Explanation as Condition Satisfaction

Paul Humphreys*†

It is shown that three common conditions for scientific explanations are violated by a widely used class of domain-independent explanations. These explanations can accommodate both complex and noncomplex systems and do not require the use of detailed models of system-specific processes for their effectiveness, although they are compatible with such model-based explanations. The approach also shows how a clean separation can be maintained between mathematical representations and empirical content.

1. Introduction. Consider the following assumptions about scientific explanation. The first, which we can call the *hierarchy assumption* is that there exists a hierarchy of sciences, each with its own domain and subdomains, the hierarchy being based on compositional relations between elements of those domains. Thus, atomic physics occurs above nuclear physics, chemistry above atomic physics, molecular biology above chemistry, cellular biology above molecular biology, and so on.[1] Given this hierarchy, the project of reductive explanation then proceeds. In the formal mode, the goal of reductive explanation is to show that the theories and models, the predicates, and the law statements within the higher-level domains follow from similar items within the reducing domain, where 'follow from' is variously taken as 'be derivable from' or 'be formally realized by', and formal realization is taken to be a relation between predicates of logically different orders, or some other relation that depends on the specifics of the theory of reduction.

*To contact the author, please write to: Corcoran Department of Philosophy, 120 Cocke Hall, University of Virginia, Charlottesville, VA 22904; e-mail: pwh2a@virginia.edu.

†Helpful comments on earlier drafts of this article were provided by audiences at the PSA 2012 meeting and the University of Bielefeld. Conversations and correspondence with Mark Bedau, James Cargile, Meinard Kuhlmann, Margaret Morrison, Charles Rathkopf, and members of an informal University of Virginia seminar on mathematical models were also important.

1. This hierarchy assumes an object ontology rather than a field ontology or some other kind of noncompositional ontology. I am not assuming the correctness of this hierarchy but simply reporting its prevalence.

Philosophy of Science, 81 (December 2014) pp. 1103–1116. 0031-8248/2014/8105-0032$10.00
Copyright 2014 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/677698 Published online by Cambridge University Press

In such cases, explanatory reduction parallels formal theoretical reduction, which in turn parallels ontological reduction, whether explanation is taken to be causal explanation, compositional explanation, deductive explanation, or some other type.

A second assumption, which we can call the *domain-specific assumption*, is that theories, and the explanations based upon them, are taken to be domain specific. The Standard Model in physics concerns elementary particles, the theory of chemical bonding concerns molecules, models of protein folding concern macromolecules, and theories of mitochondrial adenosine triphosphate production concern cells.

A third assumption about explanation, which we can call the *nomological assumption*, is that a scientific explanation requires the employment of at least one scientific law. Specific facts are explained by appeal to laws, and the law or laws cited in such an explanation can in turn be explained by appeal to more fundamental laws.

Although these three assumptions are independent, the first two tend to be associated, and acceptance of them has led to widespread acceptance of the third, because explanations of laws tend to follow the hierarchical, domain-specific pattern of reduction. Not all accounts of scientific explanation accept all three assumptions. For example, some accounts of singular causal explanation as well as the loosely clustered group of positions on mechanisms deny that the nomological assumption is correct.[2] Unificationist accounts of explanation in principle favor the rejection of the domain-specific assumption, but in Kitcher's (1989) version examples of argument patterns are all drawn from specific sciences such as chemistry and linguistics. Overall, the joint influence of these assumptions is pervasive, and, implicitly or explicitly, the structure of truth-based accounts of explanation reflect at least one and often all three of these assumptions.

I shall argue that none of these three assumptions is well suited to a widely used class of explanations in the sciences. This class contains many explanations that use mathematical models as their primary representational vehicle, and these models are not domain specific, they are not hierarchically ordered, and they do not explain by reference to laws. In this article, I draw my examples from probabilistic explanations.[3] Probabilistic explanations are well suited to my purpose because their subject matter is not restricted to the domain of a single science and it is widely acknowledged that the probabilistic models used in the examples are applicable to a highly diverse range of subject matter. This choice of examples does not by itself show that the

2. As does Woodward (2000).

3. For simplicity I also restrict myself to univariate distributions. Generalization to multivariate distributions can be made, but additional assumptions beyond those used here are required.

https://doi.org/10.1086/677698 Published online by Cambridge University Press

domain-specific assumption is incorrect because whether an appeal to probability distributions can result in an informative explanation is itself something on which there exists a variety of views. The emphasis on probabilistic explanations does, however, have the virtue that such explanations can be applied both to complex systems and to noncomplex systems.

The second line of inquiry in this article is designed to understand the relations between explanations based on these transdisciplinary representations and explanations that appeal to lower-level, domain-specific processes. In what follows, technical details are necessary to provide an accurate presentation of the formal models, but I shall, as far as is possible, provide an informal interpretation of the formal apparatus.

2. Probabilistic Explanations in Noncomplex Systems. Suppose we are asked for an explanation of the empirical fact that approximately 0.13% of people who take an IQ test receive scores greater than 130. We are told that the data from the IQ test fit a Gaussian probability distribution, that the empirical mean score in the population is 100, that the empirical standard deviation is 10, and that for Gaussian distributions, 0.13% of scores will be higher than 3 standard deviations out in the right tail of the distribution. Because we are considering features of an empirical data set, the facts cited so far are descriptive and not explanatory. To get closer to an explanation we must answer the question 'Why do the empirical data fit a Gaussian so well?'. One way to do this is to address the question 'How were the data generated?'. This second question is crucial, not the least because it affects which interpretation of probability can be used. Many, perhaps most, relative frequency interpretations of probability restrict themselves to data sets and exclude the generating conditions. This restriction was often motivated by an empiricist epistemology, even though relative frequencies are not inherently empiricist and can be applied to data of any kind.[4] In contrast, propensity accounts are more sensitive to the role of the chance setup. For now we can remain indifferent between a single case propensity approach, a long-run propensity approach, and a relative frequency approach to the probabilities that is sensitive to the origins of the data. So the question is shifted to 'Why did the conditions that generated the data produce results that are accurately described by a Gaussian distribution?'. Now we can make progress. Many probability distributions, such as the Gaussian, need not be taken as a primitive fact about a population but can be constructed from a set of assumptions about the system that is generating the data. We now need to distinguish three different situations.

4. For example, pseudorandom numbers generated by a mathematical formula have relative frequencies.

https://doi.org/10.1086/677698 Published online by Cambridge University Press

First, we have the situation already described in which a set of physical generating conditions gives rise to a finite data set that is agreed, using standard statistical tests, to conform to an empirical Gaussian distribution. Second, we can identify mathematical conditions from which it can be proved that a mathematical Gaussian distribution will result. Third, we can attempt to identify physical conditions that can be argued to satisfy the mathematical conditions. By separating these three cases, we can better understand the role played by the different types of assumptions in the candidate explanations. It is easiest to begin with the mathematical case.

A running example to keep in mind is the situation in which there are n stochastic influences on an outcome, each represented by a random variable that can have positive or negative values. In the IQ example, it is plausible that there are multiple influences on how well an individual does on such a test, such as the educational environment of the individual, the commitment of the parents to education, nutrition, genetics, whether the native language of the individual is the same as that in which the test is administered, and so on. Some of those influences will tend to raise the score, and some will tend to lower it, but we do not need to know with any degree of specificity what those factors are, only that they exist. We are interested in the net effect of these influences, and this net effect is represented by the sum of the random variables. It is a substantive empirical matter whether such factors are additive and probabilistically independent, but if they are, the contributions remain bounded and no one factor dominates, the Central Limit Theorem tells us that a large set of data generated by measuring the IQ of multiple randomly selected individuals in a population for which the conditions are true will be well described by a Gaussian probability distribution having the familiar bell shape.[5] We can summarize the conditions from which the Central Limit Theorem follows as

1. The random variables contribute in an additive manner to the outcome and are probabilistically independent.
2. They have finite mean and variance.
3. The Lindeberg condition is satisfied.

These conditions are sufficient for the sequence $Z_n = \sum_{k=1}^{n}(X_k - \mu_k)/s_n$ to converge in distribution as $n \rightarrow \infty$ to a random variable having the standard normal distribution $N(0, 1)$.

5. There is more than one Central Limit Theorem. The one given here is more general than the familiar version in which the random variables are identically distributed as well as independent. I have not used the Lyapunov condition because it is harder to motivate empirically.

https://doi.org/10.1086/677698 Published online by Cambridge University Press

It is worth spelling out the Lindeberg condition explicitly. Take a collection of independent but not necessarily identically distributed random variables $\{X_i : 1 \leq i \leq \infty\}$ defined on the probability spaces $\langle \Omega, \mathfrak{B}, P_i \rangle$, and let F_i be the distribution function associated with P_i. Let $s_n^2 = \sum_{k=1}^{n} \sigma_k^2$, where σ_k^2 is the variance of X_k. Then $\{X_i\}$ satisfies the Lindeberg condition just in case

$$\lim_{n \to \infty} \frac{1}{s_n^2} \sum_{k=1}^{n} \int_{\{|X_k - \mu_k| > \epsilon s_n\}} (X_k - E(X_k))^2 \, dF_k = 0,$$

for all $\epsilon > 0$. Informally, the Lindeberg condition requires that the contribution of any single random variable to deviations from the mean must become arbitrarily small as $n \to \infty$.

If we can argue that conditions 1, 2, and 3 have empirical counterparts in the system from which the IQ data are taken, then because those conditions are generic, we have an explanation of why the data follow a Gaussian distribution that is not subject matter specific. We are then outside the familiar hierarchy of the sciences, and any further explanation will not need to appeal to sciences that are, in the case at hand, 'more fundamental' than psychology. The very strong domain independence of the Central Limit Theorem ranges from sociology to the Brownian motion of dust particles, and it holds regardless of the kinds of probability distributions that govern the individual component random variables. The reason is easy to see. All that condition 1 requires is that there are many stochastic influences on the outcome, some positive and some negative, and these influences are, in a special probabilistic sense, acting independently of one another. We do not need to know what these influences are, only that they exist. The Lindeberg condition requires that none of these influences dominates the rest. Condition 2 merely requires that fluctuations resulting from these influences remain bounded. Finally, no laws are cited in the conditions. Once those conditions hold in a given instance, the conclusion follows mathematically. One could generalize to a universal claim, and many statements of the theorem do so, but logically, one can remain at the level of singular statements.

What now is the status of our original explanatory question 'Why did approximately 0.13% of people who took the IQ test receive scores greater than 130?'? A preliminary answer is 'The influences on the IQ performance of any given individual are a result of many independent influences that are probabilistic in nature, some of those influences raise the score and some lower it in a bounded manner, the joint effect of those influences can be represented by an arithmetically additive operation, and none of those influences dominates the rest. Because of this, the probability of a randomly chosen individual achieving a score on the test of greater than a is described

https://doi.org/10.1086/677698 Published online by Cambridge University Press

by a Gaussian distribution, and thus when $a = 130$, that probability is 0.0013. Hence, when a large number of randomly chosen people take the test, approximately 0.13% of them will score higher than 130.' This is a scientific explanation with empirical content, it appeals only to very general conditions, and it does not use a detailed model of specific causal processes.

We now have a situation in which the second and third assumptions about explanation with which we began do not hold. There are no laws cited in this explanation, and so the third assumption fails. There is nothing essentially about psychology or any other particular science in the explanation, and so the second assumption fails. We can also extract an important positive claim from the example. The explanation succeeds (a) by identifying empirical conditions that generate the probabilistic data for which the explanation is requested, (b) because there are mathematical counterparts to those generating conditions, and (c) because from those mathematical conditions a mathematical counterpart to the fact needing to be explained can be proved.

To recap, three different situations must be distinguished.

1. The data are correctly described by a Gaussian distribution within the limits prescribed by standard statistical criteria of fit. The discrete empirical distribution is a description of the data and is not an explanation of them.
2. Conditions stated in mathematical language are given from which the Gaussian distribution can be derived as a mathematical object. The fact that the distribution applies in the infinite limit, given those mathematical conditions, follows with mathematical necessity.[6]
3. Physical conditions are described, and an essential part of the explanation is a claim that if those physical conditions are satisfied then the mathematical conditions mentioned in 2 above are satisfied. Whether the physical conditions hold is a contingent matter.

A key philosophical issue is then the relation between the conditions in 2 and the conditions in 3. To begin, note that the statement of the Central Limit Theorem contains mathematical terms such as 'random variable', 'additive', 'finite variance', and 'probabilistically independent'. There is no empirical content in those conditions, the proof, or the result. The mathematical conditions remain as mathematical conditions—they are not interpreted as gaining empirical content when they are satisfied by a particular system—and the necessity contributed by the proof to the establishment of the Gaussian distribution remains a mathematical necessity. To see what is

6. The sum of the random variables converges in distribution to the Gaussian, but this type of convergence is typical for probabilistic limit theorems and does not affect the conclusions drawn here.

https://doi.org/10.1086/677698 Published online by Cambridge University Press

needed to satisfy the mathematical conditions, consider a set $\{X_i\}$ of random variables from which we form the successive partial sums X_0, $X_0 + X_1$, $X_0 + X_1 + X_2$, . . . and consider a set of physical properties $\{P_i\}$ that we successively add to the stock of physical influences on a system. We must represent each P_i by a corresponding random variable X_i and justify representing the joint contribution of two such physical properties to the outcome by the sum of two random variables. Furthermore, the physical properties must be unrelated in a way that can be captured by probabilistic independence within the joint probability space on which the random variables are defined. Each of these representational claims must be justified in terms of a relation between the mathematical object and the natural object, and we can construe that relation as in part a satisfaction relation: P_i satisfies X_i, the joint contribution of P_i and P_j satisfies $X_i + X_j$, and so on.

It is tempting to construe this satisfaction relation simply as a semantic relation. That position is too simplistic; at the very least, the relation is vastly more complicated than the standard "'dog' refers to dog" kind of case. To take just one example, the relation between probabilistic independence and physical independence is notoriously difficult to articulate. It involves in part a matter of translation between the two vocabularies of independence, but the relation is not one of a bridge law because there is no reduction involved. Somewhat differently, the joint action of the physical factors that is represented by an addition operation in the probability space requires that the physical processes involved enable an additive measurement structure to be employed.

This representation allows us to account for the role of the theorem in explaining the empirical facts in terms of why conditions 1, 2, and 3 are satisfied. Then we have two choices regarding the structure of the explanation. As just described, we can present a mathematical explanation of the fact that $P(X > a) = p$ that is grounded in conditions 1, 2, and 3 and then argue that the mathematical conditions are satisfied by empirical conditions in the system, together with some physical interpretation of the probability. In that case, all of the empirical content of the explanation is contained in the satisfaction relations between the natural and the mathematical conditions. The contingency of the explanatory apparatus results from the contingent fact that satisfaction conditions for the mathematical conditions to apply were in place at that time in that system. Alternatively, one could construe the explanation as consisting entirely of elements in the natural realm. That approach is much harder to motivate because it prevents us from appealing to the proof of the mathematical Central Limit Theorem, and it results in a different kind of explanation, one grounded in an empirical generalization.

A further philosophical conclusion can now be drawn. No appeal to 'mechanisms' is made in the explanation even though we could, if so de-

https://doi.org/10.1086/677698 Published online by Cambridge University Press

sired. The generating conditions do contain physical properties, and the additive form of the contributing factors might suggest a parallel to decomposability, but the conditions are so abstract that nothing about the specific underlying processes can be inferred. Nevertheless, they are not examples of mathematical explanations as conceived of in some recent literature (e.g., Baker 2005).

3. Financial Markets. I now show that broadly similar conclusions can be drawn about certain kinds of complex systems. So let us compare the IQ example to one in which we are asked why more than 0.13% of returns on the New York Stock Exchange are greater than 3 standard deviations from the mean return.[7] In this case there is an important difference from the IQ example because the tails of such returns follow a power law distribution rather than a Gaussian distribution, so that the percentage of returns in the region targeted by the explanation is greater than 0.13% rather than equal to 0.13% as given by the Gaussian of our previous example. The empirical distribution of tail returns over a year of trading on the New York Stock Exchange and also on the London stock exchange is described by a Pareto distribution rather than a Gaussian distribution, where the Pareto distribution is given by $1 - F(x) = (a/x)^{\alpha}$ for $x > a$ and $F(x) = 0$ for $x \leq a$, where $a,\ \alpha > 0$.[8]

Such power law distributions are characteristic features of scale-free systems, and a discussion of the connection between the two aspects is given in Kuhlmann (2014). Bedau (2014) also provides a clear example of a model that can generate a power law distribution.

Once again, if we move to examining the conditions under which the data are generated, we can begin to provide an explanation. As background, we have four core assumptions of classical economics:[9]

1. All agents are perfectly rational and have the same knowledge. This allows the use of representative agent models in which every agent has the same properties.
2. Equilibrium states are the focus of the dynamics of the system.
3. Significant changes in share prices are caused by new information that is exogenous to the system. The arrival points of this information are

7. In this area the explanandum is often a qualitative feature that is found invariantly across a number of different economic systems. Such features, often called 'stylized facts', are neither singular facts nor laws and are often only approximately true of a given system.

8. For $\alpha < 1$, the variance is infinite.

9. These assumptions are modified versions of those given in Cristelli, Pietronero, and Zaccaria (2011).

https://doi.org/10.1086/677698 Published online by Cambridge University Press

independent and distributed randomly in time. Endogenous processes play only a minor role.
4. The basic mechanism of the market is self-adjustment to the new information, and the offer and sell prices are adjusted to restore equilibrium.

Under these classical assumptions (particularly the random external information assumption) the price of an equity should be a random walk with the fluctuations (i.e., returns) following a Gaussian distribution. Because this is false for the financial markets just mentioned, we have to reject at least one of these four assumptions and at least one of the three assumptions that are sufficient for the classical Central Limit Theorem.

4. Power Laws and Explanations in Complex Systems. How can we explain the empirical fact that the returns on specifiable financial markets are not accurately described by a Gaussian distribution? There exist Central Limit Theorems for power laws, but they are both more complicated and less intuitive than the classical Central Limit Theorem. For details I refer the reader to Nolan (2013), but the important content can be conveyed semiformally. Let X_1, X_2, and X_3 be independent random variables having identical distributions F. Then F is *stable* just in case for every $c,\ d \in \mathbb{R}^+$ there exists an $a \in \mathbb{R}^+$ and some $b \in \mathbb{R}$ such that $cX_2 + dX_3$ has the same distribution as $aX_1 + b$. For specifiable parametric values of stable distributions, the tails of the distributions are asymptotically power laws. Then, one form of a generalized Central Limit Theorem says that if $X_1, X_2, \ldots X_N$ are independent and identically distributed random variables with infinite variance and there is a sequence of constants $\{a_N, b_N\}$, $a_N > 0$, and $b_N \in \mathbb{R}$, then $a_N(X_1 + \ldots + X_N) + b_N$ converges in distribution to Z, where Z is a nondegenerate stable distribution. The key feature is that this generalized theorem holds when the variables contributing to the system output have infinite variance.

We thus see that there exist sufficient conditions for power law tails to exist, and those conditions can be stated abstractly without an appeal to domain-specific models or to scientific laws. The second important point is that this feature of having abstract sufficient conditions that can serve as the basis for explanations exists both in noncomplex systems and in complex systems and across a broad cross-disciplinary class of applications. I do not need here an uncontroversial definition of what counts as a complex system. It is sufficient to note that Gaussian distributions apply to systems that are not generally regarded as complex, such as the total number of heads resulting from tossing a large number of identical coins, and that power law distributions are widely held to be a characteristic feature of complex systems, of which financial markets are a widely agreed-upon example.

https://doi.org/10.1086/677698 Published online by Cambridge University Press

Because power law distributions can be produced by a variety of models that use different processes to generate the distribution, there is pressure to drop down from the distribution itself to the specific model that applies to the market in question in order to explain the behavior. We can drop down to the level of subject matter specific processes if desired, but I have shown that this is unnecessary in the cases considered to arrive at an understanding of why the system produces the outputs that it does. In contrast, Bedau (2014) provides an interesting set of examples in which detailed processes are the starting point and the distribution of outcomes is the effect of those generating processes.

As should be obvious from the conditions for the generalized Central Limit Theorem, the way to move to more specific models in the financial case is to find processes that can generate arbitrarily large changes in the market returns. Because of space constraints, I mention, but do not describe in detail, only two such models, a modified version of the Minority Game and the well-known Lux and Marchesi models.[10]

The Minority Game is a highly simplified model of equity markets. The basic model has N agents faced with the binary decision to buy or sell at each time step. If agent m_i is in the minority with respect to the aggregate choices made in the group, then m_i receives a payoff. Conversely, if m_i is in the majority, a negative payoff occurs. This models the fact that if there are more buyers than sellers in a market, the sellers receive a premium price, and if there are more sellers than buyers, the buyers obtain the asset at a reduced price. The information I available to each agent at time t is a combination of the success of the agent's strategy up through $t - 1$ and a random number representing exogenous information.

In the modified Minority Game model of Challet, Marsili, and Zhang (2001), the N agents are divided into two groups, producers and speculators. Producers act only on information that is exogenous to the market, whereas speculators use strategies, which are maps $S : I \rightarrow \{\text{buy, sell, stay out of market}\}$. Power law tails are produced by this modified model with the exponent value depending on the parameter settings in the model. The dominant process producing the large price changes is agents entering the market en masse when the information I is favorable.

Compare the Challet et al. model with a variant of the well-known model of Lux and Marchesi (1999). This variant contains N agents divided into

10. For details of these and other models, see Cristelli et al. (2011). The Lux and Marchesi model is discussed in more detail in Kuhlmann (2014). Different types of examples of explanations within complex systems that appeal to such details and that generate different kinds of distributions are given in Bedau (2014). In this article I use the term 'processes' rather than 'mechanisms'. Use of the latter is likely to introduce philosophical biases that are distracting here. Arguments about the need for lower-level mechanisms are given in Kuhlmann (2011, 2014).

https://doi.org/10.1086/677698 Published online by Cambridge University Press

fundamentalists and speculators. The fundamentalists follow the traditional strategy of identifying a fair price on the basis of a company's economic fundamentals. Fundamentalists will buy a stock if its price is below the fair value and sell if it is above. This fair value price p_f is the same for all fundamentalist agents. Speculators believe that prices will deviate from p_f and imitate patterns of buying and selling by other agents. The numbers of agents in each category can vary as long as N remains fixed. At moderate values of N, fluctuations in the transition rate can lead to periods of self-reinforcing speculation leading to power law tails and volatility clusters. For us, the important point is that in both the Challet and the Lux and Marchesi models, there are processes—cascades of agents entering the market in the former, herding behavior in the latter—that can produce unbounded swings in prices that lead to a higher percentage of extreme returns than is predicted by a Gaussian distribution.[11]

5. Philosophical Issues. If, as many writers on probabilistic explanations have suggested, an appeal to probability distributions such as the Gaussian can serve as the basis of an explanation, what is the relation between those abstract explanations and explanations of the same facts that appeal to system-specific features and processes? The position argued here is that the tension between abstract probabilistic distributions and specific models of generating processes is in part a product of ignoring the intermediate level of construction assumptions for the distributions.[12] Those assumptions, when in place, probabilistically guarantee that the appropriate data distribution will occur in the limit, and because the assumptions are largely subject matter independent, they can be realized in many different kinds of systems. The understanding comes in two stages, by understanding, first, that any system having a specifiable set of generic features will generate data conforming to a particular distribution and then, second, that the specific system involved has those generic features.

When a general explanation can be given that is independent of the details of the specific system, we can often proceed to a more detailed explanation in terms of the specific system's characteristics that explains why the conditions that are sufficient for the general explanation are satisfied for the system under consideration. Such a move can be motivated by the need to change the processes operating in the system. In the case of financial markets, this could be to prevent flash crashes that undermine investor confidence, and such interventions are usually not possible at the level of the

11. Another such process is an avalanche of selling orders by automated trading programs that produces increasingly large deviations in price from the historical averages.

12. A discussion of the role of construction assumptions applied to the case of differential equation models can be found in Humphreys (2004), secs. 3.6 and 3.7.

https://doi.org/10.1086/677698 Published online by Cambridge University Press

more general conditions. Note that the explanatory order here moves from the general to the particular, thereby inverting the order in which unificationist accounts find unity in previously disparate explanations. Although it is not inherent in unificationist approaches to explanation that theory T unifies theories $T_1, \ldots, T_n$ only if T is more general than each of the $T_1, \ldots, T_n$, this seems to be an implicit assumption in those accounts. Our explanations also proceed in the opposite direction to Hempelian explanations. Hempelian explanations require laws, and the explanation of why a law applies to a particular situation is in turn explained by more general laws. So as explanations become 'deeper', we proceed to greater and greater generality. Indeed, the very metaphor of depth is based on the hierarchy assumption.

I note here that stable distributions, those that appear in the generalized Central Limit Theorem, turn out to be the only fixed points of renormalization group transformations, which themselves represent aggregation procedures. The fact that those methods are widely used in condensed matter physics does not require one to think that they have some kind of explanatory priority. The conditions for the generalized result to hold can be stated without reference to renormalization methods, and the fact that they can plausibly be argued to hold for economic systems is a further reason to reject the hierarchy assumption for such situations. Whatever appeal physicalism may have for other reasons, it is an unnecessary assumption here.[13]

Notice also that the relation between the abstract explanation and the more specific explanation is not one that fits into the hierarchy of subject matter specific theories connected by bridge laws. The relation is one of instancing general conditions rather than showing that the properties of one subject area can be reduced to the properties of another area.

6. Conclusion. I have shown that in the case of Gaussian and power law probability distributions we can have explanations that are not domain specific and in which no appeal to laws is made. Those examples are easily generalized, although the scope of this approach to explanations is an open question. It can be applied straightforwardly to some other probability distributions, such as the Poisson, for which explicit sufficient conditions can be given, and it can also be used for scientific models, such as differential equation models, for which an explicit construction of the model can be given from abstract conditions.[14] Such domain-unspecific explanations can

13. As Humphreys (2000) has argued, the physicalist assumption built into the causal mechanical account of explanation is far too stringent. This and the points made in the text do not lend support to disunity views of science—the theorems discussed here as well as complexity theory itself provide an enormous amount of unity across the sciences, just not in the way that contemporary unification accounts have suggested.

14. See discussion in n. 12.

https://doi.org/10.1086/677698 Published online by Cambridge University Press

be supplemented by domain-specific models both in cases in which the phenomena arise in complex systems and those in which more traditional noncomplex systems play a role. There are affinities with some other models of explanation. Railton's (1981) deductive-nomological-probabilistic account of probabilistic explanation contains elements of this approach, but it does not locate the explanatory power in the conditions needed for the probability distribution to hold. Kitcher's (1989) version of the unificationist approach, with its emphasis on argument patters, has some overlap with the approach given here, but even setting aside the fact that all of Kitcher's examples are domain specific, the unifying power of the common representational apparatus is a by-product of the reliance on true sufficient empirical conditions rather than argument patterns. Finally, the principal open philosophical question for the approach proposed here is to articulate the precise relations between the empirical conditions that are represented by the mathematical conditions and the mathematical conditions themselves. This is not a formal mapping, as it is in the semantic account of theories, nor is it a purely semantic relation as in more linguistically oriented approaches. It is, I believe, where all or almost all of the empirical content of a representation lies, but that is an argument for another occasion.[15]

REFERENCES

Baker, Alan. 2005. "Are There Genuine Mathematical Explanations of Physical Phenomena?" *Mind* 114:223–38.

Bedau, Mark. 2014. "Testing Bottom-Up Models of Complex Citation Networks." *Philosophy of Science*, in this issue.

Challet, Damien, Matteo Marsili, and Yi-Cheng Zhang. 2001. "Stylized Facts of Financial Markets and Market Crashes in Minority Games." *Physica* A 294:514–24.

Cristelli, M., M. L. Pietronero, and A. Zaccaria. 2011. "Critical Overview of Agent-Based Models for Economics." ArXiv, Cornell University. http://arxiv.org/abs/1101.1847.

Humphreys, Paul. 2000. Review of *Causality and Explanation*, by Wesley C. Salmon. *Journal of Philosophy* 97:523–27.

———. 2004. *Extending Ourselves*. New York: Oxford University Press.

———. 2008. "Probability Theory and Its Models." In *Probability and Statistics: Essays in Honor of David A. Freedman*, ed. Terry Speed and Deborah Nolan, 1–11. Beachwood, OH: Institute of Mathematical Statistics.

Kitcher, Philip. 1989. "Explanatory Unification and the Causal Structure of the World." In *Scientific Explanation*, ed. P. Kitcher and W. Salmon, 410–505. Minnesota Studies in the Philosophy of Science 13. Minneapolis: University of Minnesota Press.

Kuhlmann, Meinard. 2011. "Mechanisms in Dynamically Complex Systems." In *Causality in the Sciences*, ed. P. McKay Illari, F. Russo, and J. Williamson, 880–906. Oxford: Oxford University Press.

———. 2014. "Explaining Financial Markets in Terms of Complex Systems." *Philosophy of Science*, in this issue.

Lux, Thomas, and Michele Marchesi. 1999. "Scaling and Criticality in a Stochastic Multi-Agent Model of a Financial Market." *Nature* 397:498–500.

15. Protoarguments for this view can be found in Humphreys (2008), sec. 7.

https://doi.org/10.1086/677698 Published online by Cambridge University Press

Nolan, J. P. 2013. "Stable Distributions: Models for Heavy Tailed Data." Unpublished manuscript, Mathematics and Statistics Department, American University. http://academic2.american.edu/~jpnolan/stable/chap1.pdf.
Railton, Peter. 1981. "Probability, Explanation, and Information." *Synthese* 48:233–56.
Woodward, James. 2000. "Explanation and Invariance in the Special Sciences." *British Journal for the Philosophy of Science* 51:197–254.

https://doi.org/10.1086/677698 Published online by Cambridge University Press