

HOW TO THINK ABOUT INDISCERNIBLE PARTICLES

by

Daniel Joseph Giglio

A Thesis Submitted in
Partial Fulfillment of the
Requirements for the Degree of

Master of Arts

in Philosophy

at

The University of Wisconsin-Milwaukee

May 2013

ABSTRACT
HOW TO THINK ABOUT INDISCERNIBLE PARTICLES

by

Daniel Joseph Giglio

The University of Wisconsin-Milwaukee, May 2013
Under the Supervision of Professor Robert Schwartz

Permutation symmetries which arise in quantum mechanics pose an intriguing problem. It is not clear that particles which exhibit permutation symmetries (i.e. particles which are *indiscernible*, meaning that they can be swapped with each other without this yielding a new physical state) qualify as “objects” in any reasonable sense of the term. One solution to this puzzle, which I attribute to W.V. Quine, would have us eliminate such particles from our ontology altogether in order to circumvent the metaphysical vexations caused by permutation symmetries. In this essay I argue that Quine’s solution is too rash, and in its place I suggest a novel solution based on altering some of the language of quantum mechanics. Before launching into the technical details of indiscernible particles, however, I begin this essay with some remarks on the methodology – instrumentalism – which motivates my arguments.

TABLE OF CONTENTS

§1	METHODOLOGY		
	1.1	GENERAL REMARKS	p. 1
	1.2	WHOLESALE INSTRUMENTALISM	p. 4
§2	THE PUZZLE OF INDISCERNIBLE PARTICLES		
	2.1	INTRODUCING THE PUZZLE	p. 9
	2.2	MATHEMATICAL FORMULATION OF PIP	p. 12
	2.3	QUINE'S PROPOSED SOLUTION	p. 18
§3	AN ALTERNATIVE SOLUTION		
	3.1	COUNT NOUNS AND MASS NOUNS	p. 21
	3.2	'ELECTRON' AS A MASS NOUN	p. 25
	3.3	OBJECTIONS CONSIDERED	P. 31
§4	CONCLUDING REMARKS		p. 34
	WORKS CITED		p. 37

At length, however, a crisis occurred. For some three thousand years the Fourth Men had pursued their research with constant success, but latterly progress had been slow. It was becoming increasingly difficult to devise new lines of research. True, there was still much detail to be filled in, even in their knowledge of their own planet, and very much in their knowledge of the stars. But there was no prospect of opening up entirely new fields which might throw some light on the essential nature of things. Indeed, it began to dawn on them that they had scarcely plumbed a surface ripple of the ocean of mystery. Their knowledge seemed to them perfectly systematic, yet wholly enigmatic. They had a growing sense that though in a manner they knew almost everything, they really knew nothing.

~Olaf Stapledon, Last and First Men

§1 METHODOLOGY

1.1 *General Remarks*

My aim in this essay can be described as narrow in one respect, broad in another. That is to say, my primary focus is to establish a very esoteric solution to an equally recondite problem of quantum mechanics. Yet at the same time, looming in the background is the cherished hope of mine that the novelty and fecundity of my proposed solution will engender credence and plausibility to the methodological commitments which underpin my primary arguments. I realize that the derivation of an attractive conclusion (assuming, for the nonce, that my conclusion *is* in fact attractive) does not by itself establish the validity of the methodology invoked in deriving it. But it doesn't hurt. Moreover, the only reasonable touchstone for the validity of some given methodology is the satisfactoriness of the fruit it bears; a methodology can only be valid insofar as it actually *works* in practice, produces satisfactory argumentative results. And the claim of the previous sentence is, of course, intimately bound up with the pragmatic perspective from which I propose to argue in this essay.

But first I must back up a few steps. A methodology is, broadly speaking, a set of practices, procedures, and rules for engaging in some particular line of inquiry. I take it to be uncontroversial that argumentation presupposes methodology. In other words, it is constitutive of the very practice of exchanging justifications that one adopts some particular justificatory framework from which to barter – some protocol which dictates the dialectical norms at play. One cannot argue from methodologically-neutral ground any more than one can take a photograph from no particular perspective whatsoever.

One can, perhaps, evade being labeled as the proponent of any particular methodology by abruptly changing methodologies each time one's interlocutor makes explicit some pattern of reasoning one has tacitly made use of. But such a wily tactic amounts to nothing more than pure philosophical spontaneity, and any inquirer who argues in this way drains the meaning from 'inquirer' and 'argues' to the point that they do not deserve to be taken seriously.

Like it or not, then, to the extent that each of us argues consistently, each of us occupies some particular methodological perspective. One can (and ought to) question the validity of one's methodology quite often; but this higher-level inquiry cannot proceed freely of methodological commitments any more than ground-level inquiry can. That is because the justification of the argumentative patterns utilized in ground-level inquiry will itself take on some argumentative pattern or other. If one exercises a healthy amount of philosophical shrewdness, then this higher-level inquiry will proceed much more stringently than the ground-level inquiry.¹ But from this point the higher-level methodology can in turn be brought into question, and it becomes clear that this highly self-reflective game of bringing one's own methodology into question can continue *ad infinitum*.

Thus we are forced to come to terms with an unavoidable consequence of the practice of exchanging justifications – i.e. the realization that methodological justification must come to an end somewhere: “If I have exhausted the justifications I have reached bedrock, and my spade is turned. Then I am inclined to say: ‘This is simply what I do’” (Wittgenstein 1953, §217). However, this is not the dire situation that it may seem to be

¹ This point is nicely made by Arthur Fine (1986, p. 114 ff.) in his discussion of mathematical vs. meta-mathematical and scientific vs. meta-scientific reasoning.

at first glance. We are not left with a philosophical arena in which “anything goes,” methodologically speaking, for rival methodological foundations still have a number of ways of competing with each other. Some may fail by their own standards. Others may lead to empirically outrageous conclusions. Thus there is no reason to rule it out *a priori* that a few winners (or perhaps even one exceptionally hardy warrior) might emerge.

The methodological bedrock of this essay is, in a word, pragmatic. One consequence – perhaps the most *salient* consequence – of building a philosophical structure on this foundation is that a theory will be true precisely insofar as it works in practice. A true theory, rather than being one which accurately “maps” or “corresponds to” the way the world is, will be one upon which we can ride, one that can “carry us prosperously from any one part of our experience to any other part, linking things satisfactorily, working securely, simplifying, saving labor” (James 1907, p. 34). There is simply nothing further that ‘truth’ could have ever meant in the first place, according to the pragmatic perspective. In this way our theories become “instruments, not answers to enigmas” (*ibid.*, p. 32).

As already stated, my primary aim in this essay is to articulate a very specific solution to a very specific problem which arises in quantum mechanics. Therefore, given what has been said above regarding methodological bedrock (and the fact that, like it or not, one’s spade will turn sooner or later) I must resist the temptation to launch a full-scale apology for this pragmatic perspective from which my main argument proceeds; truth be told, I possess neither the space nor the expertise to embark on such a colossal project. Nevertheless, I cannot help myself from briefly stating my case for occupying

the pragmatic perspective, and explaining why one particularly threatening objection turns out, upon close scrutiny, to cut no philosophical ice.

1.2 *Wholesale Instrumentalism*

In the context of philosophy of science, the pragmatically-inspired methodology gestured at above is most naturally referred to as ‘instrumentalism.’

The central claim of the instrumentalist view is that a theory is neither a summary description nor a generalized statement of relations between observable data. On the contrary, a theory is held to be a rule or a principle for analyzing and symbolically representing certain materials of gross experience, and at the same time an instrument in a technique for inferring observation statements from other such statements (Nagel, 1961 p. 129).

In this way a scientific theory functions not as a representation of objective reality, but rather as “a rule or a guide for making logical transitions from one set of experimental data to another set” (ibid.). According to the instrumentalist, the goal of scientific investigation is “not to formulate theses that mirror or copy reality. It is to organize experience in ways that are useful for explanation and prediction” (Schwartz 2012, p.87). For an example of the instrumentalist viewpoint in action, consider the Rydberg Formula for Hydrogen atoms: $1/\lambda = R*(1/(n_1)^2 - 1/(n_2)^2)$. The instrumentalist does not interpret ‘ n_1 ’ and ‘ n_2 ’ in this equation as representing the energy levels occupied by some electron before and after a transition. Instead, the instrumentalist simply sees this equation as a tool for predicting the observed wavelength (λ) of spectral lines, and she does not trouble herself by worrying about the mechanism which renders the Rydberg Formula an effective tool. She may deny the reality of the underlying electron and its

transition, thus adding a fictionalist or anti-realist flavor to her instrumentalism. Or she may simply refrain from asking the question “is electron-transition *really* what’s going on here?” thereby adding a flavor of insouciance toward the realist/anti-realist debate.

Now the objection against instrumentalism that I have in mind runs as follows. You tell me (says the instrumentalist’s interlocutor) that scientific theories are nothing but tools. Very well. But now you owe me an explanation of why the tool works, why it works as well as it does, why other tools would not complete that particular task, and an entire host of related questions. Any conceivable answer to these questions will necessarily rely upon mentioning the way that the world is, independently of us theorizers. For example, if asked why a screwdriver is so successful at removing screws, any adequate story would have to mention the geometry of the screw, its tendency to strip, its material composition, certain features of the item in which it is embedded, etc. The point is that in mentioning these objective features of the screw the instrumentalist saws off the very branch on which she sits.²

A few comments are in order at this point. First, we must bear in mind that the tool/instrument analogy is just that: an analogy. All analogies must break down at some point. An analogy works by taking two categorically distinct entities (e.g., a scientific theory and a screwdriver) and highlighting some common feature (e.g., the fact that they are both used intentionally to execute some specific task) while disregarding other features. But there is not a perfect isomorphism between the two, and the mathematically-based explanations utilized in scientific theories proceed quite differently from carpentry-based explanations. So the fact that not everything which can be said about one can be

² Thanks to Robert Kraut for continually and rigorously pushing this objection on me in conversation.

said about the other does not (by itself) constitute an objection to the instrumentalist viewpoint.

However, this response of mine (though on point) is too quick and dismissive, and by no means adequate. Presumably the instrumentalist's interlocutor will happily agree about this inherent limitation of analogizing, at which point there still remains a very genuine concern about the coherency of the instrumentalist position: how can the instrumentalist help herself to her beloved toolkit without committing herself to some sort of realism regarding the items on which her tools operate?

It is at this point that I think an important distinction must be made within the wide spectrum of positions which fall under the heading 'instrumentalism.' I shall call the two positions to be distinguished "retail" instrumentalism and "wholesale" instrumentalism.³ The retail instrumentalist only invokes the theories-as-tools analogy in small quantities. She sees theoretical posits of scientific theories as *merely* instrumental, as fictional entities which (by some miraculous means) manage to lead to the right predictions without accurately mirroring reality. Thus she is happy to acknowledge the full-blown reality of the extended, colored, substantial desk in front of her; but the buzzing mass of atoms is mere fiction – a physicist's conceit – which allows her to easily calculate things like temperature flow and charge distribution across her desk. The retail instrumentalist is thus an anti-realist with respect to the microscopic objects of scientific theories (e.g. electrons, photons, and nuclei), but nevertheless a realist with respect to the macroscopic objects of common sense (e.g. voltmeters, patches of light, and nuclear explosions).

³ Cf. Schwartz (2012, p. 86 ff.).

On the other hand, the wholesale instrumentalist goes in for the theories-as-tools analogy at every opportunity. Whether the posit in question is housed in a theory of physics or a theory of everyday experience is immaterial to the wholesale instrumentalist: it's posits all the way down, so to speak. She refuses to phrase things in terms of correspondence with reality, and rejects the realist/anti-realist dichotomy from the start. All that matters to the wholesale instrumentalist is whether the theoretical posit in question helps us to make successful predictions, economize previous chunks of theory, etc. Surely enough there are some crucial differences between the two types of posit just mentioned. For one thing, the macroscopic posits of everyday life are *interpolations* from experience, whereas the microscopic posits of physical theories are *extrapolations* from experience. Accordingly, the posits of physical science, do, admittedly, "[take] us further out on a limb" than the homely posits of common sense do (Schwartz 2012, p. 84). But the salient point is that they are both posits nonetheless. This significance of this last point is nicely stated by William James:

There is no *ringing* conclusion possible when we compare these types of thinking [i.e. scientific positing vs. common-sense positing], with a view to telling which is the more absolutely true. Their naturalness, their intellectual economy, their fruitfulness for practice, all start up as distinct tests of their veracity, and as a result we get confused. Common sense is *better* for one sphere of life, science for another... but whether either be *truer* absolutely, Heaven only knows. (1907, p. 93)

For my part, I think the retail instrumentalist has no coherent response to the question posed above: namely, how can the instrumentalist explain the success of her toolkit without mentioning some feature of the way the world is in itself, independently of us theorizers? By adverting to these features, the retail instrumentalist (a professed anti-realist regarding physical theories) commits herself to an explanation which is realist

in spirit; thus she tries to have her cake and eat it too. The wholesale instrumentalist, however, is in much better shape. She (unlike the retail instrumentalist) sees no interesting distinction between the posits of fundamental physics (e.g. energy-level transitions of electrons and subsequent emission of photons) and the posits of everyday experience (e.g. thin, distinctly colored lines observed with the naked eye). In order for the objection under consideration to cut any ice, precisely this distinction would have to be maintained by the instrumentalist. The key, then, to dodging the blow of this objection, lies in the conviction that nothing is truly at stake in the dispute between realists and anti-realists.

One consequence of wholesale instrumentalism is that a certain variety of pluralism emerges. By 'pluralism' I mean the thesis that rival theories may each play useful roles in inquiry, though neither theory can be completely recast within the other, and though the two may even stand in stark opposition to one another. Under a more orthodox conception of scientific theories as representations of the ultimate laws of nature, pluralism would be a strange thesis indeed; presumably there is only one "nature," and therefore only one way of perfectly describing it in its entirety. However, when we conceive of theories as tools, quite the opposite situation obtains: just as different tools may be used to accomplish the same task, different theories may be used in pursuing the same line of inquiry. Both a jigsaw and a handsaw may suffice to make the desired cut, though perhaps one is more appropriate given the level of precision required. Similarly, both the particle and wave models of light serve to make successful predictions in the laboratory, though on any given occasion it will make more sense to use one and not the other.

I mention the pluralistic thesis now because it will play a large role later in this essay where I will suggest a novel way thinking about quantum particles which, though at odds with the traditional conception, nonetheless serves a very specific purpose in “carrying us prosperously from one part of experience to another” and is therefore a valid tool to be used in the act of theorizing, so far as the wholesale instrumentalist is concerned. Theories are, according this methodology, “geared to deal with local problems by diverse means and under diverse constraints” (Schwartz 2012, p. 73). Accordingly, my aim in the rest of this essay will be to manufacture a new intellectual instrument for the express purpose of dealing with a local problem in quantum mechanics. The utility of this specially-designed tool will, I hope, become evident by the end of the essay, for there is nothing quite as satisfying as using the proper tool for the job.

§2 THE PUZZLE OF INDISCERNIBLE PARTICLES

2.1 Introducing the Puzzle

The legendary theoretical physicist Richard Feynman once wrote, “I think I can safely say that nobody understands quantum mechanics” (1965, p. 129). This famous remark of his appeared in the context of a discussion of the confounding nature of wave-particle duality. A not unrelated, yet equally vexing, enigma concerns the indiscernibility that is inherent in collections of elementary particles of the same kind. The Puzzle of Indiscernible Particles (henceforth ‘PIP’) will receive a precise mathematical elucidation in §2.2, but for now the puzzle can be summarized in the following way: there are cases in which any two particles in some larger collection can be swapped (with respect to their

locations, say) without this yielding a new state for the overall collection of particles. Another way to state the puzzle would be to say that certain collections of particles exhibit *permutation symmetry*: two particles can be transposed with one another, but this transformation maps the entire collection back onto the same state, rather than a new state. If a collection of particles exhibits permutation symmetry in this way, then throughout this essay I will say that those particles are *indiscernible* from each other.⁴

I suppose the mere fact that there are viewpoints from which certain things are indiscernible is not by itself so confounding. For example, one rarely cares to discern between the eggs which enter into the cake which is being baked: any three of them from the dozen just purchased will do. So from the point of view of the recipe, we might say, those dozen eggs are indiscernible from each other. But of course a few of those eggs may be cracked, or perhaps one is brown while the remaining eleven are white, or maybe a mischievous kitchen-assistant has clandestinely hard-boiled half of them, in which case there is an obvious viewpoint from which they are discernible. However, the truly nefarious ramification of PIP (as will become clear in §2.2) is that there is no physical fact about the particles in question which could serve to distinguish them; no viewpoint from which we could tell one apart from the other. All elementary particles of a given type therefore seem to be indiscernible not just in *practice*, but in *principle*.

Because of the interchangeability just mentioned, the particles in question cannot be said to exhibit individuality in any robust sense, and their identity- and persistence-

⁴ It is somewhat misleading to speak of the “location” of a quantum particle, for such particles (electrons, in particular) are most accurately described as being smeared throughout a region of space, rather than being localized at specific points. We can, of course, locate the regions to a certain degree of accuracy. But in cases of excessive overlap between regions, we may consider talk of “location” as shorthand for referring to which energy level of the atom said electron occupies. In this way, swapping the “location” of two electrons would amount to swapping the first quantum number in the state-description of the two electrons.

conditions are similarly elusive. It is not clear in what sense such bizarre entities can even be called “physical objects.” In response to this very puzzle, W.V. Quine⁵ flirts with the prospect of parting ways with the notion of *physical object* altogether in order to embrace a field theory which merely posits regions of spacetime that can be instantiated with certain qualities.⁶ From this point he goes even further, suggesting that it may be possible to reduce those regions of spacetime to mere quadruples of numbers, thereby supplanting the ontology of physics with that of pure set theory. However, this would not be a picture in which “our system of the world reduces [completely] to set theory,” for the *ideology* of physics is not reducible to the *ideology* of set theory in any straightforward manner (1976, p. 503). Thus Quine ends that paper by noting that “[w]e might come to look to pure mathematics as the locus of ontology as a matter of course, and consider rather that the lexicon of natural science, not the ontology, is where the metaphysical action is” (ibid., p. 504). The conclusion of the previous sentence is electrifying, to say the least, and worthy of extended discussion. However, that is not my task in this essay. Before Quine’s intriguing insouciance toward ontology can be considered, we must first become certain that his ontological leap from physical objects to regions of spacetime is warranted. Simply put, I do not think that the leap is warranted (and certainly not necessitated), and it is my aim in this essay to explain why.

That we can circumvent PIP by jettisoning the naive conception of matter is indisputable; there can be no Puzzle about Indiscernible Particles if there are no such

⁵ See “Whither Physical Objects?” (1976); also *Pursuit of Truth* (1990, ch. 2).

⁶ It is important to note here that Quine is *not* suggesting that familiar physical objects like tables and chairs do not exist. Rather, he is exploring a revisionary claim about the ontology of particle physics: the quantifiers of fundamental physical theories ought to range over regions of spacetime rather than physical objects. The ontology of our “everyday” theories of tables, chairs, persons, dogs, etc. can go unaffected by this shift in the ontology of particle physics.

particles in the first place. However, parting ways with the homey notion of *physical object* is by all counts a drastic maneuver, one which strikes me as guilty of throwing the baby out with the bathwater. Accordingly, it is my contention in what remains of this essay to show that PIP creates a problem not for physical objects *in general*, but rather for a special subclass of physical objects – namely, those which are the referents of *count nouns*. I will then propose a way of reconfiguring some of the language of quantum mechanics in terms of *mass nouns*. This linguistic shift provides a way for us to retain the notion of *physical object* despite the difficulties raised by PIP, and despite Quine’s advice. If my reasoning is sound, then mathematics’ claim to being the “locus of ontology” will come nowhere near being “a matter of course.”

Throughout this essay I will argue from the wholesale instrumentalist perspective (to which Quine was explicitly sympathetic⁷) and I will help myself to certain resources from the Quinean framework, the most important of which being the Indeterminacy of Translation thesis. So in a way, my aim in this paper is to articulate what Quine *should have* said in response to PIP. But before launching into Quine’s solution, its shortcomings, and my alternative solution, we first need to examine exactly what PIP amounts to.

2.2 *Mathematical Formulation of PIP*

The program of statistical mechanics is to characterize a system in terms of the overall distribution of particles across a range of states; depending on what we are interested in, these may be energy states, positions, velocities, or a variety of other

⁷ See “Posits and Reality” in Quine 1966 (pp. 233-241).

quantities. For concreteness, I will focus on energy states in this section. Without ever knowing what any particular particle is doing, we use this distribution to predict the overall behavior of the system *in the long run*, which is to say when it reaches a state of equilibrium.

It is helpful to introduce a few technical terms at this point. The “occupation number” of an energy state is the number of particles that occupy that energy state. The collection of occupation numbers for every allowable energy state therefore provides a way of characterizing the entire system, and we call this the “configuration” of the system. As a toy example, consider a system with just three allowable energy states (S_1 , S_2 , and S_3) and two particles (A and B). Here is one possible configuration for the system: (2, 0, 0). What this means is that in this configuration, two particles (to wit: A and B) occupy S_1 , zero particles occupy S_2 , and zero particles occupy S_3 . Other possible configurations for the system would then be: (0, 2, 0), (0, 0, 2), (1, 1, 0), (1, 0, 1), and (0, 1, 1).

The ultimate goal in statistical mechanics is to calculate the most probable occupation number for each energy state; in other words, the goal is to predict what fraction of particles will, in the long run, settle down into each energy state. In order to perform this calculation, we must make a decision as to how we are going to individuate among overall states of the system. For example, consider the configuration characterized as (1, 1, 0). How many different ways are there for the system to achieve this configuration? The natural answer is two: the first way places A in S_1 and B in S_2 (and nobody in S_3), whereas the other way places B in S_1 and A in S_2 (and nobody in S_3).

If we work under this assumption – that interchanging two particles produces a *distinct* state – then we can derive⁸ what is known as the Maxwell-Boltzmann distribution:

$$N(\epsilon) = e^{(\mu - \epsilon)/(kT)} \quad (\text{MB})$$

What this equation tells us is that, in the long run, we can expect to find that the number of particles with some particular energy ϵ is given by some not-so-complicated function of that energy (ϵ), the chemical potential of the system (μ), Boltzmann's constant (k), and the temperature of the system (T). The MB distribution provides an astonishingly accurate characterization of the behavior of gas molecules (within a limited range of parameters), and lands us in the realm of *classical* statistical mechanics.

But suppose we had gone a different route from the beginning. Instead of saying that there are *two* distinct ways of achieving the configuration (1, 1, 0), let us now say that there is only *one* way. In other words, we are saying that swapping A and B with respect to their energy states makes *no difference whatsoever* to the overall state of the system. Such a collection would exhibit *permutation symmetry*: two particles can be transposed with one another, but this transformation maps the entire collection back onto the same exact state, rather than a new state. Not surprisingly, this assumption gives rise to an altogether different mathematical derivation. From this point, depending on whether we allow either one particle or an unlimited number of particles to occupy each state, we can derive either the Fermi-Dirac distribution:

$$N(\epsilon) = 1 / [e^{(\epsilon - \mu)/(kT)} + 1] \quad (\text{FD})$$

or the Bose-Einstein distribution:

$$N(\epsilon) = 1 / [e^{(\epsilon - \mu)/(kT)} - 1] \quad (\text{BE})$$

⁸ See Griffiths (2005, pp. 233-241) for all the gory details.

What we find experimentally is that the behavior of a system of particles with half-integer spin (fermions) is brilliantly described by the FD distribution, and that the behavior of a system of particles with integer spin (bosons) is brilliantly described by the BE distribution. The difference between these two distributions is that FD allows a maximum of one particle per state, whereas BE allows an unlimited number of particles to occupy each state.⁹ In either case, this lands us in the realm of *quantum* statistics.

So what is the upshot? The upshot is that whether or not the particles in our system are discernible from one another – i.e., whether or not swapping two of them yields a *new* state – gives way to unequivocal mathematical differences in the distribution which characterizes the behavior of that system in the long run. Given the unique shape of each distribution, we can determine whether the particles in some particular collection are discernible or not by seeing which distribution that collection obeys. If the collection obeys the MB distribution, then the particles in question are discernible; if the collection obeys either the FD or the BE distribution, then the particles in question are indiscernible. When we empirically investigate the matter, there is simply no question that all fermions of a given kind (e.g. all electrons, all top quarks) as well as all bosons of a given kind (e.g. all photons, all gluons) obey a statistics which reflects the assumption that all particles in the system are indiscernible from each other.¹⁰

⁹ For simplicity, I will ignore spin states throughout this paper.

¹⁰ Actually, Muller & Saunders (2008) have constructed an intricate framework in which fermions *are* in fact discernible (while bosons remain indiscernible). I take their position to be compelling, but it does not seem to me that the dust has settled quite yet. Therefore, for the purposes of this paper, I will occupy the more traditional camp which maintains that both the Fermi-Dirac statistics *and* the Bose-Einstein statistics describe systems of indiscernible particles (rather than just the Bose-Einstein statistics). See Huggett & Norton (forthcoming) for an insightful discussion of some highly technical issues with the apparatus of Muller and Saunders.

A nice heuristic for illustrating the difference between the classical distribution (MB) and the quantum distribution (BE) involves flipping two coins.¹¹ Suppose we are interested in measuring the probability of both coins showing up “heads.” The MB distribution predicts that this occurs $\frac{1}{4}$ of the time, whereas the BE distribution predicts this to occur $\frac{1}{3}$ of the time. Given the fundamental assumption of statistical mechanics (roughly: that each distinct state is equiprobable), we can note a very peculiar metaphysical difference between a pair of coins which obeys MB, and a pair of coins which obeys BE. If each distinct state is equiprobable, then the reciprocal of the probability for any arbitrary state will reveal the total number of distinct states accessible to the system. Therefore, the MB coins would seem to admit of four distinct states (namely: heads-heads, tails-tails, heads-tails, and tails-heads) whereas the BE coins would seem to admit of only three distinct states (namely: heads-heads, tails-tails, and a mixed state that is indifferent as to which coin is heads and which coin is tails). This last state – the *mixed* state – is the embodiment of a permutation symmetry, and precisely what gives rise to the Puzzle of Indiscernible Particles. It is of paramount importance to notice that the inability to discern between two BE coins (or particles) has nothing to do with the physical limitations of our laboratory equipment. Permuting two such coins (or particles) yields the exact same physical state, which is tantamount to saying that there is no physical fact which could even in principle serve to distinguish them. Permuting two such items *leaves no physical trace*.

The metaphysical ramifications of this aspect of quantum statistics are drastic. It is not even clear whether particles which obey FD or BE ought to be called “objects,” for

¹¹ With a few slight adjustments, the exact same point can be made regarding FD. But just for concreteness, in this example I will focus on BE and ignore FD.

it immediately follows from the above discussion they do not obey Leibniz's Law (a.k.a. the Identity of Indiscernibles), which one may reasonably take to be the touchstone of objecthood. Leibniz's Law states that no two objects can share *all* of their properties. More explicitly: $\forall x \forall y (\forall F (Fx \leftrightarrow Fy) \rightarrow x=y)$. To say that there is no physical fact which could serve to discern the particles in question is just to affirm the antecedent of Leibniz's Law, for this is just another way of saying that they have all of their physical properties in common.¹² But it would be a grave mistake to assert that the two particles are identical to each other. As David Lewis says:

Identity is utterly simple and unproblematic. Everything is identical to itself; nothing is identical to anything else except itself. There is never any problem about what makes something identical to itself; nothing can ever fail to be. And there is never any problem about what makes two things identical; two things can never be identical (1986, pp. 192-193).

Compare with a similarly pithy remark from Wittgenstein: "To say of two things that they are identical is nonsense, and to say of one thing that it is identical with itself is to say nothing" (1921, §5.5303). The example we are considering is one in which there are *two* particles present; so to say that they are identical would be to completely bungle the meaning of 'identity.' Furthermore, the claim that there are two (or three, or seventeen, etc.) particles present is directly verifiable. I have deliberately been trying to maintain the generality of the system under consideration; but if we so desired, we could flesh out the details of our system, actually construct it in the laboratory, measure the overall charge, divide this quantity by the charge of a single electron (if it is electrons that we are talking about), and at that point we'd know *exactly* how many electrons are present. We

¹² Whether the particles may possess some non-physical property (a haecceity, a "thisness") is a thorny issue which I will not have space to address in this essay. Readers interested in this line of argument should consult the writings of Paul Teller (1997, ch. 2; also 1998).

would then have an example of two electrons (or three, or seventeen, etc.) that are indiscernible from each other, yet most certainly not identical.¹³

Thus we are faced with a dilemma: either part ways with our famed *a priori* principle of metaphysics (and thus admit that there are non-identical indiscernibles), or bite the bullet and extirpate physical objects from our ontology altogether. As we are about to see in §2.3, Quine has at times recommended that we take the second horn of the dilemma. But as is often the case with dilemmas, there are ways of recasting the situation so as to avoid *both* horns; I will explore such a strategy later in the essay.

2.3 *Quine's Suggested Solution*

Quine's recommendation is that, in light of PIP, we reduce quantum particles to the mere regions of spacetime which they occupy and allow these regions to take on certain qualities. In other words, the ontology of particle physics is to contain regions of spacetime, but not particles. In building up to this conclusion, he asks us to consider what we mean by 'physical object' in the first place:

We think first of bodies, but the notion of a body is both too vague and too narrow. It is too vague in that we are not told how separate and cohesive and well rounded a thing has to be in order to qualify as a body. And it is too narrow, since for ontological purposes any consideration of separateness and well-roundedness is beside the point. Rather, let us understand a physical object, for a while, simply as the aggregate material content of any portion of space-time, however ragged and discontinuous (1976, p. 497).

In this passage Quine is searching for a gloss on 'physical object.' To see why 'body' is "too vague" to succeed, consider a lamp whose shade is not affixed to the base, and therefore falls off every time the base is slightly bumped. We would hesitate to call

¹³ See French & Krause (2006) for a nice entry point into the vast literature on how indiscernible particles relate to Leibniz's Law.

the base-plus-shade a single, unified body. Duct-tape the two together, however, and now we are unambiguously talking about a body. Intermediate cases would certainly arise – e.g., there is enough tape to keep them fastened together not all the time, but only after 50% of all bumps – in which case our ordinary usage of the term ‘body’ would not settle the question of whether the base-plus-shade indeed qualifies as a body. Whether something qualifies as a physical object, however, seems to be much more cut and dry than this. Something is either a physical object, or it isn’t; there is no middle-ground here.

Moreover, ‘body’ is too narrow to serve as a gloss for ‘physical object’ in that *anything at all* can be construed as an object for some theory or other, regardless of whether it qualifies as a body. If we so desired, we could construct a theory which quantified over catdogs (the result of adding together the material content of a region of spacetime occupied by a cat with another occupied by a dog) rather than cats and dogs separately. Though catdogs are too “ragged and discontinuous” to count as bodies, they would nevertheless be physical objects of the theory in question. For Quine, an object is simply whatever a theory quantifies over: to be is to be the value of a bound variable.¹⁴ Therefore catdogs gain object-status simply because they are the values of certain variables; and they gain their *physical* object-status in virtue of the fact that they are composed of matter. Notice that on this front we can see Quine’s wholesale instrumentalism emerge quite clearly. Contained in his criterion of ontological commitment (“to be is to be the value of a bound variable”) is the idea that no class of objects is “more real” than another; so long as the item in question is the posit of some

¹⁴ See “On What There Is,” (Quine 1953, pp. 1-19).

theory, it is just as “real” as any other theoretical posit. Surely enough certain posits may be more *useful* than others (and presumably the catdog is not one of these), but the question of utility is completely orthogonal to the question of existence; being posited *suffices* for existence.

Quine goes on to point out that even “the aggregate material content of any portion of space-time,” otherwise known as the “naive conception of matter,” similarly fails as a gloss for ‘physical object’ (1976, p. 498). This is primarily due to the difficulties posed by PIP. Whereas we would expect that transposing two physical objects would yield a new state for the system, we saw in §2.2 that transposing two electrons does no such thing; by definition, particles which exhibit permutation symmetry yield the *same* overall state when transposed. Given this confounding empirical result, Quine suggests that “we should not think of individual electrons... at all, but of states of [the regions of space-time which they occupy]” (ibid., p. 499). In short, the notion of physical object has “gone so tenuous that we find ourselves turning to the space-time regions for something to cleave to” (ibid., p. 499). On this same point, Quine writes elsewhere that “[i]t would seem then not merely that elementary particles are unlike bodies; it would seem that there are no such denizens of space-time at all, and that we should speak of places... merely as being in certain states... rather than as being occupied by two things” (1990, p. 35).

It is worth noting that these field-theoretic sentiments receive implicit corroboration from elsewhere in the Quinean corpus. In particular, the “No Entity Without Identity” slogan would also seem to bolster the shift toward a field theory.¹⁵

¹⁵ See “Speaking of Objects,” (Quine 1969, pp. 1-25).

What is meant by this slogan is that we cannot coherently posit an entity without first being able to provide a rough characterization of its identity- and individuation-criteria. The electrons lately discussed fail to satisfy this criterion insofar as two of them can be swapped, though we would be (in principle) none the wiser for it. Furthermore, to identify a given electron over time often makes little sense, and “their integration into world lines, or enduring particles, is only a matter of convenience, and arbitrary in varying degrees” (Quine 1976, p. 498). This last point is nicely summarized by physicist David Griffiths:

You can always tell the [classical] particles apart, in principle – just paint one of them red and the other one blue, or stamp identification numbers on them, or hire private detectives to follow them around. But in quantum mechanics the situation is fundamentally different. You can’t paint an electron red, or pin a label on it, and a detective’s observations will inevitably and unpredictably alter its state, raising doubts as to whether the two had perhaps switched places... It’s not just that we don’t happen to know which is which; God doesn’t know which is which, because there is no such thing as “this” electron, or “that” electron” (2005, pp. 203-204).

§3 AN ALTERNATIVE SOLUTION

3.1 *Count Nouns and Mass Nouns*

As stated in §2.1, my aim in this paper is to show that Quine’s response to PIP is too drastic, and that there is a much better option on the table – one that can in fact be constructed out of the resources in Quine’s own idiosyncratic system of thought. In order to make this argument, I will rely upon the linguistic framework utilized by Quine in many of his writings, and in particular in *Word and Object* (1960, ch. 3). What I have in mind is the distinction between count nouns and mass nouns. These two varieties of noun are distinguished according to how their referents are *quantified*. A count noun can

be modified by a bare numeral, which is merely to say that the referent of a count noun is (not surprisingly) the kind of thing that can be *counted*. To illustrate with examples: ‘apple’, ‘brick’, and ‘cat’ are count nouns insofar as phrases like ‘three apples’, ‘four bricks’, and ‘five cats’ make good grammatical sense. Mass nouns, however, *cannot* be modified by a bare numeral. Notice how phrases like ‘three water’, ‘four sugar’, and ‘five trash’ – or, for that matter, even ‘three waters’, ‘four sugars’, and ‘five trashes’ – do not make good grammatical sense. Bare numerals do not attach to mass nouns because *any* quantity whatsoever of the stuff a mass noun refers to is treated as an undifferentiated unit: “any sum of parts which are water is water” (Quine 1960, p. 91). A unit of measurement must be specified in order for a mass noun to be quantified: three *gallons* of water, four *pounds* of sugar, five *bags* of trash, etc.

But notice that in common parlance certain terms vacillate between being used countwise and being used masswise. Compare the count-use of ‘apple’ in ‘There are three apples on the table’ with the mass-use of ‘apple’ in ‘This salad needs more apple.’ And although ‘water’ is most frequently used masswise (e.g. ‘Water is refreshing’), its count-use is not unheard of (e.g. ‘Still waters run deep’). Or, to take a cute example from Quine, notice the ambiguity in ‘Mary had a little lamb.’ If ‘lamb’ is being used countwise, then we have a statement about the stature of Mary’s pet; but if ‘lamb’ is being used masswise, then we have a statement about what Mary ate for dinner. Therefore terms like ‘apple’, ‘water’, and ‘lamb’ are neither inherently mass nouns nor inherently count nouns. Rather, it is the way in which one of these words is actually *used* on a given occasion that determines whether that specific token of the word is a count noun or a mass noun.

A more sophisticated and explicit illustration of the count/mass indeterminacy concerns translating certain Japanese phrases which contain lexical items known as “classifiers.”

Commonly [classifiers] are explained as attaching to numerals, to form compound numerals of distinctive styles. Thus take the numeral for 5. If you attach one classifier to it you get a style of “5” suitable for counting animals; if you attach a different classifier, you get a style of “5” suitable for counting slim things like pencils and chopsticks, and so on. But another way of viewing classifiers is to view them not as constituting part of the numeral, but as constituting part of the term – the term for “chopsticks” or “oxen” or whatever. On this view the classifier does the individuating work that is done in English by “sticks of” as applied to the mass term “wood,” or “head of” when applied to the mass term “cattle” (Quine 1969, p. 36).

In either case, what we have is a three-word Japanese phrase translated into a two-word English phrase: for example, ‘five oxen.’ But depending on how we imagine the classifier to be functioning, the third word in that Japanese phrase can emerge either as a count noun or a mass noun. According to the first view, the classifier declines the numeral 5 (which produces a style of 5 suitable for counting animals) which then modifies the count noun ‘ox.’ But according to the second view, the classifier applies to the mass noun ‘cattle’ (which results in something akin to ‘head of cattle’) which is then modified by the bare numeral 5. The first reading “makes for more of a feeling for the Japanese idiom” while the second reading of the classifier “makes for efficient translation into idiomatic English” (ibid., p. 38). But since both fit all verbal behavior equally well, the answer to the question of which is the “correct” translation of the third Japanese word – i.e. whether it is really a count noun or a mass noun – is “indeterminate in principle... there is no fact of the matter” (ibid., p. 38).

All we have done so far is examine particular examples of count/mass indeterminacy. However, there is an even stronger claim to be made: that it is indeterminate whether *any* given noun is being used masswise or countwise. This is a claim which Quine himself seems to be on board with, for it follows directly from the Indeterminacy of Translation thesis. A case in point derives from his famous example of the field linguist attempting to translate ‘Gavagai,’ which is exclaimed by a jungle native just after a rabbit scurries by.¹⁶ ‘Rabbit’ would obviously be a good candidate for the translation of ‘Gavagai,’ as would ‘some rabbit meat.’ The former treats ‘Gavagai’ as a count noun, whereas the latter treats it as a mass noun. The only way to accept one translation and reject the other with full confidence would be to query the native along the following lines: “Is this the same gavagai as that?” while pointing to, say, a live rabbit followed by the butchered and separated remains of that rabbit. However, as a radical translator, the very thing that our field linguist lacks is a full grasp of how demonstratives, copulas, and other such lexical items function in the native’s language. So the point of Quine’s Indeterminacy of Translation thesis is that there can be no fact of the matter as to which translation is “correct.” Both can account for all utterances from the native, given certain adjustments in other portions of the translation manual that is currently being constructed.

Now there do seem to be crucial differences between the examples of count/mass indeterminacy recently considered. ‘Mary had a little lamb’ seems to be more *ambiguous* than *indeterminate*, for our confidence in homophonically translating sentences of native English is quite high; the Japanese translation took us further out on a limb, but if this

¹⁶ See *Word and Object* (1960, ch. 2).

seriously concerned us we could presumably cross-check our translation with other bits of Japanese literature translated in the past; but the translation of the jungle native seems most radical of all, and perhaps the only one which fully deserves to be called “indeterminate.” However, the fact that some of these translations are much more tenuous than others is tangential to the point being made – i.e. that *in principle* there will always be two viable (though incompatible) translations, one which produces a mass noun and one which produces a count noun. As Quine frequently emphasizes, “radical translation begins at home” (ibid., p. 46). As children learning language from scratch on mother’s knee, we must go through the same complex process of linguistico-behavioral guess-and-check as we do when translating the words of a long-time Japanese friend, or even a recently made jungle acquaintance. In either case, the translation is indeterminate through and through. Multiple (and incompatible) translations of some particular lexical item are always viable given adjustments in other portions of the language already translated. It then follows as a corollary that translation is indeterminate when we step into the laboratory and attempt to translate the physicist’s theories into plain English; and this is precisely the point which I wish to exploit throughout the remainder of the essay.

3.2 *‘Electron as a Mass Noun*

Now, finally, what is all talk of count and mass nouns in aid of? The answer lies in how these concepts connect up with PIP. *The claim that I have been driving toward, and am now ready to unpack, is that PIP creates a problem for the referents of count nouns, but not for the referents of mass nouns.* We have already seen what the problem that PIP poses for count nouns amounts to: two electrons can be swapped without this yielding a

new state. This is problematic in that electrons – when construed as separate, enduring entities with individual careers – should resist such radical interchangeability. The same puzzle would arise if we discovered, somehow, that swapping the position of two apples in a barrel did not in any way alter the state of the barrel of apples; or that we could swap a few bricks between adjacent houses without altering the overall state of the two-house system. These would be genuine puzzles if they obtained in reality. Fortunately we can make sense out of the idea that the apple-swap produces a new barrel of apples (at least in some sense), and that the brick-swap produces a new two-house system (at least in some sense). The trouble is that there seems to be no clear sense in which an electron-swap produces a new state.

But notice how the puzzle evaporates when we switch our focus from count nouns to mass nouns. Consider some particular gallon of water contained in a jug. Even after the inner constitution of the water is rearranged – swirl it, blow bubbles into it, pour it from container to container – we are still left with the *same* water. Similarly, I can rearrange the furniture in my living room and it will still be the *same* furniture. Or, to take one last example, I can neatly arrange all my spare change into layers of quarters, dimes, nickels, and then finally pennies, but it will still be the *same* change. The inner constitution of the referent of a mass noun can be rearranged, and this transformation will map the entire collection back onto itself. There is nothing mysterious or paradoxical about the fact that the referents of mass nouns exhibit permutation symmetry in this fashion. This immediately follows from the distinguishing characteristic of mass nouns: namely, that any quantity whatsoever (of the referent of a mass noun) is treated as an undifferentiated unit.

There are of course pragmatic constraints on the claim that the referents of mass nouns exhibit permutation symmetry among their constituents. The crucial detail lies in how we choose to individuate among overall states of the stuff in question. If I am severely dehydrated, then blowing bubbles into the water in the jug in front of me does not alter the state of the water; it will slake my thirst all the same. However, if instead I am properly hydrated and interested in measuring the temperature distribution across the water, blowing bubbles will indeed alter its state. Similarly, if I am only concerned with how much money I have saved up, then neatly arranging all of my spare change will not yield a new state for me; but if I have a compulsion for organization, this would make all the difference in the world. But notice how in the case of indiscernible electrons there *cannot* be a pragmatic constraint that is tighter than a description of the overall state of the system. This last point is *crucial*. Any information concerning what each individual electron is doing leaves no physical trace whatsoever. This follows from the facts that (1) such a system is brilliantly described by the Fermi-Dirac distribution and that (2) said distribution has built into it the assumption that permuting two particles yields one and the same physical state.

Therefore, if the case can be made for ‘electron’ to be construed as a mass noun, rather than a count noun, the puzzle of indiscernible particles will cease to be a puzzle altogether. Quine’s suggestion from §2.3 – that we ought to abandon the notion of *physical object* and instead adopt a field theory which quantifies over regions of spacetime – will go by the wayside. *There will no longer be any pressure to abandon the notion of physical object*. For we can still think of *electron* in terms of a physical object,

except now on a par with water and sugar, rather than apples and bananas.¹⁷ I have found that speaking of “electron-stuff” rather than “electrons” helps to combat the way in which the count-use of ‘electron’ has become entrenched in our language.

I argued in §3.1 that many paradigmatic count nouns (such as ‘apple’) admit of a mass-use, and that many paradigmatic mass nouns (such as ‘water’) admit of a count-use. We have also seen that it follows from Quine’s Indeterminacy of Translation thesis that, strictly speaking, there can be no fact of the matter as to whether a noun is being used countwise or masswise in any given utterance; adjustments elsewhere in the “translation manual” can always accommodate either reading. So at the very least there is nothing heretical about explicitly proposing a mass-use for ‘electron’ as I wish to do in this section.

What would such a linguistic shift amount to? To illustrate with a simple example, a phrase such as ‘There are two electrons present’ would have to be replaced with something along the lines of ‘There is $-3.2 \cdot 10^{-19}$ Coulombs worth of electron-stuff present.’¹⁸ The latter of these two sentences allows us to speak of a physical object (a $-3.2 \cdot 10^{-19}$ Coulomb chunk of electron-stuff) while circumventing any difficulties raised by the phenomenon of permutation symmetry. But aside from the benefit of retaining the notion of *physical object*, this shift in language actually seems *natural* given the observed behavior of electron-stuff, as I will now try to make clear.

¹⁷ To avoid any confusion: here I stipulate that I am using ‘water’ and ‘sugar’ masswise, and ‘apple’ and ‘banana’ countwise.

¹⁸ $-1.6 \cdot 10^{-19}$ Coulombs being the charge of a single electron.

In order to employ a term as a count noun, we must be able to grasp that term's criterion of individuation. If such a criterion is not forthcoming, then we can be sure that the noun is being used masswise:

To master the term 'apple' [in its capacity as a count noun] it is not enough to be able to decide whether a portion of spacetime is taken up with apples; we must also be able to determine where one apple leaves off and an adjacent one sets in. We must be able to distinguish between pointing to two apples and pointing to different places on the same apple...Thus, whereas the refrain that may be imagined to accompany the ostensive teaching of the term 'water' is 'Here is water', the refrain that accompanies the ostensive teaching of the term 'apple' should not be thought of as 'Here is an apple'; in principle it should be thought of as 'Here is the same apple as here', in accompaniment with two pointings (Quine 1972, pp. 6-7).

But if count nouns require a criterion of individuation, then it is emphatically not the case that we are in the business of individuating quantum particles (in particular, electrons). The atomic orbital model likens the arrangement of electrons around a nucleus to a cloud. Rather than being localized at particular spots, it is more apt to say that electrons are smeared throughout a region of space in the vicinity of the nucleus, with the density of the electron-cloud indicating the probability of finding an electron in that particular region if a measurement is made. Given this model, it is more or less meaningless to ask where "one electron takes off and an adjacent one sets in." Quantum particles are simply not demarcated in the unambiguous way that apples and dogs are.

Thus the linguistic adjustment I am advocating seems to be almost necessitated by the hard facts of quantum mechanics. When talk of entities at a certain level of resolution ceases to be useful (as is the case with individual water molecules when the task at hand is, say, to retrieve a bucket of water from the well), or better yet when such talk ceases to be *possible* (as is the case with individual electrons), we are permitted to ascend to a discourse with lower resolution, one which chops the world into relatively larger and

more manageable chunks: gallons of water, clouds of electron-stuff, etc. Once this linguistic shift is embraced, PIP ceases to be a puzzle altogether, and there is no need to part with the age-old concept of *physical object*.

In fact, this is precisely the course of action prescribed by Quine's maxim of the "Identification of Indiscernibles." This is the principle which states that objects which a theorist does not care to distinguish from one another (i.e., objects which are "indiscernible" for the purposes of the theorist) ought to be lumped together into an equivalence class (i.e., "identified" with each other). In Quine's own words: "references to [indistinguishable] objects should be reconstrued for the purposes of the discourse as referring to other and fewer objects, in such a way that indistinguishable originals give way each to the same new object" (1953, p. 71). Roughly put, this maxim is a recipe for *conceptual integration*. For example, consider trying to construct an economic theory out of a discourse which quantifies over persons. Suppose further that the age, height, gender, race, etc. of each person has no bearing whatsoever on the economic theory in question – all that matters is the person's *income*. It would make sense then to lump persons into income brackets, and thereby reconfigure the theory's ontology so that it only contained income brackets (and not individual persons). As a result, "[d]istinctions immaterial to the discourse at hand are thus extruded from the subject matter" (ibid.). So in the same way that we are afforded brevity and ontological tidiness when we switch from talk of persons to talk of income brackets, I am suggesting that we will be afforded economy (and, moreover, intelligibility!) by switching from talk of individual electrons to talk of electron-stuff.

3.3 *Objections Considered*

In this section I will address what I take to be the two most pressing objections to the position I am advocating. The first objection proceeds by pointing out the myriad physical situations in which reference must be made to *single* electrons. Consider the Photoelectric Effect: an incoming photon collides with some surface (a metallic plate, say), transfers its energy to one of the electrons on the surface, liberates the electron from its atomic binding, and sends it flying off into space. This is just one of many types of scattering phenomena which inevitably make reference to *single* electrons, and therefore present an obstacle to the account I am proposing. If ‘electron’ is to be construed as a mass noun, then talk of single electrons seems problematic since in the phrase ‘single electron’ we seem to be using ‘electron’ countwise.

However, I see nothing paradoxical or contradictory about treating ‘electron’ as a count noun some of the time, and a mass noun the rest of the time. The main point I argued for in §3.1 was that any term (such as ‘electron’) is neither inherently a mass noun nor a count noun; rather, it is individual *instances* of that term which appear as either mass nouns or count nouns. Both usages are viable, and we may choose one over the other so as to make the most sense out of the situation at hand. Furthermore, the Indeterminacy of Translation thesis assures us that there will always be room for adjustments in the portion of “physicist-talk” already translated that will permit either use of ‘electron.’ (And if all that is still not satisfying, then we can just paraphrase ‘single electron’ as ‘ -1.6×10^{-19} Coulomb chunk of electron stuff’ and thereby remain in the realm of mass nouns all the time.)

Thus the pluralism entailed by wholesale instrumentalism (see section §1.2) begins to play a key role in my argument. My proposal is analogous to the way in which we freely oscillate between thinking of quantum objects as waves on the one hand and particles on the other. In certain cases (such as Young's double-slit experiment) it is most sensible to attribute wave-like properties to light, while in others (such as the Photoelectric Effect) it is most sensible to attribute particle-like properties to light. Neither model could by itself explain the full range of light-related phenomena that we are currently able to explain by helping ourselves to *both* models (though never at the same time, of course). Similarly, I am proposing that we help ourselves to both the mass-use and the count-use of 'electron' in order to bolster our explanatory arsenal. Just as we are forced to adopt the pluralistic viewpoint concerning light, we ought to feel free to adopt the pluralistic viewpoint in dealing with indiscernible particles.

The second objection I wish to consider points out a crucial difference between uncontroversial mass nouns (such as 'water') and the lately suggested mass noun ('electron' or 'electron-stuff'). The referents of *true* mass nouns, the objection says, may take on a continuous range of quantities: we can have a gallon of water, two gallons of water, and any amount in between; in other words, the referents of mass nouns come in *continuous* quantities. By way of contrast, the referents of count nouns come in discrete quantities: we can have one dog, two dogs, but not any intermediate number of dogs. The point is that so-called "electron-stuff" seems to come in discrete packets, each containing -1.6×10^{-19} Coulombs worth of electric charge. Thus it begins to seem that 'electron-stuff' is actually a *count* noun, every bit as much as the traditional 'electron' is, and therefore my proposed solution hasn't even gotten off the ground.

While I admit that electron-stuff does come prepackaged in discrete amounts, and that there is a minimum amount of electron-stuff that can exist, I point out that the referents of familiar, uncontroversial mass nouns function in *exactly the same way*. Consider water. It is only a naïve misconception that any amount of water between 1 and 2 gallons is obtainable. Taking V_{\min} to be the minimum volume of water physically possible (i.e., the volume of a single H_2O molecule), we can only obtain quantities of water which are multiples of V_{\min} .¹⁹ Similarly for sugar and its own unique V_{\min} . In other words, if we were to take inventory of all physically possible amounts of water between 1 and 2 gallons, then these quantities would only *sparingly* populate the segment of the number line between 1 and 2. Were we to take seriously the claim that the referents of mass nouns must come in *continuous* quantities, mass nouns simply would not refer to anything in the real world. Therefore we would do best to drop the continuity constraint altogether.

Now there *is* a crucial difference between water and sugar on the one hand and electron-stuff on the other: namely, the difference in size between their respective V_{\min} 's and the amounts that have been recently considered. Everyday amounts of water and sugar (e.g. bottles of water and spoonfuls of sugar) are orders of magnitude above their minima. However, the amounts of electron-stuff appearing in the examples in this paper (e.g., -3.2×10^{-19} Coulombs worth) are on the same order of magnitude with their minima. But notice this is merely a difference of degree, and does not therefore distinguish 'water' and 'electron-stuff' as grammatically different in any way. If lakes and pastries were much smaller we would probably be accustomed to purchasing water and sugar by the

¹⁹ This is of course an exceptionally rough estimation. To be more precise we'd have to take into account spacing between molecules and the dissociation of H_2O molecules, but none of that would affect the point I am trying to make.

molecule (instead of by the gallon and by the pound), at which point the grammatical similarities between ‘water,’ ‘sugar,’ and ‘electron-stuff’ would be undeniable.

§4 CONCLUDING REMARKS

I think it will be helpful at this point to retrace the logical space that this paper has traversed. Empirical findings have shown us that all elementary particles of a given kind exhibit permutation symmetry. We know this on the basis that, when measured, collections of such particles unequivocally conform to the Fermi-Dirac/Bose-Einstein statistics – i.e., the statistical distributions for which swapping two particles does *not* yield a new state. This is troubling because ordinary physical objects *do* yield a new overall state when swapped. One suggested solution to this puzzle (Quine’s) would have us do away with physical objects altogether, and instead work in an ontology which only contained regions of spacetime (and, eventually, we might even reduce those regions of spacetime to mere quadruples of numbers). This clearly solves the puzzle (without any particles, there can be no puzzle *about* particles), but at a great cost. I have proposed an alternative solution – a new intellectual tool, so to speak – which preserves the notion of physical object. The key to this solution lies in pointing out that PIP only creates a problem for the referents of count nouns; the referents of mass nouns are unfazed by PIP. My arguments on this front were launched from the perspective of wholesale instrumentalism, from which an attitude of pluralism naturally followed. From this perspective, questions like “Does electron-stuff *really* exist?” are not worthy of serious attention. Electron-stuff is a theoretical posit which solves a very specific problem, and is therefore just a “real” as any other theoretical posit we might be obliged to make.

The impetus to my solution was therefore to minimize ontological revision, and to perpetuate a discourse which makes reference to physical objects composed of ordinary matter. In this respect my solution attempts to meet Quine on his own terms, insofar as he believes that in choosing a starting point for empirical inquiry “we can choose no better than the selfsame world theory which we are trying to improve, this being the best available at the time” (1966, p. 241). This prescription more or less follows from Quine’s echoing of Otto Neurath’s likening of inquirers to sailors who must rebuild their ship while at sea, plank by plank, without returning to the dock (Quine 1960, p. vii). To extirpate physical objects from scientific theorizing would, on this picture, be tantamount to removing a very large plank from our ship, which could quite possibly sink us once and for all.

It is easy to lose grasp of the meaning of ‘physical object’ in pondering these issues, and therefore important to ask precisely what is meant by ‘physical object’ in the claim that it is viable to think of electron-stuff as a physical object. Here I think the answer is deceptively simple. In saying that some particular lump of electron-stuff is a physical object, we mean precisely the same as when we say that this lump of Carbon is a physical object, or that this cloud of Oxygen is a physical object. To be a physical object is simply to be composed of matter. In the same way that “any sum of parts which are water is water,” any sum of parts which are physical objects is also a physical object. It’s matter all the way down, so to speak. Therefore if we are to retain the commonsensical claim that lumps of Carbon, clouds of Oxygen, and jugs of water (and so on) are all composed of matter, then we must be able to understand some way in which these everyday-sized physical objects could be composed of physical objects at their most

fundamental and microscopic level. My efforts in this paper have been solely directed at constructing such an account.

WORKS CITED

- Castellani, Elena (1998), *Interpreting Bodies*. Princeton University Press.
- Feynman, Richard (1965), *The Character of Physical Law*. MIT Press (1989 reprint).
- Fine, Arthur (1986), *The Shaky Game*. Chicago: University of Chicago Press.
- French, S. & Krause, D. (2006), *Identity in Physics*. Oxford University Press.
- Griffiths, David J. (2005), *Introduction to Quantum Mechanics*. New Jersey: Pearson Education, Inc.
- Huggett, Nick & Norton, Josh (forthcoming), “Weak Discernibility For Quanta, The Right Way.” *The British Journal For The Philosophy of Science*.
- James, William (1907), *Pragmatism & The Meaning of Truth*. [1975 reprint, Cambridge: Harvard University Press].
- Lewis, David (1986), *On the Plurality of Worlds*. Blackwell Publishing.
- Muller, F.A. & Saunders, S. (2008), “Discerning Fermions.” *The British Journal For The Philosophy of Science*, **59**.
- Nagel, Ernst (1961), *The Structure of Science*. Harcourt, Brace & World Inc.
- Quine, W.V. (1953), *From a Logical Point of View*. Cambridge: Harvard University Press.
- (1960), *Word and Object*. Cambridge: MIT Press.
- (1966), *The Ways of Paradox*. Random House, Inc.
- (1969), *Ontological Relativity and Other Essays*. NY: Columbia University Press.
- (1972), “Identity.” Unpublished lecture (typescript).
- (1976), “Whither Physical Objects?” [in *Essays in Memory of Imre Lakatos*, pp. 497-504. R.S. Cohen, P.K. Feyerabend, and M.W. Wartofsky (eds.), Springer Publishing.]
- (1990), *Pursuit of Truth*. Harvard University Press.
- Schwartz, Robert (2012), *Rethinking Pragmatism*. Wiley-Blackwell.
- Stapledon, Olaf (1931), *Last and First Men & Star Maker*. [citation comes from p. 162 of Dover’s 1968 reprint].
- Teller, Paul (1997), *An Interperative Introduction to Quantum Field Theory*. Princeton University Press.
- (1998), “Quantum Statistics and Haecceities” [in Castellani 1998].
- Wittgenstein, Ludwig (1921), *Tractatus Logico-Philosophicus*. [reprinted in (2009) *Major Works*, pp. 1-84. New York: Harper Perennial.]
- (1953), *Philosophical Investigations*. G.E.M. Anscombe (trans.), Basil Blackwell.