

The virtues of non-reduction, even when reduction is a virtue.

1. Introduction.

Contemporary thinkers continue to be confused about what our attitudes towards reductionism should be. It is undeniable, as Tom Nagel notes, that “Reductionism has been a creative driving force in the history of modern science” (1998: 3). It’s also very clear that in engineering, medicine, manufacturing, and all manner of practical endeavors, people use straightforwardly reductionistic assumptions in figuring out how things work. At the same time, however, ‘reductionism’ has been a dirty word in the intellectual culture for several decades now. “ ‘Reduction’ and ‘reductionism,’ ” writes Jaegwon Kim “have become common epithets thrown at one’s critical targets to tarnish them with intellectual naiveté and backwardness” (1998:89). “ ‘Reductionism’ has come in some quarters to be used as a general term of insult and abuse,” writes Patricia Churchland. “Sometimes it is used as a synonym for ‘behaviorism’ ...or as a synonym for such diverse sins as ‘materialism,’ ‘bourgeois capitalism,’ ‘experimentalism,’ ‘vivisectionist,’ communism,’ ‘militarism,’ ‘sociobiology,’ and ‘atheism’ ” (1986:278).

This paper aims to reduce the confusion about what our proper attitudes toward reductionism should be. I will begin by saying briefly why reductive explanations are generally desirable. I will then spend the bulk of the paper laying out what I consider to be the best epistemic reasons for thinking that developing non-reductive accounts is also highly desirable. I aim to show that the best arguments for the desirability of reduction, and for the desirability of non-reduction, are rooted less in any deep metaphysical principles, and more in the general nuts and bolts of information storage in cognitive agents.

For the sake of argument, I will assume, in this paper that all properties may ultimately have reductive physical explanations. For my purposes, it is important not to be distracted by issues of whether and where reduction is possible. Although many people base their opposition to reductionism on arguments that physical reduction is impossible for certain things, I want to make the stronger argument that non-reductive representations are important, even if we live in a world where reduction is always possible and always desirable.

2. Reduction's desirability

“Reductionism” is an ordinary language folk concept that can be precisified in many different ways. Indeed, one of the reasons that reductionism engenders so much opposition is that different versions of it each have different sorts of opponents in various areas of scholarship, creating a large chorus of anti-reductionists who need not agree on what they are opposing. Most reductionists, nevertheless, embrace two central assumptions. One is the idea that to understand what something is, one must look (or it is best to look, or it is good to look) at the features of

which it is *composed*. The other is that understanding the nature of a large number of properties can only be done (or is best done or is done well) by identifying them as instances and arrangements of a *smaller* set of entities or properties. The most prototypical notions of reductionism contain both assumptions. It is interesting to note, however, that these notions are distinct and one could easily embrace one without the other. One might hold, for example, that it is important to identify properties with an arrangement of component types, but believe that a *large* cornucopia of different sorts of components exist in the world (so the number of sorts of micro-properties might actually be larger than the number of types of macro-properties). By contrast, one might be termed a “reductionist” even if one has little interest in the details of composition. Someone might be a sort of reductionist by claiming “all is water” or “all is part of the mind of God” without any interest in how wholes are constructed out of parts. Such views are reductionistic in holding that one can use a smaller number of facts about the world to account for a larger number. While holders of both assumptions (and holders of many others) have been called “reductionists,” in this essay I will largely concentrate on the prototypical idea of reductionism in which both these assumptions are held.

Why would a scientist be inclined toward reductionism? That attractiveness of reduction, for many, begins with the belief that it is simply a true fact about our world that most things *just are* nothing more than a certain arrangement of other things. But this, by itself, does not tell us why a knowledge of smaller components should provide us with a much improved *understanding* of larger objects and properties. When one begins looking, however, it is not too difficult to see that there are numerous ways in which our knowledge of the world is improved by being aware of the components of which things are made. To begin with, one straightforwardly knows *more facts* about an entity or property when one knows facts about its micro structure as well as facts about its macro structure. One knows more about diamonds when one knows that diamonds are made of strings of carbon atoms. And knowing something’s micro structure often enables one to make inferences about the presence of macro structural features that one didn’t know about when one didn’t know the micro structure. Knowing that diamonds are made from carbon atoms allows us to deduce that they can be created by subjecting graphite (another arrangement of carbon atoms) to the right temperature and pressure. Knowing the micro structure of an entity often tells us why macro structural laws involving these entities have exceptions under certain circumstances (e.g., a diamond can be scratched in special circumstances) (see Hempel 1966). Now one could just *add* these exceptions to our repertoire of knowledge about the item, instead of inferring these exceptions from knowledge of the micro structure. But this points to another advantage of reduction. Adding numerous “exception clauses” to every known macro law would take up a great deal of storage space in whatever medium one is storing information about the

world. When these exceptions are derivable from a knowledge of the micro structure, however, one doesn't need any *new* storage space to house that information. Indeed, where an item is reductively identifiable with its component parts, knowing those part should enable you to derive all of the properties of that item. (Intrinsic properties are derived by combining the parts. Relational properties are derivable by showing how the presence of certain external properties, along with this combination of parts, creates certain properties). If the only properties that agents need to know of to represent *all* the world's properties are physics ones, then agents do not have to learn or store an enormous array of diverse non-physical representations of the world's properties.¹ Reduction, in other words, enables agents to save an enormous amount of representational resources. Finally, when one reduces items in different realms to a common (say physical) vocabulary, we make it more likely that information gained in one realm can provide us with information about another realm. Discoveries about diamond carbon might tell us something about certain properties of organic carbon compounds, which, in turn, might tell us something about methane gas (CH₄). Whatever difficulties we may have in coming up with reductive accounts of various entities and properties, we must admit that there are tremendous benefits that result from making reductive identifications when we know how to do so.

3. Why not be a reductionist?

3.1 Epistemic and non-epistemic virtues

With all the advantages reduction provides to researchers, what reasons could people have for not wanting to spend a great deal of our research efforts trying to make reductive identifications and to draw inferences from these? Why would reduction so often be such a dirty word in academia? There are surely numerous different reasons for the hostility to reduction. Many of these have little to do with any lack of epistemic virtue. Indeed, I suspect that for many humanities scholars it is the *presence* of certain epistemic virtues of reduction that lead to reduction's disfavor. In a number of popular psychology books, Mihaly Csikszentmihalyi argues that happiness in life is the result of finding the right compromise between excess anxiety and excess boredom. We reduce anxiety by trying to make the world as familiar and predictable as possible. We reduce boredom by trying to increase novelty and *unpredictability* in our lives. For some, the problem with reduction is not that it inhibits understanding, but that the *increased* understanding enabled by reduction removes some of the exoticness that helps decrease boredom. At times, improved understanding certainly can make things seem less interesting. Most magic tricks, for example, are less exciting to watch when one has a knowledge of how the trick is done. Humanities disciplines contain their fair share of romantics and thrill seekers who take it as their task to celebrate mysteriousness and unpredictability. For such scholars, the increased predictability and understanding that reduction affords runs contrary to what they see

as their mission -- making human behavior seem more interesting and exciting.

While some scholars object to reductive representations on more or less aesthetic grounds, others object on moral grounds. Reduction gives us an improved understanding of properties. Improved understanding often gives us improved predictive abilities; and improved prediction provides opportunities for improved control. There are many to whom improved control, over friends and colleagues, over citizens, over other nations, or over nature, is repugnant. The question of whether, and to what extent increased control is a good thing is an important one. The question of when reductively representing something makes it more or less interesting is also worth asking. Note, however, that both of these objections to reduction are premised on the implicit assumption that our *understanding* of a realm *is* improved when we gain reductive knowledge of it. The objection these sorts of opponents have to reduction is that *increases in epistemic virtues* are outweighed by supposed losses elsewhere. I don't intend to debate this issue here. What I am interested in is whether we can *epistemically* improve our picture of the world in ways other than gaining increasingly detailed reductive information.

I suggest there are, indeed, ways in which developing non-reductive representations enhances our understanding of the world. The fundamental epistemic challenge of agents is to gain more information about the world using fewer resources. Making reductive identifications improves our epistemic situation by enabling us to store more information about the world using fewer of certain types of resources. Yet there are other ways of representing the world that save other types of cognitive resources. Developing general non-reductive *subsumptive* categorization schemes, enables us to communicate more quickly, access information more quickly, save computational time and energy, and save memory space by keeping us from having to separately record similar information over and over again. Non-reductive subsumptive categories can also save memory space by avoiding storing certain details. Such representations can usually be *developed* more rapidly and at less cost than reductive representations. Let us discuss these considerations below.

3.2 Saving representational resources in non-reductive ways

If giving reductive physical descriptions enables agents to be able to derive all the properties that any given entity has, while limiting the number of properties one ever needs to know, it's hard to see how any other kinds of representations can have epistemic virtues that can compete with reductive ones. Nevertheless, there are some disadvantages to reductive representations. Consider a prototypical reductive explanation . To explain an event, we start with a set of basic entities (like basic physical particles), a description of their exact location in space at some time, and a set of laws detailing how such entities interact. At time T1, the entities will be interacting in a certain manner, in accordance with these basic laws. This means that at time T2, the entities

will have become arranged, as a result of these lawful interactions, into some new configuration. In this new configuration, they will continue to interact, in accordance with these laws -- which means that at time T3 they will be arranged in yet another configuration. We continue applying the laws to these configurations until we see exactly why the configuration we are trying to explain came to be structured as it is.

It should be immediately clear what the disadvantages and costs of such reductive explanations are. To begin with, it takes a great deal of *time to compute* how the collection of entities will be arranged at some future time. It is usually advantageous to real agents, however, to access information about the world in as little time as possible. In addition, performing such computations will take a great deal of *energy*, whatever the medium of computation used. Agents benefit when they can minimize the amount of energy used for getting the information about the world that they need. The classical reductionist strategy of reducing one's ontology to a minimum and computing on the basis of these, certainly has large costs for any agent.

3.2.1 Large memory structures vs. subsumption

How could such costs be reduced, while still retaining most of the information that a reductive account would give you? One way would be to enable agents to access information about the world in a different manner than computing from initial states. In a prototypical reductive understanding, an agent applies cognitive procedures to a set of representations in order to produce, step-by-step, a new set of representations (paralleling, in a way, the processes by which laws transform physical entities to produce new end-states). But there are less time and energy intensive ways to infer that a certain set of properties must be present. One way would be to *store* detailed "bit-map" information about what the world was like at T1, T2...etc. in some sort of "memory." Such a representational scheme, like prototypical reductive accounts, could still give us low-level detailed accounts about what physical particles were where. In this scheme, however, one infers what is where not by creating a new representation from prior information, but merely by *looking up* what exists where. This set-up could give agents detailed information about what to expect from the world, while potentially saving a great deal of the computational energy and time that classical reductive accounts take.

But storing the locations of huge numbers of low-level particles in a giant bit-map memory, has its own disadvantages and costs. One of the biggest, of course, is that any real bit-map storing agent would very rapidly run out of memory *space* for storing much detailed information. It would also likely take such an agent a huge amount of *time* to find the information needed for undertaking various tasks. If all an agent had was the "raw data" about what physical particles were where, it could only find the particular particle *arrangement* it was looking for through an enormously time consuming blind search.

How could matters be improved, while still storing all this information and still minimizing energy use? One way of making representation more efficient is to make use of what might be called “subsumption.” Subsumption is a general term for what's done when an information processor takes a larger set of signals (sensory perceptions, motor procedure signals, symbols, etc.) and converts it to a smaller set of signals. "Cluster subsumption" is what we can call a large number of different items, connected to each other in space or time, coming to be represented by a single concept. "Type subsumption" is when a number of discernibly different properties (wherever they may be in space and time) are represented by the same concept. These sorts of subsumption are often combined, and both can save a great many resources for agents. A cluster concept can save time for agents, for example, when they communicate using a label that represents a concatenated cluster, rather than explicitly naming each detail of the cluster. Consider the amount of time saved when an agent makes use of "cluster subsumption" to communicate to an agent that a particular snake, Sammy, was present. Without a subsumptive concept it might take hours for a speaker to speak and the listener to hear about the entire array of basic physical properties that compose Sammy the snake's body. Using a subsumptive label, however, Agent #1 could take a detailed reductive representation of Sammy, utter the words “Sammy the snake” and enable Agent #2 to create the same reductive representation using some type of “decoding” device. This act of communication takes a fraction of the resources required to create the representation in another agent one physical bit at a time. Subsumptive representations of clusters also save space, in that the agent can represent things in memory by using only that concept, rather than the entire array of entities and properties that that concept represents.

But how is it that having a cluster concept really saves memory space, since taking up the space in memory to note that, "henceforth, a certain arrangement will be represented by a certain node," must itself take up new memory resources? The answer is that if an agent comes to have a concept that represents a *type* of cluster, the initial outlay of memory resources will usually "pay for itself" in memory savings numerous times over. An agent acquires a “type subsumptive” concept by utilizing a symbol or signal that can be used to mark the presence, not only of a particular entity or arrangement, but also of other arrangements identical or highly similar to the first.² Memory is saved because the complicated arrangement represented by the concept need only be stored in *a single place in memory and referred back to*, instead of having to represent that whole space-hogging representation over and over again in memory. For example, an agent can mark each place that carbon is present in an area by noting that fact with some sort of *carbon symbol* instead of storing information about the complex arrangement of electrons, protons, and neutrons that make up the carbon atoms in each particular place. An agent could compactly store all the instances of the presence of the subatomic arrangement with a subsuming carbon symbol,

and extract the details that this subsuming symbol represents only when she needs them. The more stringently similar an agent requires things to be in order to merit the same symbol/signal, the more we retain all information we can have about what's in the world. If two balls must be composed of identical particle arrangements to be labeled "p-ball," then anytime an agent uses the concept "p-ball" we know exactly what arrangement is there. Using type concepts of any sort does, of course, mean that some information will usually be lost, since many different slight variations on a type of arrangement of entities are all represented by the *same* small set of signals (with no way of telling which variant it is if one tries to "decode" the compressed signal) But the memory savings that occur by recording what is there using a previously stored concept, rather than recording the entire arrangement or one of its slight variants over and over in memory can make this information loss well worth the cost.

And agents can save even more resources by using concepts that represent clusters of clusters (which I'll discuss now) and types of types (which I'll discuss in later sections). Just as one can create a "carbon" node for a certain arrangement of subatomic particles, one can create a "carbon dioxide" node for ease in storing and communicating information, rather than representing the chemical structure every time. One can, in turn, use an "atmosphere" node to represent a certain mixture of nitrogen, oxygen, and carbon dioxide. Such hierarchies of clusters of subsumptive types enable an agent to continually save ever more memory space, as well as communication time.

Given that one might be able to save large amounts of memory by recording something using a concept that represents a cluster of clusters, one might wonder what use the middle level layers in a hierarchy are. Why not create an "atmosphere" concept that directly codes for a certain complex arrangement of subatomic components and save space by eliminating the mid-level concepts? While the presence of mid-level nodes do take up memory resources, they are extremely helpful in other ways. First, given that the simpler mid-level clusters might be plentiful in the world (and the complex "cluster of cluster of cluster" top level clusters might not be), we don't want to forego all the savings memory savings we might gain from inventing and using these mid level terms, as opposed to just using the lowest and highest ones. (E.g., we might want to retain "carbon dioxide" because of its use in contexts other than describing part of the cluster making up "atmospheric air." Next, having layers of nodes for mid-level types enables an agent to save memory search time, by specifying in what "box-within-a-box" the particular information one is looking for is. With only a top and bottom layer, we are no better off than we are doing an exhaustive search of a bit-map. Finally, the existence of middle layers enables "decodings" to save energy, time, and space by only decoding down to the level of information an agent is specifically searching for. If an agent needs to know only whether carbon dioxide is present in a particular area, one can search his memory until this fact is found, without having to "unpack"

the inner structure of carbon dioxide, or the more detailed inner structure of carbon. Agents who have created hierarchical layers of clusters of clusters of subsumed types can save a great deal more of certain types of cognitive resources than agents who represent things using vast bit maps, or initial conditions and computation.

3.2.2 Subsumption and inference

The use of subsumptive concepts can help save resources in other ways. Up until now, we've been discussing how putting large organized structures in memory, rather than calculating what is there, step-by-step, saves energy and time resources -- but at the cost of taking up memory space. But trying to store lots of information about the world in memory (using bit-maps or subsumptive concepts) is problematic for several reasons. Suppose that what an agent wants to know about is a state of affairs that will obtain in the future. There are few ways that an agent can ever *obtain and store* particular information about the future (though information about universal general properties may be stored). And even if an agent managed to obtain information about the future, there will always be limits to *how much* information an agent can store in memory. Even if organized sets of information are more parsimoniously stored in memory with subsumption instead of bit-maps, they still take up memory space. Memory storage will limit not only what can be stored within a concept, but also how many concepts an agent can store. Such limits means that it is better for agents to sometimes use *inferences to create rather than store* representations. Instead of storing that structure S [composed of l-m-n] will be present at T3, an agent could save memory space by inferring this, step-by-step, by from a set of rules combined with input information. The input would most often come from perception, but could also be the end result of a previous application of rules to an input set (the "then..." of a previously used if-then rule). We've noted that one can make inferences using reductive representations, but inferences can also be made using subsumptive representations.

Making inferences using subsumptive concepts can save more memory than subsumption without inferencing, or than inferencing without subsumption. Few sorts of agents, for example, could possibly store information about future gunpowder explosions (even if they could somehow obtain that information, and even if they stored it with a subsumptive labels). On the other hand, an agent might be able to infer that an explosion will take place by using low level physical rules. These rules could show that when a mixture of potassium nitrate, carbon, and sulfur comes into contact with a flame, the particles will fly apart with tremendous force. But making inferences using such low-level entities and rules either takes a large amount of time, or a large number of separate information processors working in parallel -- a large expenditure of resources in any case. Alternatively, an agent could make inferences based on laws involving

connections between things represented by high level subsumptive concepts. Agents could save both memory and inferencing resources by storing a rule like: “Whenever ‘black powder’ (potassium nitrate, carbon, and sulfur) comes in contact with a flame, an explosion results.” The rule could be applied whenever an agent perceived black powder and flame were coming into contact. She could then use these high level rules and perceptions to *infer* where and when explosions would occur, rather than *storing in memory* where these events occur. In general, subsumptive inferencing saves memory space by using subsumptive labels to represent the beginning and end stages of a process (and perhaps a few stages in between) rather than representing all the low level details of the initial, intervening, and end-state details of an event. Subsumptive inferencing saves calculating time and energy by using high level rules that connect initial steps to end or middle steps in ‘one-fell-swoop’ leaps, rather than by tiny, step by step calculation with low level rules. Subsumption can save resources by replacing inferencing with memorized representations, but it can also save resources by using subsumptive representations to make inferences.

The question of whether or not various sciences should concentrate largely on reduction, ultimately comes down to whether or not reductive representations are always or usually the best, most efficient ways of representing what is in the world. If non-reductive representations are sometime better for achieving epistemic goals, then it will be important for agents to devote some scholarly resources to developing sciences that utilize non-reductive terms. From what’s been said up to this point, there should be no question that agents benefit epistemically by the development of subsumptive terms. Using such terms allows agents to save large amounts of memory space by keeping us from needlessly storing highly similar details over and over again. They save energy and time by keeping agents from having to make calculations on the basis of every entity in a large cluster. At the same time, such terms, hierarchically constructed out of stringent lower-level types, can contain information about what is there at the lowest levels. Developing stringent subsumptive terms can allow agents to represent fine grained details about what is in the world, while saving time and space. Saving time and space, of course, preserves resources for recording more of what is in the world. When scholars discover generalizations about the world which they describe using subsumptive categories, they are improving our understanding of the world in a different way than reduction improves it -- but it is an improvement nonetheless.³ Even if reduction is always an in-principle possibility, and even if reduction improves our understanding of the world in numerous ways, there are other ways to improve our understanding besides trying to develop reductive representations of entities properties and laws.

3.3 Less-stringent subsumption.

We've seen that a big advantage subsumptive representations can have over reductive ones is that they enable an agent to save memory space while also saving inferencing time. A natural question is whether there are ways to gain further advantages through further memory savings. In section 3.2.1, we described how an agent could save a great deal of memory space while sacrificing only a little information by using a stringent subsumptive concept. If we want to increase the memory savings that type subsumptions can provide, one way to do it is by allowing type concepts to be *less stringent* regarding the range of things the concept applies to. The stringent label "glucose," for example, refers to a certain arrangement of carbon, hydrogen and oxygen molecules. The less stringent label "sugar" refers to both monosaccharides like glucose and disaccharides like sucrose and lactose. The even more tolerant "type of type" concept "carbohydrate" refers to sugars, starches, or any molecular arrangement that follows the general $C_x(H_2O)_x$ schematic arrangement. Using less stringent terms can save memory space in a number of different ways. Take a concept mentally defined as applying to those things that have each of a long list of features (e.g. a dog has four legs, fur, a snout, etc.) Each of those features in turn might be defined by their own list of features (a snout is a protuberance from the face, with a nose on the end and a mouth running its length...). One can make a concept less stringent by removing some of its satisfaction conditions – or removing some of the conditions of the conditions.) The less stringent "rectangle" concept, for example, can be made by removing the "equal side-length" requirement from the "square" concept. A concept that is less stringent in this way takes up less memory space in its original definition storage. A less stringent concept can also take up less initial memory storage space, if the agent *simply does not store* what the lower level defining features of defining features. An agent, for example might have a less stringent concept of "brick wall" by defining "brick" as rectangular cube of burned clay but storing very little information about what it means to be a unit of clay. By making "definition trees" of smaller width or depth, an agent can save a lot of memory space.

But a less stringent concept, however, doesn't necessarily take up less storage space with an initial definition than a more stringent one (e.g., why should we suppose the mental mechanisms for representing "blue" must take up less space than those for representing "indigo"?). And some less stringent terms likely take up *more* memory space for their definitions than more stringent ones. A "strawberry daiquiri" is a drink made from two ounces of white rum, half a lime, one bar spoon of powdered sugar and certain amount of strawberry juice. The less stringent "daiquiri," on the other hand, is made with rum, lime, sugar and *fruit juice of various kinds*. Whether an agent represents "fruit juice" with a long set of disjunctions or with a set of satisfaction conditions specifying what counts as a "fruit juice," storing a definition for "daiquiri" probably takes up *more* initial memory space than storing a definition of "strawberry daiquiri."

But even though it may take more memory space to *initially define* a less stringent type, less stringent types (including types of types) can still ultimately enable agents to represent things in the world using far less storage space than reductive or stringent subsumptive representations do. All type concepts save memory space by enabling an agent to avoid having to create a new description “from scratch” of every molecular arrangement we might want to represent. But the more stringently an agent defines her type concepts, the more memory space she has to take up *defining different type concepts* for slightly different molecular arrangements. (An agent, for example, could conceivably have different concepts for five different sub-types of fructose). The use of less stringent types, on the other hand, means more things can be described using fewer concepts.

We use less memory space by *not* taking up the space to define new concepts. Consider the less stringent type “hammer.” This concept is such that numerous different arrangements count as examples of a hammer. The less stringent the concept is, the more things we can store in memory, with the description “hammer there,” (referring back to a singly-stored hammer definition), and the fewer hammer-resembling arrangements we have to describe by naming the piece-by-piece arrangements of component parts. A piece-by-piece description not only takes up lots of memory space, but also takes up more memory space each time it must be given. Now another way we could take up less memory space than using piece-by-piece descriptions is to create lots of different stringent concepts for different types of hammers. But we can use even less memory space by *not* creating these new concepts, and making due with the more generic, less stringent concept “hammer.”⁴ And if an agent seeks to be more specific about the exact types of hammers that are in various places (ballpeen hammer, claw hammer, rubber hammer), she can do it at a minimal cost in space by *adding features to an already-defined less stringent concept*, rather than having to create numerous new definitions from scratch. To summarize: using less stringent representations enables agents to get around the problem of having to use too much memory space that can be caused by either a) representing each of a large number of instances of a type with a separate detailed bit-map, or b) inventing numerous new type labels to describe recurring “clusters of clusters.” The second way of representing things saves more memory space than the first, but the use of tolerant non-stringent labels saves more memory than either.

There’s another way, as well, that non-stringent labels save memory space. Suppose that in the interest of more accurate detailed representation, agents do create more stringent labels in addition to the non-stringent ones. Agents possessing the concept “sugar” might sometimes want to use more specific concepts like “galactose,” “glucose,” and “maltose.” We discussed in section, 3.2.2, however, that we could save various amounts of space by not having everything stored in *memory*, but by having the presence of some things *inferred* on the basis of if-then generalizations. Now if we possessed only stringent concepts like “glucose,” “fructose,” or

“maltose,” we would have to separately store numerous generalizations about each of these concepts, instead of being able to store more compact generalization about “sugars.” Being able to store generalizations about “sugars” can save a large amount of space, even if agents do have some generalizations that pertain only to “glucose.” Reducing the number of generalizations we need for understanding the world is another space-saving advantage of creating non-stringent concepts.

3.3.1 Costs of less-stringent subsumptive concepts

In considering the virtues of non-stringent concepts, one should not ignore the fact that a lack of stringency comes at a cost. If the cost of stringent representations is using up memory storage space to store an array of very specific concepts for describing the world, the cost of non-stringent concepts is the agent possessing less information about each thing he represents with that concept. The concept “hammer” tells an agent less about what is there than the concept “ballpeen hammer” and the concept “tool” tells an agent less than “hammer.” Why should the space savings engendered by using non-stringent concepts be any kind of epistemic virtue when it comes at a cost of the agent knowing less about what is present at a given space and time?

The answer is that when memory space (or time, as we’ll discuss below) is freed up by *not* taking up resources to represent certain information, an agent has more resources available for possessing *other* information. The epistemic advantage in not recording certain information lies in this enhanced capacity for possessing other information. All real agents can only acquire a finite amount of information about the world. That means that tradeoffs must always be made regarding which information an agent will acquire. The tradeoffs can be made on the basis of agents' choices, or by impersonal forces acting on the agent, or her forebears. Whether the extra information or the extra resource savings is more important to the agent depends on the agent's various other goals. A highly detailed molecular representation of a particular mummy in the British museum, for example, might take more memory space than undetailed representations of *all* the museum's other items. Different agents might consider each choice the more important one, on the basis of different non-epistemic goals. But it’s not clear that on the basis of *epistemic* goals alone (whatever this might mean) that more detailed information is more important. Having *more information about fewer things* is not necessarily epistemically better than having *less information about more things*. Nothing in our concept of "epistemic advantage" or "understanding" tells us which is more epistemically advantageous. Advocates of knowing more things might debate the advocates of knowing more details about which has been of more *practical* importance to us. Those skeptical of the importance of fine-grained knowledge can point to the enormous realms of engineering marvels that humans have been able to create, all the while knowing very little about the lower level physical or chemical level micro components

involved. But more fine-grained details vs. information about more things is truly an apples and oranges comparison. Which is better for the agent depends on an agents' needs and goals. And even if an agent's needs and goals are known, this won't decide the matter, since the needs and goals of the agent can change over time, and can change as technology changes the range of options an agent has available. What a shepherd needs to know about his sheep differs depending on whether they are healthy or diseased. Often, an increased knowledge of micro structural features can enhance our abilities in a realm, when there was little initial reason to expect it would. Only a decade or so ago, researchers in counseling psychology were fond of teasing their neurologically oriented colleagues about the irrelevance of their work for treating patients who needed counseling. A short while later, a majority of counselors had come to depend heavily on other scholars' knowledge of the serotonin uptake chemistry of substances like prozac.

What should be said about less stringent subsumptive concepts is that their ability to save memory space for an agent does provide an agent some epistemic advantages. These advantages come at the cost of a certain type of information loss, but for some agents, what is gained by these costs are well worth it. These advantages should be considered when contemplating whether it would be beneficial for a science to try to develop higher level concepts or to strive to be as reductive as possible.

3.4 Acquiring subsumptive concepts.

When looking at the costs and benefits of working to develop better reductive representations vs. working to develop other types of representations, we should consider another set of advantage that developing less-stringent concepts have: the ease and cost of acquiring them. Up till now, we have been mostly discussing the advantages and disadvantages of different ways of storing information *that has already been acquired*. But we should also consider the costs involved in *coming to learn* what is there in the first place. If the costs of acquiring reductive representations are much higher than those of acquiring subsumptive representations, then scholars who are not focused on developing costly reductive representations will have more resources to devote to representing more of a domain (albeit at a lower level of detail). And as we've just said, representing more of a domain is one kind of epistemic advantage.

In order to acquire a representation that a member of a certain category will be present at a certain type and place, an agent needs to go through two stages. One stage is learning what it is *to be a member* of a certain type of category. The other is getting enough information about the world to be able to tell that a particular example of that category is present. Doing either takes time and resources. Let's imagine that some agents have no difficulty directly perceiving low level micro features. Let's assume, further, that one easy way for an agent G to acquire a representation of a cluster of properties is for another agent D, to make the initial discovery, then

to *tell* G what he has found. Now, by parallel arguments to the ones about memory space, subsumptive representations should be able to be acquired using less time and fewer resources than reductive ones. Initially, it takes a little extra time for agents to learn a subsumptive concept (as opposed to only representing things as arrangements of reductive categories they already have -- e.g., learning "carbon molecule" in addition to being able to store "14 protons and neutrons surrounded by 14 electrons"). Once agents have learned this, however, there's a tremendous time saving. Once an agent learns what the features making up the concept carbon are, the agents can learn that carbon is present at a particular place instantaneously -- just be another agent telling him so. The existence of a concept means that the discoverer does not have to take the time to relay, piece-by-piece where every particle in that arrangement is. That information is already in the receiving agent's memory from when the concept was learned. Just as subsumption makes it easier to quickly transmit lots of information (as in our earlier "Sammy the snake is present" example), using subsumptive concepts makes receiving/acquiring information much more time efficient.

Similarly, if one learns that certain perceptual cues are a reliable indicator of the set of particle arrangements associated with a certain concept (e.g., the color spectrum associated with burning carbon), then the presence of these perceptual cues can enable an agent to very quickly learn that that arrangement is present. Again, having learned the concept itself enables the agent to rapidly learn that an arrangement is present, without her having to perceive each part of the arrangement, piece-by-piece. This can enable quite a lot of time saving, considering one saves time *each time* an agent learns through perception or communication that an example of that arrangement is present.

We obviously have perceptual and communication time savings with less-stringent subsumptive concepts, as well. An agent hearing "Teddy bear on the living room floor" certainly knows what to expect much quicker than if that agent had had to hear a reductive physical description of what was there. And the process of learning less-stringent concepts would seem to have time savings that parallel the space savings such concepts enable. It's plausible that, other things equal, agents can learn concepts with fewer satisfaction conditions than concepts for which the agent is required to learn a much longer list of conditions. It's also plausible that concepts whose "definition trees" terminates very early -- concepts for which the agent doesn't store the components of lower levels components -- can be learned more rapidly.⁵ Even if such less-stringent concepts do not take any less time to learn, the existence of non-stringent concepts enables agent to get by *fewer* concepts for describing similar scenarios. Learning fewer concepts (like sugar) takes less time than learning more concepts (like sucrose, lactose, and trehalose).

Besides being learned by perception and communication, concepts can be learned by inference from prior information. Just as space can be saved by having inferencing rules that link

high level concepts to other high level concepts, instead of having to represent the entire cluster, time can also be saved by using inferences that link together high level concepts. Instead of a long physical derivation from a large set of initial conditions to a resulting cluster, an agent might come to know what is where by means of *a few* if-then rules that link together high level subsumptive concepts (stringent or non-stringent). Such rules could allow agents to rapidly calculate an end result without having to go through numerous small steps.

But the biggest resource savings that could come from utilizing subsumptive concepts are come to agents that have an *inability* to directly perceive instances of lower level micro properties. Without being able to directly perceive sub-microscopic particles, consider what an agent has to do in order to know that a particular type of low-level cluster is present in a particular place. An agent might infer from prior assumptions that a certain arrangement of entities is present. To be able to do this, an agent has to start with information about the types of composing entities involved (e.g., electrons, protons), the properties of these entities, (e.g., attraction, repulsion), the laws regarding what these entities and properties do, and specific information about which entities are where. But how can an agent get any of that starting information? The only way is through using abduction, or inference to the best explanation, to make guesses about what must be there, based on the observations. But to do abduction properly in physics requires a great deal of time and effort. As with any abduction, a myriad of perceptual observations must be recorded and agents then come up with the model that best fits all the observations -- the larger the set the better. But this takes time -- far more time than direct perception. And many things make successful abduction more difficult for physics than for other sciences. Abduction is best done when one can do fine-grained experimental manipulations to differentiate between different possible structures that might be responsible for the same observations. But it takes a tremendous amount of energy, technical equipment, and knowledge to do any fine-grained experimental manipulation of the micro physical world. And the working of micro physical world are very different from the workings of things in the everyday familiar world. This makes it more difficult to make educated guesses about physical structure based on analogies with things we are familiar with. The micro-physical world's difference from familiar structures, in other words, makes it difficult to know where to begin in making hypotheses about what accounts for our observations. Our unfamiliarity with the microscopic world hampers us in the opposite way, as well. We don't have as many clear *constraints* on which structures it is plausible to postulate as we have in realms for which we have first hand sensory experience. We might well be able to think of a *multitude* of theoretical structures possibly existing in that complete invisible world, all of which are compatible with all our observations. A number of things, then, make it extremely difficult, expensive, and time consuming to acquire reductive

physical explanations. If agents took the same amount of time and resources to acquire subsumptive representations that don't have these obstacles to their acquisition (through inferencing, direct perception, communication, and abduction) it's likely that they could represent far larger areas of space and time than one could through attempting to acquire reductive representations. The fewer resources it takes to acquire subsumptive representations, then, gives them yet another advantage over reductive representations.

4. Reduction vs. subsumption: concluding remarks.

So what should our attitudes toward reduction be? The basic epistemic goals of agents and the nuts and bolts of information storage leads us to conclude that developing reductive representations gives agents certain advantages -- but that developing non-reductive subsumptive representations provides epistemic advantages as well. Let's briefly summarize the advantages of each.

If all properties are reducible to physics, as we have been assuming in this paper, then reductive representation gives us a way of describing all the universe's properties. Indeed, it's the *only* way to represent all of them, since without reduction, certain micro properties could never be represented. At the same time, reduction can be an important way of saving memory storage space. Representing the world using reductive concepts minimizes the number of different concepts that must be learned. And an enormous amounts of storage space can be saved when agents can *infer* what must be there from a small reductive set of initial conditions and rules, rather than storing each of the numerous situations that could be so inferred.

But there are other ways of storing more information and saving resources besides reduction. Speakers can communicate what is where in less time using subsumptive rather than reductive representations. Hearers can be made aware in less time. One can make inferences that take up less time and energy by using rules connecting high level subsumptive concepts, instead of relying on the recursive application of reductive to initial conditions. Agents can also reduce inferencing time and energy by storing subsumptive representations in memory. But subsumption could also save memory storage space, in that an agent need only to record the complex arrangement of properties that defines a concept in a single place in memory and then using only that concept to represent that type of arrangement, rather than take up memory space spelling out the details of the arrangement every time.

And an agent can save even more resources when less stringent representations are used. Less stringent representations don't have to store all the details of what is present in a cluster in a "representation tree" in memory. Learning these concepts can take less time, given that the "trees" defining the concepts must have fewer features, or have these defining features

themselves defined by fewer features. *Discovering* that an example of a less stringent category is present can take less time for similar reasons (especially given that the lowest-level reductive details are often very difficult and time consuming to uncover, given their extremely small size). And agents can save resources by representing more of the world using fewer concepts. And by recording the presence of things using broader, less stringent concepts, an agent uses up less space in memory than if she had to invent and store new concepts for all of these different types. The space and time savings enabled by using (more and less stringent) subsumptive concepts can give agents the time and space to acquire and possess vast amounts of information that they wouldn't have been able to possess had the same amount of time and space been devoted to developing reductive representations.

Developing scholarly disciplines that help agents construct reductive representations, then, can give agents important epistemic advantages. But developing more and less stringent subsumptive representations can also give agents important advantages in learning, communicating, and storing information about what is there in the world. The rational way of proceeding for most agents needing to know as much about the world, using as few resources as possible, then, is to aim at developing *different types of representations* which enhance our knowledge in different ways. Some time and energy should go to trying to reductively understand the low level structural arrangements out of which the world is created. But it is also important to develop ways of conceptualizing the world using subsumptive structures that also save cognitive resources in other ways. Agents will benefit most by developing different kinds of sciences, which represent the world using both reductive representations and various sorts of subsumptive representations.

And the idea of dividing the epistemic labor between developing reductive and subsumptive representations makes even more sense, given that we live in communities where different agents have different interests and talents. If an individual's knowledge is improved by using mixed representational strategies, surely our *collective* knowledge will be improved by having different ways of representing the world. People who have a special interest in and talent for comprehending the world in terms of broad non-stringent concepts (like those used in psychology or anthropology) should be encouraged to investigate hypotheses couched in these terms, rather encouraging everyone to try add to our reductive physics knowledge. This is, of course, what we currently do, when we divide the academic world into different subject matters. We've seen now that there are important epistemic reasons for doing so, despite the many benefits of reductionism.

Many scholars oppose reductionism and speak of the importance of developing non-reductive representations. Their opposition to reduction is often on the grounds of the alleged

impossibility of physical reduction. Moral and aesthetic considerations also likely motivate many anti-reductionists. I have argued here, however, that epistemic considerations regarding information storage alone can show us why non-reductive representations are important, regardless of whatever deficiencies or advantages reduction may have. Even in a world where it may be possible to reduce everything to physics, and where we gain certain benefits by trying to do so, our epistemic situation is also improved by having large numbers of our scholars working diligently on developing representations that are non-reductive. Some of our sciences *should* center around trying to reductively representing the world. But it is also important for agents to develop sciences that center around non-reductive concepts. We have a better understanding of the world when we have both.

Notes

1. I am using the terms “representation” and “concept” very loosely in this paper, with no assumptions about the particular types of mechanisms an agent might use to navigate through the world. An agent’s representations of the world could be declarative/symbolic (as in a “language of thought”) or they could be dispositions to enact certain procedures that “presuppose” the world has a certain structure.

2. Note that one can use type-labels to describe, not only subsumptive kinds of *arrangements of* entities, but also kinds of *primitive* types that are not themselves composed of anything else. “Quark,” for example, is a type of elementary particle, while “water molecule” is a type of arrangement of elementary particles. Having primitive types is absolutely necessary for saving memory space. It would be impossible for any agent to have an individual name for every individual quark in the universe, rather than simply saying that all these are of the type “quark.”

3. Reductive identifications are often said to be explanatory. Are subsumptive descriptions explanatory? The answer, of course, depends on what theory of explanation one subscribes to. I suspect that our term “explanation” may be too vague to give us much guidance in many cases where an account's explanatoriness is in question. In any case, in this paper, I am concerned only to show why subsumptions are epistemically advantageous to an agent, whether or not they qualify as “explanations.”

4. Such broad general categories are often said to be “multiply realizable.” Much has been written about the special qualities of multiply realizable concepts, even though multiple realizability has never been very well defined in the philosophical literature (see Shapiro: 2000). The most written about feature of multiple realizability is that MR concepts are often thought to be irreducible (see Putnam: 1975, Fodor: 1993). I argue in a forthcoming paper that MR concepts are just as reducible as any other. Less stringent MR concepts, however, will reduce to very broad physically defined classes. Physical defineability, knowing the broad physical features shared by members of a category, thus, need not mean an agent has very specific information about things described by the concept.

5. This claim seems to run counter to claims made by some psycholinguists that complex terms can sometimes be learned or understood more quickly than more simply defined terms. But this is not the case. First, I am only making the relatively weak claim, having a shorter definition is *one way* that a concept could be acquired more quickly. Secondly, while science might provide us with determinate definitions for our intuitive concepts, that doesn't mean that most of our concepts are represented in our heads as definitions. Indeed, there is much evidence that they are not (see Fodor 1998). But my claim in the text concerns the comparative ease of learning concepts *when they are* represented via defining features. Thirdly, my claim is about the comparative ease of learning simple vs. complex concepts “all other things equal.” But all other things may not be equal for certain concepts. We might learn certain complex concepts more easily because these concepts contain numerous features that are intrinsically interesting to creatures like us. We

might even have certain complex concepts 'hard-wired' in us innately. This is all consistent with simpler concepts being more easy to acquire, "other things equal."

Cited References

- Churchland, Patricia: 1986, *Neurophilosophy*, MIT press, Cambridge, MA.
- Kim, Jaegwan: 1998, *Mind in a Physical World*, MIT press, Cambridge, MA.
- Hempel, K. 1966, *Philosophy of Natural Science*, Englewood Cliffs, NJ. Prentice Hall.
- Nagel, Thomas: 1998, 'Reductionism and Antireductionism' in Gregory R. Bock and Jamie A. Goode (eds.) *The Limits of Reductionism in Biology*, John Wiley and Sons. New York, 3-14.
- Fodor, Jerry: 1998, *Concepts*, Clarendon Press, Oxford.
- Fodor, Jerry: 1993, 'Special Sciences', reprinted in R. Boyd et. al. (eds.), *The Philosophy of Science*, MIT press, Cambridge, MA, 429-442.
- Putnam, Hilary: 1975a 'Philosophy and Our Mental Life', in his *Mind, Language, and Reality*, Cambridge University Press, Cambridge.
- Shapiro, Lawrence: 2000, 'Multiple Realizations', *Journal of Philosophy* 97, 635-654.