

A Philosopher's View of the Epistemic Interpretation of Quantum Mechanics

Shahar Avin
Trinity College, Cambridge*

May 21, 2010

Abstract

There are various reasons for favouring ψ -epistemic interpretations of quantum mechanics over ψ -ontic interpretations. One such reason is the correlation between quantum mechanics and Liouville dynamics. Another reason is the success of a specific epistemic model (Spekkens, 2007), in reproducing a wide range of quantum phenomena. The potential criticism, that Spekkens' restricted knowledge principle is counter-intuitive, is rejected using 'everyday life' examples. It is argued that the dimensionality of spin favours Spekkens' model over ψ -ontic models. van Enk's extension of Spekkens' model (2007) can even reproduce Bell Inequality violations, but requires negative probabilities to do so. An epistemic account of negative probabilities is the missing element for deciding the battle between ψ -epistemic and ψ -ontic interpretations in favour of the former.

1 Introduction

Theories in the foundations of quantum mechanics can be separated into three categories:

1. ψ -Complete
2. ψ -Supplemented
3. ψ -Epistemic

The first two are ontic (ontological), the third epistemic. The first is the main-stream version, taught in undergraduate courses. However, it is inconsistent with regard to measurements and determinism. The latter two try to fix the problems of measurement and determinism by assuming a reality beyond the wave-function. The ψ -supplemented

*I owe a special thanks to Dr. Jeremy Butterfield for supervising this paper, generously providing encouragement, comments and advice. I also thank Dr. Alex Broadbent and an anonymous reviewer in the History and Philosophy Department, Cambridge University for helpful comments.

approach, often labelled a “hidden variable” approach, accepts that ψ is a complete description of some part of reality, but assumes that there exist some variables beyond the wave-function that, together with the wave-function, completely describe all of reality. The ψ -epistemic approach assumes that the wave-function is not a representation of reality, not as a whole nor of any part thereof, but rather a representation of our knowledge of reality, expressed as a probability function over an underlying ontic configuration space. [12, 8]

For example, consider the game of “find the fish” in Monty Python’s *The Meaning of Life*. In the game, a fish is hidden somewhere in a cluttered room, and the audience is asked to find it. Let’s assume, without loss of generality, that the fish is under the bed - this is the full reality of the fish. Instead of a wavefunction ψ , we describe the location of the fish using a *statement*. Theories regarding the location of the fish can be divided into three categories:

Statement-Complete “The fish is under the bed.”

Statement-Supplemented “The upper half of the fish is under the bed.”

Statement-Epistemic “The fish is probably under the bed, but it might be in the fishbowl.”

The three categories of interpretations in quantum mechanics all assume an underlying reality, and are therefore versions of the ‘realist’ approach. In contrast, the ‘operationalist’ or ‘instrumentalist’ approach rejects the notion of an underlying reality, and attempts to understand quantum mechanics in terms of describing and predicting events in the lab: preparation procedures, manipulations and measurements. In this paper I will limit the discussion to the three categories of the realist approach. The only exception will be where I note instances where a realist account turns into an instrumental account.

2 Analogies to Classical Mechanics

One heuristic argument for favouring the epistemic approach over ontic approaches relies on analogies to classical mechanics, in which both types of frameworks exist – Newtonian mechanics as an ontic framework, Liouville dynamics as epistemic. As has come to light in recent years, mostly from the field of quantum information theory [4, 9], quantum mechanics behaves in a way that is more similar to Liouville dynamics than to Newtonian mechanics. [13]

An important result of quantum information theory is the no-cloning theorem. It can be proven [18] that no quantum mechanical operation, represented in quantum mechanics by a unitary transformation, can operate as a ‘cloning’ machine. Formally, there exists no unitary transformation U for which

$$U(|\psi\rangle|0\rangle) = |\psi\rangle|\psi\rangle \quad \forall |\psi\rangle \in \{|\psi_i\rangle\}$$

where $|0\rangle$ is some arbitrary initial state of the machine and $\{|\psi_i\rangle\}$ is a set of quantum states which are, in general, non-orthogonal. The proof of the theorem for a set of pure states considers the conservation of the inner product [10, p. 532], which can be considered as a kind of information distance between two quantum states. A similar proof exists for the classical domain, which shows that

universal perfect classical cloning machines violate the Liouville dynamics governing the evolution of statistical ensembles. This kind of copying process is in conflict with the conservation of the Kullback-Leibler information distance and with the linearity of the Liouville dynamics. [4]

The classical no-cloning theorem was presented explicitly as an analogy to the quantum version, and it is therefore telling that the objects of the classical theorem are statistical ensembles, which are *epistemic states* governed by Liouville dynamics, and not a single particle or a set of individual particles, which are *ontic states* governed by Newtonian mechanics. A similar process, of a classical analogue to a quantum information theory result which support the epistemic approach, is presented in [9] for the no-broadcasting theorem.

The analogy between quantum mechanics and Liouville dynamics is not restricted to quantum information theory. The analogy is also found in chaos theory. In classical chaos theory, which deals with ontic Newtonian mechanics, a chaotic system is one which is strongly sensitive to initial conditions. Formally, for two systems with initial distance Δx_0 , we define

$$\xi(t) = \lim_{\Delta x_0 \rightarrow 0} \left(\frac{\Delta x_t}{\Delta x_0} \right).$$

Generally, this function is a solution to the differential equation $\frac{\partial \xi(t)}{\partial t} = \lambda_1 \xi(t)$, and takes the form $\xi(t) = \exp(\lambda_1 t)$. In chaos theory λ_1 is called the Lyapunov exponent, and when it is positive the dynamics described by $\xi(t)$ display an exponential increase in separation between states with similar starting conditions, and the system is chaotic. This description of chaos, however, does not cross into the quantum realm. The distance between two quantum states $|\psi\rangle$ and $|\chi\rangle$, which in quantum mechanics is represented by the inner product of the states $\langle \psi | \chi \rangle$, is conserved over time, for any pair of states $|\psi\rangle$ and $|\chi\rangle$. An analogous fixed overlap between two separate states does exist classically – namely, in the overlap of two epistemic Liouville probability densities. [17]

3 Spekkens' Knowledge Balance Principle


If quantum mechanics is so similar to Liouville dynamics, how is it different? In Liouville dynamics it is possible to have states that reduce to the underlying ontic strata – delta function probabilities correspond one-to-one with the Newtonian state. This never happens in quantum mechanics. This has led Spekkens to introduce an *epistemic* foundational principle, the knowledge balance principle, as the underlying feature that

separates quantum mechanics from Liouville dynamics [13]. The framework that arises from the knowledge balance principle is explored in a “toy model” [11], where Spekkens is able to reproduce many quantum mechanical features.

According to the knowledge balance principle,

If one has maximal knowledge, then for every system, at every time, the amount of knowledge one possesses about the ontic state of the system at that time must equal the amount of knowledge one lacks. [11, p. 3]

In a framework based on this principle, the simplest allowed system is one of four ontic states, which forms the elementary system of Spekkens’ toy model. This can be represented graphically as:



where each of the numerals represents a distinct ontic state. According to the knowledge balance principle, the maximal knowledge epistemic states that are possible for this elementary ontic system are ones where exactly two ontic states are known to be possibly occupied, and exactly two are known to be unoccupied. Labelling possibly occupied ontic states with a shade, this gives rise to six epistemic states of maximal knowledge:



These states of the toy model behave in similar ways to the six spin states in quantum mechanics: ‘up’ and ‘down’ along the three axes \hat{x} , \hat{y} , \hat{z} . For example, the two toy model states on the left correspond to the z-axis spin eigenstates $|\uparrow\rangle$ and $|\downarrow\rangle$ – they are orthogonal to each other (no overlap), but each have a 50% overlap with any of the other states. In quantum mechanics, this behaviour is the result of the uncertainty principle, which is represented by non-commuting operators.

This leads us to the first successful result of the toy model – non-commutativity of measurements [11, p. 10]. Measurements in the toy model can be described in two stages. First, the outcome is produced based on the ontic state of the system just prior to the measurement. Based on the knowledge balance principle, the outcome of the measurement must present (at least) two possibly occupied ontic states and (at most) two unoccupied ontic states. Secondly, in order to maintain the knowledge balance principle, the ontic state of the system is randomised among the possible occupied states allowed by the outcome of the measurement. This randomisation necessarily leads to non-commutativity of measurements, in the following way. Assume, without loss of generality, that a system is initially in ontic state 1, and that a measurement is carried out on this system, which has the two possible outcomes: $(1 \vee 2)$ and $(3 \vee 4)$, and is therefore labelled $\{(1 \vee 2), (3 \vee 4)\}$. The outcome of that measurement is, necessarily, $(1 \vee 2)$. This measurement is followed by a $\{(1 \vee 3), (2 \vee 4)\}$ measurement, and another $\{(1 \vee 2), (3 \vee 4)\}$ measurement. Due to randomisation, the probability of getting the same result in the first and third measurement is $1/2$. However, if the order of the second and third measurements was swapped, the probability of getting the same outcome in

the $\{(1 \vee 2), (3 \vee 4)\}$ measurements would be 1 (certainty). The non-commutativity described is a perfect analogy to a quantum mechanical result, which can be observed experimentally using a series of Stern-Gerlach apparatuses.

Spekkens' toy model goes beyond the single elementary system described above. The main strength of the model lies in describing pairs of elementary systems, as this gives rise to an important quantum mechanic-like feature: entanglement. For a pair of systems there are sixteen ontic states (four for each system):


1,1	1,2	1,3	1,4
2,1	2,2	2,3	2,4
3,1	3,2	3,3	3,4
4,1	4,2	4,3	4,4

Maximal knowledge epistemic states for the pair of systems have four possibly occupied states. These can be 'product' states, in which one can read off a maximal knowledge epistemic state for each of the individual systems, such as

or 'entangled' states, in which the correlation between the pair of systems is fully known, but no information is available about the ontic states of the individual systems, such as

The toy model representation of a pair of systems allows Spekkens to reproduce the no-cloning theorem [11, p. 16], discussed in section 2. As in the quantum mechanics case and the Liouville dynamics case, the proof of a no-cloning theorem relies on the conservation of an information distance. In the case of the toy model, this information distance is fidelity, which roughly corresponds to the number of overlapping possibly occupied ontic states between two systems.

A quick recap: the no-cloning theorem states that no universal process can exist which starts with one system in a fixed initial state and another in an arbitrary state, and ends with two copies of the arbitrary state, if the arbitrary state can be drawn from a set of non-orthogonal states. In the paper, Spekkens presents the following example of two processes, which would be carried out by the same universal cloning machine:



However, the initial systems (on the left) overlap by two possible ontic states, (2,1) and (2,2), whereas the final systems (on the right) overlap by only one possible state, (2,3). Due to the knowledge balance principle the only allowed transformation in the toy model are permutations, but these can only change the *location* of the overlap, not the *number* of overlapping states. Thus, universal cloning is forbidden in the toy model.

A counter-intuitive feature of the toy model is that measurements necessarily include a randomisation (loss) of some prior knowledge. This is a result of the knowledge balance principle, because without randomisation a bayesian update will allow us to infer the state of the system with greater accuracy the more measurements we take. However, Spekkens provides a physical explanation of this randomisation, via a characterisation of the measurement device that relies on a Von-Neumann like shift of the boundary between ‘quantum’ system and ‘classical’ observer [16, pp. 418-421]. Say we want to measure position of a system. We prepare a measurement probe with a known position. Because of the limited knowledge principle, this probe necessarily has completely unknown momentum. When we interact the probe with the system, the position of the system affects a change in the position of the probe, which is registered and used to infer the position of the system. However, when the system and probe interact, the probe affects a change in the momentum of the system. Since the momentum of the probe is unknown, the momentum of the system after the measurement is randomised. [13]

4 Understanding Restricted Knowledge

Are there analogies in ‘everyday’ life to states of necessarily incomplete knowledge? We are familiar with states of ‘effectively’ incomplete knowledge – what we call ‘complicated’ or ‘chaotic’ systems, e.g a lottery machine or a tossed coin. These are precisely the classical cases in which we would choose to use Liouville mechanics rather than Newtonian mechanics. However, as computational resources increase, more and more phenomena that previously required an ‘epistemic’ approach are replaced by computational models that represent an ontic approach, e.g. in simulating spin-waves, phase transitions and fluid dynamics. Our increased ability to produce an ontic representation of complex classical systems means that we should look for analogies to ψ -epistemic quantum mechanics elsewhere, as these are fundamentally irreducible to an ontic representation.

Consider the following example. A patient has a suspected cancerous tumour near a major blood vessel. The doctor has two measurements at his disposal: a biopsy and an X-ray. The biopsy involves probing the tumour with a syringe and extracting a sample, which is later analysed to detect mutations in the cells' DNA. However, since the tumour is close to a blood vessel, there is a 50% chance that while taking the biopsy the needle will rupture the vessel and an internal haemorrhage will occur. The X-ray can be used to detect an internal haemorrhage by filming the affected area and analysing the scan, but there is a 50% chance that the scan will introduce cancerous mutations in the cells near the blood vessel. If an internal haemorrhage has already occurred, a biopsy has a 50% of detecting it, and the syringe can be used to drain the haemorrhage. If the cells already contain one mutation that makes them cancerous, a second X-ray dose has a 50% chance of introducing a second mutation that is fatal to cancerous cells, thus killing off the tumour. This example is a perfect (albeit highly artificial) analogue of a system with a knowledge balance principle. The two variables, the existence of a haemorrhage and the existence of mutations, cannot be simultaneously known.

Despite the artificiality of the example, we have a strong intuition that at any given moment, regardless of our current knowledge, there is a definite value for the existence (or absence) both of the cancerous mutations and the haemorrhage. A similar argument was used in the famous Schrödinger cat experiment to argue against a ψ -complete interpretation.

How are decisions made in a situation of restricted knowledge? Despite the strange relation between measurement and knowledge, it is possible to make rational choices in this situation. For example, if we decide that having no cancerous mutations has greater utility than having no haemorrhage, the optimal course of action would be: (a) Take a biopsy. If the biopsy shows no mutations (50%) – goal achieved. Otherwise, (b) Take an X-ray. Regardless of the X-ray result, go back to (a). In the average case the absence of cancerous mutations will be reached after the second biopsy.

Another intuition into the meaning of the restricted knowledge principle could be had from Joseph Heller's *Catch-22*. In the book, pilots are forced to fly life-risking bombing missions. The only way to avoid having to fly the missions is to be recognised as mad by the base doctor. Agreeing to fly the missions is a clear case of madness. However, once you point out to the doctor that you are mad and ask to be discharged, you display sanity, namely a desire to stop flying missions, and so you are forced to continue flying missions.

The *Catch-22* situation can be thought of in the following way. There are two variables in the system, the flying of missions and being mad. The two cannot be simultaneously known (by the base doctor)¹. In the initial situation the doctor doesn't have knowledge

¹Thanks to Dr. Alex Broadbent for pointing out to me that the doctor *does* know both simultaneously. The doctor has access to independent sources of information about madness, for example he hears Havermeyer shooting mice at night. I therefore limit my discussion to the hypothetical case where the doctor can only learn about madness by direct report, and all other information channels are forbidden.

about either. When you walk into the doctors office, he makes a measurement of sanity, and reaches a conclusion that you are insane (based on your report), without knowledge of whether you have been flying or not. However, he is then forced to make a measurement of whether you have been flying or not, and finds out that you have. Since he has complete knowledge of your flying, he can't have knowledge about your insanity, and therefore can't relieve you from flying. Applied more generally, if one of a pair of variables that are linked by the knowledge balance principle is always being measured, the second variable can never be known, and therefore can never be acted upon.

5 Dimensionality

In Spekkens' toy model, presented in section 3, the elementary system has four ontic states and six epistemic states. The epistemic states are analogous to the six spin eigenstates in quantum mechanics, of spins aligning with or against the three axes $\hat{x}, \hat{y}, \hat{z}$. However, to claim that the spin eigenvalue states are epistemic, and that the number of ontic states is lower, seems counter-intuitive, because the three-dimensional treatment originates from the Stern-Gerlach experiment. Another strong intuition for a three-dimensional treatment of spin is by analogy to orbital angular momentum, which is unquestionably a three-dimensional space phenomenon. Both spin and orbital angular momentum are handled experimentally using the Stern-Gerlach apparatus, both follow the same quantisation rules, both display the same uncertainty relations, and are both treated using very similar operator and matrix methods. It is this analogy with orbital angular momentum that gives spin its name – the phenomenon of spin is often presented as the quantisation of the angular momentum of a particle spinning around its axis. If spin is a three-dimensional space phenomenon, then dimensionality would favour a ψ -ontic interpretation with six ontic states, rather than Spekkens' ψ -epistemic interpretation with only four epistemic states.

However, the 'classical' conception of spin as a three-dimensional real space phenomenon of a particle spinning around its axis is false. The following is a short version of the proof presented in Howard Hughes' Part IB *Quantum Physics*, Cambridge NST, 2007:

If spin angular momentum S is describing real three-dimensional rotation of, say, an electron, then the associated kinetic energy would be:

$$K.E. = \frac{S^2}{2I} = \frac{s(s+1)\hbar^2}{2I} = \frac{15\hbar^2}{16m_e a^2},$$

where I is the moment of inertia, and assuming that the electron is a sphere of radius a . Experimental results give $a \leq 10^{-18}$ m, which gives a lower bound for the kinetic energy $K.E. \geq 10^{18}$ eV. This value *vastly* exceeds the measured rest mass energy of the electron, ~ 0.5 MeV, so the spin angular momentum of the electron *cannot be accounted for* by modelling it as a rotating sphere.

Therefore, the fact that Spekkens’ toy model analogy for spin has four ontic states, and not six or any other $3n$ number, should not count as evidence against this model.

The accepted solution for the origin of the two degrees of freedom of spin is provided by the Dirac equation [5]:

$$i\hbar\frac{\partial\Psi}{\partial t} = [c\underline{\alpha}\cdot\underline{\hat{p}} + \beta mc^2]\Psi.$$

The Dirac equation was derived in the paper cited above to account for relativistic quantum mechanics, in *four-dimensional* spacetime. The spin degrees of freedom arise from this four dimensional equation in the non-relativistic limit as a two component vector, *separable from the spatial component of the wavefunction*. Dirac’s result is often used to argue that spin is an intrinsic characteristic of the particle, like charge and mass, and therefore requires no further explanation. This is despite spin’s strong analogy to orbital angular momentum, which does have a physical reductive definition, $\hat{L} = \hat{r} \times \hat{p}$, in terms of three-dimensional position \hat{r} and momentum \hat{p} . Without developing the idea further, I would like to suggest that it might be possible to account for spin as a four-dimensional physical phenomena. In such a case, restricted knowledge might arise from our lack of ability to perceive time as a spatial dimension. If such an account is presented, it would give strong support to the epistemic approach to quantum mechanics, and specifically to Spekkens’ interpretation.

6 Completeness and Locality

Einstein has argued against the ‘orthodox’ ψ -complete interpretation of quantum mechanics, favouring, according to recent interpretations [8], the ψ -epistemic view.² Einstein’s arguments are based on a contradiction between the assumed completeness of description provided by the wavefunction and the physical principle of locality. The principle of locality assumes that physical systems that occupy different regions of space are separable, so that a full description of the physics of one system does not require a description of the other, and that there is no “action at a distance”, so that information from one system cannot be transferred to the other at a speed greater than the speed of light (the latter being in contradiction with Einstein’s special relativity). In the Solvay conference of 1927, Einstein argued that ψ -ontic models are either non-local or incomplete, and advocated adoption of the latter [1, pp.175-178, pp. 440-442]. In 1935 Einstein presented more rigorous arguments for the contradiction between locality and completeness, in the famous ‘EPR’ paper [6] and in private correspondence with Schrödinger. Harrigan and Spekkens have shown [8] that the latter argument is valid not only as a criticism of the ψ -complete interpretation, but also of the ψ -supplemented interpretation. Einstein examines two entangled particles A and B, which have distinct ontic states, S_A and S_B . According to quantum mechanics, one can perform different measurements on

²Quick note on sociology of knowledge: a successful enlistment of Einstein as a retrospective proponent of the epistemic view may prove useful for its future acceptance within the physics community.

A, which result in different values for ψ_B . This means that the ontic state S_B can be in agreement with several values of ψ_B . This contradicts the ψ -supplemented approach, where a single ψ can agree with several ontic states, but each ontic state can agree with only one value of ψ . Based on Einstein's arguments it would appear that the intuitive appeal of locality, and the experimental success of special relativity, would favour an adoption of the ψ -epistemic interpretation.

However, Bell's theorem shows that *any* realist interpretation of quantum mechanics, including the ψ -epistemic approach, must have non-local features, in order to reproduce the quantum statistics of entangled systems [3]. Instead of going through the algebra in Bell's paper, let us consider the following story. A TV quiz show, sponsored by a chocolate company, invites young couples to answer a series of Yes/No questions. The partners are separated from each other, and in every round each is asked either "do you love chocolate?" or "do you love your partner?". The conditions for 'winning' a round are:

- When both are asked "do you like chocolate?", both answers should be the same. This allows the sponsor to display one of two prepared adverts – "then you'll like our chocolate!" / "but you'll like our chocolate!"
- When asked different questions, both answers should be the same. This allows the sponsor to convey the correlation between chocolate and love.
- When both are asked "do you like your partner?", the answers should differ. This creates the minimal excitement necessary to keep the audience watching the aggressive advertising.

Bell's theorem states that as long as the partners can't communicate with each other, the maximal average 'winning' rate is 75%, *regardless of the strategy chosen*. However, quantum entangled particles are able to reach a 'winning' rate of about 85% [14]. The two questions in the quantum case are "what is your spin alignment along the z-axis?" and "what is your spin alignment along an axis rotated by 45° relative to the z-axis?". Based on this result, Bell argues that any account of quantum mechanics must violate the principles of separability and no action at a distance in order to reproduce the 85% result, or conversely that any account that doesn't reproduce the 85% result is not a complete account of quantum mechanics. Spekkens' toy model falls in the latter category: as there is no equivalent in Spekkens' model for the 45° rotated measurement, the maximal 'winning' rate is 75%. Bell's argument makes all three realist approaches face the same challenge equally – to reproduce the 85% result while maintaining locality in some novel way, or to provide justification for giving up the principle of locality.

Everett's Many-World Interpretation [7] provides a possible solution to Bell's challenge. Everett's interpretation starts off from an extreme ψ -complete framework: Everett assumes the wavefunction is a complete description of reality, and concludes that the co-existence of alternative quantum states must be interpreted literally – as a multiplicity of worlds. Since conscious observers are unable to occupy more than one world, the

process of quantum ‘collapse’ describes the process of an observer ending up in one of several accessible worlds, and is not a fundamental physical process. In the case relevant to Bell’s argument, all the different correlations and anti-correlations between the entangled particles exist, in alternative worlds. The process of measurement involved in yielding the 85% result is the process of the observer committing herself to one of these worlds, and since this is not a real physical process it is allowed to violate the principle of locality. The success of Everett’s interpretation in dealing with Bell’s challenge may favour adoption of the ψ -complete approach. To counter this argument, I sketch an ‘Everettian’ version of Spekkens’ toy model, as follows. In Spekkens’ model (Section 3), a measurement occurs in two stages: an epistemic result is produced deterministically based on the ontic state prior to measurement, then the ontic state is randomised among the states allowed by the epistemic state. At this second stage of randomisation, we can introduce an ‘Everettian’ twist to the model, and say that in reality all the possible ontic states occur, in alternative worlds. Such an explanation would enable an interpretation of the resulting restricted knowledge: our knowledge of ontic states is restricted because we are unable to determine in which of the alternative worlds we ended up. This sketch of an ‘Everettian’ version of Spekkens’ model should suffice to show that Everett’s interpretation does not favour the ψ -complete interpretation over the ψ -epistemic interpretation.

Spekkens’ toy model has been extended by van Enk [15] in a way that reproduces the violations of Bell’s inequalities observed experimentally in quantum entanglement, a result which Spekkens’ toy model is unable to achieve. In van Enk’s model the permitted epistemic states are represented by a less-restricted probability distribution over the ontic states. Recall that in Spekkens’ model the probability distribution was highly restricted – each ontic state had a probability of $1/N$, if it was permitted by an epistemic state with N allowed ontic states, or a probability of 0, if it was forbidden. By allowing less restricted probability distributions van Enk brings the toy model closer to the experimental results of quantum mechanics, which has states of the form $\alpha|0\rangle + \beta|1\rangle$ for arbitrary α and β as long as $|\alpha|^2 + |\beta|^2 = 1$. In van Enk’s model, the general epistemic state of a system with four ontic states would be

$$\boxed{p(1) \mid p(2) \mid p(3) \mid p(4)},$$

with the condition that $\sum p = 1$. Using this generalised epistemic state van Enk is able to reproduce correlations between two systems that violate Bell’s inequalities. The degree by which the inequalities are violated depends on a choice of a free parameter $r > -1$, which in turns determines the measure of information for the system via a Schur-convex function

$$M_r(P) = \left(\sum_x P(x)(P(x))^r \right)^{1/r},$$

where P is the probability distribution over the ontic states x [15, p. 1450]. With a range of choices for r , the ‘winning’ rate goes from the ‘local’ 75% all the way up to (almost) 100% [15, p. 1456]. A further restriction on the toy model sets $r = 1$, and reproduces the

85% ‘winning’ rate exhibited by quantum mechanics. However, this further restriction is equivalent to the introduction of negative probabilities into the model.

In fact, negative probabilities are necessary for allowing the model to violate Bell’s inequalities, even for $r \neq 1$. Negative probabilities are valid mathematical tools, as long as they can’t be interpreted as the frequency of occurrence of some observable phenomenon [2]. In van Enk’s model the negative probabilities apply to individual ontic states, which are unobservable due to the knowledge balance principle. When the probability is coarse grained for epistemic states, by summing the individual probabilities of the allowed ontic states, the negative probabilities disappear, and therefore the use of negative probabilities isn’t a problem. However, the use of negative probabilities reduces the account to a version of instrumentalism, unless it is possible to provide a realist interpretation of ontic states with negative probabilities. While such an interpretation seems highly counter-intuitive, this might not be an entirely lost cause³. Negative probabilities might relate to backward causality, or in some other way represent over-forbidden outcomes. To sum up: The question of a realist interpretation of negative probability is still open; but a satisfactory answer would provide very strong support to the ψ -epistemic interpretation of quantum mechanics.

7 Conclusion

In this paper I have looked at the ψ -epistemic interpretation of quantum mechanics, and at various reasons to favour it over ψ -ontic interpretations. Several correlations between quantum mechanics and Liouville dynamics, a classical epistemic framework, have been discovered recently, and they hold regardless of the particular epistemic model used. For a particular model, I chose to focus on Spekkens’ toy model, since it reproduces a wide range of quantum phenomena by investigating a single simple epistemic principle. I showed how Spekkens’ use of restricted knowledge could be understood using ‘everyday life’ examples, and therefore should not be rejected on grounds of being counter-intuitive or intractable. I argued that the number of ontic states in Spekkens’ model fits better with the dimensionality of spin than the number of states in ψ -ontic models. Finally, I looked at the tension between the principle of locality and realist interpretations of quantum mechanics. While Spekkens’ toy model is not able to reproduce Bell Inequality violations, van Enk’s extension of Spekkens’ model is able to do so. However, van Enk’s use of negative probabilities may reduce his model to a form of instrumentalism. A successful realist interpretation of negative probabilities, on top of the other evidence presented in this paper, will provide a compelling reason for adopting an epistemic interpretation of quantum mechanics, specifically based on Spekkens’ and van Enk’s models.

³This possibility was raised in an informal discussion with Dr. Sorin Bangu.

References

- [1] BACCIAGALUPPI, G. *Quantum Theory at the Crossroads*. Cambridge University Press, Cambridge, 2009.
- [2] BARTLETT, M. S. Negative probability. *Mathematical Proceedings of the Cambridge Philosophical Society* 41, 01 (1945), 71–73.
- [3] BELL, J. On the Einstein-Podolsky-Rosen paradox. *Physics (US) (Discontinued, 1968)* 1, 3 (1964), 195–200.
- [4] DAFFERTSHOFER, A., PLASTINO, A. R., AND PLASTINO, A. Classical no-cloning theorem. *Physical Review Letters* 88, 21 (May 2002), 210601.
- [5] DIRAC, P. The quantum theory of the electron. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 117, 778 (1928), 610–624.
- [6] EINSTEIN, A., PODOLSKY, B., ROSEN, N., ET AL. Can quantum-mechanical description of physical reality be considered complete? *Physical review* 47, 10 (1935), 777–780.
- [7] EVERETT, H., AND DEWITT, B. *The many-worlds interpretation of quantum mechanics: a fundamental exposition*. Princeton University Press, 1973.
- [8] HARRIGAN, N., AND SPEKKENS, R. Einstein, incompleteness, and the epistemic view of quantum states. arXiv:0706.2661v1 [quant-ph], June 2007.
- [9] KALEV, A., AND HEN, I. No-broadcasting theorem and its classical counterpart. *Physical Review Letters* 100, 21 (May 2008), 210502.
- [10] NIELSEN, M., AND CHUANG, I. *Quantum computation and quantum information*. Cambridge University Press, Cambridge, 2000.
- [11] SPEKKENS, R. W. In defense of the epistemic view of quantum states: a toy theory. *Physical Review A* 75 (2007), 032110.
- [12] SPEKKENS, R. W. Foundations of quantum mech. (phys 639) - lecture 6. Perimeter Institute, <http://pirsa.org/09120070/>, 7 December 2009.
- [13] SPEKKENS, R. W. Foundations of quantum mech. (phys 639) - lecture 7. Perimeter Institute, <http://pirsa.org/09120071/>, 8 December 2009.
- [14] SPEKKENS, R. W. Foundations of quantum mech. (phys 639) - lecture 8. Perimeter Institute, <http://pirsa.org/09120072/>, 9 December 2009.
- [15] VAN ENK, S. A toy model for quantum mechanics. *Foundations of Physics* 37, 10 (2007), 1447–1460.

- [16] VON NEUMANN, J. *Mathematical foundations of quantum mechanics*, 1996 translated ed. Princeton University Press, 1932.
- [17] WEINSTEIN, Y., LLOYD, S., AND TSALLIS, C. Border between regular and chaotic quantum dynamics. *Physical review letters* 89, 21 (2002), 214101.
- [18] WOOTTERS, W., AND ZUREK, W. A single quantum cannot be cloned. *Nature* 299 (1982), 802–803.