Explaining the Autonomy of the Special Sciences

Alexander Franklin[*]

May 2022

Abstract

Fodor (1997) argues that the special sciences are autonomous, but that this autonomy is mysterious and eludes explanation. Reductionist responses to Fodor tend to eliminativism about autonomy. In this paper I set out a framework for explaining autonomy. Rather than eliminating it, this establishes that the special sciences are, in fact, autonomous from more fundamental sciences, but that this is compatible with reductive explanation.

I cash this out with a case study. Nerve signals are autonomous from the individual ionic motions across the neuronal membrane. In order to explain the autonomy of the nerve signal, we ought to identify the structures at the lower level which give rise to the signal's autonomy. In this case we can do just that: the gated ion channels underwrite the autonomy of nerve signals.

---

[*] alexander.r.franklin@kcl.ac.uk

1 Introduction

> Damn near everything we know about the world suggests that unimaginably complicated to-ings and fro-ings of bits and pieces at the extreme *micro*level manage somehow to converge on stable *macro*level properties.
>
> On the other hand, the 'somehow' really is entirely mysterious … So, then, *why is there anything except physics*? … Well, I admit that I don't know why. I don't even know how to *think about* why.
>
> [Fodor (1997, pp. 160–161), original emphasis]

Fodor notes that the special sciences are autonomous, but maintains that this is mysterious and unexplained. My goal in this paper is to offer an explanation of the co-instantiation of macrolevel stability with microlevel to-ings and fro-ings and, as such, to render this kind of autonomy non-mysterious.

Importantly, the account here differs from many reductionist responses to Fodor (see e.g. Kim (2005)) in that I do not seek to eliminate autonomy. Rather, I explain, in reductive terms, how autonomy comes about. The definition of 'autonomy' is hotly contested: Fodor thinks the autonomy of the special sciences is mysterious, and this mystery is taken by some to be essential to his definition. However, to avoid circularity in the argument that autonomy cannot be explained, he'd better accept that one may resolve the mystery without eliminating the autonomy.

Instead, autonomy relates to the fact that stability comes out of complexity and change – that a unified higher level comes out of heterogeneous lower-level systems. The mystery associated with autonomy is the consequence of an explanatory lacuna, and filling in that explanatory lacuna is non-trivial and involves detailed case-by-case analysis. In this paper I develop a general framework for explaining autonomy, and work through a particular example where the mystery is thus dispelled. The reason that a fair bit of this paper is devoted to this case study is that detailed science is required to show how to-ings and fro-ings can lead to stability.

While some responses to Fodor have focussed on multiple realisation, and indeed a part of Fodor (1997) highlights that topic, I'll prioritise autonomy here. Autonomy is more general than multiple realisation: while the former can be exhibited by a single kind of system, multiple realisation requires that different kinds of system realise the same macrolevel phenomena. Note that the quote with which this paper starts is clear evidence that Fodor intends to include variations among the realisers of a single kind of system within the scope of his 'autonomy', since microlevel to-ings and fro-ings cannot take one heterogeneous realiser to another. However, as suggested below and developed in detail in Franklin (2021)[1] my explanatory framework also applies to instances of multiple realisation.

It's also worth noting that I do not take the presence of autonomy to entail anything about novel causal powers.[2] The challenge of accounting for Fodor's (non-causal) conception of autonomy seems far more pressing. While it's at least a matter of controversy whether or not there are novel causal powers at higher levels, the more basic sense of autonomy – that macrolevel stability somehow comes out of microlevel to-ings and fro-ings – is far more epistemically secure; as emphasised by Fodor, this follows from "[d]amn near everything we know about the world". As such, I'll leave questions about novel causal powers to one side.

We know that the special sciences are autonomous, because we know that the macrolevel is stable while the microlevel involves complex to-ings and fro-ings. We also know that somehow the macrolevel stability comes out of microlevel complexity, the question is how that stability comes about. Although I contend that technical detail is required fully to dispel the mystery from such autonomy, the basic idea is fairly simple. Autonomy is the consequence of processes or structures which render a subset of the microlevel's degrees of freedom irrelevant to certain macrolevel goings-on. As a result, changes in the the microlevel degrees of freedom are compatible with stability at the macrolevel.

In §2 I characterise autonomy in more detail. In §3 I outline a framework for explaining the autonomy of the special sciences in a reductionist-friendly manner. §4 features a case study from neurophysiology which allows me to exemplify such explanations of autonomy. In §5 I conclude.

---

[1] Where I also compare my proposals to the analysis in Polger and Shapiro (2016).

[2] Although aspects of the debate have focussed on this issue, Fodor (1997, fn. 6) is relatively untroubled by the causation arguments raised by Kim.

2 Autonomy

Let's start with a definition: a system is autonomous if its states are invariant with respect to some perturbations in variables at lower (more fundamental) levels. A system may be invariant with respect to bigger or smaller perturbations in the lower-level variables, which means that that system may be more or less autonomous.[3] A science is called 'autonomous' if its wordly subject matter is autonomous in this sense.

That a system is autonomous can thus explain how microlevel complicated goings-on are compatible with macrolevel stability: this would follow if the microlevel to-ings and fro-ings correspond to the perturbations with respect to which the system is invariant.

Autonomy of this kind also licences abstracting away from certain laws and regularities when testing our theories and constructing descriptions of the world. Insofar as such laws characterise the perturbations with respect to which the system is invariant, reference to such laws is not required for predictive accuracy at the higher level. It's in this way that autonomy leads to the relative independence of higher-level special sciences.

As I use the term, 'autonomy' implies the involvement of laws – therefore the regularities of autonomous systems play a role in scientific predictions and explanations.[4] Autonomy implies invariance with respect to changes in the lower-level state: some lower-level perturbations consistent with lower-level theories are thus irrelevant to higher-level autonomous systems. In other words, autonomous evolutions are compatible with various different lower-level happenings.

When talking of 'variables', I intend this plurally: variables feature in mathematised sciences, but one may also refer to aspects of less formal scientific descriptions as variables. In both cases variables refer to the degrees of freedom of concrete physical systems; see J. Wilson (2010).

Autonomy is commonplace across science. The autonomy of the subject matters of the relevant sciences explains how, for example, expert cell

---

[3] This definition of autonomy is closely related to the generalised autonomy developed in detail in Robertson (2021): in the terminology employed there, the perturbations in the lower-level variables are irrelevant to the evolution of the system conditional on the state specified at some higher level, but unconditionally relevant to that evolution.

[4] In place of 'laws', read 'dynamics', 'robust regularities', or 'mechanisms' depending on your view of dependency relations in the special sciences.

biologists may be ignorant of the physics of quarks: cell biology is invariant with respect to perturbations of quarks. That is, changes to the state of the quarks do not make a difference to what the cell does – this justifies the absence of Quantum Chromodynamics from Biology textbooks. Autonomy would be violated if we had a strong coupling between energy scales such that abstraction was impossible (at least without introducing probabilities); if cell biology weren't autonomous from quark perturbations, then abstraction away from quark dynamics by cell biologists would be illicit.

A given system may be autonomous in some circumstances but not in others. For example, a system may be autonomous with respect to some changes at the lower level at low temperatures, but not at high temperatures. An explanation of this might be that some physical structure is in place to screen off a set of lower-level degrees of freedom, but that structure melts at high temperatures.

I have phrased much of this discussion in terms of levels. It's worth emphasising that my analysis does not presuppose a strict, global stratification into distinct levels – that's because I take to heart the lessons of Potochnik and McGill (2012). Rather I presume that, given a set of interacting systems, one can describe those systems in a number of different ways, usually corresponding to distinct spatiotemporal or energy scales. Such different levels will, in general, have laws or dynamics which refer to different variables. The case studies below should also help clarify this terminology.

Before we launch into a scientific case study in §4, I'll flesh out this account of autonomy with a simple example. Bouncy balls are autonomous because their dynamics are invariant with respect to perturbations in almost all of the approximately Avogadro number of lower-level variables which describe the displacements of the balls' constituent atoms. The bouncy ball variables – viz. centre-of-mass position, mass, diameter, coefficient of restitution, and spring constant – are sufficient for most higher-level predictions.[5] In addition, the change of variables from molecular displacement variables to the bouncy ball variables allows for new explanations, which pick up on regularities not available at the lower level.[6]

What is it that makes the states of the bouncy ball invariant with respect

---

[5] See Cross (1999).

[6] See Franklin and Knox (2018) and Knox (2016) for a discussion of how change of variables leads to new explanations.

to a great range of perturbations of many of the lower-level variables? A principal factor in the answer to this question is the compression dynamics: the bouncy ball variables' evolution is insensitive to quite how the ball deforms when it hits the ground because the shock waves dissipate in such a way that the ball maintains its shape.[7]

Further explanations of the bouncy ball's autonomy would need to detail the precise structures and mechanisms which lead to such dissipation. But the detail thus far is sufficient to illustrate how autonomy may be explained. In this particular case we can explain the autonomy of the bouncy ball by showing that there are structures in place which render lower-level perturbations irrelevant to higher-level dynamical evolution.

One way to conceptualise the reductive explanation under discussion is as a difference-making explanation. So we can ask: what would happen if the shock waves didn't dissipate and the ball didn't recover its shape? If the ball splatted on the ground, then the states at later times would not be autonomous from the molecular displacement variables: the precise way in which the ball landed, and the surface on which it landed would have a significant dynamical effect. As such, the bouncy ball is not autonomous to the same degree in circumstances where the shock waves do not dissipate. This suggests that the structures responsible for dissipation form at least part of the reductive explanation of the bouncy ball's autonomy.

It's worth noting once again that more details could be given to deepen this explanation. Nonetheless, this demonstrates how explanations can account for the higher-level autonomy, at least in principle; a more fully worked out case study is given in §4.

I'll conclude this section by commenting on how autonomy on my definition relates to alternative accounts of autonomy. From the start I have been explicit that I am focussing on a conception of autonomy inspired by Fodor: macrolevel stability coupled to microlevel to-ings and fro-ings. However, my strategy for explaining autonomy is also applicable to other conceptions. Insofar as the higher-level states are invariant with respect to perturbations at the lower-level, we may cash out the higher-level description functionally – the higher level involves an abstracted relation between certain functions of lower-level variables. In addition, the irrelevance of lower-level degrees of freedom allows for higher-level variables to be re-

---

[7] M. Wilson (2006, 2017) refers to the abstraction away from such processes as 'physics avoidance'.

lated to open-ended disjunctions of lower-level variables.[8] This can be seen by considering the bouncy ball once again:

To be a bouncy ball is to act like a bouncy ball; generic bouncy balls needn't have any specific microphysical constitution – bounciness and approximate sphericality are clearly realisable in a bunch of different ways. So, bouncy balls satisfy the bouncy ball theory: describing something as a bouncy ball suggests that the height of its next bounce will be predictable from values for the five variables listed above. The autonomy of the bouncy ball – that its variables are invariant with respect to a range of perturbations of lower-level variables – thus implies that the bouncy ball corresponds to an open disjunction of lower-level realisers.

In order to satisfy the bouncy ball theory at the higher level, each bouncy ball realiser must have some structures in place which guarantee that perturbations of most lower-level variables are irrelevant to the bouncy ball dynamics. As such, the autonomy of the bouncy ball is explained, for each realiser, if we can identify the particular structures that lead to dissipation of compression shock waves in such a way that the shape is maintained.

In the bouncy ball case, structural features of each realiser can explain the autonomy that bouncy balls exhibit. Some may find it problematic that different bouncy ball realisers may then have different explanations of autonomy. Nonetheless, in every case autonomy is explained.[9]

3 Reductive Explanation of Autonomy

Having defined autonomy, and motivated the view that autonomy can be explained, in this section I characterise explanation of autonomy more abstractly.

To set the stage it's important first to distinguish between two kinds of explanation, which I'll call 'horizontal' and 'vertical' to correspond, respectively, to intra-level and inter-level explanation.

First, consider horizontal explanations. Something of a consensus in the philosophy literature has emerged that it is an explanatory virtue to leave

---

[8] Open disjunctions are discussed in Fodor (1997, p. 156).

[9] As these other conceptions of autonomy are much closer to multiple realisability see Franklin (2021) for further discussion of this matter.

out irrelevant detail; see e.g. Strevens (2008), Jansson and Saatsi (2017), Craver and Kaplan (2018), and Woodward (2018). Horizontal explanations are proportionate: they refer to just enough detail in order to identify the difference-makers for the *explanandum*. Horizontal *explanantia* are at the same level as the *explananda*. So long as some kind of proportionality is required, horizontal explanations can be characterised on any account of explanation.

Second, consider vertical explanations. These are any kind of reductive explanation whereby features of higher-level system are explained in terms of the corresponding lower-level system. The kind of vertical explanations which are the focus of this paper do not take any specific phenomenon *at* the higher level as their *explanandum*: these vertical explanations explain the autonomy *of* the higher level phenomena.[10] Henceforth, I'll only discuss this kind of vertical explanation.

The distinction between these two kinds of explanation plays a crucial role in the argument of this paper. If one has the relevant vertical explanations to hand, then this does not signal eliminativism. That's because vertical reductive explanations do not explain the phenomena at the higher level; rather, they explain the higher level phenomena's autonomy. Even where we can explain the autonomy, *explanantia* at the higher level still provide the best, horizontal, explanations of higher-level *explananda*. As noted, on many accounts of explanation, proportionate explanations are better explanations. Since higher-level explanations that refer to autonomous systems are more proportionate, such systems shouldn't be eliminated, even if their autonomy can be reductively explained.[11]

So, how do vertical explanations work? In order to give a vertical explanation, the lower-level features responsible for higher-level autonomy must be explicitly noted. The higher-level autonomy may be expressed as invariance with respect to perturbation of a class of lower-level details. In general, wherever one finds higher-level autonomy, there are structures or mechanisms which make those variables that are relevant at the lower level irrelevant at the higher level. Identifying such structures or mechanisms provides vertical, reductive explanations.

Let's call the 'natural' lower-level variables $\{l\}$ in terms of which lower-

---

[10] Batterman (2017) makes a related distinction between type I and type II explanatory questions, where, on my account, type II questions are addressed by vertical explanations.

[11] Antony (2003) and Antony and Levine (1997) make similar arguments against eliminativism of higher-level properties.

level laws are most straightforwardly expressed. The higher-level variables will, in general, not correspond directly to elements of $\{l\}$ but to elements of $\{l'\}$, which are related to $\{l\}$ by some change of variables perhaps involving functions of elements of $\{l\}$.[12] What's required by vertical explanation is to show, in terms of either $\{l\}$ or $\{l'\}$, how and why certain elements of $\{l'\}$ are relevant and other elements of $\{l'\}$ are irrelevant. Insofar as that can be achieved in lower-level terms, a vertical explanation is achieved, even if the salience of the change of variables to $\{l'\}$ was not at all evident before it was made.

Note that appeal to 'relevance' and 'irrelevance' does not imply an epistemic distinction. Irrelevant variables are those with respect to which the system is autonomous – i.e. irrelevant variables are those whose perturbations don't affect the higher-level state of the system. Relevant variables are those which aren't irrelevant.[13]

In the bouncy ball case considered above, the lower-level variables are all the positions and momenta of the molecules which constitute the ball. $\{l'\}$ are the lower-level variables which are functions of $\{l\}$ including, for instance, the centre of mass variables, along with a great many other functions of lower-level variables. The higher-level variables have a one-to-one correspondence with the relevant subset of $\{l'\}$. The autonomy may thus be stated as the irrelevance of almost all the elements of $\{l'\}$ to the prediction of, say, the height of the next bounce. Explaining autonomy thus requires identifying structures that secure the irrelevance of such elements of $\{l'\}$. As noted above, such structures will ensure that the compression shock waves dissipate in such a way that the shape is recovered.

My goal in this paper is to dispel the mystery Fodor finds in autonomy. Vertical explanation is thus reductive insofar as the putative mystery motivates Fodor's anti-reductionism. In particular, reductive explanation of the sort discussed here plays a significant role in rebutting the anti-reductionist argument that bottom-up approaches are unable to account for the explanatory value and autonomy of higher-level science. This kind of reduction satisfies Fodor's injunction to take the special sciences seriously – this is because reductive explanation of autonomy serves to establish autonomy rather than eliminating it. By explaining how autonomy comes about, I provide an extra reason for taking autonomy seriously. The autonomous

---

[12] Consider, for example, the changes of variables discussed in Franklin and Knox (2018) and Knox (2016).

[13] This terminology comes from the renormalisation group in physics.

system plays an essential role in higher-level explanation and prediction, as such it ought not to be eliminated.

While many reductionist responses to Fodor can be found in the literature, this paper's focus on reductive explanation offers a needed and distinctive contribution. That's because existing analyses, such as Antony and Levine (1997), only establish that reduction is, in principle, compatible with a certain kind of autonomy. However, Fodor focusses on the empirical phenomenon that macrolevel stability comes out of microlevel to-ings and fro-ings. That empirical phenomenon deserves a scientific explanation, and, consequently, mystery remains unless each instance of such autonomy is reductively explained. As I'll further show in the next section, the explanatory framework developed here allows existing scientific research to be repurposed to dispel that mystery. My approach should complement existing philosophical accounts, but I claim that scientific explanations are additionally required to render empirically discovered autonomy non-mysterious, and this paper shows how such explanations do that job.[14]

4 Case Study

What's gone wrong in thinking autonomy mysterious is, I claim, not a conceptual error. Rather it is a failure to pay sufficient attention to the science itself. My argument is that, if conceived in the right way, the question of how macrolevel stability comes out of microlevel to-ings and fro-ings is a question which science can and does answer. The goal of this section is to make this point by exhibiting some relevant science. To that end, I go into some detail – it's not easy to explain autonomy, but it can be done and the rest of this section shows precisely how the framework outlined above has the desired pay-off.

Consider a person's reflex movement in reaction to a pin's pricking their finger. At a higher level this is mediated by a nerve signal, which is, at a lower level, constituted by a series of ionic motions.

Nerve signals play a crucial role in many physiological explanations and predictions. Yet such signals are autonomous from a range of ionic

---

[14] Although Batterman (2000; 2018) also employs scientific explanations to account for autonomy, he argues that such explanations are non-reductive; see e.g. Saatsi and Reutlinger (2018) for a critique of Batterman's physics-based arguments.

motions: only a small subset of collective ionic motions make a difference to the nerve signal. In order to explain the autonomy of nerve signals, we need to look for structures and mechanisms which prevent the irrelevant ionic motions from affecting such signals. The salient structures on which I'll focus below are the gated ion channels – these implement a voltage threshold for ionic motion across the neuronal membrane.

4.1 Action Potential

The primary job of neurons is to transmit signals within the body. Such signals are transmitted by the sequential movement of ions between the fluid outside the neuron (the extra-cellular fluid) and the fluid inside the neuron (the cytosol or intra-cellular fluid). Such ions primarily move in one of three ways: ion channels, gated ion channels and ion pumps. Signals are triggered and move down the axon through the activation of an action potential. See e.g. Bear, Connors, and Paradiso (2007, p.80) for more details.

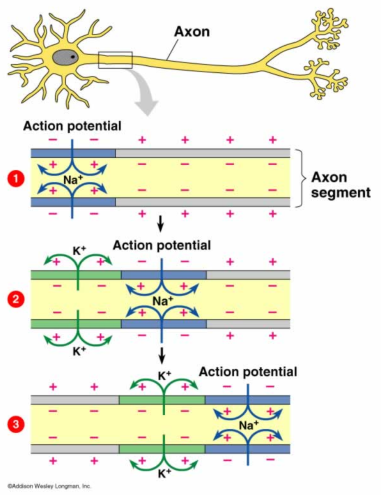

©Addison Wesley Longman, Inc.

Figure 1: Movement of ions changes the potential which leads to further movement of ions.

The neuron has a resting potential of about -65mV. When appropriately triggered – e.g. by a pin prick increasing the potential towards zero – voltage sensitive sodium gated channels will open and allow sodium ions (Na^+) to rush into the cell (① in figure 1). After the potential has increased sufficiently the gated sodium channels close (inactivation), and the neuron's resting potential is quickly restored by the opening of the potassium channels and consequent diffusion of potassium ions (K^+) along the concentration gradient (② in figure 1).

The signal is transmitted from end to end of the neuron via a relay of action potentials: a positive potential in one part of the neuron will raise the charge in nearby areas and trigger the opening (activation) of nearby (voltage sensitive) gated channels, this triggers further action potentials, and so on (③ in figure 1). The gates have an inbuilt delay after inactivation (before deinactivation) which leads to the signal's only propagating in one direction. The ions are only free to move between the extra-cellular fluid and the cytosol after the gated channels have been activated.

4.2 Explaining The Autonomy of the Signal

In order to exemplify my account of reductive explanation I needn't detail all the structures required for the autonomy of nerve signals from individual ionic motions; it will be sufficient just to describe the gated ion channels. These are part of what ensure that the signal is binary: they prevent ionic transmission except when the voltage is above some threshold.

Individual ionic motions are irrelevant to the signal while the gate is inactivated, as such they prevent gradual flow of ions across the membrane. This allows for an imbalance of ions which results in a significant flow when the gate is activated. The existence of the threshold for activation means that the signal is determined by just a few collective motion variables, rather than a multitude of individual motions. Consequently, the gate is part of the structure which ensures that signals are unaffected by individual ionic motions but are dependent upon certain collective ionic motions.

In the terminology of the previous section: the lower-level variables $\{l\}$ correspond to individual ionic motions. The autonomy of the signal is, in part, a consequence of the fact that explanations and predictions depend upon fairly few elements of $\{l'\}$ – functions of the lower-level variables. One way to choose the elements of $\{l'\}$ would be as follows: divide the neuron into columns such as those depicted in figure 1; for each column, pick a type of ion; lastly construct weighted collective variables for ionic motions. Some elements of $\{l'\}$ would then correspond to the centre of mass variable for each ion type for each column – it's these which are relevant to the propagation of the signal across the neuron: the transmission of the signal depends on a relay of such variables crossing the membrane. Following this formula there are many other elements of $\{l'\}$; these correspond, for example, to collective variables which are weighted in favour of ions much further away from the membrane – such variables will not reliably track the motions which succeed the activation of the gated channel and will consequently be irrelevant to the signal.

The fact that we can abstract away from most such elements of $\{l'\}$, and just characterise the system in terms of whether or not a few relevant collective variables have crossed the membrane, is a consequence of the signal's autonomy. Explaining this autonomy involves considering the existence of the threshold required for activating the gated channels. It is the voltage-gated ion channels which are responsible for the autonomy and the abstraction away from almost all the elements of $\{l'\}$.

The question which remains is: can we provide a reductive account of the voltage-gated ion channel? I will explain how potassium gated ion channels work; for more details see e.g. Long, Campbell, and MacKinnon (2005), for details of sodium gated ion channels see e.g. Catterall (2000).

Open (or activated) potassium gated channels allow potassium ions through at a rate 1,000 times that of sodium. They do this despite the fact that sodium ions have a third smaller radius. The mechanism for this relies on the fact that ions passing through these channels are bound to water molecules. The arrangement of oxygen atoms bound to the proteins in the gate is such that the potassium ions can swap their binding from the water molecules to the protein's oxygen atoms. The smaller sodium ion is more tightly bound to its water molecules and so cannot likewise transfer its binding to oxygen atoms in the gate. As such, specific gates select for potassium ions over sodium ions.

Importantly for the explanation of autonomy, these gates can be opened and closed. This means that only relatively few ions can leak through the membrane until some threshold voltage is reached. At that voltage the gate is activated and sufficiently many ions can pass through to constitute part of an action potential and lead to a signal.

The opening and closing of the gate is depicted in figure 2. This shows voltage sensors, labelled 1-4 on S4, which move as the voltage changes. These sensors are attached to a lever known as a 'voltage-sensor paddle unit'. The movement of these sensors lead to the closing of the pore, depicted in the centre of the diagrams in figure 2A and to the right in 2B.

Long, Campbell, and MacKinnon (2005, pp. 907–8) give more detail:

> At a qualitative level, one can understand how a transmembrane electric field, by working on the positive Arg charges on S4, can open the pore when the membrane is positive inside (pushing the charges out) and close the pore when the mem-

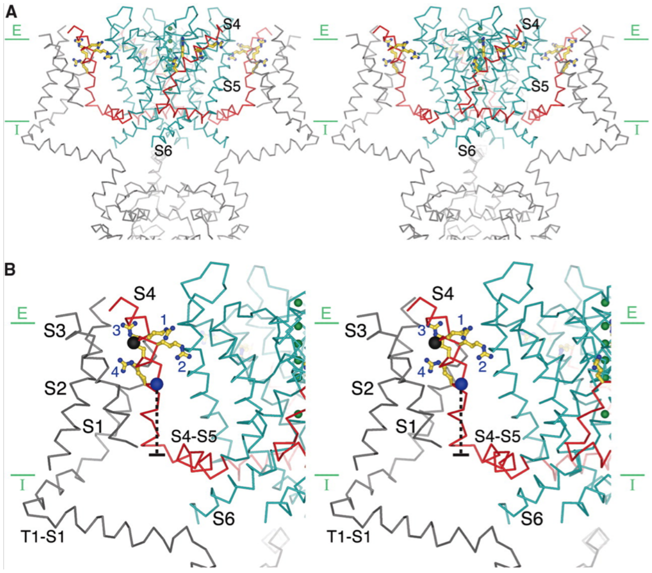

Figure 2: From Long, Campbell, and MacKinnon (2005). Voltage-sensitive ion channels. On the left the channel is open, on the right it's closed. B is an enlarged view of A. In B, 1-4 on S4 are charged and respond to the potential by moving S3 and S4 together. This acts as a 'voltage-sensor paddle unit' which closes or opens the pore in the centre of the A diagrams and to the right of the B diagrams.

> brane is negative inside (drawing the charges in) . . . We therefore suppose that S3 and S4 move together as a voltage-sensor paddle unit. We imagine that to close the channel, the paddle undergoes a motion with respect to S1 and S2, with S3 remaining "above" (on the extracellular side of S4) and S4 "below", closer to the intracellular solution.. . . In this way, S3 in the voltage-sensor paddle might serve as a recoil device, causing the voltage sensor to spring to its open conformation when the membrane is depolarized.

While research into the structure of ion gates is important for its applications to, for example, the treatment of various diseases, the understanding generated has additional upshots for the present paper: we can see part of what's required in order to make the signal autonomous from individual ionic motions. If there were no gates, then there would be no signal: the existence of the signal requires that some collective motions are relevant and others are not, and this is only possible if there are conditions when ions can travel across the membrane, and other conditions when such motion is suppressed.

The research projects to which I've referred above explicitly aim to provide a mechanism, in molecular terms, for the functioning of ion channels. These explanations thus form part of a reductive project. The existence of voltage-gated ion channels underlies the autonomy of the signal. Where the voltage is below the threshold for gate activation, individual ions may to and fro howsoever they like, but such to-ing and fro-ing will not constitute a signal; only when the gate is activated will the ionic motions make a difference to the signal. As such, the macrolevel stable signal comes out of the microlevel ionic to-ings and fro-ings as a consequence of the existence of the voltage-gated ion channel.

This project also leads to an explanation of autonomy on the open disjunction conception. Signals are described in extremely sparse terms – e.g. as a sequence of on-offs which allow for the transmission of information. The neuron may be understood as a signal transmitter because structures are in place which make certain collective ionic motion variables relevant and other variables irrelevant. Analogous structures will make it such that signals are instantiated in different kinds of system.

Say, for example, that one were to construct an artificial neuron out of plastic, or metal, or just about any kind of material with the requisite properties. In each case, there would be structures or processes which made many of the microlevel to-ings and fro-ings irrelevant to the macrolevel signal. It is a consequence of the generic availability of such structures that the higher-level signal corresponds to an open disjunction at the lower level. I've told the story of how autonomy arises for just one of the disjuncts; it's my contention that accounts could be given for other disjuncts, but such accounts may well be different.

It is no easy task to explain autonomy in any case, and any explanation requires scientific detail. Moreover, explanations of autonomy are hostage to the success of various scientific projects. The account here is, however, sufficient to demonstrate the principle that autonomy can be explained.

It should be clear that, however one wishes to characterise autonomy (as in §2, functionally, via open disjunctions, or otherwise), signals are autonomous: theories which describe signals may be quantitatively precise but apply in a great range of different cases, and signals are stable despite microlevel complexity. As such, the explanation on offer in this section is an explanation of autonomy. Whether or not the autonomy of all phenomena can be likewise explained is an open empirical question. Insofar as unexplained autonomy is evidence against reductionism, then the latter is also hostage to empirical fortune, in line with e.g. Fodor (1974).

5 Conclusion

In the quote from Fodor at the beginning of this paper he says that he doesn't "even know how to *think about* why" there isn't only physics. I have formulated a framework for thinking about why the special sciences are autonomous. If the autonomy of the special sciences is a reason to include the subject matters of those sciences in one's ontology, this should help address Fodor's quandary: there isn't only physics because there are wordly structures which make certain microlevel to-ings and fro-ings irrelevant to the macrolevel stability described by the special sciences. As such, the explanation of the autonomy of the special sciences helps explain why there are special sciences at all.

I aim to reject both eliminative reductionism and anti-reductionism. A middle road is available, where the special sciences are autonomous, but this is not mysterious. We can, in fact, explain how macrolevel stability comes out of complicated microlevel to-ings and fro-ings: such explanations involve the scientific discovery of structures which secure the irrelevance of a subset of degrees of freedom. The identification of such structures is non-trivial and it is an open empirical question whether they can be found in all cases. In those contexts where such structures cannot be found, the autonomy of the special sciences remains mysterious.

Acknowledgements

Thanks to audiences at the FraMEPhys workshop on Levels, the DIEP workshop on Emergence, the Rutgers-Bristol Workshop on the Metaphysical Unity of Science, King's College London Philosophy department and to Patrick Butlin, Donald Franklin, John Heron, Tom Hodgson, Sam Kimpton-Nye, Eleanor Knox, Tuomas Tahko, Katie Robertson, and anonymous referees for helpful discussion and comments. This work was partly supported by the European Research Council Project 'The Metaphysical Unity of Science' (ERC CoG), grant no. 771509.

References

Antony, Louise (2003). "Who's afraid of disjunctive properties?" In: *Philosophical Issues* 13, pp. 1–21.

Antony, Louise and Joseph Levine (1997). "Reduction with autonomy". In: *Philosophical perspectives* 11, pp. 83–105.

Batterman, Robert W. (2000). "Multiple Realizability and Universality". In: *The British Journal for the Philosophy of Science* 51.1, pp. 115–145.

— (2017). "Philosophical implications of Kadanoff's work on the renormalization group". In: *Journal of Statistical Physics* 167.3-4, pp. 559–574.

— (2018). "Autonomy of theories: An explanatory problem". In: *Noûs* 52.4, pp. 858–873. DOI: `10.1111/nous.12191`.

Bear, Mark F., Barry W. Connors, and Michael A. Paradiso (2007). *Neuroscience: Exploring the Brain*. 3rd. Lippincott Williams & Wilkins Publishers.

Catterall, William A (2000). "From ionic currents to molecular mechanisms: the structure and function of voltage-gated sodium channels". In: *Neuron* 26.1, pp. 13–25.

Craver, Carl F. and David M. Kaplan (2018). "Are More Details Better? On the Norms of Completeness for Mechanistic Explanations". In: *The British Journal for the Philosophy of Science*. DOI: `10.1093/bjps/axy015`.

Cross, Rod (1999). "The bounce of a ball". In: *American Journal of Physics* 67.3, pp. 222–227.

Fodor, Jerry A. (1974). "Special Sciences (Or: The Disunity of Science as a Working Hypothesis)". In: *Synthese* 28.2, pp. 97–115.

— (1997). "Special sciences: Still autonomous after all these years". In: *Noûs* 31.s11, pp. 149–163.

Franklin, Alexander (2021). "Can Multiple Realisation Be Explained?" In: *Philosophy* 96.1, pp. 27–48. DOI: `10.1017/S0031819120000285`.

Franklin, Alexander and Eleanor Knox (2018). "Emergence without limits: The case of phonons". In: *Studies In History and Philosophy of Modern Physics* 64, pp. 68–78. DOI: `10.1016/j.shpsb.2018.06.001`.

Jansson, Lina and Juha Saatsi (2017). "Explanatory Abstractions". In: *The British Journal for the Philosophy of Science*. DOI: `10.1093/bjps/axx016`.

Kim, Jaegwon (2005). *Physicalism, or something near enough*. Princeton University Press.

Knox, Eleanor (2016). "Abstraction and its Limits: Finding Space For Novel Explanation". In: *Noûs* 50.1, pp. 41–60. DOI: `10.1111/nous.12120`.

Long, Stephen B., Ernest B. Campbell, and Roderick MacKinnon (2005). "Voltage Sensor of Kv1.2: Structural Basis of Electromechanical Coupling". In: *Science* 309.5736, pp. 903–908. DOI: `10.1126/science.1116270`.

Polger, Thomas W. and Lawrence A. Shapiro (2016). *The Multiple Realization Book*. Oxford University Press.

Potochnik, Angela and Brian McGill (2012). "The limitations of hierarchical organization". In: *Philosophy of Science* 79.1, pp. 120–140.

Robertson, Katie (2021). "Autonomy generalised; or, Why doesn't physics matter more?" In: *Ergo*. Forthcoming.

Saatsi, Juha and Alexander Reutlinger (2018). "Taking Reductionism to the Limit: How to Rebut the Antireductionist Argument from Infinite Limits". In: *Philosophy of Science* 85.3, pp. 455–482. DOI: `10.1086/697735`.

Strevens, Michael (2008). *Depth: An account of scientific explanation*. Harvard University Press.

Wilson, Jessica (2010). "Non-reductive Physicalism and Degrees of Freedom". In: *The British Journal for the Philosophy of Science* 61.2, pp. 279–311. DOI: `10.1093/bjps/axp040`.

Wilson, Mark (2006). *Wandering Significance: An Essay on Conceptual Behaviour*. Oxford University Press.

— (2017). *Physics Avoidance: Essays in Conceptual Strategy*. Oxford University Press.

Woodward, James (2018). "Explanatory autonomy: the role of proportionality, stability, and conditional irrelevance". In: *Synthese*. DOI: `10.1007/s11229-018-01998-6`.