Truth vs. Progress Realism about Spin

Juha Saatsi

3.1 Introduction

Spin, a theoretical concept at the heart of quantum theory, is scientifically as firmly established as any. Spin underlies numerous explanations and predictions in physics and chemistry, as well as rapidly growing number of technological feats. Its central and multifaceted theoretical role strongly motivates a scientific realist attitude: we should be "realists about spin" — as much as we are realists about anything in physics. But what does this realism amount to? I will answer this question by distinguishing in very general terms two conceptions of scientific realism, problematising one of them, and articulating and defending the other. One of the conceptions I call *truth-content realism*. It incorporates both "standard" realism defended by, e.g., Hilary Putnam, Richard Boyd, Stathis Psillos, as well as epistemic structural realism, construed as a principled qualification of theories' truth content by reference to theories' or the world's structure. The other conception of realism — the one I favour — focuses its commitments on theoretical progress in a way that renounces typical realist claims of (approximate) truth and knowledge about the unobservable world. I will argue that this more minimal *progress realism* offers a defensible positive epistemic attitude towards a theory such as quantum mechanics, while truth-content realism problematically involves "deep" metaphysics not supported by the overall empirical evidence.

Traditionally much of the epistemological debate surrounding realism (e.g. relating to the historical record of science and underdetermination of theories by evidence) has been conducted in highly general and abstract terms (as has been customary in general philosophy of science), concerning all of "mature" science in one fell swoop.[1] Recent years have seen growing interest in anchoring

[1] These epistemological issues should be distinguished from the metaphysical question of

the debate much more firmly to local, case-specific details of particular theories and their evidential support. This support can arguably substantially vary, upon philosophical analysis, from one theory or area of science to another, as it is prima facie plausible that the balance between the evidence for and against realism gets settled differently for different areas of science. Given the vast differences in theories' explanatory presuppositions and subject matters, there is no reason to think that theory-mediated epistemic access to them is equally (un)problematic across different areas of successful, well-established science as diverse as, e.g., cosmology, genetics, geology, palaeontology, and quantum physics.[2] It thus makes sense to have a good look at the details of specific areas of science to see how the case for or against realism plays out in different scientific domains. Quantum physics is particularly interesting in this regard, as it exhibits an increasingly broadly recognised and scientifically serious case of underdetermination by evidence that seems to breathe new life to a venerable challenge to scientific realism. This underdetermination has been discussed from different angles by various authors, including many represented in this volume. In this chapter I aim to highlight some subtleties of this challenge, arguing that its implications are best appreciated in the context of the broader dialectic of the realism debate and through more general reflections on what scientific realism amounts to.

It will pay off to further zoom in on spin to bring into focus the nature of the problematic underdetermination. For standard realism 'spin' is a central theoretical term of exceedingly well-confirmed science, taken to refer to an objective feature of mind-independent world. Traditionally realist commitments to our best theories' (approximate) truth are bound up with such terms' referential status, furnishing a sense in which we can gain knowledge from these theories as representations of reality, while also expecting future revisions in these theories. Prima facie, realism about spin is extremely well motivated by the lights of usual realist arguments, given spin's various theoretical roles (cursorily reviewed in Section 2). Briefly put: spin is at the heart of various scientific *explanations* furnished by quantum physics, exemplifying the kind of explanatory reasoning upon which standard abductive arguments for realism capitalise. Spin is also *unifying* in ways that matter for many realists. Finally, the increasingly robust theoretical handle on spin has led to myriad technological developments over the past couple of decades, rendering spin increasingly *manipulable* in ways that many realists have regarded as critical to justifying realism. In the spirit of the realist's classic "no miracles" argument, we can sum this up by saying that the multifaceted success of quantum physics would be an astounding miracle if spin wasn't, well, real! Regardless of whether we endorse the letter of this argument, the prima facie case for realism about spin is forceful.

---

what is real according to a given theory (taken as approximately true). The latter question can only be answered in the context of specific theories, of course, and quantum theory has been an area of intense research in this regard. See Richard Healey's chapter in this volume for further distinctions between various realism issues.

[2] For illustrations, see the various contributions to Saatsi (2018, Part IV).

So runs a homely realist line of thought. But, as is well known, spin is also at the heart of interpretational conundrums of quantum physics, so it is immediately unclear what "realism about spin" is *about.* I will argue that underdetermination of different candidate theories of spin — correctly understood — raises a serious challenge to the received way of understanding what realism amounts to. Although this challenge arises specifically in the context of quantum physics (which is notoriously "weird" from our classical point of view), it also threatens the realist outlook on science much more broadly by demonstrating the unreliability of standard criteria for realist commitment towards indispensable theoretical notions more generally, since these criteria apply to spin just as well as they apply to any theoretical feature of science. This is one reason why it is difficult to delimit the implications of this challenge in a principled way to quantum physics, so as to give up on realism selectively in relation to the latter, while simultaneously withholding realism in relation to other areas of theoretical science.

One such strategy for limiting realist commitments has been proposed by Carl Hoefer (this volume), who adopts a selectively instrumentalist attitude towards all current *fundamental* theories of physics (including quantum physics), while retaining a realist attitude towards science for which we "cannot conceive of any way that they could be seriously mistaken, other than by appealing to scenarios involving radical skepticism." Hoefer offers various examples of epistemologically kosher objects of realist commitment from *non-fundamental* physics, including, e.g., the key properties of atoms and molecules employed in much of chemistry and statistical mechanics, and the key properties of electrons employed in much of electronics. He argues that realism about atoms and electrons (thus construed) can be underwritten by the fact that science without them is in a strong sense *inconceivable*, due to the ways in which the relevant scientific "theories and technologies are intertwined and entangled with each other." We cannot really conceive of an alternative science to ours that does without a notion of 'atom', for example, whose types and key properties are codified in the periodic table of elements.

> We [would] need to rewrite all of chemistry from the ground up, no small feat. But we also would need to rewrite much of astrophysics (what is the process that happens in stars to produce energy?), and statistical mechanics (with its theory of heat as molecular motion). We would need to revise all of microbiology as well (stories about how living cells get nutrients and oxygen, for example), and our understanding of what goes on in electron-tunneling microscopes. (Hoefer, this volume)

Thus, numerous higher-level special sciences are thoroughly intertwined with the physics of atoms and electrons — both in their theories and core technologies — rendering the potential elimination of these theoretical posits inconceivable. By contrast, alternatives to our currently accepted 'fundamental' theories of physics are clearly not similarly inconceivable. In the contrary, they are actively considered by scientists, different quantum theories (e.g. Bohmian mechanics and dynamical collapse theories) being an obvious case in point.

Hoefer has identified an interesting, even if somewhat knotty contrast. It does not, however, serve to bracket off realism about quantum physics, some central notions of which are sure to be found in the good company of atoms and electrons, on the side of Hoefer's realist commitments. Spin presents a prime example of this, given its thorough entrenchment in a vast array of theories and accompanying technologies, as indicated below. Science as empirically successful (explanatory, unified, and technologically powerful) as ours without 'spin' is just as inconceivable, as I will next discuss.

3.2 Spin

Spin, the quantum property that Wolfgang Pauli in 1925 famously described (for the electron) very minimally as a "two-valued quantum degree of freedom," has become commonly known as a particle's intrinsic angular momentum, due to the way in which it contributes to the particle's magnetic moment in analogy and in combination with its orbital angular momentum. After initial attempts (by Goudsmit and Uhlenbeck, and others) to construe electron spin in quite literal mechanical terms as a charged spinning object — an idea quickly jettisoned as incompatible with relativity — the physics community rather quickly learned to accept spin as a real quantum feature that might indeed be irreducible or "intrinsic": a property for which there is no further *dynamical* story to be told.[3] Furthermore, it is generally accepted that there is no classical counterpart to spin: unlike its energy or orbital angular momentum, say, there is no property in classical physics that corresponds to this property at the classical limit (as Planck's constant tends to zero, or whatever more sophisticated analysis of the limit one prefers).[4] In spatial terms, the issue is that if we think of spin as a particle's property, then it in some sense has direction but there is no classical vectorial quantity whose components in all possible directions belong to the same discrete set.[5]

Physicists' rapid acceptance of spin as such an irreducibly quantum property was driven by the indispensability of the new quantum degree of freedom in capturing and accounting for various phenomena, notably the anomalous Zeeman effect — observed deviations from the "normal" Zeeman effect splitting of spectral lines for some elements in a strong magnetic field — which troubled the 'old quantum theory' of Bohr and Sommerfeld. Since these early scientific

[3] In comparison, the spin of a baseball or other macroscopic body can be dynamically reduced to the orbital angular momentum of its constituent particles in a mechanical way. Whether or not *some* kind of dynamical interpretation for electron spin is possible is an open question. See, e.g., Ohanian (1986), Giulini (2008), and Sebens (2018).

[4] See, however, Giulini (2008) and Keppeler (2003).

[5] The quantum state of a spin-$\frac{1}{2}$ particle (a fermion) is represented by a two-valued spinor wavefunction, and the operators corresponding to orthogonal spin components are given by the Pauli spin matrices.

success stories the physics of spin has only gone from strength to strength, with hugely important theoretical implications in numerous subfields of physics, from elementary particle physics to condensed matter physics, optics, and physics of atoms and molecules (see Raimond and Rivasseau 2009 for a selective review). In physics without spin, a huge amount of generally accepted theory would have to be rewritten from the ground up.

Spin also matters a great deal to chemistry and molecular biology, both disciplines of which now share a substantial interface with quantum physics. For example, spin is pivotal to understanding a broad range of phenomena involving electron transfer (from one chemical species to another) and spin exchange (changing of the orientation of oppositely polarized spins), both ubiquitous in important chemical and biological processes involving isolated molecules, ions and excess electrons in solution, electrochemical systems, and so on (see Likhtenshtein 2016 for a review). Modern biological understanding of phenomena such as photosynthesis thus acknowledges the critical role of weak spin-related interactions in the relevant chemical and biological processes. At a more basic level, the periodicity of chemical elements and the notion of molecular bond — the two cornerstones of chemistry — profoundly rely on the Pauli exclusion principle, which "forbids" electrons (as fermions with half-integer spin) from occupying the same quantum state, accounting in large part for the electron shell structure of the elements.[6]

Hoefer's observation that it is impossible to conceive of chemistry or molecular biology without atoms and molecules thus also applies to spin: many achievements of these disciplines, sans spin, are equally hard to conceive of. In terms of its integration to the special sciences, spin is not relevantly different from atoms or molecules.

Spin matters greatly to vast areas of chemistry not only as a theoretical notion, but through the technologies it has led to. The spin degree of freedom of both electrons and nuclei is critical in many areas of spectroscopy, which study atoms' or molecules' interaction with radiative energy. Since its inception spin has become absolutely indispensable to the quantum physical representation of possible particle states, their energy levels, and the 'selection rules' that govern possible state transitions. These involve, for example, details of the spin-orbit coupling — typically understood as the interaction between the magnetic moments due to spin and orbital angular momentum — as well as possible transitions between nuclear and electron spin states. The two-valued character of electron spin and its coupling with the orbital angular momentum famously explains the splitting of the main lines of atoms' emission spectra into two or more components — its fine structure — and the spin of the nucleus and its

[6]At the theoretical level the exclusion principle can be connected to spin via the spin-statistics theorem, according to which integral spin particles obey Bose-Einstein statistics, while half-integral spin particles obey Fermi-Dirac statistics and the exclusion Principle. First derived by Pauli, this theorem has subsequently been proven from numerous different starting points, typically involving non-elementary presuppositions in the context of quantum field theory. See Duck and Sudarshan (1997) for a classic review.

interaction with the magnetic field produced by the orbiting electron explains the further hyperfine structure.

Many modern spectroscopic techniques have important applications outside physics laboratories. Nuclear magnetic resonance (NMR) spectroscopy presents an important application that is vital, e.g., to modern organic chemistry in offering a preeminent technique for identifying organic compounds and determining their structure and functional groups. Molecules' electronic structure and individual functional groups (viz. groups of atoms or bonds responsible for the molecules' behaviour in chemical reactions) are determined by probing the intramolecular magnetic field, which is unique (or at least highly characteristic) to individual compounds. What theoretically underlies such probing is a quantum mechanical understanding of the magnetic moment of nuclear spin, and the way it processes in external magnetic field. The rate of precession of the bulk-magnetisation due to the spins of particular kind of nuclei (e.g. ^{1}H or ^{13}C) in a given magnetic field is directly proportional to the nuclei's gyromagnetic ratio, which is a fundamental nuclear constant. A given type of nuclei in a sample can absorb energy from an electromagnetic field of appropriate (radio-length) frequency, which is dependent on the type of nuclei in question, as well as their chemical environment.[7]

Given the intimate connection between spin and particles' quantised magnetic moment, it is unsurprising that the notion of spin is central to understanding the magnetic properties of matter. Electrons' intrinsic angular momentum, in particular, is the key to ferromagnetism (of, e.g., iron fridge magnets), which is a purely kinematic consequence of spin and the Pauli exclusion principle (see e.g. Blundell, 2003). In general, magnetism at the macroscopic level is an example of collective phenomena studied in condensed matter physics, which provides understanding also of numerous more esoteric phenomena, such as spin waves and spin glasses, exhibited by ordered magnetic materials (ibid.). Theoretical advances regarding such phenomena bear promise of technological applications in quantum computing, for example, following the history of hugely important implications of spin and magnetism to electronics.[8] The steady growth of these technological achievements seems quite inconceivable without spin.

[7]This dependence can be determined on the basis of quantum mechanical understanding of nuclear spin and intra-molecular spin-spin interactions, and it offers the key to interpreting in chemical terms the spectra of magnetic 'resonance' (viz. the Larmor frequency at which nuclear spins absorb energy). The key theoretical notions involved in all this — e.g. spin precession and nuclear magnetic resonance — ground many important technological applications, and the so-called Bloch equations provide a quantum theoretic description that unifies the physics of NMR with that of electron spin resonance spectroscopy, as well as Magnetic Resonance Imaging. From its Nobel-price winning inception (Felix Bloch and Edward Purcell in 1952) the NMR theory of nuclear spins has grown to yield an incredibly powerful technology for investigating all kinds of matter, ranging from brains, bones, cells, ceramics, liquid crystals, laser-polarized gases, proteins, superconductors, zeolites, blood, polymers, colloids, and so on (Keeler 2005).

[8]'Giant' magnetoresistance (GMR) is one momentous example, underlying modern computer hard-drives, the magnetic field sensors of which are one device amongst many based on this phenomenon, accounted for in terms of electrons' spin-dependent scattering from magnetised layers (see. e.g. Blügel et al. 2009).

The increasing theoretical grasp on GMR and other spin-dependent magnetoresistance phenomena, such as tunnelling magnetoresistance, constitute the basis for *spintronics*. The development of increasingly varied spintronics devices evidences how the current theoretical handle on spin enables the control and manipulation of electric currents on the basis of electrons' spin (in addition to their electric charge). These devices utilise a "spin valve" for controlling electrical resistance on the basis of GMR by manipulating a spin-polarised electric current by changing the direction of magnetic fields, as well as other spin-related phenomena (e.g. spin Hall effect and spin-torque transfer; see *ibid.*).[9] Such ability to employ particles' spin in electronic devices renders electric currents *as quantum theoretic objects* effectively manipulable in the sense of Ian Hacking's (1983) entity realism, sloganized: "if you can spray them, they are real". Hacking's realist criterion is one way to understand the distinction that Hoefer draws between (i) electrons as quantum mechanical posits, and (ii) electrons of electronics, the manipulability of the latter being based on a broadly causal grasp of electrons' charge, which is largely independent of quantum mechanics. With the advancement of spintronics we can now see spin also fitting Hacking's realist mold, despite it being a thoroughly quantum mechanical feature of the world.

At a more theoretical level, spin is central to high-energy elementary particle physics, and efforts to understand the nature of spin itself involve deep connections between quantum theory and special relativity. The wave equation for fermions in relativistic quantum field theory — the Dirac equation — automatically incorporates spin as a kinematic property of a spinor-valued field which transforms as an irreducible representation of the Lorentz group. A group-theoretic analysis of relativistic quantum theory yields the highest level of theoretical analysis of the origin of spin in terms of symmetries that unify it with other kinematic properties of quantum states.[10]

I have surveyed aspects of spin-related physics in order to evidence the claim made above: spin presents a prime example of a notion *so thoroughly entrenched different theories and technologies that we cannot really conceive of science as successful as ours without it.* The amount of research that has consolidated the "science of spin" is colossal, and we have only selectively scratched its surface.[11] In relation to key realist criteria, spin surely ticks all the boxes, by virtue of being deeply explanatory, unifying, and even effectively manipulable. Hence, we should be realists about spin, as much as we are realists about any theoretical notion.

[9] This is even more striking in relation to recent successes to control and manipulate coherent single spin states of "quantum dots," nanometer scale semiconductor quasi-particles, which can be suitably localised and isolated from environmental disturbances to ensure sufficient coherence timescales (Hanson and Awschalom 2008).

[10] See, e.g. Zee (2016, 256), whose "mathematical pragmatist" responds to the lack of "physical understanding" of spin (cf. Morrison 2007) by saying that the electron is simply "a particle whose quantum state transforms like the 2–dimensional representation of the covering group of rotations carries spin 1, period." Ontological structural realism capitalises on this kind of understanding of spin (French, 2014).

[11] At the time of writing this the physics preprint archive arXiv.org contains nearly 38,000 articles with "spin" in their title.

But what does this realism amount to?

3.3 Two conceptions of theory realism

We get into a better position to answer the question above by first reflecting in general terms what scientific realism about theories like those of quantum physics amounts to. To this end I will first distinguish between two conceptions of realism, truth-content realism and progress realism, before problematising the former in relation to spin in the next section.

Taken at face value, realism about spin concerns the world: what exists, or is real. (Namely, spin.) This assertion makes recourse to quantum physics, of course: spin is a theoretical notion, and its cognitive content and epistemic access to it is through the theories in which spin figures. Realism about spin thus confirms or declares a positive epistemic attitude towards what our best theories say about spin-related unobservable features of the world. Presumably our current theories should not be regarded as providing a literally true description of spin, however, as we must leave room for future theoretical advances and revisions. Since today's physics should not be taken as the final word, the realist's epistemic confidence is better captured in terms of our current theories' "approximate" truth, or by reference to some other kind of representational fidelity that falls short truth simpliciter. Realism about spin (in this way of thinking about it) is hence an assertion about the world that is made in terms of thus circumscribed confidence regarding our current best theories.

Notoriously, the notion of "approximate truth" is as critical as it is difficult to spell out (particularly at the level of generality at which realist claims are typically expressed). For many realists a commitment to the referential status of central theoretical terms (e.g., 'electron', 'atom', or 'spin') has been a key to maintaining a sufficiently robust epistemic commitment towards the respective worldly features in the face of uncertainty regarding future theoretical developments, or (in other words) regarding the exact sense in which the current theoretical descriptions of those features are "approximately true". The future theories, whatever they exactly say about spin, for example, will at least be concerned with those very same entities (e.g. 'electron') and their properties (e.g. 'spin') to which our theories now refer. On this basis the realist can purportedly maintain that we have scientific knowledge about what entities exist and what these entities are like at the unobservable level: e.g., electrons exist and they have (inter alia) spin, which accounts for various observable phenomena.

So far, so familiar, at least to those acquainted with the classic writings of, e.g., Hilary Putnam (1982), Richard Boyd (1984), Stathis Psillos (1999), and others. TRUTH-CONTENT REALISM is an appropriate label for variations of this traditional realist theme, the development of which is motivated by an ancient quest for knowledge about the reality behind the veil of directly observable phenomena, whose explanatory understanding science enables by reference to their unobservable causes. There is, of course, much more to be said about the two key notions at play here — reference and approximate truth — of which realists have developed various detailed accounts. Some have offered general formal accounts of approximate truth or "verimilitude" (e.g. Niiniluoto 1999), while others have developed these notions in more informal and case-dependent terms (e.g. Psillos 1999). The relationship between approximate truth and reference can also be understood in different ways: some saddle truth-content realism with a substantial causal reading of reference, for instance, while others associate it with much more minimal commitments (see e.g. McLeish 2006b).

Yet others regard reference altogether unnecessary for characterising realist commitments. In particular, one natural reading of structural realism, according to which theories only provide us knowledge of structural features of reality, is that our current theories' truth content is limited to their correctly representing the world's "structure" (suitably construed).[12] According to many structural realists, such now-rejected theories as Newtonian gravity or Fresnel's ether theory of light can contain appropriate truth content regardless of the referential status of "gravitational force" or "ether" (Frigg and Votsis 2011).[13] Structural realism — at least in this epistemic reading — thus continues the quest of defending scientific knowledge about the unobservable reality; its novelty is in incorporating a principled limitation as to what can be known: only "the structure of reality" is knowable. Structural realism (thus construed) is hence another variant of truth-content realism.

So much for truth-content realism. There is an alternative, much less ambitious way to conceive of the realist project. One can forgo many of the central tenets of truth-content realism concerning approximate truth, reference, and even scientific knowledge of the unobservable, while upholding its *most critical tenet*: the idea that the spectacular empirical success of our best scientific theories is due to their faithfulness as representations of reality. According to this realist tenet theories of mature science, such as quantum theories, latch onto unobservable reality in ways that are responsible for their empirical successes — both predictive and explanatory — as recognised by scientists. PROGRESS REALISM is perhaps a good label for attempts to defend this realist tenet.

Progress realism has emerged in the debate revolving around the historical evidence that historicist critics of realism have presented against the realist credo that the empirical success of science would be a "miracle" if the relevant theories weren't approximately true. Various realists have responded to this criticism by showing, case by case, how it is possible to maintain that the empirical successes of past scientific theories (concerning, e.g., luminiferous ether,

[12] This way of circumscribing theories' truth content may seem like a radical departure from standard realist's idea that theories are "approximately true". But it actually only represents a natural continuation of the selective ("divide et impera") strategy already initiated by the traditional realists — especially Philip Kitcher (1993) and Stathis Psillos (1999) — for whom reference of central theoretical terms is a key component of realism.

[13] Some accounts of structural realism can be connected to the reference of theoretical terms, however. See Schurz (2009).

phlogiston, gravitational force) can be accounted for in terms of their representational relationship to unobservable reality, as we see it *from our current vantage point*. The issue at stake in these historical analyses concerns first and foremost our current theories' representational relationship to the world: do we have reason to believe that these theories' empirical success is grounded in their representational faithfulness? According to progress realism we do, in the same sense as we can from our current vantage point regard the past theories' empirical successes being due to an appropriate representational success.

We can note, as a purely conceptual point first of all, that in defending this tenet a progress realist is not making an assertion about the world, or about what we can claim to know. Indeed, one may wish to defend this tenet (as I do) without defending claims about approximate truth or knowledge of the sort that truth-content realists hanker after. Indeed, it is possible that one is only be able to argue that appropriate representational relationship *holds*, without being able to tell exactly what that relationship *is like* (Saatsi 2016). Progress realism maintains that the empirical successes of theoretical science are by and large due to theories latching onto reality in ways that ground those successes, and that there is genuine theoretical progress in science in how well theories' represent reality (ibid.).

These two conceptions of realism are hence conceptually distinct, but they are rarely distinguished. This is unsurprising, since truth-content realism implies theoretical progress of the sort defended by progress realism.[14] Hence, from the point of view of truth-content realism the two conceptions are just different sides of the same realist coin. The core commitment of progress realism in and of itself is much weaker, however. In particular, it is possible that scientific theories' empirical success is due to their representational relationships to reality, and that science progresses with respect to how its theories represent reality, without us being in a position to reliably pin down or precisify our current theories' truth-content in the absence of direct epistemic access to reality.

Whether it makes sense to capitalise on this conceptual distinction and limit one's epistemic ambitions to progress realism depends on one's perspective on the realism debate at large. In particular, one can be a lumper in relation to realist analyses of foregone scientific theories, hoping to extract from these analyses a unified story of how the empirical success of theoretical science is grounded in truths about, e.g., unobservable entities' causal powers (Chakravartty 2007), or the world's structure in some broader sense (Ladyman and Ross 2007, French 2014). Or one can be a splitter, emphasising instead the disunity and open-ended variety of different kinds of (possible) realist explanations of empirical success. If one is a splitter (as I am), it makes perfect sense to articulate and defend progress realism while forgoing the ambitions of traditional truth-content realism (Saatsi 2015). If there is no projectable unified account to be given of

[14] As realists have tended to put it: as science progresses its theories approximate the truth better and better (or, the degree of theories' verisimilitude increases), and their success is due to their approximate truth (or verisimilitude). Structural realists would capture this progress differently, e.g. in terms of theories' limiting relationships.

the representational relationships responsible for theories' empirical success, we cannot possibly hope to extract from a handful of diverse case-studies any kind of "recipe" for identifying realist commitments towards our current theories with respect to their truth content. If there is no such generalisable sense of "approximate truth" to be had, we better give up on truth-content realism.

Returning to the topic of spin, should we be satisfied with such a progress realist stance towards quantum theories? Well, in the rest of this chapter I will argue that it is the best we can have. Our overall evidence towards a quantum physical grasp of spin — incorporating both the empirical evidence for quantum theories, as well as the available evidence about the relationship between our best theories and the world more broadly — supports progress realism, but goes no further. Hence, a naturalistic philosopher who takes on board all admissable evidence should be a progress realist, but not a truth-content realist.[15] In the next section I will discuss the challenge faced by truth-content realism, before elaborating on progress realist's commitments regarding quantum physics in the final section.

3.4 Truth-content realism challenged

Truth-content realism faces a subtle problem of underdetermination. The gist of the problem, in relation to spin in particular, is that the colossal empirical evidence amassed for quantum theories does not suffice to determine what we can claim to know about, e.g., the functioning of a Stern-Gerlach apparatus, the causes of spin-orbit interaction effects in atoms, or the workings of NMR and GMR based devices. There is, of course, an extremely well-established body of theory of these and other spin-related phenomena that we can undeniably trust in many ways; the problem is not that quantum theories are not delivering what they generally claim to deliver. Truth-content realism demands, however, that this body of theory yields *knowledge* about aspects of the unobservable world, regarding spin, spin-orbit interaction, spin-valves, and so on. In order to meet this demand a realist must provide two things: (i) a true interpretation of the relevant parts of the mathematized theory, and (ii) an empirically well-grounded justification for this interpretation. A problem of underdetermination impedes the realist from achieving (ii), and hence we cannot know whether we have achieved (i). As a result the required kind of scientific knowledge about the world lies beyond our epistemic reach.

That is the problem in outline. Let me now elaborate on it and then highlight some often ignored subtleties due to the broader dialectic of the contemporary realism debate. It should be incontestable that some interpretation of the mathematical representation of spin (in terms of, e.g., spinor-valued wave-function and Pauli spin matrices) in quantum theory is required for realism. Such interpretation should, furthermore, make reference to what there is in the mind-independent world — what the world is really like. Physicists often regard

[15] Spelling out how to think about "overall evidence" here has to be left to another occasion.

this as an obsolete quest for the "meaning" of quantum theory, insisting that knowledge about, e.g., spin requires no such thing. This attitude risks deflating the content of the alleged knowledge about the unobservable world, vacillating unstably between: (i) committing to quantum states of affairs and quantum properties of entities like electrons, on the one hand, and (ii) to just using the theory as a reliable mathematical apparatus for calculating non-quantum features of the world and for "understanding" these features in some sense that is entirely ambiguous in terms of its realist commitments, on the other. In order to secure a stable realist commitment with clear cognitive content, truth-content realism about spin is unavoidably imbued with metaphysics, in the sense of being committed to quantum theory delivering us identifiable, objective truths about unobservable features of the world.

Mind you, a professed realist interpretation of the theory need not be given in macro-physical terms that are somehow readily imaginable, familiar, or visualisable to us. As Ernan McMullin (1984, 14) puts it, "imaginability must not be made the test for ontology" when it comes to microphysics: "The realist claim is that the scientist is discovering the structures of the world; it is not required in addition that these structures be imaginable in the categories of the macroworld." (*ibid.*) What is nevertheless required, however, is that a realist interpretation of the theory yields truths about the relevant "structures of the world." This follows from the fact that knowledge — scientific or otherwise — is *factive*: if something is known, it is true. (Hence, a structural realist interpretation of Pauli spin matrices and such needs to comply with this requirement as well, in as far as it defends knowledge about spin.)

The issue is not that there are no *candidate theories* that could contain the requisite truths. (I use the term "candidate theory" to refer to what is often called a "realist interpretation": a coherent formulation together with an interpretation, such as Bohmian quantum mechanics, or quantum theory as the Everettians interpret it.) The issue is that we have good reasons to think that no candidate theory as such is worthy of the realist's epistemic commitment, due to the way in which such candidates involve metaphysical assumptions that go beyond what realists should deem responsible for quantum theories' explanatory and predictive successes. I have discussed this in detail elsewhere (Saatsi 2019), introducing the notion of "deep metaphysics" to capture those theoretical assumptions that transcend what the realist should regard as accounting for the empirical success of quantum physics, given the actual scientific practice of using quantum theory to predict, manipulate, and explain things. The involvement of such "deep" metaphysics in all current candidate theories is manifested in the availability of several such candidates, which are all empirically adequate with respect to various quantum systems, whilst making radically different claims regarding spin, for example. In particular, Everettian quantum mechanics, pilot-wave theories, and dynamical collapse theories represent well-known alternative candidate theories of non-relativistic quantum mechanics, each offering a radically different account of the physical systems falling within their remit.

For illustration, consider an archetypal application of quantum mechanics: the Stern-Gerlach experiment. Why does an inhomogeneous magnetic field, as in Stern and Gerlach's experiment in 1922, appear to deflect neutral silver atoms (with zero orbital angular momentum), some "up" and some "down"? Classic textbooks account for this roughly as follows (e.g. Townsend 1992, Sakurai 1995). The silver atom have two-valued intrinsic angular momentum that is (almost) entirely due to the spin-$\frac{1}{2}$ of the "lone" 47th electron in these atoms' "outermost" electron shell. The quantum state of the atom evolves upon the magnetic moment's interaction with the magnetic field in a way that can be analysed as a (classical) force on the atom, deflecting its trajectory (or, rather, the support for the corresponding wavefunction) "upwards" and "downwards" (or into a superposition thereof), depending on the initial spin state.[16] Finally, measuring an atom's position at the end of the experiment "collapses" the superposition exhibited by the atom's quantum state.

The more foundational discussions of quantum theory that aim to paint a realist picture of this kind of experiment invariably give up on the notion of collapse-upon-measurement as codified in the textbooks' notorious "collapse" or "projection" postulate. It is agreed that such a postulate cannot as such be part of a serious realist candidate theory of quantum mechanics, as it leads to the measurement problem and theoretical incoherence. Instead, a candidate theory suitable for a realist interpretation must somehow either do without such collapses, or it must change the fundamental quantum dynamics itself so as to make it empirically consistent with the (apparent) collapses of superpositions in measurement-like situations with Born-rule statistics. The latter option leads to dynamical collapse theories (such as that of Ghirardi, Rimini, and Weber 1986, or Pearle 1989), while the former option leads to either a hidden-variable formulation (de Broglie-Bohm), or to the Everettian many-worlds theory (which aims to interpret and make sense of standard quantum theory sans the problematic collapse postulate).

These different candidate theories incorporate radically different understanding of spin. Take dynamical collapse theories first. It is far from straightforward to spell out what silver atoms and their spin physically amount to in such theories, and there is a good deal of debate about their ontology (see, e.g., Lewis 2018, Myrvold 2018, Tumulka 2018). A dynamical collapse theory can be read as one that fundamentally just describes a wavefunction living in a very high-dimensional configuration space, so structured as to give rise to effectively 3-dimensional reality, in which particle-like phenomena emerge from the fundamental wavefunction under suitable circumstances, with spin being just a feature of the (spinor-valued) wavefunction. An analysis of the Stern-Gerlach experiment looks at the magnetic field's effect on the wavefunction, which evolves according to a fundamentally stochastic dynamical law (instead of the deterministic Schrödinger equation), such that it is practically guaranteed to give rise to a determinate, randomly localised particle-like result upon its interaction with a macroscopic location-measurement device. One can try to

[16]Typical "textbook analyses" have been criticised in Hannout et al. (1998).

interpret parts of the wavefunction before its stochastic collapse as effectively realizing the (non-fundamental) ontology of a spin-$\frac{1}{2}$ particle occupying a superposition state. Or one can introduce local beables — e.g. spatiotemporal matter density or 'flashes' — as collapse theories' further, "primitive" ontology.[17] In the primitive ontology approaches spin is not a property of the theory's beables at all, but rather a nomological aspect of the world that just codifies facts about the beables' spatiotemporal positions and evolution.

Spelling out the workings of a Stern-Gerlach machine along these lines leads to specific realist accounts of what it *means* to attribute spin-$\frac{1}{2}$ to an electron or a silver atom, so as to explain the Stern-Gerlach experiment. These accounts diverge radically from a face-value textbook reading of quantum mechanics, according to which electrons have an *intrinsic (non-classical) property* of spin-$\frac{1}{2}$, which silver atoms also have due to the way in which electron spins quantum physically combine to yield the atom's total spin, which affords the atom the property of intrinsic magnetic moment that interacts with the magnetic field to yield the observed outcome (after a measurement "collapse"). Dynamical collapse theories can similarly diverge from an Everettian account of the Stern-Gerlach set-up. The latter follows the face-value reading of the "standard" quantum theory all the way up to the employment of the collapse postulate, which the Everettians jettison in favour of a deeper quantum physical account of the measurement process in terms of quantum decoherence theory (leading to the emergent branches of the multiverse; see Wallace 2012). In particular, for the Everettians spin is not relegated to a nomological feature of the world any more than spatial location is, in contrast to the primitive ontology approaches to dynamical collapse theories.

De Broglie-Bohm theory is similarly revisionary regarding spin's status as particles' property, and the Bohmians naturally regard spin as a feature of the wavefunction in way that makes attributions of spin to particles *contextual*, and dependent on the experimental arrangement (see e.g., Norsen, 2014; Daumer et al., 1996). Thus, there is no sense in which a "measurement" of spin reveals some pre-existing, intrinsic property of a particle. In this respect the de Broglie-Bohm theory is similar to the primitive ontology approaches to dynamical collapse theories. The two differ radically, however, in their analysis of quantum physical probabilities. In dynamical collapse theories the outcomes of the Stern-Gerlach experiment obey the Born rule due to the fundamental dynamics of quantum reality being stochastic in a suitable way. In the Bohmian theory, by contrast, this is due to statistical aspects of deterministic dynamics and suitable initial conditions (Norsen 2018).

This glimpse into some of the foundational analyses of quantum theory suffices to highlight the fact that different candidate theories can tell a very different story of what's going on with a Stern-Gerlach apparatus. The challenge to truth-content realism is that it seems forced to buy into "deeply" metaphysical assumptions — assumptions that are epistemologically unwarranted by the

[17]See, e.g., Lewis 2006 for an opinionated review of these options.

realist lights — in trying to spell out what we claim to know about, e.g., silver atoms in a Stern-Gerlach machine. On the one hand, without a specific candidate theory in mind the realist finds it difficult to specify substantial truths that ground her knowledge claims and support a realist explanation of the empirical success of quantum theory. On the other hand, it is hard to see how the realist can adopt any particular candidate theory on empirical grounds, given how each of them seems to be wedded to ("deep") metaphysical assumptions exactly of the sort that realists are at pains to avoid, having taken to heart the historical lessons regarding such assumptions' untrustworthiness as part of empirical science (see Saatsi 2019).

A natural reaction to this challenge is to look for a common theoretical ground at a more abstract, structural level. What substantial shared structural commitments can be found between the different candidate theories, however? We should not be content to just talk about abstract mathematical features of the quantum state attribution, since truth-content realism is concerned with the ontology *represented* by the mathematics. But without an ontological specification of what the relevant quantum states represent, it is hard to see how truth-content realism can provide a serious account of quantum theory's empirical success or the claimed knowledge of the unobservable reality. Realism about the structure of quantum reality in the spirit of truth-content realism requires that 'structure' can be explicated in a way that is compatible with all the candidate theories. It seems unlikely that this can be done by reference to these theories' dynamics or state-spaces, for example, as Ruetsche (2018, 300) notes:

> [I]t is not at all clear [these candidate theories] have in common any structure *of interest for realism.* Contender interpretations attribute QM different *types* of state spaces (for Everett, it's a Hilbert space; for Bohm, a space of particle configurations) and different *types* of dynamics (deterministic Schrödinger evolution, stochastic collapse, deterministic guidance equation).

A more metaphysical, modal characterisation of "structure" can be developed (French, 2014), but this risks opting for just another line of "deep" metaphysics, beyond the bounds of the realist's epistemological humility (cf. Saatsi, 2019).

The problem of underdetermination thus is that none of the candidate theories seems worthy of the realist's epistemic commitment, given their involvement of metaphysical assumptions that go beyond what realists should deem responsible for quantum theories' explanatory and predictive successes. Calling these assumptions "deeply" metaphysical is not derogatory, but just highlights the fact that practicing physicists who successfully deal with, e.g., spin and magnetic fields by and large do not (and seemingly need not) care about these candidates for making predictions, building instruments, or even explaining various phenomena. The theoretical details that seem unavoidable for spelling out the commitments of truth-content realism only play a role at the foundational and interpretational level of theorising, which so far has not led to any empirical successes of the sort that realists (by their own, demanding lights) should regard as eliciting a realist commitment. In as far as realism is driven by a desire to account for the established empirical and explanatory successes of science, the realist should focus her commitments on those theoretical aspects of quantum physics that can be regarded as responsible for those successes.[18] And these aspects are strikingly independent from the foundational-cum-interpretational research on quantum theory, as Healey (2017; this volume), for example, has forcefully emphasised.

There are a couple of subtleties about this underdetermination challenge that are worth emphasising. First, note that I have framed it in terms of the metaphysical nature of the existing candidate theories of quantum physics (the problem being that their metaphysical "depth" transcends the scientific realists' commitments). The alternative theories of non-relativistic quantum mechanics *manifest* the underdetermination, but the challenge to realism does not depend on the (historically contingent) fact that such alternatives have actually been developed. Even if we only had the de Broglie-Bohm theory on the table, say — not having conceived of the dynamical collapse or Everettian alternatives — a realist should want to be able to recognise upon rational reflection this theory's metaphysical "depth" in relation to the theoretical commitments that are actually required for achieving the empirical successes of quantum mechanics.[19] The challenge rather depends on there being evidence of such metaphysical "depth" in the current candidate theories, such that no substantial truth-realist commitments towards spin remain after bracketing the "deeply" metaphysical assumptions. I have argued that we have such evidence in the relative independence of the empirical and explanatory successes of actual quantum physics from the theoretical assumptions that specify each candidate theory's ontological content.

Secondly, note that I have not claimed that the current candidate theories are evidentially on a par with respect to their: (i) prima facie metaphysical plausibility; (ii) metaphysical plausibility in relation to non-quantum physics; (iii) potential involvement of ad hoc assumptions; or (iv) prospects for providing a unified interpretation of all empirically successful quantum physics.[20] Such a

[18] One should also appreciate the considerable degree of humility of contemporary realist commitments in the broader dialectic of the realism debate (cf. Saatsi, 2019).

[19] Assume, for the sake of the argument, that in the scenario envisaged here most physicists are Bohmians about quantum mechanics, and that relativistic and field–theoretic extensions of Bohmian mechanics are the hottest area of research around. Such broad allegiance to Bohmianism should not convince the realist, since the commitments of truth–content realism are not read off from scientists' beliefs. This is comparable to how an enlightened realist should want to be able to recognise the undue metaphysical depth of the ether-laden construals of Maxwellian electrodynamics in the 19th and early 20th centuries, given the role played by 'ether' in the actual physical theorising by, e.g., the Cambridge Maxwellians around the time (cf. Gooday and Mitchell 2013).

[20] Expert opinions widely differ regarding (i) and (ii). Regarding (iii), the dynamical collapse theory risks being somewhat ad hoc. And regarding (iv), arguably Everettian quantum physics gets the upper hand here, for reasons laid out by Wallace (this volume). Regarding spin, in

further claim would be a red herring regarding the challenge at stake. The challenge is *not* that we have developed alternative theories each of which the realist would happily regard as delivering truths about unobservable aspects of reality (according to their preferred realist "recipe"), were it not for the availability of an evidentially comparable competing theory for which the realist's "recipe" delivers different commitments. The challenge rather concerns the fact that each candidate theory in and of itself indispensably involves "deep" metaphysical commitments that are epistemically unjustifiable by the realist's lights.

3.5 Progress realism about spin

We can avoid getting sucked into "deep" metaphysics by defending only progress realism about spin, while giving up the epistemic ambitions of truth-content realism. Progress realism about spin defends the claim that physics' astonishing empirical success with respect to various spin-related phenomena is due to its theories' and models' representational relationships to reality, and that this area of science has progressed and continues to progress with respect to how well its theories represent the unobservable reality. Progress realism maintains that we are warranted in believing this claim on the basis of the empirical evidence enjoyed by the relevant theories and models, and what we know about physics and its history at large.

Progress realism does not defend knowledge about what spin is like or what we can claim to know about the properties of electrons or some kind of worldly wavefunction (as in the currently fashionable "wavefunction realism"; see Chen 2019 for review). What it defends is the idea that the models and theories of quantum physics stand in a robust enough representational relationship to reality to ground its empirical successes: new predictions, increasing explanatory understanding, and manipulations of quantum systems. It is an attitude towards the success-yielding theoretical practices of quantum physics in its own terms: since these do not require foundational-cum-interpretational assumptions about what the quantum wavefunction or Pauli spin matrices represent, such "deeply" metaphysical assumptions should not need to be part of the scientific realist account of the empirical successes of quantum physics either. Such assumptions belong to the *metaphysical* foundations of quantum theory, which is an extremely well-motivated and important endeavour, but one that the realist need not engage with in articulating her *epistemic* commitments.

Thus expressed, such commitments may immediately appear much too insubstantial and unsatisfying. And wholly unsatisfying they are, of course, *if* one is simply unwilling to give up on the idea that realism is a matter of defending knowledge of a certain sort and identifying theoretical truth-content

particular, it is admittedly *not* the case that for any (or even for most) spin–related phenomena there are different candidate theories on the table, since many spin-related phenomena require relativistic treatment for which no adequate extension of e.g. Bohmian mechanics has been developed.

underwriting that knowledge. As to the charge of lacking substance, what matters is how progress realism can be clearly differentiated from anti-realist or non-realist stances towards quantum physics. Progress realism aims at offering, along with physics itself, an account of the empirical successes of quantum physics, in terms of how our theories are latching better and better onto reality. Paradigmatic anti-realist positions see no role for such an account. For instance, van Fraassen's (1980) constructive empiricism maintains that a simple Darwinian story suffices, of science "red in tooth and claw" with theoretical representations' "fitness" measured entirely in terms of their degree of empirical adequacy. Likewise with instrumentalists of various stripes, such as Kyle Stanford's neo-instrumentalism, which acknowledges that there is probably some reason why "foundational" theories (e.g. in quantum physics) are empirically successful, but maintains that nothing of substance can be said of these reasons.

Such anti-realist attitudes towards the realist's preoccupation with the success of science are in striking discord with physicists' own substantial accounts of their theories' and models' empirical success and ongoing attempts to further understand them. Consider spintronics, again, whose devices usually involve magnetic materials exhibiting various spin-related phenomena involving very large numbers of constituents. Magnetism and many relevant collective phenomena can be represented in classical or semi-classical terms (e.g. the direction of magnetism or spin-current is a classical vectorial quantity), which in models of spin-valves and such appear to represent interesting new physical properties (e.g. spin-currents and spin-waves), which have led to various novel predictions and experimental manipulations. (Indeed, some physicists are calling the ongoing "quantum engineering" phase the second quantum revolution!) *How can these models be so incredibly successful?* The basic realist idea is that their success is down to their representational faithfulness. For a naturalistic philosopher this idea can be motivated just by pointing out that the successful practice of spintronics is premised on this very idea. Regardless of what exactly it is that quantum physical models capture at the "fundamental" level of individual electrons — a foundational-cum-interpretational issue on which spintronics research very rarely takes a stand on — this area of research at large is premised on the idea that there is a detailed physical theory to be given of the relationship between particles' quantum physical spin (which has no classical analog) and the collective "classical" spin phenomena at the macro-level. Physicists' answer to the above question (in italics) can be thus summed up: because the world at the "fundamental" level is quantum, involving quantum states and features that we effectively capture with e.g. the Pauli spin matrices, which are sufficiently well represented by quantum mechanics to enable the theoretical notions of spintronics, which capture collective quantum phenomena, to represent the relevant emergent features of semiconductors.

Anti-realists may be sceptical towards the notion that physics' theorising itself could be genuinely explanatory of spintronics' success. Surely the progress realist is begging the question here? It is not clear what naturalistically respectable reasons there are for such scepticism, however, in the face of the established status of spintronics research, the central notions of which go back more than 70 years. By contrast, progress realism respects physicists' authority in thinking that an account of the empirical success of spintronics can be worked out by studying the relationship of quantum physical representations of single electrons and of macroscopically large collection thereof. Although this account is still in many ways in progress, substantial scientific understanding has already been established regardless of the foundational "black boxes" that have to be drawn at the more fundamental level. As a part of this commitment a progress realist (along with the physicists) regards quantum notions like spin no less representational than any other notion in physics (while being quiet about what these notions actually represent).[21] This stands in contrast with the recent non-representationalist interpretations of quantum physics — such as quantum pragmatism of Richard Healey (2017; this volume) — which aim to account for the empirical successes of quantum physics in terms that leave no representational role for any distinctly quantum notions.

In addition to there being a clear contrast between progress realism and established anti-realist views, progress realism also places substantial constraints on the kinds of developments in science that we can rationally expect, while being open to the possibility that future science will develop in radical, presently unforeseeable directions at the level of its (currently) foundational and "deeply" metaphysical notions. In particular, it expects quantum physics to evolve in ways that make increasing sense of how current physics is latching onto the world in ways that make it successful, so as to complement and not profoundly revise our current understanding of the past theories' empirical success.[22] As to how physics can make increasingly good sense of its own empirical success, progress realism points to much studied exemplars of our best foundational studies of domain-specific inter-theoretic relations, between, e.g., classical and quantum physics, ray optics and wave optics, Newtonian and relativistic theories of mechanics and gravity, old quantum mechanics and contemporary quantum theory, and so on. These well-established exemplars of such realist understanding of how science theoretically progresses are provided from our current vantage point — where else?! — but they are not expected to be overthrown in the fullness of time. This places considerable constraints on our rational expectations regarding future scientific developments, yielding one kind of knowledge of how science progressively "latches onto" reality. This, however, is quite different from truth-content realism.

[21] This is not in any way specific to quantum physics. The progress realist does not want to commit to representational truths about, say, the electromagnetic field and its polarisation either, but just to the idea that that the relevant theoretical representations are latching onto reality in ways that are responsible for their theoretical success.

[22] This is closely related to structural realists' emphasis on inter–theoretic correspondence relations as explanatory of past theories successes. But progress realism recognises nothing distinctly "structural" in the plurality of such correspondence relations, and it denies that there is any kind of structuralist "recipe" for extracting trustworthy truth-content from our current theories.

Acknowledgements

A version of this paper was presented at 13th International Ontology Congress in San Sebastian, and I would like to thank the audience. I also benefitted from discussions and feedback from various people, including Alberto Cordero, Douglas Earl, Matthias Egg, Steven French, Tetsuji Iseda, Simon Newey, Kohei Morita, and David Wallace. The research was funded Arts and Humanities Research Council.

References

Blundell, S. J. (2003). *Magnetism in condensed matter.* Oxford: Oxford University Press.

Blugel, S., D. Burgler, M. Morgenstern, C. M. Schneider, and R. Waser (2009). *Spintronics – from GMR to quantum information.* Forschungszentrum, Zentralbibliothek.

Boyd, R. (1984). The current status of the realism debate. In J. Leplin (Ed.), *Scientific Realism*, pp. 41–82. Los Angeles: University of California Press.

Daumer, M., D. Dürr, S. Goldstein, and N. Zanghi (1996). Naive realism about operators. *Erkenntnis 45*(2-3), 379–397.

Duck, I. and E. Sudarshan (1997). *Pauli and the spin-statistics theorem.* Singapore: World Scientific.

French, S. (2014). *The structure of the world: Metaphysics and representation.* Oxford: Oxford University Press.

Frigg, R. and I. Votsis (2011). Everything you always wanted to know about structural realism but were afraid to ask. *European Journal of Philosophy of Science 1*(2), 227–276.

Ghirardi, G. C., A. Rimini, and T. Weber (1986). Unified dynamics for microscopic and macroscopic systems. *Physical Review D 34*(2), 470.

Giulini, D. (2008). Electron spin or 'classically non-describable two-valuedness'. *Studies in History and Philosophy of Modern Physics 39*(3), 557–578.

Gooday, G. and D. Mitchell (2013). Rethinking classical physics. In J. Buchwald and R. Fox (Eds.), *Oxford Handbook of the History of Physics.* Oxford: Oxford University Press.

Hacking, I. (1983). *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science.* Cambridge: Cambridge University Press.

Hannout, M., S. Hoyt, A. Kryowonos, and A. Widom (1998). Quantum measurement theory and the Stern-Gerlach experiment. *American Journal of Physics 66*(5), 377–379.

Hanson, R. and D. Awschalom (2008). Coherent manipulation of single spins in semiconductors. *Nature 453*(7198), 1043–1049.

Keeler, J. (2005). *Understanding NMR spectroscopy*. Wiley.

Keppeler, S. (2003). *Spinning Particles - Semiclassics and Spectral Statistics*. Berlin: Springer-Verlag.

Kitcher, P. (1993). *The advancement of science: Science without legend, objectivity without illusion*. New York: Oxford University Press.

Ladyman, J. and D. Ross (2007). *Everything Must Go: Metaphysics Naturalised*. Oxfords: Oxford University Press.

Lewis, P. J. (2018). On the status of primitive ontology. In S. Gao (Ed.), *Collapse of the Wave Function*. Cambridge: Cambridge University Press.

McLeish, C. (2006). Realism bit by bit: Part ii. disjunctive partial reference. *Studies in History and Philosophy of Science Part A 37*(2), 171–190.

McMullin, E. (1984). A case for scientific realism. In J. Leplin (Ed.), *Scientific Realism*, pp. 8–40. Los Angeles: University of California Press.

Morrison, M. (2007). Spin: All is not what it seems. *Studies in History and Philosophy of Modern Physics 38*(3), 529–557.

Myrvold, W. (2018). Ontology for collapse theories. In S. Gao (Ed.), *Collapse of the Wave Function*. Cambridge University Press.

Niiniluoto, I. (1999). *Critical scientific realism*. Oxford: Oxford University Press.

Norsen, T. (2014). The pilot-wave perspective on spin. *American Journal of Physics 82*, 337–348.

Norsen, T. (2018). On the explanation of born-rule statistics in the de broglie-bohm pilot-wave theory. *Entropy 20*(6), 422.

Ohanian, H. C. (1986). What is spin. *American Journal of Physics 54*(6), 500–505.

Pearle, P. (1989). Combining stochastic dynamical state-vector reduction with spontaneous localization. *Physical Review A 39*(5), 2277.

Psillos, S. (1999). *Scientific realism: How science tracks truth*. London: Routledge.

Putnam, H. (1982). Three kinds of scientific realism. *The Philosophical Quarterly (1950-) 32*(128, Special Issue: Scientific Realism), 195–200.

Raimond, J.-M. and V. Rivasseau (2009). *The spin*. Basel: Birkhauser.

Ruetsche, L. (2018). Getting real about quantum mechanics. In J. Saatsi (Ed.), *Routledge Handbook of Scientific Realism*. London: Routledge.

Saatsi, J. (2015). Replacing recipe realism. *Synthese*, 1–12.

Saatsi, J. (2016). What is theoretical progress of science. *Synthese*, 1–21.

Saatsi, J. (Ed.) (2018). *The Routledge Handbook of Scientific Realism*. London: Routledge.

Sakurai, J. J. (1995). *Modern quantum mechanics*. AAPT.

Schurz, G. (2009). When empirical success implies theoretical reference: A structural correspondence theorem. *The British Journal for the Philosophy of Science 60*(1), 101–133.

Sebens, C. T. (2015). Quantum mechanics as classical physics. *Philosophy of Science 82*(2), 266–291.

Townsend, J. S. (2000). *A modern approach to quantum mechanics*. University Science Books.

Tumulka, R. (2018). Paradoxes and primitive ontology in collapse theories of quantum mechanics. In S. Gao (Ed.), *Collapse of the Wave Function*. Cambridge: Cambridge University Press.

van Fraassen, B. C. (1980). *The Scientific Image*. Oxford: Clarendon Press.

Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford: Oxford University Press.

Zee, A. (2016). *Group theory in a nutshell for physicists*. Princeton: Princeton University Press.