CETERIS PARIBUS LAWS AND *MINUTIS RECTIS* LAWS

LUKE FENTON-GLYNN

DEPARTMENT OF PHILOSOPHY, UNIVERSITY COLLEGE LONDON

GOWER STREET, LONDON, WC1E 6BT, U.K.

Abstract. Special science generalizations admit of exceptions. Among the class of non-exceptionless special science generalizations, I distinguish (what I will call) *minutis rectis* (*mr*) generalizations from the more familiar category of *ceteris paribus* (*cp*) generalizations. I argue that the challenges involved in showing that *mr* generalizations can play the law role are underappreciated, and quite different from those involved in showing that *cp* generalizations can do so. I outline some potential strategies for meeting the challenges posed by *mr* generalizations.

1. Introduction

Many philosophers of science speak as though all non-exceptionless scientific generalizations that (appear to) play at least some aspects of the law role (counterfactual support, inductive confirmation, predictive/explanatory import) tolerably well can be classed as *ceteris paribus (cp) laws.* The following are representative quotations:

> "A nonstrict law is a generalization that contains a *ceteris paribus* qualifier that specifies that the law holds under 'normal or ideal conditions,' [...]. The generalizations one finds in the special sciences are mostly of this kind. In contrast, a strict law is one that contains no *ceteris paribus* qualifiers; it is exceptionless not just *de facto* but as a matter of law." (Lepore and Loewer 1987, 632)

> "cp lawfulness is just a species of nomological necessity, the other species of nomological necessity being strict lawfulness. What distinguishes the two

species is just that cp laws can have ... exceptions and strict laws can't" (Fodor 1991, 31–32)

"Special science laws ... are usually taken to 'have exceptions', to be 'non-universal' or 'to be *ceteris paribus* laws'." (Reutlinger et al. 2011)

In each of these quotations, the notion of a non-exceptionless law is run together with that of a cp law. It is easy to find further confirmation that this running-together is widespread (see, e.g., Schurz 2002, 351, Schrenk 2007, 221, Woodward 2002, 303–304).

The identification of non-exceptionless 'laws' with cp laws can hardly be a matter of stipulative definition. The notion of a cp law has a richness that significantly outstrips the bare notion of *a law that admits of exceptions*. For one thing, the notion of a cp law is associated with a distinctive account of how exceptions arise.

A cp law is supposed to be endowed with an implicit or explicit clause that specifies that it holds 'other things being equal', where this latter notion is usually parsed in terms of the obtaining of 'normal' or even 'ideal' conditions (see Cartwright 1983, 46, Schurz 2002), and explicated in terms of the absence of significant difference-making interference from outside the system that the law in question seeks to characterize (see, e.g., Fodor 1989, 69n, Schurz 2002, 366-370). Exceptions are taken to arise due to the non-satisfaction of this cp clause.[1]

After reviewing the notion of a cp law (Section 2), I will argue (Section 3) that it is a mistake to equate non-exceptionless laws with cp laws: there is a distinct type of non-exceptionless law – which I will call a *minutis rectis (mr) law* – which admits of exceptions that aren't explained by the non-satisfaction of a cp clause. I will argue (Section 4) that mr laws pose a distinctive set of philosophical challenges. Finally (Section 5), I will examine some potential responses to these challenges.

[1] I'm skating over some differences between the various characterizations of cp laws that appear in the literature (for an overview, see Reutlinger et al. 2011). Some (e.g. Schurz 2002) distinguish different *types* of cp law. For present purposes, it suffices to note that the notion of a *minutis rectis* law that I will distinguish below is not a type of (or variant on the notion of a) cp law.

A terminological point: talk of cp (and mr) 'laws' is rather awkward in the context of a discussion of whether, and to what extent, the exception-ridden generalizations of the special sciences play various aspects of the law-role. I'm sympathetic to the objections that some philosophers (e.g. Woodward 2005, Woodward and Hitchcock 2003a,b) have to such law-speak. Nevertheless, because law-speak is so common in the literature, I shall not try to forgo it in what follows, and I shall drop the jarring scare-quotes when I use it.

2. Ceteris Paribus Laws

In ecology, one standard equation used for predicting population growth is the Logistic Equation (LE):

$$\text{(LE)} \qquad \frac{dn}{dt} = r_c n\left(1 - \frac{n}{K}\right)$$

Here n is the number of individuals in the population, $\frac{dn}{dt}$ is the growth rate of the population (the change in n, with respect to time t), r_c is the *intrinsic per capita growth rate* of the population (the growth rate that obtains in the absence of intra-species competition for resources), K is the *carrying capacity* (the maximum sustainable population size).

LE implies that when the population n of a species in a particular habitat is very small (so that there is little intra-specific competition for resources), the actual population growth rate $\frac{dn}{dt}$ is close to the intrinsic per capita growth rate r_c multiplied by the number n of individuals. But, as the population grows, the actual growth rate declines linearly (due to increasing competition). This decline continues until the carrying capacity K is reached, at which point population growth is 0.

It is an open question whether ecological generalizations – such as LE – should be called 'laws' at all. But they do appear to play certain aspects of the law role to at least some degree. Ecologists apply LE to certain populations (especially populations that aren't subject to significant *inter*-species competition or predation) in order to make predictions, and to give explanations.[2]

LE holds only *ceteris paribus* because there are possible background conditions under which it is violated (even when applied to populations concerning which, in normal circumstances, it is predictively accurate). For example, it will not hold in the event of the population being subject to a cull, or in the event of a natural disaster that destroys (a large part of) the population. While LE may give accurate predictions about population growth *after* some such events, it won't accurately predict growth *during* such episodes. It simply doesn't include variables that represent such events.

Culls, etc., produce circumstances in which other things are *not* equal: interfering factors are present, so LE doesn't even approximately hold. An ecologist presumably wouldn't seek to model such factors, since they are not *ecological* factors. They interfere with the sorts of system that the ecologist seeks to model (viz. *ecosystems*), but come from 'outside' such systems. Perhaps this means that ecological generalizations will in-principle remain cp generalizations (compare Davidson 1970, 94, Fodor 1989, 69n).

There have been several attempts (e.g. Lepore and Loewer 1987, Fodor 1989, 1991, Woodward and Hitchcock 2003a,b, Woodward 2005) to show that generalizations like LE, despite holding only cp, can support counterfactuals, and sustain predictions and causal-explanatory relationships and thus play the law role to a non-negligible degree.

[2]Tsoularis and Wallace (2002) survey some successful applications of LE in ecology. Ecologists sometimes appeal to more complex equations than LE. The following discussion also applies to these more complex equations. In general, ecologists have an armory of equations (or systems of equations – i.e. models) for predicting population growth and other phenomena. Different models are more or less predictively successful with respect to different populations. The fact that such equations (or models) apply only to some populations – and even then only approximately – may disincline you to call them 'laws'. I'm sympathetic. (Though note that there is a nuanced literature in (philosophy of) ecology about whether there are genuine ecological laws: see, e.g., Colyvan and Ginzburg 2003, Lawton 1999, Turchin 2001.) And since this story is repeated throughout the special sciences, you may be disinclined to admit the existence of special science laws at all (except, perhaps, in a few special cases). Again, I'm sympathetic. To reiterate: the question with which I'm concerned is not whether such generalizations deserve to be called 'laws', but whether and to what extent they are able to do things like predict, explain, and support counterfactuals. As we'll see, LE is the sort of thing that Woodward and Hitchcock (2003a,b) and Woodward (2005) call an 'invariant generalization'. I have no objection to their alternative terminology.

Woodward and Hitchcock (2003a,b) argue that generalizations like LE support causal-explanatory relations because they are *invariant under a range of hypothetical interventions*.[3] For example, if we were to intervene upon the intrinsic growth rate r_c of the population (e.g. by genetic engineering to increase fertility), upon the carrying capacity K (by improving or depleting the environment), or upon the population size n (by carrying out a cull), then the actual growth rate, $\frac{dn}{dt}$, – *after the intervention episode* – would accord with LE.

The reason that LE 'supports' these interventionist counterfactuals is that, in evaluating them, we are considering the 'closest worlds' in which such interventions occur (see Hitchcock 2001, 283; compare Lewis 1979, Woodward 2005). In these worlds significant interfering factors like natural disasters don't occur.

In virtue of the fact that LE supports these interventionist counterfactuals (when it comes to the populations that it models well) it follows directly, on the account of Woodward and Hitchcock (2003a,b), that the variables on the RHS of LE causally explain the actual growth rate of the population $\frac{dn}{dt}$. So, on their account, the cp nature of LE doesn't stand in the way of its playing important aspects of the law role.[4]

3. Minutis Rectis Laws

Not all exceptions to scientific generalizations arise due to the non-fulfilment of (explicit or implicit) cp clauses. This is best illustrated w.r.t. a law that admits of exceptions, but that plausibly is not a cp law, viz. the *Second Law of Thermodynamics* (SLT), which states that the total entropy of an isolated system increases over time, until equilibrium is reached, after which it doesn't decrease.

SLT admits of possible exceptions. Given an initial non-equilibrium state of an isolated system, it is *possible* (though very 'unlikely') that the micro-state should be one that leads to a later state that is further from equilibrium. An example of SLT-violation, which is nevertheless possible (i.e. consistent with the fundamental dynamical laws), is an isolated

[3]Woodward (2005) gives a precise definition of the technical notion of an 'intervention'. For present purposes, it will suffice to think of interventions as ideal experimental manipulations of variables.

[4]The same is true on the accounts given by Lepore and Loewer (1987) and Fodor (1989, 1991), though I focus on Woodward and Hitchcock's account here.

system comprising an ice cube in hot water, in which the ice cube grows larger and colder, while the surrounding water becomes hotter.

Such exceptions to SLT *do not* arise due to failures of a cp condition to hold. SLT is not aptly construed as a cp law. Rather than a cp clause, SLT includes a precise specification of its scope of application: it applies to thermodynamically isolated systems (including the universe as a whole). Unlike LE, there's no possibility of interference from outside the systems that SLT characterizes.

Perhaps the claim that SLT is not a cp law can be disputed. Someone might, for instance, attempt to construe its appeal to an ideal isolated system as somehow amounting to a cp clause (compare Schurz 2002, 369–370). I don't need to insist that it's not a cp law. What I *do* wish to insist is that there is a type of possible exception to it that has nothing to do with the violation of any cp clause. That is, there is a class of exception that is not due to the failure of its idealizations to hold. *Even assuming an ideal isolated system*, exceptions to SLT may arise just as a consequence of certain unlikely microphysical realizations of the system's initial thermodynamic state.[5]

Laws that admit of this type of exception are what I am calling '*minutis rectis (mr)*' laws: that is, laws that hold only when the properties that they concern are realized in the right way. SLT holds only *minutis rectis* because the macro-states that it concerns are multiply realizable by points in the underlying phase space. In a non-equilibrium system, the majority of points in that space (measure ≈ 1) are on entropy-increasing trajectories. However, there are a very few (measure ≈ 0) that are on entropy-decreasing trajectories. SLT only holds if the initial macro-state of an isolated system is realized 'in the right way' – viz. by one of the 'usual' points in phase space that is on a non-entropy-decreasing trajectory.

Though I have illustrated the distinction between the notion of a cp law and that of an mr law w.r.t. a law that's an mr law but plausibly *isn't* a cp law – namely SLT – many special science generalizations hold *both* only cp *and* only mr. Such laws admit of exceptions

[5]It would be inapt to construe SLT as including an implicit cp condition that supposes away such microphysical realizations. That would be to construe SLT's implicit form as something like 'the total entropy of an isolated system is non-decreasing over time, except when the initial microstate is such that it is decreasing'. But this comes close to rendering SLT empty when clearly it isn't (compare Earman and Roberts 1999, 465).

even when their cp clauses are satisfied. Even when there is no disruptive interference from outside the systems that such generalizations characterize (so that their cp conditions are satisfied), they may still be violated just as a consequence of the properties that they concern being realized in the 'wrong' way.

LE is an example of a cp generalization that also holds only mr. I have already argued that it holds only cp. Rather trivially it also holds only mr. LE will break down if members of a population to which it normally applies start *en masse* to exhibit SLT-violating behavior: for example, if neurotransmitters suddenly stop diffusing across their synapses, or oxygen stops diffusing in their blood streams. In such a case, the growth rate of the population will not be predicted by LE. Not for nothing does Lawton (1999, 178) say that SLT is one of the "three deep universal laws that underpin all ecological systems"!

There may also be more interesting reasons why LE holds only mr. For example, the *geographical distribution* of a population can make a difference to its actual growth rate. Indeed, given that population growth can be extremely sensitive to precise initial conditions (see, e.g., May 1974), even very small perturbations of the precise, individual-by-individual initial geographical distribution of members of a population can potentially make a difference to whether the population grows according to LE or sharply declines (even where the population is well below the carrying capacity). The latter situation – in which the population is initially precisely distributed in one of those rare ways that leads to dramatically LE-violating behavior – would be analogous to a thermodynamic system's being at one of those rare points in phase space that leads to SLT-violating behavior. A population's having a certain size, n, is multiply realizable by different precise individual-by-individual geographical distributions. Only if the geographical distribution is 'right' will LE approximately hold.

4. Why It Matters

The distinction between cp and mr generalizations matters because the mr nature of a generalization poses problems for its ability to support counterfactuals and causal-explanatory relations in a way that its cp nature does not.

Consider Woodward and Hitchcock's claim that generalizations like LE are invariant (i.e., support counterfactuals about what would happen) under interventions. The argument that this is so rests upon the idea that the closest worlds in which we intervene upon (say) the population size are not worlds in which the cp condition is violated: in such worlds there is (e.g.) no natural disaster that wipes out the population immediately after the intervention. So, post-intervention, the growth rate is modeled by LE. Given Woodward's notion of an intervention (Woodward 2005, 98) and Lewis's suggested similarity measure over possible worlds (Lewis 1979, 472), this all seems plausible.

Yet the mr nature of LE appears to undercut its ability to support interventionist counterfactuals. Even concerning a population that is usually well-modeled by LE, it seems extremely doubtful that it is true that 'If the size of the population had been intervened upon, the post-intervention micro-state wouldn't have been one that leads to entropy-decreasing behavior'. After all, it seems that the post-intervention micro-state just *might* have been one of those rare entropy-decreasing ones.

It is also doubtful that, even where a population size is well below the carrying capacity, it is true that 'If the size of the population had been intervened upon, the resulting precise geographical distribution would not have been such as to lead to a severe decline in the population'. After all, it's not possible (even metaphysically speaking) to intervene on the size of the population without impacting on the precise individual-by-individual distribution (fewer or more individuals can't be distributed in the same individual-by-individual way) and, in light of the dramatic effects that slight changes in initial conditions can have on ecosystems, the post-intervention distribution just *might* have been one of those rare ones that leads to a dramatic decline in numbers.[6]

It is very doubtful that the truth of either of the counterfactuals considered in the previous two paragraphs follows from the Woodwardian notion of an intervention or the Lewisian notion of similarity among possible worlds. But if such counterfactuals aren't true, then it

[6]If we build into the antecedent of the counterfactual a specification of exactly *how* the intervention would occur (and what the resulting precise geographical distribution would be), then we might get a true counterfactual. But this is not the sort of interventionist counterfactual to which Woodward and Hitchcock appeal in their account of causal explanation.

appears that we can't reason that, if the population size had been intervened upon, then the growth rate would have subsequently followed LE.

Likewise with SLT. Consider the counterfactual 'If I had placed this ice cube into that glass of hot water, then it would have melted quickly'. SLT's mr nature appears to undercut its ability to support this counterfactual. We can't (it seems) say that if the ice cube had been placed in the hot water, then the resulting system *would not* have been in one of those rare micro-states that fails to lead to melting. The post-intervention system *might* have been in such a micro-state, and this undercuts the assertion that the ice cube would have melted. The Lewisian notion of similarity doesn't appear to make a world in which the specified post-intervention macro-state is realized by a non-entropy-increasing micro-state *more dissimilar* to the actual world than one in which it is realized by an entropy-increasing micro-state (compare Hájek (ms)).

If mr laws aren't able to sustain such counterfactuals about what would happen under interventions (i.e. if they're not invariant generalizations), then this threatens to undermine their ability to underwrite causal-explanatory relations and their predictive power. This indicates that philosophical vindications – such as Woodward and Hitchcock's – of the causal-explanatory and predictive power of cp laws are not *ipso facto* vindications of mr laws.

5. Potential Solutions

There is a range of approaches that one might take in attempting to address the problems posed by mr laws.

First, one might consider modifying the Lewisian similarity metric so that worlds in which (e.g.) I intervene on a thermodynamic system and the post-intervention system conforms to SLT come out closer than those in which the post-intervention system does not so conform. We might similarly take conformity to special science laws, like LE, to make for similarity to the actual world.

For example, in response to a worry raised by Elga (2001) about whether Lewis's similarity metric delivers an asymmetry of counterfactual dependence, Dunn (2011) suggests modifying

Lewis's metric so that, other things being equal, worlds obeying SLT and also the various special science laws come out closer to the actual world than those that don't. Such a proposal would seem to ensure that counterfactuals like 'If I had placed the ice cube in the hot water, then it would have melted quickly' come out true, so that SLT supports the counterfactuals needed to underwrite causal/explanatory and predictive relations after all.

One concern about this approach is that it appears to force upon us the truth of counterfactuals like 'If I had put the ice cube in the hot water, then the resulting system wouldn't have been in one of the rare entropy-decreasing microstates'. This counterfactual seems less plausible. But perhaps there is some room for maneuver: perhaps, for instance, one could maintain that the assertion of the latter counterfactual results in a context shift and a corresponding change in the standards of similarity (compare Lewis 1979), with the consequence that this second counterfactual utterance asserts a false proposition (while, in the original, ordinary context, the first asserted a true one). I shan't explore the prospects for such a response here.

A *second* option might be to modify the Woodwardian notion of an intervention so that (e.g.) manipulations of a population size that result in the population being geographically distributed 'in the wrong way', don't count as 'interventions' in the relevant, technical sense. A worry about this strategy is that it is not clear that the 'wrong' sort of interventions could be specified in a systematic and non-ad-hoc way. Simply specifying the relevant 'interventions' in terms of precisely those counterfactual outcomes that one wants to avoid (as in 'the "intervention" on population size must not be such as to result in a precise geographical distribution that leads to a violation of LE') is ad hoc and unsystematic. Similarly, in the thermodynamic case, one might wonder whether there is a useful notion of 'intervention' such that manipulations of a system's macro-state that happen to result in its being in a micro-state on an entropy-decreasing trajectory fail to count as interventions. I shan't explore this strategy further here.

A *third* option would be to argue that 'deterministic' mr laws are mere approximations to probabilistic laws. For example, it is tempting to say that, while SLT is a mr law, it is an approximation to a probabilistic law that is not a mr law. Statistical mechanics (SM) furnishes us with an exceptionless, *probabilistic* version of SLT.

In SM probabilities are generated by applying a uniform probability distribution (on the Lebesgue measure) to the region of phase space associated with the initial macro-state of an isolated system. Since the measure of points in this phase-space on non-entropy-decreasing trajectories is extremely high, and the measure of points on entropy-decreasing trajectories is extremely low, the result of applying the uniform distribution is an *overwhelmingly high probability* that entropy does not decrease over time (compare Albert 2000, Loewer 2001, Frigg and Hoefer 2014).

Although the probabilistic version of SLT implied by SM *does not* support counterfactuals like 'If I had placed the ice cube in the glass of hot water, then it would have melted quickly', it does support counterfactuals like 'If I had placed the ice cube in the hot water, then *the probability* that it would have melted quickly *would have been very much higher* than if (say) I had returned the ice cube to the freezer'. This appears to be precisely the sort of counterfactual that is relevant to *probabilistic* prediction, causation, and explanation.

I'm sympathetic to this proposal. It is worth noting, however, that this line of response involves construing the probabilities of SM *objectively*. Otherwise, it's hard to see how they could underwrite objective relations of probabilistic causation and explanation. Yet the view that the probabilities of SM are objective has become popular (see, e.g., Albert 2000, 2012, Frigg and Hoefer 2014, Loewer 2001, 2012).

Perhaps similar reasoning can be applied to other high-level laws – like LE – that appear to hold only mr. It might be argued that they too are merely approximations to probabilistic laws that *don't* hold only mr. Rather ambitiously, Albert (2000, 2012) and Loewer (2001, 2008, 2012) argue that SM itself actually entails probabilistic approximations of the laws of the special sciences. While this 'Statistical Mechanical Imperialism' has been criticized (Callender 2011, Frisch 2014, Weslake 2014), one might nevertheless think that there *are* probabilistic approximations to special science mr laws (perhaps derivable in some other way). If so, then these probabilistic laws may be able to support the counterfactuals relevant to probabilistic causal explanation and prediction.

For instance, in ecology the geographical distribution of populations is often modeled via a probability distribution (e.g. the Poisson distribution) over a habitat (see, e.g., Vandermeer and Goldberg 2013, 126-142). Perhaps, in general, we can get strict(er) probabilistic versions of special science generalizations via the imposition of probability distributions over the underlying state-spaces in which the properties that they concern are realized.

A lot of work needs to be done to show that this will work out. It's reasonable to wonder whether we can *always* replace deterministic special science generalizations that hold only mr with strict(er) probabilistic laws by imposing probability distributions over underlying state spaces. It's also reasonable to wonder whether we can *always* interpret the resulting probabilities objectively. That these are serious questions for both science and metaphysics shows the depth of the challenges posed by the mr nature of many special science generalizations.

6. Conclusion

The notion of a non-exceptionless law shouldn't be equated with that of a cp law. There is another important category of non-exceptionless law that ought to be distinguished, viz. mr laws. The mr nature of special science generalizations poses distinctive challenges for those aiming to show that special science generalizations can support counterfactuals, causal explanations, and predictions. Distinguishing the two categories of non-exceptionless law brings these challenges to light, but also allows us to identify possible avenues for addressing them.

Acknowledgments

Thanks to audiences at the Universities of Cologne, Wuppertal, Munich, and Cambridge, and at the 2013 BSPS Conference. I acknowledge the support of the Deutsche Forschungsgemeinschaft (Grant: SP279/15-1), the McDonnell Causal Learning Collaborative, and the Alexander von Humboldt Foundation.

References

Albert, D. (2000). *Time and chance*. Cambridge, MA: Harvard University Press.

Albert, D. (2012). Physics and chance. In Y. Ben-Menahem and M. Hemmo (Eds.), *Probability in Physics*, pp. 17–40. Berlin: Springer.

Callender, C. (2011). The past histories of molecules. In C. Beisbart and S. Hartmann (Eds.), *Probabilities in Physics*, pp. 83–113. New York: OUP.

Cartwright, N. (1983). *How the Laws of Physics Lie*. Oxford: OUP.

Colyvan, M. and L. Ginzburg (2003). Laws of nature and laws of ecology. *Oikos 101*, 649–653.

Davidson, D. (1970). Mental events. In L. Foster and J. Swanson (Eds.), *Experience and Theory*, pp. 79–101. Amherst: University of Massachusetts Press.

Dunn, J. (2011). Fried eggs, thermodynamics, and the special sciences. *The British Journal for the Philosophy of Science 62*, 71–98.

Earman, J. and J. Roberts (1999). *Ceteris paribus*, there is no problem of provisos. *Synthese 118*, 439–478.

Elga, A. (2001). Statistical mechanics and the asymmetry of counterfactual dependence. *Philosophy of Science 68*, 313–324.

Fodor, J. (1989). Making mind matter more. *Philosophical Topics 17*, 59–79.

Fodor, J. (1991). You can fool some of the people all of the time, everything else being equal: Hedged laws and psychological explanations. *Mind 100*, 19–34.

Frigg, R. and C. Hoefer (2014). The Best Humean System for statistical mechanics. Forthcoming in *Erkenntnis*.

Frisch, M. (2014). Why physics can't explain everything. In A. Wilson (Ed.), *Asymmetries of Chance and Time*. Oxford: OUP.

Hájek, A. (ms.). Most counterfactuals are false.

Hitchcock, C. (2001). The intransitivity of causation revealed in equations and graphs. *Journal of Philosophy 98*, 273–299.

Lawton, J. (1999). Are there general laws in ecology? *Oikos 84*, 177–192.

Lepore, E. and B. Loewer (1987). Mind matters. *Journal of Philosophy 84*, 630–642.

Lewis, D. (1979). Counterfactual dependence and time's arrow. *Noûs 13*, 455–476.

Loewer, B. (2001). Determinism and chance. *Studies in History and Philosophy of Science Part B 32*, 609–620.

Loewer, B. (2008). Why there is anything except physics. In J. Hohwy and J. Kallestrup (Eds.), *Being Reduced*, pp. 149–163. Oxford: OUP.

Loewer, B. (2012). The emergence of time's arrows and special science laws from physics. *Interface focus 2*, 13–19.

May, R. (1974). Biological populations with nonoverlapping generations: Stable points, stable cycles, and chaos. *Science 186*, 645–647.

Reutlinger, A., G. Schurz, and A. Hütemann (2011). Ceteris paribus laws. *The Stanford Encyclopedia of Philosophy*. URL: $\langle http://plato.stanford.edu/entries/ceteris-paribus/\rangle$.

Schrenk, M. (2007). Can capacities rescue us from ceteris paribus laws? In M. Kistler and B. Gnassounou (Eds.), *Dispositions in Philosophy and Science*, pp. 221–247. Aldershot: Ashgate.

Schurz, G. (2002). Ceteris paribus laws: Classification and deconstruction. *Erkenntnis 57*, 351–372.

Tsoularis, A. and J. Wallace (2002). Analysis of logistic growth models. *Mathematical Biosciences 179*, 21–55.

Turchin, P. (2001). Does population ecology have general laws? *Oikos 94*, 17–26.

Vandermeer, J. and D. Goldberg (2013). *Population Ecology* (2^{nd} ed.). Princeton, NJ: Princeton University Press.

Weslake, B. (2014). Statistical mechanical imperialism. In A. Wilson (Ed.), *Asymmetries of Chance and Time*. Oxford: OUP.

Woodward, J. (2002). There is no such thing as a *ceteris paribus* law. *Erkenntnis 57*, 303–328.

Woodward, J. (2005). *Making Things Happen*. New York: OUP.

Woodward, J. and C. Hitchcock (2003a). Explanatory generalizations, Part I: A counterfactual account. *Noûs 37*, 1–24.

Woodward, J. and C. Hitchcock (2003b). Explanatory generalizations, Part II: Plumbing explanatory depth. *Noûs 37*, 181–199.