Identical Quantum Particles and Weak Discernibility

Dennis Dieks and Marijn Versteegh
Institute for History and Foundations of Science
Utrecht University, P.O.Box 80.000
3508 TA Utrecht, The Netherlands

Abstract

Saunders has recently claimed that "identical quantum particles" with an anti-symmetric state (fermions) are *weakly discernible objects*, just like irreflexively related ordinary objects in situations with perfect symmetry (Black's spheres, for example). Weakly discernible objects have all their qualitative properties in common but nevertheless differ from each other by virtue of (a generalized version of) Leibniz's principle, since they stand in relations an entity cannot have to itself. This notion of weak discernibility has been criticized as question begging, but we defend and accept it for classical cases likes Black's spheres. We argue, however, that the quantum mechanical case is different. Here the application of the notion of weak discernibility indeed *is* question begging and in conflict with standard interpretational ideas. We conclude that the introduction of the conceptual resource of weak discernibility does not change the interpretational *status quo* in quantum mechanics.

1 Introduction

Both in ordinary life and in science we are accustomed to thinking in terms of *objects*, i.e. *individuals*, which differ from each other and can bear their own names. A possible point of view is that the individuality involved here should be seen as primitive, or as grounded in a metaphysical principle like haecceity or in an underlying Lockean substance. In this paper, however, we take an empiricist stance. We are interested in a notion of individuality that can be grounded empirically: we take as our starting point that at least in science entities should differ in one or more of their qualitative features, as described by the relevant scientific theory, in order to be treated as individuals. This is in the spirit of Leibniz's principle of the identity of indiscernibles (PII). Classical physics is in accordance with this empiricist motif: if PII is formulated in a sufficiently general way (including the notion of weak discernibility, needed to deal with symmetrical configurations; see below for a definition of weak discernibility), it is satisfied by the objects recognized by classical physical theory [4, Ch. 2].

In quantum theory the notion of individuality is notoriously problematic. It has recently been claimed by Saunders [10, 11], however, that quantum particles with an antisymmetric wavefunction (fermions) obey PII in the same way as classical objects in symmetrical configurations: according to Saunders, fermions are weakly discernible individuals. We oppose this conclusion. As we shall argue, the situation in quantum mechanics is essentially different from classical physics. In a nutshell, our argument is that in the case of fermion systems the assumption that there are several objects is dubious to begin with, and that the introduction of the notion of weak discernibility does not help. PII can certainly be maintained in quantum theory, but the ontology that fits the Principle is field-theoretic rather than one of many individual particles.

2 Weakly Discernible Objects

If we want to give Leibniz's principle a fundamental role we must evidently consider the putative counterexamples that have been proposed to it. These concern cases in which there appear to be indiscernible but nevertheless different individuals. Think of Max Black's spheres [2], of identical chemical composition and two miles apart in a relational space (à la Leibniz, not in Newtonian absolute space where absolute positions could label the spheres); Kant's enantiomorphic hands; or, for a mathematical example, the points in the Euclidean plane. In all these cases the objects have all their qualitative features in common: both spheres in Black's example have equal material properties and both are at two miles from a sphere; similarly, Kant's hands have the same internal geometric properties and both are mirror images of another hand; and so on. So these cases appear to demonstrate that we ordinarily employ concepts of object and individuality that are independent of the presence of distinguishing qualitative differences—in violation of PII.

As Hawley points out [6], defenders of PII can respond to such examples in a variety of ways. First, they may query whether the situations figuring in the examples are possible at all. If they are willing to concede this possibility, they can dispute that these situations are best described in terms of distinct but indiscernible individuals: they may either argue that if a correct analysis of discernibility is employed the objects are discernible after all, or they may claim that there were no distinct objects to start with.

In the just-mentioned examples the best response—so we shall argue—for friends of PII is to say that these are possible situations in which more than one objects are present; and that these objects are *not* indiscernible. Let us pay attention to the details of the discernibility issue first. We follow Saunders [10, 11], who takes his clue from Quine [8], in noting that in cases like those in the examples *irreflexive* qualitative relations are instantiated: relations entities cannot have to themselves. This irreflexivity is the key to proving that (a generalized version of) PII is satisfied after all: if an entity stands in a relation that it cannot have to itself, there must be at least two entities.

To see in more detail how this works, let us formalize a bit. In first-order formal languages we can define identity (=) as follows:

$$s = t \equiv P(s) \leftrightarrow P(t), \tag{1}$$

where P denotes an arbitrary predicate in the language, and the right-hand side of the definition stipulates that s and t can replace each other, *salva veritate*, in *any* P. This definition captures PII, and our empiricist notion of individuality, if the language is that of a physical theory, in which the predicates refer to physical properties and relations (and not to haecceities).

There can now be various kinds of discernibility ([11]). Two objects are *absolutely discernible* if there is a one-place predicate that applies to only one of them; *relatively discernible* if there is a two-place predicate that applies to them in only one order; and *weakly discernible* if an *irreflexive* two-place predicate relates them. The latter possibility is relevant to our examples. If there is an irreflexive but symmetric two-place predicate $P(.,.)$ that is satisfied by s and t, the definition (1) requires that if s and t are to be identical, we must have:

$$\forall x(P(s,x) \leftrightarrow P(t,x)). \quad (2)$$

But this is false: in any valuation in which $P(s,t)$ is true, $P(t,t)$ cannot be satisfied by virtue of the fact that P is irreflexive. It follows therefore that PII is satisfied by any two individual objects that stand in an irreflexive qualitative relation. What is needed to dispel the impression that PII is in trouble in these cases is to apply the principle not only to monadic predicates, but to all qualitative n-ary predicates (and thus to all available qualitative relations).

3 Begging the Question?

The just-sketched proof used the notion of a *valuation*. A valuation results from letting the names and bound variables in the formulas of the language refer to specific elements of the intended domain. In order to construct a valuation, we therefore have to name and distinguish the things we are talking about. However, this is an impossible task in the symmetrical configurations of the examples. Because of the symmetry we cannot uniquely refer and assign names on the basis of the given structure of properties and relations. It is impossible, for example, to single out any specific point in the Euclidean plane on the basis of the properties of the plane and its points, even if we include all relational properties. This impossibility may seem to take the edge off the above argument for the validity of PII; and therefore to jeopardize the Leibniz-style individuality of our entities after all (because individuals must be able to bear names, whereas no names can be given here on the basis of qualitative differences) [7].

But in this form the difficulty is only apparent. In order for the notions of number and names to apply to the members of a domain it is sufficient that a function *exists* that maps the domain one-to-one onto a set of labels, e.g. the set $\{1, 2, ..., n\}$ [13, p. 457]; it is not needed that we can actually *construct* such a labelling. In the examples we have been considering it was given in the description of the cases (two spheres, two hands, many points) that the required mappings exist—and this was all we needed.

Although the name-giving problem is thus dissolved, its discussion has highlighted an important issue: in the proof we needed that the domains possessed a structure underwriting the existence of bijections to sets of labels (e.g., sets of one or more natural numbers). But the existence of such a structure is obviously far from evident in the context of a discussion of the applicability of PII. As we have already seen, one of the possible responses to putative counterexamples to PII is to argue that there is no multitude of systems at all: that there is rather *one* undivided physical system. This response does not need any appeal to weak discernibility, and moreover it is ontologically parsimonious—in a way it is the simplest option. In other words, we must have good reasons to assume the existence of separate components in the first place. But how can we have such reasons without getting into a circular argument? Our starting point was the empiricist desire to ground individuality in empirically accessible qualitative features, by means of PII; but what is the use of this exercise if we must already know about the existence of individual objects before we can even apply PII? It is therefore urgent to see how we may justify the presence of an "object structure" in an empiricist way, and how it may be possible to apply the notion of weak discernibility without begging the question of individuality.

4 Relata and Relations

If it were indeed necessary to make sure that the domains we are discussing consist of distinct objects *prior* to any application of PII, this would be self-defeating for the empiricist enterprise. Why should we engage in any Leibnizean arguments if we already *know* that there are more than one objects? More importantly, it seems clear that prior to using PII we can base our judgement that there are different objects only on non-empirical grounds, involving primitive thisness, haecceity, substance or some similar metaphysical principle. As French and Krause comment in their discussion of weak discernibility [4, pp. 170-171]:

> "Doesn't the appeal to irreflexive relations in order to ground the individuality of the objects which bear such relations involve a circularity? In other words, the worry is that in order to appeal to such relations, one has already had to individuate the particles which are so related and the numerical diversity of the particles has been presupposed by the relation which hence cannot account for it".

But there is a way out here, offered by the structuralist proposal according to which the numerical diversity of the relata need not be prior to the existence of the relations ([4, p. 172]). According to this structuralist approach relata can be *determined by relations*, as a kind of nodes in a relational network. This does not necessarily mean that the relations are ontologically prior to the relata. Relata and relations may be on a par, ontologically speaking: the relata may have no other properties than specified by the relations, whereas the relations can only exist if they connect relata [3]. This position fits the empiricist outlook of PII very well if the relations are understood as the qualitative relations occurring in scientific theories. From an empiricist point of view, adopting the structuralist position appears the natural way to defuse the circularity threat.

So in our earlier examples we need not assume that there is a division of labor between "objecthood providers" and the relations that are defined on the domain. We can accept the relational structure as the only access to Leibnizean objecthood, i.e. objecthood with a qualitative grounding. The relations in our two physical examples—being at a spatial distance from each other, being each other's mirror image—indeed give us information about the presence of objects. These relations define relata that can be displaced with respect to each other, or that can be reflected and whose orientations can be compared. Such relata are actual things, objects, that differ from each other because of their mutual distance or their mirror-image relation, respectively. If irreflexive relations of this type apply, there must be more than one objects as relata.

As we shall illustrate in a moment, however, the mere possibility of speaking about a domain in terms of irreflexive relations is not enough to ensure such an objecthood structure. It is sometimes usual to employ properties or relations talk even in situations in which there are no different objects at all—in such cases considerations about the irreflexivity of the relations obviously cannot do anything to show that the "objects" are in fact weakly discernible. This means that even within the context of a structuralist outlook it has to be decided whether the concept of object is applicable in the first place, before arguments about weak discernibility based on irreflexivity of relations can make sense.

5 Relations without Objects

To see the possibility of relations without actual relata, consider the example of Euros in a bank account (not coins in a piggy bank, but transferable money in a real bank account). Imagine a situation in which by virtue of some financial regulation the Euros in a particular account can only be transferred to different one-Euro accounts. So, in a complete money transfer an account with five Euros, say, would be emptied and five different one-Euro accounts would result. In this case the Euros in the original account stand in an irreflexive relation to each other, namely "only transferable to different accounts". But this does not make them into different physical objects. We could try to exploit the irreflexive relations for the purpose of distinguishing individual Euros; e.g., we may attempt to label the Euros by means of the accounts they will end up in. However, looking at the situation *after* the money transfer clearly does not achieve anything for the purpose of distinguishing between the Euros in the account they are actually in, before the transfer. The essential point is that the relations here do not relate occurrent, actual, physical features of the situation; they do not connect actual physical relata. Rather, the relational structure is defined with respect to what *would* result if the actual situation *were changed*. There is of course no doubt that five different one-Euro accounts (with different account numbers) are five individuals; but this does not mean that it makes sense to consider the five Euros as individual objects when they are still in one common account. On the contrary, the case of more than one money units in one bank account is the standard example to illustrate absence of individuality; it is a case in which only the account itself, with the total amount of money in it, can be treated as possessing individuality [9, 12]. Although we are accustomed to using relations and things *talk* here, there is nothing in the actual physical situation that directly corresponds to this. Speaking of several Euros in a bank account is a *façon de parler*. According to our best understanding of the situation, statements like "all Euros in this account have the same value—namely one Euro" are not about individual things.

Thus, we have found an important silent presupposition in the argument for PII-based individuality in the presence of irreflexive relations. Such relations can only be trusted to be significant for the individuality issue if they are of the sort to connect actual relata.

6 Actuality

Whether relations occurring in scientific theories connect and determine actual relata is decided by their meaning, reflected by the role they play within the theory in question. In cases in classical physics without particular symmetries, a feature of such relations is that they can be used to distinguish and name the different relata. For example, in an arbitrary configuration of classical particles the mutual distances will unambiguously characterize each individual particle. Changing the configuration so that it becomes more symmetrical will change the values of the distances, but not the actuality of the objects: the description of the situation furnished by classical theories will not change as far as the actuality of the particles is concerned, when we approach a Blackean spheres-type configuration. The discernibility in asymmetrical situations thus provides us with a test for the actuality in symmetrical situations: we are justified in assuming the existence of different actual entities if breaking of the symmetry does not involve a change in the type of description given, and results in a situation with distinguishable objects (this strategy resembles the one followed by Adams, who proposes to compare Black's spheres with spheres of which one has a very slight chemical impurity [1, p. 17]). This breaking of the symmetry is analogous to the introduction of a coordinates origin in describing a figure in plane geometry. If a mapping to natural numbers exists in the presence of such a standard, our theoretical accounts will still be able to work with such a mapping when the reference point has been removed; what changes is merely the *constructibility* of the mapping. The following mathematical case provides another example. The numbers 1 and -1 share all their structural properties in the structure $< Z, + >$ (the *relational structure* of integers with addition, without individual names for the numbers), just like the spheres in our earlier example. Nevertheless, they are not identical. An indication of this is that the introduction of a standard, e.g. for being positive, makes it possible to distinguish and name these two numbers. This possibility of naming disappears when we forget about the standard—however, this does not collapse the two numbers into one. They still are the kind of entities that have, e.g., an actual numerical distance to other entities in the structure—and to each other!—even though the symmetry makes it impossible to assign names ([4, p. 265]).

Indeed, why do we feel so sure that there are two Blackean spheres and two Kantian hands? It seems obvious that this is because our mind's eye sees Black's spheres at different distances, and Kant's hands with different orientations, before us; when we break the symmetry of the configurations in these cases by introducing a point of reference, a gauge or standard, the relations make it possible to distinguish the entities with respect to this standard. Thus we can name Black's spheres via their unequal distances to a reference point and in Kant's universe we may imagine a reference hand conventionally called "left". Another example, relevant for our subsequent discussion of quantum objects, is furnished by two oppositely directed arrows in an otherwise empty Leibnizean world. If we fix a standard of being "up", we break the symmetry and the individual arrows become absolutely discernible as up or down.

These cases are to be contrasted with the case of the Euros in a bank account. Even though that situation is not symmetrical (the Euros have different destinations) this cannot be used to distinguish actual Euros in the account, according to our best available way of theoretically describing and explaining the situation.

7 The Quantum Case

A notorious interpretational problem of quantum theory concerns so-called identical particles. These are of the same kind, i.e. "they" possess the same intrinsic properties (like mass, charge, spin); e.g., electrons, protons or neutrons. It is a principle of quantum mechanics that the state of a collection of such particles is completely symmetric (in the case of bosons) or anti-symmetric (fermions). This symmetrization postulate implies that all one-particle states occur symmetrically in the total state of a collection of identical particles. It follows that any property or relation that can be attributed, on the basis of the total quantum state, to any one particle is attributable to all others as well.

The standard response is to say that identical quantum particles "lack individuality". This is an awkward way of saying that there are no different particles at all: although there is traditional talk of "many of them", this should be understood in the same way as talk about many Euros in a bank account. In principle it is better, according to this received view, to renounce talk that suggests the existence of individual particles and to reconceptualize the situation in terms of the excited states of a field (analogous to thinking of the Euros in an account as one sum of money).

However, the situation is also reminiscent of the symmetric configurations of classical objects described in the previous section. As we have seen there, symmetry is not at all decisive for proving the absence of Leibniz-style individuality: particles may well be weakly discernible individuals. Could it not be that in the quantum case there are irreflexive physical relations between particles that guarantee their individuality in the same way as they did for Black's spheres, Kant's hands and Euclid's points? This is the position adopted by Saunders, at least for the case of fermions [11]. Indeed, the anti-symmetry of the state of many-fermions systems seems to imply the existence of irreflexive relations between components of the total system: intuitively speaking, the fermions in any pair stand in the relation of "occupying different one-particle states", even though the particles do not receive individually different quantum mechanical state descriptions. It is true (and noted by Saunders) that for bosons with their symmetrical states this manoeuvre is not available, so that collections of identical bosons are still best understood as one whole. But the conclusion that standard quantum mechanics entails that fermions (these are the ordinary matter particles; bosons are quanta of interaction fields) are ordinary individuals—although only weakly discernible—is surprising and highly significant by itself.

The technical details of the argument can be illustrated by the example of two fermions in the singlet state. If $|\uparrow\rangle$ and $|\downarrow\rangle$ stand for states with spins directed upwards and downwards in a particular direction, respectively, the anti-symmetrization principle requires that a typical two-fermion state looks like

$$\frac{1}{\sqrt{2}}\{|\uparrow\rangle_1|\downarrow\rangle_2 - |\downarrow\rangle_1|\uparrow\rangle_2\}, \tag{3}$$

in which the subscripts 1 and 2 refer to the one-particle state-spaces of which the total state-space (a Hilbert space) is the tensor product. These two one-particle state-spaces are the available candidates for the description of single particles and their labels are candidate names for the individual fermions (for the moment, we are assuming hypothetically that the notion of different individuals makes sense). Now, the anti-symmetry of the total state implies that the state restricted to state-space 1 is the same as the restricted state defined in state-space 2. (The "partial traces" are $\frac{1}{2}\{|\uparrow\rangle\langle\uparrow| + |\downarrow\rangle\langle\downarrow|\}$ in both cases.) The *total spin* has the definite value 0 in state (3); that is, state (3) is an eigenstate of the operator $S_1 \otimes I + I \otimes S_2$. Therefore, it seems natural to say that the two spins are *oppositely directed*. On the other hand, we cannot assign a definite spin direction to the single particles because the up and down states occur symmetrically in each of the states in the Hilbert spaces

1 and 2, respectively.

This situation appears essentially the same as the one of the two arrows. In that case, it was not possible to designate one of the arrows as up and the other as down—but nevertheless there had to be two individual arrows in view of the oppositeness of their directions. Similarly, in the fermion case with total spin zero we apparently are dealing with two individual quantum objects with opposite spins.

8 Quantum Individuals?

On closer examination the similarity starts to fade away, however. One should already become wary by the observation that the irreflexive relations in the quantum case have a theoretical representation that is quite different from that of their classical counterparts. There the relations could be formalized by ordinary predicates that can be expressed as *functions* of occurrent properties of the individual objects (like "up" and "down" with respect to a conventionally chosen standard, or $+1$ and -1), with the correlation expressed by the fact that the sum of these two quantities has a fixed value. By contrast, in quantum theory the correlation is expressed in the following way: the state of the total system is an *eigenstate* of a linear *operator* in the total system's Hilbert space. Concomitant with this formal difference is an essential difference in interpretation: according to standard interpretational ideas[1], quantum states should be interpreted in terms of possible *measurement results* and their probabilities, rather than in terms of occurrent properties. In the case at hand, a system in an eigenstate of the total spin operator with eigenvalue 0, this means that a measurement of the total spin will with probability 1 have the outcome 0. In this special case, in which the outcome is certain (probability 1), it is—according to the same standard interpretational ideas—harmless to assume that the total system possesses the property "total spin 0" also independently of measurement; but even so, this total spin cannot be understood as being composed of definite individual spin values of the two subsystems. Although it is of course

[1] Certainly, there are also interpretations that interpret the quantum state in terms of occurrent physical properties, even if no measurements are performed. Examples are the Bohm interpretation and modal interpretations. These interpretations treat the fermions case in terms of component systems with different qualitative properties, so that they do not need the notion of weak discernibility.

possible to perform individual spin *measurements* on the subsystems (whose possible outcomes and corresponding probabilities are predictable from the total quantum state), there are no corresponding *occurrent spin properties* in the subsystems, independently of measurement. In the singlet state (3) the prediction of quantum mechanics is that individual spin *measurements* will with certainty yield opposite results, summing up to 0; but on the pain of running into paradoxes and no-go theorems it cannot be maintained that these results reveal oppositely directed spins that were already there before the measurements. This is an example of the notorious "holism" of quantum mechanics: definite properties of a composite system do not supervene on properties of its parts.

This suggests that the correct analogue to the quantum case is not provided by two oppositely directed classical arrows, but rather by a two-Euro account that can be *transformed* (upon "measurement", i.e. the *intervention* brought about by a money transfer) into two distinct one-Euro-accounts.

To investigate further whether or not the quantum relations connect actual physical systems, we can copy the strategy followed in the classical case, namely breaking the symmetry and seeing whether in the resulting situation the quantum relations can serve as name-givers. However, this cannot work as long as we stay within a many-fermions system: quantum mechanics forbids fermion systems that are not in an anti-symmetric state—it is a matter of lawlike principle that the only relations fermions can possess with respect to each other are perfectly symmetrical. This is a significant difference from the symmetrical classical cases, where the symmetry was contingent and where, e.g., the theory allowed evolutions from symmetrical to asymmetrical configurations. In quantum mechanics the mutual relations between fermions cannot serve to distinguish individual component systems *as a matter of principle*. The theory does not allow any asymmetrical situations with which to approach the symmetrical situation, and our earlier test fails.

It is true that this does not prove that there are no individual fermions—compare with the situation in a hypothetical world in which *laws* stipulate that classical spheres can only occur in symmetric configurations. In such a world we could still have good reasons to think in terms of individual spheres, because our theories could allow for an external object, serving as a point of reference that makes the spheres discernible. Our next attempt is therefore to break the symmetry in the fermion case by the introduction of a standard that is external to the fermion system itself. Quantum mechanics does not require symmetry of the total state of a fermion system plus something else, so with an external standard in hand we may hope to be able to distinguish individual fermions. This would support the idea that the quantum relations are of the kind that connect particles; it would provide evidence that it makes sense to speak about actual objects at all.

To see the inevitability of a negative outcome of any such test, consider an arbitrary system of identical quantum particles to which a gauge system has been added without any disturbance of the original system (i.e., the total state is the product of the original symmetrical or anti-symmetrical identical particles state and the state of the gauge system). Let the new total state be denoted by $|\Psi\rangle$. Any quantum relation in this state between the gauge system g and one of the identical particles, described in subspace j, say, has the form $\langle\Psi|A(g,j)|\Psi\rangle$. Here $A(g,j)$ is a hermitian operator working in the state-spaces of the gauge system g and identical particle j. We can now use the (anti)-symmetry of the original identical particles state to show that the gauge system stands in exactly the same relations to *all* identical particles. The (anti)-symmetry entails that $P_{ij}|\Psi\rangle = \pm|\Psi\rangle$, where P_{ij} stands for the operator that permutes identical particle indices i and j. Now,

$$\langle\Psi|A(g,j)|\Psi\rangle = \langle P_{ij}\Psi|A(g,j)|P_{ij}\Psi\rangle = \\ \langle\Psi|P_{ij}^{-1}A(g,j)P_{ij}|\Psi\rangle = \langle\Psi|A(g,i)|\Psi\rangle.$$

In other words, any quantum relation the gauge system has to j, it also has to i, for arbitrary values of i and j. That means that these quantum relations have no discriminating value in the situation as it actually is, without measurement interventions and the disturbances caused by them.

It must be stressed that if the situation *is changed* by a measurement interaction, distinct individual results will arise[2], just as in the case of the opposite spin results of measurements in the total spin 0 state (3). But we are here interested in the question of whether a many-fermions system *as it is* can be regarded as a collection of weakly discernible individuals; not in the question of whether such a system can be *transformed* into a collection of individuals. With regard to the first question all available evidence points into one direction: fermions behave like money units in a bank account. It does not matter what external standard we introduce, it will always possess the same relations to all (hypothetically present) entities. This leaves us

[2]The expressions $\langle\Psi|A(g,j)|\Psi\rangle$ are in this case interpreted as *expectation values*, averages over very many experimental trials.

without a good reason to suppose that there *are* any actual objects composing many-fermions systems[3]. This is an essential difference with the earlier examples of classical weakly discernible objects.

9 Conclusion

There is thus an important contrast between quantum mechanical many-fermion systems and classical collections of weakly discernible objects. In the latter case we have reason to think that there are objects that are nameable *in abstracto*, although the symmetry of the situation makes it impossible to actually assign names. Application of the concept of weak discernibility shows that nevertheless there is no conflict with Leibniz's principle. The strangeness of the quantum case runs much deeper. There is no sign within standard quantum mechanics that "identical fermions" are things at all; there is no ground for the supposition that the quantum relations "between fermions" connect any actual physical objects. The irreflexivity of relations does not help us here. Quantum relations have an interpretation not in terms of what is actual, but rather via what *could happen in case of a measurement*. Their irrelevance for the question of whether there are actual objects is illustrated by the fact that they cannot be used in any name-giving procedure, not even after the introduction of an external standard. As far as standard quantum mechanics goes, there *are* no separate individual fermions and the question of whether they are weakly discernible does not even arise. Conventional wisdom, saying that systems of identical quantum particles are best considered as one whole, like an amount of money in a bank account, appears to have it right after all.

Acknowledgment

We thank Simon Saunders for helpful comments on an earlier version.

[3]As Hawley [5] rightly observes, PII itself is in the spirit of a sober empiricism, averse to the introduction of component systems if they do not play an explanatory role within our best theoretical understanding.

References

[1] Adams, R.M. (1979), "Primitive Thisness and Primitive Identity", *Journal of Philosophy* **76**, 5-26.

[2] Black, M. (1952), "The Identity of Indiscernibles", *Mind* **61**, 153-164.

[3] Esfeld, M. and Lam, V. (2006), "Moderate structural realism about space-time", *Synthese*, to appear; DOI 10.1007/s11229-006-9076-2.

[4] French, S. and Krause, D. (2006), *Identity in Physics: A Historical, Philosophical, and Formal Analysis*, Oxford University Press.

[5] Hawley, K. (2006), "Weak discernibility", *Analysis* **66**, 300-303.

[6] Hawley, K. (2007), "Identity, Indiscernibility and Number", unpublished manuscript.

[7] Keränen, J. (2001), "The Identity Problem for Realist Structuralism", *Philosophia Mathematica* **9**, 308-330.

[8] Quine, W.V. (1976), "Grades of Discriminability", *Journal of Philosophy* **73**, 113-116; reprinted in *Theories and Things*, Harvard University Press, Cambridge, 1981.

[9] Schrödinger, E. (1952), *Science and Humanism*, Cambridge University press, partly reprinted as "What is an Elementary Particle?", in Castellani, E. (ed.) (1998), *Interpreting Bodies: Classical and Quantum Objects in Modern Physics*, Princeton University Press, 197-210.

[10] Saunders, S. (2003), "Physics and Leibniz's Principles", in *Symmetries in Physics: Philosophical Reflections*, K. Brading and E. Castellani, eds., Cambridge University Press.

[11] Saunders, S. (2006), "Are quantum particles objects?", *Analysis* **66**, 52-63.

[12] Teller, P. (1998), "Quantum Mechanics and Haecceities", in Castellani, E. (ed.) (1998), *Interpreting Bodies: Classical and Quantum Objects in Modern Physics*, Princeton University Press, 114-141.

[13] Van Fraassen, B. (1991), *Quantum Mechanics—An Empiricist View*, Clarendon Press, Oxford.