cut if Terry is in fact teaching class. Thus are meaning and information content distinguishable.

¹³ Ibid., 72.

¹⁴ Ibid., 71.

Finally, it is possible to conceive of states of affairs which fit our definition of information but which bear no semantic content at all. Imagine a waterfall with a thick cloud of mist at its bottom. Several trees grow in the mist by the river, drooping under the weight of the constant dew, dripping water off of their leaves chaotically.¹⁵ Imagine a detector is situated under one dripping leaf such that a light beam will be interrupted each time a drop falls. This detector will modulate a radio signal sent to a receiver located in orbit above the planet, which will be able to detect when the drops occur and blink a light in response.

It seems inarguable that information is present in the system; indeed, it has at its core a classical communication system following Shannon. Following Dretske, we find that uncertainty about states of affairs is reduced by watching the blinking light. We know that water droplets will fall, but we are uncertain about precisely when they will fall. We could take measurements and discern the probability of a droplet falling within any given time interval, but we would remain uncertain whether a droplet in fact fell during any actual interval except we watch the light. This is certainly information, yet it bears no meaning at all. The blinking of the light may “mean” that a water droplet is falling, but the water droplets themselves are not recognized as signs in any language and do not signify anything outside themselves. They are—if such a hermeneutically bold assertion may be made in respectable company—meaningless.

Clearly non-semantic information is the exception in the cases of interest to us. Even in the case of semantic sources, however, we may recognize the existence of

¹⁵ By “chaotically,” I mean that the timing of the drops is characterized by three properties: it is deterministic, not periodic, and not predictable in the long term. The interested reader can find a good popular-level introduction in James Gleick, *Chaos* (New York: Viking Penguin, 1987).

information when its meaning is inaccessible. As Floridi points out, hieroglyphics were considered to contain information prior to the discovery of the Rosetta Stone, and counterfactually would contain information today if the Rosetta Stone has never been discovered.¹⁶ It seems regardless of the case, despite their frequent association, information and meaning are truly distinct.

Information Transmission and Storage

The analyses of Shannon and Dretske are preoccupied with information as *event*. They discuss the transmission or mediation of information from source to destination, but do not seem to account for the case of information sitting statically on some kind of storage medium, as the loose and popular conception of information would easily allow. Can our stricter and more philosophical concept of information abide being stored?

Consider again Terry's happy news that he won free philosophy from ephilosophers.com. Of the eight possible essays he could receive, he learns that he has won a copy of "On What There Is" by W.V.O. Quine, thus reducing the eight possible states of affairs to one. The thrust of the examples in the literature concentrates, not inappropriately, on the *message* Terry receives. Messages in this sense are dynamic things, similar to events, which have clear temporal boundaries. Their temporal boundaries—and with them the sense in which they represent a reduction in the possibilities of probabilistic states of affairs—become unclear when they are committed

¹⁶ Luciano Floridi, ed., *The Blackwell Guide to the Philosophy of Computing and Information*, "Information" (Malden, Massachusetts: Blackwell Publishing, 2004), 45.

to storage media. An example may help restore clarity and release information from the would-be prison of persistent storage.

Ephilosophers.com has a certain population of users and a certain inventory of philosophy for sale. There is, during any given interval of time, a certain nonzero probability that any given user is making a purchase and that any given piece of philosophy is being purchased. In addition, there is a set of conditional probabilities that any given user is buying any given piece of philosophy, and that any given piece of philosophy is being purchased by any given user.¹⁷ A real-time display of the sales activity of the site would clearly preserve the event-oriented spirit of Shannon and Dretske, but let us consider how a printed report of past sales might do the same.

We are, a priori, uncertain of what was purchased during August, 2005. We are also uncertain of who purchased it. There are many possible states of affairs that may have obtained in ephilosophers.com in that month, and any message that reduces those possibilities contains information. If we read a printed sales report that says that Terry made a purchase in that month, our uncertainty is reduced. We could even calculate the information content of this news in bits, if we knew how many registered users could possibly have made purchases that month, how likely it is that Terry would make a purchase during August, or other factors chosen according to our purposes. (Clearly some fairly subjective judgment calls would have to be made before making such a calculation. However, that the report's information content is difficult to assay quantitatively does not suggest that it is improper to speak of it containing information.)

¹⁷ Some of these probabilities may be zero. Terry may in fact never pay money for Quine, so we may never in fact find any day in which Quine is purchased by Terry. This case does not affect the information content of any ephilosophers.com sales reports.

An interesting case arises if Terry is certain to purchase some philosophy at least once a month. In this case, the probability of Terry buying philosophy in August is one, so the message that Terry bought philosophy in August contains no information, since no reduction in possibilities has taken place. However, the message that Terry bought Willard lecture notes in August *does* contain information, since it reduces Terry's possible purchases from some presumably large number to a smaller number (in this case, one).

The sales report can continue to whittle away at uncertainty until an arbitrary limit is reached. Knowing that Terry bought Willard lecture notes in August reduces our uncertainty about the system considerably, but it leaves some ambiguity about whether he purchased them before or after Jenni bought the Frege promotional gift pack. Adding some kind of indication of sales sequence reduces this uncertainty, thus increasing the information content of the report. Even so, mere sequence leaves ambiguity as to the day on which Terry's and Jenni's purchases took place, which ambiguity could be reduced by including a date and time stamp associated with each purchase. This would further increase the information content of the report, and so on.

This example makes it clear that it is in some sense appropriate to speak of stored information alongside the instantaneous information normally envisaged by analyses in the tradition of Shannon. This opens up a significant number of new cases for analysis under the tools provided herein. In the interest of brevity, we will consider only one.

One In Many

Following our discussion of information storage, we may uncover an important insight if we consider that *how* a message is stored does not seem to have any effect on the message's information content. Recall Terry's earlier email informing him that he has won free philosophy. The missive may have contained text something like this:

- (2) Congratulations! You have won a copy of Quine's "On What There Is."

But it might very well have been rendered like this:

- (3) Congratulations! You have won a copy of Quine's "On What There Is."

Moreover, it was probably represented inside the computer system like this:

- (4) 43 6f 6e 67 72 61 64 75 6c 61 74 69 6f 6e 73 21 20 59 6f
75 20 68 61 76 65 20 77 6f 6e 20 61 20 63 6f 70 79 20 6f
66 20 51 75 69 6e 65 27 73 20 22 4f 6e 20 57 68 61 74 20
54 68 65 72 65 20 49 73 2e 22¹⁸

with no change in information represented. The eight possibilities have still been reduced to one, and indeed, to the *same* one possibility. Moreover, all three representations contain the additional information that Terry himself was a winner. Whether or not Terry was even aware of the contest, all of the registered users of ephilosophers.com were possible winners before the selection was made, and Terry's

¹⁸ This is a hexadecimal UTF-8 encoding of the characters in the messages in (2) and (3). The numbers shown are base-16 representations of the numbers that would typically be used to store the characters of the message using technology contemporary to this writing. Of course, numbers in such computers are typically "represented" by levels of charge in microelectronic memory elements or by the orientation of ferromagnetic particles deposited on the metal platters of a hard disk, but the level of abstraction pictured here by printing the numbers in hexadecimal does not misrepresent these reductive accounts.

notification that he was one of them represents a reduction in possibilities in this state of affairs.

The same thing can be said about various forms of ephilosophers.com sales reporting. A real-time sales dashboard rendered on a computer screen, showing instantaneous sales events as they occur and various histories over short time intervals in the recent past, can conceivably represent the same information as appears on printed sales reports for the same time interval. Both tell us which possible sales were made to which possible customers, despite the very different forms in which the information is rendered. Both contain the very same information—the very same *thing*.

The Ontological Assay

We are now equipped to turn our analysis to the ontology of information itself. What kind of a thing is it? What things attend its existence and transmission? According to the definition given at length above, information is a reduction in the possibilities in a state of affairs. The “state of affairs” may be analyzed as an object which has the capacity to exemplify more than one property-instance of a certain kind, but in reality actualizes only one of the possible property-instances.

The “object” can be simple or composite, concrete or abstract. Ephilosophers.com and a two-sided coin are examples of concrete composite objects. The former may have the property of awarding Terry with any of eight possible free philosophical treatises, and the latter may have the property of being on heads or tails. Each has a number of possible property-instances but in fact can actualize only one of that kind.

It is important for purposes of this analysis to consider only those possible property-instances which are of the same kind. A responsible calculation of some case of information content would not be concerned that ephilosophers.com may have the property of awarding Terry a copy of Suarez today and simultaneously that it may have the property of making \$150,000 in sales in November. Knowing that the site had \$150,000 in sales in November is not informative by virtue of Terry not having received Suarez; rather, it is so in virtue of the site not having had \$149,999 (nor \$150,001, nor \$1, etc.) in sales that month. The sameness of the kind of the property-instances may be analyzed as an identity relation holding between the essential second-order properties of the operative first-order properties whose instances are under consideration.

Abstract objects can have information-related properties also. Although the information content of mental states has not been analyzed in detail above, it seems clear that if I am thinking of a number between one and ten, and Terry guesses the number, he gains information about what I am thinking when I tell him whether he guessed correctly. While the ontology of thoughts may be controversial, they are good candidates for non-concrete objects which can be associated with information.

The reduction in possibilities can be analyzed as a temporally bounded event or process through which counterfactual property-instances are identified as not obtaining in the real state of affairs. Prior to some time t , ephilosophers.com has the possibility of having one of eight properties of the kind *essay-that-Terry-won*, but it does not exemplify any of these properties. After an essay-awarding process which completes at time t , the site begins to exemplify the property of having awarded Terry his new copy of "On What There Is." The fundamental operation in the creation of information is when the object

under analysis begins to exemplify a property of the relevant kind, and in the process *not* exemplifying the properties which were formerly possible.

Assigning a Category

Now let us consider the category in which we should regard information itself. It bears importantly on this question that information can be measured. We must know three things to make this measurement: first, how many counterfactual property-instances do not obtain; second, how many property-instances may still possibly obtain; and third, what the probabilities were of each property-instance obtaining. These quantities may be very difficult to ascertain, but a quantitative assessment of information content is in theory possible. This suggests either that information is a property, or if it is an object or a relation, then it the property of “quantity of information” must inhere in it.

It is not clear how information could be a relation. If it is, then what are its relata? The so-called “state of affairs” in the definition—the object potentially holding the various property-instances—is surely one candidate, but no other ontologically necessary entity is present in this analysis to take the place of the other. (The “state of affairs” object has at least one necessary property, but it would not do for us to be positing a relation between an object and its property other than species of the genres “has” or “is had by.”) Considering information to be an object is likewise problematic. If this were the case, then we would have to posit a relation to account for its connection to the state of affairs it describes, which is not recommended for at least two reasons. First is that it is not clear what this relation would be. Obvious examples like “informs of” or “informatively describes” seem contrived and do not increase the clarity or explanatory

power of the analysis. Second is that the addition of the relation merely for the purpose of gluing our analysans together is ontologically profligate, and should be rejected if a more parsimonious account can be given.

Considering information to be a property has its own difficulty: to wit, what shall we say it is a property *of*? There are several candidates. It can be a property of the object whose counterfactual property-instances have been eliminated (that is, the flipped coin or the essay-choosing web site). It can be a second-order property of the property-instance which has actually obtained in that object. Alternatively, it can be a property of the event or process which has eliminated the counterfactual property-instances and caused the one actual property-instance to be exemplified.

Making it a property of the analyzed object may raise troubling identity questions. If the reader's ontology will permit, let us consider that a coin maintains its identity in some sense after being flipped.¹⁹ This relatively modest metaphysical indulgence begins to presume upon good judgment if we are forced to add properties to the coin to "contain" the information generated by the flip. Is the informational property destroyed after each successive flip? If so, how? If not, and various informational properties still inhere in the coin, is the past history of its information somehow accessible to observers? Certainly the proliferation of properties—whether one at a time or increasingly many

¹⁹ The coin may exhibit the property of being on heads or the property of being on tails, but we are assuming that these can be analyzed in such a way as to do no harm to the coin's identity regardless of how it lands. This essay intentionally ignores questions of how a coin can have the property of headedness or tailedness, whether these properties involve any extrinsic relations, whether they are merely conventional, and so forth. If coins and web sites are objectionable, then simpler and even "truly" random examples (e.g., quantum phenomena, perhaps) can be adduced that should stop the mouth of even the most obstreperous reader. By comparison, however, these lack a certain panache present in shiny coins and philosophy ecommerce sites.

over time—would seem to violate the identity relation the coin has with itself. This is sufficiently counterintuitive as to cast this option into disfavor.

Making it a second-order property of the realized property-instance causes similar identity problems among the analyzed object's property-instances. Prior to realizing one of the property-instances, the analyzed object can be said to *have potentially* the possible property-instances in question. After one of the property-instances is realized, if the realized property assumes the second-order property of information, it loses its identity. It then has a second-order property that it did not have before the information was realized, so it will fail the test of the Indiscernibility of Identicals. Thus the object can never realize any of its potential property-instances, since in this scheme they would change identity as soon as they are exemplified.

This leaves us with one option for the location of the property: the event which brings about the reduction in possibilities in the probabilistic state of affairs. Since the event is not a temporal continuant (at least it is not in any important sense), the identity objections raised above do not have the same force here. This event has the property of being information-generating, or "is an information-generating event" by virtue of its agency in eliminating the counterfactuals and causing to be exhibited the one actual property. Put differently, the information property's inherence is grounded in the elimination of counterfactual property-instances and the realization of a single property-instance.

However, this exposes the analysis to two difficult considerations. First, and not least, is an ontologically disputable account of events. This analysis presumes a more or less Realist understanding of them, since they are represented as entities which stand in

the relation of predication to the proposed property of information. Whether the overall account of information contained herein stands on its own and thereby commends a Realist ontology in some small part, or whether it can be accommodated to nominalist thinking without too much damage to its distinctives, is a question for further dialogue and will not be addressed in this space.

The second problem with predicating information of an event is the temporal nature of events themselves, and the implications this nature has for information storage, which earlier we deemed to be unproblematic. Regardless of what ontological account is given of events, it is unlikely that one can be defended that regards them as temporal continuants in any meaningful sense. It is very near certainty that events are a kind of thing with clear temporal boundaries, coming into being and immediately (or at least very quickly) ceasing to exist, yet above we spoke confidently of the transmission and storage of information. There is a clear tension here, one which can be resolved by analyzing the storage of information in more detail.

To begin with, consider that an image of the Lincoln Memorial cold-forged into copper does not necessarily contain any information. Even if we know that it is not just any image of the Lincoln Memorial, but is in fact a penny on tails, and we know that pennies are common means of generating information in small doses, it is still the case that merely regarding the penny sitting tails-up on a table does not by itself mediate information. The penny may have been placed there by some agent who made certain it would be on tails, thus setting the probability of all counterfactual property-instances (i.e., heads) to zero beforehand. For this reason and for those given above, information cannot be considered to reside in the coin or in any of its properties.

Unfortunately for the metaphysician, the question is not whether information *qua se* can continue to exist past the event which generated it, for the answer seems clear: it cannot. It can indeed be said that information really *does* exist in the counterfactual-destroying event, and its existence does not depend on whether it is grasped by any cognizer. However, for information to be “remembered,” some memory of the event and its associated counterfactuals must also be preserved. This requires the introduction of property modifications and, unfortunately, memory.

It is not objectionable to require that the information-generating event modify the properties of some temporal continuant in order for information to be “stored.” The event may without controversy be said to cause the coin to have the property of being on tails, or of ephilosophers.com to have the property of having awarded Quine to Terry. The coin and the web site are things which have properties and persist in some sense over time, so they are capable of mediating these properties for observation after the event has ceased to be. It is more objectionable (at least to the metaphysician who would seek to cleanse his or her analysis of epistemic considerations) to write cognizers into the account, but in the case of information storage no other solution presents itself.

Information can be stored only insofar as the outcome of the information-destroying event (i.e., which property-instance was realized) can be stored along with the counterfactuals it eliminated (i.e., which property-instances were not realized). Thus past information, which in the strictest ontological sense no longer exists, can be known by calling to mind the persistent account of what happened in light of what might have happened. It is worthwhile to note that this becomes in part an epistemic consideration, not a purely metaphysical one. This presents information as a thing which may seem

uncomfortably transitory, but in fact having knowledge of it is no more problematic than having knowledge of what side a coin landed on or what essay you received for free from your favorite web site.

This finally gives us the tools we need to understand the loose and popular language normally used to describe information. Terry might tell us that he received an email that carried the information that he won a free copy of Quine. What this really means is that an essay-choosing process inside ephilosophers.com had the property of having some amount of information. That the information concerned an essay, and the essay was written by Quine and was called “On What There Is,” and that Terry was its winner, and that it occurred on such-a-such a day in November, 2005 is not essential to the *being* of the information. The informational property had by that event is grounded in all of these facts, but the information is not these facts themselves. Information that *s* is *F* is merely information had by the event that brought about *s* being exemplified by *F*. Surely that *s* is *F* may be of note to interested cognizers—perhaps even more so than the information entailed by it—but it is not information itself. Information is an attribute of the event that caused *s* (and not something else) to be exemplified by *F*.

Conclusion

The proposal that information is fundamentally a property of an event suggests that the loose and popular use of the term is, in the strictest possible sense, ontologically misleading. When we speak of contemporary life being saturated by information—having instant access to weather reports, news of plane crashes in other parts of the world, instantaneously updated business sales reports, and so on—what we may actually mean is

that we have an unprecedented access to *facts*, or more modestly, to *knowledge claims*.

These may be the ground of truly informative events, but they are not themselves information. Information is not an interesting fact, but a quality of an event that caused one thing, and not some other possible thing, to be.

This should not make us uneasy in the continued use of the term in its loose and popular sense. Common language frequently diverges from strict and philosophical language with no proper harm to understanding or public virtue. However, this proposal is submitted in the hope that it is of service both to further metaphysical analysis of information and in consideration of its broader role in contemporary human affairs. A careful ontological assay is surely not the end of this investigation, but it should provide a good beginning.