

Quantum- vs. Macro- Realism: What does the Leggett-Garg Inequality actually test?

O.J.E Maroney*

Faculty of Philosophy and Wolfson College, Oxford

C.G Timpson†

Faculty of Philosophy and Brasenose College, Oxford

October 14, 2013

Contents

1	Introduction	2
2	The Leggett-Garg Inequality	5
3	What is macrorealism? First pass	10
4	An operational formulation of the question	15
4.1	Operational non-disturbance implies the Leggett-Garg Inequality	18
5	Operational non-disturbance vs. Ontic non-invasiveness: The Ontic Models framework	20
5.1	The ontic models framework	20
5.2	Noninvasiveness as operational non-disturbance	22
5.3	Noninvasiveness as ontic noninvasiveness	24
5.4	Summary	25
6	What is macrorealism? Second pass	27
7	What does Leggett-Garg Inequality violation in fact rule out?	29
7.1	Macrorealism 1: Operational Eigenstate Macroscopic Realism . .	30
7.2	Macrorealism 2: Preparation-support Macroscopic Realism . . .	31
7.3	Macrorealism 3: Extra-preparation Macroscopic Realism	32
8	Macrorealism and null result measurements: Macrorealism does not imply non-invasiveness	33
9	Conclusions	36

*owen.maroney@philosophy.ox.ac.uk

†christopher.timpson@bnc.ox.ac.uk

1 Introduction

The Leggett-Garg inequality (Leggett and Garg, 1985) was introduced as a means of putting to experimental test a world-view which Leggett and Garg called *Macroscopic Realism*. According to this view, and in explicit contrast to what quantum theory allows—indeed, more strongly, in contrast to what quantum theory would sometimes seem to *require*—macroscopic objects must always be in some one determinate macroscopic state or another at any given time. No funny-business of quantum superposition is permitted at the macroscopic level. Leggett-Garg inequalities bear strong formal analogies to Bell-inequalities, except that whereas in a Bell-inequality one considers measurements occurring on two (or more) systems at spacelike separation, in a Leggett-Garg inequality, one considers repeated measurements, at different times, of a single observable, on a single system: a timelike, rather than a spacelike separation between measurements. For this reason, Leggett-Garg inequalities have often come to be called *temporal* Bell-inequalities, and, as with the Bell-inequalities proper, the intention is to rule out by experiment (or at least, to put to experimental test) a broad, interesting, and well-defined class of theories which might seem naturally appealing in some way, but which, on due reflection, have experimental implications which can be shown to conflict with the predictions of quantum mechanics.

Leggett-Garg inequalities have been a source of considerable current interest, having been the subject of a range of new experimental proposals, and, recently, of actual experimental tests (Knee et al., 2012; Dressel et al., 2011; Palacios-Laloy et al., 2010; Wilde and Mizel, 2012; Williams and Jordan, 2008; Jordan et al., 2006). Indeed, in the recent experiment of Knee et al. (2012), a test satisfying some of the important original strictures of Leggett and Garg has finally been achieved for the first time.¹ However, notwithstanding the contribution of some careful early commentators (Clifton, 1991; Foster and Elby, 1991; Elby and Foster, 1992; Benatti et al., 1994), there still remains considerable controversy about what exactly would be shown by violation of a Leggett-Garg inequality (Ballentine, 1987; Leggett and Garg, 1987; Leggett, 1988, 2002a,b). The key question, still in contention, is whether macroscopic realism would in fact be ruled out by experimental violation of the inequality.²

We will seek to gain clarity on this question by developing an adequately clear statement of what the position of macroscopic realism actually is. We will also ask how natural—and how interesting—the position thus arrived at might be as a realist response to the quantum behaviour of the micro-world. Our preliminary conclusions will be that the view that Leggett and Garg seem originally to have had in mind—essentially, ordinary quantum mechanics with a macroscopic superselection rule—is not a terribly natural realist position to

¹It should be noted that *no* test to date would plausibly be construed as involving macroscopic quantities, however. They have all been performed on microscopic systems. However the intention is to provide a proof of principle, with the hope of scaling-up to larger systems in due course. It should also be noted that (perhaps unsurprisingly) in all tests so far, the predictions of quantum mechanics have been borne out.

²To anticipate: The trouble is not with potential loopholes in the experiment, whether—in parallel to the Bell-case—with experimental imperfections such as noise and detector efficiency (though of course any good experimental implementation and analysis will need to attend to these important features), or with ‘logical’ loopholes such as retrocausality, conspiracy or superdeterminism (‘no free will’ etc., cf. Bell (1977)). Rather, the trouble lies deeper.

adopt, and we will compare it with some better-motivated realist alternatives. However, we will also show that by adopting a suitably general operationalist framework as a starting point, one can reformulate the Leggett-Garg inequality in a natural manner, and also go on to introduce a clear and operationally well-motivated macroscopic realist position.

As we shall presently see, in order to derive a Leggett-Garg inequality, one needs not only the assumption of macroscopic realism, but also the assumption of the existence of suitable *noninvasive measurements* (Leggett and Garg, 1985). We will clarify the notion of noninvasiveness by distinguishing between two distinct ideas: what we shall call *operational non-disturbance*³ and *ontic noninvasiveness*, respectively. Ontic noninvasiveness implies operational non-disturbance, but not conversely; whilst operational non-disturbance on its own is sufficient to derive the Leggett-Garg inequality. As we shall explain, therefore, what violation of the Leggett-Garg inequality primarily demonstrates is failure of one's measurements to be suitably operationally non-disturbing.

With this analysis in place, and by proceeding within an appropriately general framework, we will go on to demonstrate conclusively that violation of the Leggett-Garg inequality does not imply the falsity of macroscopic realism. We will show that there exist three broad classes of macroscopic realist theories, whilst it is only the first of these three—what we term *operational-eigenstate macroscopic realism*—which cannot allow violation of a Leggett-Garg inequality. The other two classes of macroscopic realist theories each permit violation of the inequality in a natural way which does not so much as go against the spirit (any more than it goes against the letter) of the macroscopic realist view that macroscopic quantities must always be determinately of one value or another (equivalently, that macroscopic systems must always be determinately in one macroscopic state or the other). At best, then, experimental violation of a Leggett-Garg inequality does not show the falsity of macroscopic realism *per se*, but only the falsity of operational-eigenstate macroscopic realism. This more narrow result remains far from trivial, however, for, as we shall explain, operational-eigenstate macroscopic realism is quite a natural and an interesting view, and ruling it out indicates certain fundamental limits in our ability directly to control the world experimentally, if one conceives of things within a macroscopic realist framework. Put another way, our results show that it is perfectly possible to remain a perfectly respectable macroscopic realist, whilst still accounting for the results which one might naturally think could not be encompassed within this conception.

A recurring theme of our discussion will be the question of the extent to which the case of the Leggett-Garg inequalities really parallels the more familiar case of the Bell-inequalities *methodologically*, as opposed merely to *formally* or *mathematically*. What is crucial to the significance of Bell inequalities is that their violation really does rule out a large, well-defined, and interesting class of theories *by experiment*. In order for this to be the case, the class of theories under test has to be defined in a suitably general manner—one which does not make too many ancilliary theoretical assumptions—and in particular, it needs to be the case that whether or not a given theory does, or ought to, satisfy the crucial conditions, can be determined in a suitably model-independent manner. Thus, in the Bell case, the general framework is one of *beable* theories (Bell, 1976),

³Clifton (1991) earlier called this notion *statistical noninvasive measurability*.

against an independently given background of relativistic causal structure, and the crucial condition of *local causality* (Bell, 1976, 1990b) can both be defined and *motivated* wholly independently of the details of any particular candidate theory which might be under test; in particular it can be defined independently of quantum theory.⁴ It is for these reasons that experimental violation of a Bell-inequality shows something significant directly about the world, namely, that no locally-causal theory of the world could be empirically adequate.⁵

It is our contention that, to date, discussion of the Leggett-Garg inequality has not been pursued in a suitably general and model-independent manner. In particular, discussion surrounding the crucial condition of noninvasive measurability and the question of its relation, if any, to macroscopic realism *per se* has not been, and—we shall argue—*cannot be*, made suitably model independent. In our view, therefore—and importantly—the methodological parallel with Bell's theorem fails.

⁴By way of example, this marks an important contrast with another recent question, that of the possibility or otherwise of experimental tests of contextuality, so-called. The standard notion of contextuality derives from the Kochen-Specker theorem (Kochen and Specker, 1967) but it is *defined* in wholly quantum-mechanical terms. One must assume that a great deal of quantum theory obtains in order even to deploy the definition. Whether a suitably theory-neutral notion of contextuality can be arrived at to give adequate sense to the notion of a true experimental test of whether or not nature is contextual remains a highly controversial question (see Barrett and Kent (2004); Spekkens (2005); Hermens (2011); and references therein, for discussion).

⁵Modulo, of course, some of the niceties to do with loopholes mentioned in fn.2. One should also note, of course, that failure of local-causality by no means automatically indicates the *presence* of non-local causation or action-at-a-distance (cf. Timpson and Brown (2002), for example).

2 The Leggett-Garg Inequality

Leggett and Garg begin their discussion by noting that:

“Despite sixty years of schooling in quantum mechanics, most physicists have a very non-quantum-mechanical notion of reality at the macroscopic level, which implicitly makes two assumptions.

- (A1) Macroscopic realism: A macroscopic system with two or more macroscopically distinct states available to it will at all times *be* in one or other of those states.
- (A2) Noninvasive measurability at the macroscopic level: It is possible, in principle, to determine the state of the system with arbitrarily small perturbation on its subsequent dynamics.

A direct extrapolation of quantum mechanics to the macroscopic level denies this.” (Leggett and Garg, 1985)

They invite us to consider the measurement on a given system—call it S —of a quantity Q which can take on one of two distinct values, $+1$ or -1 . If Q is some suitably macroscopic quantity, then it will follow from macroscopic realism (A1) that at any time, the system will have a definite value of Q , whether of $+1$ or -1 . Now consider what might happen if one were to make *pairs* of measurements of Q on S , with some time-lapse between the measurements. We might, for example, having prepared S in some standard way, measure Q at time t_1 and then at a later time t_2 . We might instead, following that preparation, have measured Q on S at time t_1 , and then at time t_3 , later than t_2 . Or equally, we might, as another alternative following the initial preparation, have measured Q on S at t_2 and then again at t_3 . (We may imagine that there is potentially some non-trivial dynamics operating on S as time evolves $t_1 \leq t \leq t_3$.) Call the value of Q at time t_i , Q_i .

Now, if the quantity Q *always* takes on a definite value, as will be the case if Q is suitably macroscopic and macroscopic realism obtains, then irrespective of whether or not Q is measured on S at a given time t_i , S will have a definite value of Q at t_i , that is, the Q_i belonging to S will all be well-defined and of value ± 1 for all $i \in \{1, 2, 3\}$, whether a measurement is performed at t_i or not. Under the assumption that the Q_i for S are well-defined for all $i \in \{1, 2, 3\}$ we can consider certain functions of them, such as the following:

$$Q_{LG} = Q_1Q_2 + Q_1Q_3 + Q_2Q_3. \quad (1)$$

This quantity Q_{LG} will have different values depending on what values of Q S actually has at the respective times, i.e., it will take on various values for different possible histories of S . A little thought reveals that the possible values for Q_{LG} are $+3$, and -1 , as follows (see Table 1).

Now, how might these possessed values of Q relate to values for Q one might measure in the lab, in particular, to the values that might be obtained in our three distinct pairs of measurement scenarios mentioned above? It is at this point that the assumption of noninvasive measurability, (A2), will come to the fore. Before that, however, let us first assume that our measurements of Q are good measurements, in the sense that if a system S has a definite value of Q

Q_1	Q_2	Q_3	Q_{LG}
+1	+1	+1	3
+1	+1	-1	-1
+1	-1	+1	-1
+1	-1	-1	-1
-1	+1	+1	-1
-1	+1	-1	-1
-1	-1	+1	-1
-1	-1	-1	3

Table 1: Possible values for Q_{LG} .

at the time of measurement, then this value is accurately revealed in the measurement.⁶ We will consider making measurements of Q on many copies of our given system, all of which have been subject to exactly the same preparation procedure. Each copy will also be subject to the same time evolutions in the intervals *between* possible measurements: $t < t_1, t_1 < t < t_2, t_2 < t < t_3$. Notice that, in general, even if we have prepared all our systems in the same way, and each has been subject to the same time-evolution before measurement, we may still get a statistical spread of results ± 1 when we measure Q . This is because, even on the assumption of macroscopic realism, and even given that our measurements are good measurements, it could well be that there are uncontrolled underlying variables at play that pertain to our systems, but which have not been fixed by our preparation procedure (we shall see much more of this notion as we proceed). Therefore, among the experimentally accessible and important quantities will be expectation values (evaluated for our ensemble of identically prepared systems) of the form $\langle Q_i \rangle$ and $\langle Q_i Q_j \rangle$. We will use subscripts M_i to denote expectation values determined when a measurement of Q at time t_i is actually performed. Thus $\langle Q_1 Q_2 \rangle_{M_1 M_2}$ denotes the expectation value for the product quantity $Q_1 Q_2$ which is found when a measurement M_1 of Q_1 and a subsequent measurement M_2 of Q_2 is actually made on each element of the ensemble. If our measurements are good measurements (and on the assumption that Q values are always definite), then an experimentally determined expectation value $\langle Q_i \rangle_{M_i}$ is telling us about the statistical spread of *possessed* values of Q in our ensemble at t_i , and an experimentally determined product expectation value $\langle Q_i Q_j \rangle_{M_i M_j}$ is telling us about the correlations between the *possessed* values at different times t_i and t_j in our ensemble.

Suppose we were to perform *three* rather than two successive measurements of Q , at times t_1, t_2 , and t_3 , on our ensemble. This would allow us to determine experimentally the three quantities $\langle Q_1 Q_2 \rangle_{M_1 M_2 M_3}$, $\langle Q_1 Q_3 \rangle_{M_1 M_2 M_3}$, and $\langle Q_2 Q_3 \rangle_{M_1 M_2 M_3}$. This would allow us, furthermore, to determine the quantity:

$$\langle Q_{LG} \rangle_{M_1 M_2 M_3} = \langle Q_1 Q_2 \rangle_{M_1 M_2 M_3} + \langle Q_1 Q_3 \rangle_{M_1 M_2 M_3} + \langle Q_2 Q_3 \rangle_{M_1 M_2 M_3}. \quad (2)$$

Since the possible values of Q_{LG} for a single system, are, as we know, either +3 or -1, then the expectation value for the ensemble must be bounded by these values:

$$-1 \leq \langle Q_{LG} \rangle_{M_1 M_2 M_3} \leq 3. \quad (3)$$

⁶Clifton (1991), following Redhead (1987) calls this property *faithful measurement*.

An important subtlety enters at this point. When all three measurements M_1, M_2, M_3 are actually successively performed on each element of the ensemble, then we do not need to assume that Q has definite possessed values prior to each measurement, which are then revealed in measurement, in order to know that inequality (3) obtains. We can equally well think (mathematically) of the entries in Table 1 as representing the *outcomes of actually performed measurements*, as opposed to *prior possessed values*. On either understanding of the quantities involved, Q_{LG} will be well-defined, though the *meaning* of the quantities is crucially different in the two cases.

The Leggett-Garg inequality finally comes into view when we introduce the assumption of noninvasive measurability. When all three measurements M_1, M_2 and M_3 are performed, since each measurement has a definite outcome of either +1 or -1, inequality (3) *must* always hold. The bound will be satisfied in quantum theory, in any macroscopic realist theory, and indeed in any other theory *at all* which defines joint probabilities of outcomes for a sequence of actually performed measurements! (This is a *very* large class of theories.) But now imagine that both macroscopic realism obtains *and* that it is possible to perform the first and second measurements of Q noninvasively. If both M_1 and M_2 are merely, if performed, accurately revealing a previously possessed value, and in so doing, are not at all affecting the future evolution of that value, then it will not make any difference to the possessed values later accurately revealed whether or not these measurements M_1 or M_2 were in fact performed. It follows that under these assumptions, the expectation value $\langle Q_1 Q_3 \rangle$ experimentally obtained when only M_1 and M_3 are performed should be the same as that which would have been obtained if all of M_1, M_2 , and M_3 had been performed; and similarly, the expectation value $\langle Q_2 Q_3 \rangle$ obtained when only M_2 and M_3 are performed should be the same as that which would have been obtained if all of M_1, M_2 and M_3 had been performed.⁷ Finally, by appealing to the fact that M_3 , if performed, is after M_2 , and so it shouldn't make a difference to things that happened before it, one can conclude that there should be no difference in the value of $\langle Q_1 Q_2 \rangle$ experimentally obtained if one were just to perform M_1 and M_2 , and leave out M_3 .

With these equalities in place, we derive from (3) the conclusion that:

$$-1 \leq \langle Q_1 Q_2 \rangle_{M_1 M_2} + \langle Q_1 Q_3 \rangle_{M_1 M_3} + \langle Q_2 Q_3 \rangle_{M_2 M_3} \leq 3. \quad (4)$$

This is the Leggett-Garg inequality, in one of its standard forms. Unlike the trivial (3) which brings out an entirely straightforward feature of a *single* experimental set-up, (4) relates in a highly non-trivial way three *distinct* experimental set-ups: the case in which we measure Q at t_1 and t_2 , the case in which we measure Q at t_1 and t_3 , and the case in which we measure Q at t_2 and t_3 . There is no reason at all to suppose that this inequality will hold in general theories, indeed, every reason to suppose it will not. But as we have seen, it *will* be satisfied if the special conditions of macroscopic realism and noninvasive measurability

⁷Clifton (1991) notes, applying to the Leggett-Garg case the very careful approach which Redhead developed in the parallel case of the Bell inequalities (Redhead, 1987), that there are one or two more further properties that need to be assumed to make the argument in this form watertight, but these needn't concern us at present. In the approach to derivation of the inequality which we shall adopt later-on, in so far as any issues remain, they are dealt with automatically by the framework we adopt.

hold.⁸ Leggett and Garg (1985) go on to show that it can readily be violated in quantum mechanics, and values less than -1 obtained.

The kind of physical system which Leggett and Garg originally proposed as an interesting candidate for investigating whether or not their inequality held was an rf SQUID (superconducting quantum interference device), i.e., a ring of superconductor (some millimeters in diameter) with a single Josephson junction, immersed in an external electromagnetic field. In such a device one can have the current in the superconductor circulating clockwise (label it +1) and one can have it circulating anti-clockwise (label it -1). Since the current (of the order of a few microamperes) involves a very large number of charge carriers in motion, it can plausibly be construed as in some relevant sense macroscopic. Quantum mechanics assigns orthogonal states $|+1\rangle$ and $|-1\rangle$ to these two distinct current states, and it will of course allow superpositions of such states: $\alpha|+1\rangle + \beta|-1\rangle$. Macroscopic realism, Leggett and Garg suggest, would require that the SQUID is only ever in one or other of the two distinct current states. The recent tests of the Leggett-Garg inequality mentioned above have explored other kinds of systems.

Leggett and Garg's derivation of their inequality, as above, uses, as we have seen, both the assumption of macroscopic realism, and the assumption of non-invasive measurability. Whilst Leggett and Garg note that noninvasive measurability does not logically follow as a consequence of macroscopic realism, they try to argue that it is nonetheless 'extremely natural and plausible' (Leggett and Garg, 1985) and moreover, they repeatedly assert—more strongly—that it is such a natural corollary of macroscopic realism *per se* that 'the latter is virtually meaningless in its absence' (Leggett, 2002a, p.R449), cf. also Leggett (1988, p.949). They motivate this thought by inviting us to consider *ideal negative result* or *null-result* measurements.

Suppose that there is some way of measuring Q which proceeds as follows: one has a measuring device which one knows will interact with the target system only if the target system is definitely in some one particular state of Q —say, only if $Q = +1$ —and it will not interact at all with the target system otherwise. If the system does have $Q = +1$, then the measuring device will reliably indicate this (it is a *good* measuring device, in the sense used before). Then suppose that macroscopic realism obtains and that one's target system to be measured enters this measuring device. If the device registers an outcome, we know that we are accurately revealing a pre-existing possessed value of Q (since we are assuming macroscopic realism, and the measurement is a good one) but our measurement may well disturb that value and its subsequent evolution. However, if the device (assumed to be working properly) *does not* register an outcome, we can infer that the value of Q is -1, for if it had been +1 the device would have registered, yet it did not; and +1 and -1 are the only two options. Yet the value of Q cannot have been disturbed in this measurement, since we know *ex hypothesi* that no interaction took place. This is a *null-result* measurement. The fact that *nothing happened* allows us to infer the value of Q , whilst the fact that *nothing happened* also guarantees that our learning this value does not affect the target system, thus does not affect the value of Q , or its subsequent evolution.

⁸Note for completeness that the derivation of (4) does not fundamentally require that Q be a macroscopic quantity. This only enters if the condition that Q must always have a definite value is to be derived from the assumption of macroscopic realism. One might proceed instead simply by assuming directly that Q always takes on definite values.

Notice that in this argument, the possibility of measuring Q noninvasively is derived from the *conjunction* of macroscopic realism with the assumption of the existence of suitable null-result measurements for Q . We shall return to examine in detail these arguments surrounding the connection between macroscopic realism and noninvasive measurability.

3 What is macrorealism? First pass

We have seen how the Leggett-Garg inequality (4) is supposed to function as a test for macroscopic realism. With noninvasive measurability a natural corollary of macroscopic realism, and the pair of macroscopic realism and noninvasive measurability together entailing the Leggett-Garg inequality, experimental violation of the inequality would defeat the macroscopic realist picture. But what exactly *is* macroscopic realism, and how realistic a realism is it, in fact? These are the questions to which we now turn.

At first blush, the statement of macroscopic realism (A1) above might seem perfectly clear—given a macroscopic system with two or more macroscopically distinct states available to it, it will at all times be in one or other of those states. But further reflection must give us pause. One might well be concerned with what exactly is meant by, or what would count as being, ‘macroscopic’, whether a macroscopic system, or a macroscopically distinct state. This concern, whilst real enough, is not our immediate object, however. Rather, our first concern is the perhaps more subtle question of what is meant by *state* in ‘macroscopically distinct states’. What conception of the state of a system are we operating with? Usually, the notion of state comes as part of a theory, or as part of a general framework of theories. What background theory, or framework of theories, are we considering here?

Leggett and Garg’s discussion strongly suggests that their background framework is simply quantum mechanical Hilbert space states. *Macroscopic states are the quantum states that one would assign to macroscopic, or collective, degrees of freedom.* Thus, in a SQUID, one does not trouble to assign a (massively entangled) multi-particle quantum state to the *enormous* number of individual microscopic charge-carriers, rather one simply assigns a single state to the collective degree of freedom, the direction of the current, e.g., $|+1\rangle$ or $|-1\rangle$. The content of macroscopic realism is then that the only permissible states of the SQUID are the quantum states $|+1\rangle$ and $|-1\rangle$ (and statistical mixtures thereof), quantum superpositions of these two states being disallowed.

This reading, according to which macroscopic realism is simply quantum theory subject to a macroscopic superselection rule which forbids superposition in certain regimes, emerges perhaps most clearly in the following passage of Leggett’s:

“...it seems clear that most if not all [criteria for macroscopicness] will have the property that the specific properties of a measuring apparatus as such are irrelevant to its allocation to the ‘macroscopic’ side of the divide; what is relevant is that the different final states of the apparatus are *macroscopically distinguishable*. Thus, we would expect that such theories would have the general feature that Nature, while known to tolerate linear superpositions at the atomic level, *cannot tolerate quantum superpositions of macroscopically distinct states*, whether or not these have anything to do with that small class of physical objects designed by human beings to act as measuring apparatus, but rather always selects a definite macroscopic state. Let us call this hypothesis for brevity ‘macrorealism’.” (Leggett, 1988, p.943)

He goes on:

“the hypothesis that Nature does not tolerate linear superpositions of macroscopically distinct states (‘macrorealism’) is in principle subject to experimental test.” (Leggett, 1988, p.944)

It seems that we are in the realm of applying quantum states to describe physical systems, it’s just that not all the usual states in the Hilbert space are permitted as physical. (It may be that we do not always get a new physical state by adding together two of the old ones; but the old ones—the items one is *contemplating* adding together—are themselves quantum states.)

However, if macroscopic realism is simply quantum theory with a macroscopic superselection rule, then we must note the following essential point:

There is nothing realist about denying the existence of superposition, macroscopic or otherwise.

Denying the physical possibility of quantum superposition (micro *or* macro) is neither necessary nor sufficient for realism, and it is not necessary in order to account for our determinate experience of the independently existing macroscopic world which surrounds us. Let us flesh this thought out.

To clarify terms, we take realism in the philosophy of science to be the familiar view—in brief—that our theories seek to give us a literally true description of what the world is like, both in its directly observable and its non-directly observable features; that the empirical success of our theories gives us good reason to believe (defeasible, but nonetheless *prima facie* good, reason to believe) that our theories are true or approximately true; and that the statements of our theories are true or false in virtue of mind-independent facts about the world. As applied to quantum theory, realism requires that we interpret the quantum formalism (or at least some significant part of it) as directly representing facts about the physical world. Moreover, since quantum mechanics is a fundamental theory, its scope, according to standard realism, should be considered as universal: it is apt, in principle at least, to describe the *whole* physical world, seamlessly, and all in one go. Furthermore, we should adopt an unassuming physicalism: reference to observers should play no role except in so far as we model these entities as physical systems within the theory;⁹ and more generally, we should grant big things to be made of little things, and insist that the laws, or other robust generalisations (if any), governing the behaviour of big things be consistent with the laws stipulated for little things: in this instance, consistent with the quantum laws governing the behaviour of the microscopic.

Now: One could deny superposition and yet fail to be a realist simply because one’s general view was not a realist one. For example, perhaps one maintains that one’s theory (quantum theory with a superselection rule) is not descriptive of *anything* apart from what the results of experiments would be, characterised at the level of the directly observable. (This would be an *instrumentalist* view, according to which one’s theory is just an algorithm for organising observable data, rather than a set of claims about how the world fundamentally is below the level of immediate observation). Or again, perhaps one might endorse the truly radical view that mind-independence fails, leading to some form of idealism or phenomenalism. Or one might deny some other component of the standard realist picture.

⁹Leggett effectively makes this same point in the first part of the quotation above (Leggett, 1988, p.943). Cf. also, famously, Bell (1990a).

More significant than this failure of sufficiency, however, is the fact that denial of superposition is not *necessary* for realism: one can seek to incorporate superposition, including macroscopic superposition, into one's realist, descriptive, account of how the mind-independent world is—incorporate it, moreover, in such a way as to recover the determinate nature of our experience, and of the macroscopic world. There are, of course, a number of well-developed approaches to interpreting quantum mechanics which adopt just this approach, and which are fully realist in the ways described above. We will focus initially on two key examples.

Consider first, then, the de Broglie–Bohm theory (de Broglie, 1927; Bohm, 1952; Bohm and Hiley, 1993; Holland, 1995). This is a theory in which the wavefunction for the total system (up to and including the entire universe) always evolves unitarily (thus superposition, and macroscopic superposition, is rife) but it is supplemented with the specification of definite positions at all times for particles. Momenta for these particles are also well-defined at all times (thus, all particles enjoy continuous and deterministic trajectories), and these momenta are determined by taking the gradient of the phase of the many-body wavefunction; in other words, the wavefunction is directly involved in pushing particles around. This is a deterministic hidden variable theory (the hidden variables are the particle positions, or more generally, configurations for physical degrees of freedom) and it is empirically equivalent to quantum theory, on the assumption (which may be stipulated or derived) that the probability distribution over initial configurations (understood merely as *ignorance* of the initial configuration) is given by the Born rule. (This is the assumption of the obtaining of *quantum equilibrium*, as per Valentini (1991a,b), for example.) In this theory, the results of measurements are uniquely determined given the initial configuration of the total system, the initial wavefunction, and the interaction Hamiltonian between measuring device and system measured. In this theory, one can have a wavefunction which is in a superposition, yet the values for physical quantities for the system take on definite values. For example, one's system of interest might be in a superposition of being here and there (where *here* and *there* are potentially a macroscopically large distance apart)—that is, the wavefunction has non-zero support in both these two widely distinct regions of configuration space—yet the system is definitely in one or other of the two places, as stipulated by the additional specification of the particle's (or particles') position(s).

The de Broglie–Bohm theory, then, is troubling from the point of view of our initial characterisation, following Leggett and Garg, of macroscopic realism. For this is a theory which allows macroscopic superposition, but yet, at the same time, also allows that the associated macroscopic physical variables take on definite values at all times. We conclude two things immediately, and notice a third.

First, one can evidently be a realist without denying superposition: the de Broglie–Bohm theory illustrates how superposition in the quantum state (including at the macroscopic level) need not lead to indefiniteness of physical variables. Second, and consequently, there must be something wrong with Leggett and Garg's conception of macroscopic realism as being the denial of the possibility of superposition (equivalently, being the assertion of quantum theory with a superselection rule). For, by one perfectly good rendering of the notion of 'being in a definite macroscopically distinct state', the de Broglie–Bohm the-

ory will allow that one’s system is in a definite (macroscopically distinct) state, viz., being either *here* or *there*, while yet it insists that one’s system is also in a quantum superposition. Third, we notice that the de Broglie–Bohm theory will evidently violate a Leggett–Garg inequality whenever quantum theory does, for de Broglie–Bohm theory is empirically equivalent to ordinary quantum theory (given quantum equilibrium obtains), yet it also, on the broader notion of macroscopic realism just delineated (definite values for macroscopic physical quantities), counts as being a macroscopic realist theory. So, the relationship between macroscopic realism and the necessity of satisfying a Leggett–Garg inequality is thrown into doubt.

Now let us consider another very well-developed version of quantum mechanical realism: the Everett interpretation (Everett, 1957; Saunders et al., 2010; Wallace, 2012). In this theory, as in de Broglie–Bohm, one entertains a wavefunction for the total system (up to and including the whole universe), again evolving purely unitarily. Here, however, one does not introduce further supplementary quantities, but argues instead that the unitarily evolving wavefunction on its own, construed realistically as directly representing a component of reality (a kind of non-separable field on four-dimensional spacetime, or a separable field on a very much higher-dimensional physical space (Wallace and Timpson, 2010)) is enough (given the contingent details of typical dynamics) to underpin emergent macroscopic definiteness: an emergent plurality of effectively non-interacting—thus independent—macroscopic worlds of determinate character; and correspondingly, is enough (on its own) to underpin the emergence of concrete observers, each having definite experiences, within these macroscopically definite (up to decoherence) worlds (Wallace, 2012).

The Everett interpretation shows us that we can have quantum realism without the denial of superposition. It shows us, moreover, that we can have quantum realism without macroscopic realism, even in the *broader* notion of macroscopic realism which consideration of the de Broglie–Bohm theory encouraged. Everett gives us realism about the macroscopic *without* macroscopic realism—in either the narrow Leggett–Garg sense of quantum mechanics with a superselection rule, or in the broader sense of macroscopic physical quantities always taking on a definite value.¹⁰

These reflections lead us to proffer a further slogan:

Macroscopic realism is not equivalent to realism about the macroscopic.

The de Broglie–Bohm theory shows us that the narrow initial Leggett–Garg characterisation of macroscopic realism as quantum theory with a superselection rule falls short of capturing the full notion of realism about the macroscopic (not to say, realism at all levels of description), whilst the Everett interpretation (in its modern formulations, due primarily to Saunders and Wallace) shows that even the *broadened* notion of macroscopic realism (definite values for macroscopic physical quantities) falls short of capturing the notion of realism about the macroscopic.

Regarding the Leggett–Garg inequality: it is obvious that the Everett interpretation allows violation of the inequality, however at the same time, it delivers determinate macroscopic worlds and determinate experience of those worlds.

¹⁰Everett offers us only the *relative states* being definite, or only *relative quantities* taking on a definite value, not the quantities *toute court*.

An alternative approach to quantum-mechanical realism is to assert the reality of the quantum state, as in de Broglie–Bohm theory and Everett, but to deny that the dynamics of the total system is always purely unitary. Theories in this class are realist wavefunction collapse theories, which postulate some additional term (or terms) in the law describing the evolution of microscopic systems, whose intended effect is to kill-off superposition between terms when it might seem that that’s required in order to give a definite result of an experiment—or more generally, to leave the macroscopic world in a determinate (non-superposed) state—but which have very little effect on microscopic superpositions. The Ghirardi-Rimini-Weber (GRW) theory (Ghirardi et al., 1986) is a well-known example of such an approach, one of the earliest to be constructed.

Concrete examples of realist collapse theories, such as GRW, will fall under the heading of Leggett and Garg’s original notion of macroscopic realism, and indeed, it is quite natural to take this class of theories to be just what Leggett and Garg originally had in mind. In these theories, the additional dynamical terms explain why the superselection rule—no macroscopic superposition—obtains. However, at this point, attention focusses back on what exactly is meant by ‘macroscopic’, and by ‘macroscopically distinct’. Different ways in which one seeks to implement the collapse mechanism in one’s collapse theory will lead to different stories about exactly which kinds of superposition are killed-off. Thus in many ways, Leggett and Garg’s ‘macroscopic realism’, as stated, fails to pick-out a natural kind of theories. Which basis (or approximate basis) will one achieve collapse to? Equivalently: with respect to which basis is superposition forbidden (at least on suitable timescales)? These questions can’t be answered unless one is presented with the detailed proposed theory.

By way of illustration, consider what the GRW theory would say about Leggett and Garg’s favoured example of the distinct macroscopic current states in a SQUID. As was shown by Rae (1990), even though GRW will quickly entail collapse to determinate states for macroscopic pointer variables of a measuring apparatus (spatially distinct states), it will not entail collapse to one or other of the distinct $|+1\rangle$ or $|-1\rangle$ current states of the SQUID. This is because the GRW theory works by spontaneous spatial localisation occurring to microsystems, but since the two distinct current states of the SQUID *wholly overlap* spatially, the localisation process makes almost no difference to either of these states, and certainly doesn’t distinguish between them. At best, the effect of GRW localisations will occasionally be to split-up a Cooper pair, and thus increase somewhat the normal dissipativeness of the superconductor; but it will not effect collapse. Thus GRW is a macroscopic realist theory which readily allows that the Leggett-Garg inequality could be violated for measurements on SQUIDs. In general, we will be unable to tell whether a given realist collapse theory will or will not require satisfaction of a Leggett-Garg inequality in a given experimental set-up unless we are told the details of the collapse process. It would seem that if a Leggett-Garg experiment is performed and violation of the inequality is shown, it could still be that a macroscopic realist theory in Leggett and Garg’s original sense is the true theory of nature, it is just one which does not impose collapse in the particular basis (or approximate basis) which has been deployed in the experiment. We conclude that so far, Leggett and Garg’s macroscopic realism has not, unlike Bell’s local causality, picked out a particularly natural class of theories which may be tested directly against experiment.

4 An operational formulation of the question

So far we have seen i) that denial of the possibility of macroscopic superposition cannot be the correct way, as it stands, of formulating macroscopic realism; ii) that there can be realist understandings of quantum theory which account for the determinate nature of the macroscopic world which surrounds us and, consequently, which account for the determinate nature of our experience, yet which readily admit (indeed, embrace) macroscopic superposition; and iii) that even having broadened the notion of macroscopic realism away simply from ‘no macroscopic superposition’, one can achieve realism about the macroscopic without having to be a macroscopic realist even in this broader sense. Moreover, we have seen that the question of whether or not macroscopic realism really does forbid violation of the Leggett-Garg inequality is unclear in general; and even in the relatively narrow case of the family of realist collapse theories, whether or not satisfaction of the inequality should hold for a given experimental set-up will depend on the details of the collapse mechanism, which (in principle at least) could vary really quite widely. It is time to go back to the drawing board.

Let us introduce a fairly familiar kind of operational formalism (cf. Spekkens (2005)) as a means of framing the general questions surrounding macroscopic realism and the Leggett-Garg inequality. Thus we will consider dividing-up any given experimental arrangement in the lab into three components, a *preparation* process, E , a *transformation* process, T , and a measurement process, M , having distinct outcomes taking values $Q = q_i$. We consider these three different components, from the point of view of the formalism, simply as black-boxes, with no assumption as to their internal workings. These different elements are to be identified operationally, and in the formalism, the experimental arrangement as a whole is characterised by a probability distribution $P_{(E,T,M)}(Q = q_i)$, being the probabilities that, given the preparation E , followed by the transformation T , the measurement M will have such-and-such outcomes. These probabilities are assumed to be measurable from the long-run statistics displayed by the experimental apparatus. Occasionally, for brevity, and where it will not lead to confusion, we may suppress certain indices. Note also that it is somewhat arbitrary how one divides a given experimental arrangement up—one might include a transformation as part of a preparation, for example, rather than treating it as a separate process; equally, one might include a transformation as part of the measurement procedure; or again, one might include a measurement’s having had a certain outcome as part of a preparation process, perhaps.

Preparations, transformations and measurements will naturally divide-up into equivalence classes. (Indeed, this is important in the very notion of separating out an experimental arrangement into three general kinds of process.)

1. Two different preparation procedures E_1 and E_2 will be operationally equivalent ($E_1 \simeq E_2$) *iff* for any transformation, and for any measurement, the same probability distribution over outcomes obtains whichever preparation procedure was performed:

$$E_1 \simeq E_2 \leftrightarrow \forall q_i, T, M P_{(E_1,T,M)}(Q = q_i) = P_{(E_2,T,M)}(Q = q_i).$$

2. Two different transformation procedures T_1 and T_2 are operationally equivalent ($T_1 \simeq T_2$) *iff* for any preparation and for any measurement, it makes

no difference to the probabilities for measurement outcomes which transformation took place:

$$T_1 \simeq T_2 \leftrightarrow \forall q_i, E, M P_{(E, T_1, M)}(Q = q_i) = P_{(E, T_2, M)}(Q = q_i).$$

3. Two different measurement procedures M_1 and M_2 will be operationally equivalent ($M_1 \simeq M_2$) *iff* their respective sets of outcomes $Q_1 = q_i$ and $Q_2 = q_j$ can be put into one-to-one correspondence, and whatever the preparation and whatever the transformation, there would be agreement on the probabilities assigned to outcomes for these measurements:

$$M_1 \simeq M_2 \leftrightarrow \forall q_i, E, T P_{(E, T, M_1)}(Q_1 = q_i) = P_{(E, T, M_2)}(Q_2 = q_i).$$

In this framework, an equivalence class of measurements naturally corresponds to some physical quantity or other.¹¹

Now let us go back to consider the following experimental set-up, familiar from our previous discussion. At time t_0 , a fixed preparation process E is performed on system S . At time $t_1 > t_0$, a measurement M_1 having two possible outcomes $Q_1 = \pm 1$ is performed on S . At time $t_2 > t_1$, a measurement M_2 , having two possible outcomes $Q_2 = \pm 1$ is performed on S . And finally, at time $t_3 > t_2$, a measurement M_3 , having two possible outcomes $Q_3 = \pm 1$ is performed on S . Between t_1 and t_2 , S is subject to the time evolution T_1 ; and between t_2 and t_3 , S is subject to the time evolution T_2 .

This whole arrangement is described by the joint probability distribution:

$$P_{(E, M_1, T_1, M_2, T_2, M_3)}(Q_1 = q_i, Q_2 = q_j, Q_3 = q_k). \quad (5)$$

Notice that at this stage we have made no assumption that the measurements M_1, M_2 and M_3 all belong to the same equivalence class, i.e., intuitively, that they are all measurements (perhaps in different ways) of one and the same physical quantity.

When a well-defined joint probability distribution exists, then we can of course derive the probability distribution for a smaller number of the same set of variables, simply by summing-out the remaining variables (i.e., by taking the marginal distribution). For example, if we just wanted to know the probabilities for the outcomes of the first measurement, we could derive it from (5) as follows (suppressing the indices for the initial preparation and the intermediate transformations, for brevity):

$$P_{(M_1, M_2, M_3)}(Q_1 = q_i) = \sum_{q_j q_k} P_{(M_1, M_2, M_3)}(Q_1 = q_i, Q_2 = q_j, Q_3 = q_k).$$

Similarly, if we were interested in the correlations between the outcomes of the first and second measurement in this set-up, we could calculate them from (5) as follows:

$$P_{(M_1, M_2, M_3)}(Q_1 = q_i, Q_2 = q_j) = \sum_{q_k} P_{(M_1, M_2, M_3)}(Q_1 = q_i, Q_2 = q_j, Q_3 = q_k);$$

¹¹Indeed, formally, though not metaphysically, one would want to *identify* a physical quantity with an equivalence class of measurement procedures, in this kind of framework.

and so on for other pairs of variables of interest. Expectation values for the outcomes of the measurements are then simply given by weighting the value of the outcomes by their probability, so, for example:

$$\langle Q_1 Q_2 \rangle_{M_1 M_2 M_3} = \sum_{q_i q_j} P_{(M_1, M_2, M_3)}(Q_1 = q_i, Q_2 = q_j) q_i q_j; \quad (6)$$

and so on.

Now return to reflect again on Table 1, and the quantity $Q_{LG} = Q_1 Q_2 + Q_1 Q_3 + Q_2 Q_3$. As noted before, when a joint probability distribution for the trio of variables Q_1, Q_2, Q_3 exists, as it does given (5), then evidently

$$-1 \leq \langle Q_{LG} \rangle_{M_1 M_2 M_3} \leq 3.$$

(Again, notice that here we have not, as we assumed before, required M_1, M_2 and M_3 all to be measurements of the same quantity, i.e., be in the same equivalence class. We are considering them all to be performed, however.)

As before, we will now seek to compare the behaviour of the experimental set-up where one performs all three of the measurements with the behaviour of the trio of distinct sub-experiments in each of which only two of the stated measurements are performed (though the preparation E and transformations T_1 and T_2 remain the same in all three sub-experiments). Of course, in a general operational probabilistic theory in which a joint probability distribution like (5) is defined, we can infer nothing about the joint probabilities for the pairs of values when only two experiments are performed. Since we are considering quite distinct experimental arrangements, the joint probabilities for the pairs of values in experiments in which only two measurements are performed are *not* constrained by the joint distribution for the triples of values when all three measurements are performed. Given that these are distinct experimental arrangements, the former quantities are not to be derived by taking the marginal distributions of the latter.

However, if the following highly non-trivial conditions on the probabilities obtain (again, supressing the indices for the preparation and the transformations), then the Leggett-Garg inequality as in (4) will follow (though here for the slightly more general case of possibly differing binary measurements M_1, M_2 and M_3):

$$P_{(M_1, M_2)}(Q_1 = q_i, Q_2 = q_j) = \sum_{q_k} P_{(M_1, M_2, M_3)}(Q_1 = q_i, Q_2 = q_j, Q_3 = q_k) \quad (7)$$

$$P_{(M_2, M_3)}(Q_2 = q_j, Q_3 = q_k) = \sum_{q_i} P_{(M_1, M_2, M_3)}(Q_1 = q_i, Q_2 = q_j, Q_3 = q_k) \quad (8)$$

$$P_{(M_1, M_3)}(Q_1 = q_i, Q_3 = q_k) = \sum_{q_j} P_{(M_1, M_2, M_3)}(Q_1 = q_i, Q_2 = q_j, Q_3 = q_k). \quad (9)$$

It is crucial to notice that, unlike in the standard Leggett-Garg proof, as above, we have here assumed nothing about the existence or otherwise of definite possessed values of physical quantities prior to measurement. We are working directly with the probabilities alone, which are themselves directly measurable by experiment.

4.1 Operational non-disturbance implies the Leggett-Garg Inequality

What do the conditions (7–9) above amount to? They relate the joint probabilities for the case where only *pairs* of measurements are performed to the marginal probabilities for pairs of variables in the case where all *three* measurements are performed.

The first condition, (7), is a very natural, and perhaps obligatory, constraint to do with time order: it says that whether or not one performs a measurement later-on should not affect the probability distribution over the outcomes of previously performed measurements. It is a requirement of *no signalling backwards in time*, so it can be assumed fairly uncontroversially.

The second two conditions, (8) and (9), are stronger than the time-ordering requirement. They are conditions that one's measurements M_1 and M_2 are respectively, in a sense to be made more precise in a moment, *operationally non-disturbing*. That is (roughly) it doesn't make any difference to your statistics whether or not you perform the measurements. This is clearly a very significant assumption to make.

Consider a preparation process E , followed by a measurement M , followed by a further measurement M' . As we shall define it, the measurement M is *operationally non-disturbing for the preparation E and subsequent measurement M'* iff it is not possible to tell, based upon the observed statistics of E and M' whether or not M_2 was performed. This will hold iff:

$$P_{(E,M')} (Q' = q_j) = \sum_{q_i} P_{(E,M,M')} (Q = q_i, Q' = q_j). \quad (10)$$

The case where you do not perform the intermediate measurement is statistically the same as the case in which the intermediate measurement does happen, but you do not observe the outcome.

One can then define stronger notions as follows: M might be operationally non-disturbing, given a prior preparation E , for *any* subsequent measurement M' ; M might be operationally non-disturbing for a fixed subsequent measurement M' for *any* prior preparation E . Most strongly of all, M might be operationally non-disturbing for *any* preparation E and *any* subsequent measurement M' ; in which case it will be said to be operationally non-disturbing *tout court*. This is evidently a very strong constraint indeed.

Condition (8) states that M_1 is operationally non-disturbing given the preparation E (index suppressed in the above) and the subsequent measurement M_2 . Condition (9) states that M_2 is operationally non-disturbing for the subsequent measurement of M_3 , given the prior preparation process which is constituted by the concatenation of E followed by M_1 's being performed and having a definite outcome.

Given the unexceptionable time-ordering condition (7), then, a Leggett-Garg inequality immediately follows if the measurements M_1 and M_2 are suitably operationally non-disturbing. Violation of the inequality, therefore, shows that one or other, or both, of these measurements must in fact be operationally disturbing. We note that we have had to say nothing here of macroscopic realism, pre-possessed values, or null-result measurements. The logical relation between operational non-disturbance and the Leggett-Garg inequality is simple and direct. Violation of the inequality means that at least one of the first two

measurements was operationally disturbing. The crucial question, to which we shall return in due course, is whether or not there is anything about macroscopic realism which suggests that operational non-disturbance should hold for the pertinent measurements.

5 Operational non-disturbance vs. Ontic non-invasiveness: The Ontic Models framework

With the condition of operational non-disturbance we have a very simple characterisation of a condition which is sufficient to entail that a Leggett-Garg inequality should be satisfied. Importantly, it is possible to establish directly by experiment whether or not a measurement M is operationally non-disturbing for a given preparation and subsequent measurement, simply by comparing the statistics in the case in which M is performed with those in the case in which it is not.

How does the notion of operational non-disturbance relate to Leggett and Garg's notion of noninvasive measurability which, recall, was crucial to the way that Leggett and Garg derive their inequality (as explained in Section 2)? It transpires that two distinct ways of making precise the notion of noninvasive measurability offer themselves. According to the first, noninvasive measurability simply equates to operational non-disturbance. According to the second, it amounts to a very much stronger condition. Both approaches seem to figure in some measure in Leggett and Garg's thought. In order to distinguish the two, we shall need to introduce some further formal apparatus which will provide the context for our discussion in the rest of the paper.

5.1 The ontic models framework

The framework we wish to introduce is due to Spekkens (2005) and is known as the *ontological models* or *ontic models* framework. (See also Harrigan and Spekkens (2010) and Leifer and Spekkens (2004).) It bears strong affinities to some ways in which hidden variable theories are often discussed, but it is somewhat more general and is usefully flexible.¹²

The idea is to supplement an operational probabilistic theory of the kind introduced above with some account of what gives rise to the observed probabilities. Thus a system has associated with it an *ontic state*, denoted λ , belonging to a space of states, Λ , for the system, where this ontic state captures the real physical properties of the system: properties which the system actually possesses independently of observation or measurement. For example, in n -body classical particle mechanics, the ontic state of the system would be the position in phase space; in views of quantum theory which are realist about the quantum state, the quantum state itself would be the ontic state.

The framework also associates with a system a probability distribution $\mu(\lambda)$ over the system's set of ontic states Λ . Thus, in particular, a given preparation process E (in the operational sense above) for a given system will produce a certain probability distribution $\mu_E(\lambda)$ over the system's set of possible ontic states ($\int_{\Lambda} \mu_E(\lambda) = 1$). If the preparation is very precise, this probability distribution may be a delta function—i.e., one has managed to ensure that the system is definitely in one and only one of its possible ontic states. But it may very well be that no such maximally fine preparations are possible, and that the best one can do is to produce a spread of possible ontic states, with some probabil-

¹²It is somewhat more general, for example, as the framework naturally incorporates wave-function realist versions of quantum theory, which are not usually thought of as hidden variable theories.

ity distribution over them. It will be convenient for later-on to introduce the notation

$$\text{supp}(\mu_i) := \{\lambda | \mu_i(\lambda) > 0\}$$

for the support of the probability distribution $\mu_i(\lambda)$, i.e., for the set of states $\lambda \in \Lambda$ for which the probability distribution $\mu_i(\lambda)$ is non-zero. A mixture (convex combination) of preparations E_i will give us another preparation: $\sum_i w_i \mu_{E_i}(\lambda) = \mu(\lambda)$, for $w_i \geq 0$, $\sum_i w_i = 1$.

Here are some familiar examples to illustrate the idea: 1. In statistical mechanics with the microcanonical distribution, we have a probability distribution $\mu(\lambda)$ over the ontic states of the system—our system is really lying somewhere on the fixed energy hypersurface, but we don't know where, and each possibility gets equal probability. 2. In realist quantum theory with collapse, our system is always really in some pure state, but it may well be that we don't know which (e.g. following a non-selective measurement process); the system is said to be in a *proper mixture* and we assign a probability distribution over the options.

Let us now consider measurement processes. Measuring devices M *respond* in some way to the system's being in the ontic state λ at the time of measurement. This is characterised by the *response function* $\xi_M(Q = q | \lambda)$, being the probability that the measuring device M will indicate the outcome value $Q = q$ given that the system is in ontic state λ on measurement. In general Q might take on continuous values; for our present purposes we need only consider discrete values q_i . The response of the measuring device to a particular λ might be deterministic, in which case the range of the response function would be $\{0, 1\}$, or—more generally—the response might be stochastic, in which case the range of the response function would be the full interval $[0, 1]$. Evidently, $\sum_{q_i} \xi_M(Q = q_i | \lambda) = 1$. Notice that even if two measurements M and M' belong to the same operational equivalence class ($M \simeq M'$) they may well not respond in the same way to a given ontic state λ , i.e., they may have different response functions. In this case, we have *contextuality*, a dependence on the *way* in which a given quantity is measured. (Spekkens (2005) calls this form of contextuality *measurement contextuality*.)

Transformation processes T are represented by stochastic maps from ontic states to ontic states: $\tau_T(\lambda | \lambda_0)$, giving the probability distribution over subsequent ontic states given that one started in the earlier ontic state λ_0 . Note that a preparation followed by a transformation gives us another preparation:

$$\mu_{(E,T)}(\lambda) = \int d\lambda_0 \mu_E(\lambda_0) \tau_T(\lambda | \lambda_0).$$

Putting all of preparation, transformation and measurement together, we now have:

$$P_{(E,T,M)}(Q = q_i) = \int d\lambda_0 d\lambda_1 \mu_E(\lambda_0) \tau_T(\lambda_1 | \lambda_0) \xi_M(Q = q_i | \lambda_1).$$

We also wish explicitly to note the effect that performing a measurement, and its having had a particular outcome, may have on the underlying ontic state. We formalise this using a particular transformation map which we associate with the measurement M : $\tau_M(\lambda | Q = q_i, \lambda_0)$. In other words, the probability distribution over ontic states post-measurement may depend on the particular measurement M performed, on the particular outcome $Q = q_i$ that occurred,

and on what the pre-measurement ontic state λ_0 was. Once more, even if two measurements belong to the same operational equivalence class, it may well be that they do not affect the ontic states of the system in the same way.

As we have said, realist collapse interpretations of quantum theory can be described in the ontic models framework; seeing this will help illustrate how the framework functions. In this case the ontic states λ are the pure states $|\psi\rangle \in \mathcal{H}$, where \mathcal{H} is the Hilbert space of the system. If one has a pure-state preparation procedure preparing the system in the state $|\phi\rangle$, then the probability distribution over ontic states generated by the preparation is simply $\mu_{|\psi\rangle}(\lambda) = \delta_{|\phi\rangle, |\psi\rangle}$. A mixed-state preparation is just a convex sum of these. The measurement response functions are as follows, for the simple case of projective measurement. For a measurement M which is associated with a particular orthonormal basis $\{|q_i\rangle\}$ of \mathcal{H} , we have $\xi_M(Q = q_i|\lambda) = |\langle q_i|\psi\rangle|^2$. If M is a projective measurement ‘of the first kind’ then $\tau_M(\lambda|Q = q_i, \lambda_0) = \delta_{|\psi\rangle, |q_i\rangle}$, i.e., the measurement doesn’t affect the system if it is already in an eigenstate of the quantity measured, and it leaves it in an eigenstate of the quantity measured otherwise.

A further illustration, which will be useful later-on, is given by the de Broglie–Bohm theory. Here the ontic state of a system is not given by the quantum state $|\psi\rangle$ alone, but by the ordered pair of quantum state and position in configuration space: $\lambda \in \{(|\psi\rangle, X)\}$. Now for a pure-state preparation process, the resultant probability distribution is

$$\mu_{|\phi\rangle}(\lambda) = \delta_{|\phi\rangle, |\psi\rangle} |\langle \phi|X|\phi\rangle|^2.$$

That is, even for the most refined preparation process one can achieve in the theory, there is still a non-trivial (in fact, Born rule) probability distribution over the ontic states. The measurement response functions in the de Broglie–Bohm theory are all *deterministic*, $\xi_M(Q = q_i|\lambda) \in \{0, 1\}$, but, importantly, they are *contextual*; the value assigned to the outcome of the measurement depends on which particular measurement process it is, not just on the equivalence class that the measurement belongs to.

With the ontic models framework now in hand, let us return to analysing the two distinct notions of noninvasive measurability.

5.2 Noninvasiveness as operational non-disturbance

One quite natural way of reading Leggett and Garg’s (1985) characterisation of noninvasive measurability at the macroscopic level (A2) is as follows. We are certainly used to thinking of the large-scale macroscopic objects we encounter in our daily lives as being composed of very large numbers of microsystems, the vast majority of whose detailed features make no difference at all to the coarse-grained properties and behaviour of these macroscopic objects which surround us. In a thermodynamic kind of way, the microscopic details wash-out: there are large numbers of different microstates a given object could be in, but exactly which it is in makes no difference at all to the object’s large-scale macroscopic features. The space of microstates naturally partitions into distinct sets of microstates, where the same macroscopically observable features obtain whichever of the microstates from a given such set the system happens to be in.

Noninvasive measurability at the macroscopic level might then seem intuitively plausible, since one might suppose that—granted one may very well dis-

turb the *microstate* of the system when one observes it (exchange of momentum from photons bouncing off, or what have you)—but one typically does not change the *macroscopic* features one is measuring, one only moves the microstate around inside the set of macroscopically-equivalent microstates. It might then also seem plausible (but more tenuously so, perhaps) to assume that the microscopic change one has induced will not affect transition probabilities between macroscopic states of the system in the future. In this case, noninvasive measurability is essentially a statistical notion and is equivalent to our notion of operational non-disturbance.

To help formalise this, let us introduce the notion of an *operational eigenstate*. Take an equivalence class of measurements $\tilde{M} = \{M' | M' \simeq M\}$, and call the corresponding physical quantity, which each $M' \in \tilde{M}$ can be thought of as measuring, \tilde{Q} . We define an operational eigenstate of \tilde{Q} to be a particular equivalence class of preparation processes. A preparation E is in such an equivalence class *iff*, following E , any measurement $M' \in \tilde{M}$ would give the same $\{0, 1\}$ probability distribution over the outcomes $Q' = q_i$ of the measurement. That is, if one's system has been prepared in the q_i operational eigenstate of \tilde{Q} then any measurement $M' \in \tilde{M}$ will return the value q_i with probability 1 (and conversely).

Measurement of some macroscopic quantity can be performed noninvasively in the sense just articulated if there is some measurement M in the operational equivalence class corresponding to the macroscopic quantity in question which is operationally non-disturbing when the system begins in an operational eigenstate of the macroscopic quantity, and given a subsequent measurement of that quantity.

Leggett (1995) seems to have had this statistical notion of noninvasive measurability (equivalent to our notion of operational non-disturbance) in mind when he phrased the condition as follows:

“(P3) (‘noninvasive measurability’): it is in principle possible to measure the value $Q(t)$ on an ensemble without altering the *statistical properties* of that ensemble as regards subsequent measurements.” (Leggett, 1995, p.104, emphasis added.)

This formulation would seem to allow that there could be an effect on microscopic features when we are considering a macroscopic system, but that the observable statistics should not change.

Two things to note. First, the requirement of noninvasive measurability, understood in the way we have been suggesting, does not require that *every* way of measuring a quantity \tilde{Q} be operationally non-disturbing, but only that there be *some* measurement M of \tilde{Q} which is. Second, whilst we have said there is some intuitive plausibility in the story about the possibility of noninvasive measurability and its connection to the notion of the macroscopic (one might make changes at the microscopic level, but these may not affect things at the macroscopic level) this cannot be taken too far. Perhaps in certain circumstances one hasn't noticed anything suggesting that one's observations of macroscopic objects have changed them, or most especially, changed their future behaviour, but very rarely has one actually checked in any detail to see.

5.3 Noninvasiveness as ontic noninvasiveness

The second way of understanding noninvasive measurability is as a much stronger condition, what we shall call *ontic noninvasiveness*. A measurement M is ontically noninvasive if it does not change at all the ontic state of the system that obtained prior to measurement, whatever that state might have been. M will be ontically non-invasive for a given outcome $Q = q_i$ iff the transformation τ_M associated with the process of measurement is as follows:

$$\tau_M(\lambda|Q = q_i, \lambda_0) = \delta_{\lambda, \lambda_0}.$$

If M is ontically noninvasive for any outcome value q_i , then it is ontically non-invasive *tout court*.

Ontic noninvasiveness is a very strong condition on a measurement process. It is evidently enough to entail operational non-disturbance, and thence the Leggett-Garg inequality, as follows. If M is ontically noninvasive then for any preparation E and for any subsequent measurement M' , there will be no difference to the statistics for M' whether or not M was performed, that is, M 's being ontically noninvasive entails *complete* operational non-disturbance for M (i.e. operational non-disturbance *tout court*). Therefore M will satisfy the weaker operational non-disturbance conditions (8–9) above (we consider performing the same type of measurement M at the two different times t_1 and t_2), and the Leggett-Garg inequality will follow. But the converse implication doesn't hold. The two weaker operational non-disturbance conditions (8–9) certainly do not entail that M is ontically noninvasive, and even M 's being completely operationally non-disturbing does not entail that M is ontically non-invasive, unless it is always possible to prepare arbitrary (and in particular, arbitrarily sharp) probability distributions over one's ontic states, and this will certainly not hold in general in theories. Therefore, operational non-disturbance is clearly a significantly logically weaker notion than ontic noninvasiveness.

Now suppose that the measurement M has only two possible outcomes, and that it is only ontically noninvasive for one of these, say $Q = +1$ (the argument extends in the obvious way for the more general case of a larger number of outcomes). If there is another measurement, M' , in the same operational equivalence class as M , which also happens to be ontically noninvasive, but this time only for $Q = -1$, then these two measurement processes, along with a post-selection procedure can be combined to produce a scenario in which the total effect is that of complete ontic noninvasiveness. The procedure is simple. On an ensemble of identically prepared systems, half the time one should perform M and half the time one should perform M' . If the outcome of M was $+1$ one keeps the result, and knows that the ontic state of the system on that run of the experiment was not affected by the measurement; if the outcome of M' was -1 one keeps the result, and again knows that the ontic state of the system on that run of the experiment was not affected. In all other cases, the data is discarded. The effect of this post-selection procedure is to select-out a sub-ensemble of the total ensemble where the ontic state of each of the members of the ensemble is guaranteed not to have been affected by the process of measurement. Since M and M' belong to the same operational equivalence class, they can of course be thought of as both measuring the same physical quantity.

Leggett and Garg's central, explicit, argument that noninvasive measurability should hold in the case of macroscopic realism proceeds, recall, by invoking

the notion of null-result measurements. The idea is to post-select just those results where no result of the measurement was observed, so where it is believed that nothing happened, so *a fortiori* it is believed that nothing happened to the ontic state of the system on that run, on the assumption that the system must have one or other definite value of the quantity being measured before measurement. In the manner just described, by combining a suitable pair of null-result measurements and post-selecting, the result will be that the ontic states of the systems in one's ensemble, under these assumptions, will not have been affected by the measurement, thus it will be the case that complete ontic noninvasiveness holds for the process as a whole, when one includes post-selection.

The structure of Leggett and Garg's argument from null result measurements shows clearly that they certainly have ontic noninvasiveness in mind as the relevant notion of noninvasiveness, at least at certain crucial junctures of their thought.

5.4 Summary

To summarise some of our conclusions from the last two Sections, then:

1. The specific operational non-disturbance conditions for M_1 and M_2 , that M_1 is operationally non-disturbing for the preparation/measurement pair (E, M_2) , and that M_2 is operationally non-disturbing for the preparation/measurement pair $((E, M_2), M_3)$, suffice (given time-order) to entail the Leggett-Garg inequality:

$$\text{OPND}_{\text{specific}} \longrightarrow \text{LGI}.$$

N.B. the converse does not hold: satisfying the inequality is not sufficient to show that the measurements are operationally non-disturbing.

2. Complete operational non-disturbance for the measurements M_1, M_2 , entails specific operational non-disturbance (but not conversely), so:

$$\text{OPND}_{\text{complete}} \longrightarrow \text{OPND}_{\text{specific}} \longrightarrow \text{LGI}.$$

3. Ontic noninvasiveness entails complete operational non-disturbance, but not conversely. Measurements which are only optically noninvasive for one measurement outcome can be combined with post-selection to achieve complete ontic noninvasiveness.

$$\text{ONI} \longrightarrow \text{OPND}_{\text{complete}} \longrightarrow \text{OPND}_{\text{specific}} \longrightarrow \text{LGI}.$$

It is clear therefore that what Leggett-Garg inequality violation most immediately shows is that one or other, or both, of M_1 and M_2 were operationally disturbing for their respective preparations and measurements. These are the weakest conditions that suffice to entail that the inequality should hold. One would need specific reason to believe that one of the stronger conditions in fact held in nature in order for there to be a reason to appeal to (potential) Leggett-Garg inequality violation to rule it out. Notice once more that in none of the chains of inference listed above have we had to invoke macroscopic realism to

arrive at the inequality. If we take Leggett and Garg's notion of noninvasiveness simply to be (specific) operational non-disturbance (as we have seen there is some reason to), then the inequality follows directly from that premise alone, without requiring that macroscopic realism holds. Leggett and Garg suggest that macroscopic realism combined with the existence of null result measurements allows us to infer that the very strong condition of ontic noninvasiveness holds. In that case, violation of the inequality would disprove the *conjunction* of macroscopic realism and the existence of suitable null result measurements. We defer assessing this argument until after we have a workable statement of macroscopic realism on the table.

6 What is macrorealism? Second pass

We saw in Section 3 that Leggett and Garg’s notion of macrorealism did not seem to be capturing any particular natural class of theories which might be considered to instantiate forms of macroscopic realism. Leggett and Garg’s suggestion seems to be based upon the idea of a superselection rule, preventing linear superpositions of macroscopically distinct quantum states. Unfortunately this suggestion is far too tied to the quantum formalism to make an appropriate model independent criterion in itself. We will now try and refine the notion in a suitably model independent form.

We approach this through Leggett’s comment that “what is relevant is that the different final states of the apparatus are *macroscopically distinguishable*” (Leggett, 1988, p.943). If two states are macroscopically distinct, there must be a macroscopically observable difference between them. A macrorealist, then, would believe that these macroscopically observable properties always have determinate values, at all times.

Macroscopically observable properties are the outcomes of macroscopic observations: the chair is in such and such a place, the car is some particular colour. This suggests rephrasing macrorealism as:

A macroscopically observable property with two or more macroscopically distinct values available to it will at all times determinately possess one or other of those values.

The state of the world is captured by the ontic state λ . If M is the macroscopic observation of a given property, \tilde{Q} , macrorealism requires that the response function $\xi_M(Q = q_i|\lambda) \in \{0, 1\}$ for all possible ontic states.

This must hold for all such measurements, and these measurements must all agree on the value of the macroscopic property: so the physical location of the chair does not depend upon how we look at it. For any two operationally equivalent macroscopic observations, $M \simeq M'$, any given ontic state must give the same response: $\forall \lambda \xi_M(Q = q_i|\lambda) = \xi_{M'}(Q = q_i|\lambda)$.

However, the conjunction of these two criteria is well known in quantum theory, as it means that $Q = q_i$ is non-contextually value definite. This might raise a concern that macrorealism is already ruled out by the Kochen-Specker theorem (Kochen and Specker, 1967), which forbids any non-contextually value definite formulation of quantum theory. Fortunately for the macrorealist, the Kochen-Specker theorem only entails that *some properties* are contextual¹³. The macrorealist need not be committed to the idea that all the properties of a system are determinate, only the macroscopically observable properties. To avoid confusion, we will use the term ‘macrodefinite’ for ‘non-contextually value definite for all macroscopically observable properties’.

We have not attempted here to characterise what, exactly, are the macroscopically observable properties, though they may be expected to be a very coarse-grained mutually commuting subset of the quantum observables.

Finally, we may now go on to characterise two states as macroscopically distinct if, and only if, there is at least one macroscopically observable property to which the two states have a zero probability of assigning the same value.

¹³Though it does entail that this must be the case for every possible ontic state.

With this definition, macrorealism expressed about observable properties recovers macrorealism expressed about states.

7 What does Leggett-Garg Inequality violation in fact rule out?

The arguments surrounding macroscopic realism and the Leggett-Garg inequality can be boiled down into consideration of the following simplified case.¹⁴ The macroscopic realist regarding a macroscopic quantity \tilde{Q} believes that all ontic states λ are non-contextually value-definite for \tilde{Q} (are macrodefinite for \tilde{Q}). Now consider a simple experimental sequence of: preparation—measurement of \tilde{Q} —measurement of \tilde{Q} , i.e., an (E, M_1, M_2) sequence, where $M_1 \simeq M_2 \in \tilde{M}$, \tilde{M} being the operational equivalence class of measurements which corresponds to the quantity \tilde{Q} .

Suppose our macroscopic realist has carefully tested the features of an M_1 measurement and has established conclusively by experiment that for any operational eigenstate preparations of \tilde{Q} , M_1 is operationally non-disturbing, for any subsequent measurement $M_2 \in \tilde{M}$. This amounts to saying that, whenever \tilde{Q} is prepared as having a definite value, M_1 will not disturb the statistics for later measurements of \tilde{Q} . Evidently, if a measurement is operationally non-disturbing for a given set of preparations, it will also be non-disturbing for convex combinations of those preparations; so M_1 is operationally non-disturbing for statistical mixtures of operational eigenstates of \tilde{Q} too. If M_1 is (E, \tilde{M}) -operationally non-disturbing, then as we show in the Appendix, an abbreviated form of the Leggett-Garg inequality also holds.

Now suppose that we perform the experiment (E, M_1, M_2) many times, and compare the resultant statistics with those for the sequence (E, M_2) . (The time evolutions, other than at measurements, are kept the same in the two cases.) Since the macroscopic realist already believes that M_1 is operationally non-disturbing whenever \tilde{Q} is prepared as having either one value or the other, and since the macroscopic realist believes that \tilde{Q} *must* always take on one value or the other, they will be very surprised if our experiments now reveal that M_1 was *in fact* disturbing—that the statistics for the two sets of experiments differ after all. How could this be, given that we have already checked that the measurement will *not* be disturbing when we've prepared a definite value for \tilde{Q} ? Must we conclude that it *cannot in fact be the case* that \tilde{Q} always takes on a definite value, i.e., must we conclude that macroscopic realism is false? If there are ontic states which are not macrodefinite for \tilde{Q} then this would explain how it can be that the measurement was in fact disturbing, when we'd already established that it *would not be* if the systems being measured were prepared as having a definite value of \tilde{Q} . This is the core argument in the whole discussion of macroscopic realism.

Suggestive as it is, this argument is not valid as an argument against macroscopic realism, as we shall now see. Careful reflection shows that macroscopic realism can come in a range of sub-varieties, and it is only *one* of these sub-varieties which can be refuted in this way, by showing that one's measurement was after all operationally disturbing—equivalently, by showing that a Leggett-Garg inequality is violated.

¹⁴We clarify this formally in an Appendix below, but it should be intuitively clear enough in what follows, in any case.

7.1 Macrorealism 1: Operational Eigenstate Macroscopic Realism

The macroscopic realist (for a macroscopic quantity \tilde{Q}) believes that all ontic states $\lambda \in \Lambda$ are macrodefinite for \tilde{Q} . One might also believe that one has (in principle at least) the possibility of exceedingly fine control in one's experimental preparation processes. That is, one might believe that there is no probabilistic state $\mu(\lambda)$ that a system can be in, which cannot be directly and controllably prepared (in principle at least) in the lab. (This would be the case in standard quantum theory, for example.¹⁵) In this case one will think that any possible $\mu(\lambda)$ can be given by a lab preparation E , so is equal to $\mu_E(\lambda)$ or to a convex combination (statistical mixture) of such preparable distributions. Putting together this appealing idea of the possibility of experimental control with macroscopic realism for \tilde{Q} , we reach the following view:

The only possible states of a system S are operational eigenstates of \tilde{Q} and statistical mixtures thereof.

Call this view *operational eigenstate macroscopic realism*. It is clearly a very natural way for a macroscopic realist to think: Every state of the world is one I can prepare, and every state of the world is one in which \tilde{Q} really has a definite value.

Putting things a little more formally, a preparation E_{q_i} belongs to the q_i operational eigenstate equivalence class \tilde{E}_{q_i} iff $P_{(E_{q_i}, M)}(Q = q_i) = 1$, for all $M \in \tilde{M}$. Call the probability distribution prepared by E_{q_i} , $\mu_{E_{q_i}}(\lambda)$. Then for all λ and for all $M \in \tilde{M}$, $\lambda \in \text{supp}(\mu_{E_{q_i}}) \leftrightarrow \xi_M(Q = q_i|\lambda) = 1$. Denote by $\nu_{q_i}(\lambda)$ an arbitrary mixture (convex sum) of operational eigenstate preparation distributions $\mu_{E_{q_i}}(\lambda)$. Then according to the operational eigenstate macroscopic realist, every physically possible state $\mu(\lambda)$ is given by a sum $\sum_{q_i} w_{q_i} \nu_{q_i}(\lambda)$. Every ontic state λ in the support of ν_{q_i} is non-contextually value-definite, with value $Q = q_i$, and every ν_{q_i} is reliably preparable in the lab, as is every convex sum of such distributions. On this view, the full ontic state space Λ is exhausted by the union for all q_i of the supports of the distributions ν_{q_i} which arise from the operational eigenstate preparations.

Now, if operational eigenstate macroscopic realism obtains, then if one has checked that a measurement M_1 is operationally non-disturbing for operational eigenstate preparations, it will be *quite impossible* for the statistics for (E, M_1, M_2) to differ from those for (E, M_2) , i.e., for M_1 to be operationally disturbing in this configuration, i.e., for a Leggett-Garg inequality to be violated, for the simple reason that *there are no states available which are not operational eigenstates or mixtures thereof*. Thus should M_1 prove to be operationally disturbing, or equivalently, a Leggett-Garg inequality to be violated, we must reject operational eigenstate macroscopic realism.

¹⁵This shows the force of the 'in principle' clauses. In practice, it is extraordinarily difficult reliably to prepare in the lab all the states that quantum theory in principle allows that one can. But these extraordinary and interesting practical difficulties are not usually taken to suggest that it is not in-principle possible. Hence the effort directed towards understanding and constructing quantum computers, for example.

7.2 Macrorealism 2: Preparation-support Macroscopic Realism

Now consider a different view. Suppose it remains the case that every ontic state is macrodefinite for \tilde{Q} —so the view is macroscopic realist—and suppose that it remains the case that the full space Λ of ontic states is still given by the union of the supports of the distributions which can be arrived at by taking convex sums of operational eigenstate preparations, just as in operational eigenstate macroscopic realism. Thus no ontic states exist which cannot be accessed by an operational eigenstate preparation (i.e., which do not fall into the support of some $\mu_{E_{q_i}}(\lambda)$). However, suppose that the set of possible probability distributions $\mu(\lambda)$ is larger than the set of convex combinations of operational eigenstate preparation distributions, so that whilst we have the same ontic state space as in operational eigenstate macroscopic realism, the allowed distributions over that space differ. Finally, suppose in addition that it is not the case that it is always possible to *prepare* an arbitrary probability distribution over the set of macrodefinite ontic states, i.e., that the set of preparable probability distributions is strictly a subset of the set of all physically possible distributions over ontic states: $\{\mu_E(\lambda)\} \subset \{\mu(\lambda)\}$. Thus, there are some restrictions on which from the physically possible distributions can reliably be prepared in the lab.¹⁶ Call this view *preparation-support macroscopic realism* (since all the ontic states are in the support of some operational eigenstate preparation.)

In such a theory it is quite clear how a Leggett-Garg inequality can be violated/a measurement turn out surprisingly to be operationally disturbing, while yet the theory remains fully macrorealist. One might have demonstrated that a measurement M_1 is operationally non-disturbing for every *preparation* that has been performed, without it being the case that it is operationally non-disturbing for every probability distribution $\mu(\lambda)$ that might arise during the course of the experiment. So for example, even if one starts with an operational eigenstate preparation (or a convex combination of them)—for which one has indeed established that M_1 would be operationally non-disturbing—it could be that before M_1 takes place, the $\mu(\lambda)$ has evolved away from this initial operational eigenstate distribution, and it could well be that the *new* distribution *is* disturbed by M_1 . This would be something that we simply hadn't checked for. Thus, the effect of M_1 may definitely be *ontically invasive*: it may shuffle the ontic states within the support of the initial distribution around, but that it does so in such a way that the overall probability distribution does not change, when that distribution is a mixture of operational eigenstate distributions. However, once the the distribution has evolved away from being a mixture of operational eigenstate distributions, then the shifting around of the ontic states which M_1 induces is such as to disturb the distribution in a way which may be observed.

In preparation-support macroscopic realism, what allows a Leggett-Garg inequality to be violated is not the existence of ontic states which are not macrodefinite for \tilde{Q} , but the existence of novel probability distributions over the ontic states which are not given by mixtures of operational eigenstates. One's measurements of \tilde{Q} might be operationally disturbing for the former, without being operationally disturbing for the latter.

An example of a theory of this kind is given by the original Kochen-Specker

¹⁶Of course, this entails that there are also some restrictions on what transformations can reliably be implemented.

non-contextual hidden variable model for a two-state system (Kochen and Specker, 1967).¹⁷

7.3 Macrorealism 3: Extra-preparation Macroscopic Realism

In the final variety of macroscopic realism, it of course once more remains the case that all ontic states are macrodefinite for \tilde{Q} , but that here, in contrast to both the operational eigenstate and the preparation-support versions of macroscopic realism, there are (macrodefinite) ontic states λ which do not fall into the support of any operational eigenstate preparation. When one reliably prepares a system in a definite state of \tilde{Q} (when, that is, one prepares the system in an operational eigenstate) one is only accessing *part* of the macrodefinite-for- \tilde{Q} ontic state space. There are other states λ which are still macrodefinite for \tilde{Q} but which one can't reliably get one's system into. They may only arise when one has some non-zero probability of obtaining each of the possible outcomes $Q = q_i$ of a measurement $M \in \tilde{M}$.

This view we can call *extra-preparation* macroscopic realism, as there are macrodefinite ontic states which are not contained within the support of any operational eigenstate preparation.

We saw earlier that the de Broglie–Bohm theory caused considerable trouble for Leggett and Garg's discussion of macroscopic realism and violation of the Leggett–Garg inequality, since the de Broglie–Bohm theory certainly seemed, by one perfectly good measure, to be macroscopic realist, but yet it allowed violation of the Leggett–Garg inequality just as ordinary quantum mechanics does. The de Broglie–Bohm theory is precisely an example of extra-preparation macroscopic realism, where the macroscopic quantity \tilde{Q} is (possibly coarse-grained) position. Every ontic state $\lambda = (|\psi\rangle, X)$ is non-contextually value-definite for position, but when $|\psi\rangle = \alpha|q\rangle + \beta|q'\rangle$, the distribution $\mu(\lambda)$ over ontic states again cannot be given as a convex sum of operational eigenstate distributions, here because novel—but still macrodefinite—ontic states are being accessed. The Leggett–Garg inequality can readily be violated in the de Broglie–Bohm theory, then, as, once more, we simply have not checked—when checking that our measurement M_1 was operationally non-disturbing for operational eigenstate preparations—that our measurement M_1 is operationally non-disturbing for all distributions over macrodefinite ontic states that there can be.

¹⁷Granted: it is well known that such a model cannot be extended to higher dimensions, when a sufficient number of distinct, non-compatible, observables are introduced. But this does not entail that such models cannot work for a small enough set of observables. (Macroscopic realism is not the doctrine that one's ontic states must be non-contextually value-definite for *all* quantities.) Indeed, for macroscopic quantities, it is quite plausible that all observable quantities should be *compatible*, so the issue is not, so far, especially troubling.

8 Macrorealism and null result measurements: Macrorealism does not imply non-invasiveness

Let us at long last come to consider in detail Leggett and Garg’s arguments that macroscopic realism entails—or at least very strongly supports—noninvasiveness, in the strong sense of ontic noninvasiveness. If this were so, then Leggett-Garg inequality violation really would cause trouble for macroscopic realism *per se*, rather than merely for the weaker form of operational eigenstate macroscopic realism. A significant part of our argument so far has been that weaker conditions than macroscopic realism *per se* are sufficient on their own to entail the Leggett-Garg inequality so that macroscopic realism itself need not be impugned by violation of the inequality. Again, if Leggett and Garg’s argument for ontic noninvasiveness works, then since ontic noninvasiveness entails the weaker conditions, which entail the Leggett-Garg inequality, macroscopic realism itself would be in danger, rather than (merely) the logically weaker conditions.

The argument for ontic noninvasiveness as a corollary of macroscopic realism goes as follows. Suppose we have available a pair of null-result measurements for \tilde{Q} : we believe that $M'_1 \in \tilde{M}$ will interact with our system S iff S has a definite value $+1$ (say) of \tilde{Q} , and we believe that $M''_2 \in \tilde{M}$ will interact with S iff S has a definite value -1 of \tilde{Q} . If macroscopic realism obtains (but not otherwise), then when M'_1 does not fire, we may infer that the value of \tilde{Q} was -1 , and that there can have been no affect on S , since there was no interaction. *Mutatis mutandis* for when M''_1 does not fire. Then by post-selection on cases of no-firing, we can be assured of complete ontic noninvasiveness.

Let us tighten this argument up. In the framework we have adopted it is not possible to formalise the notion of a null-result measurement directly, in part for the very good reason that it is not at all a notion on which one may get direct operational grip. That one might believe a measurement M'_1 only to involve interaction with S in certain circumstances depends on what one’s detailed model of the physical interaction is. Different physical models of the interaction might disagree about whether that was the case.

Nevertheless, one might postulate as follows. Suppose we simply stipulate that $M'_1 \in \tilde{M}$ is such that every λ which is non-contextually value-definite for the -1 outcome is not affected by M'_1 , and that M''_1 is such that every λ which is non-contextually value-definite for the $+1$ outcome is not affected by M''_1 . If (and only if) every ontic state is macrodefinite for ± 1 , then following post-selection on the -1 outcome for M'_1 and on the $+1$ outcome for M''_1 , we will have a process which overall is guaranteed to be completely ontically noninvasive. We do not talk of a ‘null-result’ here, as any measurement $M \in \tilde{M}$ is always symbolised as having a result $Q = \pm 1$; however we can still call these measurements the representations in our formalism of what Leggett and Garg have in mind. So we have shown that in our formalism, with this characterisation of null result measurements, macroscopic realism and the existence of null result measurements like M'_1 and M''_2 do indeed entail complete ontic noninvasiveness, and thence on down the chain:

$$(\text{MR and NRM}) \rightarrow \text{ONI} \rightarrow \text{OPND}_{\text{complete}} \rightarrow \text{OPND}_{\text{specific}} \rightarrow \text{LGI}.$$

Of course, one immediate remark to make on this argument is that ontic noninvasiveness does not follow from macroscopic realism alone, but only from macroscopic realism with the additional *stipulation* of measurements which are ontically noninvasive for given outcomes for certain of the macrodefinite ontic states. And that there should exist such measurements is clearly no part at all of the notion of macroscopic realism; one could deny it and lose nothing. The claim that all ontic states are macrodefinite does not require, in order for it to be sure of meaning, that some kinds of measurements of the macroscopic quantity should be partially ontically noninvasive (cf. Leggett (1988, p. 949), Leggett (2002b, p.R449)). To be sure, and for all its importance, we do not feel it likely that Leggett or Garg would be much inclined to dispute this point. Rather, it seems to us, their main thought is that *given* that one has a set of measurements in mind which would function as null-result measurements, then it would be all but impossible to maintain macroscopic realism if those measurements were to turn out to be operationally disturbing. This is of a piece with the idea that one cannot hope to turn to just *any old* set of measurements, even of macroscopic quantities, in order to set-up an interesting test of the Leggett-Garg inequalities: one had better have some good reason to believe in the first place that the macroscopic realist will think that the measurements ought to be operationally non-disturbing, for to repeat, their claim is *not* that *every* measurement need be noninvasive, whether in the sense of operationally non-disturbing or in the sense of ontically invasive. If one can find *some* measurements that the macroscopic realist might be inclined to think should be noninvasive, then one's Leggett-Garg test can begin to get off the ground.

However, what remains as a point of fundamental importance is that there is nothing *model independent* that can be appealed-to to establish whether or not one should think of one's measurements as being partially ontically non-invasive in the way described above. If one happens to have certain views as to how the detailed physics of the interaction between system and measuring apparatus goes, then one might very well believe that one had a pair of measurements apt for null-result, ontically non-invasive, measurement. But it could well be that one's model is wrong, rather than that it is macroscopic realism which is at fault. Here we find a very fundamental difference with the case of the Bell inequalities. That one's theory should be locally causal followed in a model-independent way from the setting of relativistic causal structure: spacelike separation automatically motivated a certain kind of independence condition. Nothing of the sort holds here—there is no general principle, or independent court of appeal, which might suggest that one's measurements should be partially ontically non-invasive. This can only be a model-dependent hypothesis.

Leggett and Garg (1987) and Leggett (1988, p.950) suggest that it may be possible to get an independent grip on the required property of ontic noninvasiveness directly by experiment 'at least up to a point' (Leggett, 1988, p.950): they consider (what is in our terminology) testing operational non-disturbance for operational eigenstate preparations. It is certainly true that this is an essential move when beginning discussion of when a macroscopic realist should believe a Leggett-Garg inequality should hold. Without first settling that the measurements in question are operationally non-disturbing for operational eigenstates *no* macroscopic realist is compelled to believe that a Leggett-Garg inequality should follow. But as we have seen, establishing such operational non-disturbance falls enormously short of establishing anything like ontic noninvasiveness, and whilst

the operational eigenstate macroscopic realist would be defeated if a subsequent test were then to show operational disturbance, we have seen that there are two other respectable macroscopic realist positions which have no trouble at all in incorporating operational disturbance for measurements which have previously been shown to be operationally non-disturbing for operational eigenstates.

In a telling passage, Leggett and Garg maintain:

“Should anyone wish to interpret the results of our proposed experiment (assumed for present purposes to agree with QM) by saying that the macroscopic object is indeed in a definite macroscopic state but is *nevertheless physically affected by an interaction which we know could have occurred only if it had been in a different macroscopic state* he or she is free to do so; we leave it to the reader to judge whether such an interpretation in any way diminishes the force of the quantum measurement paradox, or the significance of our proposed experiment.” (Leggett and Garg, 1987, emphasis added.)

But the crucial point is: how would one know what one is supposed to know here, if the point is to carry any force? How do we *know* that an interaction could have occurred only if the system had been in the other state? (How do we know that the measurement was optically noninvasive for the outcome recorded?) The answer is that we cannot, in a model independent manner. We can only *assume*, and our assumption may well be wrong.

9 Conclusions

We have seen that macroscopic realism should be understood not as the claim that certain kinds of quantum superposition are not possible, but as the claim that all ontic states are non-contextually value-definite for a macroscopically observable quantity \tilde{Q} . We have shown that macroscopic realism would not be impugned by a Leggett-Garg inequality violation involving measurements of \tilde{Q} . Within the notion of macroscopic realism *per se* we have seen that there are three distinct broad kinds of theories: operational eigenstate macroscopic realism, preparation-support macroscopic realism, and extra-preparation macroscopic realism. It is only the first of these which is unable to account for potential Leggett-Garg inequality violation. Nevertheless, even if Leggett-Garg inequality violation does not refute macroscopic realism, it would still remain an interesting result, since operational eigenstate macroscopic realism follows from macroscopic realism proper when combined with the idea that one is able experimentally (in principle at least) to prepare every possible probabilistic state which the world allows. If macroscopic realism is true, then, Leggett-Garg inequality violation would show that we are subject to fundamental limits in our ability directly to control and manipulate the world.

Regarding the Leggett-Garg inequality itself, we showed, first, that the inequality could be derived directly from some simple conditions pertaining to whether or not one's measurements were operationally disturbing. This derivation needed to make no mention of macroscopic realism, or of any notion of non-invasiveness stronger than non-disturbance at the statistical level. We analysed Leggett and Garg's arguments from the possibility of null-result measurements to the conclusion that noninvasiveness in a stronger sense than operational non-disturbance for operational eigenstate preparations of \tilde{Q} was a natural corollary of the notion of macroscopic realism, and found that no such argument could be maintained without appeal to potentially tendentious and model-dependent assumptions. It follows that, for all their mathematical similarity, the Leggett-Garg inequalities and the Bell-inequalities are not methodologically on a par. Unlike local causality, ontic noninvasiveness cannot be motivated as a general feature which should exist for at least some measurements when macroscopic realism obtains.

Appendix: The role of operationally detectable disturbance in violating the Leggett Garg Inequality

To draw out the role of operationally detectable disturbances in violating the Leggett Garg Inequality, it helps to introduce a few simple expressions. Starting with the marginal statistics for the experimental arrangement when all three measurements M_1, M_2, M_3 are performed:

$$\begin{aligned} P_{(M_1, M_2, M_3)}(Q_1 = q_i, Q_2 = q_j) &= \sum_{q_k} P_{(M_1, M_2, M_3)}(Q_1 = q_i, Q_2 = q_j, Q_3 = q_k) \quad (11) \\ P_{(M_1, M_2, M_3)}(Q_2 = q_j, Q_3 = q_k) &= \sum_{q_i} P_{(M_1, M_2, M_3)}(Q_1 = q_i, Q_2 = q_j, Q_3 = q_k) \quad (12) \\ P_{(M_1, M_2, M_3)}(Q_1 = q_i, Q_3 = q_k) &= \sum_{q_j} P_{(M_1, M_2, M_3)}(Q_1 = q_i, Q_2 = q_j, Q_3 = q_k) \quad (13) \end{aligned}$$

it is useful to define:

$$\begin{aligned} D_{1, q_j, q_k} &= P_{(M_2, M_3)}(Q_2 = q_j, Q_3 = q_k) - P_{(M_1, M_2, M_3)}(Q_2 = q_j, Q_3 = q_k) \\ D_{2, q_i, q_k} &= P_{(M_1, M_3)}(Q_1 = q_i, Q_3 = q_k) - P_{(M_1, M_2, M_3)}(Q_1 = q_i, Q_3 = q_k) \end{aligned} \quad (14)$$

to quantify the change in marginal statistics of M_2, M_3 , depending on whether M_1 is measured, and in the marginal statistics of M_1, M_3 depending on whether M_2 is measured.

If the introduction of the M_1 measurement cannot be detectable by an observer who only has access to the marginal statistics of M_2, M_3 , then all of $D_{1, q_j, q_k} = 0$. Similarly, if the introduction of the M_2 measurement is undetectable with only the marginal statistics of M_1, M_3 , then all of $D_{2, q_i, q_k} = 0$. An equivalently defined D_{3, q_i, q_j} must all be zero to avoid signalling backwards in time.

Now, looking at each term in the Leggett Garg Inequality, when all three measurements are performed:

$$\begin{aligned} \langle Q_1 Q_2 \rangle_{M_1 M_2 M_3} &= P_{(M_1, M_2, M_3)}(Q_1 = +1, Q_2 = +1) + P_{(M_1, M_2, M_3)}(Q_1 = -1, Q_2 = -1) \\ &\quad - P_{(M_1, M_2, M_3)}(Q_1 = +1, Q_2 = -1) - P_{(M_1, M_2, M_3)}(Q_1 = -1, Q_2 = +1) \\ \langle Q_1 Q_3 \rangle_{M_1 M_2 M_3} &= P_{(M_1, M_2, M_3)}(Q_1 = +1, Q_3 = +1) + P_{(M_1, M_2, M_3)}(Q_1 = -1, Q_3 = -1) \\ &\quad - P_{(M_1, M_2, M_3)}(Q_1 = +1, Q_3 = -1) - P_{(M_1, M_2, M_3)}(Q_1 = -1, Q_3 = +1) \\ \langle Q_2 Q_3 \rangle_{M_1 M_2 M_3} &= P_{(M_1, M_2, M_3)}(Q_2 = +1, Q_3 = +1) + P_{(M_1, M_2, M_3)}(Q_2 = -1, Q_3 = -1) \\ &\quad - P_{(M_1, M_2, M_3)}(Q_2 = +1, Q_3 = -1) - P_{(M_1, M_2, M_3)}(Q_2 = -1, Q_3 = +1) \end{aligned} \quad (15)$$

A straightforward rearrangement yields the very simple form

$$\langle Q \rangle_{M_1 M_2 M_3} = 4 \left(P_{(M_1, M_2, M_3)}(Q_1 = +1, Q_2 = +1, Q_3 = +1) + P_{(M_1, M_2, M_3)}(Q_1 = -1, Q_2 = -1, Q_3 = -1) \right) - 1 \quad (16)$$

showing clearly that the Leggett-Garg Inequality cannot be violated.

When only pairs of measurements are performed:

$$\begin{aligned}
\langle Q_1 Q_2 \rangle_{M_1 M_2} &= P_{(M_1, M_2)}(Q_1 = +1, Q_2 = +1) + P_{(M_1, M_2)}(Q_1 = -1, Q_2 = -1) \\
&\quad - P_{(M_1, M_2)}(Q_1 = +1, Q_2 = -1) - P_{(M_1, M_2)}(Q_1 = -1, Q_2 = +1) \\
\langle Q_1 Q_3 \rangle_{M_1 M_3} &= P_{(M_1, M_3)}(Q_1 = +1, Q_3 = +1) + P_{(M_1, M_3)}(Q_1 = -1, Q_3 = -1) \\
&\quad - P_{(M_1, M_3)}(Q_1 = +1, Q_3 = -1) - P_{(M_1, M_3)}(Q_1 = -1, Q_3 = +1) \\
\langle Q_2 Q_3 \rangle_{M_2 M_3} &= P_{(M_2, M_3)}(Q_2 = +1, Q_3 = +1) + P_{(M_2, M_3)}(Q_2 = -1, Q_3 = -1) \\
&\quad - P_{(M_2, M_3)}(Q_2 = +1, Q_3 = -1) - P_{(M_2, M_3)}(Q_2 = -1, Q_3 = +1)
\end{aligned} \tag{17}$$

and it can now be easily seen that if

$$\langle Q \rangle_{LG} = \langle Q_1 Q_2 \rangle_{M_1 M_2} + \langle Q_1 Q_3 \rangle_{M_1 M_3} + \langle Q_2 Q_3 \rangle_{M_2 M_3} \tag{18}$$

then

$$\langle Q \rangle_{LG} - \langle Q \rangle_{M_1 M_2 M_3} = \left(\sum_{q_j = q_k} D_{1, q_j, q_k} + \sum_{q_j = q_k} D_{2, q_i, q_k} \right) - \left(\sum_{q_j \neq q_k} D_{1, q_j, q_k} + \sum_{q_j \neq q_k} D_{2, q_i, q_k} \right) \tag{19}$$

This can be simplified by noting that $\sum_{q_j, q_k} D_{1, q_j, q_k} = \sum_{q_i, q_k} D_{2, q_i, q_k} = 0$, and it follows that

$$\begin{aligned}
\langle Q \rangle_{LG} &= 4 \left(P_{(M_1, M_2, M_3)}(Q_1 = +1, Q_2 = +1, Q_3 = +1) + P_{(M_1, M_2, M_3)}(Q_1 = -1, Q_2 = -1, Q_3 = -1) \right) \\
&\quad + 2 \left(\sum_{q_j = q_k} D_{1, q_j, q_k} + \sum_{q_j = q_k} D_{2, q_i, q_k} \right) - 1
\end{aligned} \tag{20}$$

For the Leggett Garg Inequality to be violated it is a necessary condition that at least one of the set $D_{1, q_j, q_k}, D_{2, q_i, q_k}$ be non-zero. At least one of the M_1 or M_2 measurements must introduce a disturbance which is detectable in the marginal statistics of the other two measurements.

The expression above suggests that a simpler form is possible, obtained by preparing the system to actually be in a operational eigenstate of M_1 , with the value $Q_1 = +1$. In this case all $D_{1, q_j, q_k} = 0$. As $P(Q_1 = -1) = 0$ for any measurements on such an initial preparation, it will also be the case that $P_{(M_1, M_2, M_3)}(Q_1 = -1, Q_2 = -1, Q_3 = -1) = 0$ and that $D_{2, -1, q_k} = 0$. This produces

$$\langle Q \rangle_{LG} = 2D_{2, +1, +1} + 4P_{(M_1, M_2, M_3)}(Q_1 = +1, Q_2 = +1, Q_3 = +1) - 1 \tag{21}$$

Acknowledgements

We thank George Knee, Erik Gauger, Andrew Briggs, Simon Benjamin, and Rob Spekkens for helpful discussion. This work was supported by a grant from the John Templeton Foundation.

References

- Ballentine, L. E. (1987). Realism and quantum flux tunneling. *Physical Review Letters*, 59(14):1493–1495.
- Barrett, J. and Kent, A. (2004). Non-contextuality, finite precision measurement and the Kochen-Specker theorem. *Studies in History and Philosophy of Modern Physics*, 35:151–176.
- Bell, J. S. (1976). The theory of local beables. *Epistemological Letters*. Repr. in *Dialectica* 39:86 (1985) and in Bell (1987), Chpt. 7.
- Bell, J. S. (1977). Free variables and local causality. *Epistemological Letters*. Repr. in *Dialectica* 39:103 (1985) and in Bell (1987), Chpt. 12.
- Bell, J. S. (1987). *Speakable and Unspeakable in Quantum Mechanics*. Cambridge University Press. Second Edition (2004).
- Bell, J. S. (1990a). Against ‘measurement’. *Physics World*. August, pp.33-40. Repr. in Bell (1987) Second Edition.
- Bell, J. S. (1990b). La nouvelle cuisine. In Sarlemijn, A. and Kroes, P., editors, *Between Science and Technology*. Elsevier. Repr. in Bell (1987) Second Edition, Chpt. 24.
- Benatti, F., Ghirardi, G., and Grassi, R. (1994). On some recent proposals for testing macrorealism versus quantum mechanics. *Foundations of Physics Letters*, 7(2):105–126.
- Bohm, D. (1952). A suggested interpretation of the quantum theory in terms of hidden variables, I and II. *Physical Review*, 85:166–179;180–193.
- Bohm, D. and Hiley, B. J. (1993). *The Undivided Universe: An Ontological Interpretation of Quantum Theory*. Routledge, London.
- Clifton, R. (1991). Noninvasive measurability, negative-result measurements and watched-pots: Another look at Leggett’s arguments for the incompatibility between between macro-realism and quantum mechanics. In Lahti, P. and Mittelstaedt, P., editors, *Symposium on the Foundations of Modern Physics 1990*. World Scientific.
- de Broglie, L. (1927). The new dynamics of quanta. Trans. in *Quantum Mechanics at the Crossroads: Reconsidering the 1927 Solvay Conference*, Bacchiagaluppi G and Valentini A editors, Cambridge University Press 2009.
- Dressel, J., Broadbent, C., Howell, J., and Jordan, A. (2011). Experimental violation of two-party Leggett-Garg inequalities with semiweak measurements. *Physical Review Letters*, 106:040402.

- Elby, A. and Foster, S. (1992). Why squid experiments can rule out non-invasive measurability. *Physics Letters A*, 166:17–23.
- Everett, III, H. (1957). “Relative state” formulation of quantum mechanics. *Reviews of Modern Physics*, 29:454–62.
- Foster, S. and Elby, A. (1991). A squid no-go theorem without macrorealism: What squid’s really tell us about nature. *Foundations of Physics*, 21(7):773–785.
- Ghirardi, G. C., Rimini, A., and Weber, T. (1986). Unified dynamics for microscopic and macroscopic systems. *Physical Review D*, 34:470–491.
- Harrigan, N. and Spekkens, R. W. (2010). Einstein, incompleteness, and the epistemic view of quantum states. *Foundations of Physics*, 40:125–157.
- Hermens, R. (2011). The problem of contextuality and the impossibility of the experimental metaphysics thereof. *Studies in History and Philosophy of Modern Physics*, 42(4):214–225.
- Holland, P. R. (1995). *The Quantum Theory of Motion: An Account of the de Broglie-Bohm Causal Interpretation of Quantum Mechanics*. Cambridge University Press, first paperback edition.
- Jordan, A., Korotkov, A., and Buttiker, M. (2006). Leggett-Garg inequality with a kicked quantum pump. *Physical Review Letters*, 97:026805.
- Knee, G., Simmons, S., Gauger, E., Morton, J., Riemann, H., Abrosimov, N., Becker, P., Pohl, H., Itoh, K., Thewalt, M., Briggs, G., and Benjamin, S. (2012). Violation of a Leggett-Garg inequality with ideal non-invasive measurements. *Nature Communications*, 3:606.
- Kochen, S. and Specker, E. (1967). The problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics*, 17:59–87.
- Leggett, A. (1988). Experimental approaches to the quantum measurement paradox. *Foundations of Physics*, 18(9):939–952.
- Leggett, A. (1995). Time’s arrow and the quantum measurement problem. In Savitt, S., editor, *Time’s Arrows Today: Recent Physical and Philosophical Work on the Direction of Time*, pages 97–106. Cambridge University Press.
- Leggett, A. (2002a). Probing quantum mechanics towards the everyday world: where do we stand? *Physica Scripta*, T102:69.
- Leggett, A. (2002b). Testing the limits of quantum mechanics: Motivation, state of play, prospects. *Journal of Physics: Condensed Matter*, 14:R415–R451.
- Leggett, A. and Garg, A. (1987). Comment on “Realism and quantum flux tunneling”. *Physical Review Letters*, 59(14):1621.
- Leggett, A. J. and Garg, A. (1985). Quantum mechanics versus macroscopic realism: Is the flux there when nobody looks? *Physical Review Letters*, 54(9):857–860.

- Leifer, M. and Spekkens, R. (2004). Logical pre- and post-selection paradoxes, measurement-disturbance and contextuality. [arxiv.org://quant-ph/04121](https://arxiv.org/abs/quant-ph/04121).
- Palacios-Laloy, A., Mallet, F., Nguyen, F., Bertet, P., Vion, D., Esteve, D., and Korotkov, A. (2010). Experimental violation of a Bell's inequality in time with weak measurement. *Nature Physics*, 6:442–447.
- Rae, A. I. M. (1990). Can GRW theory be tested by experiments on SQUIDS. *Journal of Physics A*, 23:L57–L60.
- Redhead, M. (1987). *Incompleteness, Nonlocality and Realism: A Prolegomenon to the Philosophy of Quantum Mechanics*. Oxford University Press.
- Saunders, S., Barrett, J., Kent, A., and Wallace, D. (2010). *Many Worlds? Everett, Realism and Quantum Mechanics*. Oxford University Press.
- Spekkens, R. (2005). Contextuality for preparations, transformations and unsharp measurements. *Physical Review A*, 71:052108.
- Timpson, C. G. and Brown, H. R. (2002). Entanglement and relativity. In Lupacchini, R. and Fano, V., editors, *Understanding Physical Knowledge*, pages 147–66. University of Bologna, CLUEB, Bologna. [arXiv:quant-ph/0212140](https://arxiv.org/abs/quant-ph/0212140).
- Valentini, A. (1991a). Signal-locality, uncertainty, and the sub-quantum H -theorem. I. *Physics Letters A*, 156(1, 2):5–11.
- Valentini, A. (1991b). Signal-locality, uncertainty, and the sub-quantum H -theorem. II. *Physics Letters A*, 158(1, 2):1–8.
- Wallace, D. (2012). *The Emergent Multiverse: Quantum Mechanics according to the Everett Interpretation*. Oxford University Press.
- Wallace, D. and Timpson, C. G. (2010). Quantum mechanics on spacetime I: Spacetime state realism. *British Journal for the Philosophy of Science*, 61(4):697–727.
- Wilde, M. and Mizel, A. (2012). Addressing the clumsiness loophole in a Leggett-Garg test of macrorealism. *Foundations of Physics*, 42(2):256–265.
- Williams, N. and Jordan, A. (2008). Weak values and the Leggett-Garg inequality in solid-state qubits. *Physical Review Letters*, 100:026804.