Mechanisms, Resources, and Background Conditions*

Colin Klein
Department of Philosophy
Macquarie University
Sydney, Australia
colin.klein@mq.edu.au

Abstract

Distinguishing mechanistic components from mere causally relevant background conditions remains a difficulty for mechanistic accounts of explanation. By distinguishing *resources* from mechanical parts, I argue that we can more effectively draw this boundary. Further, the distinction makes obvious that there are distinctive resource explanations which are not captured by a traditional part-based mechanistic account. While this suggests a straightforward extension of the mechanistic model, I argue that incorporating resources and resource explanations requires moving beyond the purely local account of levels that some mechanists advocate.

Acknowledgements
Thanks to Peter Clutton, Liz Irvine, Annelli Janssen, Meg Ivory, Chris Lean, Marco Viola, Elizabeth Schier, and participants at the 2014 Cognitive Science Kioloa workshop for helpful comments on earlier drafts. Thanks to Yellow Car, and indirectly to Karen Bennett, for inspiring many of the automotive examples. This research was supported by Australian Research Council grant FT140100422

*October 2017 draft

1 Parts and Resources

Back in the day, my car began to overheat. Potential explanations fell into two categories. First, a mechanical failure: a stuck thermostat, a failing pump, a blocked radiator. Second, a problem with the coolant itself: contamination, age, leaks. Any of these, if true, would be part of a perfectly respectable mechanistic explanation of my car's overheating. Yet *qua* explanations, they also have importantly different properties.

The first set of explanations appeals to what I'll call *mechanical parts*: discrete entities with well-defined roles to play. Mechanical parts are individually present or absent. They work or they don't. The second set of explanations, by contrast, appeals to a *resource* that is used by mechanical parts in the course of performing their roles. Resources usually have a quantity and a quality: we ask *how much* coolant the car has, and what *degree* of contamination it might show. Resources often play multiple roles. The very same stuff is involved in cooling the engine and heating the passenger compartment.

This paper will do two things. First, it will clarify and argue for a distinct notion of resource explanations. Resource explanations are a kind of mechanistic explanation. Anything that shows up in the explanation of car troubles seems like it ought to be. Yet resources have some features that make them look more like background conditions. By separating them off and detailing their role, we thus find it easier to draw the distinction between background conditions and true mechanical components (which include both mechanical parts and resources). Hence by showing how resource have distinct features—while still remaining under the broadly mechanist umbrella—we make the problem of carving off background conditions much easier.

That said, the move to resources ends up highlighting a problem for mechanisms. As noted, mechanistic explanation as a research strategy has focused on local decompositional hierarchies that are drawn relative to explanations of particular phenomena. There is no requirement to coordinate parts across different explanations; insofar as coordination *is* done, it is merely convenience. I think the existence of resource explanations make this position untenable. Resources are *causally promiscuous*: they regularly appear in explanations of different phenomena. Further, in some cases it is important to identify the same resource across different explanations, for only then can we explain the role of things like resource competition. So, secondly, I will argue that the local levels thesis is untenable. Recognizing this, I argue further, lets us make sense of some otherwise problematic cases for the mechanistic program, while preserving a broadly mechanistic spirit.

2 Two Categories of Mechanistic Component

2.1 Background conditions

To motivate my position, I focus on a persistent problem for the mechanist: that of distinguishing mechanical components from background conditions.

Suppose we explain how a car engine works. That explanation will assume (usually tacitly) that the engine is properly lubricated, in an environment with an ambient temperature in a certain range, and so on: these are mere background conditions that are necessary for the mechanism to work, but which are contrasted with the mechanical components that actually do the explaining. One would like that distinction to be principled, and so a good account of mechanistic explanation requires a criterion to distinguish background conditions from mechanistic components.

Distinguishing background conditions runs into two distinct problems that ought to be distinguished. First, mechanistic explanations are a species of causal explanations, and all causal explanations must distinguish background conditions from the meatier *explanans*. This is a well-known problem for any causal theory of explanation, one that has attracted a lot of attention (Menzies, 2007). For that reason, I set it aside. I am interested in a second problem that is specific to the mechanistic project.

Mechanistic explanations, at least in contemporary formulations, appeal to *local decompositions* of target phenomena (Craver, 2007; Bechtel, 2008). Explaining the activity of a target object involves appeal to the activity of some of its spatiotemporal parts. Not all parts are equally relevant, though. If a wombat crawls into my engine compartment and goes to sleep, that doesn't make it a mechanistic component of my engine—even if it has a variety of causal interactions with the engine. The criterion that rules out the wombat thus needs to rule out spatiotemporal parts that aren't parts of the mechanism even if they are causally related to the mechanism—and that needs to include background conditions as well.

This is arguably a distinct, and harder, problem than the one that all causal theories face. The standard problem of background conditions starts with a well-delineated set of potential causes, and asks which we should prefer. The mechanist, however, must also find a way to delineate the potential causes in the first place. That is the second, and more pressing, problem of background conditions. I focus on the account of mechanistic explanation given by Carl Craver. A mechanism, in Craver's (2007, 5) formulation, is "a set of entities and activities organized such that they exhibit the phenomenon to be explained." Mechanistic explanations are fundamentally multi-level. They work by demonstrating that the behavior of a whole can be captured in terms of the organized activities of its parts (2007, 139).

Craver offers the relationship of *constitutive relevance* in order to distinguish the true mechanistic components. Constitutive relevance is in turn cashed out in terms of mutual manipulability. Some proper part x of S is constitutively relevant to S's ψ-ing if there is an intervention on x's ϕ-ing that would change

S's ψ-ing, and if there is an intervention on S's ψ-ing that would change x's ϕ-ing (Craver, 2007, 153). In other words, it must be that one is able to reliably manipulate the behavior of a constitutively relevant component to manipulate the behavior of the whole, and vice-versa. The behavior of the fuel injectors can be manipulated by speeding up the engine, and the engine can be manipulated by changing the speed of the injectors; that is why the injectors count as constitutively relevant. By contrast, there's no (reliable) change I can make in the interloping wombat by manipulating the engine, and poking the wombat has an extremely inconsistent effect on the engine.

Craver also argues that the relationship of mutual manipulability fails for mere background conditions (Craver, 2007, 153). Lubricants are part of the causal chain between pushing the gas pedal and making the car go. But the manipulability relation is asymmetric and non-specific. Pressing the pedal makes no difference to the presence of lubrication. Changing whether there's sufficient lubrication makes a difference to the target phenomenon, but only by making everything break down.

Mutual manipulability does seem to carve off very general background conditions. Yet problem cases remain. Consider: Forming long-term memories via long-term potentiation uses glucose, which is largely provided by astrocytic glycogen. Depleting glycogen experimentally leads to amnesia. The elderly often have memory problems, and also tend to have lower levels of glucose. Supplementing them with glucose improves their memory (Messier, 2004; Newman et al, 2011). Hence forming memories depletes astrocytic glycogen, and depriving astrocytes of glycogen causes amnesia.

Hence astrocytic glycogen passes Craver's test for mutual manipulability. This seems like a problem. Intuitively, glycogen doesn't quite seem like part of the mechanism of memory formation: it's the fuel that drives the mechanism.

The problem may simply be that mutual manipulability is an insufficiently sensitive tool for discerning mechanical components. I think there is a more interesting solution. For there is an important sense in which astrocytic glucose *isn't* like lubrication, ambient temperature, and so on. Instead, then, I claim that mechanical components are more diverse than they originally seemed. I argue we ought to distinguish at least two types of component: mechanical parts and resources.

2.2 Mechanical parts

Fuel injectors, hearts, sodium channels, and syntactic modules appear to have, on the face of it, little in common—other than that they are all cited in mechanistic explanations. Abstracting away a bit, though, we can see that they share several important features. Call components with these features *mechanical parts.*

First, mechanical parts *persist* over relatively long timescales—at least over the timescale of the behavior of the whole that we wish to explain. The different engine components are present the entire time the engine is running. That is an important feature of the mechanistic explanation: we implicitly assume that when the fuel rail sends gas to the injectors, the injectors are always there and ready to receive it.

Conversely, the failure of a mechanical part is often the explanation for the failure of the whole mechanism to perform. The Challenger exploded rather than went into space because an O-ring failed. Tetrodotoxin kills by blocking the activity of sodium channels. Mechanical parts must thus be available and functional throughout the timescale of the explanation. Failure to do so explains deviations from the norm.

Second, mechanical parts are *individually important*. In addition to the availability of component-types, the identity of a component over time is often crucial to the explanation. My car starts hard in the morning. Why? It is too old to take ethanol-augmented gasoline. The ethanol destroyed the check valve in the fuel rail. As it sits overnight it loses pressure and makes for a hard start. Note that for this explanation to work, it is important that the same entity was affected by the ethanol, leaked overnight, and therefore explains what happens the next day.

Thus mechanical parts don't just persist, but can persist as the very same entity over a series of activities. This is often obscured when we give mechanistic explanations because we often care about explaining *types* of mechanisms rather than individual tokens. But for any individual token mechanism, it is usually assumed that the same components persist, and persist as the same components, over time.

Third, mechanical parts are typically portrayed as *causally conservative*. That is, each mechanical part interacts with only a limited subset of the other components in a small number of ways. The fuel rail does one thing, the intake another, the injectors a third; further, the fuel rail interacts only with the fuel line and the injectors, the injectors with the rail and the chamber, and so on. Causal conservativeness allows for *modularity*, which is a feature of complex mechanisms (Simon, 1996; Bechtel and Richardson, 2010). Further, mechanistic decomposition is itself aided by the causal conservativeness of the individual components. Most of the processes we care about are complicated. We explain a complicated process by looking at the components involved and the limited sets of relationships between them. Note that the components need not be *simpler* than the whole. Rube Goldberg machines show that it's possible to build up a simple process from complex components. What's important, rather, is that each mechanical part affects something *less* than the whole, and so we can partition out the effects of mechanical parts more easily.

Fourth and finally, our mechanistic explanations are often *indifferent to the composition* of the mechanical parts. Many mechanistic explanations care only about the location and characteristic activities of the mechanical parts. Mechanistic explanations often involve *functional* decompositions (Cummins, 1975, 1983), and functional decomposition cares only about the activities of the components and their relation to one another.

Note that this is a weaker claim than saying that mechanical parts are multiply realizable, at least in the classic sense. Mechanical parts might well be constrained, and rather severely, by the environment in which they must work

(Shapiro, 2005). I'm inclined to think that myself. The claim is rather that for many mechanistic explanations, we don't care about how the components are composed: the fuel rail is just something that manages to get fuel from the pump to the injectors. The details aren't that important, at least when we're explaining hard starts.

2.3 Resources

Not all explanations focus on mechanical parts. My engine stops running. Why? Because the car ran out of gasoline. Without gasoline, the engine doesn't go. That seems, intuitively, like a mechanistic explanation—certainly, it's on a par with the explanation of my engine's failure where I cited a bad check valve. (Many troubleshooting procedures for engine failure begin by checking whether you ran out of gasoline.) Furthermore, gasoline seems to pass the constitutive relevance condition: by changing how much gas there is in the car I can change whether the engine goes, and running the engine changes how much gas there is.

Yet gasoline is very much unlike the rest of the engine. Gasoline is transformed by the action of the engine. The engine persists through those transformations. Gasoline thus gets *used* by engines, while engines do the using.

Gasoline is one example of a *resource*. There are many resources that show up in explanations, across a variety of fields. The grass in the paddocks and water in the streams are resources for the livestock. The catalytic enzyme argininosuccinate synthase is the rate-limiting step in the synthesis of arginine. Some resources are more abstract. City planners must consider tax base, water supplies, and volume of low-cost rental units. Each are resources that a city needs to run efficiently, but each is spread out over time and space. Computational complexity affects the amount of memory and processing time that an algorithm needs to run, and is in turn affected by the number of processors available to implement the algorithm. The performance of an internet connection depends on resources like bandwidth and downstream cache. Cognitive scientists appeal to a variety of resources like working memory capacity, attention, and willpower.

Each of these in turn appear in resource explanations. We explain characteristic patterns of psychological performance and deficit by showing how psychological resources can be blocked, competed for, and otherwise made unavailable. Insufficient housing leads to negative health outcomes. Having too many browser tabs open at once uses up memory and slows down the rest of the computer.

Examples are easily multiplied. I have already presented a diverse bunch. What do they have in common? I will suggest a cluster of features, each of which complements those of mechanical parts.

First, tokens of resources need not persist across the timespan of explanations in which they feature. Cows eat grass. Engines burn gasoline. These activities irreversibly transform token chunks of resource, and that transformation is necessary for the proper functioning of the mechanism.

More broadly, resources can be made *available* or *unavailable* for transformation. The availability of resources is often a key feature in explanations. Working memory has a certain capacity, and holding items in working memory diminishes that capacity. Streaming a movie makes it harder to download system updates, because each takes a certain amount of bandwidth.

Second, and along the same lines, token chunks of resources are *not individually important.* That is, it doesn't (typically) matter which bit of resource gets used, so much as there's enough of it to go around. The car factory doesn't care which ingot of steel it uses; it doesn't (usually) matter. Similarly, it does not matter (for most applications) whether you use the first sector on the hard drive or the ten thousandth.

Mechanical parts are individually important precisely because the tokens do such different things. Resources, by contrast, admit of equipotent divisions. Every bit of gas is basically as good as every other. These divisions might be effectively continuous (as with water supplies), or they might come naturally chunked (as with working memory). Each mechanistic part of the computer, by contrast, does something individually important—you can't interchange the hard drive for the memory. This distinction is roughly reflected in a linguistic division. Resources are typically picked out using mass nouns. In order to talk about some collection of them, we need to discretize them using a measure word that specifies the relevant divisions (three *cups* of water, twelve *kilobytes* of memory). mechanical parts, by contrast, are usually referred to by count nouns, and so can take plurals without further divisions specified.

Third, most resources are *causally promiscuous.* (I suspect that all are potentially so.) That is, resources can interact with many different components of a complex system, and many different components have an interest in using that resource. All of the grass in the paddock is available to any of the cows that want it. Even relatively restricted resources—like, say, modality-specific cognitive buffers—are available to any process that might use that resource type.

Causal promiscuity is important because many resource explanations involve competition for a limited amount of resource. Conversely, many complex mechanisms have portions devoted to controlling access to resources. Multitasking operating systems spend a great deal of time mediating conflicts over access to memory and processor time. Even when resources aren't actually promiscuous, it is often because of deliberate strategies taken by mechanisms to *control* access, rather than the intrinsic properties of the resources themselves.

Fourth, we typically *do* care about the realization of resources, and otherwise similar resources can't be substituted willy-nilly. Car engines need *gasoline*; it would be nice if something else worked, but nothing works quite as well. Even small deviations can be problematic—recall the issues caused by ethanol-augmented gasoline . Food chains are vulnerable to collapse precisely because organisms need specific resources. Ion channels might work in a variety of ways, but it's important that they let in only one specific sort of ion. Unlike mechanical parts, then, we often *do* care what a resource is made of.

Even substitutions that maintain functional similarity in a broad sense can cause a variety of problems. Ethanol-augmented gasoline in your old car, or a substandard batch of steel in your factory, may not shut things down—but that substitution is what you'll have to cite when you want to explain why things break, fail, and otherwise fall short of expectations.

3 Resource Explanations

Resources appear to be the primary *explanans* in a variety of mechanistic explanations. That is, there are explanations that account for the activity of the whole in more or less the same way that canonical mechanistic explanations do, but which appeal to resources rather than (or at least in addition to) mechanical parts.

Most obviously, many explanations *must* cite resources if they're to be complete. You don't know how a car factory works unless you know about the materials it uses and the machines that do the using. An explanation of how a heart works is obviously incomplete without mentioning the blood that gets pumped around. In many such explanations, resources might not play the most *interesting* causal-explanatory role, but they certainly play an *ineliminable* one.

Further, many explanations give resources pride of place. Some of these invoke various ways in which resources can be depleted, blocked, competed for, or otherwise made unavailable. Availability of resources thus shapes the behavior of the whole. Sometimes this relationship is crude: no gas, no driving. But the relationship can be more complex. Forming long-term memories via long-term potentiation uses glucose, which is largely provided by astrocytic glycogen. Depleting glycogen experimentally leads to amnesia. The elderly often have memory problems, and also tend to have lower levels of glucose. Supplementing them with glucose improves their memory (Messier, 2004; Newman et al, 2011).

Crucially, the relationship between glucose and LTP is not a simple on-off sort of relationship: the relationship is a specific one, in the sense specified by Woodward (2010). That is, there are many distinct possible levels of glucose, and a (roughly) one-one mapping between those levels and the efficiency of LTP within a certain range.

Computational complexity theory similarly posits specific relationships between an algorithm's properties and the amount of time, space, and so forth it can use. Consider: many low-level assembly language algorithms are tuned to avoid accessing RAM, which can be orders of magnitudes slower than accessing in-CPU cache (Duntemann, 2011). The performance of instances of the algorithms can be complex, but largely explained by the differences in access time between two distinct types of memory resource. Along the same lines, some enzymes are the rate-limiting step in reactions, and so the norm of reaction within the mechanism is largely explained by reference to the bulk availability of that enzyme.

Appeal to resources can also be evidentially important, because they serve to invalidate (or at least cast suspicion upon) certain types of mechanistic explanation. Tim Shallice (1988) relies on considerations of resources in his argument that single dissociations offer only weak evidence for the distinctness of two cognitive processes. Suppose I have a lesioned patient who can spell regularly-spelled English words and regular nonwords, but not irregular words. I might conclude that there are two routes to spelling: a preserved phonetic one that involves spelling-to-sound correspondences and a damaged lexical one that retrieves stored information about spelling. But as Shallice points out, there is another possible explanation: spelling irregular words might just be *harder*. That is, spelling might have only one route, but that route might require a general resource that is stressed more by harder tasks than by easier ones—intelligence or working memory or attention, say. Partial damage to this resource will only affect more difficult tasks. Ruling out these *resource artifacts* is the primary reason why double dissociations are so prized in neuropsychology, and why the combination of lexical and phonological agraphia is necessary to establish a dual-route model of spelling (Caramazza, 1986; Shallice, 1988).

Further, many of the techniques for distinguishing and exploring resources involve the same methods that picking out mechanical parts do. Just as it is not always obvious which spatiotemporal parts count as mechanical parts, it is also not obvious which resources a mechanism actually uses. The very same techniques for exploring mutual manipulability that Craver details in his (2007) can be applied to resources. So for example, to determine whether something is a resource you can manipulate it, block access to it, change the whole to measure the amount of the resource, and so on. So it seems that scientists do very similar things to investigate both mechanical parts and resources. The two thus deserve explanatory parity.

Finally, differences between resources can often explain differences between mechanisms. Sometimes these involve different resources as inputs. Different steelmaking processes, for example, vary importantly in what resources they use (e.g. forced air in the Bessemer process vs pure oxygen in the contemporary basic oxygen process). Differences in resources often require different mechanisms: diesel engines don't have spark plugs and have heavier combustion chambers, because diesel must be detonated rather than merely ignited. Explanations of why the two types of mechanism differ thus appeals to the different resources those mechanisms are meant to use, even if they are functionally similar at the level of the whole.

At a more abstract level, resources themselves can vary in a variety of ways. As an incomplete taxonomy: resources can be limited or unlimited, they can admit of discrete or continuous divisions, they can be permanently transformed or merely blocked, usable in parallel or in serial, and so on. These distinctions make a difference to the systems that use them. Computers with random-access memory have different properties to those with older serial-access memory. Omnivores have a different ecological profile than carnivores. Breakbulk shipping has a different timescale and efficiency curve to container shipping. The differences between these complex mechanisms is explained by the differences in the resources they use. Resource explanations are thus important when we *compare* mechanisms as well as when we try to account for the behavior of single mechanisms.

In sum, resources constitute their own class of *explanans*, distinct from mechanical parts, and sometimes form the primary or even the sole explanatory variable for the behavior of mechanisms. They are thus worthy of focus in their own right.

4 Mechanistic Components and Background Conditions

4.1 The distinction

The preceding two sections focused on the differences between mechanical parts and resources. Yet they also have a fair bit in common, enough that both can function in mechanistic explanations. I suggest further that resources have some features in common with ordinary background conditions. Overlooking resources as a distinct category has been one of the reasons why it has been difficult to give a precise delineation between mechanistic components and background conditions. Properly distinguished, however, resources play an *explanatory* role like that of mechanical parts.

I summarize the features of each of the three categories in table 1.

	Mechanical Parts	**Resources**	**Background Conditions**
Causally...	Conservative	Promiscuous	Promiscuous
Realization	Irrelevant	Relevant	Relevant
Individual?	Yes	No	No
Actual Variation?	Yes	Yes	No
Specfic?	Yes	Yes	No
Systematic?	Yes	Yes	No
Mutual Manipulability?	Yes	Yes	No

Table 1: Contrasting components and resources

The entries in table 1 ought to be seen as claims about the typical or usual features of the respective kinds—as I'll show in section 4.2, the boundaries between the classes are (acceptably) vague. The line down the center of the table indicates where resources are grouped with either mechanical parts or background conditions.

On the one hand, there are several respects in which resources are similar to background conditions. Both tend to be causally promiscuous: engine oil lubricates and cools a wide variety of components, the sun rises on the just and the unjust alike, and so on. Both facilitate mechanisms in realization-specific ways: oxygen is a background condition necessary for the struck match to light, and pretty much only oxygen will do. And both tend to be important in bulk but not individually.

Yet there are a variety of ways in which resources are more like mechanical parts. For one, both mechanical parts and resources can vary in their actual properties and activities over the timescale of explanations. This makes them, in Kenneth Waters' (2007) term, *actual* difference-makers. By contrast, background conditions enable a process by remaining constant over the timescale of the explanation. Similarly, as Woodward (2010) notes, causes might have more or less *specific* relationships to their effects. When X is a relatively non-specific cause of Y, then "we are more likely to regard X as a mere enabling (or background) condition for Y" (Woodward, 2010, 317). Conversely, Campbell (2007; 2010) notes that explanatorily interesting variables let one make well-defined changes to a target phenomena, a relationship that Klein (2017) dubs "systematicity." Background conditions don't seem to have systematic effects either. The classic background condition makes everything break down in unsystematic and uninteresting ways when it fails.

Finally, true background conditions don't typically pass the *mutual* manipulability test. I might be able to affect my engine's functioning by changing the local gravitational field, but I can't change gravity by running the engine (Craver, 2007). By contrast, resources and mechanisms do pass the mutual manipulability test.

Thus while resources have much in common with background conditions, it is the cluster of properties they have in common with mechanical parts that are most relevant for *explanation*. With that in mind, return to the problem case with which I began. Recall that astrocytic glycogen is a crucial energy source for long-term memory formation. Forming memories depletes astrocytic glycogen, and depriving astrocytes of glycogen causes amnesia. So astrocytic glycogen satisfies Craver's mutual manipulability criterion. Further, there is a specific and systematic relationship between glycogen and performance on memory tasks (Newman et al, 2011): a little depletion causes a little deficit, more causes more, and so on. So astrocytic glycogen passes these tests as well. Finally, astrocytic glycogen actually varies in populations we care about, and that actual variation is part of the explanation of actual differences between mechanisms in that population. Of course, like most resources, glycogen is available to many processes (and sub-processes within memory formation itself), is relization-specific, and is important only in bulk. That explains why it seems like a mere background condition, and why the line was difficult to draw. Indeed, I suspect that some of the difficulty in coming up with an adequate characterization of background conditions (at least for mechanistic explanations) stems from the fact that mere background conditions often resemble resources. All the more reason, then, for philosophers to distinguish the two.

4.2 Vague and relative boundaries

You might be concerned that the boundaries between mechanical part, resource, and background condition are not as clear-cut as one might like. For one, the boundaries seem *vague*. Engine oil degrades, but only very slowly; is it a background condition or a resource? Processing cores on a chip seem like mechanical parts until you have a lot of them, and then they start to seem a bit more like a resource. For another, the boundaries are *relative*. RNA is a mechanism for transcription, a cellular resource for viral load, and a background condition for mitosis. Jeeps are complex mechanisms to the motor pool, matériel to the general. And so on.

Neither vagueness nor relativity are problematic so long as they are well-behaved. Both are plausibly context-dependent. So long as *within* a given context there is both a lack of relativity and clear cases on either side of a vague boundary, then the distinction between mechanical part, resource, and background condition is on no worse footing than most philosophical concepts. Indeed, it is a commonplace in discussions of causation that contextual factors can determine what counts as background condition and what counts as cause (Menzies, 2007). Further, as noted above, mechanistic components are explicitly defined relative to the explanation of a particular phenomenon. The trick, as ever, is to find principled ways to delineate contexts.

As with all mechanistic accounts, the primary tool for delineating contexts is appeal to the interests of people in explanatory contexts. There is no great mystery here. The criteria in table 1 themselves admit of a certain amount of interest-relativity. The butcher doesn't care which kidney is which—so long as they're fresh—and sells them by the pound. The anatomist makes finer distinctions, cares more about functional connections, and studies them in place and individually. There is also a straightforward relativity to the phenomenon being explained. Engine oil is a resource *relative to cooling*: its cooling capacity varies with engine heat, and can be overwhelmed. Conversely, it is a background condition *relative to lubrication*: it is perpetually present in normal conditions, and its absence causes widespread catastrophe.

Similarly, the criteria in table 1 also allow for a certain amount of vagueness. The mechanic can treat engine oil as a background condition on timescales of days and weeks. But engine oil slowly degrades, and so at longer timescales behaves more like a resource. The distinction between a short and a long timescale is a vague one, but not problematically so—once a day is definitely too frequently to change the oil, once a decade definitely too long. Similarly, differences in spatial scale might lead to vague boundaries: the rolling prairie is a resource to be divided for the farmer and a stable background condition for the ecologist. Each are concerned with a fundamentally different spatial grain, which gives rise to different explanations (Bechtel and Mundale, 1999).

Similar vagueness might arise from different levels of abstraction. RAM is a mechanism for a chip designer, a resource for a programmer, and a background condition for an accountant. The gaps hide intermediate cases: a low-level assembly-language programmer might need to know at least some details about physical memory and its implementation, and thereby treat RAM sometimes as a resource, sometimes as a mechanism.

In sum, nothing counts as a mechanical part, or resource, or background condition *simpliciter*. Mechanical components are already defined relative to a decomposition for the mechanist; the present account shows how one might further accommodate various sorts of vagueness and relativity.

5 Extended Mechanistic Explanation

So much for the defense of resource explanations. The distinction itself is of interest for those of us who care about the nitty-gritty of scientific explanation. I think, however, there is a broader consequence for the mechanist project as a whole.

Philosophers of science used to talk about different sciences as corresponding to different *levels* of nature (Oppenheim and Putnam, 1958). Physics covered the smallest bits. Chemistry was a bit "above" that, biology further still, and so on. On this *global* picture of levels, there is some level-making relation R that gives a partial ordering on the entities of the world. It does so in a way that creates a set of equivalence classes of things that are not R-related to one another. Each equivalence class corresponds, more or less, to a distinct science. Questions about intertheoretic relationships between sciences can equally well be recast as questions about metaphysical relationships between levels: whether one science reduces to another depends on the details of R.

Yet actual sciences don't seem to respect these distinctions. That is, any simple R (size, composition, complexity) doesn't seem to give an ordering that lines up neatly with the natural sciences (Wimsatt, 1976; Kim, 2002; Craver, 2007; Klein, 2014). Physics talks about atoms and about planets; planets are large, slow, materially complex, and made up of atoms.

In response, mechanists have advocated a move to *local* levels. While mechanistic explanation requires decomposition, that decomposition is always made relative to a particular system behaving in certain ways. According to Craver, "levels of mechanisms are not monolithic divisions in the structure of the world" (2007, 190); hence "it makes no sense to ask if my heart is at a different level of mechanisms than my car's water pump because there is no mechanism containing the two" (2007, 191). Similarly, Bechtel notes that "[l]ocal identification of levels is sufficient for understanding levels in a mechanism and for capturing how mechanistic explanation is reductionist" (2008, 147).[1]

Local decomposition arguably does fit better with how individual scientists investigate a particular phenomenon. Yet the presence of resource explanations shows that local decomposition cannot be the whole story. For two distinct phenomena can interfere with one another just in case they each utilize the same limited resource. The irrigation system and the hydroelectric generator both draw from the same river; more cotton means more brownouts, and more power means less cotton. The mechanistic explanation of how the dam works includes a resource—the flow of water in the river—that is the same token resource that appears in mechanistic explanations of how the irrigation system works. Similarly so when taking dictation and reading aloud compete for the

[1]See also Craver and Darden (2013, 88) and Bechtel and Richardson (2010, xix). A local sense of level arguably predates the mechanists. Marr's (1982) famous tripartite division of levels is drawn relative to a particular explanatory task; the lower levels are investigated in the context of sorting out the details of the computational level. Similarly, homuncular functionalism is concerned with functional decomposition relative to a starting capacity (Lycan, 1990).

same orthographic representations (Botvinick and Cohen, 2014).

There is an explanation about how the irrigation system works, and an explanation about how the dam works. I assume that the mechanist doesn't want to insist that these are insufficient unless we conjoin them into an unwieldy whole (which would thereby bring the river under the purview of a single mechanistic decomposition). But if that's the case, then strictly local levels won't do. The appropriate constraints on a mechanistic decomposition are not *just* considerations like mutual manipulability. In many cases, we choose parts of a system because they are a resource that is part of some *other* system as well, and competition is key.

Indeed, I think that the mechanists implicitly rely on such a principle even when delineating mechanical parts. Why do I chose the *thermostat* as the appropriate part, rather than (say) an entity consisting of all of the thermostat save for the top few layers of atoms? Why do I pick the wire *including* the insulation, when it is only the copper core that carries the electricity relevant to the explanation? In general, there are many functionally equivalent spatiotemporally carvings of objects (Haugeland, 1998), and we don't usually pick the minimal one.

The reason, I suggest, is because we try to pick mechanical parts in such a way as to give them a certain generality: thermostats and wires show up in a variety of different explanations, and that is why it is useful to pick *them* rather than something minimal. Resource explanations make the point more strongly, because the same token resource is often involved in several distinct phenomena, and resource explanations rely on this overlap. Indeed, once we take this position, we see that background conditions are simply a limiting case of the same phenomenon: the same things typically appear as background conditions in a great number of explanations, though they are actual difference-makers in none of them.

Ultimately, mechanistic explanation as a *practice* tries to draw from the same store of parts that appear in other explanations. Indeed (I suggest), this is precisely what distinguishes mechanistic explanation when it is available, rather than an unadorned causal/interventionist story.

References

Bechtel W (2008) Mental mechanisms: Philosophical perspectives on cognitive neuroscience. Taylor & Francis, New York

Bechtel W, Mundale J (1999) Multiple realizability revisited: Linking cognitive and neural states. Philosophy of Science 66:175–207

Bechtel W, Richardson R (2010) Discovering complexity: Decomposition and localization as strategies in scientific research. MIT Press, Cambridge

Botvinick MM, Cohen JD (2014) The computational and neural basis of cognitive control: Charted territory and new frontiers. Cognitive science 38(6):1249–1285

Campbell J (2007) An interventionist approach to causation in psychology. In: Gopnik A, Schulz L (eds) Causal Learning: Psychology, Philosophy and Computation, Oxford: Oxford University Press, pp 58–66

Campbell J (2010) II—Control variables and mental causation. In: Proceedings of the Aristotelian Society, vol 110, pp 15–30

Caramazza A (1986) On drawing inferences about the structure of normal cognitive systems from the analysis of patterns of impaired performance: The case for single-patient studies. Brain and Cognition 5(1):41–66

Craver C (2007) Explaining the brain. Oxford University Press, New York

Craver CF, Darden L (2013) In search of mechanisms: Discoveries across the life sciences. University of Chicago Press, Chicago

Cummins R (1975) Functional analysis. Journal of Philosophy LXXII(20):741–765

Cummins R (1983) The nature of psychological explanation. MIT Press, Cambridge

Duntemann J (2011) Assembly language step-by-step: Programming with Linux. John Wiley & Sons, New York

Haugeland J (1998) Mind embodied and embedded. In: Having Thought: Essays in the Metaphysics of Mind, Harvard University Press, Cambridge, pp 207–240

Kim J (2002) The layered model: Metaphysical considerations. Philosophical Explorations 5(1):2–20

Klein C (2014) Psychological explanation, ontological commitment, and the semantic view of theories. In: Sprevak M, Kallestrup J (eds) New Waves in Philosophy of Mind, Palgrave Macmillan, New York, pp 208–225

Klein C (2017) Brain regions as difference-makers. Philosophical Psychology 30((1–2)):1–20

Lycan W (1990) The continuity of levels of nature. Mind and Cognition: A Reader pp 77–96

Marr D (1982) Vision: A computational investigation into the human representation and processing of visual information. WH Freeman, New York

Menzies P (2007) Causation in context. In: Corry R, Price H (eds) Causation, physics, and the constitution of reality: Russell's Republic Revisited, Oxford University Press, Oxford, pp 191–223

Messier C (2004) Glucose improvement of memory: A review. European journal of pharmacology 490(1):33–57

Newman LA, Korol DL, Gold PE (2011) Lactate produced by glycogenolysis in astrocytes regulates memory processing. PloS One 6(12):1–11

Oppenheim P, Putnam H (1958) Unity of science as a working hypothesis. Minnesota Studies in the Philosophy of Science 2:3–36

Shallice T (1988) From Neuropsychology to Mental Structure. Cambridge University Press, Cambridge

Shapiro L (2005) The mind incarnate. MIT Press, Cambridge

Simon HA (1996) The Sciences of the Artificial, 3rd edn. MIT Press, Cambridge

Waters CK (2007) Causes that make a difference. The Journal of Philosophy 104(11):551–579

Wimsatt WC (1976) Reductionism, levels of organization, and the mind-body problem. In: Globus G, Savodnik I, Maxwell G (eds) Consciousness and the Brain: A Philosophical Investigation, Plenum Press, New York, pp 205–267

Woodward J (2010) Causation in biology: Stability, specificity, and the choice of levels of explanation. Biology & Philosophy 25(3):287–318