

Patternhood and Generality: A Peircean Approach to Emergence

Jimmy Aames
Osaka University Graduate School of Human Sciences
asonosakan@gmail.com

Abstract

A distinction often drawn in the literature on emergence is that between *epistemological emergence* and *ontological emergence*. I argue that the position that there are instances of ontological emergence in the world (*radical emergentism*) and the position that all instances of emergence are epistemological (*ontological reductionism*) are both problematic. My aim in this paper is to outline a third form of emergence, inspired by the philosophical ideas of Charles S. Peirce and Daniel Dennett, which will allow us to avoid both radical emergentism and ontological reductionism. Since what emerges in this kind of emergence is a *real pattern*, I call it *real pattern emergence* (*RP emergence*).

A defining characteristic of a pattern is that it is more *general* than the elements or processes instantiating it. That is, the same pattern can be instantiated by different individual elements or processes. This generality of patterns makes them perfect candidates for applying Peirce's rich theory of generals, including his pragmatism and Scholastic realism. I will argue that a real pattern is what Peirce calls a *real general* (or *real Third*), and like a real general, it is autonomous from the elements or processes instantiating it, because it supports predictions about not only what *will* happen in a given situation, but also what *would* happen in an indefinite variety of possible micro situations that are not actualized. This autonomy of real patterns will be the key feature that sets RP emergence apart from mere epistemological emergence, without, on the other hand, collapsing it into ontological emergence. The difference between RP emergence and Mark Bedau's related notion of *weak emergence* will also be clarified.

Keywords: Charles S. Peirce, Daniel Dennett, real patterns, pragmatism, Scholastic realism, weak emergence

Get rid, thoughtful Reader, of the Ockhamistic prejudice of political partizanship that in thought, in being, and in development the indefinite is due to a degeneration from a primal state of perfect definiteness.

Charles S. Peirce, “Some Amazing Mazes, Fourth Curiosity”

Introduction

A distinction often drawn in the literature on emergence is that between *epistemological emergence* and *ontological emergence*.¹ Epistemological emergence is a kind of emergence that is only in the eyes of the beholder. An epistemologically emergent phenomenon is one that can in principle be reduced to—predicted or derived from—its underlying elements or processes, but is in practice unpredictable or underivable from these elements or processes due to its complexity and the epistemic limitations on the part of the observer (limitations in computational resources, knowledge about initial conditions, etc.). An ontologically emergent phenomenon, on the other hand, is one that is unpredictable or underivable from its underlying elements or processes even in principle. Epistemological emergence seems to be the conception of emergence most widely accepted among philosophers and scientists, whereas it remains controversial whether there are any instances of ontological emergence in the actual world.

Suppose there is no ontological emergence, and that every instance of emergence in the world is epistemological. Combined with the assumption that every emergent feature in the world ultimately emerges from the entities and laws of fundamental physics, this implies that every emergent feature in the world can in principle be reduced to the entities and laws of fundamental physics, even if reference to emergent phenomena is indispensable in practice. Let us call this view *ontological reductionism*. Although widely endorsed by philosophers and scientists, I find this view problematic.² Let me explain why by citing a thought experiment put forth by Daniel Dennett.

In his essay “True Believers,” Dennett asks us to imagine that beings of vastly superior intelligence, say Martians, descend upon us (Dennett 1987:25). Suppose, he says, that they are “Laplacian super-physicists, capable of comprehending the activity on Wall Street, for instance, at the microphysical level. Where we see brokers and buildings and sell orders and bids, they see vast congeries of subatomic particles milling about” (Dennett 1987:25). According to the ontological reductionist, these Martians would know everything there is to know about the world. However, Dennett

¹ Of course, different authors prefer different terminology. What I am here calling “epistemological emergence” has also been called “weak emergence” (Chalmers 2006; not to be confused with Mark Bedau’s version of weak emergence, which I will discuss in §4), “conservative emergence” (Seager 2012), and “conceptual emergence” (Humphreys 2016), while what I am here calling “ontological emergence” has also been called “strong emergence” (Chalmers 2006), “radical emergence” (Seager 2012), and “brute emergence” (Strawson 2006). It should be noted that Chalmers’ notion of weak emergence is broader than what I call epistemological emergence, since he defines weak emergence in such a way that cases of strong emergence are also cases of weak emergence, whereas I prefer to define epistemological and ontological emergence to be mutually exclusive.

² The view is held, for example, by Anderson (1972), Weinberg (1987), and Seager (2012).

points out that even if the Martians were able to comprehend and accurately predict everything that happens on Wall Street using their Laplacean methods, they would be missing something perfectly real if they did not also see us as intentional beings, that is, if they did not also see the *patterns* in human behavior that we describe in intentional terms, such as *believing that p* or *desiring q*. As he puts it:

Take a particular instance in which the Martians observe a stockbroker deciding to place an order for 500 shares of General Motors. They predict the exact motions of his fingers as he dials the phone and the exact vibrations of his vocal cords as he intones his order. But if the Martians do not see that indefinitely many *different* patterns of finger motions and vocal cord vibrations—even the motions of indefinitely many different individuals—could have been substituted for the actual particulars without perturbing the subsequent operation of the market, then they have failed to see a real pattern in the world they are observing. (Dennett 1987:26)

Although Dennett's argument is couched in terms of intentional patterns, the same argument can also be made with respect to any pattern outside the purview of fundamental physics—the Martians would be just as blind to fingers and vocal cords as they are to the intentional behavior of humans. This is a powerful thought experiment that shows us the inadequacy of ontological reductionism. It is inadequate because it fails to recognize that there is more to reality than what can be described at the level of fundamental physics. The patterns that we observe at everyday scales, as well as those studied by the various special sciences, are not mere epistemic crutches that can be dispensed with by hypothetical Laplacean super-physicists; they are just as real and fundamental as the patterns studied by fundamental physics.

Now if we are to deny ontological reductionism, it seems we ought to embrace the existence of ontological emergence.³ However, the view that there are instances of ontological emergence in the world—which I will hereinafter call *radical emergentism*—is also problematic. The problem is that an ontologically emergent phenomenon is by definition utterly inexplicable, in the sense that there is absolutely nothing about the underlying elements or processes by virtue of which it should emerge, and should have the features that it has. It is, in a word, sheer magic: it simply pops into existence without any why or wherefore. But to use a Peircean turn of phrase, to posit something utterly inexplicable is to set up a roadblock to inquiry. The synechistic philosophy demands that we do not introduce such brute discontinuities into the fabric of being.

Galen Strawson has gone further and argued that the notion of ontological emergence (which he calls “brute emergence”) is incoherent:

If it is really true that Y is emergent from X then it must be the case that Y is in some sense wholly dependent on X and X alone, so that all features of Y trace intelligibly back

³ One could also attempt to deny the assumption, mentioned earlier, that every emergent feature in the world ultimately emerges from the entities and laws of fundamental physics. However, I will simply accept this as a plausible assumption.

to X (where ‘intelligible’ is a metaphysical rather than an epistemic notion). *Emergence can’t be brute*. It is built into the heart of the notion of emergence that emergence cannot be brute in the sense of there being absolutely no reason in the nature of things why the emerging thing is as it is (so that it is unintelligible even to God). For any feature Y of anything that is correctly considered to be emergent from X, there must be something about X and X alone in virtue of which Y emerges, and which is sufficient for Y. (Strawson 2006:18)

If both ontological reductionism and radical emergentism are untenable, then it seems we are at an impasse. I suggest that the problem lies in the assumption that epistemological emergence and ontological emergence exhaust all conceivable forms of emergence. If there is a further, third form of emergence, then denying the existence of ontological emergence will not entail that every instance of emergence is epistemological. My task in this paper will be to outline such a third form of emergence, inspired by the philosophical ideas of Charles S. Peirce as well as those of Dennett, that will enable us to steer a path between the Scylla of ontological reductionism and the Charybdis of radical emergentism. Since what emerges in this form of emergence is a *real pattern* (where the precise sense of “real” will be clarified later), I will simply call this form of emergence *real pattern emergence*, or *RP emergence* for short.⁴

A defining characteristic of a pattern is that it is more *general* than the elements or processes that instantiate it. By this I mean that the same pattern can be instantiated by different individual elements or processes. A melody played on the piano, for example, will retain its identity even if each of its individual notes is raised or lowered an octave. The generality of patterns makes them perfect candidates for applying Peirce’s rich theory of generals (more commonly known as *universals*), including his pragmatism and Scholastic realism. What I want to explore in this paper are the implications that such an application will have for our understanding of emergence.

This paper is organized as follows. I begin in the first section with some preliminary observations on the notion of “levels,” often presupposed in discussions of reduction and emergence, and offer my reasons for avoiding talk of levels. In the second section, drawing on Dennett’s theory of real patterns (Dennett 1991), I will introduce the distinction between *real patterns* and *non-real patterns*, and I will further argue that patterns are generals. This will set the stage for our study of Peirce’s theory of generality—specifically, his modal analysis of generality—in the third section. The central point I want to make here is that a real pattern is what Peirce calls a *real general* (or *real Third*), and like a real general, it is autonomous from the elements or processes instantiating it, because it supports predictions about not only what *will* happen in a given situation, but also what *would* happen in an indefinite variety of possible micro situations that are not actualized. This autonomy of real patterns will be the key feature that sets RP emergence apart from mere epistemological emergence, without, on the other hand, collapsing it into ontological emergence. Finally, in the fourth section I will try to throw what I call RP emergence into sharper relief by comparing it with Mark Bedau’s related notion of *weak emergence* (Bedau 1997, 2002).

⁴ The notion of “real pattern” is taken from Dennett (1991).

1 Levels

The notion that nature is organized in a hierarchical structure of “levels” is deeply embedded in discussions of emergence. One often speaks, for example, of “higher-level” entities possessing properties lacked by “lower-level” entities. However, it is by no means evident what levels are—are they objective features of nature, or do they somehow reflect the way we choose to describe nature?—or whether there are such things as levels at all. Indeed, James Ladyman and Don Ross have denied the existence of levels, arguing that talk of “levels” is a metaphor to which contemporary science gives no interesting content (Ladyman & Ross 2007:53–57). Despite the elusiveness of the levels concept, most discussions of emergence in the philosophical literature simply assume that there are such things as levels, without addressing the issue of what they are and whether they actually exist.⁵ It therefore behooves us to undertake a preliminary examination of the notion of levels before delving into our discussion of emergence.

While I share Ladyman and Ross’s skepticism regarding the notion of levels, it seems to me that their denial of the existence of levels needs to be qualified in at least two respects. In the first place, the relative strength and range of the four fundamental forces (the strong force, electromagnetic force, weak force, and gravity), together with the kinds of matter upon which they act, give rise to a separation of three “natural” levels or regimes, which I shall call the *subatomic regime*, *electromagnetic regime*, and *gravitational regime*.⁶ Both the electromagnetic and gravitational forces have an infinite range and act on all size scales, but the electromagnetic force only acts on electrically charged matter, while gravity acts on every object in spacetime. Furthermore, in large objects the positive and negative electric charges tend to cancel each other out, making the object as a whole electrically neutral. This is why the influence of gravity tends to dominate at large scales, such as the scale of stars and galaxies. On the other hand, the electromagnetic force is much stronger than gravity, and so it tends to dominate at smaller scales, such as the scale of atoms, molecules, and what we regard as medium-sized objects. This is how the separation between the electromagnetic regime, dominated by the electromagnetic force, and the gravitational regime, dominated by gravity, arises. The boundary between these two regimes is by no means sharp, as is evinced by the fact that we experience the effects of both the electromagnetic force and gravity at everyday scales. The separation of the subatomic and electromagnetic regimes likewise arises from the fact that the influence of the weak and strong forces is confined to very small distances (of the order of 10^{-15} m, roughly the size of an atomic nucleus).

Thus, contrary to Ladyman and Ross’s claim that levels do not exist, there are at least three levels or regimes in nature corresponding to three different size scales. However, levels identified solely in terms of size do not provide a sufficient basis for discussions of emergence. Some examples of emergence often discussed in the literature are the emergence of life from chemical processes, the

⁵ Some rare exceptions, in addition to Ladyman & Ross (2007), are Wimsatt (1976:237–263), Wimsatt (1994), and Humphreys (2016:120–26).

⁶ The discussion in this paragraph owes much to Reiji Sugano’s study of the hierarchical structure of nature (Sugano 2013, Chapter 3). I also want to thank Taksu Cheon for his insightful suggestions on the topic in personal correspondence.

emergence of mental phenomena from neural processes, and the emergence of social behavior patterns (such as crowd behavior) from interactions between individual humans. The problem is that there are no size scales that uniquely characterize organisms, mental phenomena, or social behavior patterns, and hence if these are to be identified as levels, this identification cannot be made solely in terms of size. As pointed out by William Wimsatt (1994:236), a bacterium could have the same size as a black hole, but we would hardly consider the two as belonging to the same level, as they would behave in radically different ways in similar circumstances. As for mental phenomena and social behavior patterns, it is not even clear whether these could have sizes at all.

Evidently, size cannot be the sole factor in terms of which we identify levels. But there are other factors that enable us to do so, and this brings us to the second respect in which Ladyman and Ross's rejection of levels needs to be qualified. Consider, for example, the level of individual organisms. How do we identify this as a distinct level? I suggest that it is by focusing on a cluster of recurrent *patterns* or regularities that we observe in nature, such as reproduction, metabolism, and homeostasis. The level of individual organisms can be regarded as "higher" than the level of chemical processes because these patterns are lacking at the chemical level. In general, we can say that identifying a level involves picking out a set of patterns from the phenomena we observe in nature, and that a given level *A* is "higher" than another level *B* if the patterns in terms of which we identify *A* are lacking in *B*. This implies that what we identify as levels (apart from the three regimes mentioned above) depends to some extent on what patterns we detect and choose to focus on. Ladyman and Ross's rejection of levels has thus been qualified in two respects: in the first place, there are at least three regimes in nature corresponding to three different size scales; and in the second place, while it may be true that levels other than these three regimes do not exist apart from pattern-detecting agents, we can nonetheless speak intelligibly of these levels as long as we keep in mind their observer-dependent character.

Despite all of this, in this paper I will avoid talk of levels and speak instead of emergent patterns and their underlying elements and processes. The main reason for this is that the language of levels tends to give the impression that there is a fixed hierarchy of levels that somehow exists independently of the act of detecting and picking out the patterns that characterize them. This, however, is not the case: as I argued above, it is the act of detecting and picking out a certain set of patterns that gives rise to a distinct level; the level does not exist independently of this act. This is closely related to a point made by Paul Humphreys. A basic distinction between types of emergence (orthogonal to the epistemological/ontological distinction) is that between *synchronic emergence* and *diachronic emergence* (Humphreys 2008, 2016). In synchronic emergence, the emergent phenomenon is considered to exist simultaneously with its substrate, as when mental phenomena are considered to emerge from neural processes. In diachronic emergence, on the other hand, the emergent phenomenon is considered to develop over time from prior states of a system, as when complex patterns are generated in cellular automata. Humphreys argues that contemporary discussions of emergence in the philosophical literature have been overly focused on synchronic emergence (perhaps due to the circumstance that most philosophical discussions of emergence have taken place in the context of the philosophy of mind) and have neglected diachronic forms of emergence (Humphreys 2016). The notion of levels is one manifestation of this overemphasis of synchronic emergence, for

as Humphreys points out, “the levels imagery is shot through with synchronic concepts” (Humphreys 2016:121). The levels framework is ill-suited to dealing with diachronic emergence. Again, the problem is that the language of levels tends to suggest that the “higher” level exists independently of the act of detecting and picking out the patterns that characterize it, whereas in diachronic emergence, where the emergent patterns come into being over time, it hardly makes any sense to say that there is a “higher level” when the patterns that allow us to identify it have not yet emerged. For these reasons I believe that the levels concept tends to distort rather than facilitate our understanding of emergence, and so I will avoid speaking of levels in this paper, except when presenting the views of other authors who use the concept.

With these preliminaries in place, let us turn to our main topic, RP emergence.

2 Real Patterns

My task in this section will be twofold: first, to specify what it means for a pattern to be *real*, and second, to show that patterns are *general* in the sense of being instantiable by different individual elements or processes.

2.1 Algorithmic Compressibility

To begin with, it will be useful to have a definition of the key term *pattern*. Here I will be relying on the theory of real patterns proposed by Dennett in his paper “Real Patterns” (Dennett 1991). In the most general terms, a pattern is a regularity in some data, where *data* is construed in the broadest possible sense as something that is observed or may be observed. Consider, for example, an endless random string of 0’s and 1’s. There is no regularity in this data. On the other hand, consider an endless string of alternating 0’s and 1’s: 010101010 ... etc. What we should notice is that this data can be compressed into a program that commands: “generate an endless string of alternating 0’s and 1’s.” There is no way of compressing the random string of 0’s and 1’s—the only way this data can be transmitted to another person is to send the *bit map*, which identifies each digit *seriatim* (the first place value is 0, the second place value is 0, the third place value is 1, etc.). In more general terms, a bit map is a zero-compression encoding, where each bit of information in the initial data is mapped one-to-one to a distinct bit in the encoding.

Gregory Chaitin, one of the founders of algorithmic information theory, gives the following definition of *randomness*: “A series of numbers is random if the smallest algorithm capable of specifying it to a computer has about the same number of bits of information as the series itself” (Chaitin 1975:48). Reversing this idea, Dennett proposes the following criterion for the presence of a pattern: “A pattern exists in some data—is real—if there is a description of the data that is more efficient than the bit map, whether or not anyone can concoct it” (Dennett 1991:34). That is, there is a pattern in some data if there is an algorithm that reproduces the data using a smaller number of bits

than the data itself (when there is such an algorithm, we say that the data is *algorithmically compressible*).⁷

An interesting aspect of pattern recognition is that not all observers are able to discern the same pattern in the same data, and even the same observer may discern different patterns in the same data on different occasions. The famous duck-rabbit illusion is a prime example of the latter. As an example of the former, suppose that an image file—say a jpg image of a human face—is translated into binary notation, pixel by pixel. The pattern is still there, but it would be impossible for the human eye to discern it visually. Other creatures with different sense organs may readily perceive patterns that are imperceptible to us (Dennett 1991:34). Hence Dennett’s proviso that the presence of a pattern should not depend on whether or not anyone is actually able to concoct a compression algorithm: there is a pattern in some data if the data is *in principle* compressible by a potential observer.

2.2 Predictive Power

Dennett’s algorithmic compressibility criterion specifies a necessary and sufficient condition for the *presence* of a pattern, but it does not by itself guarantee that the pattern is *real*. It is true that he formulates the criterion using the term “real”: “A pattern exists in some data—is real—if there is a description of the data ...” (Dennett 1991:34). However, I believe this is simply due to his ambiguous use of the word “real” in this paper, for he is clear throughout that there are non-real as well as real patterns, and he goes on to specify a further criterion for distinguishing between the two.

To get an idea of what a real pattern is, let us first consider what a non-real pattern is. A non-real pattern is one that is due to pure chance. Suppose, for example, that we are to throw a fair die one thousand times and record the outcome of each trial. Let us further suppose that after having done ten throws, we find that in all of the first ten throws the die turns up six. This, of course, is a pure accident, but if we were to show the results of the first ten trials to someone without telling her that the numbers were generated by throwing a die, and we were to ask her to predict the next number, she would most likely predict that it will also be a six. She may happen to be right, but in that case we ask her to predict the next number, and the next, and so on. Then her prediction is bound to fail sooner or later. A pattern is undoubtedly present in the results of the first ten trials, and indeed, the data for the first ten trials can be compressed into a program that commands: “generate ten sixes.” Nonetheless, further accumulation of data will eventually reveal the accidental nature of this pattern.

Conversely, we can say that a pattern is real if by discerning it we are able to make better-than-chance predictions about future events. A real pattern is one that is projectible into unobserved instances. Or as Dennett puts it, a pattern is real if you can get rich by betting on it (Dennett 1991:36). Dennett himself is not altogether clear on the relation between this predictive power criterion and

⁷ As we will see below, in the same paper Dennett gives another criterion for the reality of patterns, according to which a pattern is real if by discerning it one can make successful predictions about future events. The relation between this predictive power criterion and the algorithmic compressibility criterion cited here, however, is unclear. It seems to me that this unclarity derives from Dennett’s ambiguous use of the term “real” in the quoted passage. The unclarity can be resolved if we take Dennett’s algorithmic compressibility criterion as a criterion simply for the *presence* of a pattern, regardless of whether or not it is real (in the sense of supporting predictions about future events). I will touch upon this point again below.

the algorithmic compressibility criterion. The two criteria are clearly not equivalent, as can be seen in the above example of throwing a die. There will be no obscurity, however, if we simply take the predictive power criterion as a necessary and sufficient condition for the *reality* of a pattern, and the algorithmic compressibility criterion as a necessary and sufficient condition for the *presence* of a pattern, regardless of whether it is real or not. As we shall see below in §3, Dennett's idea that the reality of patterns consists in their predictive power is strikingly similar to one of Peirce's arguments for Scholastic realism.

The notion of predictive power brings us back to our earlier consideration, that not all observers are able to discern the same pattern in the same data, and that even the same observer may discern different patterns in the same data on different occasions. This means that patterns are in some sense *observer-dependent*. Dennett explicates this notion in terms of prediction: patterns are observer-dependent in that they can be discerned only from the point of view of an observer that adopts a certain predictive strategy, or *stance*, to use Dennett's terminology. For instance, Dennett calls the predictive strategy from which intentional states—beliefs, desires, and the like—can be discerned the *intentional stance* (Dennett 1987:17). Likewise, there can be predictive strategies for discerning any kind of pattern whatsoever: the Newtonian mechanics stance, the cellular biology stance, the microeconomic stance, etc. The idea is that patterns are not simply “out there,” naked in the world; on the contrary, the recognition of a pattern must always involve an element of active participation on the part of the observer, namely the adoption of a certain predictive strategy. This should not be taken to mean that the act of adopting a predictive strategy is always a conscious, deliberate act: the decision of which predictive strategy to adopt is dictated to a large degree by the structure of our sense organs, our genetic makeup, and the evolutionary history of our culture (Dennett 1991:36).

Patterns thus have an observer-dependent being; but at the same time, they are in another sense *observer-independent*. They are observer-independent in that the facts about the success or failure of our predictive strategies do not depend on what we may think or will them to be; they are completely out of our control. It is this uncontrollability of the outcome of our predictions that imparts to some patterns—namely, those whose discernment leads to successful predictions—a real being.

2.3 Multiple Instantiability

An important characteristic of patterns is what I shall call their *multiple instantiability*, that is, the capacity for the same pattern to be instantiated by different individual elements or processes. In the introduction I mentioned the example of a melody played on the piano: the melody will retain its identity even if each of its individual notes is raised or lowered an octave. Our body is also multiply instantiable in this sense, since it remains the same even though the cells that compose it are constantly being replaced by new ones. I contend that any pattern one could think of has this property: the same flocking pattern can be exhibited by different individual birds; a wave propagating through a gas or liquid persists even as the molecules that constitute the wave constantly change; the solubility of salt in water manifests itself every time the relevant conditions are fulfilled; etc.

There are two points of clarification that I want to make here. First, a difference between the flocking pattern and the wave is that in any given instance of flocking behavior, the birds constituting the overall flocking pattern remain fixed, whereas in the case of a wave, the molecules constituting the wave pattern change over time. Humphreys (2008:437–38) refers to patterns of the first kind as *micro-stable patterns* and those of the second kind as *micro-dynamic patterns* (and further distinguishes three sub-types of the latter). Our body is also an example of a micro-dynamic pattern. This difference, however, does not make the flocking pattern any less multiply instantiable than the wave or body is.

Second, it may be argued that by saying that our body or a wave is multiply instantiable, I am confusing the instantiation relation with a part-whole relation. The cells that compose our body, and the molecules that constitute a wave, are in a part-whole relation with the overall pattern, and the part-whole relation is not the same as the instantiation relation. My reply is that the constituent elements are in *both* a part-whole relation and an instantiation relation with the overall pattern. It seems undeniable that my body at a time t_1 and at a time t_2 a year later, say, are the same pattern instantiated at different points in time by different constituent elements; and the elements' being in a part-whole relation with the overall pattern does not prevent them from also being in a relation of instantiation with that same pattern.

We have thus established that patterns are multiply instantiable. But this is simply another way of saying that patterns are general, for the traditional definition of a general (or universal) ever since Aristotle is: a general is that which can be predicated of many things; or in other words, a general is that which is multiply instantiable.⁸ That patterns are general is not surprising if we consider the fact that to discern a pattern is to discern a certain *form*, and a form is general. Samuel Alexander, one of the fathers of British emergentism, explicitly draws a connection between patterns and generals, arguing that the “quality” that emerges in any instance of emergence is at once a pattern and a universal: “To adopt the ancient distinction of form and matter, the kind of existent from which a new quality emerges is the ‘matter’ which assumes a certain complexity of configuration and to this pattern or universal corresponds the new emergent quality” (Alexander 1920, 2:47). Given that patterns are generals, Dennett’s theory of real patterns can be said to be a revival, in modern garbs, of Scholastic realism—the doctrine that there are real generals—although I doubt that Dennett himself views his theory in this way.⁹ The generality of patterns makes them perfect candidates for applying Peirce’s rich theory of generals, in particular his modal analysis of generality. To this we shall now turn.

⁸ Aristotle, *De Interpretatione*, VII, 17a38.

⁹ There is, however, a residue of nominalism in Dennett’s theory, namely his retention of Hans Reichenbach’s distinction between *illata* (concrete physical objects) and *abstracta* (abstract objects), and his characterization of the latter as “lossy compression[s]” (Dennett 2000:360). Ross (2000), rightly in my opinion, argues that Dennett should abandon the *illata/abstracta* distinction.

3 Peirce's Modal Analysis of Generality

We have seen that according to Dennett's criterion, a real pattern is one whose discernment allows us to make successful predictions about future events. What I want to show in this section is that an emergent pattern's capacity to support predictions about what will happen in the future is not enough for it to be autonomous from its underlying elements and processes. In order for a pattern to have this kind of autonomy, it must support predictions about not only what *will* happen in the future, but also what *would* happen in conceivable situations that are not actualized. Or as Peirce would say, it must be not merely a real *will-be*, but a real *would-be*. Peirce's modal analysis of generality, according to which generals have the modality of *would-be's*, will, I hope to show, give us a clear understanding of the autonomy characteristic of RP emergence.

3.1 The Pragmatic Maxim

Peirce's modal analysis of generality is intimately connected with his pragmatism, so we shall begin by considering his pragmatism. Peirce himself was not always clear on the difference between *will-be's* and *would-be's*. In his early years he tended to vacillate between using the indicative mood (*will be*) and subjunctive mood (*would be*) in stating pragmatic clarifications.¹⁰ It was only in his later years, in the 1900s, that he became explicit about the difference between the two and began stressing the importance of using the subjunctive mood in stating pragmatic clarifications. A good way to understand the difference between *will-be's* and *would-be's*, therefore, is to trace the development of Peirce's own ideas on the subject.

Peirce's pragmatic maxim appeared in public form for the first time in his 1878 paper "How to Make Our Ideas Clear" (HTM), published as part of the *Illustrations of the Logic of Science* series in the *Popular Science Monthly*. It was formulated as a logical principle for clarifying the meaning of ideas, for attaining the "third grade of clearness" of apprehension, the first two being the traditional criteria of *clearness* and *distinctness* as formulated by Descartes and developed by Leibniz (EP 1:124–27, W 3:257–61, 1878).¹¹ The famous statement of the maxim runs as follows: "Consider what effects, which might conceivably have practical bearings, we conceive the object of our conception to have. Then, our conception of these effects is the whole of our conception of the object" (EP 1:132, W 3:266, 1878).

To use the well-known example that Peirce himself gives, the meaning of the concept of *hard* is that an object to which *hard* can be applied as a predicate, such as a diamond, would not be scratched even if we apply pressure to it with, say, a knife-edge. That is, the meaning of the concept of *hard* can be analyzed into general laws expressible in conditional propositions of the form:

If you were to apply pressure to *X* (where *X* is an object to which *hard* can be veritably applied as a predicate), then *X* would resist the pressure.

¹⁰ See Lane (2018:51–58) for an account of this vacillation.

¹¹ See the list at the end of this paper for an explanation of the abbreviations used in referring to Peirce's writings.

Notice that I have formulated this conditional in the subjunctive mood (*would*) rather than the indicative mood (*will*); this is in accordance with Peirce's later, more considered view. Now let us consider the general case. The pragmatic maxim states that the meaning of any concept can be analyzed into general laws expressible in conditional propositions of the form:

(*) If you were to do *m* to *X* (where *X* is an object to which the concept in question can be veritably applied as a predicate), then you would have an experience of type *n*.

We perform a certain kind of operation on nature (for example, we apply pressure to a diamond), and nature gives us a certain kind of response (the diamond resists the pressure). The *general law* that operations of a certain kind are always (or with a certain degree of probability) followed by responses of a certain kind is what constitutes the meaning of the concept in question.

What, then, is the difference between formulating pragmatic clarifications in the indicative mood and in the subjunctive mood? Indicative-mood clarifications analyze the meaning of a concept into laws that dictate what *will* happen in the future; whereas subjunctive-mood clarifications analyze the meaning of a concept into laws that have reference not only to what will happen in the future, but also to what *would* happen in possible situations that are not actualized. This difference is well-illustrated by a thought experiment that Peirce poses in HTM. There, he supposes that a diamond is burned up before its hardness could be tested, and asks whether it would be false to say that the diamond was soft. His answer, at the time of writing HTM, is "no":

[T]here would be no *falsity* in such modes of speech. They would involve a modification of our present usage of speech with regard to the words hard and soft, but not of their meanings. For they represent no fact to be different from what it is; only they involve arrangements of facts which would be exceedingly maladroit. This leads us to remark that the question of what would occur under circumstances which do not actually arise is not a question of fact, but only of the most perspicuous arrangement of them (EP 1:132, W 3:267, 1878).

In his 1905 paper "Issues of Pragmaticism," however, Peirce admits that this conclusion was a mistake, and argues that a diamond's hardness is a real property, regardless of whether it is actually put to the test or not (EP 2:356–57, 1905). The crucial difference between Peirce in the 1870s and Peirce in the 1900s lies in his conception of what kind of modality general laws possess. For the earlier Peirce general laws are *will-be's*, that is, they have reference only to events that we know will occur in the future, while for the later Peirce general laws are *would-be's*, that is, they have reference also to counterfactual situations.

While the pragmatic maxim was originally formulated as a rule for clarifying the meaning of *concepts*, it can be extended into a principle for identifying the intellectual purport of any kind of general, anything capable of functioning as a predicate, including patterns. Together with Peirce's mature conception of general laws as *would-be's*, this implies that the intellectual purport of any

general, and hence any pattern, is a *would-be*.¹² Thus, to judge that something has a certain general form, or that something exhibits a certain pattern, is to judge that it is governed by general laws expressible in subjunctive conditionals of the form (*). From this it further follows that to discern a pattern in some system is to make predictions about how the system would behave under certain conceivable conditions. As we saw in the previous section, this is precisely what Dennett's theory asserts, except that here the predictions have reference not only to future events, but also to counterfactual situations.

3.2 Scholastic Realism

So far we have seen that according to the pragmatic maxim, to predicate a general of some object is to judge that the object is under the governance of certain *would-be*'s. Not all *would-be*'s, however, are real. Recall the example of throwing a die that we considered in §2.2. To say that the die is *fair* is to ascribe to the die a real *would-be*, namely, that if it were thrown many times, each face would turn up with a relative frequency close to 1/6, the deviation from 1/6 becoming smaller and smaller as more throws are made. But if we were to say, after having observed only the first ten throws, that the die is *loaded* in such a way that it only turns up sixes, then the *would-be* we ascribe to the die—that it would turn up a six every time it is thrown—would not be a real *would-be*, but a hasty generalization based on insufficient data. This difference can also be expressed by saying that the property of *being fair* is a real property of the die, whereas the property of *being loaded* is not. The pragmatic maxim analyzes the intellectual purport of any general predicated of a given object into a set of *would-be*'s, but it does not by itself determine whether the *would-be*'s are real or not; nor does it determine whether the general is a real property of the object or not. We therefore need some other criterion by which we can determine the reality of a general predicated of a given object. This criterion is provided by one of Peirce's arguments for Scholastic realism, the doctrine that there are real generals. As I mentioned in §2.2, this argument is intimately related to Dennett's formulation of the reality of patterns in terms of predictive power.

But before we go into the argument, a note on Peirce's use of the term "real" is in order. He defines *real* as follows: the real is "that whose characters are independent of what anybody may think them to be" (EP 1:137, W 3:271, 1878). That is, something is said to be real if its being such as it is cannot be (or could not have been) altered by the mere act of thinking it to be otherwise. Thus, Prince Hamlet is not real, since his characters (that he is the Prince of Denmark, that he sees his father's ghost, etc.) could have been different if Shakespeare had conceived of him differently; but the very fact that Shakespeare wrote *Hamlet* is real, since nothing about that fact can be (or could have been) altered by the mere act of thinking it to be otherwise. While this definition of "real" is clearly

¹² It is not quite clear whether Peirce himself intended the pragmatic maxim to be a rule about the intellectual purport of any general. In a later formulation, the maxim is framed as a rule for identifying the "intellectual purport of any symbol" (EP 2:346, 1905), but not all generals are symbols, since in Peirce's famous ten-fold classification of signs, there are legisigns (general signs) that are not symbols (EP 2:289–99, 1903). On the other hand, legisigns are general laws that are signs, and hence have the modality of *would-be*'s. Therefore, assuming that all generals are legisigns, it seems that for (the later) Peirce all generals have the modality of *would-be*'s.

different from Dennett's formulation of the reality of patterns in terms of predictive power, we shall see that a pattern's being projectible into the future is a good indication of its being real in Peirce's sense.

The argument I want to take up appears in the "Seven Systems of Metaphysics," the fourth of Peirce's 1903 Harvard Lectures on pragmatism. There, he takes a stone in his hand and announces to the audience that he will perform an experiment: he will let go of the stone and see whether it will fall to the floor (EP 2:181, 1903). The experiment, of course, is meaningless, since everybody knows what will happen. But the deeper meaning of the experiment lies in the very fact that it is meaningless. How is it that we do not have to actually perform the experiment in order to know its result? The answer must be that the stone is governed by a real law operative in nature. If the law were only a mental formula and not real, there would be no way of explaining why future events will conform to it (and we know that they will), unless we were to suppose that the mind had some kind of miraculous power of prognosis. If, on the other hand, we suppose that the stone is governed by a real law dictating what would happen in certain kinds of situations, then witnessing the actual instantiations of the law will be no wonder. Now since laws are generals, it follows that there are real generals. Thus runs Peirce's argument for realism from our experience of anticipation.

Now according to the (extended) pragmatic maxim, the intellectual purport of any general consists in laws expressible in subjunctive conditionals of the form (*). In the case of a stone, judging that something is a stone involves identifying the laws that it conforms to, one of which can be expressed as: "if you were to let go of the stone, it would fall," and the conjunction of all such laws constitutes the entire intellectual purport of the general form or pattern of *stoneness*. Therefore, insofar as we know those laws to be real, then so is the general form or pattern whose intellectual purport these laws constitute. We are thus able to see why a pattern's being projectible into the future is a good indication (but not conclusive proof—there is no such thing as a conclusive proof in matters empirical) of its being real in Peirce's sense. It should be noted that although the predictions associated with a given pattern have reference to counterfactual situations, the reality of the pattern can be assessed only by testing whether the predictions it affords are actually fulfilled. It is only after we have established the reality of the pattern that we are justified in further generalizing the predictions to counterfactual situations.

3.3 Real Pattern Emergence

We are now in a position to define RP emergence:

A pattern P of a system S is said to be an *emergent real pattern* if it satisfies the following three conditions:

- (1) P is an emergent pattern of S , that is, P is a pattern that arises in S but cannot be manifested by the constituent elements of S ;
- (2) P can in principle be predicted or derived from knowledge about the constituent elements of S ;

(3) *P* is real in the sense that it supports predictions about not only what *will* happen given a certain state of *S*, but also what *would* happen in an indefinite variety of possible micro situations of *S*.

Any instance in which a real pattern emerges in the above sense is an instance of *real pattern emergence* (*RP emergence*).

Condition (1) is simply a statement of the traditional notion of an emergent property as a “higher-level” property that cannot be possessed by “lower-level” entities, except that here it is framed in terms of patterns rather than properties. By condition (2), RP emergence is similar to epistemological emergence and differs from ontological emergence, in that the emergent pattern can in principle be derived from or explained in terms of its underlying elements or processes. This stipulation is necessary in order to prevent RP emergence from having the occult character of ontological emergence. On the other hand, by condition (3), RP emergence differs from epistemological emergence in that an emergent real pattern is in a strong sense *autonomous* from its underlying elements or processes. Recall that the Laplacean super-physicist Martians from Dennett’s thought experiment have no need to appeal to epistemologically emergent phenomena in order to make their predictions. Everything that epistemologically emergent phenomena might allow them predict they can also predict from the laws of fundamental physics. Epistemologically emergent phenomena are nothing more than epistemic crutches that make computations more tractable for beings like us who have limited computational resources and power; the Martians have no need for them.

This, however, is not the case with RP emergence. Insofar as the Martians are blind to emergent real patterns, there are phenomena which they will be unable to predict, but which we who have access to the patterns *can* predict. Suppose, for example, that Bob is ill-tempered. This is a real pattern in his behavior, and we shall assume that it somehow emerges from the physiological and neural processes that take place in his body. On the basis of this pattern, I am able to predict, not only what he *will* do in a given circumstance, but also what he *would* do in an indefinite variety of possible circumstances. The Martians too will be able to predict what Bob will do in a given circumstance. Suppose that I decide to put some *wasabi* in his dessert. The Martians will be able to predict Bob’s ensuing fit of anger, but they can do so only by tracking every physical condition that could conceivably have an effect on the outcome, such as the temperature, the direction of the wind, and perhaps whether a butterfly had been fluttering in a specific location in Brazil the day before. But they cannot predict what Bob *would* do in a merely hypothetical circumstance, because they are unable to isolate the factors relevant to the prediction (such as Bob’s being ill-tempered) from those that are irrelevant. Since I know that Bob is ill-tempered, I can predict how he would act in an indefinite variety of possible circumstances, without having to specify conditions such as the temperature, the direction of the wind, etc. But the Martians cannot make their prediction without all of this information, which in fact is irrelevant to the prediction, and this is because they do not have access to the real pattern that Bob is ill-tempered. In order to know what factors are relevant to making a prediction, one must have access to the appropriate patterns.

Many philosophers seem to think that an emergent pattern embodies less information than descriptions of its underlying elements or processes. In other words, they think that even if a “higher-

level” description of a system is useful or perhaps indispensable in making predictions, it is merely the result of shaving away some of the information contained in the “lower-level” description. Even Dennett, whose theory of real patterns goes a long way towards offering a modern vindication of Scholastic realism, seems to lapse into this position when he characterizes abstract objects as “lossy compression[s]” (Dennett 2000:360). An emergent real pattern, however, can embody *more* information than descriptions of its underlying elements or processes, in that it supports predictions which the latter do not.¹³ And this is precisely because patterns are more *general* than the elements or processes instantiating them. One is tempted to see in the view that emergent patterns are mere “compressions” of their underlying dynamics the “Ockhamistic prejudice” that Peirce refers to in the passage quoted in the epigraph: “Get rid, thoughtful Reader, of the Ockhamistic prejudice of political partizanship that in thought, in being, and in development the indefinite is due to a degeneration from a primal state of perfect definiteness” (CP 6.348, 1907).

4 Bedau on Weak Emergence

In this section I want to throw what I call RP emergence into sharper relief by comparing it with Mark Bedau’s related notion of *weak emergence* (Bedau 1997, 2002). Bedau distinguishes three kinds of emergence: nominal, weak, and strong. A *nominally emergent* property is “a macro property that is the kind of property that cannot be a micro property” (Bedau 2002:9). Thus, properties of water such as fluidity and transparency are nominally emergent properties, since they cannot be possessed by the individual molecules that constitute water. Nominal emergence is the barest and broadest notion of emergence, and encompasses both weak and strong emergence as special cases. Note that it corresponds to condition (1) in my definition of RP emergence. A *strongly emergent* property is one which, in addition to being nominally emergent, is a “supervenient propert[y] with irreducible causal powers” (Bedau 2002:10). These macro causal powers have a determinative influence on both the macro and micro levels, and in the latter case it is called *downward causation*. Although Bedau’s notion of strong emergence is defined differently from what I have been calling ontological emergence, it is similar to the latter in that the causal powers associated with it are “brute” natural powers that arise inexplicably from the micro elements or processes (Bedau 2002:11).

Finally, a *weakly emergent* property of a system is one which, in addition to being nominally emergent, can be derived only through a step-by-step simulation of the system. Bedau’s definition is as follows: “Assume that P is a nominally emergent property possessed by some locally reducible system S . Then P is weakly emergent if and only if P is derivable from all of S ’s micro facts but only by simulation” (Bedau 2002:15). A *locally reducible* system is, roughly, a system whose macro properties are all structural properties—they are wholly constituted by the states and locations of the system’s micro entities—and whose micro dynamics is context-sensitive in the sense that a micro entity’s state depends on the states of its micro-level neighbors (Bedau 2002:14). Since weakly

¹³ Interestingly, Erik Hoel (2017) has arrived at the same conclusion, that a macro-level description of a system can contain more information than a micro-level description, through an application of information theory to the analysis of causal structures; see Hoel (2018) for a non-technical exposition.

emergent properties are macro properties of a locally reducible system, they are wholly constituted by the states and locations of the system's micro entities. In other words, they are ontologically reducible to micro phenomena: "their existence consists in nothing more than the coordinated existence of certain micro phenomena" (Bedau 2002:12). This is what makes weak emergence a weaker form of emergence than strong emergence, which involves irreducible causal powers. On the other hand, the context-sensitivity of the micro dynamics of locally reducible systems entails that understanding how the micro entities behave in isolation or in certain simple contexts does not, in general, enable us to predict how they will behave in more complicated contexts (Bedau 2002:14). Locally reducible systems thus possess a certain kind of unpredictability, and a weakly emergent property is a macro property of a locally reducible system that is unpredictable in a specific sense: it is undervivable from knowledge about the micro entities except by explicit simulation of the system.

Although Bedau frames his definition of weak emergence in terms of *properties*, we can also define weak emergence in terms of *patterns*: Let S be a locally reducible system, and let P be a nominally emergent pattern of S , that is, P is a pattern that arises in S but cannot be manifested by the constituent elements of S . Then P is a weakly emergent pattern if and only if P is undervivable from knowledge about the constituent elements of S except by explicit simulation of S . Hereinafter I will speak of weakly emergent patterns rather than properties.

What deserves emphasis is that the impossibility of deriving weakly emergent patterns except by explicit simulation is not a merely practical impossibility that might be overcome some time in the future, or by beings with greater computational power than humans. Weak emergence has nothing to do with the epistemic limitations of the human mind or lack of available computational resources. Rather, "it involves the formal limitations of any possible derivation performed by any possible device or entity" (Bedau 2002:17). To dramatize this point, Bedau considers a Laplacean supercalculator—not unlike the Martians we have been considering so far—whose computational speed and accuracy are not bounded by any human or hardware-related limitations. He insists that even such a being would not be able to derive weakly emergent patterns except by direct simulation (Bedau 2002:17). This is because the process leading up to the emergence of weakly emergent patterns is, to use the terminology of Wolfram (1985), *computationally irreducible*, that is, as a matter of principle there can be no short-cut derivations of these patterns that are simpler than the natural computational process by which they are generated.

A good way to understand the notion of computational irreducibility is to consider a process that is *computationally reducible*. Suppose, for example, that we throw a stone straight up into the air. Let us assume that the stone is subject only to gravity and an air resistance proportional to its velocity. Given the initial position and velocity of the stone, we can determine, using Newton's laws of motion, the position and velocity of the stone at any desired time t after it has been thrown. Newton's laws thus provide us with a short-cut derivation of the system's state at time t : we do not have to actually go through the entire evolution of the system leading up to time t in order to determine the system's state at that particular time. This is what it means to say that a process is computationally reducible. A computationally *irreducible* process, on the other hand, cannot be bypassed in this way. It is of such complexity that in order to determine the state of the system at some time t , we must explicitly follow the entire evolution of the system leading up to t . The emergence of patterns

in the Game of Life, which we will discuss in detail below, is an example of a computationally irreducible process (and hence of weak emergence).

Even if the impossibility of deriving weakly emergent patterns except by explicit simulation is an impossibility in principle rather than an impossibility in practice, one might still urge that the impossibility is merely an epistemological one, and that weak emergence is therefore a form of what I have been calling epistemological emergence. The argument would go something like this: since, by definition, weakly emergent patterns are nothing more than aggregations of the micro phenomena that constitute them, they do not have any real explanatory power: “all the explanatory power resides at the micro level and the macro phenomena are merely an effect of what happens at the micro level” (Bedau 2002:37). Hence, even if weakly emergent patterns have explanatory autonomy, this autonomy is a merely epistemological one—it amounts to nothing more than “our inability to follow through the details of the complicated micro causal pathways” (Bedau 2002:38)—and does not reflect any autonomous and irreducible feature of reality. Bedau’s reply to this line of argument throws interesting light on the relation between weak emergence and RP emergence. In response to the argument, he makes a distinction between cases of weak emergence for which the argument is sound, and cases for which it is not. He grants that in some cases, weakly emergent patterns are indeed mere effects of what happens at the micro scale, and their explanatory autonomy is merely epistemological (Bedau 2002:38).

As an example, consider John Conway’s Game of Life. This is a cellular automaton consisting of an infinite, two-dimensional lattice of square cells, each of which is in one of two possible states, dead or alive. Time in the Game of Life flows in discrete steps. At each time step, each cell updates its state according to a simple function of its own state and the states of its eight neighboring cells in the previous step. The update rule for the Game of Life is as follows:

- (1) A living cell stays alive if either two or three of its neighbors were alive in the previous step; otherwise it dies.
- (2) A dead cell becomes alive if exactly three of its neighbors were alive in the previous step; otherwise it remains dead.

Given a suitable initial configuration of living and dead cells, the above rule will generate various enduring patterns in the playing field; whether and what patterns appear depend on the initial configuration. For example, there is a particular pattern called a “glider” that moves diagonally across the field, shifting one cell along the diagonal every four time steps (Fig. 1). There are also various kinds of “glider guns” that periodically shoot gliders; one particular type of glider gun, known as the Gosper glider gun, is shown in Fig. 2.




Fig. 1 Glider (the black cells represent living cells and the white cells represent dead cells)




Fig. 2 Gosper glider gun

Even if there is no glider gun, gliders can also be produced accidentally from interactions among other patterns or patternless clusters of living cells. Bedau asks us to consider a “configuration in the Game of Life that accidentally ... emits a stream of six evenly spaced gliders moving along the same trajectory” (Bedau 2002:38). He argues that the emergence of this glider stream is an example of merely epistemological weak emergence, because it has no overarching explanation: “The explanation for the glider stream is just the aggregation of the causal histories of the individual cells that participate in the process” (Bedau 2002:39). The glider stream is similar to the accidental succession of sixes in our example of throwing a die (§2.2): just as there is no real law or *would-be* governing the succession of sixes, so there is no real law or *would-be* governing the glider stream.

On the other hand, suppose there is a glider gun shooting a stream of gliders, as in Fig. 2. As in the previous case, this glider stream can be explained by the aggregation of the causal histories of the individual cells participating in the process. However, there is more to this second glider stream: it is produced by a glider gun, which provides an overarching macro explanation of the stream. This macro explanation is applicable not only to the case at hand, but also to any other instance in which a glider gun shoots a stream of gliders. The aggregate micro explanation omits this information (Bedau 2002:39). Furthermore, the instances in which the macro explanation is applicable include counterfactual situations:

The same glider stream would have been produced if the first six gliders had been destroyed somehow (e.g., by colliding with six other gliders). Indeed, the same glider stream would have been produced if the configuration had been changed into any number of ways, as long as the result was a gun that shot the same kind of gliders. Any such macro gun would have produced the same macro effect. (Bedau 2002:39)

The glider gun is thus autonomous from its underlying micro dynamics, because it is a macro pattern that “supports counterfactuals about what *would* happen in an indefinite variety of different micro situations” (Bedau 2002:41–42, emphasis mine). In other words, it is autonomous because it is an instance of RP emergence.

The accidental glider stream and glider gun are both instances of weak emergence, insofar as they are both underivable without actually going through the Game of Life step-by-step.¹⁴ Nonetheless, the former does not have the autonomy of the latter, and this is because the latter is projectible into counterfactual situations while the former is not. What this shows is that emergent patterns that are autonomous from their underlying elements or processes are so by virtue of their being instances of RP emergence, not by virtue of being instances of weak emergence.

The accidental glider stream is an instance of weak emergence that is not an instance of RP emergence. Hence, not all instances of weak emergence are instances of RP emergence. Conversely, not all instances of RP emergence are instances of weak emergence. As an example of RP emergence that is not an instance of weak emergence, consider a thought experiment put forth by Hilary Putnam in “Philosophy and Our Mental Life” (Putnam 1975:295–97). Suppose we have a rigid board with two holes, a circle one inch in diameter and a square one inch high, and a cubical peg slightly smaller than one inch in each dimension. We want to explain the fact that the peg passes through the square hole but not the round hole. One way of going about would be to regard the board and peg as lattices of atoms, and attempt to calculate all the possible trajectories of the peg from the laws of quantum electrodynamics (if this sounds infeasible, we can suppose that the calculation is carried out by the Martians from Dennett’s thought experiment). We could say that we have attained our explanation if we are able to deduce that the peg never passes through the round hole, but there is at least one trajectory in which it passes through the square hole.

There is, of course, a much simpler explanation. We simply note that both the board and peg are rigid, the round hole is smaller than the peg, and the square hole is larger than the cross section of the peg. This is an explanation that appeals to the *shape* of the holes and peg, which is a pattern that emerges from the way the atoms composing the board and peg are arranged. There is presumably nothing about this pattern that makes it underivable from knowledge about its underlying elements except by explicit simulation; it is a trivial result of the atoms being held together in a certain configuration. It is therefore not an instance of weak emergence. On the other hand, there is an indefinite variety of possible trajectories by which one could attempt to make the peg pass through either of the holes, and the shape of the holes and peg allow us to predict, in the case of any of these possible trajectories, whether the peg would pass through (as well as explain why it passes through or not). Furthermore, the same kind of prediction (and explanation) will hold for any set of objects with the relevant geometrical features, regardless of their size, the material they are made of, etc.

¹⁴ The underivability of patterns in the Game of Life without explicit simulation is a consequence of the fact that the Game of Life is Turing complete, that is, it can be used to simulate an arbitrary computer program. Suppose there is a general algorithm that allows us to accurately predict the behavior of the Game of Life for an arbitrary initial configuration, without going through it step-by-step. Since the Game of Life is Turing complete, this algorithm will also be able to determine the behavior of an arbitrary program with any possible input, including whether it will halt or not. But this contradicts the undecidability of the halting problem, so there can be no such algorithm.

Just like the glider gun, the shape of the holes and peg supports predictions about what would happen in an indefinite variety of possible micro situations, including situations that are not actualized. It is thus an instance of RP emergence.

Concluding Remarks

So far we have developed an account of RP emergence, building on Dennett's notion of real patterns and Peirce's modal analysis of generality. We have seen how RP emergence allows us to avoid both radical emergentism, which holds that there are instances of ontological emergence in the world, and ontological reductionism, which holds that all instances of emergence in the world are epistemological. Finally, I have attempted to throw the concept of RP emergence into sharper relief by comparing it with Bedau's related notion of weak emergence. My aim has been to illustrate how Peirce's theory of generality, as encapsulated in his pragmatism and Scholastic realism, offers us a novel perspective on the issues surrounding emergence, while at the same time bringing Peirce's Scholastic realism up-to-date in the form of a realism about patterns.

Acknowledgements

I wish to express my gratitude to Tatsuya Higaki, Kuniyoshi Morita, and Yoriyuki Yamagata for their insightful comments on earlier versions of this paper. The main ideas in this paper were presented at the Japan Association for Philosophy of Science conference held on June 15–16, 2019 at Kanazawa University; I want to thank the participants of the session for their valuable questions and feedback. This work was supported by JSPS KAKENHI Grant Number JP18J20438.

Abbreviations

CP $x.y$ = *Collected Papers of Charles Sanders Peirce*, volume x , paragraph y .

W $x:y$ = *Writings of Charles S. Peirce: A Chronological Edition*, volume x , page y .

EP $x:y$ = *The Essential Peirce: Selected Philosophical Writings*, volume x , page y .

References

- Alexander, Samuel. 1920. *Space, Time, and Deity: The Gifford Lectures at Glasgow, 1916–1918*, 2 vols. London: Macmillan.
- Anderson, P. W. 1972. "More is Different: Broken Symmetry and the Nature of the Hierarchical Structure of Science." *Science* 177 (4047): 393–96.
- Bedau, Mark A. 1997. "Weak Emergence." *Noûs* 31 (s11): 375–99.
- . 2002. "Downward Causation and the Autonomy of Weak Emergence." *Principia* 6 (1): 5–50.

- Chaitin, Gregory J. 1975. "Randomness and Mathematical Proof." *Scientific American* 232 (5): 47–52.
- Chalmers, David J. 2006. "Strong and Weak Emergence," in *The Re-Emergence of Emergence: The Emergentist Hypothesis From Science to Religion*, eds. Philip Clayton & Paul Davies. Oxford: Oxford University Press.
- Dennett, Daniel C. 1987. "True Believers," in *The Intentional Stance*. Cambridge, MA: MIT Press.
- . 1991. "Real Patterns." *The Journal of Philosophy* 88 (1): 27–51.
- . 2000. "With a Little Help from My Friends," in *Dennett's Philosophy: A Comprehensive Assessment*, eds. Don Ross, Andrew Brook, & David Thompson. Cambridge, MA: MIT Press.
- Hoel, Erik P. "When the Map Is Better Than the Territory." *Entropy* 19 (5): 188.
- . "Agent Above, Atom Below: How Agents Causally Emerge from Their Underlying Microphysics," in *Wandering Towards a Goal: How Can Mindless Mathematical Laws Give Rise to Aims and Intention?* eds. Anthony Aguirre, Brendan Foster, & Zeeya Merali. Springer.
- Humphreys, Paul. 2008. "Synchronic and Diachronic Emergence." *Minds and Machines* 18: 431–42.
- . 2016. *Emergence: A Philosophical Account*. Oxford: Oxford University Press.
- Ladyman, James & Don Ross. 2007. *Every Thing Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press.
- Lane, Robert. 2018. *Peirce on Realism and Idealism*. Cambridge: Cambridge University Press.
- Peirce, Charles S. 1931–1958. *Collected Papers of Charles Sanders Peirce*, vols. 1–6, eds. Charles Hartshorne & Paul Weiss (1931–1935); vols. 7 & 8, ed. Arthur W. Burks (1958). Cambridge, MA: Harvard University Press.
- . 1982–2009. *Writings of Charles S. Peirce: A Chronological Edition*, ed. Peirce Edition Project. Bloomington, IN: Indiana University Press.
- . 1992–1998. *The Essential Peirce: Selected Philosophical Writings*, vol. 1, eds. Nathan Houser & Christian Kloesel (1992); vol. 2, ed. Peirce Edition Project (1998). Bloomington, IN: Indiana University Press.
- Putnam, Hilary. 1975. *Mind, Language and Reality: Philosophical Papers, Vol. 2*. Cambridge, MA: Harvard University Press.
- Ross, Don. 2000. "Rainforest Realism: A Dennettian Theory of Existence," in *Dennett's Philosophy: A Comprehensive Assessment*, eds. Don Ross, Andrew Brook, & David Thompson. Cambridge, MA: MIT Press.
- Seager, William. 2012. *Natural Fabrications: Science, Emergence and Consciousness*. Berlin: Springer.
- Strawson, Galen. 2006. "Realistic Monsim: Why Physicalism Entails Panpsychism," in *Consciousness and Its Place in Nature: Does Physicalism Entail Panpsychism?* ed. Anthony Freeman. Exeter: Imprint Academic.
- Sugano, Reiji. 2013. *Fukuzatsukei Kagaku no Tetsugaku Gairon* (An Introduction to the Philosophy of Complex Systems Science). Tokyo: Hon no Izumi.
- Weinberg, Steven. 1987. "Newtonianism, Reductionism and the Art of Congressional Testimony." *Nature* 330: 433–37.

- Wimsatt, William C. 1976. "Reductionism, Levels of Organization, and the Mind-Body Problem," in *Consciousness and the Brain: A Scientific and Philosophical Inquiry*, eds. Gordon G. Globus, Grover Maxwell, & Irwin Savodnik. Boston: Springer US.
- . 1994. "The Ontology of Complex Systems: Levels of Organization, Perspectives, and Causal Thickets." *Canadian Journal of Philosophy* 24 (sup1): 207–74.
- Wolfram, Stephen. 1985. "Undecidability and Intractability in Theoretical Physics." *Phys. Rev. Lett.* 54 (8): 735–38.