Revenge of Ontology by Linking*

Special Extended PDF Edition

Christopher von Bülow†

December 2007

Abstract

This is another, somewhat more successful attempt to clarify (also to myself) what the new approach to ontology I want to propose consists in. I briefly characterize what I call "ontology by founding" – 'foundationalist' analytical ontology – and what I dislike about it, e. g., its lack of concern with epistemological matters (how can we know, say, redness?). Then I try to elucidate the alternative approach I want to propose, an antifoundationalist, Neurathian 'ontology by linking', where basic ontological concepts aren't taken as primitive but rather explained via something like implicit definitions involving also everyday, scientific and teleological concepts. More specifically, I suggest explaining ontological categories by recourse to physical agents, their goals/functions and how they recognize features of the world: their ways of carving chunks out of nature are the predecessors of our ontological concepts. In a digression, I argue against the common linguistic approach to conceptual questions. I prefer to analyze concepts by investigating what their functions in the information-processing of agents are, especially which sorts of phenomena (under which circumstances) constitute their intended domains of application. This is to be done by successively answering three types of questions, which I illustrate by example applications: to (parametrized) states of affairs, kinds of ordinary particulars, and properties/relations. In the course of this, there are remarks on gerrymandered objects, objects with fuzzy boundaries, the Problem of the Many and the Ship of Theseus. I suggest conceiving of properties as 'recognizabilia' and specifying them via specifying types of corresponding recognition apparatuses (taking into account the agent's goals), thus obtaining a useful notion of a property's intension. Finally, there is a lengthy footnote on what the places of relations may be on my approach, an attempt to explain why Goodman's "grue" is so unnatural, and brief remarks on primary vs. secondary properties, on the notion of existence (particularly for properties) and on justifications for different logics.

standard analytical ontology (Lewis, Armstrong): **ontology by founding** – is what I oppose

I am dissatisfied with the way ontology is standardly done in analytic philosophy[1] – so far as I am acquainted with it. (The ontological methods of other schools of thought will not even be considered here.) The way of doing ontology which I reject is exemplified by David Armstrong and David Lewis. Ontologists like them try to

*This is really just "Ontology by Linking 2" (cf. my 2006), but that sounded *so* boring. – This text was slightly modified on September 5, 2013.

†eMail: Christopher.von.Buelow@uni-konstanz.de; Website: www.uni-konstanz.de/FuF/Philo/Philosophie/philosophie/88-0-NoName.html.

[1]See my 2003 and 2004b.

specify a foundation of everything that is (and of everything that is the case, or might be). (That's why I call this "ontology *by founding*".) What each one of them puts forth as that foundation are entities[2] belonging to certain categories, standing in certain 'relationships' and having certain 'properties'. Examples of categories which may turn up in such an account are ordinary (concrete, physical) particulars, universals, tropes, sets/classes, facts, states of affairs, possible worlds, functions, events, and processes (did I leave anything out? Almost certainly). The 'relations' and 'properties' one encounters in such accounts may not be classified as relations and properties by the respective account; but sometimes they are. Among their number are, e.g., exemplification, similarity, set-membership, and the part–whole relation.

So, a foundationalist ontological account proposes some family of concepts which are then used to describe a foundation for everything. These concepts – concepts for the categories and the 'relations' and 'properties' used in the account – are illuminated to some degree via axioms that govern them, via suggestive metaphors, and perhaps by delineating which of our ordinary words refer to which (non)entities of the account. These concepts cannot be defined explicitly; at least, not all of them can, because they are already the most basic, most general concepts there are, according to that account. The undefined concepts are declared to be primitive and unanalyzable.

Quine's 'naturalistic' ontology – not what I want either

There is another standard paradigm in analytic ontology, namely, Quine's (and Carnap's?) so-called naturalistic way of dealing with ontological questions. There, these questions (those which aren't senseless) are answered by looking at what the accepted scientific theories quantify over. I don't know this kind of ontological method at all well, so I will not treat of it here. Suffice it to say that I do not subscribe to this method either. Anyway, I use ontology of the Lewis–Armstrong type, rather than the Quinean one, as a landmark to distance myself from, in order to locate and motivate my own approach.

my sources/influences

My own ideas concerning ontology are adopted from or inspired by Daniel Dennett's works, especially by the approach to ontology I assume behind his distinction of the physical, design, and intentional stances (1971) and by his paper "Real Patterns" (1991). Though Dennett is an analytic philosopher, his views in these matters certainly are nonstandard. Another influence is Ruth Millikan. – The more I learn (in passing) about the Erlangen school of constructivism, the more it seems to me that there may be some similarities between *its* project and mine. These similarities are, however, merely superficial. As I understand it, this school's starting-point is shared cultural and linguistic practices of societies, whereas mine is arbitrary agents (including stupid ones like thermostats) and their functions or goals.

(analytic ontology: merely *pragmatic* questions?)

I am dissatisfied with analytic ontology on the presupposition that it wants to find out what the world is like 'at bottom'. If its true goal is merely to explore the pragmatic pros and cons of different formal characterizations of the world, then I guess analytic ontology has been successful to some degree; but then it is also much less interesting than I imagined it to be. Seeing how the proponents of competing ontological theories dispute the truth and falsity of their respective ontologies, however, it seems that most of them believe themselves to be debating something more substantial than the pragmatic virtues of various formalisms.

(formalization is nice, but informativeness is important too)

I am not opposed to formalization in itself. It can be great for achieving precision and clarity of expression. It seems to me, however, that in ontology a much more informative, content-rich account can be given than the ones produced by standard analytic ontology, though it may be less amenable to formalization, i.e., to precision.

[2] And/or 'nonentities' like merely possible entities; cf. Spohn 2007.

ontological accounts should be useful for epistemology

One aspect of my dissatisfaction concerns analytic ontology's low regard for epistemic matters. However beautiful the ontological (or ontologico-semantical) story you can tell – if it doesn't help explaining (or, even worse, makes it a mystery) how we can come to know the things we do know, then something is wrong with it.

what redness is *not*

Suppose the property of being red is really the set R of all red things. Then what are we doing when we recognize a tomato as being red? Do we search a mental list of the members of R for a designation of that tomato? Hardly. To 'know' the stupendously big set R at all, we need to understand what's distinctive about its members, viz., redness. That is, before we can grasp the set of red things, we have to grasp redness-the-property (in some feasible other way); hence taking recourse to sets is in general no help in understanding our capability to handle properties. (The suggestion that redness is a *mapping* which has sets of objects as its values and takes possible worlds or instants in time as arguments obviously makes the grasping of redness even more impractical.)

(*modelling* redness as something isn't saying what it *is*)

I do not at all want to deny that it is useful for some purposes (e.g., in model theory or in formal semantics) to *model* properties as, or to *represent* them by, this or that kind of more quotidian entities, like sets. But those projects are altogether different from the one ontology pursues.

being red as exemplification of *red*

Suppose that being red means exemplifying the universal *redness*. Then recognizing a tomato as red consists in noticing that the tomato and the universal are 'related' by exemplification. If these acts are one and the same then this description at least doesn't add any new mystery: recognizing this fact about exemplification just *means* recognizing that the tomato is red. But neither does this description aid our understanding of acts of recognition. Certainly we don't have easier or more direct access to exemplification than we have to redness.

what *is* **ontology by linking?**

I call my own way of doing ontology "ontology *by linking*". What does this mean? What is it that I want to link? One problem I see in 'founding' is that the basic concepts of foundationalist accounts aren't supplied with very much content. They are 'implicitly defined' by axiom systems (like 'the numbers' are implicitly defined by the Dedekind–Peano axioms), i.e., they get a kind of structural characterization in terms of *each other*. What we would of course prefer to such an implicit definition is an *explicit* one. But there are no more basic, more general concepts we could use to formulate such a definition. So there can't be any way of making ontological matters completely explicit. Whatever you use as basic must remain undefined to some degree; or, speaking ontically instead of linguistically again: it must remain unexplained to some degree. More precisely, your basic concepts cannot be explicitly defined in terms that are already perfectly clear and understood; and the basic kinds of 'facts' your account recognizes must be brute.

implicit and explicit definitions

different ways of implicitly defining ontological concepts

It follows that the best we can do is, give an implicit definition of the basic concepts of ontology. It does not follow, however, that the only possible way of doing this is the one taken by foundationalist ontology; an implicit definition of ontological terms needn't characterize them *in terms of each other*. This is where epistemic (and teleological) matters come in.

agents and information-processing

I want to illuminate the ontological concepts by looking at how they arise from the information-processing constraints arbitrary agents labor under. By "agents" I mean people, other organisms, robots, and other automata – physical systems with some degree of autonomy, which can be viewed from the design or intentional stance.[3] That

[3]Maybe also idealized agents which are physically impossible, in order to deal with those entities or

is, these systems can be usefully conceived as having a certain function or purpose, or as having certain goals and beliefs.[4] In order to reach their goals (or fulfil their purpose), agents must successfully negotiate their environment; to do that, they must react adequately to any given situation; and to do that, they must recognize what their situation is at each moment and decide what to do about it. This recognizing and decision-making by agents is their processing of the information that can be gleaned from the physical symptoms of the goings-on in their environment which impinge upon them.

carving nature at its (more or less salient) **joints** (because we can't swallow it whole)

How do ontological categories figure in agents' information-processing? I don't want to say that agents in general think about ontological matters. Not at all. But their information-processing has to proceed by representing the world in certain categories, by 'carving nature' at certain 'joints'. It must portion off the environment in manageable chunks, leaving out ('abstracting from') the infinitude of further detail that is there. I believe these chunks are predecessors of our concepts, and these *ways* of carving nature are predecessors of our *ontological* concepts.

This part, about information-processing and the 'units' it is done in, will be highly contentious, I expect. I suppose these ideas might get some support from psychology, cognitive science, cybernetics or related quarters. But as yet I cannot adequately defend these assumptions against sceptics, so I must ask my readers to indulge me in this for the time being. Hopefully the fruits of these labors will go some way toward redeeming the approach.

ontological concepts mirror the kinds of chunks we (must) carve nature into

So, I want to clarify and explicate the ontological concepts we use, both in academic philosophy and in everyday reasoning, by bringing out what their respective functions are, what kinds of phenomena they are there to capture. I believe these concepts are[5] abstractions from the ways we and other kinds of agents parcel out the world. Furthermore I believe that these ways of carving nature aren't arbitrary. However, they need not be forced on us by the structure of the world, i.e., they needn't be nonarbitrary in the sense that the world actually *is* structured like we construe it. For all I know, the world may be one indivisible, irreducible whole, like a blanket where the folds we see aren't really distinct parts but merely aspects of the shape of the whole. Perhaps any description of the world that doesn't capture *every* aspect of the world at once, down to the last nuance, must be incorrect to some degree; perhaps there are no true joints in nature. On this subject, I don't have an opinion. When I say that our ways of conceptualizing the world aren't arbitrary I mean they are forced on us by our limitations as finite agents: by the finite resources for information-processing we have at our command and by the economics of information-processing.[6]

linking ontological concepts to everyday, scientific, and cognitive concepts

Back to linking: Our ontological concepts aren't very well understood. I want to make them better manageable and their contents clearer by linking them to concepts we understand better, viz., our everyday and scientific concepts, and to some other concepts which we possibly do not understand a whole lot better, namely, "agent" and "information-processing" or "pattern-recognition". We take the latter three groups of concepts as our reasonably-well-understood conceptual basis (or 'foundation', if you like) and elucidate the ontological concepts from there. *Then*, when we have gained a

structures of mathematics which are beyond the reach of any physically possible, finite agent.

[4] For less rudimentary accounts of Dennett's stances, read Dennett 1971 or my 2004a, sect. 2. I have decided against putting scare quotes around all the 'mentalistic' or 'intentional' terms (e.g., "react", "recognize", "decide"). In predicating them of arbitrary agents, I do *not* presuppose consciousness or intelligence in those agents. I use these terms in the very broad sense which allows us to say about a computer things like, "Now he wants you to tell him where to put the file."

[5] With a few exceptions, like tropes.

[6] Cf. Ladyman and Ross 2007 about compressibility.

somewhat better grasp on the ontological concepts than we had before, we can use the greater clarity of their explicated versions to improve our understanding of what we have formerly used as our conceptual basis; what acted as foundation before, now becomes the object of our efforts at refinement. And when we're through with that, maybe roles are reversed again, and again, in a virtuous circle. (a circle of explications)

There is no absolutely secure foundation, there are just different degrees of security and clarity. So we have to use some reasonably firm ground (hopefully, *common* ground) to reinforce some other territory from (some other deck of Neurath's ship), which can then in turn be used for further improvement measures. Just because we are dealing with (in some sense) maximally basic matters doesn't mean that we have to take these matters as inexplicable. Or, turning to definitions again: instead of an implicit definition of the ontological concepts in terms of merely each other, we try to obtain a more comprehensive definition, involving also everyday and scientific concepts as well as cognition-related information-processing concepts. The new definition may *look* like an explicit definition of the ontological concepts in terms of everyday–scientific–cognitive concepts and may be so interpreted; but more precisely it would be construed as part of an implicit definition of all these kinds of concepts together. Neurath's hips

I see this as analogous to the situation in fundamental physics: At the 'bottom' level, there are concepts of weird elementary particles and other strange stuff. These get some content by a description of how the entities falling under them interact – kind of an implicit definition. But not just any class of entities which satisfy this structural characterization can be the elementary particles etc. *The* elementary particles are special in that they are what is measured by the physicists' instruments and manipulated by their experimental apparatuses – and what everyday phenomena consist of, in the end, even if we can't perceive it by everyday means. They are unique among entities interrelated as the basic laws of physics say, because of their causal correlation with macroscopic phenomena, including everyday phenomena like tables and cats. (Some of these correlations are what is described by 'correspondence rules', I suppose.)[7] (foundations in physics)

To summarize: instead of trying to describe an absolutely safe foundation where none exists, I propose to *link* conceptual domains (or domains of phenomena) which hitherto we did not relate (or at least did relate much less strongly), thereby heightening clarity and understanding on all sides.[8]

Clarifying concepts is customarily done in analytic philosophy by looking at language: at how we use these concepts in talking about the world, and at our intuitions concerning what would or wouldn't sound reasonable in various circumstances. To be sure, the linguistic perspective has the big advantage that the objects of scrutiny – words, sentences, utterances and linguistic intuitions – are relatively well discernible. The 'informational' perspective, by contrast, seems rather speculative, and what we see from it, rather ill defined. In which categories do we think? Aren't the categories we think in just *words* again? I believe this is at least not the whole truth, though it probably is an important part of the truth about people. It can't be the whole truth because infants and many nonlinguistic animals do also think; and if we look at nonconscious information-processing (as performed, e. g., by current computers) then we have many more examples where language isn't involved: all animals, for instance, Digression: **the linguistic approach** – not the *ultima ratio*

[7] I also see a parallel to Dennett's intentional and design stances, or think these can be understood along the lines of what I have said about explicit and implicit definitions; cf. my 2006, p. 2.

[8] Cf. my 2004a, Sect. 7.1.3.

and also automata. In the absence of language use, here there is still recognition and representation of, as well as adaptation or reacting to, features of the environment.

Language matters for ontology *insofar* as its basic structure reflects the basic structure of the world. But how far is that? Which of language's features are indicative of features of the world, and which are merely extraneous, contingent features of one or all human languages, that could just as well have been different, e. g., in Martian languages? Or is reality perhaps language-dependent, so that world and language (*which* language?) are isomorphic? To most readers this will not be very plausible, and neither is it to me. I believe language reflects (to some degree) the way we think, and this in turn reflects (to some degree) what the world is like. Hence, language reflects (to some degree) features of the world. But our concepts and our language needn't mirror the world faithfully. All an agent needs is a working model of the features of the world that are (sufficiently) relevant to its goals (or its function). Where bending the truth a little works better for its purposes, the truth gets bent, e. g., by simplification/idealization. So we mustn't put too much faith in the categories of our thought (or our language). They aren't absolutely reliable indicators of the structure of the world.

ontological categories as natural kinds

"Particular", "property" and their kin are in some sense the most comprehensive, most general *natural kinds* there are. On the other hand, they may not be as 'natural' as the categories used by, say, physics. Our everyday reasoning doesn't involve a curved space-time or probability waves, but nevertheless these may be considered as candidates for the most basic features of the world. So, the ontological categories most natural for us needn't coincide with the fundamental categories of science. In what sense then are they natural kinds, if at all? – I will say more on this question when I deal with 'primary' and 'secondary qualities' on page 18 below.

explication of the method

I want to get clearer about how I conceive my ontological method before I try to present it in action.[9] The method consists of three parts, which build on each other and can be characterized by three questions:

1. Given some ontological category C (say, states of affairs, or ordinary particulars, or universals), when should we interpret an agent as *treating* an arbitrary sort of phenomenon *as a* C?

2. We inspect a certain ontological category C and a certain kind of traits of phenomena, t (for example, being chemically rather homogeneous or rather heterogeneous; having rather sharp or rather fuzzy boundaries; moving in a continuous trajectory or not; being a rather fleeting or a rather lasting phenomenon). Then we ask the following question: for arbitrary types of agents, α, and arbitrary sorts of phenomena, P, suppose we (or rather, evolution) contemplate(s) constructing an apparatus that will allow α to recognize P-phenomena, treating them as C's; *how will the trait* t (or its various manifestations) *influence the cost–benefit ratio*[10] for such apparatuses?

[9] However, the following explication attempt will be hard to comprehend before having read the illustrative example investigations that come below.

[10] In evolutionary terms, what I refer to by "cost" and "benefit" are matters of the following sort: How much of the apparatus is already in place, having evolved for other reasons, and how much must be 'evolved from scratch'? How many fortuitous mutations will it take? How much time and effort will it take a particular organism and its parents on average to obtain the additional food necessary to grow the new features needed for the apparatus? How much will the apparatus enhance the agent's reproductive success?

3. What *is* there (which of the putative C's exist, are real), and in what sense? (to which degree?)

Giving an informative, useful answer to *Question 3* – "What exists?" – is of course the main goal of the ontological enterprise. To do it, we must get a better grasp on the ontological concepts, including "existence" itself. That means we must find reasonably good explications of these concepts.

Question 3: what exists?

That's what *Question 2* is for. To explicate a concept, one must find out what its purpose in cognition and communication is. A concept's purpose or function is, to capture and represent certain kinds of phenomena (or patterns, or regularities) in the world. What kinds of phenomena is our concept of C-entities intended to capture?[11] For example, is the concept of ordinary particulars (this 'way of carving nature') intended[12] to be applied foremost to phenomena with sharp boundaries (e.g., to crystals rather than clouds)? If so, that would be a reason, I suppose, to conclude that having sharp boundaries is a central part of our concept of ordinary particulars.

Question 2: what makes a phenomenon 'suitable' for treatment as a C?

To find out what an ontological concept's – say, C's – function is (especially, what the domain of phenomena it is intended to be applied to is) we must ask why it – or rather, the corresponding way of cognitively carving nature, the ability to recognize phenomena, treating them as C's – *evolves* in organisms generally. Let t^+ and t^- be mutually exclusive traits phenomena can have (e.g., sharp vs. fuzzy boundaries). Then we can ask: does it make a difference for the '*profitability*'[13] of corresponding recognition apparatuses whether phenomena of type P have t^+ or rather t^-; and if so, which of the two traits *enhances* for P-phenomena the 'profitability' of being able to recognize them, treating them as C's, and to which degree? Suppose t^+ strongly enhances in P-phenomena their 'evolutionary suitability' for treatment as C's, while t^- renders their suitability low or negative. Then I would infer[14] that our C-concept is *intended* to capture P-phenomena *with* t^+, not with t^-; and P-phenomena *with* t^+ are good candidates for being C's, are *paradigmatic C's*, while those with t^- are at best borderline cases. A good explication of the C-concept must then take this into account.

In answering Question 2 we will have presupposed that we know reasonably well what it means to say that an agent treats something as a C. So, to get off the ground at all with Question 2 we must first answer *Question 1*. Here, we invest part of our C-concept into our account: you cannot judge when to interpret an agent as treating some phenomena as C's if you don't have some preconception of what C's are like. In doing this we run the risk of getting out of our account only what we have put into it in the first place. However, the risk is quite small, I think. My example investigations should demonstrate that we can reap a lot more than we have sown.

Question 1: what's 'treating as'?

I'm afraid my distinction between *treating* phenomena as C's (e.g., as particulars or as universals) and *recognizing* them as such and such (e.g., as prey or professors) will have caused some confusion. I try to make the distinction clearer: Anteaters can recognize ants. But what do they *treat* them as, cognitively? Do they treat each

(treating vs. recognizing)

[11] This is more or less the question of the intended area of application, or the "*Definitionsbereich*", of the given concept. For which kind of phenomena will it work, for which not? Where will it work best, where less well, where will it break down? (Of course there's a difference between what it is intended for and what it does; these needn't coincide.)

[12] For evolved organisms my use of "intended" is of course merely metaphorical. A trait of an organism is 'intended' to do p iff it evolved for that 'purpose', i.e., iff its ability to do p gave its bearers the selective edge which brought about that they displaced from the population the organisms *not* having this trait.

[13] Cf. footnote 10 on page 6 above.

[14] Entitled by what I'm not sure.

ant as a particular? I don't think so: they do not gather and use information about individual ants, e.g., ant Anna's coloration and habits. Nor do they treat ants as a *kind* of particulars: an anteater's behaviour is not directed at single ants individually. Rather, he licks them up like we lick on an ice-cream cone. So, anteaters treat ants like a *stuff*, like water or ice-cream. Do anteaters *recognize* ants *as* a stuff? No. That kind of feat is usually only accomplished by ontologists and other similarly sophisticated agents, I suppose. And these agents probably *wouldn't* categorize ants as a stuff. (However, they do characterize *water* as a stuff, even though it consists of individual H_2O-molecules.) In short: anteaters can recognize a stuff, namely, ants (they can recognize ants as *ants*; i.e., they have something like a concept of the ants-stuff), but they can't recognize the ants-stuff *as* a stuff, because they don't have anything resembling a concept of *stuff*.

Now, let's ask Question 1 for different categories. We start with states of affairs (*Sachverhalte*).

what's 'treating s.th. as a *state of affairs*'?

A state of affairs ("s.o.a.") either obtains or doesn't. Accordingly, an agent *treats* some phenomenon as an s.o.a. in recognizing it as so-and-so if he has some kind of internal switch[15] 'representing' that phenomenon. That is, the agent can flip between two states, one of which results in behaviour adequate for the phenomenon's obtaining, and the other, for its *not* obtaining; and the change is triggered by the agent's registering of symptoms which are diagnostic of either the one or the other.[16] An automatic door-opener constitutes an example: either there is someone who wants to enter or leave (an s.o.a. which can also be expressed by "open the door and keep it open!"), or there isn't ("keep the door closed, or if it's open, close it!"). The door-opener doesn't register any further details, like, how many people want to pass through, how old they are, what sex they are, how fast they walk. He doesn't even register whether he is opening the door for people, for dogs, or for out-of-control baby-carriages rolling past. All he registers is massive bodies in motion, and for those he opens the door.[17]

graded s.o.a.'s

However, already for s.o.a.'s things are not quite so simple as some might hope. Consider thermostats, which can, in their own humble way, recognize whether it's warm enough in a room. Maybe there are very simple ones which can only switch on and switch off the heating. But the usual thermostat is able to administer stronger and weaker heating, depending on how far below the preset temperature the room is. So, does such a thermostat recognize just a pair of s.o.a.'s, or lots of s.o.a.'s, or maybe some more complicated kind of entity? I tend to say it recognizes a *graded* s.o.a., one which can obtain to a higher or lower degree, which can be more or less urgent. And there are other ways s.o.a.'s may be 'parametrized': one kind of primitive organism may be able to distinguish merely between danger ("move fast!") and no danger; another may be able to discriminate additionally between danger from behind ("move forward fast!") and danger from in front ("move backward fast!"). Philosophers will prefer to conceive of these different *aspects* of one s.o.a. as really different *s.o.a.'s*, but I think this would be rather artificial in my example cases: is every degree of deflection of a thermostat's bimetallic strip designed to

[15] Or a boolean variable, in computer-science terms.

[16] A not-too-exciting consequence: An agent who can recognize *one* s.o.a. can also recognize another one, viz., the former's *not* obtaining.

[17] Future door-openers may of course be more sophisticated. On the upper end of the sophistication spectrum we will perhaps one day see doors which can truthfully say, "Thank you for making a simple door very happy" (Adams 1995, 80).

represent one particular temperature (or degree of 'temperature suboptimality')? – Should we however decide that it isn't okay to use the term "state of affairs" for these phenomena, then we can still introduce "parametrized s.o.a." as a technical term for a new ontological category. A parametrized s.o.a. would be a family or spectrum of true s.o.a.'s, I suppose.

what's 'treating s.th. as *a kind of ordinary particulars*'?

When does it begin to be more appropriate to speak of an agent's treating some kind of phenomenon as a *kind of particular* instead of as a state of affairs? Part of the answer is: when recognizing that 's.o.a.' essentially includes tracking *locations* such that reacting to this 's.o.a.' is always done directed at those locations. For example, suppose the s.o.a. "predator(s) present" is recognized in such a refined manner by the prey organism that it keeps track of an indefinite (though not too large) number of places where this s.o.a. is especially 'salient' or 'urgent' (viz., what would be considered as the predators' locations from the ordinary point of view). Then, while one might go on categorizing this as an s.o.a., just 'parametrized' in an especially complex way, it would be more elegant and natural to say that what the prey organism does is, to recognize *predators*, i.e., a certain kind of ordinary particulars.[18]

what is 'treating s.th. as a *relation*'?

When should we say that an agent can recognize a *property* or (non-zero-place) *relation*, i.e., that he treats some kind of phenomenon as a property/relation? For one thing, a property draws a line among (a kind of) *particulars*, between those which have it (or are related by it) and the rest. So, when an agent α can recognize some kind of particulars (or maybe a stuff) *and* distinguishes among its instantiations according to some criterion, treating the ones this way and the others that way, then we *might* perhaps say that α can recognize a certain property here. Then again, maybe being able to recognize a property or relation takes more than that. Maybe α can truly recognize a property only when he can recognize its instantiations among *different* kinds of particulars, reacting to these instantiations differently, depending on which kind they belong to. For example, just because a prey organism can keep track of where his predators are we probably wouldn't say that he can recognize properties like "being in *this* place" and "being in *that* place" (even if these location properties are conceived as merely picking out certain locations relative to the agent's own body). What the agent recognizes here are only different 'parametrizations' of predator occurrences. But when an agent interprets redness in *apples* as "good to eat", and in other agents' *faces* as "angry", maybe then it's justified to say he treats redness as a property. And maybe it takes even more; I don't know.

I have avouched that my approach could help clarifying ontological concepts. Let's ask Question 2 now and see in what measure I can redeem that promise. In the process, it should also become clearer what my approach consists in.

states of affairs

To be an agent, you must be able to *do* something. That something may be extremely trivial, like an automatic door-opener's opening or closing of a door. Also, you must do it at the right times, i.e., you must be able to recognize occasions (*Anlässe*) for doing it. Thus, all agents are capable of recognizing some *states of affairs*. A functioning door-opener can recognize the state of affairs "somebody wants to enter or to leave" or, what amounts to the same in his case, "I should open the door

[18] In computer-science terms what happens in the prey agent could be described as follows: As soon as the agent notices that the s.o.a. "predator(s) present" obtains he starts generating an adequate (though not too large) number of temporary variables in which to store the varying location(s) of these predators (and possibly other information pertinent to each corresponding predator, if he is more intelligent). The changing locations stored in the variables are then treated in the manner appropriate for predators: the prey agent must move away from them or, if need be, attack them.

now". So, from my agent-centered point of view I would characterize states of affairs as *occasions for doing* (*or refraining from doing*) *something* (or for behaving a certain way?).[19]

Objection: Aren't there lots of s.o.a.'s which just don't matter, e.g., to us, and which therefore *aren't* occasions for anything? For example, when a pink balloon is flying past, that's not a reason for doing, or omitting to do, anything. We just recognize that this s.o.a. obtains, and probably forget about it again soon, and that's all. – My answer has three parts: (a) You never know whether it might not come in handy one day to be able to recognize even some ordinarily completely irrelevant s.o.a. Some day it might be a sign of something very important happening. Then recognizing the boring s.o.a. would be a means of recognizing an important one, and could thus be the recognition of an urgent occasion for action. In the case of big-brained creatures like us, it has been easier for evolution to furnish us with wholesale recognition capabilities than with lots of single-purpose recognition mechanisms for just those s.o.a.'s which ordinarily matter for us. (b) Even in pink-balloon cases there often is something we do: we remember – if only for a while. That is, we modify our *internal* structure such that, should the need arise, we will act *outwardly* in accordance with the s.o.a. of a pink balloon's having flown past, in whatever way is appropriate.[20] (c) Finally, I could cede the point and admit that occasions for action are merely the *root* of the s.o.a. category, the paradigmatic cases (at least for the agent-centered perspective), and with the advent of very sophisticated agents the category has come to include phenomena whose recognition doesn't occasion any action even in the slightest degree.

kinds of ordinary particulars (objects, things)

I will go on with kinds of ordinary particulars or 'objects'. Belonging to this category are stones, lakes, cats and tables. People, considered as living creatures, are on a par with cats, though considered as persons they may be not quite so ordinary particulars. Other borderline cases are cascades like the Falls of the Rhine, candle flames, holes and shadows. What do these things (or phenomena or whatever) have in common? Or, to rephrase the question in accordance with my specific approach: When they are worth reckoning with for an agent, *why* are they so? In still other words: Which traits of phenomena make them more suitable for treatment as kinds of particulars?

a simple nature

For one thing, they should have some homogeneous (or 'simple') nature throughout their spatial extension, and be *maximal* with that property. What I mean is, when an agent stumbles upon, e.g., a candle, it will in general be more economical for him information-processing-wise if he treats the *whole* candle as an object, rather than treating as two independent objects its upper and its lower half. For most purposes, in most circumstances, thinking about the candle as a whole will deliver the same results as thinking about its two (or more) parts separately (because one long piece of wax will in most respects behave just like the two conjoined shorter pieces of wax it

[19] Since (some) s.o.a.'s are the first phenomena any agents can recognize at all one might say that s.o.a.'s are prior to all other ontological categories, including particulars and properties/relations. On the other hand, one could perhaps consider s.o.a.'s as zero-place relations, so we shouldn't make too much out of this.

[20] *Internal* behaviour may not be your idea of an action. This broad sense of "action" is appropriate here, however, because it yields a more elegant picture of how agents work (I think). And yes, when we remember, it isn't *us* who modify our brains; our brains or bodies do it by themselves. We can conceive 'ourselves' as agents in different ways: categorizing us as persons, or as animal bodies (maybe there are more options besides). Here, I sometimes do the one, sometimes the other. What the person wants needn't be identical with what the body is designed to do. Much of what we as persons do (believing, wanting and remembering included) depends on our bodies' doing things we usually don't notice.

consists of), but the former will be faster and use up less memory. On the other hand, if the agent construes the 'object' before him as *bigger* than the candle, including some adjoining volume of air, he will be dealing with a heterogeneous 'object', part wax, part air. Because of its bipartite chemical nature, this 'object' will be more difficult to deal with than a homogeneous piece of wax. So, the agent will do best if the 'object' most salient to him has as its extension just the extension of the candle, no more, no less.

This is kind of a derivation of the maxim exhorting us to carve nature at its joints, as applied to ordinary particulars. Just as 'objects' like "this table today, my mother tomorrow, Mount McKinley on the day after" consist of very heterogeneous temporal parts (each of which has far more natural possible 'expansions' or 'extrapolations' in time), so the candle + air 'object' consists of very heterogeneous spatial parts.

The 'natures' which segregate the more natural ordinary particulars from their surroundings can be of different types. There is chemical nature or substance/material, as in the case of the candle, but an agent might also differentiate some kinds of objects by purpose or by origin or still other respects. This will depend on the agent's goals/function.

a *stable* nature

If, during its existence, an object wildly changes those of its properties which are relevant for the agent, bothering about it will again not be worthwhile for him. To be worth reckoning with, an object's relevant aspects must be reasonably stable (or change according to a reasonably simple rule, like caterpillars into butterflies; cf. Carle 1969). If an object changes too fast or too randomly there is no knowing how to deal with it, because a behavior towards it that was adequate one moment may be useless or detrimental the next. If an object tends to change *radically* there may not even be a way of reidentifying it for the agent.[21] This is why the table–mother–Mt. McKinley 'object' is weird.

no bilocation

It seems to be important that the entity in question not be (wholly) in two or more places at the same time. When you have to worry that the opponent you are fighting in front of you may close in on you from behind, then you shouldn't treat him as *one* ordinary object but rather as two very similar ones, or as two parts of a not-so-ordinary object. If you knock out the opponent in front of you, you shouldn't have to worry that behind you he may keep fighting.[22]

continuous movement vs. 'teleporting'

Are there constraints with regard to movement in space which must be satisfied by an entity or phenomenon to qualify as a reasonably ordinary, natural particular? It does not seem to be essential for that entity to move only continuously, instead of 'teleporting' from place to place sometimes, as long as its 'nature' stays more or less the same. The table–mother–Mt. McKinley 'object' teleports around (from, say, the kitchen where the table stands, to wherever my mother is during the day on which she constitutes that 'object', to the mountain ranges of Alaska) *and* in so doing radically changes its 'nature': from being a piece of furniture, to being a person, to being a mountain. If, however, everything about an object except its location stays more or less the same in teleportation then it would be appropriate to *treat* it as the same (except when 'it' would sometimes be in two places at the same time as

[21] Even the shape-shifters in science fiction or fantasy tales have a stable core consisting of their personality, their beliefs, goals and memories. After radically altering all of their physical aspects they keep on acting to further the goals of their old self; that's how you can still recognize them.

[22] In a time-travel scenario, however, it may be reasonable to say that Anna today, and Anna's future self who has traveled backward in time to reach the present day, are in some sense the same person – though not in every sense: the two Annas will differ in some of their properties at the same time, e.g., they may wear different sets of clothes. – Similarly for duplication of people by other means.

a result). When an object moves too fast for an agent to track it, it doesn't matter very much for him whether it does so in a continuous manner or not. It does make a difference, though, with regard to the possibility of the object's motion being impeded (or predicted): a continuous mover can be stopped by a wall or caught in a cage, a teleporter cannot.

continuous existence? – not so important

A phenomenon may be better suited to be an ordinary particular when it is located *somewhere* at all times from its origin to its termination. But, again, this does not appear to be of the essence: usually it doesn't make much of a difference for an agent whether some object is merely somewhere else (where it can't interfere or be interfered with) or whether it doesn't exist at all. The unimportance of continuous existence is reflected in parts of our everyday reasoning: a lake can dry out for long periods and still when its basin fills up again we usually[23] consider it the same lake – even though while the lake was dried out there *was no* lake.

sharp vs. fuzzy boundaries

Do ordinary particulars need sharp boundaries in space and/or time? I can't see why. How does an object's (sharp or fuzzy) boundary feature in an agent's behaviour? Mainly in three ways, it seems: it plays a role in the recognition, the localization, and the manipulation of the object by the agent. An agent may recognize an object by its shape, and that is usually a function of its boundary (although some objects' distinctive shape may be internal). Whether an object is in a certain place or area will usually also depend on its boundary, viz., on *its* location relative to the given area.[24] And to manipulate an object, the agent has to make sure that his efforts are directed at the object's location as precisely as necessary. But in all of these activities it seems irrelevant whether the object's boundaries are sharp or fuzzy. It may often be preferable for the agent if the object's boundary is sharper instead of fuzzier, but it is almost never important whether the object's boundary is sharp at the level of, e.g., single atoms.[25] The boundaries of ordinary objects are never sharp in the sense some philosophers (e.g., Unger 1980) prefer, but their slight fuzziness doesn't matter in the least, because they are usually sharp enough for our purposes. If you want to grab a cat's neck, its hairs are unimportant. And a cumulus cloud's boundary may be extremely fuzzy, but it is still sharp enough to shoot a rocket loaded with silver iodide at the cloud in order to make it rain.

(Unger, Lewis, and 'the many')

The belief that nice objects need sharp boundaries leads Peter Unger (1980) and David Lewis (1993) into rather awkward positions: Unger says that cats or tables do not exist at all; and Lewis says that for any given cat, there are really millions of very similar cats roughly in the same place. For me, these consequences amount to a *reductio* of the belief engendering them. Why would someone want to cling to that belief? For one thing, it may seem inherent in our concept of ordinary particulars that they have sharp boundaries. But what is inherent in that concept is at best the boundaries' sharpness in the everyday, not the microcosmic, sense. Then, a philosopher may depend on an object's boundary to *individuate* it, i.e., to constitute it as an object in the first place, by delimiting it and separating it from the rest of the world. An agent, however, will individuate objects (at least ordinary ones) by the *kinds* of objects they are. That in its turn is, I suppose, a matter of their instantiating certain

[23] If the whole terrain hasn't changed too much.

[24] This is not to say that the agent must recognize the object's boundary as such, and locate *it*, before he can locate the object.

[25] Yes, that *is* important when the agent is a physicist using a scanning tunnelling microscope (*Rastertunnelmikroskop*). But even in exceptional cases like this one fuzziness on the order of small fractions of a nucleus's diameter is again irrelevant.

patterns or properties, and to do that, they usually don't need sharp boundaries.[26]

When an object has sharp boundaries in *time*, there are two points in time such that the object's existence begins at the one and ends at the other – precisely.[27] (We neglect complications due to possible temporary nonexistence.) It will be unsettling for some philosophers to imagine (if they can) an object that doesn't have absolutely sharp boundaries in time, because then there are moments in time for which it isn't determinate whether the object exists or not, or at which the object exists just a little or exists mostly, but isn't either definitely there or definitely not there (and the degree to which it is there is itself vague, with no specific moment where it changes from zero to nonzero or from below 100 % to fully 100 %). I believe it is like that for most or all ordinary particulars, if you look closely enough: the pattern or property that individuates a particular object (in the beginning – for this may perhaps also change with time) doesn't get instantiated in one fell instantaneous swoop, from zero to one in literally no time. You may ask: "What is there to instantiate that pattern 'to some degree' when the object itself doesn't yet fully exist?" The same thing that instantiates, at bottom, the pattern when the object does fully exist: the world, or, more specifically, that (fuzzy) region in the world where the object is or comes into being.

boundaries in time

Let's try to put these ideas to use. As a sample puzzle, consider the ship of Theseus. Theseus starts traveling around in a wooden ship s. From time to time, this or that worn-out plank is replaced with a new one until, after a few years, all the original parts of ship s have been supplanted. During all this time a used-ship dealer has been following Theseus and has collected all the discarded parts. These are not yet so rotten as to be completely unusable, and the ship salesman reassembles them in just the way they were originally assembled in ship s. Now there are two ships: the one Theseus is traveling around in, call it s_t ("t" for "traveling"), and the one reassembled from the parts of the original ship s, call it s_r ("r" for "reassembled"). Now, is s identical with s_t or with s_r or with neither? (Obviously s can't be identical with both, for then the clearly distinct ships s_t and s_r would have to be one and the same.) On the one hand, the identity $s = s_t$ is plausible because the successive replacing of one small part after the other, always keeping intact the structure, function and use of the whole, does not seem to change the identity of the whole. Ship s_t is, as it were, s after extensive repairs. On the other hand, $s = s_r$ is plausible as well because s_r consists of the same material (save for some degradation) the original ship s consisted of, assembled in the same arrangement. It is as if s had been taken apart and put together again, obtaining s_r.

the ship of Theseus

Which identity statement is correct? The biggest mistake one can make here is believing that a correct answer must exist. Our concepts (in this case, "ship", and "identity" as applied to ships) are tools for reasoning; and like other tools, they aren't built to function under all circumstances whatsoever. Ships and other artifacts are individuated by a handful of traits which make up a particular artifact's 'nature', among these continuity of use and sameness of material constitution. Under normal conditions (i.e., in the kinds of circumstances or contexts for which our concept of ship is intended) these different characteristics all point in the same direction: usually there is at best one plausible candidate for being identical to some ship that once

[26] Along these lines I may arrive at something like Aristotelian substances.

[27] Let's say a certain object exists from noon to midnight. If we take these temporal boundaries as sharp in the philosopher's strong sense then it is a reasonable question whether this object does already exist at noon, precisely, or *only* at all the moments after noon (at noon$+\varepsilon$ for arbitrarily small $\varepsilon > 0$), until midnight (again in- or exclusively). But isn't this distinction in reality rather weird?

was; all criteria point at it, or at least none point elsewhere. But in some contrived imaginable situations, which are very unlikely ever to occur,[28] these criteria disagree, and designating some of them as unimportant in order to recover a unanimous 'vote' would be *ad hoc*. This is what happens in the ship-of-Theseus story: the criterion of continuity of use suggests $s = s_t$ while sameness of material constitution suggests $s = s_r$. Expecting a definite answer here is like expecting a compass to work inside an MRI[29] scanner.

properties

Next, I want to look at properties. How do they enter into an agent's information-processing? First, the agent must be able already to recognize some kinds of particulars, or else there is nothing he can recognize as *having* a certain property. You can't recognize quadrupedality if you can't 'see' any objects whose legs you can count. In this sense particulars are 'prior' to properties.[30] "But," a sceptic may argue, "we can have impressions of, e. g., redness, wetness or heat without having the faintest idea of what kind of object it is that exemplifies these properties in a given case." I think, however, that these are not examples for recognition of properties. What is being recognized here are rather certain *states of affairs*, which might be characterized, say, as "There's something red in front of me" or "There's something wet (or hot) close to my skin."

properties as recognizabilia

What are properties, then? They are, I would say, what a particular can be recognized as having or not having. Insofar as this characterization has any information content at all it seems to get something wrong: contrary to what it says, some properties may be impossible to recognize. That's a valid point, against which I want to say for the moment only that the agents doing the recognizing maybe needn't be physically possible agents (cf. footnote 3 on page 3).

But there is more information to be milked from this characterization than meets the eye; it doesn't merely imply that properties are what objects have or don't have. The capability of recognizing a certain property (in objects of some kind) isn't an unexplained feature magically present (or assumed) in an agent; recognizing a property

recognition apparatuses

has to be done by some *means*. An agent is able to recognize a property of objects because of some *recognition apparatus* he has for that property, i. e., because of certain parts of his sensory apparatus, certain methods of veri- or falsification and/or certain inductive or deductive procedures. For example, we can see colours because of how our eyes, our optic nerves and the visual centers in our brains work.

intensionality of properties

The best way of specifying a property, I think, is by specifying the corresponding recognition apparatus. In this manner we capture the intensional aspect of properties: the recognition apparatus for a property p embodies not only the information of which objects *actually* exemplify p, but also, which objects *would* exemplify p, given

Sinn & Bedeutung

certain counterfactual (or past or future) circumstances. Thus we get a concretization of Fregean senses: *Sinn* and *Bedeutung* of some unit of language can both be given by specifying a recognition apparatus determining which objects fall under a certain concept (or whether what a certain sentence says is the case, or which entity is the denotation of a certain name). A (type of) recognition apparatus embodies a way for a property (or, more generally, an entity) to be given.

[28] Assuming a not-too-advanced level of technological development.

[29] Magnetic resonance imaging.

[30] I don't want to assert a fundamental difference between particulars and properties/relations, however. Maybe these aren't definite disjoint categories at all; maybe one property's instantiating object is another object's property. But if you have a kind of objects and a family of properties/relations applicable to these objects then the ability to recognize the objects is a prerequisite for recognizing (instantiations of) the properties/relations.

This account of properties has a built-in connection to epistemology: Properties aren't *identified* with recognition apparatuses, but they are what is recognized *using* such apparatuses. We don't grasp a property by *knowing* the corresponding recognition apparatus but by *having* it. This account of properties doesn't really tell you what properties *are*, it doesn't give an explicit definition of "property". Nevertheless, the kind of implicit definition given here is more informative than the other characterizations that have been proposed up to now – or so I hope.

identity criteria

My approach also reflects the difficulties with identity criteria for properties:[31] When are properties p and p′ identical? When the corresponding recognition apparatuses are the same. "The same" doesn't mean "identical" here, of course – we want to allow for different agents' being able to recognize the same property without having to take recourse to one and the same apparatus token. No, for p and p′ to be identical, the corresponding apparatuses need only be of the same type. But whether they are of the same type depends on how fine-grained you conceive types of recognition apparatuses. Must they be atom-for-atom replicas of each other? Then no two actual apparatuses would characterize the same property. It seems to me that even if degrees of fine-grainedness for apparatus types can be specified, any fixing of such a degree will be arbitrary. Still, by specifying properties via recognition apparatuses, we have given some substance at least to claims of *similarity* between properties. Even identity statements for properties can still be made, but they would always have to be qualified by saying with respect to which conception of *apparatus type* they are made.

mistaken recognition

There seems to be one problem in this account of properties: the possibility of a recognition apparatus's making mistakes. Even the best-designed apparatus for recognizing (instances of) the property p will sometimes fail to register, or falsely register, that an object has p (especially under somewhat 'unnormal' conditions, where its efficacy is diminished). Can this be, when the responses of the recognition apparatus are supposed to *fix* the intended property? When Otto sees a red tomato as purple (e. g., due to unusual lighting conditions), should we say that the tomato *is* purple – in the sense specified by his visual apparatus? Surely not: Otto and his visual apparatus are *mistaken*. But then how can purpleness be characterized by recourse to Otto's visual apparatus (assuming him to be normal-sighted)? This seems to refute my account, but it is really an occasion to give a more sophisticated rendering of the account. When an agent has a recognition apparatus r for a property p, the extension of p isn't delineated by the responses apparatus r will or would output under various circumstances, but rather by the responses r is *intended* to output (interpreting r and the agent from the design stance). So, what matters isn't what r is disposed to do but what it is designed to do, taking into account the agent's goals/function.[32]

[31] If you expect an account of properties to yield precise identity criteria then you will see this as a drawback of the account. I prefer seeing it as a result.

[32] Cf. my 2004a, 13–16. – The recognition-apparatus account of properties can also give an answer to the question of what the arity (degree, adicity, *Stelligkeit*) of a relation is, and perhaps even tell us what a relation's places – whose number is the relation's arity – are (if such things exist). Usually, a relation's arity is explained via recourse to the way that relation is expressed in a language: "greater than" is a two-place relation because its expressions, e. g., "x is greater than y", have two argument places. But *why* do they have two argument places instead of, say, one or three? Could they perhaps have 17 places in a Martian language? Presumably not, but why is that? We could talk about the specific unsaturatedness (Frege's *Ungesättigtheit*) of the relation, but what would *that* mean?

Now consider some agent's recognition apparatus for "greater than". It must use information about some objects x and y to deliver a verdict on whether x is greater than y or not, so it will contain one 'input channel' for the information about x and another for the information about y. Thus, "greater than" is dyadic because it is a 'two-channel property'. You could feed information into these two channels which picks out the same object in both cases; thus, this object gets compared to itself. However, this recognition

grue Let's take a look at the property *grue* from Goodman's New Riddle of Induction (Goodman 1954; cf. von Kutschera 1975). All emeralds hitherto observed (call them $e_1, \ldots, e_n$) have been green. We feel quite secure in concluding inductively that *all* emeralds are green. However, all emeralds hitherto observed have also been *grue*, which is defined as follows (say, for material objects x):

$$x \text{ is grue} \quad :\Leftrightarrow \quad \begin{cases} x \text{ is green}, & \text{if } x \in \{e_1, \ldots, e_n\}; \\ x \text{ is blue}, & \text{otherwise.} \end{cases}$$

We wouldn't infer that all emeralds are grue, and we would think anyone irrational who did. Why? What is the relevant difference between the two properties *green* and *grue*?

(non)qualitative predicates One attempt to solve Goodman's paradox begins by distinguishing 'qualitative' from 'nonqualitative' predicates (cf. von Kutschera 1975, 59–60). The predicate "grue" is *nonqualitative* because its definition crucially involves reference to some specific objects, namely $e_1, \ldots, e_n$. Any sensible definition of "green", by contrast, would presumably *not* involve any such reference; rather, it would characterize greenness in some completely general, abstract, i.e., purely *qualitative*, way. Qualitative predicates like "green", so the proposal goes, are adequate for inductive inferences, nonqualitative ones like "grue" are not.

qualitativeness is language-relative But then a problem arises: Predicates (and the corresponding properties) are qualitative or nonqualitative only with respect to a particular language. In this case, that language is ordinary English, where "green" is rather basic and "grue" must be defined by taking recourse to $e_1, \ldots, e_n$, as exhibited above. However, there could be languages in which these roles are reversed. The English language could have evolved in such a way that "grue" and "bleen" – the latter defined in ordinary English by

$$x \text{ is bleen} \quad :\Leftrightarrow \quad \begin{cases} x \text{ is blue}, & \text{if } x \in \{e_1, \ldots, e_n\}; \\ x \text{ is green}, & \text{otherwise –} \end{cases}$$

were rather basic predicates, and "green" had to be defined using "grue", "bleen", and reference to $e_1, \ldots, e_n$:

$$x \text{ is green} \quad :\Leftrightarrow \quad \begin{cases} x \text{ is grue}, & \text{if } x \in \{e_1, \ldots, e_n\}; \\ x \text{ is bleen}, & \text{otherwise.} \end{cases}$$

Since the distinction between qualitative and nonqualitative predicates was intended to be an absolute one, not one which yielded different results depending on which

apparatus may as well be described in a different way. We can conceive of the input channels for the information characterizing x and y as *one* input channel for information about a different kind of object (the ordered pair (x, y) – whatever that is). That would make the relation a monadic one. Or else maybe this recognition apparatus utilizes specifications of the objects x and y via spatial coordinates x_1, x_2, x_3 and y_1, y_2, y_3. From that perspective, "greater than" would appear to be a 6-place relation. So, does considering recognition apparatuses just engender confusion instead of understanding? I don't think so, of course. I hope that if we stick to the most natural description of that recognition apparatus and take into account its function for the agent, the characterization of "greater than" as dyadic will emerge as the best one by far.

And now what are the 'places' of a relation? Are they the 'input channels' of a corresponding recognition apparatus? of a *token* apparatus?? Surely not: the identity and arity of a relation do not depend on token recognition apparatuses but only on types of these. Does a *type* of apparatus have input channels? I guess one would rather say that it *specifies* them. So maybe the places of a relation are abstractions from corresponding (types of) recognition apparatuses.

Tuples, e.g., ordered pairs, if they can be construed as objects at all, are perhaps abstractions from attempted 'acts' of recognition.

language you consider, it thus seems like this distinction doesn't help solving the riddle.

When I first heard this, I remained quite certain that the qualitativeness solution was on the right track. The idea that there might be a natural language in which "grue" was more basic than "green" seemed to me implausible in the extreme. It would be a very unnatural language in which unnatural properties like *grue* have more basic predicates than properties like *green*. Why is *green* more natural than *grue*, though, if we mustn't explain this by referring to the naturalness of languages but rather *vice versa*?

evolution of colour vision ...

Imagine a species of animal which is about to evolve some kind of colour vision. Which recognition capabilities are animals likely to evolve first: the ability to recognize *blue* and *green*, or the ability to recognize *grue* and *bleen*? They are astronomically likely to become able first to recognize whichever properties are *easier* to recognize, i.e., those whose corresponding recognition apparatuses are simpler.[33] It is vastly easier to evolve (or build) a mechanism for recognizing *green*, than for *grue*, because the former will be much simpler than the latter; and that is because a recognition apparatus for *grue* will have to include recognition apparatuses for *green*, for *blue*, and for the particular objects $e_1, \ldots, e_n$. So, barring circumstances which are astronomically unlikely, animals will evolve the capability of recognizing gruesome properties, if at all, then only much later than recognition capabilities for corresponding natural properties. (This is especially so if, as in our case, during the animals' evolution the majority of the emeralds $e_1, \ldots, e_n$ hasn't even been unearthed yet.)

... and of language

Now suppose further that these animals evolve language abilities. Again a predicate like "green" will be 'prior' to one like "grue", in the sense that it will crop up and stay in the animals' language with a much higher likelihood, and thus also much earlier. This is because (a) probably these animals can't recognize grueness at all, (b) if they can, grueness probably is terribly irrelevant for their practical concerns, compared with greenness, and (c) probably their first grasping of grueness will be by linguistic means anyway, i.e., they will first conceive that property by means of a definition based on a predicate for greenness. So a predicate for greenness will very likely be much more basic, much better entrenched, in their language than a predicate for grueness, if they have one at all. For these reasons, in the overwhelming majority of natural languages which contain predicates for both *green* and *grue*, the one for *grue* will be nonqualitative, or at least more so than the one for *green*.

(who knows what apparatuses are possible?)

Readers who are sceptical of evolutionary (or design-stance) arguments may object that I cannot be sure about the relative simplicities of these two types of recognition apparatus: perhaps the emeralds $e_1, \ldots, e_n$ have some special physical feature in common which any emeralds to be observed in the future do not share, and the physics of greenness and blueness are such that ... whatever, and so, finally, recognition apparatuses for *grue* might be physically possible which would actually be *simpler* than any possible recognition apparatus for *green*. Against that, I would bring forward two arguments: First, while this possibility is of course conceivable for a trained philosopher, still I think, physics being what it is (or what it seems to be to date), there isn't a shred of evidence for it; and I expect physicists and engineers to underwrite my position. My position *may* of course be wrong; in that case I would perhaps suggest we switch our inductive practices to prefer *grue* over *green*. Secondly, a recognition apparatus of this type wouldn't respect the intension of "grue" as

[33] It will also matter which sets of properties are more *relevant* for the animals' purposes/function. This aspect of the question, however, will point at the same conclusion as the ease-of-recognition aspect, so I will not consider it here.

outlined in its definition, so it wouldn't be a recognition apparatus for *grue* at all, but for some other, coextensional property.

primary/secondary qualities

Some properties of material objects, e.g., colours, smells, and tastes, seem to reside less in the objects themselves than in the eye (or nose or other sense organs) of the beholder. These are called *secondary* (or subjective) *qualities*. Properties like mass, shape, and movement, by contrast, seem to be independent of their being perceived by subjects; they are called *primary* (or objective) *qualities*. It is contentious, I believe, whether the intended difference is real; if so, whether the one or the other class of properties isn't empty; and if there are secondary qualities, whether they are nevertheless real properties of objects or rather figments of our perception processes, robust perceptual illusions, as it were. My approach to properties can shed some light on these matters.

Primary qualities are more natural properties than secondary ones because the latter are derived from the recognition apparatuses of a *particular* kind of agent, whereas the former would presumably be categorized in exactly the same way by *all* kinds of agents able to recognize them.[34] Bees, like us, can perceive colours, but the colours they perceive are different from ours.[35] For one thing, bees can see part of the ultraviolet spectrum, which is invisible to our unaided eyes. I suppose they also divide up the part of the spectrum they share with us in a different way from ours. Still, 'human-red' objects *are* (human-)red even when they don't seem so, because they will seem to be otherwise only when the conditions for human colour perception are unnormal. When normal conditions obtain, both in- and outside a given normal-sighted human agent, human-red objects *will* look red to her. These objects may also have a colour for bees, and humans can recognize their bee-colour by investigating the visual apparatuses of bees and how they would react to these objects. Therefore secondary qualities are real.

existence

Next: existence. How does the (non)existence of various kinds of entities figure into the dealings of agents? If unicorns exist then it is in principle possible that I might *run into* a unicorn one day (or it into me). Maybe I would have to live in the far past or on another planet for this to happen (if unicorns existed only then, or there); but if I did, it might happen. Maybe this is enough for a first approximation to the meaning of "existence": entities of a certain kind *exist* (in a certain 'place') iff it is in principle possible to *encounter* them (in that 'place').

What does this mean, e.g., for properties? We can understand this question in two ways: (a) What does it mean to say that some *particular* property exists, or doesn't? (b) What does it mean to say that properties *in general* do, or don't, exist? As to (b), certainly we can't stub our toes on properties, but that would be expecting too much 'reality' of them anyway. The nonexistence of properties in general would have to

[34] There might, however, be agents for whom a primary quality like, e.g., mass gives rise to qualitative differences comparable to our colours. Imagine giant extraterrestrial – indeed extraplanetary – organisms (consisting perhaps of interstellar gas) who can perceive distortions of space-time as they are caused by massive bodies like planets and suns. For such beings, distinctions between kinds of bodies like ordinary suns vs. neutron stars vs. black holes might be both natural and relevant. Beyond the merely quantitative differences in mass, there might be, for them, qualitative similarities and dissimilarities such that all black holes 'felt' the same, and all neutron stars 'felt' the same, but 'felt' different from black holes. These qualitative differences in mass would then be secondary qualities with respect to something that for us is a primary quality. (I'm sorry I couldn't think of any less outlandish natural categories for mass.)

[35] I am *not* talking about qualia here but about kinds of colour receptors and ways of processing the information they deliver.

mean that the property concept is incoherent, I suppose. I haven't yet seen arguments to that effect, and until I do I shall keep on believing in properties.[36]

Regarding (a), the first and most simple answer presenting itself is: a property exists iff it is *exemplified*. Running into an instantiation is very satisfying as a way of 'encountering' a property. Thus, the property of being a unicorn *doesn't* exist because there are no unicorns. However, denying that the property of being a unicorn exists strikes me as similar to denying that, say, Angela Merkel exists.[37] Sherlock Holmes really doesn't exist, in the sense that while we have a *concept* of a detective bearing that name, nobody actually matches that concept. The property of being a unicorn, by contrast – I think I know quite well which property that is, so, if the very idea of properties itself isn't incoherent then that property exists, because I have grasped it. I would thus tend to say that we 'encounter' properties by grasping what it would mean for something to instantiate those properties, i.e., by having or conceiving recognition apparatuses for them. If, on the other hand, recognition apparatuses for some putative property are impossible then there is no such property, but merely an (empty) concept of a property.[38]

justifying logics

One last remark. I believe my approach to ontology might also serve to deliver a (substantial, not too circular) foundation or justification for logic (*Logikbegründung*) – for various logics, actually. If you think about which kinds of inferences an agent should and shouldn't draw in thinking about the world in general, in order to be successful, then I suppose you would probably obtain classical logic. If, however, you think about the inferences an agent should draw when she is reflecting on the *beliefs of another agent*, α (or on the theorems of a formal system), and suppose she uses "p" to represent α's belief that p, then I suppose you would get a logic *without tertium non datur*: just because α doesn't believe p that doesn't mean that α believes $\neg$p. And there may be other areas of cognition, where still other logics may be appropriate.

References

Adams, Douglas. 1995. *The Hitch Hiker's Guide to the Galaxy: A Trilogy in Five Parts*. London: Heinemann.

Carle, Eric. 1969. *The Very Hungry Caterpillar*. Putnam/Hamish Hamilton.

Dennett, Daniel C. 1971. "Intentional Systems." *Journal of Philosophy* 68 (4): 87–106 (February 25). Reprinted in Dennett 1981, 3–22.

———. 1981. *Brainstorms: Philosophical Essays on Mind and Psychology*. Cambridge (Mass.)/London: MIT Press (A Bradford Book). Original edition 1978; sixth printing 1993.

———. 1991. "Real Patterns." *Journal of Philosophy* 88 (1): 27–51 (January). Reprinted in Dennett 1998, 95–120.

[36] I wouldn't be surprised, of course, to see some philosopher elaborating a concept of properties that put such stringent demands on them that they ceased to exist.

[37] Of course, as a good Fregean I abhor in any case the application of "existence" to singular terms.

[38] I have to be careful here about the structures of mathematics: though recognition apparatuses for some legitimate structures are physically impossible, I wouldn't want to deny the existence of these structures and thereby restrain mathematics to the realm of the constructive. I guess I would allow stretching the meaning of "possible". The corresponding structures then may be 'less real' than the constructive ones.

———. 1998. *Brainchildren: Essays on Designing Minds*. Cambridge (Mass.): MIT Press (A Bradford Book).

Goodman, Nelson. 1954. *Fact, Fiction, and Forecast*. London: University of London (Athlone Press).

Ladyman, James, and Don Ross (with David Spurrett & John Collier). 2007. *Every Thing Must Go: Metaphysics Naturalized*. Oxford/New York: Oxford University Press.

Lewis, David. 1993. "Many, but Almost One." In *Ontology, Causality and Mind: Essays in Honour of D. M. Armstrong*, edited by J. Bacon, K. Campbell, and L. Reinhardt, 23–38. Cambridge.

Pinker, Steven. 2007. *The Stuff of Thought: Language as a Window into Human Nature*. New York/Toronto/London/Dublin/Camberwell/New Delhi/Auckland/Johannesburg: Viking.

Spohn, Wolfgang. 2007. "How are Mathematical Objects Constituted? A Structuralist Answer." *Proceedings des 6. Kongresses der Gesellschaft für Analytische Philosophie, Berlin, 11.–14. September 2006*. 106–19.

Unger, Peter. 1980. "The Problem of the Many." *Midwest Studies in Philosophy* 5:411–67. Studies in Epistemology.

von Bülow, Christopher. 2003. Bauen auf Nichts? Warum ontologisches ‚Fundieren' nicht klappt. `www.uni-konstanz.de/FuF/Philo/Philosophie/philosophie/files/ontologie.pdf`.

———. 2004a. Auf Neuraths Schiff um die Welt! Vorschlag für eine ‚anbindende' Ontologie. `www.uni-konstanz.de/FuF/Philo/Philosophie/philosophie/files/neurath.pdf`.

———. 2004b. Dasselbe in Grün: Ein Streifzug durch den Zoo der Ontologien. `www.uni-konstanz.de/FuF/Philo/Philosophie/philosophie/files/ontozoo.pdf`.

———. 2006. Ontology by Linking. `www.uni-konstanz.de/FuF/Philo/Philosophie/philosophie/files/alles.pdf`.

von Kutschera, Franz. 1975. "Nelson Goodman: Das neue Rätsel der Induktion." In *Philosophie der Gegenwart III: Moore, Goodman, Quine, Ryle, Strawson, Austin*, edited by Josef Speck, Grundprobleme der großen Philosophen, 51–86. Göttingen: Vandenhoeck & Ruprecht.