

Article

Information Physics—Towards a New Conception of Physical Reality

Philip Goyal

Department of Physics, University at Albany (SUNY), 1400 Washington Avenue, Albany, NY 12222, USA; E-Mail: pgoyal@albany.edu; Tel.: +1-518-442-2610; Fax: +1-518-442-5260

Received: 10 September 2012 / Accepted: 26 September 2012 / Published: 17 October 2012

Abstract: The concept of information plays a fundamental role in our everyday experience, but is conspicuously absent in framework of classical physics. Over the last century, quantum theory and a series of other developments in physics and related subjects have brought the concept of information and the interface between an agent and the physical world into increasing prominence. As a result, over the last few decades, there has arisen a growing belief amongst many physicists that the concept of information may have a critical role to play in our understanding of the workings of the physical world, both in more deeply understanding existing physical theories and in formulating of new theories. In this paper, I describe the origin of the informational view of physics, illustrate some of the work inspired by this view, and give some indication of its implications for the development of a new conception of physical reality.

Keywords: quantum theory; information; information theory; information physics

1. Introduction

The concept of information plays a fundamental role in our everyday experience. We generally view ourselves as embodied finite beings immersed in a physical world of which we possess limited knowledge. Each of us constantly directs our attention to aspects of the world about which we wish to gain more information, and that information is supplied by our sensory systems in a virtually endless stream and is dynamically shaped by our mental faculties into an ever-changing predictive model of the world. That our predictive model provides us with limited information about the future is evidenced by our constant experience of making predictions (consciously or unconsciously) which—even in the simplest situations which we habitually encounter—rarely precisely coincide with what actually

happens. According to this view, the concept of information enters in two distinct ways: (1) At any moment, the information available to us through our sensory systems about the external world is limited; and (2) the information provided by our predictive model about the future is also limited.

Despite its ubiquity in everyday experience, the concept of information is absent in the mechanical conception of reality that underlies classical physics. This conception, which underpinned the development of physics for over two hundred and fifty years—from the seventeenth-century mechanics of Galileo and Newton to the nineteenth-century electrodynamics of Faraday and Maxwell, articulates a highly-idealized view of reality and our relationship to it. According to this view, in essence, the totality of all that exists in the physical world consists of matter and fields which sit on the fixed stage of space, and which evolve according to universal laws in step with a universal time, with no fundamental limits on how well an ideal agent can know the state of all that exists.

In this mechanical conception of reality, the concept of information is redundant in both of the senses identified above. First, an agent, in principle, has complete information about the state of reality, so there is no sense in maintaining a distinction between the state of reality and an ideal agent's knowledge of that state. Second, in principle, such an agent who is also in possession of a description of the universal laws of motion can exactly predict any future (or past) event, and thus has complete information about the state of reality at all other times, past or future.

Over a period of over two centuries, the unprecedented success of the physical theories built within the classical framework—chiefly Newtonian mechanics and Maxwellian electrodynamics—naturally gave rise to the belief that the assumptions and idealizations underlying that framework were basically sound. In particular, it gave rise to the belief that the issues of limited access to information and limited predictive information which are characteristic of our everyday existence have no bearing on our quest to precisely formalize nature's innermost workings, these issues simply being nuisances which we (as non-ideal agents) must live with. The assumptions and idealizations of the classical framework certainly did not go unchallenged, either within physics community itself or within the wider intellectual culture that had been profoundly affected and challenged by the mechanical world-view spawned by classical physics. Nevertheless, largely owing to the fecundity of the classical framework, these assumptions essentially held sway in the physics community until the early part of the twentieth century.

By the 1890s, classical physics was in a highly developed state. However, there were some clouds on the horizon, curious experimental facts—such as the spectrum of frequencies of light emitted by heated bodies—that stubbornly resisted explanation within the existing theories of classical physics. The explanation of these facts, initiated by Max Planck, would, over the next thirty years, bring about the development of quantum theory—an entirely new theoretical framework for physical theories that would replace the classical framework. Within a few years after its creation, it was clear that the quantum framework departed radically from the framework of classical physics, and was in many respects at odds with the assumptions and idealizations underlying the mechanical conception of reality. In particular, in the quantum framework, an agent has, *in principle*, an informationally-limited view of reality, both of the state of any physical system, and of the outcome of almost all future measurements performed on the system. Hence, as described within the quantum framework, the relationship of an agent to the physical world resembles that in our everyday experience, and, accordingly, the concept of information acquires a non-trivial role.

Since the creation of quantum theory, it has been widely debated whether these non-classical features, particularly the radically different status of the measurement-performing agent, are fundamental, reflecting the very structure of physical reality, and so ought to be regarded as generic features of physical theories, or are only apparent, emergent from some yet-to-be-discovered more fundamental theory closer in nature to the theories of classical physics which lacks most (if not all) of these features.

Over the last few decades, bolstered by many developments in physics (such as black hole physics) and other disciplines (such as Shannon's theory of information), the view has arisen amongst a growing number of physicists that the concept of information may indeed have a fundamental role to play in our understanding of quantum theory and, more generally, the physical world, a role as fundamental as that occupied by such concepts as space, time, matter and energy in classical physics. In recent years, this *informational view* of the physical world has been increasingly fruitful. It has given rise to new disciplines (quantum information and quantum computation) that have shown how quantum theory can be harnessed to store, transmit and compute information in ways that are difficult or impossible according to classical physics, and has inspired new efforts to derive some of the predictions of quantum theory—and even the mathematical structure of quantum theory (long regarded as enigmatic) itself—from informationally-inspired physical postulates.

The purpose of this paper is to describe the origin of the informational view of physics, to illustrate some of the recent work inspired by this view, and to give some indication of its broader implications for the development of a coherent new conception of physical reality that is capable of providing a viable alternative to the mechanical conception of classical physics and of guiding the future development of physics.

The remainder of this paper is organized as follows. In Section 2, I provide an overview of quantum theory, using the concept of information to illuminate some of the key differences between quantum and classical physics. In Section 3, I give a brief overview of the developments in physics and related areas which have contributed to the emergence of the informational view. Then, in Section 4, I describe in particular how the informational point of view has reinvigorated the long-standing quest to understand quantum theory by providing new starting points from which to derive the mathematical formalism of quantum theory. Finally, in Section 5, I conclude with some observations on how the efforts to derive the quantum formalism from informational postulates, and the informational view more generally, is contributing to the development of a new conception of physical reality.

2. The Role of Information in Quantum Physics

In order to understand the role of information in quantum theory, it is helpful to first summarize the essential ideas underlying classical physics.

2.1. Classical Physics

Classical physics can be usefully divided into three components, namely:

1. A conception of the physical world, namely the mechanical conception of reality;
2. A precise conceptual and mathematical framework, which formalizes this mechanical conception;

3. Classical physical theories—in particular Newtonian mechanics and Maxwellian electrodynamics—which are built within this framework.

As mentioned in the introduction, classical physics is underpinned by a mechanical conception of reality. According to this conception, the physical universe is a vast machine whose state at any time is precisely describable in every detail, and which evolves in time according to quantitatively precise laws. These laws operate in the same manner at all times and at all places within the universe, and yet themselves remain unchanged. Furthermore, according to this conception, ideal agents can make measurements to learn about the state of the universe as precisely as they wish without disturbing it, and can build up a knowledge of that state by aggregating knowledge gained about spatially-disjoint regions of the universe.

The core of the mechanical conception is formalized in the framework of classical physics, which provides the conceptual and mathematical framework for the theories of classical physics. In this framework, one speaks abstractly about a *physical system*, its *dynamics*, and *measurements* performed upon it. The physical system may be the physical universe as a whole, but is usually some subset thereof. At any given moment in time, t , the *physical state* of a physical system is described by a mathematical object, S , simply called the *state*. The state S completely describes the physical state insofar as predictions of the outcomes of all possible measurements performed on the system are concerned. The state is a n -tuple of real numbers since it is assumed that the properties it encodes (which are in turn revealed by performing measurements on the system) have a continuum of possible values. The *state space* of the system is the set of all possible states of the system. The temporal evolution of the state of the system during the interval $[t, t + \Delta t]$ is represented by a continuous bijective map, $M_t(\Delta t)$, over state space.

Every measurement is assumed to yield an outcome which is determined by the state of the system, but without affecting the state of the system itself. Accordingly, such a measurement is represented by a map from state space to a space of possible outcomes. In principle, there exists an informationally-complete measurement from whose outcome it is possible to infer the state of the system. Such a measurement is represented by a bijective map, so that there is a one-to-one correspondence between states of the system and outcomes of the measurement. An ideal agent, who by definition can perform such an informationally-complete measurement, will have complete knowledge of the physical state of a system. On the basis of such a measurement, an ideal agent would therefore be able to assert that, with certainty, the system is in a definite state, also known as a *pure* state. In this case, then, there is no distinction between the state of the system on the one hand, and the ideal agent's knowledge of that system on the other [1].

Finally, if a system is composed of sub-systems, the state of the system is determined by the states of its sub-systems. This means that agents can build up knowledge of a system by aggregating knowledge about its sub-systems.

A particular classical physical theory is built within this framework by specifying the state, S , and the dynamical map $M_t(\Delta t)$. For example, in Newton's theory of gravity, matter is composed of mutually-gravitating particles represented by geometrical points in Euclidean space, moving in geometrically-precise trajectories. The state of a system of N such particles is given by a list of the real-valued positions and velocities of these N particles, and the dynamical map is obtained from

Newton's laws of motion by taking into account the gravitational forces between these particles via Newton's law of gravitation. In principle, there exists an informationally-complete measurement whose outcome is simply a read-out of the positions and velocities of the N particles that constitute the physical system.

2.2. Quantum Physics

In the framework of quantum physics, just as in the classical framework, a model of a physical system consists of (a) a complete theoretical description (the mathematical *state*, or simply the *state*) of its physical state; (b) a mapping from the space of states to itself, which represents the dynamics; and (c) a model of the process of measurement. In this framework, an ideal agent who has prepared a quantum system using a repeatable measurement is able to assign a definite state (a *pure state*) to the system immediately following the measurement.

For a composite system composed of subsystems, the framework also specifies the pure state of the composite system when the subsystems are in known pure states. However, in quantum theory a composite system can be in a pure state even when the subsystems are not.

In the quantum framework, just as in the classical framework, the dynamical mapping is bijective, representing deterministic and reversible temporal evolution. However, the quantum and classical frameworks differ sharply in two ways—in their models of the measurement process and in their descriptions of composite systems.

2.2.1. Quantum Model of the Measurement Process

As described above, in the classical framework, an agent can perfectly and completely learn the state of a physical system, and furthermore can do so without disturbing the system to any significant degree. In contrast, the quantum framework posits a model of the measurement process that has four distinct, non-classical features:

1. *Discreteness.* The number of possible outcomes of a measurement may be finite or countably infinite;
2. *Probabilistic Outcomes.* The outcome of a measurement performed on a physical system is only predictable on a probabilistic level;
3. *Disturbance.* A measurement almost invariably changes the state of the system upon which it is performed;
4. *Complementarity.* A measurement only yields information about some of the parameters needed to specify the state of the system, at the expense of the others.

The first of these features, discreteness, is that, when certain measurements are performed on physical systems, the number of possible outcomes can be finite or countably infinite, where this discreteness is not due to boundary conditions but is intrinsic to the system. This stands in contrast with the classical assumption that all physical quantities (such as the position of a particle) can take a continuum of possible values. Hence, discreteness challenges the classical idea that the continua of space and time are the fundamental bedrock of physical reality.

The second of these features is that, in contradistinction to one of the basic tenets of the mechanical conception of reality, it is *not* true that every detail of every event is determined by universal laws. Instead, the quantum framework asserts that only the *probability* that a measurement will yield a particular result is predictable; the outcome that will be obtained in a particular run of an experiment, in general, is not.

The quantum formalism, as the classical framework, is concerned with idealized measurements which are repeatable—that is, measurements which, when immediately repeated, yield the same outcome with certainty. If one considers such measurements, then it follows as a direct consequence of the fact that measurements are probabilistic that, in general, they disturb the state of the system upon which they are performed. In fact, they disturb it almost completely—almost no trace of the pre-measurement state of the system is left in the post-measurement state.

The fourth feature, complementarity, can be precisely expressed in a number of different ways. Perhaps the simplest is to say that, unlike the situation in classical physics, one cannot perform a repeatable measurement on a system which yields information about all of the degrees of freedom of the state of the system.

Taken together, the probabilistic nature of measurement outcomes and complementarity severely constrain an agent who wishes to learn about the state of a physical system which has been prepared by an ideal agent. First, due to complementarity, an observer cannot rely upon one type of measurement, but must use more than one. Second, due to the probabilistic nature of measurement outcomes, the agent must perform many measurements on identical copies of the system. Yet, after a finite number of measurements, the agent's knowledge about the state will still be imperfect. It is only in the unattainable, idealized limit of an infinite number of measurements that the agent's knowledge of the state becomes complete. Therefore, unless an agent has directly prepared a system herself, or knows precisely how a system was prepared, there is always an informational gap between the theoretical description of the underlying reality—the quantum state of the system—and her knowledge of that state.

In order to make the above remarks more concrete, consider Stern–Gerlach measurements performed upon a particle with a magnetic moment. Classically, one can represent the magnetic moment (or *spin*) of a particle as a three-dimensional vector, μ . A Stern–Gerlach measurement consists of a pair of magnets oriented in some specific direction, $\mathbf{r}$, which deflect the silver atom up or down by an amount that, classically, depends upon the component of μ in the direction $\mathbf{r}$. Accordingly, in the classical framework, a Stern–Gerlach measurement oriented in the z -direction records μ_z , the component of μ in the z -direction. Since measurements are ideally non-disturbing in the classical framework, subsequent Stern–Gerlach measurements oriented in the x - and y -directions can be made on the spin to record the x - and y -components of μ . In this way, an experimenter can, in principle, precisely determine μ by performing a sequence of Stern–Gerlach measurements on a single spin.

Quantum mechanically, the situation is radically different. First, the state of the spin is represented as a two-dimensional complex vector,

$$v = \begin{pmatrix} \sqrt{p_1}e^{i\phi_1} \\ \sqrt{p_2}e^{i\phi_2} \end{pmatrix} \quad (1)$$

A Stern–Gerlach measurement in the z -direction performed on the spin will yield just one of *two* possible outcomes, conventionally labelled *up* ($\uparrow$) and *down* ($\downarrow$). This discreteness does not arise due to a discretization of a real-valued quantity, but, rather, is—according to the quantum formalism—an inherent feature of certain measurements on quantum systems.

The Stern–Gerlach measurement will yield *up* with probability p_1 and *down* with probability p_2 . If the measurement is now repeated immediately afterwards, then, due to the property of repeatability, the measurement will (at least ideally) yield the same outcome with certainty. Suppose, for the sake of argument, that the first measurement yields the outcome *up*. Then, in order that the second measurement *also* yields *up*, the state of the spin in between the two measurements be

$$\mathbf{v} = \begin{pmatrix} e^{i\phi'_1} \\ 0 \end{pmatrix} \quad (2)$$

where ϕ'_1 is arbitrary. Hence, the first measurement *disturbs* the state of the system. In fact, the disturbance is almost total: the post-measurement state carries almost no trace of the degrees of freedom in the pre-measurement state. As a result, after the first measurement is made, no further information about the pre-measurement state can be obtained by performing additional measurements on the system.

Now, the Stern–Gerlach measurement oriented in the z -direction yields information about p_1 and p_2 , but no information about ϕ_1 and ϕ_2 , and we have already established that subsequent measurements on the same system yields no additional information. Thus, in choosing to perform this particular Stern–Gerlach measurement (oriented in the z -direction), *the experimentalist has effectively chosen to learn about p_1, p_2 at the expense of ϕ_1, ϕ_2* . In other words, she learns about *one-half* of the degrees of freedom in the state at the expense of learning nothing about the other half [2]. This trade-off—which is one way of expressing the idea of *complementarity*—is not something which can be alleviated by performing a more sophisticated repeatable measurement, but is a fundamental limitation built into the quantum framework [3]. If the experimentalist wishes to learn about ϕ_1, ϕ_2 , she would need to perform Stern–Gerlach measurements in different directions (for example, in the x - and y -directions) on other identical copies of the spin.

Taking stock of the above, we can identify two natural ways in which the agent is informationally constrained:

(1) *Limited information about future measurement outcomes.*

Given the state of the system and the measurement to be performed, the experimenter lacks information about the outcome that will be obtained. In the above example, prior to performing a Stern–Gerlach measurement in the z -direction on the spin, the experimenter does not know which outcome (*up* or *down*) will occur, but only that the probabilities of the two possible outcomes are p_1 and p_2 , respectively.

Quantitatively, prior to performing the measurement, the experimenter has uncertainty $H(p_1, p_2)$ about which outcome will be obtained. The H -function here is an uncertainty function, such as the Shannon entropy function, $H(p_1, p_2) = -\sum_i p_i \ln p_i$. After performing the measurement and obtaining a definite outcome, the experimenter's uncertainty has been removed. Hence, the

uncertainty, $H(p_1, p_2)$, can be interpreted as the amount of information the experimenter lacks prior to performing the measurement about which the outcome will be obtained.

(2) *Limited information about the unknown state of a physical system.*

If an experimenter is presented with a system in an unknown state and wishes to learn what that state is, the quantum framework imposes two kinds of fundamental limits. First, due to the probabilistic and disturbance features of measurements, the outcome of a single measurement performed on the system provides scant information about the state of the system. In practice, in order to build up any useful knowledge of the state, the experimenter must perform a large number of measurements on identically-prepared copies of the system. Furthermore, due to complementarity, a single type of measurement only provides access to one-half of the degrees of freedom of the state of the system, so that the experimenter must perform other types of measurement in order to build up information about all of the degrees of freedom in the state.

In the electron spin example, the experimenter wishes to learn about the unknown state, $\mathbf{v} = (\sqrt{p_1}e^{i\phi_1}, \sqrt{p_2}e^{i\phi_2})$. The experimenter's information about the outcome probabilities, p_1, p_2 , prior to performing the Stern–Gerlach measurement is encoded in the Bayesian prior probability $\Pr(p_1|I)$, where I symbolizes the experimenter's prior state of knowledge. After the experimenter has performed n identical Stern–Gerlach measurements on n identically prepared copies of the system, obtaining data which can be summarized in the data string $D = \text{"}\uparrow\downarrow\uparrow\uparrow\dots\text{"}$ of length n , the prior can be updated to the posterior, $\Pr(p_1|D, I)$, using Bayes' rule:

$$\Pr(p_1|D, I) = \frac{\Pr(D|p_1, I) \Pr(p_1|I)}{\Pr(D|I)} \quad (3)$$

If the experimenter obtains f_1 cases of outcome *up* and f_2 of outcome *down*, then

$$\Pr(D|p_1, I) = \frac{n!}{(nf_1)!(nf_2)!} p_1^{nf_1} p_2^{nf_2} \quad (4)$$

The amount of information the data thus provides the experimenter about p_1, p_2 can readily be quantified using, for example, the continuum form of the Shannon entropy,

$$\Delta H = \int \Pr(p_1|D, I) \ln \frac{\Pr(p_1|D, I)}{\Pr(p_1|I)} dp_1 \quad (5)$$

which is finite for finite n . Hence, for finite n , the experimenter only has finite, imperfect information about p_1, p_2 . Only in the practically unattainable limit as $n \rightarrow \infty$ does the experimenter gain perfect knowledge of p_1, p_2 .

Furthermore, the data string, D , provides no information about the ϕ_i . In order to obtain information about the ϕ_i , the experimenter needs to perform Stern–Gerlach measurements in other directions.

2.2.2. Quantum Description of Composite Systems

One of the basic premises of classical physics is that one can conceptually *decompose* the whole of physical reality into spatially disjoint *parts*, describe each of these parts separately, and then combine these partial descriptions together to form a description of the whole. This premise is formalized in the classical framework by the postulate that the state of a composite system is simply a list—a concatenation—of the states of its subsystems. For example, the state of a system, $(\mathbf{r}^{(1)}, \dot{\mathbf{r}}^{(1)}; \mathbf{r}^{(2)}, \dot{\mathbf{r}}^{(2)})$ of two particles is a list of the two states, $(\mathbf{r}^{(1)}, \dot{\mathbf{r}}^{(1)})$, $(\mathbf{r}^{(2)}, \dot{\mathbf{r}}^{(2)})$ of the particles considered separately.

In essence, in the classical framework, a description of the whole is a simple aggregate of the description of its parts. As a consequence, an agent who in practice is necessarily restricted to studying spatial regions of the physical universe of limited extent at any one time can nonetheless aggregate information gained about these regions to form a description of larger regions. According to the classical framework, by proceeding in this manner, an agent can aspire to an arbitrary precise description of an arbitrary large part of the physical universe.

In contrast, the quantum formalism asserts that physical reality is *not* constituted in this way—the description of a composite system is, in general, *not* determined by a description of its parts. For example, there are states of a composite system composed of two subsystems that cannot be specified by giving the state of each subsystem considered separately. Such states are known as *entangled* states, and vastly outnumber states that are not entangled. Thus, entangled states are generic—not the exception but the rule.

More generally, quantum theory implies that a physical system becomes entangled with other systems as it interacts with them. For example, if a proton and an electron interact with one another locally, they become entangled with one another. Furthermore, this entanglement persists if they are subsequently separated from one another, no matter how widely. Consequently, a typical physical system as it is found in nature is generically entangled with many other physical systems, including those at great distance from it. However, an agent studying that system has no way of determining the precise nature of these entanglements by studying the system alone. Furthermore, when an agent chooses to perform measurement on any given system, the agent breaks any entanglement between that system and other systems elsewhere. Thus, an agent who singles out a particular system for study gains limited information about its entanglements to other systems, and furthermore disturbs the system by irrevocably breaking those entanglements.

To make these comments more concrete, let us consider a bipartite system consisting of two particles, each with spin. In the simplest case, the spin component of each particle can be represented by a two-dimensional complex vector, and one possible state of the composite system that represents the spin components is

$$\mathbf{v}^{(12)} = \frac{1}{\sqrt{2}} \left[\begin{pmatrix} 1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right] \quad (6)$$

where $\otimes$ denotes the tensor product operation, and where the state has been written down with respect to Stern–Gerlach measurements in the z -direction. This state is entangled, meaning that it cannot be expressed in the form $\mathbf{v}^{(1)} \otimes \mathbf{v}^{(2)}$, where $\mathbf{v}^{(1)}$ and $\mathbf{v}^{(2)}$ are the spin states of the respective particles.

Now, if an agent performs a Stern–Gerlach measurement in the z -direction on the first particle, the above entangled state yields equal probability of obtaining *up* or *down*. The quantum formalism further predicts that a Stern–Gerlach measurement in *any* direction performed on the first particle has equal probability of yielding the two possible outcomes. This prediction holds true even for a *different* entangled state such as

$$\mathbf{v}^{(12)'} = \frac{1}{\sqrt{2}} \left[\begin{pmatrix} 1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 0 \end{pmatrix} - \begin{pmatrix} 0 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right] \quad (7)$$

Hence, the agent is incapable of distinguishing between distinct entangled states on the basis of the statistics of measurements performed on a single particle. Furthermore, if two agents were to perform Stern–Gerlach measurements in the z -direction separately on the two particles then, for both states $\mathbf{v}^{(12)}$ and $\mathbf{v}^{(12)'}$, the agents would each have equal probability of obtaining “*up, up*” or “*down, down*”, and zero probability of obtaining “*up, down*” or “*down, up*”, and hence would not be able to distinguish between them [4].

We also note that, after the above-mentioned Stern–Gerlach measurement is performed on the system in the state $\mathbf{v}^{(12)}$, the post-measurement state is

$$\mathbf{v}^{(12)} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \otimes \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad \text{or} \quad \mathbf{v}^{(12)} = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \otimes \begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad (8)$$

depending upon which outcome is obtained. Thus, the post-measurement state is unentangled.

It is not obvious by inspection of the above example (nor by casual inspection of the quantum formalism itself) that, when the measurement on the first particle is made and the state of the composite system is thereby changed (becoming disentangled), there is an *actual physical influence* which passes from one system to the other. In fact, were there such an influence, it would need to have an infinite speed of propagation in order to be consistent with the quantum framework. Such an influence would, however, be in conflict with one of the basic tenets of classical physics, namely that space is the mediator of all interactions between physical systems and, moreover, that the influences mediating these interactions propagate through space at a finite speed, an assumption known as *locality*. Nevertheless, in 1964, John Bell [5] showed that, if the quantum predictions are valid, and if one accepts that there is no strange conspiracy between the prepared entangled state of a bipartite quantum system and the measurements that experimenters later perform on each of the subsystems, then locality must be false—there must exist such instantaneous influences.

Bell’s argument concerns the correlations between the outcomes of measurements performed on different sub-systems of a composite system. The quantum framework predicts that if these sub-systems are entangled, then these measurement outcomes are correlated. For instance, in the above example, as we have already mentioned, if two experimentalists were to perform separate Stern–Gerlach measurements in the z -direction on the two particles in the spin state $\mathbf{v}^{(12)}$, their respective outcomes would be perfectly correlated—they would either both get *up* or both get *down*.

Now, correlations such as these are also possible in the classical framework. For example, if Alice and Bob are each given a classical spin and told that the spins are both pointing up or both pointing down, then, when they subsequently make a measurement in the z -direction on their respective spins,

they will find that the spin directions they measure are perfectly correlated with one another. However, a pair of quantum systems admits correlations which are much richer, and Bell showed that, if a pair of quantum systems is allowed to interact and are then separated by an arbitrarily great distance, and the two agents are then allowed to freely perform *different* measurements upon the two subsystems, then the outcomes of their measurements can be correlated in a way that *cannot* be accounted for if we assume that reality behaves in accordance with the mechanical conception of classical physics (being local, in particular). An excellent and accessible self-contained exposition of Bell's theorem requiring only high school algebra and no prior knowledge of quantum theory can be found in [10]. In short, Bell showed that entanglement indeed leaves an experimentally detectable fingerprint—*non-local correlations*—which cannot be accounted for within the classical framework.

Bell's theorem implies that there is *some kind of connection* between physical systems that have interacted in the past, a connection that endures irrespective of their distance, and that enforces subtle but very real correlations between them. This connection is of a type never seen in classical physics. Unlike the forces of gravity, electricity or magnetism, the connection is unattenuated by distance, and is specific to the systems that previously interacted. Consequently, an agent who performs a measurement on a physical system unavoidably influences all other systems with which it is entangled, thus disturbing not only the system of interest but also affecting the other systems with which it is entangled.

Remarkably, although this non-local connection exists, the quantum framework implies that this connection cannot be used by two agents to send *information* to one another instantaneously. As physicists sometimes say, there is a peaceful coexistence between the properties of non-locality and no-signaling.

In view of its extraordinary implications, the interpretation of Bell's theorem remains controversial. For example, various authors have claimed that, either explicitly or implicitly, Bell's theorem makes one or more additional assumptions (such as assuming the validity of counterfactual reasoning) beyond those stated, and argue that it is more reasonable to deny these additional assumptions than to accept them and the concomitant conclusion that the validity of quantum predictions implies the violation of locality (for recent examples, see the papers by Blaylock [6] and Griffiths [7], and rebuttals by Maudlin [8,9]). However, in recent years, the conclusion of Bell's work has been bolstered by the demonstration that entanglement is a powerful physical resource that can be used to carry out various information-processing tasks such as dense coding [11], quantum teleportation [12], and cryptographic key distribution [13], which are impossible or very difficult according to classical physics.

3. The Rise of the Informational View

As described in the Introduction, over the last few decades the view has arisen amongst a growing number of physicists that the concept of information may have a fundamental role to play in our understanding of quantum theory and in our search for new physical theories. Quantum theory, with its non-classical features (described in the previous section) and unprecedented empirical success over the span of close to a century, has played a key role in the emergence of the informational view. However, one can identify several key stages in which other developments in physics and related subjects have also played a crucial role. In order to put the informational view in an appropriately broad context, and,

reciprocally, to be able to appreciate the possible broader impact of this view, the key developments will now be briefly described in approximate chronological order.

3.1. Mach and the Primacy of Experience over Concepts

As mentioned in the Introduction, in the two centuries following its articulation, the assumptions and idealizations underlying mechanical conception of classical physics were challenged from many directions. One early important source of challenge came from physicist-philosopher Ernst Mach (1838–1916). Mach developed the notion that a physical theory should not be regarded so much a description of how the world *is*, but is rather a compact summary of a related group of experimental facts:

“The goal which it (physical science) has set itself is the simplest and most economical abstract expression of facts.”

In expressing such an attitude, Mach was seeking to undermine the emphasis hitherto placed on the conceptual framework of a physical theory (the mechanical framework of classical physics in particular), arguing that experimental facts, not concepts, ought to be regarded as the primary reality, while the concepts which we use to frame these facts ought to be regarded with a certain suspicion. He also asserted that, in order to place concepts on the firmest possible foundation, ideally every concept in the conceptual framework of a physical theory ought to be definable in terms of definite experimental operations, a view that came to be known as *operationalism*. For example, through painstaking conceptual analysis of Newton’s mechanics, Mach showed that Newton’s concept of absolute space was an inadequate basis for Newton’s explanation of accelerated motion, and suggested that absolute space was actually simply a proxy for a rest frame defined by the “fixed stars”.

Mach also argued that abstract physical theories, with the uniformity that their so-called universal laws promise, are in fact *approximations* to our experience of an ever-changing reality, and not the converse:

“In mentally separating a body from the changeable environment in which it moves, what we really do is to extricate a group of sensations on which our thoughts are fastened and which is of relatively greater stability than the others, from the stream of all our sensations. Suppose we were to attribute to nature the property of producing like effects in like circumstances; just these like circumstances we should not know how to find. Nature exists once only. Our schematic mental imitation alone produces like events.”

Mach’s emphasis on the primacy of experimental data and the concomitant demotion of the metaphysical assumptions and idealizations of physical theory to a secondary status had a strong impact on the development of physics in the early twentieth century. The theory of relativity—both special and general—developed by Einstein bears testimony to the freedom of thought (from undue attachment to the assumptions underlying classical physics) fostered by Mach’s point of view. For example, Einstein’s key insight leading to the special theory of relativity was that the apparent conflict between Maxwell’s equations and Galilean invariance could be resolved by giving up the Newtonian notion of absolute time (which had also been previously criticized by Mach as not operationally well-grounded) in favor

of the notion that two observers in relative motion do not, in general, agree on the amount of time that passes between two events.

Heisenberg, one of primary contributors to quantum theory, was—via Einstein—also influenced by Mach’s ideas. While trying to understand the regularities in the frequency spectra of light emitted by excited atoms, he recognized that the assumption hitherto made that an electron orbits an atomic nucleus in the same way that a planet orbits a star was empirically not well-founded and so set it aside, instead seeking to directly capture the numerical regularities seen in the observed frequency spectra.

Although Mach did not speak about information *per se*, his view of a physical theory as, first and foremost, a mental construction by which we economically organize our sense-impressions (possibly mediated by experimental extensions thereof) placed great emphasis on careful analysis of what an agent actually experiences, how the agent infers features of the world from those experiences, and how those experiences are thereby gradually abstracted and idealized. Thus, he placed the agent—together with the processes by which the agent learns about the physical world by abstraction from his experience—center stage.

3.2. Thermodynamics, Statistical Mechanics, and Maxwell’s Demon

Thermodynamics is concerned with the transfer of heat energy between physical systems, and the interconversion of heat energy and mechanical energy (“work”). In thermodynamics, a physical system is described in terms of such large-scale variables as volume, temperature, and pressure. This mode of description is intentionally coarse—thermodynamics is, by design, tailored to the practical concerns and constraints of human agents interacting with physical systems. Although it intentionally does not adopt a fundamental mode of description, thermodynamics is nevertheless a quantitative, precise science with conceptually elegant and mathematically precise laws that concern the definition of temperature (the zeroth law of thermodynamics), the mechanical equivalent of heat energy (the first law), and the transfer of heat between bodies at different temperatures (the second law).

Statistical mechanics, as developed in the last few decades of the nineteenth century, sought to build a bridge capable of connecting the thermodynamic macro-description of a physical system with its micro-description supplied by classical physics. Its major accomplishment was the recognition that the thermodynamic state variable known as *entropy* could be understood, from the microscopic standpoint, as a measure of the number of microstates consistent with a given thermodynamic macrostate. Thus, as would later be shown more clearly by Jaynes (see Section 3.3, below), thermodynamic entropy could be interpreted as a measure of the amount of information about the microstate of a system that an agent lacks if he knows only the macrostate of the system.

In one of its simplest statements, the second law of thermodynamics asserts that heat does not spontaneously flow from a cooler to a hotter body. As a statement of our everyday experience, this statement is unobjectionable. However, it is at odds with a classical mechanical description of the situation, which (as described in Section 2.1) asserts that all dynamics is reversible, so that any change between two given states is possible no matter which of these states is placed earlier in time than the other. This lack of congruency between the laws of thermodynamics and classical mechanics caused considerable disquiet and discussion. The dominant view, formulated by Maxwell, asserted that the second law was only *statistically* valid: although the fundamental dynamics are indeed reversible, it is

overwhelmingly more likely that one will observe certain transitions from one macrostate to the other rather than the other way around.

In the process of articulating his point of view, Maxwell formulated an ingenious thought-experiment. Maxwell asks the reader to consider a microscopic creature who is a “very observant and neat-fingered being”, capable of following the positions and velocities of all of the molecules in a box of gas. Such a being, Maxwell argued, would be able to sort the molecules to create a macroscopic temperature difference between two sides of the box without the expenditure of any work, thereby generating a violation of the second law of thermodynamics.

In his argument, Maxwell tacitly assumed, as per the idealizations of classical physics, that it is possible in principle for an agent to obtain arbitrarily precise information about the state of a physical system without in any way affecting the physical world. In 1929, Leo Szilard proposed an ingenious solution to Maxwell’s demon (as Maxwell’s microscopic creature came to be known) that brought this assumption into question [14]. Szilard realized that if we were to assume that such a demon could *not* in fact perform the measurement-and-sorting operation envisaged by Maxwell without any physical cost, then it might be possible to “save” the second law of thermodynamics from the activity of such a being. In particular, Szilard argued that the demon’s activity would need to be accompanied by the generation of thermodynamic entropy (of at least $k \ln 2$ if he determines whether a molecule is in one half or the other half of a box, and acts appropriately on this information) in order to offset the reduction in thermodynamic entropy of the system due to his activity.

Taken together, thermodynamics and statistical mechanics played an important role in the evolution of physicists’ thinking in at least three distinct respects. First, thermodynamics demonstrated that it was possible to construct a rigorous, mathematical theory to describe the physical world from the point of view of non-ideal agents, so that the construction of a physical theory did not *require* that the physical world be described from the point of view of an all-seeing ideal agent. Second, the tension between the macroscopic and microscopic descriptions (respectively belonging to thermodynamics and to classical mechanics) of a physical system provided the first quantitative study of a situation where there is a gap between what is known by an agent about a physical system and the actual state of the system. Third, Szilard’s proposal that even an ideal agent (Maxwell’s demon) might, in fact, be subject to a new physical constraint was perhaps the first time that the classical notion of an ideal agent—capable of arbitrarily precise probing of the state of a physical system with negligibly small disturbance thereof—was brought into question in a quantitative manner, and the first time that it was proposed that *information acquisition or processing* had an necessary associated *physical cost*.

3.3. Shannon’s Theory of Information, and Its Applications

In 1948, the electrical engineer Claude Shannon introduced a measure of the information gained when one learns of the outcome of a probabilistic process [15]. For concreteness, suppose that one is presented with a standard coin, its sides labeled *heads* and *tails*, and that prior to tossing the coin, one is uncertain about what side the coin will show once it has come to rest on the table. Once the coin is tossed and reveals a particular side, one’s uncertainty is reduced to zero, which one can regard as a gain of information. Intuitively, labeling the outcome probabilities p_H (heads) and p_T (tails), one’s gain of

information is maximal when these probabilities are both $1/2$, is zero when one of these probabilities is zero, and varies smoothly between these extremes for intermediate probability distributions.

Shannon showed that it was possible to formalize these and a few other intuitive ideas in the form of mathematical postulates, and derive from these postulates a definite expression for the information gain applicable to any probabilistic process. More precisely, if a probabilistic process has N possible outcomes with probabilities $p_1, p_2, \dots, p_N$, he showed that the information gain is given by

$$H(p_1, \dots, p_N) = -k \sum_i p_i \ln p_i \quad (9)$$

where k is an arbitrary positive constant. At the suggestion of John von Neumann, Shannon decided to call $H(p_1, \dots, p_N)$ the *entropy* of the probability distribution, in conscious acknowledgment of the connection to the work of Szilard described in Section 3.2.

The establishment of a precise, quantitative measure of information associated with learning the outcome of a probabilistic process raised the question of whether this measure could provide new understanding into the formalism of existing theories of physics that involve probabilistic processes. In 1957, Jaynes showed that, indeed, it was possible to reinterpret the derivation of the formalism of statistical mechanics due to Gibbs at the turn of the century as an application of a new principle of inference, namely the Principle of Maximum Entropy, which states [16,17]:

Principle of Maximum Entropy: In assigning a probability distribution $\mathbf{p} = (p_1, \dots, p_N)$, select the distribution which has maximum entropy, $H(p_1, \dots, p_N) = -k \sum p_i \ln p_i$ subject to the normalization constraint $\sum_i p_i = 1$ and any other constraints on $\mathbf{p}$.

For example, the canonical distribution of statistical physics can be derived by imposing the constraint that the expected energy of a system, $\langle E \rangle = \sum_i E_i p_i$, is known, where E_i is the energy of the i th microstate of the system. Thus, the Shannon entropy of the canonical distribution can be regarded as a precise measure of the amount of information an agent lacks about the microstate of the system if he knows only the average energy of the system.

3.4. Black Hole Physics

In 1916, Schwarzschild published an exact solution of Einstein's field equations of general relativity for a point mass of mass M . The solution showed a coordinate singularity along a sphere of radius $2GM/c^2$ centered on the mass. Subsequent work showed that this coordinate singularity represented a kind of real physical boundary—once in-falling matter passes the boundary, neither it nor the light it emits is able to penetrate the boundary. Hence, according to general relativity, no detailed information about matter within the spherical boundary can be obtained by an observer external to the boundary, a situation which led Wheeler to name this object a *black hole*. General relativity also implies that a black hole, no matter how it is formed, is completely characterized by just three parameters, namely mass, angular momentum, and electric charge. Thus, whatever the internal constitution of a black hole may be, the only information available to an external observer about its internal constitution are these three real-valued parameters.

In the early 1970s, it was established that, when quantum mechanical effects are taken into account, a black hole can be treated within the framework of thermodynamics—it has an entropy (for example,

$S = 4\pi kGM^2/\hbar c$ for a non-rotating, chargeless black hole of mass M) and a temperature, and obeys four laws analogous to the four laws of thermodynamics [18]. Viewed from the point of view of statistical physics, the entropy of a black hole quantifies an external observer's lack of information about the internal structure of the black hole. The interpretation of the Bekenstein–Hawking expression for the entropy of a black hole (which is derived using thermodynamic arguments [18]) from a microscopic point view is supported by calculations from a microscopic point of view within the frameworks of string theory [19] and loop quantum gravity [20] that both yield the Bekenstein–Hawking expression. For example, for a non-rotating, chargeless black hole, the information that an external observer lacks can be written down in a binary string of length given by A/a_p , where A is the area of the event horizon and a_p is the area $(4 \ln 2)l_p^2$ with l_p being the Planck length (approximately 10^{-35} m).

However, according to general relativity alone, the internal structure of the black hole involves a non-denumerably infinite number of degrees of freedom, so that one would expect this lack of information to be infinite. The fact that it is actually *finite* suggests that the degrees of freedom are not non-denumerably infinite. This essential idea has led to the proposal—the *holographic principle*—that a finite amount of information is sufficient to completely describe *any* region of space.

Black hole physics has been of relevance to the emergence of the information view of physics in a number of ways. First, it provided the first example of a physical system whose precise physical state is, according to general relativity, *in principle* inaccessible to an external observer. Second, the finiteness of black hole entropy suggests that the fundamentality of the continuum description of matter and space-time assumed in classical physics may need to be reconsidered, and that the continuum description is itself a coarse approximation of an ultimately discrete underlying structure.

3.5. Computation as a Physical Process

The elementary unit of information is a bit (short for binary digit), which takes the value zero or one. Abstractly, a computation can be regarded a deterministic map from one string of bits to another string of bits. In 1936, Turing proposed a minimal abstract computing device—a *Turing machine*—capable of enacting this transformation in a step-by-step manner [21]. In a Turing machine, the machine acts on a tape divided into square spaces, in each of which a symbol is printed. The machine has a discrete internal state, and a “head” that, at any time, is positioned opposite a particular space on the tape. At each time-step, according to the internal state, the head either reads the symbol printed on the space, or erases it and writes a symbol to the space; and the machine then moves the tape one unit to the left or right or leaves it in the same place. The machine's internal state is also modified in a deterministic way. All of our modern computers, however sophisticated, are essentially realizations of a Turing machine.

Turing's characterization of a computing machine presupposes the essential idealizations underlying classical physics, and can indeed be regarded as a formalization of the notion of a mechanical process that underlies classical physics. In particular, direct analogs to the the classical notions of state, deterministic dynamics, and ideal measurement are all clearly visible in Turing's construction.

In proposing a concrete realization of the notion of a mechanical computing machine, Turing took a crucial step in bringing physics and computation together: If a computation is to actually be carried out step-by-step by a machine of our imagination that we could potentially construct, then the design of that machine must conform to the known laws of physics, and its capacities will necessarily be constrained

by those laws. Any change in those laws, or indeed any disagreement about the content of those laws, has the potential to impact any putative computing machine. There have been two important developments in this direction.

First, as mentioned in Section 3.2 above, in order to “save” the second law of thermodynamics from an ideal classical agent (Maxwell’s demon), Szilard postulated that the acquisition and processing of one bit of information is necessarily associated with the generation of a quantity $k \ln 2$ of thermodynamic entropy. Turing’s machine did not take into account any such fundamental physical cost of reading, erasing, and writing a symbol to a tape. In 1961, following Szilard, Landauer postulated that there indeed exists such an entropic cost, associated with the irreversibility of certain logical operations involved in a computation [22], and Bennett subsequently argued that the entropic cost is specifically associated with the *erasure* step [23]. A computing device within infinite memory could side-step this entropic cost since it could operate indefinitely without performing an erase operation. However, a device with finite memory would eventually fill up all available working memory and, from that point onward, need to perform erase operations, and would thus eventually be forced to start paying this entropic cost.

Second, Turing’s computing machine presumes the framework of classical physics. It is natural to wonder whether information processing fundamentally differs when embodied in the framework of quantum theory and, in particular, whether quantum theory imposes new constraints on how the information encoded in quantum systems can be manipulated or whether new things are possible when the full richness of quantum reality (such as the phenomenon of entanglement described in Section 2.2) is harnessed.

Let us first consider how information is encoded in, and read from, a quantum system. The elementary unit of quantum information is not the bit but the *qubit*, a two-state system with the state given in Equation (1). Information is encoded in the state of the qubit, which can be visualized by a point on the surface of a sphere. As detailed in Section 2.2, according to quantum theory, a measurement performed on the qubit will yield scant information about the state of the qubit. In particular, the measurement will probabilistically yield one of *two* possible outcomes. Thus, there is a vast disparity between the state of the qubit (which can take on a two-fold continuum of possible values) and the result of a measurement (which can take only two possible values). Furthermore, the act of measurement unavoidably brings about a change in the state of the qubit. Thus, due to the distinctive features of quantum measurement described in Section 2.2, it follows that, unlike the state of a classical bit, the state of a qubit cannot be passively “read”.

Now, one might imagine that there might at least exist a process which is capable of *copying* the state of a qubit without actually reading it. One of the most striking early findings of quantum information in the early 1980s is that, according to quantum theory, this is impossible [24]. That is, unlike the situation with a classical bit, there is no general device which, when fed a qubit in an unknown state, can output two qubits in the same unknown state.

But quantum theory does not only entail restrictions. Remarkably, over the last thirty years, it has been found that, when information is embodied in quantum systems, it is possible to carry out information processing tasks which are impossible or very difficult using classical information processing. For example, one can in principle build a *quantum computer* which can quickly solve certain problems (such as the problem of factoring large numbers) that are of great practical importance and yet intractable

on classical computers. One can also carry out tasks using quantum information processing that are classically *impossible*. For example, using quantum theory in what is known as *quantum cryptography*, two parties can communicate a message encoded in qubits, and yet to be able to detect whether there is an eavesdropper listening to their communication. This sensitivity to eavesdropping relies essentially on the fact that, in quantum physics, unlike classical physics, a measurement (the eavesdropping) is an *active* process that, in general, affects what is being measured.

Most of the innovations allowed by quantum information processing depend crucially upon the use of *entanglement*. For this reason, in quantum information processing, entanglement has come to be viewed as an indispensable resource. At a foundational level, this has stimulated the study of entanglement in its own right, and has persuaded many physicists that entanglement is a very real facet of physical reality.

4. Information Physics

Over the last few decades, there has arisen a growing belief amongst many physicists that the concept of information may have a critical role to play in our understanding of workings of the physical world. This view has been perhaps best articulated by John A. Wheeler under the slogan “*It from Bit*” [25,26]:

“ ‘*It from bit*’ symbolizes the idea that every item of the physical world has at bottom—at a very deep bottom, in most instances—an immaterial source and explanation; that which we call reality arises in the last analysis from the posing of yes-no questions and the registering of equipment-evoked responses; in short, that all things physical are information-theoretic in origin, and this in a participatory universe.”

and

“What we call reality consists of a few iron posts of observation between which we fill an elaborate papier-mâché of imagination and theory.”

As described in Section 3, many developments in physics and other disciplines over the course of the last century have paved the way for the emergence of this informational view. One can discern a number of key stages in this emergence:

1. **Shift from the view of a physical theory as a description of *reality in itself* to a description of *reality as experienced by an agent*.** Mach’s emphasis on the primacy of the experience of an agent over the concepts of a physical theory, thermodynamics as a theory explicitly constructed to interrelate the macro-variables accessible to limited agents, and quantum theory with its highly non-trivial model of the measurement process have all helped to shift the focus of physical theory from being a description of *reality in itself* to a description of *reality as experienced by an agent*.
2. **Breakdown of the classical notion of an ideal agent who has unfettered access to the state of reality.** Szilard’s proposal that there is a physical entropic cost associated with measurement (itself arising from the tension between the second law of thermodynamics and the reversible dynamics of classical mechanics), the quantum model of measurement (with its features of discreteness, probabilistic outcomes, disturbance and complementarity), and the limited access to information of the internal constitution of a black hole by an external observer all point to the *breakdown of the*

classical notion of an ideal agent who has unfettered access to the state of reality without bringing about any reciprocal change as a result, and all suggest that the interface between an agent and the physical world is a highly non-trivial one governed by precise rules which we can come to know.

3. **Recognition that the informational view might lead to new quantitative understanding of the reality described by existing physics.** Szilard's quantification of the entropy generation associated with information acquisition, Shannon's quantification of the information gain resulting from learning the outcome of a probabilistic process, Jaynes' derivation of statistical mechanics on the basis of Shannon's information measure, the quantification and statistical interpretation of black hole entropy, and the constructive use of the mathematical tools of information theory to discover and explore unexpected phenomena in the quantum world have all lent support to the hope that the informational view might lead to *new quantitative understanding of the reality described by existing physics* and that it is important that information be taken into account in the development of new physical theories.

As described in the previous section, the quantum framework brings about a relationship between an agent and the physical world which differs in a number of respects from that which exists in the classical framework, one in which the concept of information naturally plays a key explanatory role. This fact, together with the insights into the nature of quantum reality gained through an exploration of quantum information processing, naturally gives rise to the question of whether specific predictions of quantum theory and the mathematics of quantum theory—whose exact physical origin has long been regarded as obscure—can be derived from information-theoretic principles. As I will elaborate below, work carried out over the last thirty years has shown that this indeed appears to be the case.

4.1. Understanding Quantum Theory

In order to better describe the difficulty involved in understanding quantum theory, it is helpful to first consider its origin. As described in Section 2.1, the mathematical framework of classical physics was, by and large, arrived at through a mathematization of a clear conception of reality. In contrast, the mathematical framework of quantum theory did not come into existence in such a straightforward and principled manner, but through an indirect, rather obscure path consisting of somewhat *ad hoc* modifications of classical physics, guided by heuristic ideas (an example of such idea, due to de Broglie, is that each particle has an associated “guiding wave”) and by ingenious mathematical guesswork about what aspects of the existing mathematical structures of classical physics needed to be changed and what aspects could be retained.

The first quantum theory was non-relativistic quantum mechanics, developed in distinct forms by Schroedinger and Heisenberg in 1925–1926. Its key achievement was accounting in quantitative detail for the light spectrum of hydrogen in a variety of different circumstances. Over the course of the subsequent few years, a general and rather beautiful mathematical formalism—the *quantum formalism*—emerged, which provided a new mathematical framework (replacing the mathematical framework of classical physics) within which to build physical theories.

Owing to the indirect process by which the quantum formalism was created, the physical origin of many of the mathematical aspects of the formalism was not clear. For example, the states of quantum

systems are represented by complex vectors, but it was not clear precisely why complex numbers were necessary. As described in Section 2.2, some features—such as probabilistic measurement outcomes, complementarity, disturbance, and entanglement—of the reality described by the quantum formalism could nonetheless be safely read off from the formalism, and these were sufficient to establish (within a few short years after its creation) that it departed radically from the framework of classical physics. In the period shortly after the formulation of quantum theory, physicists responded to these non-classical features of quantum theory in essentially two distinct ways.

The first view, held by many of the founders of quantum theory, such as Bohr, Heisenberg and Pauli, was that the non-classical features of the formalism, such as the statistical nature of its predictions, reflected the very structure of physical reality, and that the classical mechanical view of physical reality therefore had to be replaced by something fundamentally new.

In order to illuminate these non-classical features of the quantum formalism, some of these physicists attempted to identify related concepts in the existing philosophical literature or to develop new concepts of which these features could be regarded as particular instantiations, thereby placing these non-classical features in a broader philosophical light. For example, Bohr developed the concept of complementarity, which he expressed as meaning that the process of coming to know anything about some aspect of reality is an active process that unavoidably has the effect of bringing into existence some aspect of it at the expense of simultaneously rendering inaccessible some other aspect of it [27]. Bohr believed that this concept had general validity beyond the field of physics itself, in areas such as psychology and sociology, for instance [28]. In his view, the impossibility in quantum theory of performing a repeatable measurement that provides complete knowledge about a physical system (but only one property of it at the expense of another) was simply a special case of the general principle of complementarity. Similarly, Heisenberg and Pauli both suggested that the Aristotelian notions of potentiality and actuality could be useful in understanding the relationship between the state of a system and the outcomes of measurements performed upon the system, and in understanding why the classical mode of thought led to inconsistencies when applied to such simple experimental situations as Young's double slit experiment [29,30].

The second view, held by Einstein and some other physicists, maintained that notwithstanding the manifestly non-classical features of quantum theory, the classical mechanical view of reality did not require the revision of its fundamental tenets. Einstein, for instance, argued that the statistical nature of quantum predictions was simply an indication that the quantum description of reality was incomplete, and spent a significant part of his later life searching for a classical field theory that was capable of underpinning quantum theory. De Broglie, in particular, supported this point of view by showing that the quantum theory of an ensemble of particles could be re-written in a form closely akin to classical mechanics, albeit with some curious non-classical features.

In the intervening eighty years, various attempts have been made to provide a physically intelligible interpretation of the quantum formalism. However, although these interpretations—such as the Copenhagen interpretation, the de Broglie–Bohm interpretation and the Many Worlds interpretation—paint extraordinarily diverse pictures of the quantum world, there appears to be little possibility of objectively choosing between them as each interpretation appears both internally consistent and consistent with the known experimental facts.

Over the last thirty years, it has become increasingly recognized that interpretations by themselves do not provide an adequate understanding of quantum theory [31,32]. One reason is the above-mentioned difficulty of choosing between interpretations. The other reason is that none of the existing interpretations account for the precise mathematical structure of quantum theory itself—all of the standard interpretations take the majority or entirety of the quantum formalism as a given. Thus, none of the standard interpretations is capable of accounting for the mathematical structure of quantum theory itself, including its many peculiar features alluded to above. It has thus become increasingly recognized that, in order to make progress in understanding quantum theory, it is necessary to first *derive* or *reconstruct* the mathematics of quantum theory from a set of postulates with a clear physical meaning.

4.1.1. Information-based Reconstruction of Quantum Theory

Attempts to reconstruct the mathematics of quantum theory have a long history and can, in fact, be traced back to the early 1930s. Indeed, Heisenberg, one of the creators of quantum theory, recognized that it would be highly desirable if the quantum formalism could be reconstructed using his uncertainty principle as a key axiom. Nonetheless, broadly speaking, reconstructive attempts prior to the 1980s tended towards highly abstract, intricate systems of axioms, and consequently made little impact [32,33].

Wootters' derivation of Malus' law. One of the first attempts to reconstruct quantum theory using a small number of physically compelling postulates was carried out by Wootters [34]. Wootters considers the amount of information gained by an experimenter about the parameters of the state of a quantum system when n identically-prepared copies of the system are analyzed, and postulates that, if one holds the experimental situation fixed but imagines varying the physical laws which govern the situation, then the actual laws are those which maximize the amount of information that the experimenter gains about the system.

Making some additional assumptions which will be described below, Wootters is able to show that, if the outcome probabilities, p_1, p_2 of a two-dimensional quantum system are a function of the angle θ between the preparation and analysis stages of an experiment, then $p_1(\theta) = \cos^2(m(\theta - \theta_0)/2)$, where $m \in \mathbb{Z}$, which is a generalized form of Malus' law, a well-verified prediction of quantum theory.

In outline, the argument runs as follows. Consider an experimental arrangement consisting of a preparation stage and an analysis stage. The preparation stage consists of a Stern–Gerlach apparatus oriented at (θ, ϕ) with a block in the negative channel, which prepares incoming electron spins. The analysis is performed by a vertically-oriented Stern–Gerlach apparatus. Suppose that the experimenter does not know the angles (θ, ϕ) because, for example, the preparation stage is physically removed from the laboratory. The aim of the experimenter is to learn about (θ, ϕ) from the data he obtains from performing measurements on n identically-prepared spins.

As already mentioned in Section 2.2, the quantum mechanical state of the prepared spin can be written

$$\mathbf{v} = \begin{pmatrix} \sqrt{p_1} e^{i\phi_1} \\ \sqrt{p_2} e^{i\phi_2} \end{pmatrix} \quad (10)$$

where p_1 and p_2 are the probabilities that the analysis will yield the outcomes *up* and *down*, respectively. If n spins identically prepared by this arrangement are analyzed, the experimenter can learn about the

values of the parameters p_1, p_2 from the frequency data that he obtains. According to quantum theory, the p_i are functions of θ , and quantum theory permits the precise determination of these functions. Using these functions, the experimenter can transform his knowledge about the p_i into knowledge about θ . Viewed in this way, the n electrons are “transmitting” information about the preparation parameter θ to the experimenter. Wootters’ strategy is to try to fix the function $p_1(\theta)$ by requiring that the information conveyed to the experimenter is maximized.

Upon each analysis, the experimenter obtains one of two possible outcomes. After n analyses, the experimental data, D , thus consists of the frequencies f_1, f_2 , where f_i is the frequency with which outcome i ($i = 1, 2$) has been obtained. The posterior probability, $\Pr(\theta|D, n, I)$ can be calculated from Bayes rule,

$$\Pr(\theta|D, n, I) = \frac{\Pr(D|\theta, n, I) \Pr(\theta|n, I)}{\Pr(D|n, I)} \quad (11)$$

and the likelihood can be written in terms of the function $p_1(\theta)$ as

$$\Pr(D|\theta, n, I) = \frac{n!}{(nf_1)!(nf_2)!} p_1^{nf_1} (1 - p_1)^{nf_2} \quad (12)$$

The amount of information obtained about θ itself depends upon the value of θ . The expected amount of information gained about θ is given by

$$\Delta H_{ave} = \sum_D \Pr(D|n, I) \{H[\Pr(\theta|n, I)] - H[\Pr(\theta|D, n, I)]\} \quad (13)$$

which is the expected change in H , the average being taken over all possible data, D , obtainable in a fixed number, n , of detections. The evaluation of ΔH_{ave} requires the specification of the prior probability $\Pr(\theta|n, I)$ which is equal to $\Pr(\theta|I)$ since the choice of n has no bearing on θ . Wootters assumes that $\Pr(\theta|I)$ is a constant on the grounds that, for any ϕ , all θ are equally likely. In the limit $n \rightarrow \infty$, one finds that the maximization of ΔH_{ave} yields a generalized form of Malus’ law,

$$p_1(\theta) = \cos^2 \left(\frac{m(\theta - \theta_0)}{2} \right) \quad (14)$$

where $\theta_0 \in \mathbb{R}$ and $m \in \mathbb{Z}$ remain undetermined. Wootters’ information maximization principle thus becomes:

Wootters’ Information Maximization Principle: The laws of quantum physics are such that the expected gain in the Shannon information about the state of a quantum system after analysis of a large number of identically-prepared systems is maximized. The average is taken over all possible data that can be obtained in a given number of analyses.

Wootters then attempts to extend this principle to an N dimensional quantum system, but the derivation invokes additional assumptions which presuppose prior knowledge of rather abstract features of the quantum formalism itself. Nevertheless, the derivation of Malus’ law is striking in its simplicity

and economy of means, and clearly demonstrates that it is not unreasonable to hope to be able to reconstruct the quantum formalism from information-theoretic principles.

Recent Reconstructions of Quantum Theory. Following Wootters' pioneering work, numerous attempts have been made to derive specific predictions of quantum theory, or to reconstruct the quantum formalism itself, from postulates inspired by an informational view, with much of the recent work inspired by the new perspective afforded by advances in the field of quantum information [35,36]. Almost without exception, these derivations take as a given the probabilistic nature of measurement outcomes, and formulate informationally-inspired postulates within this probabilistic framework.

The postulates that have been employed are diverse. One class of postulates concerns bipartite systems. For example, one early investigation [37] centers on the remarkable fact that, although quantum reality is non-local in the sense shown by Bell, quantum theory also implies that this non-locality cannot actually be used for instantaneous signaling, so that there is a peaceful coexistence between the properties of non-locality and no-signaling. Accordingly, the authors propose the postulate that such peaceful coexistence holds, and then explore to what extent this is capable of accounting for the structure of quantum theory. Although the authors concluded that that this peaceful co-existence is not by itself sufficient to account for the structure of quantum theory, it has proved to be important inspiration for later work. For example, it has very recently been shown that it is possible to account for the maximum *degree* of non-local correlation permitted by quantum theory (the so-called Tsirelson bound) by means of a new information-theoretic principle (the "principle of information causality") that can be regarded as a strengthening of the no-signaling condition [38].

Another example of a postulate of this type is due to Barrett [39], who elevates to the status of a postulate the observation that, in quantum theory, the state of a bipartite system can always be reconstructed from the statistics of joint measurements performed separately on the two sub-systems [40]. On the basis of this postulate, Barrett is able to account for at least some of the structure of the quantum formalism, such as the tensor product rule for determining the state of a composite system when its subsystems are in known pure states. This postulate plays an important role in recent reconstructions of quantum theory, in particular that due to Hardy [41,42] and Chiribella *et al.* [43], wherein it is allied to additional assumptions to obtain the quantum formalism.

Another class of postulates concern the behavior of individual systems. For example, in recent work, I have shown that it is possible to construct the core of the quantum formalism (namely Feynman's rules of quantum theory) by suitably formalizing the notion of complementarity and by postulating that certain measurements that yield no useful information about a physical system also do not disturb its state in any detectable way [44,45]. Remarkably, the complex structure of the quantum formalism, together with the rule that determine the outcome probabilities, are completely determined.

5. Towards a New Conception of Reality

As described in Section 2.1, classical physics is underpinned by a mechanical conception of reality. This conception is multi-faceted, but its essential presumption is that all particular events that occur in the physical world are, in their finest details, completely quantifiable and determined by physical laws. In particular, the mechanical conception posits that the totality of all that exists in the phenomenal world is *matter* moving in *space* according to *universal laws of motion*. Furthermore, in principle, agents can

probe this matter as precisely as they wish without disturbing its nature or its motion, and can use the knowledge thereby gained to make arbitrarily precise predictions of the behaviour of the entire universe.

In contrast, the informational view of physics places the agent center stage, confers non-trivial properties (such as indeterminacy and complementarity) on the interface between the agent and the physical world in which the agent is immersed, and suggests that such a basic construct as space is, in fact, an approximation that neglects both the pervasive entanglement of widely-separated physical bodies and the postulate that only a finite amount of information is required to completely describe any region of space. As outlined above, all of these ideas are now underpinned by a considerable body of theoretical and experimental evidence. At one level, one can regard the contents of the informational view simply as summarizations of experimental observations. But, at a deeper level, one would like to place these separate ideas into a coherent conception of physical reality within whose framework these separate ideas appear natural or even expected and thus related. The formulation of such a conception is important not only to unify our present understanding for its own sake, but also to provide a reliable guide to the future development of physics.

To illustrate the kind of understanding that is sought, consider the probabilistic nature of measurement outcomes postulated by quantum theory. It is one thing to accept that measurements are indeterministic as an operational principle (that is, as a summary of what we find in our experiments), but quite another to accept at a philosophical level the idea that events *cannot* in actual fact be predicted in their finest detail based on any prior available knowledge, an idea that is in clear violation of the determinism assumed in the classical framework. How are we to get a philosophical handle on such an idea? Is there any sense in which such an idea is reasonable? Similarly, in the case of complementarity, why is it that a repeatable measurement can only access *one-half* of the degrees of freedom of the state of a system? And, furthermore, is there some deeper reason why Nature might possess *both* of these features—indeterminacy *and* complementarity—and not just one or the other?

In order to make progress, it is helpful to recognize (as alluded to above) at the outset that the mechanical conception of reality is itself somewhat removed from our actual everyday experience of the world. As brought to light by the penetrating analyses of Hume, Mach, and others, the manifold abstractions upon which the classical framework is based are not as strongly rooted in our experience as we might sometimes suppose. For example, as has been long observed, the deterministic and reversible dynamics postulated by classical physics leads to model of reality in which there is no longer any fundamental distinction between past, present, and future. Such a model is in principle incapable of accounting for basic aspects of our experience, for example why we do not experience all times at once and instead experience reality unfolding gradually, and why we experience the past as fixed and the future as open and malleable [46].

If, however, one accepts the probabilistic nature of measurement as fundamental, then one cannot conceive of the future as objectively existing, and the sharp contrast between the past and future is naturally restored. Furthermore, as described in [45], probability can be regarded as a quantification of the degree to which the present state of reality (encoded in a logical proposition) implies any future state of reality (also encoded in a logical proposition), and the entirety of probability theory can be systematically built up on this basis. On this interpretation, then, probability simply quantifies the degree to which the present is *informed* about future events.

One of the most pressing challenges is the interpretation of the notion of measurement itself. Since the process of measurement as described by the quantum formalism is an active process which generates actual physical change in the system under observation, the formalism implies that there are in fact two distinct physical processes that a system can undergo: deterministic evolution and measurement. The question then arises: When is a physical process to be regarded as a measurement? In the classical framework, where a measurement passively registers the state of a system, this question is moot. However, in the quantum framework, due to the active nature of measurement, the question cannot be avoided. The quantum formalism itself provides no ready answer to this question. Thus far, it has been possible to apply the quantum formalism on the assumption that the process the formalism refers to as “measurement” only occurs when an agent performs a real measurement on a system. However, the question naturally arises whether, say, an electron interacting with a large protein molecule might undergo a physical process we call a measurement. At present, although there is no definitive answer to this question, several so-called spontaneous (or objective) collapse models have been proposed which imply that measurement-like processes occur spontaneously in nature, quite independently of probings initiated by macroscopic agents [47]. It is hoped that experimental advances may enable such models to be subject to experimental test in the not too distant future.

If one accepts that physical systems can undergo two distinct types of processes—deterministic evolution and a measurement-like probabilistic process—is there some way this can be understood at a deeper level? As mentioned earlier, Heisenberg and Pauli both suggested that there is a close parallel to the Aristotelian notions of potentiality and actuality, an idea that has been developed by various authors such as Whitehead [48].

Finally, if entanglement is indeed generic, then the status of space is fundamentally altered: Space is no longer a fundamental entity which mediates all interactions between material bodies, but is rather a useful approximate construct that neglects the fact the bodies also interact via an ever-changing web of inherently non-spatial connections. If this is so, a number of questions arise: Precisely *how* does space arise, *why* is space such a good approximation, and *when* does the approximation break down? Thus far, these are very much open questions with little in the way of plausible answers.

At present, then, it is far from clear how to combine the various facets of the emerging informational view into a coherent conception of physical reality. Nonetheless, what seems quite clear is that these various facets are now sufficiently well underpinned by theoretical and experimental work that they ought to be taken seriously as descriptions of how nature works, and the creation of such a conception is a vital next step.

References and Notes

1. More generally, the knowledge that an agent (be the agent ideal or non-ideal) possesses about the state of a system can be represented by a probability distribution over the state space of the system. This distribution itself is often also referred to as “the state” of the system. If the distribution picks out a single state, as would it be in the case of an ideal agent, it is said to be *pure*.

2. This characterization holds true for an N -dimensional quantum system. In that case, the state is represented by $\mathbf{v} = \sum_i \sqrt{p_i} e^{i\phi_i} \mathbf{v}_i$ where $\mathbf{v}_i$ is the i th eigenstate of measurement operator $\mathbf{A}$, and measurement A will yield information about the p_i (which constitute $N - 1$ independent degrees of freedom since $\sum_i p_i = 1$) at the expense of information about the ϕ_i (which constitute $N - 1$ independent degrees of freedom since the overall phase of the state is predictively irrelevant).
3. If one is willing to sacrifice repeatability, then it is possible to perform measurements—known as *informationally-complete* measurements—which are capable of accessing all of the degrees of freedom of a quantum state.
4. It is, however, possible for the two agents to distinguish between these two entangled states if they allowed to perform a sufficient number of *different* measurements on many identically-prepared copies of the two spins.
5. Bell, J.S. On the Einstein Podolsky Rosen paradox. *Physics* **1964**, *1*, 195–200.
6. Blaylock, G. The EPR paradox, Bell's inequality, and the question of locality. *Am. J. Phys.* **2010**, *78*, 111–120.
7. Griffiths, R.B. EPR, Bell, and quantum locality. *Am. J. Phys.* **2011**, *79*, 954–965.
8. Maudlin, T. What Bell proved: A reply to Blaylock. *Am. J. Phys.* **2010**, *78*, 121–125.
9. Maudlin, T. How Bell reasoned: A reply to Griffiths. *Am. J. Phys.* **2011**, *79*, 966–970.
10. Maudlin, T. *Quantum Non-Locality and Relativity*, 3rd ed.; Wiley-Blackwell: Malden, MA, USA, 2011.
11. Bennett, C.; Wiesner, S. Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states. *Phys. Rev. Lett.* **1992**, *69*, 2881–2884.
12. Bennett, C.H.; Brassard, G.; Crépeau, C.; Jozsa, R.; Peres, A.; Wootters, W.K. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Phys. Rev. Lett.* **1993**, *70*, 1895–1899.
13. Ekert, A. Quantum cryptography based on Bell's theorem. *Phys. Rev. Lett.* **1991**, *67*, 661–663.
14. Szilard, L. On the decrease of entropy in a thermodynamic system by the intervention of intelligent beings. *Z. für Physik* **1929**, *53*, 840–856.
15. Shannon, C.E. The mathematical theory of communication. *Bell Syst. Tech. J.* **1948**, *27*, 379–423.
16. Jaynes, E.T. Information theory and statistical mechanics I. *Phys. Rev.* **1957**, *106*, 620–630.
17. Jaynes, E.T. Information theory and statistical mechanics II. *Phys. Rev.* **1957**, *108*, 171–190.
18. Bekenstein, J.D. Black hole thermodynamics. *Phys. Today* **1980**, *33*, pp. 24–31.
19. Strominger, A.; Vafa, C. Microscopic origin of the Bekenstein-Hawking Entropy. *Phys. Lett. B* **1996**, *379*, 99–104.
20. Rovelli, C. Black hole entropy from loop quantum gravity. *Phys. Rev. Lett.* **1996**, *77*, 3288–3291.
21. Turing, A.M. On computable numbers, with an application to the Entscheidungsproblem. *Proc. Lond. Math. Soc.* **1936**, *2*, 230–265.
22. Landauer, R. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.* **1961**, *5*, 183–191.
23. Bennett, C. The thermodynamics of computation—a review. *Int. J. Theor. Phys.* **1982**, *21*, 905–940.
24. Zurek, W.H.; Wootters, W.K. A single quantum cannot be cloned. *Nature* **1982**, *299*, 802–803.

25. Wheeler, J.A. It from Bit. In *Proceedings of the 3rd International Symposium on the Foundations of Quantum Mechanics*, Tokyo, Japan, 1989.
26. Wheeler, J.A. Information, physics, quantum: The search for links. In *Complexity, Entropy, and the Physics of Information*; Zurek, W.H., Ed.; Addison-Wesley: Boston, MA, USA, 1990.
27. Bohr, N. Causality and complementarity. *Philos. Sci.* **1937**, *4*, 289–298.
28. Pais, A. *Niels Bohr's Times*; Oxford University Press: Oxford, UK, 1991.
29. Heisenberg, W. *Physics and Beyond*; HarperCollins Publishers Ltd.: Hammersmith, London, UK, 1971; Translated from the German original.
30. Pauli, W. *Writings on Physics and Philosophy*; Springer-Verlag: Berlin, Germany, 1994.
31. Fuchs, C.A. Quantum mechanics as quantum information. Available online: <http://arxiv.org/abs/quant-ph/0205039> (accessed on 27 September 2012).
32. Grinbaum, A. Reconstructing instead of interpreting quantum theory. *Philos. Sci.* **2007**, *74*, 761–774.
33. Grinbaum, A. Reconstruction of quantum theory. *Br. J. Philos. Sci.* **2007**, *58*, 387–408.
34. Wootters, W.K. The Acquisition of Information from Quantum Measurements. Ph.D. thesis, University of Texas at Austin, Austin, TX, USA, 1980.
35. Fuchs, C.A. Quantum mechanics as quantum information, mostly. *J. Mod. Opt.* **2003**. Available online: <http://perimeterinstitute.ca/personal/cfuchs/Oviedo.pdf> (accessed on 27 September 2012).
36. Brassard, G. Is information the key? *Nat. Phys.* **2005**, *1*, 2–4.
37. Popescu, S.; Rohrlich, D. Causality and nonlocality as axioms for quantum mechanics. Available online: <http://arxiv.org/abs/quant-ph/9709026> (accessed on 27 September 2012).
38. Pawłowski, M. Information causality as a physical principle. *Nature* **2009**, *461*, 1101–1104.
39. Barrett, J. Information processing in generalized probabilistic theories. *Phys. Rev. A* **2007**, *75*, 032304:1–032304:21.
40. Bergia, S.; Cannata, F.; Cornia, A.; Livi, R. On the actual measurability of the density matrix of a decaying system by means of measurements on the decay products. *Found. Phys.* **1980**, *10*, 723–730.
41. Hardy, L. Quantum theory from five reasonable axioms. Available online: <http://arxiv.org/abs/quant-ph/0101012> (accessed on 27 September 2012).
42. Hardy, L. Why Quantum Theory? Available online: <http://arxiv.org/abs/quant-ph/0111068> (accessed on 27 September 2012).
43. Chiribella, G.; Perinotti, P.; D'Ariano, G.M. Informational derivation of quantum theory. *Phys. Rev. A* **2011**, *84*, 012311:1–012311:47.
44. Goyal, P.; Knuth, K.H.; Skilling, J. Origin of complex quantum amplitudes and Feynman's Rules. *Phys. Rev. A* **2010**, *81*, 022109:1–022109:12.
45. Goyal, P.; Knuth, K.H. Quantum theory and probability theory: Their relationship and origin in symmetry. *Symmetry* **2011**, *3*, 171–206.
46. Norton, J.D. Time really passes. *Humana. Mente* **2010**, *13*, 23–34.
47. Ghirardi, G.; Rimini, A.; Weber, T. Unified dynamics for microscopic and macroscopic systems. *Phys. Rev. D* **1986**, *34*, 470–491.
48. Whitehead, A.N. *Process and Reality*; Free Press: New York, NY, USA, 1929.

49. Weart, S.R.; Szilard, G.W. *The Collected Works of Leo Szilard: Scientific Papers*; MIT Press: Cambridge, MA, USA, 1978.
50. Rosenkrantz, R.D. *E.T. Jaynes: Papers on Probability, Statistics, and Statistical Physics*; Kluwer Boston: Dordrecht, the Netherland, 1983.

© 2012 by the author; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>).