



# VAGUENESS IN EMERGENT LEVELS

*VAGHET I EMERGENTE NIVÅER*

VÅREN 2017

UNIVERSITETET I BERGEN

INSTITUTT FOR FILOSOFI OG FØRSTESEMESTERSTUDIER

MASTEROPPGAVE I FILO350

*Øyvin Higdem Waaler*

*Veileder: Gunnar Karlsen*

## Contents

Abstract .....	3
<b>Introduction.....</b>	<b>4</b>
<b>Chapter One: Vagueness .....</b>	<b>11</b>
1.1 Sorites .....	11
1.2 The Core Problem .....	13
1.2.1 Higher Order Vagueness .....	15
1.3 Solving the Problem .....	15
1.3.1 Infinite Value Logic .....	16
1.3.2 Supervaluation.....	17
1.3.3 Epistemic Vagueness.....	18
<b>Chapter Two: Emergence .....</b>	<b>22</b>
2.1 The anatomy of emergent properties.....	22
2.1.1 Novelty.....	23
2.1.2 Non-derivability .....	24
2.1.3 Non-reducibility .....	26
2.1.4 Additional characterizations .....	29
2.1.5 Types of Emergence.....	29
2.2 The scope of application .....	32
2.2.1 Weak Emergence in Cellular Automata.....	33
2.2.2 Strong Emergence of the Mental and the Sparse view.....	36
2.2.3 Downward Causation.....	37
2.3 The Case for a Layered Nature.....	41
2.3.1 Outline for an emergent hierarchy of special science .....	42

<b>Chapter Three: Vagueness in emergent hierarchies.....</b>	<b>51</b>
3.1 Vagueness in the characteristics of emergence.....	51
3.1.1 Novelty.....	52
3.1.2 Non-derivability.....	55
3.1.3 Non-reducibility.....	57
3.2 Vagueness in emergent levels.....	59
3.2.1 Criticism.....	61
3.3 Consequences.....	64
<b>Summary .....</b>	<b>67</b>
<b>References.....</b>	<b>69</b>

## *Abstract*

In this thesis, I will present a case for philosophical emergence, and my argument that it is too vague. To do this, I will provide an overview to both vagueness and emergence, and focus on the issue of emergent levels. I will argue that emergence and its characterizations are vague, and that this is a general point, not specific to one single definition. Lastly, I will show why this is problematic for theories about emergent levels.

I denne tesen, så skal jeg presentere filosofisk emergens, og argumentere for at det er for vagt. Jeg vil gi et overblikk i både vaghet og emergens, med fokus på teorier om emergente nivåer. Jeg vil så argumentere for at emergens og dets karakteriseringer er vage, og at dette er et generelt poeng, ikke spesifikt til en bestemt definisjon. Til slutt vil jeg vise hvorfor dette er problematisk for teorier om emergente nivåer.

## Introduction

Our use of language tells us that we humans are quite ontologically entitled beings. We would like to know what kinds of birds there are, how many grains of corn equates to a heap and what types of radiation that are the most dangerous. All the while we implicitly assume that there are such things as kinds, heaps and types to begin with, and not just a spectrum of instances or a solid whole. We both divide and group; it seems either to make the world more bite-sized and digestible, or because it really represents nature, a world of groups and divisions.

To further frame this, we should look at the attempts to structure our world in science and philosophy. If we accept a vaguely realistic position, in the sense of *scientific realism*, a clear aim of science would be to figure out the truths about nature. If we are a bit more secure in our conviction that science can attain good knowledge about reality, we would also want to assign some sort of grid to this knowledge so that we can differentiate between *kinds*, and not just *instances* (analogous to types and tokens). To substantiate this kind of claim, a strong strategy should be to get some insight into differentiating properties. Properties that can be used as points of division. Where anything south of this point is type A, and north is type B. Insects have six legs, arachnids eight.

Historically categorizations of this kind have proven to be difficult, like Aristotle's famously problematic claim that humans are rational animals. These divisions have come a long way, and are paramount to us, in both a scientific everyday sense. We need to be able to talk about types of things, and to say what kinds of categories this and that belong to, to compress the large amounts of information we are dealing with. We cannot go around having different words to refer to each instance of something different. The borders between terms and categories are not easily assigned, however. We would like the divisions to be based on good reason, rather than to be arbitrarily allotted. If all divisions and groupings had certain property differences, then everything could be neatly categorized. If we could say that all the phenomena of a certain type of interaction for instance, were the domain of physics, and all the phenomena of another type of interaction was the domain of chemistry, our problems would be solved. Unsurprisingly,

this ideal has a lot of obstacles to overcome before we can do this with philosophical safety, as most properties are not unique to a single phenomenon.

Another issue about defining and dividing is about whether something can be composed of other things, and how we separate between the things it is composed of and the thing in the full scale. This is essentially divisions of complexity. We divide between levels of complexity all the time. Take for instance the difference between a group of atoms and a table that the atoms constitute. There seems to be some various clear-cut ways of looking at this table. There is the level of regarding the table as its atoms, there is the level of molecules, the level of wood, and the level of table. There might be more levels, like the level of being an object of some sort (like if you wanted to count all the objects in a room, and didn't care about the fact that it was a table, and not, say, a lamp). An essential realization about this, is that the table is not simply an abundance of atoms, but it is as a system very different than its components. So, what if we could separate between different phenomena and scientific areas of study on the complexity of the phenomenon? There seems to be a change in the kind of objects we are dealing with, at different levels of complex description.

The philosophical theory that embodies dividing phenomena by their level of complexity, is emergence. It can be characterized by the maxim that “More is different”. Although this is an oversimplification, it reflects the spirit of emergence. It was a prevalent way of thinking in the mid 1800's and up until the early 1900's among a wave of thinkers that have been later called the British Emergentists. Among these were C.D Broad, J.S. Mill, Samuel Alexander and Lloyd Morgan (McLaughlin, 2008, p. 19). This was an essentially materialist position, that viewed all the phenomena of the world, as being various levels of complexity of some base matter, self-organizing and systemized by the laws and mechanics of nature (McLaughlin, 2008). This, in turn, provided them with a system or hierarchical organization in nature, as laws of biology emerged from laws of chemistry, that emerged from the laws of physics. In addition to this hierarchical view of natural phenomenon and the sciences that described them, Alexander went even further, and postulated a fixed grid of emergent phenomena that existed in the world, space and time being the most basic, and each emerging from the other culminating in

‘Deity’ emerging from the base level of all the collected human minds. The nature of ‘Deity’ would be as unknown to us as the nature of the mind is unknown to the physical aspect of our brain (McLaughlin, 2008). While this seems outlandish, it follows from the basic principles of the philosophy, and as we will see is not radically different from views on emergence by thinkers like George F. R. Ellis, except for the more difficult to swallow metaphysics.

In recent years, there has been a resurgence of interest in emergence, but with a near complete rehabilitation of emergence, which largely does away with its dubious metaphysics, and frames emergence in a more scientific perspective. However, the ideas about hierarchical structure of complexity have survived. In this thesis, I will only deal with the modern debate about emergence, and not with British Emergentism. My focus will be with emergent claims about levels and hierarchies, and idea that we can properly differentiate between emergent levels.

Emergence looks like a strong contender for dividing and grouping our nature neatly, by using *complexity* as the standard property of doing so. The use of emergence in academia is usually loose and widely applicable. Whenever phenomena appear out of seemingly dissimilar components, emergence can be invoked to classify it. When we are looking at cases of emergence, to help inform our philosophical debate there are many examples to choose from. There are wisdom of crowds examples, which claim that if you have a large collection of educated guesses at a subject (e.g. farmers guessing the weight of oxen or taxi drivers’ estimation of downtown traffic at a given time), the mid-point between all the guesses will tend to be more correct than any single guess. So, that we could say that no single guess was right, but the guess of the system as a whole was right and therefore an emergent property. In biology, phenomena like the complex behavior of ants has been cited as emergent, because the ants by themselves are very basic and non-complex, compared to the organization of the ant hill that can perform tasks like farm fungus, move home and have designated garbage disposal areas, without any of these things being attributable to any single ant. In modern science, complexity, and with it emergence, is gaining conceptual ground. In computation, a hugely influential scientific area, various kinds of complexity are being developed, investigated and used, and the

language we use about this reflects a turn towards emergence, as a way of viewing these systems. Ellis further nominates evolution, earthquakes, sociology and many more subjects, besides computation, as being emergent (2006, p. 80, p. 83). The examples are many and varied, but as this philosophical discourse demands, we must narrow down the requirements for emergence. In chapter two I will investigate the conditions that are required for stating that something emerges from something else, and the kinds of things that can emerge.

The paradigmatic case of modern weak emergence, which is the type of emergence that is the focus of this thesis, are cellular automata. These are discrete, abstract mathematical systems, but ones that can be realized in various “real life” systems. A cellular automaton consists of a lattice of cells, each with a state value assigned to it, from a finite set of states. This grid of cells can be one, two or three-dimensional. The set of rules for this lattice of cells is one that decides the state of each cell on the grounds of the states of its neighborhood of cells, so that there are no actions at a distance. At each time step, the value or state of each cell is changed or kept because of the state of its neighboring cells. This system is completely deterministic, but still its cells can form surprising and different macro structures that are not readily derivable or reducible. The reason why we should be interested in these systems when discussing emergence, is that they give us arguably the most easily dissectible case of weakly emergent phenomena, that will prove useful when trying to identify the conditions of macro phenomena emerging from a micro base level. We will see why these are considered weakly emergent in section 2.2.1.

In philosophy, there are certain problems that, once you notice them, they never seem to go away. For me, this is the case with the problems of division and definition. It is an exceedingly difficult prospect, to define exactly what it means for something to be high, for instance. Firstly, we must disambiguate the term, and decide what kind of high we are talking about, whether it is a drug induced high, or if the thing is simply tall. If it is the latter, a mountain would probably not be high if it was as high as a high bookshelf, so it seems context dependent. How high does a mountain need to be to be high anyway, and how many rocks do we need to stack on top of a hill before it becomes a



mountain? This last problem is a version of the heap paradox, and it pertains to the philosophical study of vagueness. A vague term or predicate is a term or predicate that have an unknown extension; meaning that we don't know, and probably cannot know, the exact cases it is true for. There is nothing an all-knowing geologist could tell you that would make you realize when a hill goes from being a hill to being a mountain, other than just defining some arbitrary cut off point like 'rock formations of 500m and up are mountains'. No amount of research about hills and mountains would tell you either. I claim that emergent levels are vague in this exact way, and I hold that this is a problem for the doctrine of emergent levels.

The philosophical study of vagueness is primarily a linguistic one, insofar as we take the vagueness of certain terms to be a property of the language that they exist in. There are those that would argue that even these kinds of examples of vagueness are not created by language, but rather represent a real divide that, in theory, could be traced to a set point. Someone with an epistemic view of vagueness, would say that there is an actual number  $n$  drops of red paints that makes a bucket of yellow paint turn red, and an exact point  $n$  of height in meters that differentiates hills from mountains. The problem is just that in vague cases, we have no way of knowing what that exact number  $n$  is. Vagueness, in this view, is an epistemic problem, rather than a problem of language. I will show that no matter the theory of vagueness one accepts, they will be problematic for the emergence doctrine.

To best make my argument, I want to look at is whether there are good philosophical grounds for establishing emergent levels, if these kinds of levels are as clear cut and simple as they seem, and whether we can ever reliably identify them. I also want to argue that emergent accounts of these levels are susceptible to vagueness. Lastly, I want to find out whether these levels reflect some sort of natural hierarchy. To do this, I will first give an overview of different accounts of vagueness in chapter one. Then, in chapter two, I will give an overview of modern emergence, dissecting the terms and its different characterizations, scope of application, types of emergence and lastly some accounts of emergent levels. In the closing chapter, I will stake out my argument against the weak notion of emergence, specifically that it is too problematic to identify the emergent

levels, because the extension of what levels are emergent is incurably vague. Incurable because all the possible reasonable characteristics of emergence are vague, and therefore any definition of emergence that include these will either be too weak, too vague or too arbitrary. Explicitly, my argument is this:

1. The novelty condition of emergence is either too weak or too arbitrary
2. The non-reducibility condition is either vague or too strong.
3. The non-derivability condition is vague
4. Because of (1)(2) and (3), any functional definition of emergence is vague.
5. Show that Emergent hierarchical levels are inconsistent with vagueness.
6. Because of (4) and (5), we must accept one of the following:
  - a. Accept a nihilistic conclusion
    - i. This is in the case that you reject the possible accounts for vagueness.
  - b. Accept epistemic vagueness
    - i. This is in the case that you want to keep emergent hierarchical levels.
  - c. Deny emergent hierarchical levels.
    - i. This is in the case that you cannot accept epistemic vagueness.

In additions to the three consequences sketched out here, I will sketch out my reasons that the denial of emergent hierarchical levels is the least problematic of the three. I will then defend my argumentation against several possible lines of critique. In summary, this thesis argues my claim that if the candidate properties that are meant to separate emergent levels from one another are vague, the system cannot be as rigid, and ontologically sound or epistemically, as the emergentists would like. My thesis is not that the philosophical project of emergence is of no use - instead I think that having a system and a language to dissect phenomena and problems through the lens om emergence can be a tremendous scientific tool, at it has shown to be in computation. In the end, my aim is that the vague analysis of emergent levels will motivate us to refrain from the ontological commitments of emergence, not because it is necessarily untrue, but because it is too philosophically empty.



## Chapter One: Vagueness

The issue of vagueness made its modern debut in 1994 with Timothy Williamson's aptly named book "Vagueness". It re-introduced an ancient discussion about terms that don't seem to have clear boundaries, into the journals of philosophy. His treatment, and surprising solution to the issue, would bring about a wave of contemporary writings in philosophy of language and logic that tried to tackle the challenges that vagueness brought with it. Regardless of this, vagueness is quite often ignored in discussion where it could be invoked, as I will attempt to show later, in the case of emergence.

### 1.1 Sorites

As briefly mentioned in the introduction, vagueness is an ancient problem. Its true age is unknown but it is attributed to the Greek philosopher and maker of riddles, Eubulides of Miletus. He posed the following dialectic puzzle: If you have a single grain of corn, is this then a heap? Obviously, the answer is no. How about if we add one more grain to it, making it two grains, does this now constitute a heap? Still, I would say, the answer is no. This question is then repeated, adding one more grain of corn each time the reply is no. The crux of the riddle is two-fold. First, the person answering the questions might at some point say yes to one of these propositions. Let's say at seven grains of corn. If you give an answer like this, then you must admit that this seems arbitrary. Why not six or eight? There is no real reason to say that seven grains of corn constitute a heap, while six doesn't, other than just deciding on seven to have the riddler stop bothering you.

The second part of this argument is that the former realization that stopping at any specific point is arbitrary, will lead us accept that there is no point  $n$  where  $n$  is not a heap, while  $n+1$  is a heap; or simply that one single grain of corn can never make the difference between a heap and a non-heap. After all, it would always seem like an arbitrary division, no matter where  $n$  is situated on the spectrum. This assumption would then, in turn, make you vulnerable for a reduction ad absurdum. If no such point exists, and you start out at one grain, then you could keep going up in increments of one indefinitely, and no matter how many grains of corn you have, it would never amount to a heap. To put it explicitly:

1. One grain of corn is not a heap
2. If  $n$  is not a heap, then  $n + 1$  is not a heap.

3. Because 1 is not a heap, 2 is not a heap (Follows from 1 and 2)
4. Because 2 is not a heap, 3 is not a heap (Follows from 3 and 2)
5. ....
6. Because 999 is not a heap, 1000 is not a heap.

This conclusion, that a 1000 grains of corn is not a heap, is clearly false. The two first assumptions lead to an unacceptable answer. This is a heap paradox, also known as the sorites paradox. The name is derived from the ancient Greek word for heap, “*sōrós*”, and the terms heap paradox and sorites are used interchangeably. Before we dissect this and see which, if any, assumption we should dismiss to account for this apparent paradox, there are a couple of important properties of this argument that needs our attention. The first of which is that this is in no way an argument specific to heaps. It seems to be a common feature of terms where a kind of incremental progression is possible. Terms like *lake*, *coffee* and *many*, are obvious contenders for vagueness. You could add one drop of water to a crater and in every instance, ask when it turns into a lake. You could have a cup of hot water, and put in some tiny amount of instant coffee in increments, then proceed by asking when the liquid turns from simply being hot water into coffee. With a lot of these examples, you could simply change out the word “heap” in the heap example with some other predicate, and you’ve got a new variation. How many pages do you have to add to a pamphlet before it turns into a book, or how many hairs do you need to lose, before you are bald? We should also notice that these can be made to go either way. Just as it works from non-heap to heap, it will function the same way if we start off with a clear case of a heap, like a thousand grains of corn, and then make the claim that there is no number of grains  $n$  where  $n$  is a heap, while  $n-1$  is not.

Vagueness is sometimes seen as a problem of disagreement rather than logic, because if we did some statistical survey of what people think a heap should be, what is stopping us from just using the result from the poll to help us define the exact point where a heap stops being a heap? If we equated the true extension of a predicate with some notion of statistical conformity, we might be able to avoid the problem altogether. This line of answering the problem is misleading because we want to be able to say that people are wrong or right in stating propositions, and the statistical view won’t allow for this because some people might answer the poll without knowing anything about a heap, and

the things that make a heap, and thus skew the results in the wrong direction. The decision of accepting this as true, would then simply be arbitrary. It would just amount to the sum of a group arbitrarily deciding the boundaries, instead of just one person doing so. If we do accept this kind of statistical determination of predicates, and meaning is relegated to statistical consent, then the reasoning behind acceptance of a proposition becomes a circular one: ‘there is consent in that proposition  $p$  is true, because there is consent that  $p$  is true’ (Williamson, 1994).

Now at this point it is important to disambiguate vagueness from something being inaccurate, imprecise or unspecified. The latter is not much of a challenge; vagueness is not unspecificity, even though they are often used interchangeably in normal speech. If someone asks me to meet him or her somewhere, at no particular time, I could say that was a bit vague, even though philosophically speaking, it would not be vague, but merely unspecified. (Even though both “somewhere” and “particular time” are vague in the philosophical sense). The difference being that unspecificity is not providing enough information, but vagueness is there, no matter how much information you have. There is nothing anyone could tell you that would make the differentiating between a heap and a non-heap any easier. We should also set vagueness apart from misinformation, because if someone tells you to meet them at one o’clock and they don’t show up until two, it is not them being vague as to when you should meet up, but they’re simply dealing in misinformation. Similarly, them saying ‘meet me when the clock strikes a number that ends in two’, is not a case of the statement being vague, the extension of the statement is merely too big to be practical. For every time  $t$ , we can easily check if  $t$  ends in the number 2, so there are no borderline cases. Even if you allow for infinite decimal points on the clock, for every number  $t$  it will always be non-vague whether it ends on a two or not.

## 1.2 The Core Problem

To return, from the specifics of what cases count as vague, to its wider consequences in philosophy, it turns out that the sorites seems like a paradox may lie with the principle of bivalence. This principle states there is only two possible truth values, true and false. That for every statement, there is a definite truth to that statement, or not, regardless of what we know about it. Bivalence, also known as the law of excluded middle, is one of

the main pillars of classical logic. Without it, it would be challenging indeed to define the principles of logical operators, and to keep our well established laws of logical validity. If we then present a proposition to which neither truth or falsity could reasonably be attributed, it would spell trouble for this principle, and by extension, classical logic. Vague propositions do just that.

Consider the following proposition: “Six grains of corn make a heap”. This sentence is impossible and unreasonable to conclude as either true or false, as the number six cannot be attributed to the extension of either heap nor non-heap (Incidentally, if you disagree with my portrayal of six as a vague case of heap, then just change the number to something that is clearly vague to you. This exercise will also hopefully underline the point that these terms are hard to classify). We can make similar cases with all vague predicates. To spare bivalence of this criticism we would have to, for all heap-paradoxes, be able to state that every line of the argument (the repeating incremental question pattern) is either true or false. This cannot be done, without some good explanation, like the one provided by Timothy Williamson which we will look at later in this chapter, who sees this loss of bivalence as a great loss to logic:

“To reject bivalence is to lose both the account of the meaning of basic logical operators and the mechanical test of validity for inference forms involving them provided by the two-valued truth-tables.” (Williamson, 1994, p. 99)

The problem is not contained only by the principle of bivalence, as it seems to permeate all  $n$ -value logical systems. We could postulate a three-valued system of logic, where the third value represented something being meaningless, like the Swedish logician Sören Halldèn attempted (Williamson, 1994, pp. 103-108) so that for all propositions, there are three and only three possible answers, True, False and Meaningless. Say we identify our earlier proposition about six grains of corn making a heap as meaningless. It is neither true nor false, so it should be labeled as such. Although this solves the problem of the excluded middle at first glance, as we now should be able to damn the pesky borderline cases to just be meaningless. However, it does not take much prodding before we see this quick fix tearing at the seams. Because just as easy as we can make sorites examples with two valued logics, we could do so with three, there is still no hard border between

true and meaningless, or meaningless and false. When does something go from being meaningless to being true?

### 1.2.1 Higher Order Vagueness

These kinds of attempt to clarify vagueness are trying to invent some second order language solution to a first order language problem. The assumption is that if we could just clarify the language we use to describe the vague predicates, they would no longer be vague. This point is made clearer if we introduce a new operator, D which will signify a clear case of something when operating with vague predicates. If we say that F is a vague predicate, we can use the D operator to signify something being clearly a case of F like so:  $DFx$ , clearly not being the case  $D\sim Fx$ , and not clearly being the case  $\sim DFx$ . This move to the second order allows us to make room for more cases of F, there might be some that were not clearly F that we could take account of without the D operator.

Although we allow for more kinds of cases of F, and relationships to its extension, there is still the problem of there being gaps in between these cases where vagueness can hide, this is second order vagueness. There could well be cases that are neither clearly vague, clearly not vague or not clearly vague, and thus the extension of F is still not certain. Let us then, move up another level, a third-order description, and add the D operator to these three cases. We'll then end up with the following nine cases:  $DDFx$ ,  $DD\sim Fx$ ,  $D\sim Dfx$ ,  $D\sim DFx$ ,  $D\sim D\sim Fx$ ,  $D\sim\sim Dfx$ ,  $\sim DDFx$ ,  $\sim DD\sim Fx$  and  $\sim D\sim Dfx$ . Now, besides the complaint that it seems arbitrary that the third order of logic should succeed in ridding us of vagueness, where the second order failed, there is nothing new on this level. We can just as easily make a sorites paradox with the predicate "clearly a clear case of heap" as we can with "heap" itself; the problem remains the same: there will seemingly always be gaps in between the different possible cases of F, and it seems to be doomed to vagueness, no matter the order of it. If the metalanguage used to describe the language of the vague predicate is vague, the vagueness will stay with us<sup>1</sup>.

### 1.3 Solving the Problem

In summary, what can we say about the problem of vagueness itself? I would say at least this: the solution to this problem must either use some metalanguage, or order of

---

<sup>1</sup> For a more in depth analysis of the logic of the D operator that there is room for here, see Rosanna Keefe's excellent account of it (Keefe, 2000).



abstraction that does not itself contain any vagueness, declare the world as not vague or ultimately give up. A tall order either way. The attempted solutions to vagueness align with this insight closely. Starting with the latter, we could accept its paradoxical nature, and thereby accept a sort of nihilism. This alternative is not particularly interesting to pursue, for obvious reasons, and is not a place most philosophers would like to end up. We would have to throw out logic straight away, because if a hundred corns of grains can be a heap and not a heap at the same time, we would have nothing to build logic on, because acceptance of self-contradictory, paradoxical claims, is not exactly solid ground. So even though it is useful to see where one might end up, if we cannot make up our minds regarding the problem of vagueness, it is not the place to start our search for a viable solution.

The second approach to the problem of vagueness is to attempt to create a space for vagueness in logic, and this is one we've already touched on. The attempt of any n-value logic does not seem fruitful at first glance, but it shouldn't be written off so quickly. I want to analyze the consequences to vagueness of accepting another variant of multiple value-logic: an infinite-value system. I also want to investigate an alternate way forward, through Supervaluation.

### 1.3.1 Infinite Value Logic

Vagueness seems to be a property of terms that would like to be more binary, and solutions like higher order descriptions of vagueness try to make it even more binary, and solve the problem that way, by the route of extreme defining. But what if the solution is found at the other end of the spectrum, what if truth and falsity is not best measured in binary true and false, but is a more fading, analog thing, where something can be more or less true in every possible degree, avoiding the problem all together. Seven grains of sand might not be either a heap or not a heap, but maybe it is half true that it is a heap, maybe it is 0.0342 true on a scale from 0 to 1. This is the kind of intuition that leads us to thinking infinite value logic might be viable as a solution to the problem of vagueness. There is progress to be had here, moving from specific n-valued systems. The greatest asset of the system is that we can reflect the truth of a vague expression to the perfect degree at face value, at least in theory. Infinite value logic does not assume that we can reliably attribute the exact degree of truth, or discriminate

between something being say 0.767 true and 0.768 true. The continuum of truth values just makes room for the possibility (Williamson, 1994, p. 113). Armed with this, we can quite trivially account for the sorites paradox. For each increment of grains added to the would be heap, the truth of the statement ‘is n grains of corn a heap’ goes up to a corresponding value.

In the end, even this kind of infinite continuum fails the challenge of vagueness, because we can still find some refuge of sharp boundaries. The difference between something being 0.9999999999 true and something being true to a complete degree of 1 (remember, the scale is 0 to 1), is only a fraction of a degree, yet there is a substantial difference between being as close as you can get to true without being true, and being completely, no holds barred true. The same goes for the difference between almost false and completely false. We could still make a sorites that says, when does something go from being a true heap, to an almost completely true heap. There are many additional considerations when looking at a new value system to replace bivalence, other than the way it deals with the vagueness or sorites, of course. There are discussions that can be had about whether it preserves modus ponens, and how truth-functionality is can be formalized when the operators need to deal in degrees, but none of this is of further importance for our later application<sup>2</sup>.

### 1.3.2 Supervaluation

Supervaluation is one of the most prominent way, and certainly the most successful, to solve the problems vagueness poses to logic, without resolving to an epistemic view. It introduces a new, stronger notion of logical truth, called supertruth and defines this using “precisifications”. Rosanna Keefe states that a theory that tries to avoid the epistemic view needs to be able to fulfill the following three objectives: (i) It should allow for cases of vague predicates to be neither true nor false. (ii) It should “avoid the commitment to sharp boundaries with vague predicates”. (iii) It should be “tried and tested logic”, that does not break too harshly with our intuitions or lead us into strange unforeseen consequences (Keefe, 2000, p. 153).

---

<sup>2</sup> For an extensive account of many-valued logics and infinite value logic see (Williamson, 1994).

In brief, Supervaluation states that Truth (that is the proper truth with a capital T), is supertruth. Conversely, falsity is superfalsity. Things are supertrue or superfalse, if and only if, they are true or false on all possible precisifications. A precisification is a certain interpretation of the extension of a vague predicate. When talking about something being big, like a clearly vague predicate, we would say that some things might be said to be big, and might be said to be small, leaving big to have an unclear extension. If we call a football big, then that might be true on some precisifications, while untrue on others, rendering it indeterminate. This also works for ontological vagueness, as in the proposition “There is an odd number of houses in Bergen”. This statement, on all precisifications, will be either true or false; that is, no matter the division of what count as Bergen or not, and no matter the way we count a tree as being inside the border or not, it will always be true or untrue, and therefore, in Supervaluation, be indetermined. An example of something being supertrue would be “1 is a number”, because it will be true on every precisifications.

What is important to take away from this discussion in the context of my later argument, is that Supervaluation deals with vagueness in such a way that it stays vague, it does not remove vagueness from ontology in the way the epistemic view does. It is not a denial of vagueness in language, rather it accepts it, but conforms logic to fit its challenges. Consequently, any claims about a predicate being vague cannot be “solved” by calling on Supervaluation, because it cannot do anything other than saying that certain extensions of a predicate is supertrue or superfalse, or indetermined. To provide us with ontological bivalence, you would need an epistemic view that gives borderline cases an “actual” ontological status of either true or false. Supervaluation cannot help you out of any vagueness-related ontological problems, because it simply rescues logic from vagueness’ consequences for bivalence, it does not give you any tools to deal with the ontology of vague predicates, even theoretically. For this, you need an epistemic view.

### 1.3.3 Epistemic Vagueness

If none of these solutions are satisfactory, then the last refuge of non-nihilism is in an epistemic account of vagueness. This is most famously put forward by Timothy Williamson, and I will use his account of it as found in *Vagueness (1994)* as representative

for this kind of argument. It differs as a possible solution from the abovementioned, in the way that it doesn't struggle as much with explanatory power, as it does in motivation for this kind of explanation; meaning that this way of viewing vague predicates is not as intuitive.

The argument begins with the denial of the second premise in a sorites argument; If  $n$  is not a heap, then  $n + 1$  is not a heap. For Williamson, this is false. This point exists, yet we simply do not have epistemic access to know the value. If we accept this, then we can keep the principle of bivalence, because for every proposition including a vague predicate, it is either true or false (Williamson, 1994, p. 217). There is no case of "is  $n$ -number of corn a heap", where the answer is truly a borderline case, but there are still plenty where it is either unknowable, uncertain, or both. On this account, vagueness is merely a manifestation of ignorance.

If vagueness is not truly there, and it is not in our language, then how is our epistemic processes to blame? There needs to be some sort of epistemic filter that translates non-vague phenomena, to vagueness in a non-vague language. Williamson gives two basic insights that explain this. Firstly, we can have inexact knowledge. Secondly, there is no reason for why we should know the borders of our own terms (Williamson, 1994, p. 216). It is quite reasonable to assume that humans do not have perfect insight into ourselves, or the way we view the world around us. So, the presumption that the borders of the terms we use are not available to us, is not a problematic one. But while it is easy to explain our ignorance, the case of our knowledge is more demanding. This inexact knowledge, while still constituting as true or untrue, is so about non-exact matters. Williamson gives the example of someone surveying a large mass of people, trying to estimate the number of people in the crowd (Williamson, 1994, pp. 217-226). The crowd-watcher estimates there being around 20.000 people there. Now whatever this estimated number of people may be, there is no way of verifying that it is true or not. This does not, however, disqualify the viewer from having gained some knowledge about the number of people present. It is for example, not only one person present, and there is certainly not 10 billion people in the crowd. But the viewer cannot have the knowledge that there is 20.000 people in the crown, even if the estimation is correct. This is because the viewer would not be able to differentiate between 20.000 and say

19.000 or 20.001 ( $n + 1$ , or  $n - 1$ ). No amount of discussion or analysis could get you to a true knowledge about the exact number of people in the crowd. Indeed, for every possible number  $n$  where  $n$  is a guess about the size of the crowd, we could say that the viewer would not be able to differentiate between a crowd of  $n$  and a crowd of  $n - 1$  or  $n + 1$ . This knowledge seems to be, just in the nature of it being the way it is, unknowable, while it still seems to be the case that the observer has some knowledge about the crowd.

One of the thickest veins of critique against Williamsons position is centered around question of whether we have a good motive for viewing vagueness as an epistemic phenomenon, rather than one rooted in language. There are other paradoxes, like self-referring paradoxes, that would be hard-pressed to be presumed to be of an epistemic nature, rather than an artifact of logic or language. Why should the case of vagueness be any different? Regardless if this account of vagueness is true or not, the main take away for our later application, is that epistemic vagueness, while not forfeiting bivalence, and providing us with at least the possibility of a non-vague extension of terms, the vagueness does not dissipate from our epistemology. Even though epistemic vagueness retains classical logic, the non-vague references of our language are not readily available to us, which is why vagueness remains a factor for us, even though it might not exist in the world.



## Chapter Two: Emergence

The modern account of emergence is often one that goes hand in hand with some sort of scientific discipline. Where it used to come from philosophers, this idea is now widespread in certain emergent minded scientific communities; The kinds of sciences where wholes seem to have operational modes all to themselves. Biology, chemistry and computation are all bringing their expertise and cases of study to the field of emergence. Consequently, the meaning and use of emergent vocabulary, and concepts will typically be different between programmers, biologists and physicists, although there is some set of core ideas that repeat across all the uses.

The term is lightly used in most cases, and is used to describe everything from some unintentional side effect, or interaction dynamics. A former case could be if a garden was designed to be serene and silent, but inadvertently turned in to a blooming social arena. Or that faking injuries on the football field to get free kicks has become a part of the game, despite it never being intended as such. The second is when describing different systems interacting in such a way that new system appears. An example of this would be that you could run, but you could also jump, and if you combine the two, you would achieve a forward moving jump that none of the two other systems alone would be able to give you. These examples are only emergent in the most superfluous sense, as the supposedly emergent property was there all along (and who is to say that running forwards is not an emergent property of forward jumping whilst not jumping?).

The way in which this philosophical analysis will treat emergence is, like vagueness, predictably a bit different. To understand emergence, we first need to pick it apart by its seems, and identify the characteristics of the phenomenon.

### 2.1 The anatomy of emergent properties

When talking about emergence, we usually focus on properties. We operate with properties in such a way that a certain compound of materials, in a certain structure has a certain, definite set of properties. For instance, a compound of materials can be set up in such a way that it only reflects red light off the surface, giving the property a red color

when viewed from the outside. This would count as a basic property, along with all the other, listable features of the compound materials in question.

An emergent property, is a property of a system of physical materials, that has attained a type of complexity, such that it starts to display different, novel features, which are not found in its constituent properties. This deserves some unwrapping. This does not entail that the connection between the lower level parts, and the emergent properties, is a causal one, but as we will see in section 2.1.4, it probably needs to be supervenient.

Our definition of what the emergent property is should be a functional one, meaning that we view the functional aspects of the emergent phenomena as the thing that emerges. In the case of the mental being emergent from the physical, a functional description of the mental phenomena is what we take to be the higher-order emergent property.

### 2.1.1 Novelty

A central claim in the literature arguing for emergent understandings is that some cases of complexity can display *novel* properties to those things the complexity constitutes of. A notable example of this can be found in the claim that the mental is an emergent property from its physical base. As little as we know about the physical and the mental, we at least know that they are not the same, functionally speaking. The emergentist hypothesis that this relationship is emergent says that a certain type of complex configuration of relatively simple physical material, arranged in the form of a human brain in a live human, produces a consciousness. This mental then is not like anything you could have in a physical form, it is *novel*. The novelty characterization is twofold, the higher order property must be different from the lower level, and it must be the kind of property that could not be a lower level property. Jaegwon Kim argues that this feature is a necessary feature for emergence for it to succeed in its aspirations (Kim, Making Sense of Emergence, 2008). Without it, emergence will be left without the theoretical punch it needs, because the emergent property would simply be additive or identical.

The novelty condition is one that is usually defined quite intuitively. I hold the account of novelty as given by S. Rasmussen, N. A. Baas, B. Mayer and M. Nillson to be one of the more extensive and robust. They state that at heart, the novelty condition can be



defined like this: “A property that applies at a given level is emergent if it does not apply at any lower level.” (Rasmussen, Baas, Mayer, & Nillson, 2008). I take this as not including emergent properties that are in essence a higher degree of some shared lower level and higher order property, compared to the lower level. Like some components being slightly blue, and the whole being very blue. But that the features of the macro structure are different in some respect, in addition to not applying, to the micro structure. I assume this because I think it is a breach of the definitions intention if cases where a macro structure merely has a more prominent version of a feature, are counted as emergent. I also take this to not include negative features, like some macro structure not having the properties of its micro structure. This too because it goes against what I take to be the intention of novelty, that something is added, not novelty in subtracting features.

### 2.1.2 Non-derivability

Emergent properties should not be predictable, or derivable, from its component parts, such that the emergent property is cannot be attributed directly to its base. The reason why we would want this as condition for something being emergent, is that it fits our intuitions about what emergence should be like. When we say, something is emergent from something else, it entails that the higher order property is somehow separate from the lower level ones, if they were not separate, then it would just be a direct property of the lower level. Without the separation between the lower and higher level, the relationship between them would be describable as straightforwardly causal, and higher order phenomenon would be epiphenomenon, improbable and uninteresting as causal powers themselves. Non-derivability provides us with this wall of separation in the form that even if we are given full knowledge about the base properties of some domain, we would not be able to derive its emergent properties from it.

The unpredictability characterization of emergence is related to novelty. Because a lot of the novelty of a phenomenon stems from its non-derivability (Kim, Making Sense of Emergence, 2008). If some atomic property, like a carpet fiber, is soft, then we could predict that a square 1x1m carpet would also be soft, maybe in an unusual way, but still soft. This property is both predictable and not new. If, on the other hand, the 1x1m

carpet would turn out to be able to fly, then that would be unpredictable, and thereby also novel.

Non-derivability can be defined quite straight-forwardly like this: something B is non-derivable from something else A if and only if no amount of knowledge about the conditions and interactions of A, and no amount of analytical capacity would allow you to derive B with certainty. This type of non-derivability is strict to the extreme, we allow for no probability calculus. This is not “actual” non-derivability, as defined by the current abilities we have at deriving something, but rather the kind of theoretical non-derivability, that an infinite all-knowing Laplacian demon would not be able to get past. Anything that is even remotely theoretically derivable, by the most time consuming, unpractical way, does not count.

A way to argue that non-derivability is possible, or at least as common as to accommodate a range of emergent phenomena, is to invoke multiple realizability. The argument, in short form, is that some types of phenomena, like mental phenomena types, could be actuated by not only one, but several types of lower level phenomenon. Say for instance someone is in pain in form of a headache, and this same kind of pain can be caused by distinct types of chemical imbalances, even though the sensations of pain is interchangeable. This type of Headache, in this example, would be multiply realizable, because several types of physiological conditions, would realize the same type of. This point of multiple realizability in the non-derivability of emergence will not be discussed further in my thesis, partly because it is not a required aspect of the non-derivability condition, and because this is a stronger point with the strong emergence debate, which is not my focus, and lastly because the debate about it is too large and unwieldy to represent properly while still retaining my own focus<sup>3</sup>.

There are weaker ways of defining the non-derivability characterization that are more directly useful for our purposes, like Mark A. Bedaus non-derivability except by simulation (Bedau, 2008). In this article, he attempts to avoid the strictness of proper non-derivability, to make room for a weaker sense of emergence. In his sense, something is non-derivable except by simulation *if and only if* the only way of deriving a result from

---

<sup>3</sup> For further discussions about multiple realizability in emergence, see (Fodor, 2008) and (Kim, Making Sense of Emergence, 2008).

some set of initial conditions, is to step by step, replicate the functions of the system step by step, in a one to one scale. This entails that there are no theoretical shortcuts to what set of emergent properties will spawn from a certain set of component parts. The only way of predicting it is to do an experiment, observe the emergent results and then inductively predict that the same will happen again with the same set of component parts, *ceteris paribus*. The reason emergent properties need to be unpredictable in this sense, is that otherwise a strong point could be made that nothing emerged, it was all there in the component parts all along. This reasoning behind this variant is best understood considering the cellular automata, which will be discussed in detail in section 2.2.1.

Emergent phenomena could be *inductively predictable*, in that we know what the same or similar conditions led to the last time, and assume it will do so again. The unpredictability of emergence is a purely deductive claim, that with all possible knowledge of the component parts, we will not be able to predict the emergent outcome with certainty, avoiding inductive reasoning.

### 2.1.3 Non-reducibility

Another important aspect of emergent properties is that they should be causally distinct, and not reducible to the causal powers of the component properties. If this is not granted, then emergent properties would be epiphenomena; causal dead ends. This commitment to distinctive causal powers for emergent properties, in turn, leads the emergentist to have to accept downward causation, because if emergent properties can have causal powers on other emergent properties, and the state change of emergent properties entail a state change in component properties, then higher order causality entail downward causation. This non-reducibility means that emergent phenomena should neither be reducible in the way described above, but also that they are not explainable in terms of its component properties.

This non-reducible aspect of emergence can, and has been, critiqued for being a negative characterization, and therefore uninteresting. Jaegwon Kim argues that things that have a certain characterization, like things that are rectangular, or green, constitutes a group of things to a greater degree than things that have a negative characterization do, like being not-rectangular or not-green (Kim, *Being Realistic about Emergence*, 2006, p. 201).

Members of the positive characterization could be things like houses and books, grass and green pencils, while the members of the negative characterization could be democracies and words in both cases, as neither are rectangular nor green. The characterization might as well have been non-characters in Sherlock Holmes or members of the British parliament in 1947. The members of the positive characterization could be said to have something in common, saying that about the negative would be a stretch. Because of this, negative characterizations are often too weak for them to be relevant in categorization, and not being reducible certainly fits the bill of a negative characterization. Nonetheless it is included in most definitions of emergence in the literature, so it would be misguided to cast it aside before starting, but this critique will be revisited when analyzing Kim's arguments on emergence.

If emergence with this kind of non-reducibility is true about the human mind, this would mean that you cannot reduce or explain the experience of consciousness from facts about the brain. This might seem counterintuitive, because we do this all the time. We know some things about the causal relationship between the brain and the mind, and we can explain some mental diseases for instance, by explaining certain chemical imbalances in the brain, that correlate to the conscious states in some way. The problem here is that these psychological phenomena are multiply realizable. This means that there are multiple arrangements, compositions and variants of lower level constitute parts, that will end up producing the same higher level emergent phenomena. We can find a clear analogy to this concept in programming, where the macro features of a software can be realized through several different implementations. George F. R. Ellis gives us a helpful indicator for why this could be:

“Higher-level behaviour is based on throwing away vast amounts of information, selecting what is relevant from a vast flow of incoming information, storing it, analyzing it in a broad existential context, differentially amplifying it, and utilizing it in feedback control systems” (Ellis, 2006)

However, the arguments from multiple realizability are not, like most other philosophical areas, undisputed.

The basic idea of reductionism is that some phenomena is explainable, or constituent of, nothing more than something else; that A is reducible to B iff A is *nothing but* B. There

are several distinct types of reduction available to us, John Searle provides us with five categories of reductions (Searle, 2008, ss. 70-71) :

1. Ontological reduction – Objects of one type constitute nothing but objects of another type.
2. Property ontological reduction – Properties types of objects constitute nothing but properties of another object types
3. Theoretical reduction – The laws of one theory, can be arrived at from nothing more than the laws of another theory.
4. Logical or definitional reduction – Reduction from general to specific case of a word and vice versa.
5. Causal reduction – One causal power is explainable in terms of another causal power.

Of these only number 1, 2, 3 and 5 are useful to us, and the logical definitional reduction would not apply here. Proponents of emergence usually argue against various kinds of reductionism. For example, P.W. Anderson argues that while he holds ontological reductionism to be true, he argues against theoretical reduction in science, in the sense that there are various methodologies and theoretical tools that are needed and used at the different scientific levels, although they are essentially nothing more than the thing they consist of, ontologically speaking. He posits a kind of hierarchical structure of science, because ‘The elementary entities of science X obeys the laws of science Y’ (Anderson, 2008, s. 222). Solid state or many-body physics obeys the laws of elementary particle physics, cell biology obeys the laws of molecular biology, all the way up to that psychology obeys the laws of physiology and social sciences obey the laws of psychology (Anderson, 2008).

Non-reducibility as a characterization for emergence is for the above reasons a very attractive proposition. But is not so readily used as non-derivability is, because of the problems that come with the downward causation non-reducibility entails, as we will see in section 2.2.3.

### 2.1.4 Additional characterizations

Lastly, it is important to recognize that this is not necessarily a two-level hierarchical structure. Emergent properties can emerge upon other emergent properties, depending on the scope of assessment. You could argue that H<sub>2</sub>O has emergent properties that neither the single hydrogen, or neither of the two oxygen atoms have on their own, like its ability to quench human thirst. Meanwhile a huge body of water, has different properties from a single H<sub>2</sub>O molecule, like the ability of having a boat float on it. This example is not meant to argue the for or against of water as an emergent property, but simply to give a clearer picture of layered emergent properties.

In summary, emergent properties should be unpredictable, causally non-reducible and novel, in varying degree according to the type of emergence we're hunting, but these are the general hallmarks. Even though the concept of emergence, like many other terms in philosophy, is a rather slippery one, I hold that this sums up the general requirements for emergence, although I admit there being variation.

### 2.1.5 Types of Emergence

A lot of the variation in what counts as emergent comes down to the strength one requires of the fulfillment of the conditions we have discussed, rather than the difference in conditions. Mark A. Bedau suggests separating between three degrees of emergent strength: nominal, weak and strong emergence (Bedau, 2008, p. 158), while Andrew Assad and Norman H. Packard provides us with at least three: Weakly emergent, strongly emergent and maximally emergent (Assad & Packard, 2008), while staying open to the possibility of more intermediate definitions. The three being approximately analogous to the other (Nominal emergence for Bedau being roughly equivalent to Weak emergence for Assad and Packard and so on), but with slight but crucial differences in definition. I will go by the naming scheme proposed by Bedau, and treat the definitions of Assad and Packard and later David Chalmers as alternate versions of Bedaus definition, to keep the language consistent.

The nominal version of emergence, in Bedau's definition, is a broad sense of the concept, ranging from mere aggregate resultants, but also encompassing properly emergent phenomena. Nominal emergence applies to any phenomena that cannot be a micro property. These nominal emergent phenomena are usually philosophically

uninteresting and straightforward, like the fact that two individual shoes, make up a pair of shoes, even though none of the shoes had the quality of being a pair on their own. This kind of emergence is dependent on the lower level properties, but is still autonomous from them in the sense that the lower level properties do not share the emergent properties. This is the weakest form of emergence of this definition. On the other hand, Assad and Packard view the most basic kind of emergence to be when “Behavior is deducible in hindsight from the specification *after* observing the behavior” (Assad & Packard, 2008, s. 232). So, in their understanding, it would be a system that is surprising, but still fundamentally rule following, so that we can presumably make good inferences about the systems behavior after observing it.

To filter out the properly emergent from the merely resultant or aggregates, we need to tighten the net of conditions somewhat. Weak emergence describes those cases where the supposedly emergent phenomena is different in kind from its component parts; that is that the condition of novelty is strong to such a degree that it is not a result of mere aggregability (ref). The definition of weak emergence is one that fluctuates somewhat, relative to the definitions of nominal and strong emergence. Weak emergence simply includes the cases that are properly emergent, but not strongly emergent. The most important aspect of both Bedau’s, and Assad and Packard’s definitions of weak emergence (what Assad and Packard call strong emergence), is the level of non-derivability that is found in these cases. For all of them, weakly emergent phenomena are phenomena that are immensely challenging to deduce, while still being theoretically derivable. As mentioned earlier, Bedau’s definition of weak emergence only allows for one type of derivability, and that is *by simulation* (Bedau, 2008, ss. 161-164), while Assad and Packard doesn’t specify it further than it being “prohibitively difficult” (Assad & Packard, 2008, s. 232).

Strong emergence is, as the name implies, the kind of emergence which demands the most of its conditions. It is so stringent in fact, that a lot of the debate on strong emergence revolves around whether there are any phenomena that are properly strong at all. According to Bedau, strong emergence are cases of emergent phenomena that are properly emergent, while still being causally autonomous. They are “[...] supervenient properties with irreducible causal powers.” (Bedau, 2008). In addition, they are non-

derivable in the most stringent sense. One aspect that both weak and strong accounts of emergence might share, that we've not yet touched upon, is supervenience. Kim makes a compelling case of why supervenience should be admitted as a clear constituent of any emergent phenomena. He argues that supervenience simply implies that emergence has upward necessitation, in that it states that if two different sets are alike in their lower level properties, they will be alike in their higher-level properties. Kim puts it in explicit terms like so:

Supervenience: if property M emerges from properties  $N_1, \dots, N_n$ , then M supervenes on  $N_1, \dots, N_n$ . That is to say, systems that are alike in respect of basal conditions,  $N_1, \dots, N_n$  must be alike in respect of their emergent properties. (Kim, *Emergence: Core ideas and issues*, 2006, p. 550)

It is important to note that this does not necessarily imply that the supervenient connection is a causal one, it could be statistical or stochastic, but that the supervenient condition needs to be fulfilled. The reason for this is that without this kind of supervenient relation, the relationship between lower level properties and higher level ones, would be rather dubious. If the higher order phenomena do not necessarily occur, when its lower level enabling properties occur, who is to say that the former emerges from the latter, and not from something else? Therefore, there are strong reasons to believe a good theory of emergence must include a supervenient relation between the higher and lower level. This does not exclude multiple realizability, because the supervenience claim only is not an if and only if statement, and the consequent can be realized without the antecedent being realized, but not the other way around. The supervenience characterization of emergence point is not uncontested, as expected, because it proves problematic through the famous causal exclusion problem that we will investigate later. If accepted, this supervenience view of emergence would be a part of an alternate definition of emergence. Kim defines emergence this way, but in addition to supervenience he argues that emergence is also unpredictability. Here we return to his argument that we discussed earlier, where a negative characterization might be a problematic one.

Finally, there is sometimes made a distinction between *ontological emergence* and *epistemic emergence*. The former variant being one that claims that these emergent properties are



actual entities in a metaphysical sense, separate from its lower level. The latter tries to hold back judgements about the ontological nature of emergent phenomena, and rather regard them just as emergent, because they seem emergent. The epistemic view is the more dominant in modern debates, and most if not all the accounts of emergence covered here are epistemic (Bedau, P.W. Anderson, Chalmers etc. all have epistemic views). The shift from the ontological to the epistemic is partially the use for emergence in science, and partly because of the failures of the British Emergentists.

We will explore whether there are any good cases of strong emergence in section 2.2.2, but before we do, let us again turn to weak emergence. This kind of emergence does not, according to Bedau, require irreducible causal powers, but it remains supervenient and is unpredictable (Bedau, 2008). Bedau goes even further with his analysis, concluding that weak emergence is “Underivability Except by Simulation” (Bedau, 2008). This means that weak emergence is not “shortcut-predictable”; that there are no epistemic shortcuts in predicting the resultant emergence, other than running the case scenario, step by step. The clearest example of a possible emergent system with no shortcuts is cellular automata, to which we will return shortly.

## 2.2 The scope of application

A question in need of answering, that has already been touched upon in the introduction, is what cases are eligible for emergent phenomena. After all, the kinds of phenomena we want to describe bears a lot of the motivation for the whole project of emergence in the first place. We should strive to designate as many type cases of possible emergent as either emergent or non-emergent as we can. As I alluded to in the previous section, cellular automata are a huge contender for harboring weakly emergent phenomena. Even if this is disputable, it certainly is the paradigmatic case of weak emergence in a closed system. When investigating strong emergence, however, the dominant case in point is the mind body relation. Specifically, that the mental emerges on the physical.

In addition to the question of what cases might be emergent, there are some other, broader questions relating to its scope. Some of these are: What kind of entities can be emergent? What areas of science has valid emergent phenomena (nominal, weak or strong)?

### 2.2.1 Weak Emergence in Cellular Automata

Cellular automata are a primarily abstract group of theoretical machines with a specific mode of operation. They are cell-based, and each cell follows certain rules according to the state of its neighboring cells. The term is used for a wide group of these kinds of systems, but in the literature, it most often refers to a two dimensional, two value cellular automata. This means that they operate on a two-dimensional board of n-size, like that of a chessboard, and that the cells have two states, on and off. These cells toggle on and off, per the state of their neighbors, following a finite set of rules. The most famous set of rules, are those set up by John Conway in 1970. This notorious zero-player game is called The Game of Life. Let us go through the functionality of this game, and then explore the case for weak emergence in this closed system.

As mentioned, this game requires a board, specifically a two-dimensional lattice. Each square on the lattice can be either alive or dead. The rules for the game of life are as follows:

*Game of Life* A living cell remains alive if and only if either two of its neighbors were alive at the previous moment; a dead cell becomes alive if and only if exactly three of its neighbors was alive at the previous moment (Bedau, 2008, p. 166).

These rules are essentially a programmatic update function, repeated for every time step. Using the definition of the board and the rules of the game, we can start a game by giving the game its initial conditions, like in the first frame (T<sub>0</sub>) of Figure 1. For each frame, then, we can calculate which cells will be born, die, keep living and keep being dead. Keep doing this to see the game playing out. Even though it probably wouldn't be a hit game to bring on rainy camping trips, it is interesting to explore the behavior of the game under different circumstances.

A lot of the initial configuration of cells in this game, just leads to a short fizzle of energy, and then it dies out, with no remaining living cells. Other configurations lead to endless repeating patterns, the same cells blinking in and out of existence ad infinitum. Some configurations lead to systems that continually generate new cells, expanding forever.

There are a lot of repeating local patterns that occur with certain behavior, like the one in Figure 1, which incidentally has been dubbed a “glider” by the Game of Life community (Yes, there is a Game of Life community). It is called a glider because it repeats itself, but only after shifting its position 1x1 on the lattice (in this case south east), as though it’s gliding through the frame.




FIGURE 1: THE EVOLUTION OF SET OF CELLS, FROM T0 TO T4

You could make a glider that moved in three other diagonal directions. There are plenty of other repeating patterns in a lot of configurations of the Game of Life, some that interact with others in peculiar ways. Take the so called “Eater” in Figure 2, going from the leftmost T0 forwards, we see a configuration of cells, and a glider approaching it. When they collide, there is some commotion, but after all is said is done, the glider is destroyed, but the eater remains, no worse for wear. It’s easy to see where the Game of Life got its name from, because not only are the cells “alive” and “dead”, but there is some sort of circle of life going on in this strange plane of existence<sup>4</sup>.




FIGURE 2: A GLIDER BEING DEVoured BY AN EATER, FROM T0 TO T4

There are some notable features of this system. First, it is capable of computation. One could devise a Turing machine using this system, and thereby, per the Church-Turing

---

<sup>4</sup> The vocabulary of the Game of Life is betraying at least the hope of finding proper emergent phenomena here. Indeed, the people who mastered the art of identifying entities in Life like the glider and the eater, have been called zoologists.

thesis, it can compute any solvable algorithmic problem (Bedau, 2008, s. 167). So, it seems that we have at least one good contender property for emergence, this systems computability, or even more interestingly it could play any game a computer can play, like chess. Another feature of cellular automata in general, is that they are unpredictable. This means that for any distribution of live cells on the Game of Life lattice, the only way of deductively predicting the outcome of this after  $T_n$  time-steps, is to do the calculation, step by step,  $n$  times. There is, however, a layer of predictability to this. Daniel Dennett points out that if we do an “ontological shift” upward from the world of cells, and to the world of gliders and eaters, we can do some predictions (Dennett, 2008, s. 201). Seeing that the glider is a repeating pattern, we can assume that it will keep repeating *ceteris paribus*. If there is an eater in the way of a glider, then we can predict that the glider is going to be destroyed when they meet, like in Figure 2. These kinds of predictions are vulnerable to faulty reasoning, not considering all the factors, and incomplete knowledge of the dynamics of these entities. But on the level of cells, there is no predicting of this kind. There is only prediction by simulation, with the consequence being that you cannot theoretically predict an outcome, before you simulate or observe the outcome. It is understandable why thinkers like Mark Bedau gives this feature a lot of weight in his understanding of weak emergence. It is close to pure theoretical unpredictability. A real-world analogy would be to say that if the lottery numbers are predictable only in the case that every possible causal aspect of the future lottery draw is accounted for in a simulation, then they are in essence, unpredictable.

“To say that their salience is considerable is to say that one can, with some small risk, ascend to this design level, adopt its ontology, and proceed to predict—sketchily and riskily—the behavior of larger configurations or systems of configurations, without bothering to compute the physical level.” (Dennett, 2008, s. 199).

So, we can have different levels of description, on different levels of ontological consideration, with different modes of prediction.

Lastly, a feature of Cellular Automata that is important for us to consider is that they are divergent and convergent. The convergence means that there are different initial conditions that lead to the same, or inseparably similar results. This in turn makes it, in most cases, impossible to go backwards from some stage  $T_n$ , and reason our way to the

original conditions. If our universe was a cellular automaton, we would not be able to reason our way back to the big bang. If we again turn our ontological heads upward, to the level of gliders and eaters, dubbed the “design level” by Dennett, we could do some reasoning based on our knowledge of the entities on this level and their interactions. We could assume that if a glider was all there was in a system, that if we go back in time it was moving backwards, but since gliders can appear from non-glider configurations, we would make these predictions on shaky grounds. The divergence, similarly, shows itself when we have two extremely similar configurations, that because of the seemingly insignificant difference, has widely different outcomes after some timesteps  $T_n$ . This tells us that there is a degree of chaos in cellular automata, as this kind of divergence is known as the chaotic “butterfly effect” in popular culture. What this means in terms of emergence, though, is that the waters for predicting and reducing, is further muddied, showing us the difficulty to predict systems, even if we predict it on the “design level”.

Can we then find evidence of a type of weak emergence in the Game of Life then? As I have covered, it seems to show behavior on higher levels, that it does not display on the lower level. For instance, you could configure a system that plays chess, like any other sufficiently powerful computer with the required program could, but the individual cells cannot. The (highly theoretical) ability to play chess, is not predictable through other means than prediction or inference on a higher ontological level, or by simulation. Lastly, this system would not be readily reducible, in the way that we could infer our way from the chess playing, to how the machine works on a cell basis, because such a machine is multiply realizable. Weak emergence, at least initially, can be found in cellular automata, so long as we accept unpredictability except by simulation as a valid form of unpredictability.

### 2.2.2 Strong Emergence of the Mental and the Sparse view

The patterns we get in cellular automata are surprising, but they are derivable, strictly speaking. If they were not, then they would be strongly emergent. It is David Chalmers’ view that the biggest difference between weak and strong emergence is that weak emergence is in principle derivable from the lower level, but surprising, while the strong is underivable even in theory (Chalmers, 2006, s. 244). The stringency of this definition can lead us to think that there are no emergent phenomena that can truly be “strong”.

This view is called a sparse view of emergence. If emergence exists, the sparse view says that it is either an exceedingly rare, or nonexistent phenomena.

But there is one paradigmatic case of strong emergence, and that is the phenomena of consciousness, as an emergent property of the physical base, the body. David J. Chalmers holds that this is the one and only case of strong emergence (Chalmers, Strong and Weak Emergence, 2006). A system is conscious for Chalmers, ‘when there is something it is *to be* that system’ (Chalmers, Strong and Weak Emergence, 2006, s. 246). He does not, in my reading, explicitly deny the possibility of other strongly emergent properties, but does not expect there to be any other either. The reason for this is that consciousness is the only thing we know about which is demonstrably different from physical things. It is not controversial to state that (almost) no matter what stance you might have on the relationship between consciousness and the body, there is a difference in being a something without consciousness than it is being something that has consciousness. Chalmers argues this positions through a couple of, by now, well known arguments. The first is Frank Jackson’s thought experiment of Mary the super-scientist (Jackson, 1982). The argument states that there are phenomena related to consciousness, like the perception of color and pain, that are not deducible from perfect knowledge of physical facts, even though they might correlate (supervene). The other is Chalmers’ own zombie argument: there is no logical impossibility in a physical identical world to our world, would have no mental features in it (Chalmers, 1996), (Chalmers, 2006, p. 247). The ‘people’-equivalent in this world would then amount to nothing but non-conscious zombies.

It is useful to remember that a view that there exist strongly emergent phenomena, and a view that there are weakly emergent phenomena, are in no way mutually exclusive. The definitions of strong emergence are usually weak emergence plus something else. In the case of Chalmers, weak emergence plus actual strict non-derivability. Because of this, an account of strong emergence entail weak emergence, but not the other way around.

### 2.2.3 Downward Causation

Downward causation is one of the most hotly contested points of debate in emergence theory. Because emergence theorists require emergent phenomena to be novel and causally autonomous, it not only opens for, but requires the possibility of higher level

phenomena causing same level, or lower level phenomena, without us being able to trace that causality back to the lower-level phenomena that the higher-level emergence upon. Armed with downward causation, emergence can provide a lot more explanatory power, and motivation for needing emergence in scientific theory in the first place. If downward causation is a real aspect of emergence, it provides higher order phenomena with causal autonomy, so that to properly explain the causal relations in nature, we need something like a concept of emergence to do so.

As the paradigmatic case in point for downward causation, the claim that the mental can have causal effects on the physical is one that is accepted by philosophers such as Ellis (Ellis, 2006, s. 102). Human decision-making is an oft-cited case where the mental state of wanting to do something, is taken to be what causes the action of doing that something. My mental state of wanting to pick up a pen, causes me to pick up a pen. My preference for hotdogs causes me to choose to buy a hotdog in favor of a hamburger, and so on. This kind of causal efficacy for higher level properties, in this case the mental, is often regarded as a characteristic of emergence in itself (Kim, *Making Sense of Emergence*, 2008, s. 141).

In physics, downward causation is also a discussed problem. If you have some complex interaction, where the whole seems to change how the low-level interactions, some theoretical problems appear. Paul C. W. Davies frames it this way:

“How can wholes act causatively on parts if all interactions are local? Indeed, from the viewpoint of a local theory, what is a ‘whole’ anyway other than the sum of the parts?”  
(Davies, 2006, p. 40)

He regards this kind of ‘whole-part causation’ to be different from the kind of ‘level entanglement’ of the mind body relation (Davies, 2006, pp. 41-43). For the whole-part causation, it is the kind of causal explanatory view of saying that a ball rolling down a hill, causes the individual parts of the ball to spin around. Davies remarks that this is a sort of quasi-autonomy where the downward causation is trivial, because the ball is not causing the atoms it consists of to act in a certain way, this system is perfectly understandable by only its local interactions. The causation is best understood the other, more traditional way, that forces are acted on interconnected parts in such a way that a whole is perceived. The same goes for whirlpools and other phenomena that seem to

have a holistic feature that is not represented in the atomic parts, but that is completely derivable. In the emergent vocabulary, these phenomena would be best described as nominally emergent, although Bedau seems to regard them as weakly emergent, and non-derivable except by simulation (Bedau, 2008, p. 177).

The other kind of downward causation, level-entanglement as Davies dubs it, and is more related to the conceptual levels. Both mental causing physical and software causing hardware are cited as examples (Davies, 2006, p. 42). It is not as though we think that there are any new physical forces acting from the mental to the physical, or from the software to the hardware, none of the actions caused on the lower level from the higher level are insufficiently explained from the lower level conditions. Rather Davies thinks that this kind of downward causation is a case of “[...] the global system harnesses existing local forces.” (Davies, 2006, p. 43). He likens it to a computer with an attached robot arm, which can more freely dependent on the software that controls it. If the software tells it to move the robot arm over its own circuitry, making changes, the software will effectively have caused a change in the hardware, in a sort of software-hardware feedback exchange (Davies, 2006, p. 43). Whether any of these are truly representative of downward causation, is not a given, as other accounts, like the one given by Bedau, does not consider the whole-part causation type to be as trivial from the lower-level as Davies does.

There are even further complexities in taking on downward causation, as Kim suggests that we should separate between two accounts of downward causation; synchronic downward causation, and diachronic downward causation, both of which are reflexive. The two are separated by their differences in time relation between cause and effect, the former being instantaneous and the latter having a delay. The reflexivity just means that the causal relation is a two-way street, the lower causing the higher, and the higher causing the lower. The mind body relation is surely a diachronic relation, as the time scale for the mental cause and the physical effect is the same. It is not as though we do the mental act of pulling a lever, and then the pulling of the lever happens after, if it is anything, it is simultaneous.

George F.R. Ellis also makes a case for downward causation in his model, that we will discuss in further detail in section 2.2.1, in a sort of trivial way. He gives an example of



an emergent computer system, where you can have same-level actions in the operating system. His point is that these same-level operations can be regarded in their own right, without referring to the lower level computation, because *"it can be analyzed without knowledge of the underlying lower-level interactions"* (Ellis, 2006, s. 89). My take on this is that he's not really talking about top down, or same level action. He is talking about top down descriptions, and same level descriptions. That is that the model works without the assumption of lower level parts. Describing things in this manner, gives us the opportunity to use shorthand in science, to abstract away from lower level and achieve clarity in describing the phenomenon a certain descriptive level. As he writes:

"10<sup>24</sup> nuclei and the associated electrons moved simultaneously in a coordinated manner so as to decrease the volume available to 10<sup>24</sup> gas molecules', which requires about 10<sup>36</sup> bits of information for a full description, actually describes the phenomenon 'the piston moved and compressed gas'. Indeed, this is the reason that we develop and use higher-level language and mathematical descriptions in the first place." (Ellis, 2006, s. 89)

It seems that same level causality, for Ellis, occurs when you don't have to invoke lower levels principles to explain something. In other words, when it safe to ignore the lower level properties. If we look at the higher level, we can safely ignore the lower level, because the system is perfectly coherent without it, and sometimes, like in the example Ellis gives, way more coherent. While it's undoubtedly true that you could coherently describe a car without going into details about the chemistry of fuel or metallurgy of the chassis or the engineering of the engine, these aspects are not actually causally impotent, just because we ignore it when analyzing the reasons of a car crash. These factors played their part, but the police can safely ignore it in their investigation. Instead of attributing this to emergence, or downward causation, it could be explained by contextual language. Removing irrelevant factors from the equation, is an important technique of critical assessment. I do not think that this same-level coherence of explanation, is interchangeable with proper downward causation.

Although the application of downward causation in emergent properties lends it some hefty explanatory power, it also comes with some issues. The most famous argument against it in philosophical debate might be Jaegwon Kim's variant of the exclusion argument, in which he argues that downward causation does not provide higher order

properties with the causal autonomy that it supposes. The problem, is comes about because of the relationship between the lower and the higher-level properties. If we have some higher order property M2, acts causally on some lower level property, say P, we would have downward causation. However, as Kim points out, M being an emergent property cannot be so with being reliant on having instantiated base properties.

Therefore, to instantiate M2, we first need to instantiate its lower level properties M1. What then, stops us from saying that M2 did not really cause P, rather it was caused by M1, because M2 could not be instantiated without M1. On this argument, the higher-level property, M2, become causally inefficient, something of an epiphenomenon. This would be disastrous for any emergent theory that aims to have causally autonomous emergent properties, or even emergent properties that should be accounted for in a causal theory at all. Because not only is M2 not causally efficient, but the causal model works perfectly fine without it. We could just say that M1 causes P, one the grounds of transitivity. The other option is that both M1 and M2 causes P, which would be overdetermination.

Hong Yu Wong challenges Kim's causal exclusion argument by saying that an emergentist may still choose systematic overdetermination, concluding that Kim therefore must condemn it, because it is coherent with emergentism (Wong, 2010, s. 13). Whether or not downward causation can survive the causal exclusion problem, it remains a big part of the language used to describe emergence.

Downward causation and the exclusion argument are not of significant importance for my later argument, other than being a different argument that echo my reservations. But still it remains a very central concept to emergence, a concept which in many ways defines it. If we can allow for downward causation, without resulting in problematic epiphenomenon like Kim suggest it does, it will greatly strengthen the claims about emergent causal autonomy and its non-reductive property.

### 2.3 The Case for a Layered Nature

One of my main points of contention with the emergentist position is whether there is any good reason to assume that the ontological layers that emergence suggests, exists at all, or if it is useful enough that it could reasonably be employed, regardless of its ontological status. We could imagine a state of the world where we have pragmatically

good grounds to accept that nature is layered, or at least that our model of nature should be, even in the case that it is not. This gives us reason to investigate both the case for the truth of such a layered nature, and what it would look like, as well as its usefulness.

I consider an emergent view of science to imply anti-reductionism, although this is not always a conceded point. To show why I believe this to be true, we first need to explain what is meant by reduction.

Anderson also holds that this kind of reductionism does not entail what he calls constructionism “The main fallacy in this kind of thinking is that the reductionist hypothesis does not by any means imply a “constructionist” one: The ability to reduce everything to simple fundamental laws does not imply the ability to start from those laws and reconstruct the universe” (Anderson, 2008, s. 222). He argues this is the case, because there are macro-structures that are practically impossible to see from the perspective of the lower level. This seems to mean you can be both reductionist, and have a layered emergent view on special sciences.

But in general, emergence is anti-reductionist in some sense, because for the special sciences to be emergent, they need to be causally separate, and therefore causal reduction cannot hold. If this was not true, then the causal powers of phenomena on a social level would be determinable from physics, and not something that could be acceptable in a theory of emergent science. Emergentists must hold higher order phenomena to be causally distinct from its base.

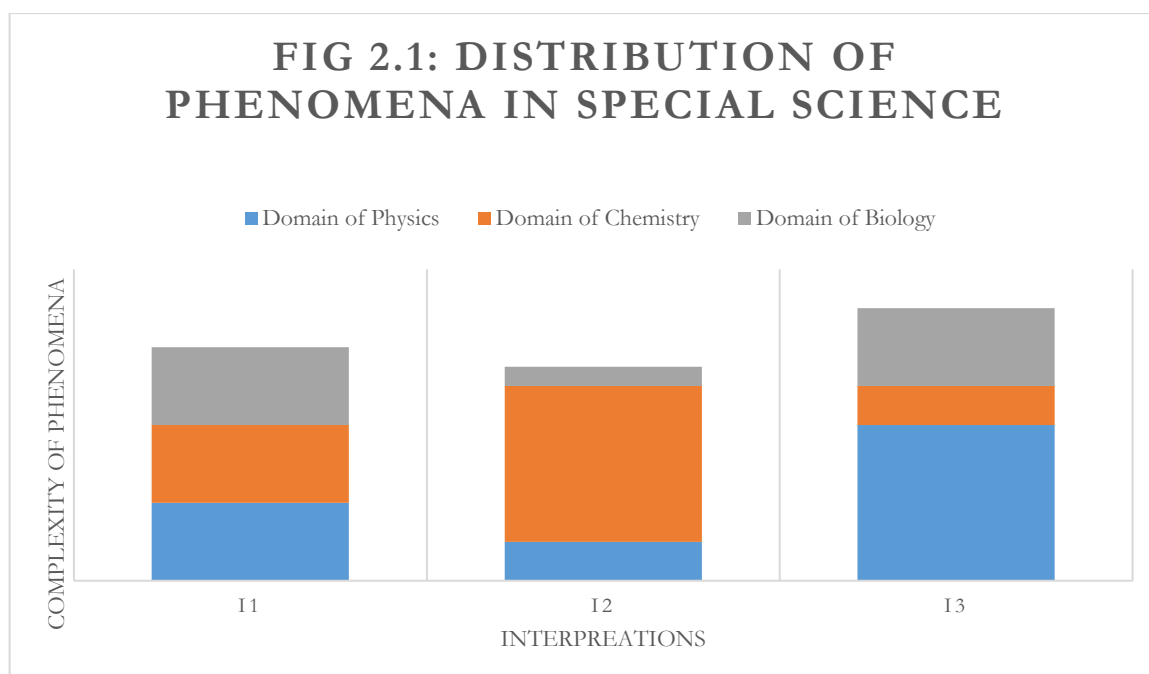
### 2.3.1 Outline for an emergent hierarchy of special science

The view of nature and science as emergent structures, is not one that is extensively developed in the main works of emergence, and for my use, does not need to be. I would, however, like to try my hand at a simple outline that an emergent structure of special sciences would need to have, inferred from other assumptions of emergence, and using the few sources I have of similar systems, most notably that of P.W. Anderson (2008) and George F.R. Ellis (2006), we will return to these two later. I propose a simplified model as consisting of assuming the following:

1. The most basic special science is physics.

2. Each special science starts where the higher order properties of the phenomena of the science before it begins, except the most basic level.
3. The levels all describe different (novel) kinds of phenomena, although they might describe the same object(s) they differ in order of complexity.

When we have these three steps, we only need a working definition of what counts as emergent phenomena, and it should effectively divide all natural phenomena neatly into being the domain of its respective special sciences. That is, it takes an interpretation of emergence, and produces the extension of the different scientific domains. This should be true for both ontological and epistemic accounts of emergence. The scientific domains that result are emergent sciences, because if the emergentists are right, this would be the only proper way of dividing the natural sciences. A natural kind science has the scope of complexity of solving scientific problems of a certain kind, with no overlap in subject matter with the other natural kind sciences. Different sets of interpretations of what constitutes an emergent relationship would produce different divisions of what belongs to which special science, as visualized in fig 2.1, with physics, chemistry and biology.



In fig 2.1, the first interpretation (I1) is a division of the three special sciences in question, that is fairly even, all three having similar scope in number of phenomena within its domain. The second interpretation (I2), has a sparser domain of biology and

physics, but a large domain of chemistry. The third interpretation (T3) is the opposite of the second. The models in fig 2.1 is not meant to realistically portray a model of scientific hierarchy, or try to reflect any real division in domains, but rather to give an indicator to the dynamics of the boundaries that are generated with different interpretations of emergence. Neither are any of the interpretations T1, T2 or T3 real interpretations. The true model, if there is one, would presumably be more complex, and perhaps branching into various areas of science, but this does not affect this discussion much, because rather than focusing on what the models look like in the end, I am interested in how their internal borders would form, so this basic model only serves the purpose of illustrating that.

The relationship would be such that all phenomena covered by the special science of biology would emerge from the phenomena in chemistry, and all chemistry from physics. On this view, we exclude the possibility of having an emergent phenomenon with lower level and higher level properties being covered by the same domain of special science. Such that all emergent phenomena have their base in one special science, and the higher level in another. In the case of the game of life then, all the identifiable levels of emergence, should be considered members of different special sciences subject matter. This might be problematic, because it might produce too many special sciences. This could be avoided if we allow for two types of emergence, one that has the base level and the higher level in the domain of different special sciences, and one that has both in the domain of the same special science. I will however not concern myself with this, and assume that all emergent phenomena are of the former kind.

The different domains of the special sciences that are produced by the definitions of emergence, should be non-overlapping, and that the union of the extensions of all the special sciences should include every possible phenomenon, on every possible level of description. The various levels are also presumable mainly different in levels of description, as they can describe the same set of materials. As the borders between the domains of the sciences is non-overlapping, it follows that they all have a clear, non-vague extensions of the phenomena they describe. It would be incoherent to assume that we could have domains that covered all phenomena on all levels of descriptions, never overlap, but still do not have a clear extension. It is a trivial point that there can be no

phenomena on any level of description that is either both a member of a scientific domain, and isn't, or is not a member of any domain. Every phenomenon must be a member of one, *and only one*, special scientific domain on a specific level of description.

I still have not touched upon whether this emergence of sciences should aim to be a weak or a strong emergent, ontological or epistemic, and from the inclusion of a non-reductionist claim in my definition, this can be taken to mean that I think this emergent view of science is or should be a strongly emergent ontological one. I don't think this is the case, because the causal autonomy of higher order phenomena is a requirement for both weak and strong emergence, and that causal non-reduction is required to obtain this autonomy, I hold that this would be a requirement in both a strong and weak view of the emergence of special sciences.

The requirements for natural kind sciences to be strongly emergent would be its underivability, even in theory, from its base. I think a case could be made, and it has been made, that some sciences are strongly emergent on others. P.W. Anderson, for instance, shows us that in his field of many body physics, it is not possible to derive the level of symmetry in a complex chemical system from the physics of its components (Anderson, 2008). Supposing this is true<sup>5</sup>, then we might have a strongly emergent relationship between many-bodies physics and chemistry. Although the same kind of strong emergent relationship need not be true between say psychology and social science. A strongly emergent hierarchy of special science would then probably be sparser than a weakly emergent one, because it would probably include fewer relationships as emergent.

In any case, the best course of action is to take weak emergence as the model of emergence to use for the purposes of emergent hierarchies. I have two main reasons for this. The first is that I think there is good reason to believe, along with Chalmers argumentation, that strong emergence would be sparse, and that the only candidate for strong emergence is consciousness. If this is the case, then the claim that emergence provides us with a hierarchy of phenomena that lines up with the special sciences is trivially false, as physics and consciousness are not the two only areas of special science.

---

<sup>5</sup> Andersons argument for this is too technical and specific to the field of physics for me to deal with, so I will let it stand uncontested. My own argumentation does not rest on its truth.

The other reasons are that even if strong emergence did provide us with what we needed, our model of how hierarchies are formed, and what they look like, would not need to be altered much. This is because strongly emergent phenomena surely also are weakly emergent. So, any critique raised against the hierarchies formed by weak emergence, can be similarly raised against the hierarchies formed strongly.

When talking about these layers being emergent upon each other, it is using emergence in a slightly different way than when talking about individual emergent phenomena. Because, what would it mean that chemistry emerges from physics? Does it have to be true that all the properties of chemistry are novel and unpredictable from physics, or just that certain hallmarks of chemistry are. Or that the laws of the one, emerge from the other. When I say that one special science emerges from another I mean that the phenomena being described by that special science, emerge from the properties in the special science below.

The view of science as a series of emergent layers in a hierarchy is fully embraced and developed in George F. R. Ellis 'On the Nature of Emergent Reality', in which he supposes that there is a given structure of special sciences, and that this is emergent. This system is first and foremost the hierarchy of natural systems, and the order of the sciences is made to match. The difference between this account and that of P.W. Anderson is that Ellis suggest that this hierarchy is there in an ontological sense, while P.W Anderson only remarks that we *could* organize science in this way. Ellis puts forward the following emergent hierarchy of structure, going from the highest level to the lowest (Ellis, 2006, s. 80):

*Sociology/ Politics/ Economics*

*Animal Behavior/ Psychology*

*Botany/ Zoology/ Physiology*

*Cell Biology*

*Biochemistry/ Molecular Biology*

*Molecular Chemistry*

*Atomic Physics*

*Nuclear Physics*

*Particle Physics*

These levels have different modes of description, relational hierarchy and phenomenology, meaning that we cannot describe the higher-level properties in the language of the lower, and vice versa. Politics is not describable on the level of description of Cell Biology, Particle Physics in Botany and so on. The characteristics Ellis provides for the emergence that produces this kind of hierarchy is not so simple to get a hold of, and not very important to get into, but they consist largely of not being straight forwardly reducible, meaningful top-down actions and goal seeking system properties, all arranged into five different degrees of emergent strength (Ellis, 2006, ss. 99-100). The reason it is not that important to get into, is that I want to focus this part on the hierarchical systems themselves, and not so much the interpretation that provide us with them. Ellis's model is supposed to have clear extensions of the phenomena that are included at each level in the hierarchy. The language used to describe the relationship between these levels is largely that of computer science, more specifically object oriented programming. On Ellis's model the emergent structures are governed by 'abstraction, encapsulation and inheritance' (Ellis, 2006, s. 86). A higher order of emergence, is something of a higher level of abstraction, such that the lower level details are washed away through coarse-graining.

An immediate suspicion that can be raised against this kind of model is that it seems to take for granted that a lot of the language and methodologies that is used in artificial systems such as object oriented programming is applicable to nature in any way greater than analogy. The takeaway from Ellis's approach in the later context of vagueness should be that there is a great deal of ontological commitment happening on the grounds of methodologies and modeling borrowed from computing and physics. The connection between the seeming emergent phenomena we observe, and their ontological truth is simply assumed, and the models of computing are used to create sharp divide between these levels.



S. Rasmussen, N. A. Baas, B. Mayer and M. Nillson gives us a more in depth, yet simpler, variant of the kind of emergent hierarchical structure proposed above in their “Ansatz for Dynamical Hierarchies”. The aim of their model is more specialized to molecular biology, but they do not exclude that it can be a guide to a more general hierarchy for other systems of complexity, although they do not claim that it is. In their study, they focus on three orders of emergence, from monomer, to polymer to micelle. Each having a functional emergent description of their own, as seen in Fig. 3 (Rasmussen, Baas, Mayer, & Nillson, 2008, p. 324).

<i>Level of Description</i>	<i>Molecular Structure and Order</i>	<i>Observed (Emergent) Functionality</i>
Level 3	micelle (3rd order)	inside/outside, permeability, self-reproduction
Level 2	polymer and water (2nd order)	elasticity, radius of gyration

FIGURE 3, FROM (RASMUSSEN, BAAS, MAYER, & NILLSON, 2008, P. 324)

These layers of emergence are causally efficient both up and down, the lower level functional properties have causal power over the higher level, and the functions on the higher-level causes changes in the lower. This, along with the observation that these systems are self-organizing, is why they term this hierarchy to be dynamical. Another feature worth mentioning about this hierarchy is that it is simulation predictable, and part of the project is to recreate these systems in simulations.

Now the interesting part for our discussion is how this model separates between the levels, and the manner of definition about the functional descriptions of the emergent properties. The only operational characterization of emergence in their model is that it is novel, so this explains why the emergent levels are very dense, because novelty alone is an easier condition to satisfy than novelty together with non-derivability and non-reducibility. The only thing required for emergence here is that it “[...] does not apply at any lower level” (Rasmussen, Baas, Mayer, & Nillson, 2008, p. 309), this is then refined in an explicit formal definition.

I think this model reflects the kind of bigger emergent hierarchy that the emergentists might want to work, and is the inspiration for my own take on this. The notable difference is that the emergence they operate with, is what we could characterize as nominal emergence, the weakest sense of emergence possible. If a model of emergent

hierarchies were to reflect some sort of ontological emergence in nature and special science, I think it would have to be at least a strong variant of weak emergence for it to have the theoretical punch required to justify the model, because a mere nominal emergent hierarchy is still both reducible and deducible, and not the kind of paradigm shift we would like. I do not hold that these emergent hierarchies are possible in the strict sense of them referring to some ontological natural order of special science akin to the kind proposed by Ellis, but I see them as in the case of the more limited model suggested by Rasmussen et al., as a scientific tool for simulating and understanding complex behavior. I will argue for this position in the next chapter, now that the stage is set.



## Chapter Three: Vagueness in emergent hierarchies

So far, we have discussed the features and issues found in the literature on both vagueness and emergence. In the following chapter, however, we will depart from the known debates and venture into more unknown territory, with my own application of vagueness on emergent theories. Although the suspicion that there is vagueness in emergence is not discussed in any depth in the literature, it is still not entirely untouched, as both Jaegwon Kim and Mark A. Bedau acknowledge to some extent that the boundaries between higher level and lower level phenomena be vague. The arguments that I put forward here, however, have not been put forward before to my knowledge.

### 3.1 Vagueness in the characteristics of emergence

If one has prior knowledge of vagueness when reading about emergence, as I hope anyone reading this now has, one might have a nagging question in the back of one's mind. Could emergent phenomena be vague? After all, a lot of the "magic" of emergence revolves around that something is suddenly something else. This great divergence, as we've discussed in chapter two, seems to insinuate the exclusion of vagueness. Make a minor change in the right place in a brain, and we might not get consciousness at all, suggesting that there is a point where it goes from heap, to non-heap, in just one seemingly insignificant step in the lower level properties.

My aim is to show that emergence is vague. To do this, we need to establish vagueness in emergences defining characteristics. As established in the first chapter, the most common defining features for emergence are non-derivability, novelty and non-reducibility, in various combinations and levels. Vagueness, as shown in the second chapter, is a term that applies to predicates with an unclear extension. If I am right, both emergence and its compound features are vulnerable to sorites paradoxes, as they all have unclear extensions. Then, by going further, I want to show that emergence is vague, because its defining features are vague. Emergent levels generated by emergence are then bound to be vague, and emergence is therefore not useful or suitable as a guide to layering nature. So, in this chapter I will discuss the alluded vagueness in non-derivability, novelty and non-reducibility, and lastly the vagueness of emergence itself. Lastly, I will show what this means for emergent theory.

My investigations into this matter is not an attempt to refute the possibility of emergent properties, or reveal some internal inconsistency to the ontology and epistemology of them, but rather to show that it is naïve to expect the philosophical analysis of emergence to provide us with any ontological or epistemic tools for uncovering what is and what isn't emergent.

### 3.1.1 Novelty

Novelty seems to be the weakest of the possible characteristics of emergence, because it is difficult to say that any aggregate, iterative or complex structural property of some basic property is not novel. Imagine the following case: we have a set of base features F1 from which emerges an emergent function description F2 where the description of F2 is not identical to F1. This would be a clear case of novel emergent property. For it to be untrue, the description of F2 and F1 must be identical. For the novelty of the set of properties F2 to be vague, it need to be unclear whether the properties of F2 belongs in the extension of “novel from F1”, or “different than F1”. If we think about this in terms of cellular automata, a structure made from cells, would have to show features that are not present in the lower level. Let us try the sorites method on cellular automata, in the form of the repeating question.

1. Does a single cell in the game of life have novel features that a single cell does not have?
2. Does a system of two cells, arranged next to each other have novel features that a single cell does not have?

The answer to the first question is undoubtedly no, as a single cell should not be able to have different properties than itself, all else being equal. The answer to the second question though, interestingly, is a bit more debatable. Two cells arranged next to each other in the game of life, does have a different functional description than a lone one has. For one, they form a line, and not just a point. And no matter how they are arranged, you could always draw a straight line between them. This feature is novel in the way that is different from the features of only one cell, but it's a quite mundane feature. It's not like two cells next to each other have consciousness, while the single cell does not. The type of novelty is much less radical, the single cell also occupies geometric space, but in a less line-ish way that the double configuration does. It does not however

have different functionality in the sense that in T1 (Timestep 1), both one cell and a system of two cells will be gone, because in the game of life a cell needs two or three neighbors to survive to the next timestep, which none of the cells in neither of these configurations have. From this, we see that there is no novelty on the “design level” of gliders and eaters, merely in the geometric properties of the first timestep. It seems we can make the case of it being both weakly emergent (systemic geometric properties) and non-emergent (design-level properties), depending on if we view weakly emergent phenomena as diachronic, synchronic, and the kinds of differences we require to fulfill the condition. If we continue this increment to three:

3. Can three cells, arranged in any way, have novel features that a single cell does not have?

Here we can observe a similar geometrical novel feature. If you draw lines between them, they will form a triangle. But more systemically than that, and on the design level, three cells in the game of life is the least number of cells that you could have that can be arranged in such a way that they do not immediately die in the next time step. If we arrange them closely together, they would, in the next timestep, cause a fourth cell to come to life, creating a 2x2 block of cells which then would stay alive forever, assuming no other cells encounter it. Or we could arrange them in a line, which would cause it to seemingly flip the row in a 90° angle from the middle, even though what happens at the cell level is that the cells at the left and right end die, while the cells at the top and bottom of the middle cell, get born, while the middle cell survives. This will alternate for all foreseeable future, *ceteris paribus*. So again, we have novelty, in that we have a mechanical feature in the game of life. But are these different emergent levels, and are they emergent in an equivalent way?

In fact, it seems that for every step of complexity we add, there will always be novelty. This feature of emergence turns out to be extraordinarily weak, in that we can always argue that something that is not identical to something else, has novel features, meaning different but not identical features. If this is true, then novelty is so weak that it’s problematic to find any claim about emergence where the novelty condition is not satisfied. Taken like this, we can deem it to general to be problematically vague, as it would be like having every number of grains, other than zero, be a heap.

But I think that this weakness of definition can relate to the problem of vagueness. For novelty to have a proper function in a definition of emergence, it needs to be strengthened. To do this, we can specify the kinds of things that count as new. Let's say, in the case of the Game of Life, that we could only like to have layers of description about the kinds of functionality that are unique to the Game of life, like the "design level" emergence of gliders and eaters, and not other kinds of novelty, like geometrical properties and the like. This decision would be useful in clearing up the kinds of phenomena that we should investigate in the Game of Life, and exclude the other things that are not unique to this system, but it is also an obviously arbitrary decision to include the one, and exclude the other, because novelty gives us no such decision-making power on its own, as we have seen. The reason for this arbitrariness is the vagueness of the specification of our definition of novelty, in this instance, "the kind of functionality that are unique to the Game of Life". The extension of the predicate F, where F is "the kind of functionality that is unique to the Game of Life" is unclear, because there are borderline cases of it. An instance of a borderline case of F is the kinds of geometrical properties that we discussed previously in this section. The phenomena of geometrical shape from point based structures is not unique, but the application of it to the Game of Life is, and so it remains as not a clear case of F, and neither a clear case of not F. I believe that vagueness of specifications of this kind are quite general, and should be considered at least partly the reason for the requirement of arbitrary decisions to what counts as something, in this case novel, and what doesn't.

It should also be noted that the specification of observable properties is somewhat arbitrary. This is quite analogous to the arbitrariness involved in the specification of the primitive objects and their interactions. So, a certain kind of phenomena can be viewed as Nth-order emergent only relative to arbitrary choices about what counts as first-order objects and properties, what counts as second order objects and properties, and so forth. (Rasmussen, Baas, Mayer, & Nillson, 2008, p. 311)

This admittance is not of significant importance for their project of dynamical hierarchies in molecular biology, because it is used as an epistemic tool for understanding structure in nature, and not as an argument for an emergent ontology. We should also remember that the only operational characteristic of emergence in this

model, is novelty, so this arbitrariness-problem does not imply that more complex characterizations with non-derivability, non-reducibility and supervenience have the same problem.

Whether I am right in my suspicions that vagueness help cause this need for arbitrary decision-making or not, the condition of novelty for emergent phenomena seem to have a two-fold problem. Either the power of the condition is trivially weak, such that every phenomenon, except identical ones, are deemed to be emergent. If this is the case then novelty is as useless in the definition of emergence, as “being a something” is for defining a watch, because they would both be trivially true in every case that matters. Or, on the other hand, if we try to make the novelty conditions more rigid, it is bound to be too contextually arbitrary for it to be ontologically plausible.

### 3.1.2 Non-derivability

When analyzing the vagueness of unpredictability or non-derivability, we should first make sure that our usage of underivability is compatible with the use in emergent discussions. One clear case of underivability comes from Mark. A Bedau as mentioned in chapter two. This is unpredictability (or underivability as he puts it) except by simulation (Bedau, 2008). In this understanding, a phenomenon would be perceived as unpredictable from the knowledge of its component parts, in the case that the only way of doing so would be to go through with the process, step by step, in a simulation, without being able to skip ahead. The other is the strictly non-derivable variant, associated with strong emergence. Considering this being a discussion about emergent hierarchies, and my position that weak emergence is the more relevant variant of the concept in that context, I will focus on non-derivable except by simulation as the core of the non-derivability characterization. The industrial strength non-derivability used in Chalmers’ strong emergence, is as we have seen, so strong that there is only one case that satisfies it, and that is consciousness. Non-derivability to this degree is therefore not useful if we want to allow for more variety of emergent phenomena. So, I will rather investigate Bedau’s non-derivability in weak emergence.

To determine if something could be vaguely non-derivable is a challenging task. But a good place to start is to try to stress the borders of the terms extension. We have seen how this kind of non-derivability works in a cellular automata system, with the only



exception to the non-derivability being simulation. An important aspect to investigate is what kinds of simulations that are valid exceptions. In the case of cellular automata, the simulation is extraordinarily accurate. It is a one-to-one simulation of the exact conditions of another cellular automaton. There are no differences in the structure or content of the simulation versus the event being simulated. This is not normally the case in simulations. Simulations tend to abstract away the unimportant aspects of the process and focus on functionality. Often, a lot of processes gets simplified in a simulation, while leaving the overall mechanics intact. For instance, if we were to simulate the impact of a car crashing into a brick wall. It would not be necessary to simulate all the individual molecules that the gasoline consists of, for the purposes of this simulation we only need to assume that the car is moving forwards, without concerning ourselves with the specific hydrocarbons included in the gasoline mix.

The kind of simulation proposed by Bedau is more of a duplication rather than a simulation. Everything is completely identical, down to the last cell. If we follow through with this understanding of simulation, then running a computer program downloaded from a website, would be to simulate of the original program. When he talks about simulations then, I will take this to mean simulation by duplication. Once we see the simulation in this light, a lot of other phenomena can count as being underivable except by simulation (by duplication). Deriving the swaying of a tree and its branches in the wind from the structure of the trees fibers and the movement of the air around it might well be impossible by other means then replicating these exact conditions in a computer simulation or in physical form, which might prove to be practically impossible. The waving of the branches would also count as a clear case of novelty from the base level, and I think also non-reductive, in the emergentist sense. Yet, it does not seem intuitively emergent in the strong weak sense that we would like. It might be nominally emergent, as that doesn't take much other than that the waviness of the tree could not be a feature of the tree fibers and the environmental conditions, and even this might be wrong. That the condition of underivability except by simulation is vague is not lost on Bedau however:

“I have been speaking of underivability except by simulation as if there were a sharp dividing line separating weak emergent properties from merely resultant properties, but

this is an oversimplification [...] One can define various sharp distinctions involving underivability except by simulation, but focusing on one to the exclusion of the others is somewhat arbitrary” (Bedau, 2008)

He notes that you can have emergent phenomena that are strictly shortcut-derivable, but needs to be simulated, underivable except by finite simulation and underivable except by impossible or infinite simulation (Bedau, 2008).

The vagueness in the non-derivability characterization is not really a case of the actual derivability, as in the strict derivability, being especially vague. But if we want to provide any special condition to limit the derivability, like “except by simulation”, that is connected to the epistemics of derivability, we quickly encounter problems with vagueness. The extension of things that are simulatable, is not a clear one, and that rubs off on the non-derivability.

Because of the points mentioned above, I think it is fair to say that defining weak emergence in terms of its underivability except by simulation, does not shield it from vagueness, on the contrary it seems to invite it in. As this is by far the strongest unpredictability theory for weak emergence, I conclude that unpredictability in weak emergence is vague.

### 3.1.3 Non-reducibility

The cousin of non-derivability, non-reducibility is a more problematic condition for the emergence project overall, because of the exclusion problem (covered in 2.4), and the lack of it in weak forms of emergence. As something to guarantee the causal autonomy of higher order phenomena, it is supposed to provide emergence with causal efficacy, allowing higher order phenomena to provide causality that is separate from its base level. Nevertheless, whether there are viable solutions to the causal exclusion of higher order properties or not, the non-reducibility characterization is too vague. To show this, we first need to disambiguate what type of non-reducibility we are looking for.

As discussed in 2.1.3, there are multiple ways in which something can be reducible, and therefore also non-reducible. I will take the theoretical reducibility to be of most importance to us, because it is what provides us with the wall of explanatory separation between the different special sciences. We should also consider causal reduction as a

strongly in the scientific theory context that we are focusing in on. To reiterate: causal reduction in emergent phenomena is when a higher order causal property is definable in terms of its lower level causal properties, just like Kim supposes in his causal exclusion argument. Theoretical reduction is when a higher order law or theory, like for instance Newtonian physical laws, are reducible to its lower level properties, the laws of quantum mechanics. If the higher order law or theory could be removed, without removing any theoretical ability, then it is reducible, it is non-reducible if this cannot be done.

In the case of the first type, causal reduction, is a feature of most accounts of strong emergence, but not weak emergence. Bedau's account of weak emergence does not include causal non-reducibility, neither does P.W. Anderson. Yet the latter account is represented in both. For the second type, the theoretical, or explanatory reduction, we can again turn to Jaegwon Kim for some helpful analysis. In *Making Sense of Emergence* he describes the process of this kind of reduction for the purposes of identifying the relation between the base and the emergent property, as we discussed in 2.1.2. To do these kinds of functional definitions of emergent properties, we will need to make certain definitory decisions based on our, probably inexact, knowledge about it. If we have some emergent phenomena E, we must decide where it begins and where it ends, for us to make a functional description of it. Let us suppose we have some emergent property E from base property P that we need to describe, so that we can work out whether E is reducible to P or not. Because E is in some way or other *made up of* P, our description of E should be functional, to separate the two, as I have argued before. The systemic properties of P might have many possible functional descriptions, that might be variously degrees of emergent on P, and some that are not emergent at all, but we would like to isolate a specific emergent property. We then must decide on some cut-off point somewhere to include all the aspects of the apparent phenomena properties that belong to E, and exclude all the ones that do not, based on our best knowledge. These decisions are based on what we want to explain, the scientific context and our current knowledge, but to some extent also arbitrary. Kim acknowledges this point:

“We should keep in mind that such conceptual decisions can be and often are based on empirical knowledge, knowledge of the causal/nomic relations in which E is embedded, and can be constrained by theoretical desiderata of various sorts, and that in practice the

boundary between what's conceptual and what isn't is certain to be a vague and shifting one." (Kim, *Making Sense of Emergence*, 2008, s. 133).

After these decisions are made, we can see if this functionalization D of E is reducible to P. The role of vagueness in this model is easy to see. When giving the emergent phenomena some functional name, we imply a certain, probably vague, extension of phenomena to be included in this definition. If we take the glider in the cellular automata as an example. Our functional definition of the emergent property realized by a certain initial condition, could be that it moves diagonally in some direction, depending on the orientation of the glider. This is not the full extent of its causal powers though, as it can also interact with other configurations of cells in a practically infinite number of variations (it can collide with a x configuration at y place in z time etc.). All of these should not be included, as mentioned before, but some of them probably should. The extension then, of "the function of E" is very unclear, and thusly vague, which is why we must take the context of our specific investigation and other desiderata into account, for us to make it a viable functional definition, but these are in a very real ontological sense, arbitrary, based on our interests as observers. I want to underline the point that it is not just the case that the extension of "the functions of E" is incredibly large and unwieldy, although it certainly is that. We have no way of knowing, without it being arbitrary, whether "if a glider meets an eater in this configuration, at this place in t234" should be included in its extension or not.

This point also goes for the previous sections on novelty and non-derivability, as the functional description of the emergent property is required in all three characterizations, to define or encapsulate what the emergent property consists of. To work out whether something is a novel property of something else, we first need to establish the functional description of the thing that is supposed to be novel, as well as the description of the thing it needs to be different than. The process of describing or defining emergent properties lays bare its inherent vagueness, as the emergent phenomena we observe have vague extensions of what is included in the phenomena and what isn't.

### 3.2 Vagueness in emergent levels

Following my claims of vagueness, there are some problems that must be accounted for. The version of dividing complex phenomena into the domain of different special

sciences that I gave in chapter 3, is one that aims to produce specific extensions. If you have an interpretation  $S_1$ , it needs to produce a hierarchy of phenomena that have no borderline cases to be useful, because it separates between the extensions of the special sciences with emergence, producing no overlap in their domains. So, no matter the interpretation of emergence we are operating with, the aim is for it to produce non-vague extensions of the domains of the special sciences. It is crucial that these borders can be clear.

I have shown that all the prominently featured characteristics of weakly emergent properties are vague, and therefore I hold weak emergence to be itself vague. This point is more general than just being about a specific set of defining characteristics of weak emergence being vague, like for instance Bedau's underivable except by simulation (Bedau, 2008), or Chalmers derivable but surprising condition (Chalmers, Strong and Weak Emergence, 2006). Both of these, along with every other configuration of characteristics of weak emergence, will be vague, because dividing up phenomena in the way that they attempt to do, invariably produce unclear borders.

The core problem, that causes vagueness to be a recurring issue in the definition of emergence, is that the whole concept of emergence revolves around its divergence, which is the same place the sorites paradoxes get their confusing nature from. Take the Turing-machine that has been created in the Game of Life. This construction, is realized by the exact nature of its base properties. If we remove only one cell, a fraction of a fraction of the whole puzzle, the contraption fails, and the universal computer stops emerging. The system will still have a lot of other high level features, like eaters and gliders on the design level, but this more impressive computing feature is gone. When we try to identify the difference of the system  $S$  that contains a universal Turing-machine, and the system  $S-1$  that contains a universal Turing-machine, minus one piece (assuming all the pieces are essential for the machine to function), we naturally think that there must be a major change in the nature of the system from  $S$  to  $S-1$ , because how else could such a substantial change on the high level come from so little a difference on the low. Suddenly, it stops being a heap, so to speak.

The strongest reason to consider emergence to be vague is that the functionalization process that is required to deal with these lower and higher level properties is inherently

vague, as discussed in 3.1.3, so no ontological, nor epistemic clarity should be expected to come of the process. This shows us that while the vagueness in the emergent characterizations are relatively mild, the vagueness related to the higher order phenomena themselves is inescapable. While identifying which phenomena in the world are emergent and which are not, seems quite simple and intuitive, the formalization process is not as simple. It is when we try to determine the true extension of some causal prowess or functional property of a glider in the game of life, or a school of fish moving in unison, that we realize that the emergence is not so easily captured.

Going back to the case of special science: it is quite clear that if we have no non-arbitrary way of determining the extension of emergent phenomena, it will be impossible to determine what phenomena pertains to the different special sciences, merely on the grounds of emergence. The hope was that with a rigid system of emergent levels, we could be able to allow for a specific extension of the domains of the special sciences, on different interpretations of emergence. But if the vagueness of emergence, as I have argued, is not specific to any definition, but rather to the definition/decision process of what the higher-order property is, then this would be impossible on any interpretation. Seeing as the vagueness is inherent in the process of identifying emergent, such a non-vague scientific hierarchy of emergent levels will not be possible, and therefore will every interpretation of emergence, yield vague levels. The consequence is quite clear, we cannot have a hierarchy of special science based on emergence, if that system requires clear extensions. Either we must do away with the requirement of clear extensions of the levels, or reject with the idea altogether.

### 3.2.1 Criticism

There are several possible critiques that could be raised against my collected arguments against emergence that I would like to discuss. As I see it there are four main lines one could take:

*Rejecting functional descriptions of emergent properties.* As I have shown, the process of formalizing emergent properties in terms of its functional content is problematic because it involves arbitrary cutting off fuzzy borders. This seems to be solvable by changing our model of descriptions for emergent properties, or simply specifying what kinds of functional descriptions we require, in a non-arbitrary way.

If we don't want to change our model of description, but simply try to amend the weaknesses of the functional one, I think an effective way to start would be to narrow down the kinds of functionalities that we want to include in our description of an emergent property. The root of the problem is that there is no non-arbitrary way of discriminating the kind of functional description we want, and one that is too large. We could give a much stronger definition like this: 'The functional definition of an emergent property should include all the properties that allow us to derive all of the emergent properties causal relations, and no more'. Defined like this, we exclude all the relational and causal functions that bloated up our first definition. In the case of the glider in the cellular automata that we used before, the functional description would be something like the movement of the glider, plus the placement of every cell in all the time-steps of a gliders four phases. Armed with some formalization of these, we can work out all the possible interactions with other configurations, if we know their functional description.

There are several problems with this model, though. The first is that the interactions of the system are no more derivable from our new functional description, than the emergent property is from the base conditions. This being a cellular automaton, the interaction between one functional model and another is still non-derivable except by simulation. Also, this definition would essential allow us to ditch the whole functional aspect, we could merely include the second part with the placement of all the cells in the four phases of the glider, which is essentially the base condition for the emergent property. If we dismiss non-derivability except by simulation it would further lead us to not being able to discriminate between what counts as the base property of the glider, and what counts as the emergent. In the end, we are left with more problems than we started with.

Even though my reply to this is specific to a specification of the functionalization process that has not been proposed by anyone as far as I know, I think the points we can draw from it remain valid. It is very problematic to specify what counts as an emergent function without it being arbitrary or dependent on our investigative interests.

*The linguistic critique about the vagueness of the process has no bearing on the ontological truth of emergent layers in nature or special science.* This objection can safely be bundled together with the claim that if we a consistent account of epistemic vagueness, my reservations about

emergent levels should dissipate. Objections of this kind are valid, in some sense, but I think misses the point of my argument. I am not arguing that because of the inherent linguistic vagueness in defining and using these terms to refer to specific non-vague emergent phenomena and basal relations, that these relations cannot exist. Rather my argument says that we have no reason to think that they do, and even if they did, we have very good reasons to think that our theoretical understanding of the phenomena would reveal them to us.

If we attempt to solve our theoretical difficulties by accepting epistemic vagueness à la Timothy Williamson, I think the theoretical power this provides would be somewhat disappointing. Even though, in an epistemic view, there such a thing as bivalence, and every proposition is metaphysically speaking, always true or false, vagueness is still prevalent in conceptual representation (Williamson, 1994, p. 249). Thinking that epistemic vagueness can provide us with the ability to separate emergent from non-emergent from basal conditions would be an ungrounded belief. The only true refuge it can provide is the ontological possibility that these exists, because there can be no such thing as ontological indeterminacy.

This does not, however, doom us to having to accept some sort of nihilistic conclusion, that we cannot know anything because our descriptive language by which we categorize is inherently vague. I think there are theoretical definitions of phenomena that are not vague, and some that are not problematically vague.

*The vagueness is not problematic on a purely epistemic view of emergence.* I have held back on separating too much between the epistemic and ontological views of emergence in my argumentation partly because the points I have made are applicable to both positions, and second, I don't think the difference between them is that big. To the first point, as I have already mentioned, the vague extension of emergent phenomena is not only a problem for the ontology of the emergent theory that produced the extension. It is also very much an epistemic problem, as we cannot refer to the specific higher level properties without running into issues with vagueness, leaving the epistemic aspect severely limited. It does not matter that the theory does not claim that these levels or phenomena exists, but merely that they seem to exist, if we cannot provide any non-arbitrary extension for those levels and phenomena. In fact, the fact that we cannot



reliably tell what is emergent and what isn't, leaves us with very little motivation to keep any emergent model at all, if the ontological truth of emergence is not assumed or argued for. If, however, the epistemic account for emergence is intended to be completely void of positive content, and is just intended as a non-ontologically committed abstract concept, with as little root in natural phenomena like statistical or economical models, then it could be scientifically viable as long as the model is not confused for the world.

Lastly, I would like to address the point that *accepting Supervaluation solves vagueness, and therefore also solves the vagueness argument against emergence*. This line of critique is similarly to the suggestion that accepting epistemic vagueness would help, misguided. The result in Supervaluation logic to solve vagueness, is simply to make room in logic for indeterminacy. This means that, contrary to epistemic vagueness, it is perfectly possible for something to be neither emergent and non-emergent. This is not at all helpful if we want to establish the theoretical groundwork for weakly emergent phenomena, as a lot of the wanted extension of it would be undetermined, and therefore would not have the kind of power we would want. If we repeat the example of a glider in the Game of Life, and ask “a glider cell-configuration in the Game of life has a weakly emergent property”, this proposition would be indeterminate, because it would be true on some precisifications, like Bedaus understanding of non-derivability and simulation and so on, and untrue on others, like Chalmers sparse view of emergence.

I think that the sparse view of emergence is perhaps consistent with Supervaluation, as you could argue that “Consciousness is a strongly emergent property of the brain” is true on all precisifications, but it might not be the case. If a predicate is indeterminate, it does not mean that it is theoretically useless, it just means that it will not refer to a specific extension. This means that it is not useful in working as a guide to refer us to all the cases of emergence.

### 3.3 Consequences

All this will lead us to having to make some decisions. If emergence is vague, like described in the previous sections, we can put forth some challenges to a hierarchical understanding of emergence. Vague predicates have unclear borders, and there is no reason to say that emergence happens at a certain point of complexity. We can easily

construct a sorites example with emergence, the same way we did with a heap, showing us that we do not have these clear separations between non-emerged and emerged. Physics, chemistry, biology and the other special sciences just fade into each other. This is not to say that my arguments commit us to an ontological denial of emergent levels, as I covered.

On the one hand, we could accept this vagueness as an obstacle in language, and go from there. This will lead us to solutions in the vain of some n-value logic or the stronger option, Supervaluation. Yet none of these, even if accepted fully, would help us solve the problems we are facing. They would only help us make due with vagueness in logic; how to work around it. They would not provide us with the theoretical backbone we need to make emergent levels suddenly exist at a certain point. To make room for this, the emergentist might have to accept some sort of epistemic view of vagueness.

Although vagueness as discussed here is a linguistic matter, it seems that it points us to the inability to arrive at sharp ontological boundaries, because we start out on vague linguistic grounds. I think that this linguistic analysis of a process of determining ontology has epistemic and ontological consequences in the sense that it limits what positions we can accept in good faith.

Recall that in the explicit form of my argument from the introduction, I sketched out three possible responses to my argument, given acceptance of the premises: Accepting some sort of scientific nihilism, accept epistemic vagueness and lastly denying emergent levels. This is slightly misleading, because accepting the second might also lead you to rejecting emergent hierarchies, or at least rejecting that we can have epistemic knowledge of it. The first option seems like an overreaction, because why would this linguistic problem with defining emergent phenomena lead us to scientific nihilism? It is because nihilism about bivalence is a response to vagueness in general, and not just this specific case. So, if we respond to the challenge of vagueness with refusal to believe that anything can be or not be, and that everything is vague and problematic, then there is not much we could say about science, and nihilism in the one would lead to nihilism in the other. Lastly there is the denial of epistemic levels, which is not as anti-emergent as it sounds, at least it does not have to be. The denial is about the ontological claims of emergent

levels, and probably a denial of the epistemic possibility of identifying even non-ontological emergent layers in a non-vague way, but still it is not any more than that. While denying emergent levels, we can still use emergence as a methodological tool in science, and as a model for reality, but that does not necessitate accepting it as ontological truth. Similarly, we are not required to think the atom looks like the Bohr model we have for it in physics, with the sphere like atom in the middle and the neutrons and positrons revolving around it like moons around a planet. It is simply a model made for easier understanding, not for naturalistic representation, and so we can keep the models of emergence, as explanatory tools, but must reject its claims of being a true representation of the phenomena it models.

Further the idea that there is some hierarchical nature of emergent special sciences is highly problematic. The sort of hierarchical emergent structure of nature, as suggested by Ellis should be seen as confusing the model with reality (Ellis, 2006). As I have shown, there are major obstacles when we want to define the kind of functional descriptions of the emergent levels. These obstacles, show us that there is vagueness and arbitrariness in including and excluding phenomena as emergent from something. This has the consequence that we cannot decisively say what kinds of phenomena count as the domain of one special science, and what kinds of phenomena count as another, purely on the grounds of our theoretical analysis of emergence; the extensions of the domain of the special sciences is bound to be vague, on an emergent theoretical model of division.

Looking at sciences as the plateaus of emergent phenomena is misleading. As the defining features of these layers are vague, there are any number of plateaus that are not accounted for, that fades into each other, it is confusing the discussion even to talk about them as plateaus, because no case can be made for them existing without accepting some account of epistemic vagueness. If not, then we need to give a logically sound explanation to what it is for these conditions to be undetermined. Either way, this divide in levels would not be determinable in the former case, and non-existent in the latter. Because of this, the doctrine that emergence shows us the *true special sciences*, or even eludes to where they lie in the scientific landscape, is untrue. Even graver than this, the process of going from theoretical understanding of some phenomena, into

generalizing the conditions for its occurrence, and use that to generate new examples of the same type of phenomena, is riddled with problems.

For these reasons, we should not be so eager to accept emergentist levels so wholeheartedly. There is more reason to believe that if emergence exists in any meaningful way, then its shape is not one with clear definable levels, but rather a gradual continuum.

## Summary

In this thesis, I have used the methodology and theoretical research provided by Timothy Williamson and Rosanna Keefe on vagueness, to shed light on some of the issues of emergence. I have presented various versions of emergent strength, their characteristics, the main objections to emergence and the scope of emergent application. Lastly, I have presented my argument for vagueness in emergence in general, but more specifically, as a problem for emergent levels.

I think the modern emergence project is somewhat undecided on central points. On the one hand, they want to move away from the objectionable ontologies that in the end proved to be the downfall of British Emergentism, and focus more on the epistemic, usable emergence for scientific methodologies. On the other, there seems to be a hope that ontological emergence could be recovered, and that emergence with a rehabilitated ontology would provide us with a framework of the natural world, that would lead to new understanding. I think modern emergentists do their view a disservice if the latter is attempted, as the ontologically non-committal path of emergence as utility seems a much more fruitful one, although its expectations needs to be managed. I have two main reasons for this claim: ontological emergence runs into problems with causal over/underdetermination, as shown in the causal exclusion argument, and emergent phenomena are bound to be vague, leaving the ontology problematically impotent. My argument, as I hope I have shown, does not say that the ontological commitments are false, but that even if they were developed with no internal problems, like the causal exclusion problem, it would still not provide us with any useful tools for identifying emergent properties, their relations, or separating between supposed emergent levels,

because emergence is too vague. I have also made the point that even the epistemically focused variant of emergence is problematic, if its expectations are not managed.

What is clearly apparent is that emergence is a modern and dynamic field of philosophical research and debate, with major contributions being very recent, by the standards of philosophy.

## References

- Anderson, P. (2008). More is Different: Broken Symmetry and the Nature of Hierarchical Structures of Science. In M. A. Bedau, & P. Humphreys, *Emergence: Contemporary Readings in Philosophy and Science* (pp. 221-229). Cambridge: The MIT Press.
- Assad, A., & Packard, N. H. (2008). Emergence. In M. A. Bedeau, & P. Humphreys, *Emergence: Contemporary Reading in Philosophy and Science* (pp. 231-234). Cambridge: The MIT Press.
- Bedau, M. A. (2008). Downward Causation and Autonomy in Weak Emergence. In M. A. Bedau, & P. Humphreys, *Emergence: Contemporary Readings in Philosophy and Science* (pp. 155-188). Cambridge: The MIT Press.
- Chalmers, D. J. (1996). *The Concious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.
- Chalmers, D. J. (2006). Strong and Weak Emergence. In P. Clayton, & P. Davies, *The Re-Emergence of Emergence: The Emergentist Hypothesis from Science to Religion* (pp. 244-254). Oxford: Oxford University Press.
- Clayton, P. (2006). Conceptual Foundations of Emergence Theory. In P. Clayton, & P. Davies, *The Re-Emergence of Emergence: The Emergentist Hypothesis from Science to Religion* (pp. 1-31). Oxford: Oxford University Press.
- Crane, T. (2001). *The Significance of Emergence*. London: Cambridge University press.
- Davies, P. C. (2006). The Physics of Downward Causation. In P. Clayton, & P. Davies, *The Re-Emergence of Emergence: The Emergentist Hypothesis from Science to Religion* (pp. 35-52). Oxford: Oxford University Press.
- Dennett, D. C. (2008). Real Patterns. In M. A. Bedau, & P. Humphreys, *Emergence: Contemporary Readings in Philosophy and Science* (pp. 189-205). Cambridge: The MIT Press.

- Ellis, G. F. (2006). On the Nature of Emergent Reality. In P. Clayton, & P. Davies, *The Re-Emergence of Emergence: The Emergentist Hypothesis from Science to Religion* (pp. 79-107). Oxford: Oxford University Press.
- Fodor, J. (2008). Special Science (Or: The Disunity of Science as a Working Hypothesis). In M. A. Bedau, & P. Humphreys, *Emergence: Contemporary Reading in Philosophy and Science* (pp. 395-409). Cambridge: The MIT Press.
- Jackson, F. (1982). Epiphenomenal Qualia. *Philosophical Quarterly* 32, 127-136.
- Keefe, R. (2000). *Theories of Vagueness*. Cambridge: Cambridge University Press.
- Kim, J. (2006). Being Realistic about Emergence. In P. Clayton, & P. Davies, *The Re-Emergence of Emergence: The Emergentist Hypothesis from Science to Religion* (pp. 188-202). Oxford: Oxford University Press.
- Kim, J. (2006, August 09). *Emergence: Core ideas and issues*. Retrieved from Springer: DOI: 10.1007/s11229-006-9025-0
- Kim, J. (2008). Making Sense of Emergence. In M. A. Bedau, & P. Humphreys, *Emergence: Contemporary Readings in Philosophy and Science* (pp. 127-153). Cambridge: The MIT Press.
- McLaughlin, B. P. (2008). The Rise and Fall of British Emergentism. In M. A. Bedau, & P. Humphreys, *Emergence: Contemporary Readings in Philosophy and Science* (pp. 19-59). Cambridge: The MIT Press.
- Rasmussen, S., Baas, N. A., Mayer, B., & Nillson, M. (2008). Ansatz for Dynamical Hierarchies. In M. A. Bedau, & P. Humphreys, *Emergence: Contemporary Readings in Philosophy and Science* (pp. 305-334). Cambridge: The MIT Press.
- Searle, J. (2008). Reductionism and the Irreducibility of Consciousness. In M. A. Bedau, & P. Humphreys, *Emergence: Contemporary Reading in Philosophy and Science* (pp. 69-97). Cambridge: The MIT Press.
- Williamson, T. (1994). *Vagueness*. London: Routledge.

- Wimsatt, W. C. (2008). Aggregativity: Reductive Heuristics for Finding Emergence. In M. A. Bedau, & P. Humphreys, *Bedau, Mark A.; Humphreys, Paul;* (pp. 99-110). Cambridge: The MIT Press.
- Wong, H. Y. (2010). The Secret Lives of Emergents. In A. Corradini, & T. O'Connor, *Emergence in Science and Philosophy* (pp. 7-24). New York: Routledge.