Mixed-Effects Modeling and Nonreductive Explanation

Wei Fang*†

This essay considers a mixed-effects modeling practice and its implications for the philosophical debate surrounding reductive explanation. Mixed-effects modeling is a species of the multilevel modeling practice, where a single model incorporates simultaneously two (or even more) levels of explanatory variables to explain a phenomenon of interest. I argue that this practice makes the position of explanatory reductionism held by many philosophers untenable because it violates two central tenets of explanatory reductionism: single-level preference and lower-level obsession.

1. Introduction. Explanatory reductionism is the position that holds that, given a relatively higher-level phenomenon (or state, event, process, etc.), it can be reductively explained by a relatively lower-level feature (Kaiser 2015, 97; see also Sarkar 1998; Weber 2005; Rosenberg 2006; Waters 2008).[1] Although philosophers tend to have slightly different conceptions of the position, two central tenets can still be extracted:[2]

*To contact the author, please write to: Research Center for Philosophy of Science and Technology, Shanxi University, 92 Wucheng Road, Taiyuan, Shanxi 030006, China; e-mail: wesleyfang@outlook.com.

†I am thankful to Xiang Huang, Chuang Liu, and Wei Wang for comments on earlier drafts of this article. This research was supported by the Ministry of Education of China's Humanities and Social Sciences Youth Fund.

1. According to Sarkar (1998), explanatory reduction is an epistemological thesis that is distinguished from constitutive (ontological) and theory reductionism theses. Kaiser further distinguishes two subtypes of explanatory reduction: (*a*) "a relation between a higher-level explanation and a lower-level explanation of the same phenomenon" (2015, 97) and (*b*) individual explanations; i.e., given a relatively higher-level phenomenon, it can be reductively explained by a relatively lower-level feature (97). This article will focus on the second subtype. When referring to levels, I mean either hierarchical organizations such as universities, faculties, departments, etc., or functional organizations such as organs, tissues, cells, etc. When referring to scales, I mean spatial or temporal scaling where levels are not so clearly delimited.
2. A similar summary of the position can be found in Sober (1999).

Philosophy of Science, 86 (December 2019) pp. 882–894. 0031-8248/2019/8605-0006$10.00
Copyright 2019 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/705449 Published online by Cambridge University Press

Single-level preference: A phenomenon of interest can be fully explained by invoking features that reside at a single, well-defined level of analysis (e.g., molecular level in biology).

Lower-level obsession: Lower-level features always provide the most significant and detailed explanation of the phenomenon in question, so a lower-level explanation is always better than a higher-level explanation.

Philosophers sometimes express these two tenets explicitly in their work. For example, Rosenberg holds that "there is a full and complete explanation of every biological fact, state, event, process, trend, or generalization, and that this explanation will cite only the interaction of macromolecules to provide this explanation" (2006, 12). Weber expresses a similar idea in his explanatory hegemony thesis, according to which it is always some lower-level physicochemical laws (or principles) that ultimately do the explanatory work in experimental biology (2005, 18–50). Bickle (2003) attempts to motivate a 'ruthless' reduction of psychological phenomena (e.g., memory) to the molecular level.

However, many philosophers have questioned the plausibility of the position on the basis of scientific practice (Hull 1972; Craver 2007; Bechtel 2010; Brigandt 2010; Hüttemann and Love 2011; Kaiser 2015). To counter that position, some authors have pointed to the relevance of an important practice that has not received sufficient attention before: multiscale or multilevel modeling, or what is sometimes called an integrative modeling approach, where a set of distinct models ranging over multiple levels or scales—including the macrophenomenon level/scale—are involved in explaining a (often complex) phenomenon of interest (Mitchell 2003, 2009; Craver 2007; Brigandt 2010, 2013a, 2013b; Knuuttila 2011; Batterman 2013; Green 2013; O'Malley et al. 2014; Green and Batterman 2017). Often these models work together by providing diverse constraints on the potential space of representation (Knuuttila 2011; Green 2013).

This multilevel modeling surely casts some doubt on explanatory reductionism, for it seems unclear what reductively explains what—all those facts in the set of models ranging over different levels/scales are involved in doing some explanatory work. However, there is a species of multilevel modeling that has slipped away from most philosophers' sights:[3] mixed-effects modeling (MEM; also called multilevel regression modeling, hierarchical linear modeling, etc.), wherein a single model incorporating simultaneously two (or even more) levels of variables is used to explain a phenomenon. For a mixed-effects model to explain, features of the so-called reducing and reduced

3. Bourrat (2016) is an exception.

https://doi.org/10.1086/705449 Published online by Cambridge University Press

levels must be simultaneously incorporated into the model; namely, they must go hand in hand.

MEM deserves special attention because it sheds new light on the reductionism-antireductionism debate by showing that (*a*) a mixed-effects model violating the two central tenets of explanatory reductionism can provide successful explanation and that (*b*) a single mixed-effects model without integrating with other epistemic means can also provide such successful explanation. Therefore, MEM first further challenges the explanatory reductionist position and second offers a novel perspective bolstering the multilevel/multiscale integrative approach discussed by many philosophers.

The article proceeds as follows. Section 2 discusses the challenges faced by the traditional single-level modeling approach and examines the reasons why the MEM approach is preferable in dealing with these challenges. Section 3 describes a MEM practice using a concrete model. Section 4 elaborates on the implications of MEM for the explanatory reductionism debate. Finally, section 5 considers potential objections to my viewpoint.

2. Challenges to Reductive Explanatory Strategies. In many fields (e.g., biological, social, and behavioral sciences) scientists find that the data collected show an intrinsically hierarchical or nested feature. Consider a simple example: we might be interested in examining relationships between students' achievement at school (A) and the time they invest in studying (T).[4] In conducting such research, we might collect data from different classes (say five classes in total), with each class providing the same number of samples (say 10 students in each class). The data collected among classes might be assumed to be independent. Then we may use certain traditional statistical techniques such as ordinary least squares (OLS) to analyze the data and build a linear relationship between A and T.

However, this single-level reductive analysis can lead to misleading results because it ignores the possibility that students within a class may be more similar to one another in important aspects than students from different classes. In other words, each group (class) may have its own features relevant to the relationship between A and T that the other groups lack. Hence, the data collected from the students are in fact not independent; that is, the subjects are not randomly sampled because the individuals (students) are clustered within groups (classes). In technical terms, we say our analysis may fall prey to the *atomistic fallacy*, where we base our analysis solely on the individual level (i.e., we reduce all the group-level features to the individuals). Therefore, traditional OLS techniques such as multiple regression cannot

4. For scientific studies of this kind, see Schagen (1990), Wang and Hsieh (2012), and Maxwell et al. (2017).

https://doi.org/10.1086/705449 Published online by Cambridge University Press

be employed in this context because the case under consideration violates a fundamental assumption of these techniques: the independence of observations (Nezlek 2008, 843).

Conversely, we may face the same problem the other way around if we fail to consider the inherently nested nature of the data. Consider the student-achievement-at-school case again. We may observe that in classes where the time of study invested by students is very high, the achievements of the students are also very high. Given such an observation, we may reason that students who invest a lot of time in studying would be more likely to achieve more at school. However, this inference commits the *ecological fallacy* because it attributes the relationship observed at the group level to the individual level (Freedman 1999). The individuals may exhibit within-group differences that the single group-level analysis fails to capture. In technical terms, this inference is flawed because it reduces the variability in achievement at the individual level to a group-level variable, and the subsequent analysis is solely based on the group's mean achievement results (Heck and Thomas 2015, 3). Again, traditional statistical techniques such as multiple regression cannot be employed in this context.

In sum, a single-level modeling approach that disrespects the multilevel data structure can commit either an atomistic or an ecological fallacy. Confronted with these problems, one response is to 'tailor' the traditional statistical techniques by, for example, adding an effect variable to the model that indicates the grouping of the individuals. However, many have argued that this approach is unpromising because it may give rise to enormous new problems (Luke 2004; Nezlek 2008; Heck and Thomas 2015). Alternatively, scientists have developed a new framework that takes the multilevel data structure into full consideration, that is, the MEM approach, to which we now turn.

3. Case Study: A Mixed-Effects Model. Depending on different conceptual and methodological roots, we have two broad categories of MEM approaches: the multilevel regression approach and the structural equation modeling approach. The former usually focuses on direct effects of predictor variables on (typically) a single dependent variable, while the latter usually involves latent variables defined by observed indicators (for details, see Heck and Thomas 2015). For the purpose of this article's arguments, I concentrate on the first kind.

Consider the student-achievement-at-school example again. Since students are typically clustered in different classes, a student's achievement at school may be both influenced by her own features (e.g., time invested in studying) and her class's features (e.g., size of the class). Hence, we have two levels of analysis: the individual level (level 1) and the group level

https://doi.org/10.1086/705449 Published online by Cambridge University Press

(level 2), and individuals ($i = 1, 2, \ldots, N$) are clustered in level 2 groups ($j = 1, 2, \ldots, n$).[5] Now suppose that students' achievements at school are represented as scores they get on an exam. The effect of time invested in studying on scores can be described as follows:

$$Y_{ij} = \beta_{0j} + \beta_{1j}X_{ij} + \varepsilon_{ij}, \tag{1}$$

where Y_{ij} refers to the score of individual i in the jth group, β_{0j} is a level 1 intercept representing the mean of scores for the jth group, β_{1j} is a level 1 slope (i.e., different effects of study time on scores) for the predictor variable X_{ij}, and the residual component (i.e., an error term) ε_{ij} is the deviation of individual i's score from the level 2 mean in the jth group. Equation (1) looks like a multiple regression model; however, the subscript j reveals that there is a group level incorporated in the model. It can also be seen from this equation that both the intercept β_{0j} and slope β_{1j} can vary across the level 2 units; that is, different groups can have different intercepts and slopes.

The most remarkable aspect of MEM is that we treat both the intercept and slope at level 1 as dependent variables (i.e., outcomes) of level 2 predictor variables. So here we write the following equations expressing the relationships between the level 1 parameters and level 2 predictors:

$$\beta_{0j} = \gamma_{00} + \gamma_{01}W_j + u_{0j} \tag{2}$$

and

$$\beta_{1j} = \gamma_{10} + \gamma_{11}W_j + u_{1j}, \tag{3}$$

where "β_{0j} is the level-1 intercept in level-2 unit j; γ_{00} is the mean value of the level-1 dependent variable, controlling for the level-2 predictor, W_j; γ_{01} is the effect (slope) of the level-2 predictor, W_j; and u_{0j} is the error, or unmodeled variability, for unit j" (Luke 2004, 10). Also, "β_{1j} is the level-1 slope in level-2 unit j; γ_{10} is the mean value of the level-1 slope, controlling for the level-2 predictor, W_j; γ_{11} is the effect of the level-2 predictor, W_j; and u_{1j} is the error for unit j" (10).

Equations (2) and (3) have specific meanings and purposes. They express "how each of the level-1 parameters are functions of level-2 predictors and variability" (Luke 2004, 10). They aim to explain variations in the randomly varying intercepts or slopes by adding one (or more) group-level predictor to the model. These expressions are based on the idea that the group-level characteristics such as group size may affect the strength of the within-group effect of study time on scores. This kind of effect is called a *cross-level interaction*, for

5. Note that, for instructive purposes, our case involves only two levels; however, the MEM approach can in principle be extended to many more levels.

https://doi.org/10.1086/705449 Published online by Cambridge University Press

it involves the impact of variables at one level of a data hierarchy on relationships at another level. We discuss this in detail in the next section.

Now we combine equations (1), (2), and (3) by substituting the level 2 parts of the model into the level 1 equation. We finally obtain the following equation:

$$Y_{ij} = (\gamma_{00} + \gamma_{10}X_{ij} + \gamma_{01}W_j + \gamma_{11}X_{ij}W_j) + (u_{1j}X_{ij} + u_{0j} + \varepsilon_{ij}). \quad (4)$$

This equation can be simply understood as that Y_{ij} is made up of two parts: the fixed-effect part expressed by the first four terms and the random-effect part expressed by the last three terms. Note that the term $\gamma_{11}X_{ij}W_j$ denotes a cross-level interaction between level 1 and level 2 variables, which is defined as the impact of a level 2 variable on the relationship between a level 1 predictor and the outcome Y_{ij}. We have seven parameters to estimate in (4); they are four fixed effects (intercept, within-group predictor, between-group predictor, and cross-level interaction) and two random effects (the randomly varying intercept and slope, and a level 1 residual).

Now a mixed-effects model has been built, and the next step is to estimate the parameters of the model. However, we will skip this step and turn to explore the philosophical implications of the modeling practice relevant to the explanatory reductionism debate.

4. Implications for the Explanatory Reductionism Debate. Looking closely into the MEM practice, we find that a couple of important philosophical implications for the explanatory reductionism debate can be drawn.

4.1. All Levels Are Indispensable. The first, and most obvious, feature of MEM is that it routinely involves many levels of analysis in a single model, and all these levels are indispensable to the model in the sense that no level can be reduced to or replaced by the other levels. These levels consist of both the so-called reducing level in the reductionist's terminology (typically a lower level that attempts to reduce another level) and the reduced level (typically a higher level to be reduced by the reducing level). In our student-achievement-at-school case, for example, a reductionist may state that the group level will be regarded as the reduced level, whereas the student level is the reducing level.

The indispensability of each level in the model can be understood in two related ways. First, because of the nested nature of data, only when we incorporate different levels of analyses to the model can we avoid either the atomistic or ecological fallacy discussed in section 2. As discussed in the student-achievement-at-school example where students are clustered in different classes (in the manner that students from the same class may be more similar to one another in important aspects than students from different classes), reducing

https://doi.org/10.1086/705449 Published online by Cambridge University Press

all the analyses to the level of individual students can simply miss the important information associated with group-level features and thus lead to misleading results. Although it is true that the problem might be partially mitigated by tailoring traditional single-level analytical techniques such as multiple regression, it is also true that this somewhat ad hoc maneuver can simply bring about various new vexing and recalcitrant issues (Luke 2004; Nezlek 2008; Heck and Thomas 2015).

Second, the problem can also be viewed from the perspective of identifying explanatory variables. In building a mixed-effects model, the main consideration is often to find a couple of variables that may play the role of explaining the pattern or phenomenon observed in the data. Here a modeler must be clear about how to assign explanatory variables; for instance, she must consider whether there are different levels of analyses and, if so, which explanatory variables should be assigned to what levels, and so on. These considerations may come before her model building because of background knowledge, which paves the way for her to develop a conceptual framework for investigating the problem of interest. However, without such a clear and rigorous consideration of identifying and assigning multilevel explanatory variables, an analysis can fail simply because it confounds variables at different levels.

Respecting the multilevel nature of explanatory variables has another advantage: "Through examining the variation in outcomes that exists at different levels of the data hierarchy, we can develop more refined theories about how explanatory variables at each level contribute to variation in outcomes" (Heck and Thomas 2015, 33). In other words, in respecting the multilevel nature of explanatory variables, we get a clear idea of how, and to what degree, explanatory variables at different levels contribute to variation in outcomes. If these variables do contribute to variation in outcomes, as it always happens in MEM, then the situation suggests an image of *explanatory indispensability*: all the explanatory variables at different levels are indispensable to explaining the pattern or phenomenon of interest.

Given these considerations, therefore, one implication for the explanatory reductionism debate becomes clear: it is not always the case that, given a relatively higher-level phenomenon it can be reductively explained by a relatively lower-level feature. Rather, in cases in which the data show a nested structure or, put differently, the phenomenon suggests multilevel explanatory variables, we routinely combine the higher level with the lower level in a single (explanatory) model. As a result, one fundamental tenet of explanatory reductionism is violated: single-level preference.

4.2. Interactions between Levels. Another crucial feature of multilevel modeling is its emphasis on a *cross-level interaction*, which is defined as "the potential effects variables at one level of a data hierarchy have on

https://doi.org/10.1086/705449 Published online by Cambridge University Press

relationships at another level. . . . Hence, the presence of a cross-level interaction implies that the magnitude of a relationship observed within groups is dependent on contextual or organizational features defined by higher-level units" (Heck and Thomas 2015, 42–43). Remember that there is a term $\gamma_{11}X_{ij}W_j$ in our mixed-effects model discussed in section 3, which indicates the cross-level interaction between the group level and the individual level. More specifically, this term can be best construed as the impact of a group-level variable (e.g., group size) on the individual-level relationship between a predictor (e.g., study time) and the outcome (e.g., students' scores).

The cross-level interaction points to the plain fact that an organization or a system can somehow influence its members or components by constraining how they behave within the organization or system. This does not necessarily imply top-down causation (sec. 5.2 will return to this point). Within the context of scientific explanation, however, it does imply that it is not simply that characteristics at different levels separately contribute to variation in outcomes but rather that they interact in producing variation in outcomes. In other words, the pattern or phenomenon to be explained can be understood as generated by the interaction between explanatory variables at different levels. Therefore, to properly explain the phenomenon of interest, we need have not only a clear idea of how to assign explanatory variables to different levels but also an unequivocal conception of whether these explanatory variables may interact.

Different models can be built depending on different considerations of the cross-level interaction. To see this, consider the student-achievement-at-school example again. In some experimental setting we may assume that there was no cross-level interaction between group-level characteristics and the individual-level relationship (between study time and scores). In such a situation, we kept the effect of individual study time on scores the same across different classes; that is, we kept the slope constant across classes. In the meanwhile, we treated another group-level variable (i.e., intercept) as varying across classes; that is, different classes have different average scores. So, this is a case in which we have a clear idea of how to assign explanatory variables but no consideration of the cross-level interaction. Nonetheless, in a different experimental setting we may assume that there existed cross-level interaction, and hence the effect of individual study time on scores can no longer be kept constant across different classes. At the same time, we treated another group-level variable (i.e., intercept) as varying across classes. Hence, this is a case in which we have both a clear idea of how to assign explanatory variables and a consideration of the cross-level interaction. Corresponding to these two different scenarios, two different mixed-effects models can be built, as shown in figure 1.

Given such a cross-level interaction, therefore, the explanatory reductionist position has been further challenged. This is because any reductive

https://doi.org/10.1086/705449 Published online by Cambridge University Press

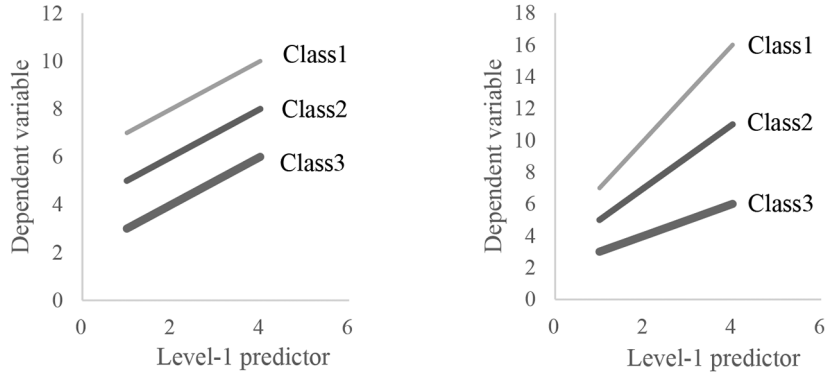


Figure 1. Two different models showing varying intercepts (*left*) or varying intercepts and slopes (*right*). Lines represent three classes. Adapted from Luke (2004, 12). Color version available as an online enhancement.

explanation that privileges one level of analysis—usually the lower level—over the others falls short of capturing this kind of interaction between levels. If they fail to do so, then they are missing important terms relevant to explaining the phenomenon of interest. As a consequence, a mixed-effects model involving interactions between levels simultaneously violates the two fundamental pillars of explanatory reductionism: first, it violates single-level preference because it involves multilevel explanatory variables in explaining phenomena, and second, it violates lower-level obsession because it privileges no levels—all levels are interactively engaged in producing outcomes.

5. Two Potential Objections

5.1. In-Principle Argument. One argument that resurfaces all the time in the reductionism-versus-antireductionism debate is the in-principle argument, the core of which is that even if reductive explanations in a field of study are not available for the time being, it does not follow that we will not obtain them someday (e.g., Sober 1999; Rosenberg 2006). Therefore, according to some reductionists, the gap between current science and future science is simply a matter of time, for advancement in techniques, experimentation, and data collecting can surely fill in the gap.

However, I think the argument fails. To begin with, advancement in techniques, experimentation, and data collecting is not always followed by reductive explanations. For example, in our MEM discussed in section 3, even if the data about the individual level are available and sufficiently detailed, it is not the case that we explain the phenomenon of interest in terms of the data from the individual level alone. Consider another example: in dealing with problems associated with complex systems in systems biology, even

https://doi.org/10.1086/705449 Published online by Cambridge University Press

though large-scale experimentation (e.g., via computational simulation) can be conducted and high-throughput data arranging over multiple scales/levels can be collected, a bottom-up reductive approach must be integrated with a top-down perspective so as to produce useful explanations or predictions (Green 2013; Green and Batterman 2017; Gross and Green 2017).

Nevertheless, reductionists may reply that the situations presented above only constitute an in-practice impediment, for it does not undermine the possibility that lower-level reductive explanations, typically provided by some form of 'final science', will be available someday. The possibility here may be construed as a *logical possibility* (Green and Batterman 2017, 21; see also Batterman 2018). Nonetheless, if it is merely logically possible that there will be some final science providing only reductive explanations, then nothing can exclude another logical possibility that there will be some 'mixed science' providing only multilevel explanations. After all, since both reductionism and nonreductionism are logically possible, it is doubtful that logical possibility alone could offer anything insightful in helping us understand how science proceeds. As Batterman puts, "Appeals to the possibility of *in principle* derivations rarely, if ever, come with even the slightest suggestion about how the derivations are supposed to go" (2018, 869).

Another interpretation of possibility may be associated with real possibilities, referring to the actual cases of reductive explanations happening in science. Unfortunately, I do not think the real scenario in science speaks for the reductionist under this interpretation. Although it is impossible to calculate the absolute cases of nonreductive explanations occurring in science, a cursory look at scientific practice can tell that a large portion of scientific explanations proceeds in a nonreductive fashion, as suggested by multilevel modeling (Batterman 2013; Green 2013; O'Malley et al. 2014; Green and Batterman 2017; Mitchell and Gronenborn 2017). Moreover, even in areas such as physics, which was regarded as a paradigm for the reductionist stance, progressive explanatory reduction does not always happen (Green and Batterman 2017; Batterman 2018). In sum, we have shown that the in-principle argument fails, for it neither offers help in understanding how science proceeds if it is construed as implying a logical possibility nor is in tune with scientific practice if it is construed as implying real possibilities.

5.2. Top-Down Causation. In section 3 we showed that there is a cross-level interaction taking the form that higher-level features may affect lower-level features. A worry arises: Does this imply top-down causation?

My answer to this question is twofold. First, it is clear that this short article is not aimed to engage in the philosophical debate about whether, and in what sense, there exists top-down causation (see Craver and Bechtel 2007; Kaiser 2015; Bechtel 2017). Second, what we can do now is to show that the cross-level interaction is a clear and well-defined concept in multilevel

https://doi.org/10.1086/705449 Published online by Cambridge University Press

modeling. It unambiguously demonstrates the constraints on the lower-level processes exerted by the higher-level parameters (Green and Batterman 2017). In our multilevel modeling discussed in section 3, we showed that group-level features may affect some individual-level features through the way that each group possesses its own feature relevant to explaining the differences at the individual level across groups. This idea is incorporated into the mixed-effects model by assigning some explanatory variables to the group level and a cross-level interaction term to the model.

The idea of cross-level interaction as a constraint is widely accepted in multilevel modeling broadly construed, where constraint is usually expressed in the form of initial or boundary conditions. For example, in modeling cardiac rhythms, because of "the influences of initial and boundary conditions on the solutions of the differential equations used to represent the lower level process" (Noble 2012, 55; cf. Green and Batterman 2017, 32), a model cannot simply narrowly focus on the level of proteins and DNA but must also consider the levels of cell and tissue working as constraints. The same story happens in cancer research, where scientists are advocating the idea that tumor development can be better understood if we consider the varying constraints exerted by tissue (Nelson and Bissel 2006; Shawky and Davidson 2015; cf. Green and Batterman 2017, 32).

6. Conclusion. This article has shown that nonreductive explanations involving many levels predominate in areas where the systems under consideration exhibit a hierarchical structure. These explanations violate the fundamental pillars of explanatory reductionism: single-level preference and lower-level obsession. Traditional single-level reductive approaches fall short of capturing systems of this kind because they face the challenges of committing either the atomistic or ecological fallacy.

REFERENCES

Batterman, Robert. 2013. "The 'Tyranny of Scales.'" In *The Oxford Handbook of Philosophy of Physics*, ed. Robert Batterman, 255–86. Oxford: Oxford University Press.

———. 2018. "Autonomy of Theories: An Explanatory Problem." *Noûs* 52 (4): 858–73.

Bechtel, William. 2010. "The Downs and Ups of Mechanistic Research: Circadian Rhythm Research as an Exemplar." *Erkenntnis* 73:313–28.

———. 2017. "Explicating Top-Down Causation Using Networks and Dynamics." *Philosophy of Science* 84:253–74.

Bickle, John. 2003. *Philosophy and Neuroscience: A Ruthlessly Reductive Account*. Dordrecht: Kluwer.

Bourrat, Pierrick. 2016. "Generalizing Contextual Analysis." *Acta Biotheoretica* 64 (2): 197–217.

Brigandt, Ingo. 2010. "Beyond Reductionism and Pluralism: Toward an Epistemology of Explanatory Integration in Biology." *Erkenntnis* 73 (3): 295–311.

———. 2013a. "Explanation in Biology: Reduction, Pluralism, and Explanatory Aims." *Science and Education* 22:69–91.

https://doi.org/10.1086/705449 Published online by Cambridge University Press

———. 2013b. "Integration in Biology: Philosophical Perspectives on the Dynamics of Interdisciplinarity." *Studies in History and Philosophy of Biological and Biomedical Sciences* 44:461–65.
Craver, Carl. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Oxford University Press.
Craver, Carl, and William Bechtel. 2007. "Top-Down Causation without Top-Down Causes." *Biology and Philosophy* 22:547–63.
Freedman, David. 1999. "Ecological Inference and the Ecological Fallacy." In *International Encyclopedia of the Social and Behavioral Sciences*, vol. 6, ed. Neil Smelser and Paul Baltes, 4027–30. New York: Elsevier.
Green, Sara. 2013. "When One Model Isn't Enough: Combining Epistemic Tools in Systems Biology." *Studies in History and Philosophy of Biological and Biomedical Sciences* 44:170–80.
Green, Sara, and Robert Batterman. 2017. "Biology Meets Physics: Reductionism and Multi-Scale Modeling of Morphogenesis." *Studies in History and Philosophy of Biological and Biomedical Sciences* 61:20–34.
Gross, Fridolin, and Sara Green. 2017. "The Sum of the Parts: Large-Scale Modeling in Systems Biology." *Philosophy, Theory, and Practice in Biology* 9 (10).
Heck, Ronald, and Scott Thomas. 2015. *An Introduction to Multilevel Modeling Techniques*. 3rd ed. New York: Routledge.
Hull, David. 1972. "Reductionism in Genetics: Biology or Philosophy?" *Philosophy of Science* 39 (4): 491–99.
Hüttemann, Andreas, and Alan Love. 2011. "Aspects of Reductive Explanation in Biological Science: Intrinsicality, Fundamentality, and Temporality." *British Journal for the Philosophy of Science* 62 (3): 519–49.
Kaiser, Marie. 2015. *Reductive Explanation in the Biological Sciences*. Dordrecht: Springer.
Knuuttila, Tarja. 2011. "Modelling and Representing: An Artefactual Approach to Model-Based Representation." *Studies in History and Philosophy of Science* A 42:262–71.
Luke, Douglas. 2004. *Multilevel Modeling*. London: Sage.
Maxwell, Sophie, Katherine Reynolds, Eunro Lee, Emina Subasic, and David Bromhead. 2017. "The Impact of School Climate and School Identification on Academic Achievement: Multilevel Modeling with Student and Teacher Data." *Frontiers in Psychology* 8:2069.
Mitchell, Sandra. 2003. *Biological Complexity and Integrative Pluralism*. Cambridge: Cambridge University Press.
———. 2009. *Unsimple Truths: Science, Complexity, and Policy*. Chicago: University of Chicago Press.
Mitchell, Sandra, and Angela Gronenborn. 2017. "After Fifty Years, Why Are Protein X-Ray Crystallographers Still in Business?" *British Journal for the Philosophy of Science* 68 (3): 703–23.
Nelson, Celeste M., and Mina J. Bissell. 2006. "Of Extracellular Matrix, Scaffolds, and Signaling: Tissue Architecture Regulates Development, Homeostasis, and Cancer." *Annual Review of Cell and Developmental Biology* 22:287–309.
Nezlek, John. 2008. "An Introduction to Multilevel Modeling for Social and Personality Psychology." *Social and Personality Psychology Compass* 2 (2): 842–60.
Noble, Daniel. 2012. "A Theory of Biological Relativity: No Privileged Level of Causation." *Interface Focus* 2 (1): 55–64.
O'Malley, Maureen A., Ingo Brigandt, Alan C. Love, John W. Crawford, Jack A. Gilbert, Rob Knight, Sandra D. Mitchell, and Forest Rohwer. 2014. "Multilevel Research Strategies and Biological Systems." *Philosophy of Science* 81:811–28.
Rosenberg, Alex. 2006. *Darwinian Reductionism; or, How to Stop Worrying and Love Molecular Biology*. Chicago: University of Chicago Press.
Sarkar, Sahotra. 1998. *Genetics and Reductionism*. Cambridge: Cambridge University Press.
Schagen, I. P. 1990. "Analysis of the Effects of School Variables Using Multilevel Models." *Educational Studies* 16:61–73.
Shawky, Joseph, and Lance Davidson. 2015. "Tissue Mechanics and Adhesion during Embryo Development." *Developmental Biology* 401 (1): 152–64.
Sober, Elliot. 1999. "The Multiple Realizability Argument against Reductionism." *Philosophy of Science* 66:542–64.

https://doi.org/10.1086/705449 Published online by Cambridge University Press

Wang, Yau-De, and Hui-Hsien Hsieh. 2012. "Toward a Better Understanding of the Link between Ethical Climate and Job Satisfaction: A Multilevel Analysis." *Journal of Business Ethics* 105:535–45.

Waters, C. Kenneth. 2008. "Beyond Theoretical Reduction and Layer-Cake Antireduction: How DNA Retooled Genetics and Transformed Biological Practice." In *The Oxford Handbook of Philosophy of Biology*, ed. Michael Ruse, 238–62. New York: Oxford University Press.

Weber, Marcel. 2005. *Philosophy of Experimental Biology*. Cambridge: Cambridge University Press.

https://doi.org/10.1086/705449 Published online by Cambridge University Press