Luca Possati

The Paradox of Quantum Information and Its Ethical Consequences

1. Introduction

> The word "information" has been given different meanings by various writers in the general field of information theory. It is likely that at least a number of these will prove sufficiently useful in certain applications to deserve further study and permanent recognition. *It is hardly to be expected that a single concept of information would satisfactorily account for the numerous possible applications of this general field*. (Shannon 1993, 180; emphasis mine)

Will quantum technologies change the way we understand information? This paper is an attempt to answer this question. The central thesis of the paper is that the measurement problem in quantum mechanics prevents the use of a semantic information theory for quantum computing and information (QI), and this has deep technological and ethical consequences. Here I am interested in seeing whether Floridi's theory of information can be appropriately extended or adapted to the quantum domain. I claim that it doesn't seem possible for Floridi's theory to be applied to quantum systems because it requires a process of validation whose requirements are incompatible with quantum mechanics.

The second section of the paper analyzes Floridi's theory of semantic information. This theory is neither the only nor the best possible theory of semantic information. Rather, I see Floridi's semantic information theory as a reference model through which to explore the transformation of the concept of information in a quantum context. The third section analyzes the problem of semantic information in QI. This section is divided into two parts. The first gives a non-mathematical description of the measurement problem. The second shows how the measurement problem exists in the intersection between computation and information. In the quantum field, the veridicality thesis, that is, the core of Floridi's approach, fails because there isn't full accessibility to the system. *If*, following Floridi, we define factual semantic information as a set of data that is (a) well-formed, (b) meaningful, and (c) truthful, *then* in QI we can only get (a) and (b), but we cannot get (c). The third and fourth sections show the consequences of this negative result.

Scenari, n. 17, 2022 • Mimesis Edizioni, Milano-Udine Web: www.mimesisjournals.com/ojs/index.php/scenari • ISSN (online): 2785-3020 • ISBN: 9791222301167 • DOI: 10.7413/24208914137

2. The Theory of Strongly Semantic Information

In this section, I intend to analyze some aspects of Floridi's theory of semantic information (Floridi 2011, 2019), considering it one of the main models of information philosophy available today. The goal of this section is not to develop a history of information theories that is as comprehensive and complete as possible, nor to defend Floridi's theses, nor to present an original theory of semantic information. The theory of semantic information in Floridi is only a starting point for reflecting on the concept of semantic information in QI.

Generally, by information we mean what is transmitted by messages carried through physical channels. The first to apply this notion in a quantitative way were Wiener (1948) and Shannon (1993; see also Shannon and Weaver 1949). Both thought that information was connected to the reduction of uncertainty or entropy. Therefore, any content is more or less informative in relation to the type and amount of uncertainty it reduces. Shannon's theory also inspired philosophical theories of information (Dretske 1981, Floridi 2011; see Adriaans 2013 for a brief review).

We can distinguish between theories of non-semantic information and theories of semantic information. Shannon's theory is about non-semantic information. It excludes any reference to semantics; a message is simply any physical structure (for example, a string of signs) connected to a certain probability. Informativeness is connected to the greater or lesser probability of the structure. Theories of semantic information, on the other hand, deal with the reduction of a different uncertainty, which is not formal but concerns a certain state of affairs: "We call semantic information the information a signal carries by reducing uncertainty about some state of affairs. In this case, semantic aspects are crucial: what information the signal carries is constitutively related to what the signal stands for" (Piccinini and Scarantino 2010, 241).

A natural starting point for introducing the notion of semantic information is Carnap and Bar-Hillel's paper "An Outline of a Theory of Semantic Information" (1953). Carnap and Bar-Hillel criticize and extend the theory of mathematical information developed by Claude Shannon. Floridi criticizes this model and proposes a theory of strongly semantic information (TSSI). His central thesis is that the TSSI can resolve the paradox introduced by Carnap and Bar-Hillel.

Carnap and Bar-Hillel (1953) calculate the amount of informativity (i.e., how informative some information is) encoded in a sentence of a particular language. In their case, the language in question is monadic predicate logic. The philosophical details are grounded on an idea that has come to be known as the *inverse range principle* (IRP). The IRP states that the amount of information encoded by a sentence is inversely proportional to the likelihood of the truth of that sentence. In other words, there is an inversely proportional relationship between the probability and the informativeness of a sentence. The less likely a sentence is to be true, the more informative it is.

In a nutshell, Bar-Hillel and Carnap have developed a poor concept of semantic informativeness; people may learn or be informed about necessary truths. Bar-Hillel and Carnap measure of informativeness does not capture that notion.

Bar-Hillel and Carnap's position implies two paradoxes. The first is the Bar-Hillel-Carnap paradox (BCP), according to which contradictions are the most informative sentences because they are impossible, that is, not probable at all. "There is more information in a contradiction than in a contingently true statement" (Floridi 2011, 109). On the other hand, logical truths are the least informative sentences because they are obvious – as Wittgenstein argued: "All the propositions of logic say the same thing, viz nothing. They are tautologies" (*Tractatus*, 4.46, 6.1). According to Floridi, this idea is paradoxical because it clashes with our intuitive conception of information that information is something true – it cannot be false: contradictions are always false, so they cannot be information. In other words, the thesis of Bar-Hillel and Carnap imposes a rigid separation between the truth and the informativeness of a sentence. The more informative a sentence is, the less likely it is to be true it is. This is paradoxical.

The second paradox is called the deduction scandal (SOD) and is a direct consequence of BCP. The SOD states that an inference does not provide any additional information to our knowledge. "Since tautologies carry no information at all, no logical inference can yield an increase of information, so logical deductions, which can be analysed in terms of tautological processes, also fail to provide any information" (Floridi 2011, 130). This means that the information carried by the conclusion of the deduction must be already contained in the information carried by (the conjunction of) the premises. Logic and mathematics turn out to be utterly uninformative. Sequoiah-Grayson (2008) have shown that Hintikka's (1973) attempts to resolve this paradox (through the distinction between superficial and deep information) fail.

Floridi (2011) proposes to solve BCP and SOD through an alethic approach to information, that is, a theory of strongly semantic information (TSSI). Floridi's central thesis is that only an approach to information that connects information and truth solves the paradoxes and is closer to our intuitive understanding of information. Therefore, semantic information is defined as a set of well-formed, meaningful, and truthful data. This is the formal definition (Floridi 2011, 84):

σ (an infon) is an instance of semantic information if and only if:

1. σ consists of n data (d), for $n \geq 1$;
2. the data are well-formed (*wfd*);
3. the *wfd* are meaningful;
4. the *d* are truthful.

It is important to briefly note two key aspects of Floridi's information theory from a philosophical point of view: (a) the lack of any reference to the concept of intentionality and, therefore, to the human subject; (b) the connection to the communication process, in which information is always something communicated, transmitted. The first aspect implies that information can also be generated, distributed, and stored by non-human agents. The second suggests a connection to media theory.

Now some brief clarification about the formal definition:

– "well-formed" means that the data are clustered together following the rules that govern the chosen system, code, or language being analyzed. Syntax here must be understood "broadly (not just linguistically), as what determines the form, construction, composition, or structuring of something" (Floridi 2011, 84);

– "meaningful" means that the data must comply with the meanings of the chosen system, code, or language in question. In this case, "let us not forget that semantic information is not necessarily linguistic. For example, in a map, the illustrations are such as to be visually meaningful to the reader" (Floridi 2011, 84). In other terms, "meaningful" means that the data must also be representative, that is, they must concern the object of the information, i.e., what the information is about.

– Floridi's formal definition establishes that factual semantic information *encapsulates truth*: Matthew is informed that milk contains calcium if and only if Matthew holds that milk contains calcium and it is true that it does. Floridi justifies this thesis through (a) the action-based semantic theory, which explains how data acquires meaning and interpretation, (b) the correctness theory of truth, which explains how well-formed and meaningful data may become truthful, and (c) the logic of being informed. The core of his strategy is the so-called "veridicality thesis," which can be summarized as follows: "the quantity of strongly semantic information in a proposition p is calculated in terms of the distance of p from a situation z (where situations are partial or incomplete worlds) that p is supposed to model" (Sequoiah-Grayson and Floridi 2022). From this point of view, false information and tautologies are not information (Scarantino and Piccinini [2010] criticized Floridi's argument that false information is not information).

It is important to briefly note two key aspects of Floridi's information theory from a philosophical point of view: (a) the lack of any reference to the concept of intentionality and, therefore, to the human subject; (b) the connection to the communication process, in which information is always something communicated, transmitted. The first aspect implies that information can also be generated, distributed, and stored by non-human agents. The second suggests a connection to media theory.

3. Quantum Mechanics and Semantic Information

The problem of measurement

Let us first try to give an elementary description of the problem of measurement (Carroll 2020, 26). The wave function in quantum mechanics evolves deterministically according to the Schrödinger equation as a linear superposition of several states. However, the actual measurement of it always finds the physical system in a defined state. Any future evolution of the wave function is based on the state in which the system is found when the measurement is made, which means that the measurement has changed the system and that this change is not obviously a consequence of evolution according to Schrödinger. The problem with measurement is describing what that change is, that is, how a superposition of many possible values becomes a single measured value. In other words (Weinberg 1998, 2005), the Schrödinger wave equation determines the wave function in the times following the measurement. If the observers and their measuring apparatus are themselves described by a deterministic wave function, why can we not predict the precise results of the measurements but only probabilities?

As Hagar (2003) explains, it can be said that at the core of the measurement problem lies the mutual inconsistency of the three following claims:

1) The wave function of a system is complete; the wave function specifies all the physical properties of a system.

2) The wave function always evolves in accordance with a linear dynamical equation (the Schrödinger equation).

3) Measurements of, for example, the spin of an electron always (or at least usually) have determinate outcomes, that is, at the end of the measurement, the measuring device is either in a state that indicates spin up (and not down) or spin down (and not up).

A caveat is necessary. There are many measurement techniques in QI (perfect measurements, erroneous measurements, measurements perturbed by the environment, etc.), and we cannot cover them all. Most devices capable of detecting a single particle and measuring its position strongly modify the particle's state in the measurement process; for example, photons are destroyed when striking a screen. Less dramatically, the measurement may simply perturb the particle in an unpredictable way; a second measurement, no matter how quickly after the first, is then not guaranteed to find the particle in the same location. A "quantum nondemolition measurement" can be used in which the uncertainty of the measured observable does not increase from its measured value during the subsequent normal evolution of the system (Yoshikawa et al. 2008). However, none of these measurement techniques avoid what physicists call "the collapse of the wave function," that is, the third statement.

Let us now examine two aspects of the measurement problem. The first concerns the status of the measurement itself. What does it mean to "measure" a quantum state? When and how do we apply Born's rule to extract probabilities? Now, a measurement is not just any interaction with a physical system, "but an interaction so designed as to yield information about *features the system had antecedently to the interaction*" (Maudlin 2019, 96). For example, when I measure my weight by climbing on a scale, I consider the measurement result to be something related to my pre-existing weight. Thus, when we measure the position of an electron on a screen, the measure (the sign that indicates its position) would suggest that the electron had that position *prior* to the measurement. However, this is exactly the problem: "We know by observation that a mark was created. We do not know by observation whether it corresponds to an antecedently existing position of the electron" (Maudlin 2019, 97). Born's rule can be used to calculate a probability for each outcome, but it does not specify how or when the mark was made on the phosphorescent screen.

The second problem concerns the "wave function collapse." In quantum mechanics, the collapse occurs when, during a measurement, a wave function – initially in a superposition of several states – is reduced to a single state due to interaction with the external world. There are many different interpretations of this phenomenon, but this is not the place to examine them (for more technical explanations, see Vermaas 1999, Chapters 10 and 14; Maudlin 2019, Chapter 4; Ney and Albert 2013). Our problem is that:

a) In quantum mechanics, the nature and effects of collapse are still poorly understood (Maudlin 2019, 75).

b) All information about the (probability) amplitudes "is destroyed upon measurement" (Grumbling and Horowitz 2019, 71). Measurement fundamentally disrupts a quantum state: "it 'collapses' the aspect of wave function that was measured into a single observable state, resulting *in a loss of data*" (Grumbling and Horowitz 2019, 57; emphasis mine). In QI, we know the results of measurements, which are probability distributions on the system. However, we have no direct knowledge of the evolution of the process, just several possible interpretations. In other words, we deal with not observable, or noumenal, states of affairs (Vermaas 1999, 211) as a "thing known by the mind as against the senses." In other terms, in quantum mechanics, observations and measurements made on systems do not give immediate knowledge of those systems. We need an interpretation, that is, a model, to *infer* something about the state of a system from observations and measurements.

Now, if we examine some important QI textbooks (De Lima Marquezio et al. 2019; Hughes et al. 2021; Rieffel and Polak 2011; Kaye et al. 2007; Kurgalina and Borzunov 2021), the problem of measurement and collapse are only hinted at, or they are not dealt with at all. The measurement of the qubit is reduced to a mathematical formalism to be applied to the algorithmic process, but it does not pose a problem in itself. For example, Kaye et al. (2007, 55) state that "given a single instance of an unknown single-qubit state, there is no way to determine experimentally what state it is in; we cannot directly observe the quantum state. It is only the results of measurements that we can directly observe" and then introduce mathematical formalism to deal with the measurement.

The measurement does not imply anything mysterious. It shows us the probabilistic and indeterminate nature of quantum systems. In our case, we need to understand how the measurement problem impact on the veridicality thesis.

The veridicality thesis

According to Floridi, semantic information is characterized by the so-called "truth encapsulation." This point is justified through the correctness theory of truth, which explains how well-formed and meaningful data may become truthful. The veridicality thesis has also been subscribed to by Dretske (1981) and Graham (1999). This notion has been heavily criticized (Colburn 2000; Fetzer 2004; Piccinini and Scarantino 2010). I do not want to deepen the debate here – that is not the purpose of this paper. However, I believe that all these criticisms affect Floridi's thesis only to some extent so far. Just a brief consideration: Floridi does not speak of "semantic information" in general, but of "semantic information functional for epistemic purpose." This means two things: (a) obviously, there are propositions that may be false and informative, but which do not have an epistemic value; b) semantic information with epistemic value is always an interface between an agent that processes data and models related to a system that were developed in order to understand that system better and better. The veridicality thesis is therefore based on three crucial factors: (a) the need to avoid the BCP, (b) the justification of the phenomenon of semantic erosion, and (c) the possibility of connecting information and knowledge. To invalidate the thesis, it is necessary to eliminate all of these factors, but the price to pay for doing so is very high (Floridi 2007).

My claim is that QI undermines the veridicality thesis and forces us to reformulate the theory of semantic information. The core of the veridicality thesis is the process of validation. Floridi (2011) explains this process through two key concepts: interaction and accessibility. Both presuppose one element: the informee, that is, the agent who seeks information. Floridi writes:

> The sort of accessibility at stake here is a matter of pragmatic or factual interaction, which provides an exogenous grounding of correctness. It is the one that we find specified in computer science, where accessibility refers to the actual permission to read (technically, sense and retrieve) and/or write (again, technically modify and record) data as a physical process. The result is that *a*'s proximal access to *m* commutes with *a*'s distal access to *s* if and only if *a* can read/write *s* by reading/writing *m*. (Floridi 2011, 197)

Validation, that is, the examination of correctness, takes place in two steps. First, the informee has immediate access (writes/reads, modifies) to the model and thus has remote access (writes/reads, modifies) to the system, and vice versa. By modifying the model, the informee modifies the system, just as by modifying the system, the model is modified. The scope is the increasing adequacy of the model to the system. Therefore, *x* is true semantic information if and only if it is fully verified and validated, that is, if and only if (a) it meets the criteria of the inquiry (the context, the level of abstraction, and the purpose), and (b) it is better than other models we can develop based on the data. In other words, the validation process implies the *traceability* of the constitution of data and information from the model to the system and from the system to the model.

In QI, the accessibility relationship fails. Floridi (2011, 196) explains the accessibility relationship through the concept of a "commutative diagram." In category theory, a commutative diagram is a diagram in which all directed paths with the same start and end points lead to the same result. In other words, the accessibility relationship to the model and the system is commutative. Proceeding from the model to the system is equivalent to proceeding from the system to the model. In QI, on the other hand, proceeding from the model to the system (i.e., measuring the qubit, then reading/writing data) is not the same as proceeding from the system to the model. A commutative relationship is not possible because the measurement transforms the system – we cannot know the system regardless of the measurement. We can only go from the model to the system in the act of measuring. However, the measurement eliminates all the data on the evolution of the system before the measurement itself – it transforms the system. This situation does not arise when the information concerns a classical physical system.

Therefore, we face a paradoxical situation. On the one hand, QI gives us good semantic information functional for epistemic purposes. A quantum algorithm (e.g., Shor's algorithm, Grover's algorithm, etc.) allows us to solve very complex problems and gives us true and epistemically relevant information. On the other hand, the same algorithm does not allow a full validation of this information because of the lack of accessibility. We can go from the model to the system, but not from the system to the model. From an engineering point of view, this works great. From the point of view of semantic information theory, the traceability of data and processes is not possible.

In summary, *if* we define factual semantic information as a set of data that is (a) well-formed, (b) meaningful, and (c) truthful, *then* in QI we can only get (a) and (b), but we cannot get (c).

4. The Ethical Consequences of the Paradox

The measurement problem does not affect QI from a technical point of view; the results of the computation work, and this is enough from an engineering point of view. The measurement problem becomes instead a very serious obstacle on the ethical level because it makes processes opaque and prevents the justification of the behavior of a quantum algorithm. Transparency is not an "ethical principle in itself but a pro-ethical condition for enabling or impairing other ethical practices or principles" (Floridi and Turilli 2009, 105). Transparency and explainability are particularly important in QI. As is noted in the Australian strategy for the quantum revolution, "a precondition for the social debate about quantum technology is that all participants have a reasonable understanding of the technology and its implications" (Brennen et al. 2021, 10).

I draw on the concept of opacity from the AI literature:

> Opacity refers to the epistemic barrier between the engineer and its creation. The created systems will become so complex, and self-modifying (via machine learning) and the engineering teams are so big, that no single person can comprehend it fully. Hence, there is *a lack of visibility*, or there is

> the presence of opacity in a sense from a human point of view. This is despite the fact that in the case of software, details are knowable down to the last bit. (Héder 2020, 3; emphasis mine)

Opacity in AI is an epistemic concept, not an ontological one, in the sense that it concerns the ability of engineers to know and explain the behavior of a system. The behavior is knowable in itself, except that the level of complexity is too high for a single person to fully analyze it. Factors contributing to the overall lack of algorithmic transparency include the cognitive impossibility of humans interpreting massive algorithmic models and datasets; a lack of appropriate tools to visualize and track large volumes of code and data; code and data that are so poorly structured that they are impossible to read; and ongoing updates and human influence over models (Tsamados et al. 2021). In the case of neural networks, humans know what AI is doing (for example, they know the code of the AI and know how it works in general), "but in another sense they don't know (they cannot explain a particular decision), with the result that people affected by the AI cannot be given precise information about what made the machine arrive at its prediction" (Coekelbergh 2021, 117). The impossibility of opening the "black box" is a practical and cognitive limit, that is, it is linked to the limits of human capabilities.

Now, Héder (2020) points to five main barriers to transparency in AI systems:

1. The emergent behavior of machines. (In the upper layer, at any point in the architecture, things happen that are within the boundaries set by the lower level but not entirely governed by them. The exact interplay of the two layers cannot be predicted.)
2. The embodiment effects (the issues that can arise from the material realization of an artifact and that cannot be predicted).
3. Material layer effects on hardware. (The effects of the dissipation of heat between the layers of the system, for example, can affect the whole performance of the system.)
4. Statistical knowledge (the environment's unpredictability + the probabilities generated from past experiences distilled from the training data).
5. Human factors (the unpredictable reactions of the humans, e.g., malevolent acts against the system, and the cognitive impossibilities).

Quantum computing amplifies all these opacity factors and thus makes it even more difficult to solve the problem of opacity and trust. What constitutes a good explanation in QI? What are the differences between explanations and reasons, and can machines provide any of these? Here the risk is not only of manipulation and domination of the technology by capitalists or technocratic elites, creating a highly divided society – the so-called "quantum divide" (de Wolf 2017). The deeper risk that looms here is a high-tech society in which even those elites no longer know what they are doing, and in which nobody is accountable for what is happening in the system. If in AI the "black box" might be opened in theory (Coeckelbergh 2021, 120), in QI this possibility does not exist. If in AI the lack of transparency is a practical fact, *in QI it is an ontological fact*.

The measurement problem makes the traceability of the algorithm's behavior problematic, that is, the understanding of why a quantum algorithm decided in a particular way and not in another. This problem arises on three levels:

a) a technical level, that is, the impossibility, due to the measurement problem, of knowing directly and justifying the evolution and decisions of the quantum algorithm;
b) an interpretative level, that is, the difficulty of interpreting data and constructing semantic information; for example, see all the different interpretations of the collapse;
c) a communicative level, that is, the difficulty of communicating the why and how of the decisions of a quantum algorithm to an audience of non-experts, for example, politicians, entrepreneurs, or citizens, due to the general counter-intuitiveness of quantum mechanics.

Artificial intelligence ethics can provide some help in formulating design solutions suitable for QI. Making AI systems explainable and transparent is an important ethical principle (Watson and Floridi 2020; Gultchin et al. 2021). The explanation of a decision in AI must always be contextualized in such a way as to be adequate and protect the autonomy of the recipient of the explanation. Furthermore, explanation and communication in an AI system builds trust and fosters cooperation in the system. It has been shown that an increase in the complexity and opacity of a system corresponds to an increase in the de-responsibility of the human beings who use it (Floridi 2022, 184–186; Buchholz et al. 2020). This can easily lead to the spread of crimes. However, the explanation and communication must be established in correct terms. According to Wachter et al. (2017), they must concern (a) the functioning, (b) the logic used, (c) the reasons why a certain decision was made, and (d) the objective of the system. Nevertheless, it must also be remembered that transparency is not always the best choice; there could be very good reasons why AI system designers decide not to inform users of the goals or logic of their software (e.g., security reasons or the need to preserve the scientific value of a project). For this reason, it is essential for the design of an

AI system to assess what level of transparency (what type, for whom, and on what) the system intends to adopt in relation to its overall goal and the context of the implementation.

In QI, these tasks are even more essential. Designers must therefore provide rational and persuasive arguments in line with the general purpose of the system and suitable for the various types of stakeholders (customers, institutions, universities, companies, etc.). The internal communication of the system between the different components becomes essential. Furthermore, given the complex nature of quantum computational processes, the interpretation of the results acquires an even more decisive role. For these reasons, it is essential to involve physicists and philosophers of physics in the public debate on the social impact of quantum technologies together with stakeholders and other actors (Vermaas 2017). Designers should therefore create graphical interfaces capable of clarifying as much as possible the quality of the data used and the nature of the processes carried out in QI systems. In this case, data visualization (Rodighiero 2021) can play a key role. How to visualize QI is a great challenge. It is an essential condition for a transparent use of QI, and therefore also for the design of QI systems capable of transmitting and protecting values in relation to a specific social context.

5. Conclusions

The aim of this paper was not to increase the sense of strangeness around quantum theory and quantum technologies. As Carroll (2020) claims, and I agree with him, "quantum mechanics is unique among physical theories in drawing an apparent distinction between *what we see* and *what really is*" but "it isn't hopelessly mystical or inexplicable. It's just physics" (17). I think that there is a profound connection between the lack of clarity on the foundation of quantum mechanics and the difficulty of popularizing this theory.

In this paper I have provided several arguments demonstrating that the measurement problem in quantum theory makes a semantic information theory for QI more challenging – not impossible.

As I claimed, the problem of the construction of semantic information in quantum technologies is not simply a theoretical problem. It also directly affects how we attribute meaning to quantum technologies and thus the expectations and fears associated with them. Perhaps there is a connection between the opacity of quantum technologies and the absence of great narratives about them. The emergence of quantum technologies is not accompanied, at first glance, by the development of powerful narratives such as those we have seen for AI or nanotechnologies

(Grunwald 2014). From this point of view, quantum technologies pose a hermeneutic problem, that is, the hermeneutic appropriation of quantum technologies.

The difficulty of building a quantum semantic information theory also has an impact on software engineering, that is, the development of suitable software for quantum computer. In fact, software is not just a set of algorithms. Software is primarily a translation process. Software translates a problem and a solution to that problem from the plane of human understanding to that of the machine. Our computers can accomplish certain tasks because we have succeeded in translating the way of accomplishing those tasks into binary strings understandable to a Turing machine (Possati 2022). A semantic information theory is essential in this undertaking.

There are two limitations to my thesis. The first concerns the information theory model to which I have referred, namely Floridi's philosophy of information. There are obviously other models, and they could give different results. However, I believe that Floridi's model is a good model that is useful for identifying and analyzing a problem. The second limitation relates to the purpose of this paper. Highlighting the measurement problem in quantum mechanics and its impact on QI does not mean that I want to emphasize the enigmatic nature of quantum phenomena and so compromise the public debate. I think that if we want to develop the public debate on quantum technologies, we do need to make these technologies more transparent, and to do so, we need to identify the problems that prevent this transparency on ethical levels. For this reason, physicists and philosophers of physics are an essential part of this endeavor.

References

Adriaans, P. 2013. "Information." In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Fall 2013 Edition). http://plato.stanford.edu/archives/fall2013/entries/information/

Brennen, G., Devitt, S., Roberson, T., Rohde, P. *An Australian Strategy for the Quantum Revolution*. Available online: https://s3-ap-southeast-2.amazonaws.com/ad-aspi/2021-05/Quantum%20revolution-v2.pdf?VersionId=tST6Nx6Z0FEIbFDFXVZ2bxrGp2X8d.iL

Bub, C. 2005. "Quantum information and computation" *arXiv*:quant-ph/0512125

Bar-Hillel, Y., and Carnap, R. 1953. "An Outline of a Theory of Semantic Information." In Bar-Hillel, Y. 1964. *Language and Information: Selected Essays on Their Theory and Application*. Reading, Mass.; London: Addison-Wesley, pp. 221-274.

Carroll, S. 2020. *Something deeply hidden: Quantum world and the emergence of spacetime*. New York: Dutton.

Coekelbergh, M. 2021. *AI ethics*. Cambridge, MA: MIT Press.
Colburn, T. R. 2000. "Information, thought, and knowledge" *Proceedings of the world multiconference on systemics, cybernetics and informatics*, pp. 467–471.
De Lima Marquezino, F., Portugal, R., and C. Lavor. 2021. *A Primer on Quantum Computing*. Dordrecht: Springer.
De Wolf, R. 2017. "The potential impact of quantum computers on society" *Ethics Inf Technol* 19: 271–276.
Dretske, F. I. 1981. *Knowledge and the Flow of Information*. Oxford: Blackwell.
Fetzer, J. H. 2004. "Information: Does it have to be true?" *Minds and Machines* 14(2): 223–229.
Floridi, L. 2007. "In defense of the veridical nature of semantic information" *Eujap* 3(1): 32-41.
Floridi, L. 2011. *The philosophy of information*. Oxford University Press.
Floridi, L. 2019. *The logic of information*. Oxford University Press.
Floridi, L. 2022. Etica dell'intelligenza artificiale. Milano: Raffaello Cortina.
Floridi, L., and M. Turilli. 2009. "The ethics of information transparency" *Ethics and Information Technology* 11 (2): 105-112.
Fuchs, C., and A. Peresh. 2001. "Quantum Theory Needs No Interpretation" *Physics Today* 53(3): 70-71.
Fuchs, C. 2002. "Quantum Mechanics as Quantum Information (and only a little more)" arXiv; https://arxiv.org/abs/quant-ph/0205039
Grumbling, E., and Horowitz, M. (eds) 2019. *Quantum Computing Progress and Prospects.* Washington, DC: The National Academies Press.
Hagar, A. 2003. "A Philosopher Looks at Quantum Information Theory" *Philosophy of Science* 70 (4): 752-775.
Hagar, A., Hemmo, M. 2006. "Explaining the *Un*observed – Why Quantum Mechanics Ain't Only About Information" *Found Phys* 36: 1295–1324.
Héder, M. 2020. "The epistemic opacity of autonomous systems and the ethical consequences" *AI & Society* https://doi.org/10.1007/s00146-020-01024-9
Hintikka, J. 1973. *Logic, Language-Games and Information: Kantian Themes in the Philosophy of Logic*. Oxford: Clarendon Press.
Hughes, C., Isaacson, J., Perry, A., Sun R.F., Turner, J. 2021. *Quantum Computing for the Quantum Curious.* Dordrecht: Springer.
Lundgren, B. 2019. "Does semantic information need to be truthful?" *Synthese* 196: 2885–2906. https://doi.org/10.1007/s11229-017-1587-5
Macaulay Ferguson, T. 2015. "Two paradoxes of semantic information" *Synthese* 192(11): 3719–3730.
Kaye, P., Laflamme, R., Mosca, M. 2007. *An introduction to quantum computing*. Oxford University Press.
Kurgalin, S., and S. Borzunov. 2021. *Concise guide to quantum computing*. Dordrecht: Springer.
Maudlin, T. 2013. "The nature of the quantum state." In A. Ney and D. Z. Albert (ed), *The Wave function. Essays on the metaphysics of quantum mechanics*. Oxford University Press, pp. 126-153.
Maudlin, T. 2019. *Philosophy of Physics: Quantum Theory*. Princeton University Press.

Ney, A., and D. Z. Albert (ed). 2013. *The Wave function. Essays on the metaphysics of quantum mechanics*. Oxford University Press.
Piccinini, G., and A. Scarantino. 2010. "Computation vs information processing: why their difference matters to cognitive science" *Studies in history and philosophy of science* 41: 237-246.
Perrier, E. 2022. "Ethical quantum computing: A roadmap" *arXiv*, arXiv:2102.00759.
Possati, L. 2022. *Software as hermeneutics*. Dordrecht: Springer.
Rieffel, E., and W. Polak. 2011. *Quantum computing. A gentle introduction*. Cambridge, MA: MIT Press.
Rodighiero, D. 2021. *Mapping Affinities: Democratizing Data Visualization*. English edition. Geneva: Métis Presses
Sequoiah-Grayson, S. 2008. "The scandal of deduction: Hintikka on the Information Yield of Deductive Inferences" *Journal of philosophical logic* (37): 67-94.
Shannon, C. E. 1993. *Collected Papers.* ed. by N. J. A. Sloane and A. D. Wyner. New York: IEEE Press.
Shannon, C. E., and Weaver, W. 1949 rep. 1998. *The Mathematical Theory of Communication.* Urbana: University of Illinois Press.
Tsamados, A., Aggarwal, N., Cowls, J. 2022. "The ethics of algorithms: key problems and solutions" *AI & Society* 37: 215–230.
Vermaas, P. E. 1999. *A Philosopher's Understanding of Quantum Mechanics. Possibilities and Impossibilities of a Modal Interpretation*. Cambridge University Press.
Vermaas, P.E. The societal impact of the emerging quantum technologies: a renewed urgency to make quantum theory understandable. *Ethics Inf Technol* 19, 241–246 (2017)
Yoshikawa, J., Miwa, Y., Huck, A., Andersen, U., Van Loock, P., and A. Furusawa. 2008. "Demonstration of a Quantum Nondemolition Sum Gate" *Phys. Rev. Lett*. 101: 25050. https://doi.org/10.1103/PhysRevLett.101.250501
Watson, D. S., and L. Floridi. 2020, "The explanation game: A formal framework for interpretable machine learning" *Synthese* 20: 1-32."
Weinberg, S. 1998. "The Great Reduction: Physics in the Twentieth Century." In M. Howard & W. R. Louis (eds.). *The Oxford History of the Twentieth Century*. Oxford University Press.
Weinberg, S. 2005. "Einstein's Mistakes" *Physics Today* 58 (11): 31–35.
Wiener, N. 1948. *Cybernetics or Control and Communication in the Animal and the Machine*. 2nd edn. Cambridge, Mass.: MIT Press.
Wootters, W.K., and W.H. Zurek. 1982. "A single quantum cannot be cloned" *Nature* 299: 802-804.

The Paradox of Quantum Information and Its Ethical Consequences

This paper explores the impact of the measurement problem in quantum mechanics on quantum information theory. It intends to see whether Floridi's theory of information can be appropriately extended or adapted to the quantum domain. The claim is that it doesn't seem possible for Floridi's theory to be applied to quantum systems because it requires a process of validation whose requirements are incompatible with quantum mechanics. The second section of the paper analyzes the notion of semantic information through the lens of Floridi's theory of semantic information. The third section examines the problem of semantic information in quantum computing. This section is divided into two parts. The first gives an elementary description of the measurement problem. The second part shows how in the quantum field the veridicality thesis, that is, the core of Floridi's approach, fails. The fourth section investigates the ethical consequences of this thesis. These consequences mainly concern the opacity of quantum algorithms, that is, their lack of transparency and the difficulty of explaining and justifying their decisions.

KEYWORDS: *information, quantum technologies, uncertainty, semantic*