

The Structure of Asymptotic Idealization

Michael Strevens

Department of Philosophy
New York University
strevens@nyu.edu

ABSTRACT

Robert Batterman and others have argued that certain idealizing explanations have an asymptotic form: they account for a state of affairs or behavior by showing that it emerges “in the limit”. Asymptotic idealizations are interesting in many ways, but is there anything special about them as idealizations? To understand their role in science, must we augment our philosophical theories of idealization? This paper uses simple examples of asymptotic idealization in population genetics to argue for an affirmative answer and proposes a general schema for asymptotic idealization, drawing on insights from Batterman’s treatment and from John Norton’s subsequent critique.

ACKNOWLEDGMENTS

For invaluable discussion and feedback, thanks to Bob Batterman and John Norton, Kate Vredenburg, the Research Triangle Philosophy of Science Reading Group, the audience at ISHPSSB 2013, and the anonymous reviewers at *Synthese*.

1. Asymptotic Idealization

In *The Devil in the Details*, Robert Batterman discusses a class of idealized models that provide what he calls “asymptotic explanations” (Batterman 2002). Such idealizations, he claims dramatically, though they rest on infinities that do not exist in the physical world, are essential for explaining certain physical behaviors; they do this by revealing emergent qualitative features of the situation that dominate the relevant actual physics and so “constrain and largely determine” what happens, or at least those aspects of what happens that are to be explained.¹ Other writers, such as Belot (2005), Wayne (2012), Pincock (2012), and Norton (2012), have disputed Batterman’s views, arguing that asymptotic idealization is not the sole route to understanding the phenomena in question (Belot and Norton) or that idealization does not necessarily imply the emergence of new properties (Wayne and Pincock).

In order to adjudicate this debate, it is surely important to have a clear picture of the role that asymptotic idealizations play in model-building—that is, a philosophical account of the rules that guide scientists in their construction and interpretation of models that purport to use idealizing limit-taking processes to predict and explain. This essay attempts to make a contribution to the topic by posing the following question: is there any interesting difference between the rules that govern idealization in such canonical non-infinitary cases as frictionless planes and ideal gases, on the one hand, and Batterman’s asymptotic cases, on the other? I follow Batterman in focusing on the enterprise of explanation; although I believe that much of what I have to say goes for prediction, I do not have the space to justify that assertion here.

I will lay out a straightforward characterization of the rules that apply to the non-infinitary cases—the rules of what I will call “simple idealization”—and then ask whether such rules can be used to deal with every case involving infinities. The answer is that they can cope with some, but not all, infinities.

1. Quoting Batterman (2010), p. 22.

Batterman-style idealizations, in particular, demand something more. I attempt to provide that something more, offering an account of the structure of asymptotic explanatory models, while remaining ecumenical about the purpose of idealization and the nature of explanation.

My nominal opponent throughout is Norton (2012), who argues that Batterman's examples require no new thinking about idealization: in such cases, "far from being ineliminable, there are no infinite idealizations employed" (p. 208). But I should add that my thinking is in many respects rather closer to Norton's than to Batterman's. For reasons of space, I will stop short of engaging with Batterman's claims about essentiality and emergence, and indeed I will avoid Batterman's own interesting but complicated examples, resting my views on a simple example from population genetics that can be grasped with only elementary mathematical and scientific knowledge.

2. Simple Idealization

Take some pattern of behavior that you wish your theory to explain—the *target phenomenon*. It might be the approximately inverse relationship of pressure and volume in a dilute gas at equilibrium asserted by Boyle's law, the pattern of concentrated light (a "caustic") created by sunlight refracted through a glass of water, or the evolution of melanism in *Biston betularia* moths during the Industrial Revolution. Then go about your explanatory task, by building a model representing the system that exhibits the target phenomenon—for example, a model of a gas, of the refraction of light by water, or of the ecology of British woodlands.

Such models characteristically represent the system inaccurately. The ideal gas model represents gas molecules as being infinitely small, for example, and the ray model of refraction represents light as consisting of one-dimensional beams rather than waves. These fictions are called (at least when deliberate rather than inadvertent) *idealizations*.

Idealization tends to reduce the complexity of a model, by comparison

with a more accurate model that foregoes the idealization. When long-range forces in a gas are set to zero, the physical facts about the nature of long-range attraction—the details as to how much force one particle exerts on another and by what underlying process—are rendered irrelevant, and they are typically omitted from the idealized model altogether. The forces are “idealized away”.

This streamlining hints at one important function of idealization: it serves the pragmatic end of keeping calculations simple, either by making the mathematics more tractable or by reducing the empirical demands of the model (for example, lowering the number of and precision of the initial conditions and parameters whose values must be determined in order to derive consequences from the model).

I have argued elsewhere that idealization has an important function in addition to simplification: idealizations call attention to factors that are explanatorily irrelevant to the target phenomenon. The deliberate falsification of certain details, then, is a way of communicating forcefully the idea that those details do no explanatory work (Strevens 2008, chap. 8). The infinitely small size of molecules in an ideal gas tells us, for example, that intermolecular collisions (absent in an ideal gas) make no contribution to the explanation of Boyle’s law, while the erasure of wave-like properties in the ray model of light tells us that those properties are irrelevant to the structure of caustics.

I will put aside these matters while I develop my story about the asymptotic approach, however, returning to the question of the function of idealization only at the end of section 5. My characterization of the nature of asymptotic idealization will, as a consequence, be compatible with any of a number of views about idealization’s scientific role: the pragmatic view, my irrelevance-signaling view, and many others.²

I will introduce asymptotic idealization by contrasting it with what I

2. For a great deal more on idealization and various approaches to understanding its role in explanation and elsewhere, see McMullin (1985); Weisberg (2007); Elgin (2007); Potochnik (2017).

will call *simple idealization*. Because asymptotic idealization is characterized with reference to the process by which an idealized model is created—that is, some sort of limiting process—it will be illuminating to offer a procedural characterization of simple idealization too, identifying it not so much by the structure of the end product, a certain deliberately misrepresenting model, but by the way in which the structure is derived.

To that end, I say that a simple idealization notionally begins with a non-idealized model of the target phenomenon—a veridical model, representing the system in question entirely accurately. (Real science almost never constructs such models, of course; hence “notionally”. What is important here is the relation between the idealized model and the veridical model; the conceit that the veridical model comes first is a dramatization of this relation for philosophical purposes, to emphasize and explore the connection between the idealized model and the truth.)

Now an idealization is made—the veridical model is altered in some way that makes it no longer veridical. In a simple idealization, this is achieved by the straightforward operation of setting some parameter or parameters in the model to non-actual values, often zero. A model of projectile trajectories, for example, might be idealized by setting the force due to air resistance to zero, so falsely representing air resistance as entirely absent. A molecular-level model of a gas might be idealized by setting the long-range intermolecular force to zero, so falsely representing long-range forces as absent. If the idealization is a valid one, the new model is, like the old veridical model, a good model of the target phenomenon, by whatever criteria determine such things.³ Further, the new model will do whatever idealized models are supposed to do: it will ease calculations, or it will effectively spotlight irrelevant factors.

The salient characteristic of a simple idealization, then, is that a simple operation, namely, a *change in the parameters* of the veridical model, creates a new, partly fictional model that continues to model and indeed to explain the

3. Weisberg (2013) provides a treatment of the question.

target phenomenon.

In asymptotic idealization, by contrast, a fiction is introduced by taking some sort of limit. Rather than setting a parameter to zero, for example, a limit is taken as that parameter approaches zero. In some cases, taking a limit in this way is merely a roundabout path to a destination that can be more easily reached by simple idealization. You could create an ideal gas model by taking the limit of the veridical model (in a certain sense to be characterized later) as the long-range force strength and molecular size go to zero, but you would merely be showing off. The same model can be built by simply substituting zero for the veridical values of those parameters.

Asymptotic idealization is an interesting proposition, then, only in those cases where a simple substitution cannot be performed, which is to say only in those cases where a veridical model for mathematical reasons falls apart or otherwise behaves badly at the limiting value. The hope, very roughly, is that we can derive some explanatory insight from the model by observing its behavior as it approaches the limit, if not at the limit.

John Norton (2012) has recently argued, however, that in these “interesting” cases, an asymptotically derived structure cannot be regarded as a genuine explanatory model of the system in question. At best, such structures must be interpreted as mere aids to calculation, used to approximate the behavior of models that are not asymptotically derived and that are not in any sense infinitary—models that assign values to the parameters in question for which the veridical model is well behaved. All purported cases of asymptotic idealization, then, can either be achieved through a more straightforward simple idealization (in cases where the veridical model is well behaved for the idealized values of the parameters in question), or must fail altogether as attempts at asymptotic modeling (though they may have practical uses as mathematical devices).

To assess Norton’s argument, we need a more precise understanding of asymptotic idealizing explanation. Batterman’s principal examples of this

enterprise—of the explanation of fringe patterns on caustics (such as rainbows’ “supernumerary bows”) and of critical points in phase transitions—are rich and complex. The richness is welcome; the complexity less so. Rather than reviewing Batterman’s cases, then, I will philosophize about asymptotic idealization using a far simpler example. I do not claim that my example has all the interesting features of Batterman’s cases, but as a paradigm of asymptotic explanation it is sufficiently deep for my purposes: it can be used to counter Norton and to affirm Batterman, showing that not all idealization is simple idealization.

3. Population Genetics and Determinism

Models in population genetics represent changes in a population’s gene frequency due to natural selection, mutation, and the effects of stochastic processes that can “by chance” affect a gene’s flow from one generation to the next.

A veridical population genetics model would represent all of these processes in the utmost detail. Naturally, no one attempts to build such representations, but some models are far less realistic than others. Many models, for example, omit the effect of mutation, and an especially important class of idealized models leaves out altogether not only mutation but every other stochastic aspect of gene transmission, resulting in a purely deterministic model of the change of gene frequencies. It is the derivation and rationalization of these deterministic models that is my topic in what follows.

The stochastic aspect of gene transmission, putting mutation aside, is called genetic drift.⁴ What the deterministic models omit, then, is mutation and drift. Such models have purely pedagogical uses, but they are also use-

4. Some writers reserve “drift” for the effects of the third class of processes enumerated in the opening paragraph of this section, thus not for the stochastic element of natural selection. My use of the term in the main text to include variability in the effects of selection is the more standard one in population genetics.

ful for explaining phenomena to which mutation and drift are irrelevant. A typical and well-known example is the explanation of the persistence of the hemoglobin sickling gene in human populations exposed to malaria—surprising because two copies of the gene cause sickle-cell anemia, a severe physiological disorder. The reason the gene is not driven extinct is that having a single copy conveys considerable resistance against malaria; it follows that the gene pool will, in malaria-prone areas, find an evolutionary equilibrium where the sickling gene is present in significant numbers (though its proportion cannot rise too high). Neither mutation nor drift plays a role in this explanation; only the deterministic effects of selection are explanatorily relevant. The idealized deterministic model, then, is up to the explanatory job.

An appropriate version of the veridical model is also up to the job, of course (or so I believe). But the deterministic model is often preferred by scientists for the various reasons given by philosophers of idealization: to simplify calculation, to signal clearly the irrelevance of mutation and drift, and perhaps other reasons as well. At this stage, I simply assume that the idealized model is useful for some reason or other; my question is how it is built.

Mutation may be idealized away by setting a parameter, the probability of mutation, to zero; this is what I called in the previous section a simple idealization. To idealize away drift is not so straightforward, since it is not represented by a single number or set of numbers in the veridical model. Consider meiosis in female mammals, for example. The cell that undergoes meiosis (the oogonium) splits into a “polar body” and a “daughter cell”; the daughter cell then splits into another polar body and a cell that becomes (when circumstances are right) a mature ovum. The ovum has half the mother’s cells, selected roughly at random. The process is therefore stochastic, and its stochasticity, in conjunction with the stochasticity of fertilization itself, can (especially in small populations) result in some genes being better represented

than others in the next generation simply by “good luck”. Meiosis contributes in this way to random genetic drift. We want to expunge drift from the model in the cause of idealization. But there is no parameter or set of parameters in the model’s representation of meiosis that can be set to zero in order to remove its stochasticity—drift is, as it were, built into the mechanics of meiosis itself.

The same conclusion can be reached if we take as the starting point not a full veridical model that represents meiosis explicitly, but the characteristic stochastic models of population genetics. These models say nothing about the mechanics of meiosis; they represent its effects rather by putting a probability distribution over a gene’s being transmitted to the next generation. In the absence of selection and mutation, this distribution is a simple Bernoulli distribution that ascribes to each of the alleles at a given locus a 50% chance of being transmitted to any particular offspring.⁵ Mathematically, it is precisely the same probability distribution that characterizes tosses of a fair coin.

In order to “idealize away” drift simply, what is wanted is a parameter in such a probability distribution that can be set to zero (or some other value) to obtain a deterministic process. But there is, in general, no such parameter. There is no simple idealization, then, that transforms a more realistic, hence stochastic, population genetics model (let alone a veridical model) into a deterministic model.

One way forward is to build a deterministic model from scratch, and to rationalize the use of the model to explain the fixation of a phenotype or some other evolutionary phenomenon by pointing out that its behavior closely tracks the behavior of the real system in the relevant respects. This strategy is what Norton calls “approximation”; I will discuss its features and disadvantages in the next section.

An alternative strategy employed by some population geneticists is, by

5. Individual offspring are not always represented in the models of population genetics; many models keep track only of the composition of the gene pool as a whole. More comprehensive models keep track of genotypes, and thus in effect of individuals. In what follows it will be convenient to suppose that we are dealing with the more realistic genotype models.

contrast, explicitly asymptotic. Deterministic models should be regarded, these theorists say, as derived from more realistic models under the assumption that population size is infinite—since in an infinite population, there is no drift. (See, for example, Gillespie (2004), p. 21.) Drift cannot be set to zero directly, then, but it can be zeroed indirectly by taking population size, another parameter in stochastic models, to be infinitely large.

That may sound like a simple idealization. As I will show you in the next section, however, it is not: the population size in a stochastic model cannot be set to infinity, not even by allowing the structure of the model to creep up to infinity by way of a limiting process. Are population geneticists making a mistake? That is Norton (2012)'s contention (concerning related examples). I will argue for an alternative construal of the infinite population idealization, on which it is well founded and has an important scientific role to play, especially in explanation. Indeed, the infinite population strategy provides an especially simple arena for thinking about what it is, exactly, that asymptotic idealization amounts to.

4. Interpreting the Infinite Population Idealization

Consider a very simple case of natural selection: selection on a single locus of a haploid organism with two alleles.⁶ Suppose that the two alleles correspond to two possible variants of a trait, one of which provides a relative adaptive advantage over the other. The evolutionary process can in this case be modeled in terms of organisms rather than genes, thereby eliminating the genetic jargon, which I choose to do for simplicity's sake. A simple stochastic model of the process will represent two sub-populations, one for each variant, and a reproductive probability distribution. The distribution takes as its input the

6. In such a population, an organism passes on all its genes to its offspring, so there is no meiosis-like source of genetic drift. Stochasticity comes in many other ways, however, such as environmental accidents that kill one but not the other of two genetically identical individuals. There is plentiful drift to be found, then, in finite populations of the sort considered here.

relative frequencies of the variants in the present generation, and yields the probabilities for various possible frequencies in the next generation. Told that 50% of individuals in the present generation have variant *A* and 50% have variant *B*, for example, the distribution will determine the probability that the percentage of *A* individuals in the next generation is 45%, or 50%, or 55%, and so on. Such a model might be used to explain both the direction and the approximate rate of evolutionary change, as well as an ultimate outcome such as one variant's entirely taking over the population.

The model's reproductive probability distribution has, let me specify, just a single adjustable parameter, namely, total population size. As the size increases, the variance of the probability distribution will (given some assumptions of very general validity) decrease. Suppose, for example, that in a population with 10 individuals, and given a 50/50 split of variants in the present generation, the probability distribution for the frequency of variant *A* in the next generation has the Gaussian (that is, normal) form shown in Figure 1a. (As you can see, the expected value for the frequency—the mean of the distribution—is 60%, so the *As* will more likely than not undergo an increase in population relative to the *Bs*.) The larger the size of the population, the narrower the distribution will tend to be. For a starting population of 25 individuals, for example, the distribution over the next generation might be as shown in Figure 1b, for 50 individuals as shown in Figure 1c, and so on for greater numbers, as shown in figures 1d through Figure 1f.

The variance of the distribution roughly corresponds to genetic drift: it represents the “chance” elements of the stochastic process of gene transmission in virtue of which the frequencies in the next generation stray from the expected frequencies (Ramsey 2013). As population size increases, then, the magnitude of drift gets smaller and smaller. In the limit, variance disappears altogether: the probability distributions shown in Figure 1 converge on a distribution that ascribes a probability of one to the expected outcome—that is, a next-generation frequency of 60% for variant *A*—and a probability of zero

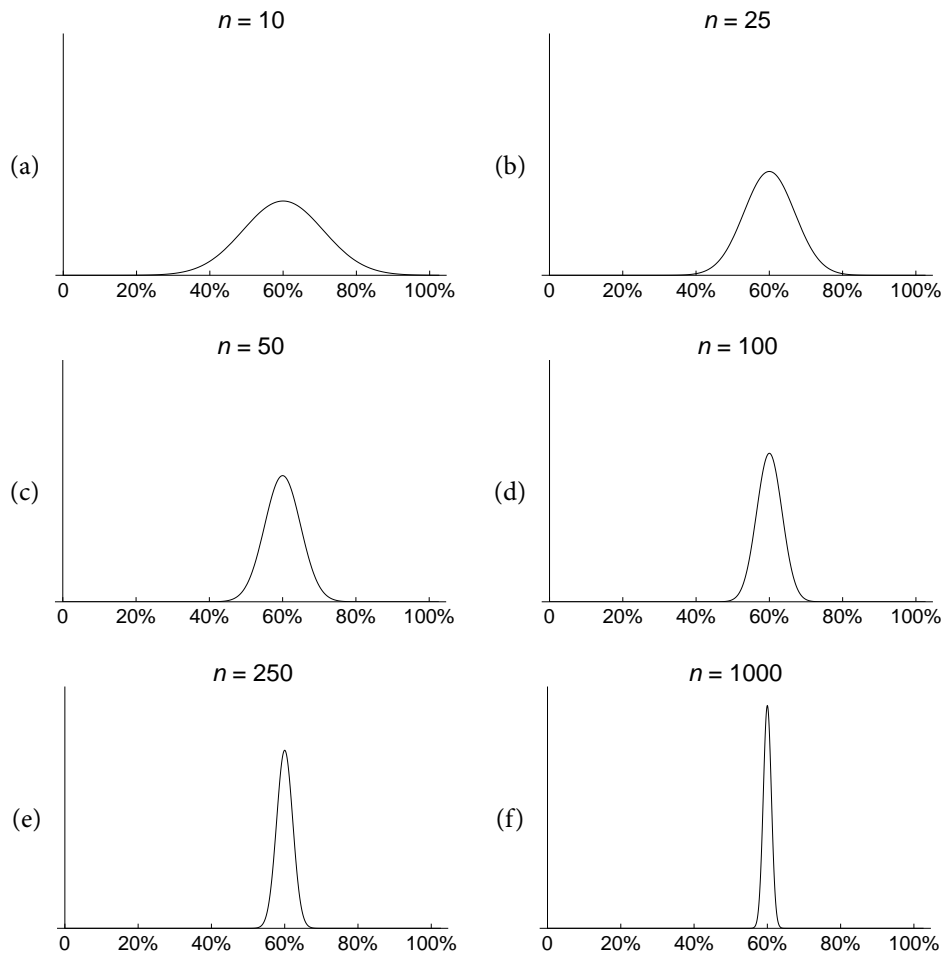


Figure 1: A probability distribution over the distribution of a trait variant in the next generation given various total population sizes: (a) 10; (b) 25; (c) 50; (d) 100; (e) 250; (f) 1000. The expected frequency of 60% remains the same, but the variance decreases as the population size increases. (As you can see by comparing the areas under the curves, some scaling has been performed.)

to any other frequency. This does not quite imply that outcomes other than a 60% frequency are impossible, but I will not quibble about zeroes; let me say for the sake of the argument that what emerges in the limit is deterministic change, thus evolution with no drift at all. Here you have the nub of the idea that drift can be eliminated from a model of population change—it can be set to zero—by specifying an infinite population size.

And so it is that population geneticists say that under the assumption of infinitely large populations, there is no drift and so changes in gene or genotype frequency are (putting aside mutation) deterministic, dictated by selection alone.

That suggests that a simple idealization is being made: the deterministic models that appear in population genetics can be obtained, you might suppose, from more realistic stochastic models by setting the population size parameter to infinity—or, less crudely, by allowing the population size to approach infinity and taking as the “no drift” model whatever structure the stochastic model takes on in the limit. As I have said, however, things are not so straightforward. Probabilistic population genetics models do not survive the passage to the limit: when the population size becomes infinite, the structure over which the models assign probabilities disintegrates.

The reason is pointed out by Abrams (2006). Relative frequencies (other than zero and one) are defined only for finite populations: if there are infinitely many *As* and infinitely many *Bs*, there is no determinate ratio of one to the other. (Consequently, anything that depends on that ratio, such as the probability of a randomly sampled organism’s being an *A*, is equally indeterminate.)

In an infinite population, probability mathematicians therefore turn to limiting frequencies. But a limiting frequency is defined only relative to an ordering of outcomes. Suppose that I have an infinitely big pile of *As* and an infinitely big pile of *Bs*. I can alternately take one from each pile and put them in an infinitely long sequence. That gives me a limiting frequency of one-half

for the *As*. But I can equally well build an infinitely long sequence by taking two *As* for every one *B*, giving me a limiting frequency of two-thirds for the *As*. (There is no prospect of running short of *As*; there are infinitely many.) So the limiting frequency depends on how the objects in question are ordered in a sequence.

It makes sense to talk about the limiting frequency of heads in an infinite series of coin tosses because the tosses are ordered; it is possible in such a case to define a probability distribution over different limiting frequencies and to show—using the strong law of large numbers—that with probability one the limiting frequency of heads in an infinite series of tosses is equal to the expected frequency (for a fair coin) of one-half.

In a population of organisms (or genotypes, or genes), by contrast, there is no ordering, and so in an infinite population there is no fact of the matter about limiting frequencies—no fact of the matter about, say, the frequency of organisms with variant *A*.⁷ Indeed, for any rational number you like there exists an ordering relative to which the limiting frequency of *As* is equal to that number.

Without a preferred ordering, stochastic models of population genetics lack the structure to represent limiting frequencies of variants or alleles; *a fortiori*, they lack the structure needed to apply the strong law of large numbers so as to derive a population dynamics for the infinite case in which the frequencies in future generations are equal to the expected frequencies, as the infinite population idealization appears to require. Consequently, the infinite population idealization cannot be treated as a simple idealization in which the population is set to infinity.

How, then, to understand it? There are three options:

1. It is an error. Biologists ought to renounce talk of infinite populations

7. I assume that the infinite population contains infinite numbers of both *A* and *B* variants. If not, then there are well-defined relative frequencies of zero and one but there is no dynamics to model.

and find some other rationale for the use of deterministic models in population genetics.

2. Biologists are tacitly building additional structure into population genetics models, enabling them to represent limiting frequencies and thus probability distributions over changes in limiting frequencies.
3. The infinite population idealization is a successful idealization, but it is not a simple idealization.

Let me discuss these in turn. The third option will lead to my proposal concerning the structure of asymptotic idealization.

The first option is to regard biologists' talk of "infinite populations" as a (possibly harmless) mistake. Rather than saying that evolution is deterministic in infinite populations, biologists should say that deviation from the expected value is, in very large populations, almost certain to be very small.

What, then, is the rationale for the use of literally deterministic models in population genetics? Edwards (1977, vii), avoiding talk of infinite populations, writes that deterministic models are appropriate where "populations [are] large enough for stochastic variation to be neglected". Perhaps the deterministic model ought simply to be regarded as a stand-in, then, for a stochastic model with a large population size, its use justified by its making predictions that are close to those made by the stochastic model. It is, in that case, strictly speaking a mere calculation device, a formal structure that has heuristic use as an estimator of the behavior of an explanatorily legitimate model. Talk of "deterministic models" in population genetics is therefore misleading; they are not models at all.

This, perhaps, is Abrams's (2006) view. It is also more or less the treatment advocated by Norton (2012) for cases in physics where models disintegrate in the limit. Norton writes that when such breakdowns occur, we should regard talk of infinite populations and the sort of informal reasoning exemplified by Figure 1 not as specifying a well-defined model, but as suggesting an "inexact

description” of the target system. More precisely (focusing on the population genetics case), the asymptotically derived probability distribution should be regarded not as the correct probability distribution for a system containing infinite populations, but rather as an incorrect, though approximately correct, proxy for the actual, finite-population distribution.

The justification for using such a proxy has nothing to do with asymptotic thinking, on this approach; what matters is simply that it is close enough to the actual probability distribution to provide a convenient estimate of the actual distribution’s implications. Consequently there is no infinitary model, merely a piece of infinitary mathematics that derives certain behaviors of a non-infinitary model. (This is why Norton says that in these cases “there are no infinite idealizations employed”: an infinite idealization by its very nature goes by way of an infinitary model.)

Perhaps the least satisfying feature of this revisionary approach is its implication that no evolutionary events can be properly understood using a deterministic population genetics model—since such a construct is not itself a legitimate model of the system, but is rather only a mathematical heuristic to calculate what a certain legitimate model will predict. That attitude fails to explain the importance of deterministic models as genuine explainers in population genetics: many evolutionary changes (namely, those to which drift made or makes no difference) can be understood deterministically. Population geneticists do not use the deterministic theory as a mere pedagogical or calculating device, then; they use it to construct models that are explanatory in their own right.⁸

On to the second strategy for dealing with the breakdown at the limit.

8. For example, Gillespie (2004), p. 1 promises that the contents of his first chapter, which is based entirely on deterministic models, will provide “true insights”. Anyone who has looked knows, however, that scientists very rarely comment, in uncontroversial cases, on which aspects of their theories and models are explanatory. Indeed, the term “explanation” turns up only occasionally in the official scientific literature, although many scientists will be happy to tell you in their popular writing and elsewhere that explaining the world is science’s ultimate aim—as in Weinberg (1992)’s Chapter Two, with its insistent chain of *whys*.

Limiting frequencies are well defined only relative to an ordering. Why not, in that case, build an ordering into the model? Treat the organisms or instances of genes in a population genetics model like football players—give them identifying numbers. The frequencies are then defined relative to the ordering implicit in the assignment of numbers. This numbering scheme does not, of course, represent any biologically relevant properties. It may not represent any real property at all. The point is solely to give the model sufficient structure to sustain limiting as well as relative frequencies.

As far as I can see, such a strategy is viable. But I do not think that it is the route actually taken by population geneticists. They do not countenance additional structure in their models; nor do they supply, or think that they need to supply, the mathematics that will take advantage of such structure to provide the “strong law” conclusion that evolution will with probability one match expected frequencies.⁹

The first two interpretations of the infinite population idealization, then, are inadequate. If my reading is correct, then population geneticists such as Gillespie and Edwards treat deterministic models as genuine models that can be used (when circumstances are right) to explain real events. Further, many of them—Gillespie, for example—take seriously the idea that such models chart the course of evolution in systems with infinitely large populations. But they have no interest in putting into the models the mathematical structure that is

9. There is of course something absurd about the numbering scheme. John Roberts suggested to me a less outré, though more complex, way to represent limiting frequencies in an infinite population: imagine the ecosystem as consisting of infinitely many tiles each containing a finite population. If the relative frequencies on each tile are the same, then however the tiles are ordered, the limiting frequency will be the same and equal to the relative frequency. One might feel obliged, however, to model the population dynamics of each tile separately, which would result in each tile having a stochastic dynamics and thus in different tiles exhibiting different relative frequencies at any given time. How to aggregate these relative frequencies to obtain the overall expected frequency? A determinate answer to that question requires an ordering of tiles, just as the original question required an ordering of organisms. It is a delicate matter—and one concerning which population geneticists have no interest whatsoever.

required to represent infinite populations by way of a simple idealization.

We must therefore embrace the third option presented above: the infinite population idealization is a genuine idealization, but not a simple idealization. It is an asymptotic idealization.

5. Infinite Populations as Asymptotic Extrapolations

I now present an account of asymptotic idealization—conceived as a practice that sits alongside, but is distinct from, simple idealization—using the infinite population idealization as a paradigm.

Simple idealization is achieved by a simple act of parameter substitution: the scientist takes the veridical model and alters the value of one or more parameters. Asymptotic idealization is achieved by a more complex process that I will call *asymptotic extrapolation*. Like simple idealization, the process notionally begins with a veridical model, or at any rate a relatively veridical model such as, in population genetics, a stochastic model, and performs an operation that produces the idealized model.

I aim to satisfy several desiderata. First—of course—*asymptotic extrapolation* applied to a stochastic population genetics model should be capable of producing a deterministic model without violating any logical or mathematical precepts. Second, *asymptotic extrapolation* should achieve something that simple idealization cannot; in other words, it should not be merely an overly complicated way of implementing simple idealization. Third, *asymptotic extrapolation* should be applicable in a wide variety of cases across the sciences where philosophers assert the existence of “*asymptotic idealization*” or “*asymptotic explanation*”. Fourth, the account of *asymptotic extrapolation* ought to be compatible with just about any philosophical approach to explanation, idealization, or the nature of scientific models. And fifth, the operation should more or less capture the intuitive line of thought, associated with Figure 1, that leads population geneticists to say that in infinite populations there is no drift.

To expand on this last desideratum, asymptotic extrapolation ought to be regarded as a rational reconstruction of what is going through population geneticists' minds when they advance the infinite population idealization, much as Hempel's deductive-nomological model was supposed to be a rational reconstruction of what is going through scientists' minds when they advance scientific explanations in general. I am not saying, then, that scientists would spell out the process of asymptotic extrapolation if they were asked what asymptotic idealization is all about, or that every element of asymptotic extrapolation passes through their heads explicitly when they produce or make use of such idealizations. I do claim, in the Hempelian spirit, that my characterization of asymptotic extrapolation both rationalizes and captures key features of asymptotic idealization as it is manifested in scientific practice.

To perform an asymptotic extrapolation on a stochastic population genetics model with respect to population size n —the “extrapolation parameter”—take the original model and assign a fixed value to n (say, 10). This yields a new model in which n cannot change. In this new model, delete any reference to n itself while preserving the model's ability to predict the behavior relevant to the kind of explanandum you have in mind, that is, some change in organism or gene frequencies. The obvious way to do this is to replace the model's representations of absolute numbers with representations of relative frequencies. The new model will represent the way that relative frequencies change, then, in a system whose population has a certain fixed size.¹⁰ It will achieve this by using as dynamic principles the kinds of probability distribution shown in Figure 1, which given the relative frequencies in one generation yield the probability for any given relative frequency in the next, assuming a certain population size.

You can do this (notionally) for any value of n . Call the resulting models *extrapolation models*. They are put to work to create an idealized model

10. In many systems the population size does not naturally stay fixed. The derived model therefore cannot model such systems, but no matter: that's not what it's for.

using asymptotic reasoning as follows. Construct a sequence of extrapolation models for increasing values of n . Given certain further assumptions, these structures will, as n increases without bound, converge on a limiting structure. In the case of the infinite population idealization, the limiting structure is a probability distribution that, for each possible present-generation frequency, assigns probability one to the expected next-generation frequency and (therefore) probability zero to every other frequency. Make the obvious identification: the limiting structure is the idealized model, said to represent the case where n is infinitely large. It is by way of this extrapolation process, then, that I claim asymptotic idealization is realized.

The process of idealization through extrapolation avoids the difficulties of simple idealization because it does not require the existence of a stochastic model with an infinitely large population. It uses stochastic models with finite populations to define extrapolation models in which population is not explicitly represented, from which it then extracts the idealized model. The structure of the idealized model is therefore determined solely by stochastic models with *finite* populations (Figure 2), and it produces what is recognizably

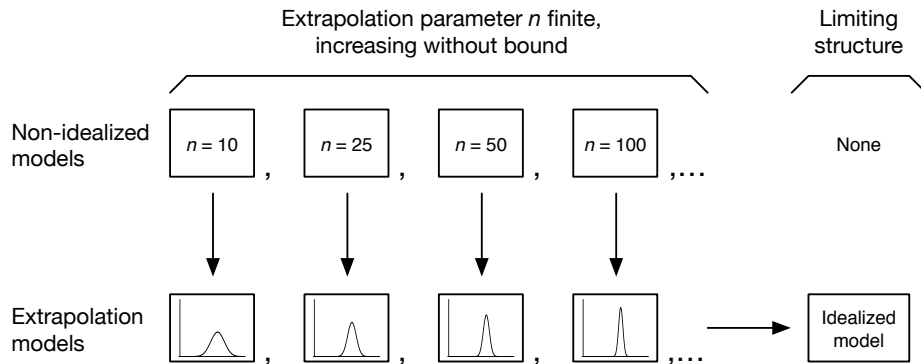


Figure 2: Asymptotic extrapolation. The idealized model is determined solely by (relatively) non-idealized models with finite values for the population size n . The construction does not require anything to occupy the top right corner—it does not require a stochastic model with infinite population size.

the core of the deterministic models: a probability-one transition function from the frequencies in one generation to the frequencies in the next.

Although the extrapolation models contain less information about the target system than the stochastic models from which they are drawn, they nevertheless specify a complete dynamics of gene frequencies. Consequently, the idealized model also specifies a complete gene frequency dynamics, a set of quasi-deterministic rules telling you, for any initial frequency, what the frequency in the next generation will be, and thus by extension what the frequency will be in any future generation.

Further, the same limiting operation can be carried out for a stochastic dynamics that is considerably more interesting than the very simple evolutionary dynamics described above. The generation-to-generation probability distributions can incorporate frequency-dependent selection; they can be sensitive to the age structure of the population; they can trace the complexities of diploid genetics. All of these sophistications will be passed on, through asymptotic extrapolation, to the idealized deterministic model. That is why the properties of population genetics' deterministic models are rich enough to fill chapters and even entire books.

Presented with even a moderate degree of philosophical rigor, asymptotic extrapolation may at first seem rather far from the casual glosses that population geneticists provide when they appeal to infinite populations to rationalize the use of deterministic models. Look past the formality, however, and what's there corresponds very closely to population geneticists' actual, very straightforward reasoning from infinite populations to evolutionary determinism. It captures rather well, in other words, what's meant when an expositor points to something like Figure 1 and says: look, as the population increases, drift's contribution to frequency dynamics decreases, and in the limit as the population goes infinite, drift disappears altogether and you get fully deterministic change in frequencies.

I foresee two objections from the skeptic about asymptotic idealization:

that the operation of taking the limit is not so easily defined with complex dynamic structures as it is with simple properties, and (essentially Norton's complaint) that my idealized model does not, whatever the population geneticists say, contain anything that can be identified as an infinite population. These concerns have merit; I will use them to develop further aspects of the account of asymptotic extrapolation.

First, it is quite correct to observe that the sequence of extrapolation models is not sufficient in itself to determine a limiting structure and so an idealized model. Also required is what you might call an extrapolation space, a mathematical framework in which to define and compute the limit. The space provides the means to represent both the extrapolation models and the limiting structure, and it provides a metric—in the infinite population idealization, a measure of similarity between probability distributions—with respect to which the limit is defined. Population geneticists may not concern themselves with these formalities when they point to Figure 1 and say: “See how it converges on the expected frequency”, but they implicitly assume the obvious extrapolation space, namely, the space of all probability densities over frequency together with an intuitive measure of similarity. (In this case, the mathematical details of the similarity measure do not much matter.)¹¹

The second objection concerns the notion of an infinite population itself. Asymptotic extrapolation works because the infinite parameter, population size, does not appear explicitly in an extrapolation model, which spells out frequencies but not absolute numbers of genes or organisms. For precisely that reason, however, nothing in the idealized model straightforwardly proclaims

11. What counts as an appropriate metric when constructing the extrapolation space? That is determined by the target phenomenon, the behavior that you ultimately wish to predict or explain. A good metric is one that picks out limiting structures such that the target phenomena predicted by the terms in an extrapolation sequence converge on the target phenomenon predicted by the sequence's limiting structure. What metric to use, then, to evaluate the convergence of the target phenomena? That depends on your theoretical aims—for example, on what you consider to be a prediction that is close to the actual behavior—but appropriate choices are usually quite obvious.

itself to be an infinite population—there is nothing that is identified as a population of size n , with n infinitely large. The idealizer declares the limiting structure to be a model of the behavior of infinite populations, but there is no intrinsic structure in the model to represent the infinitude. (This is, I should emphasize, not an idiosyncratic consequence of my own particular conception of asymptotic extrapolation. The end product of the infinite population idealization, everyone agrees, is the sort of deterministic model described above, and it is quite clear that such a model, though it represents frequencies in a population, does not explicitly represent the population as having any particular size.)

The question, then: in virtue of what ought we to regard such models as representing the behavior of infinite populations? I propose a conventionalist answer: the assumption of infinitude is simply built in to the practice of asymptotic idealization. Whenever a model is constructed by asymptotic extrapolation, then, it is taken by fiat to represent the behavior of systems in which the extrapolation parameter—the quantity that is allowed to tend to infinity in the course of the extrapolation—takes on an infinite value. This is hardly an arbitrary stipulation. The process of asymptotic extrapolation is in some intuitive sense clearly designed to investigate the behavior of systems in which the parameter is infinitely large. But it is a stipulation nonetheless.

A related question: Are the frequencies in the idealized model to be understood as relative frequencies, limiting frequencies, or something else? There is no philosophically or empirically significant fact of the matter. They are just the frequencies in “population geneticists’ infinite populations”. Their real-world explanatory relevance lies entirely in their ability to represent relative frequencies in actual populations. As to their intrinsic mathematical nature there is nothing further to say.

I see no technical difficulties with the convention that asymptotically idealized models represent systems in which the extrapolation parameter goes infinite. It is, nevertheless, possible to imagine a Nortonizing voice

chiding: What is the point of such a convention? Why insist that population geneticists' deterministic models falsely assert the population's infinitude, rather than taking the alternative view that they say nothing about population size, or (a very similar view) that they represent populations of large but finite size?

One rationale for insisting on the link to infinitude is simply to do justice to the way that population geneticists such as Gillespie talk. My project all along has been to construct a translation manual for this way of talking, and in particular for understanding what population geneticists mean when they say that deterministic models represent the dynamics of infinite populations. Nothing about this project demands that the deterministic models must be interpreted in this way; the point is rather to clear space for the possibility of such an interpretation in a way that avoids imputing, as Norton does, an embarrassing mistake to the scientists in question.

Nevertheless, it would be illuminating to show that the interpretation has some sort of practical or theoretical payoff, that is, that understanding the deterministic models as representing infinite populations serves a useful function in scientific thinking. That I will now try to do.

I must as a consequence reopen the question of the function of idealization in explanatory models, which I have until now put aside. One goal of idealization, mentioned in section 2 above, is to simplify modeling by making calculation or calibration easier. That goal is, however, realized by a deterministic population genetics model regardless of whether it is understood as representing populations as infinite. So it cannot motivate the infinitary interpretation.

My own favored account of the primary function of idealization is more helpful in this regard. Idealization, I suggest, is a rhetorical strategy to underline a factor's explanatory irrelevance. (The view was sketched in section 2; see also Strevens (2008), chap. 8 and Strevens (2017). I equate explanatory relevance with "difference-making", but that aspect of the view is not needed

here.) Rather than printing the factor in red ink, or putting a big 'X' through its name, an idealization blatantly falsifies the way things are with respect to that factor. The falsehood, properly interpreted, says of the factor which it misrepresents: this factor is not a part of the explanation. Stating that the air does not resist a cannonball is a way of saying, indirectly, that air resistance is irrelevant to the ball's approximate trajectory (a claim whose correctness depends, of course, on what properties of the trajectory are under discussion). Stating that economic actors are perfectly rational is a way of saying that the many humanly common departures from narrow instrumental rationality are irrelevant to the economic phenomenon that is to be explained. Stating that there is no genetic drift is a way of saying that drift is irrelevant to the outcome of the evolutionary explanandum—to the fixation of an allele, say.

Mere fiction is, however, not always sufficient to alert the scientific community to an idealization—the falsehood might, after all, be an inadvertent error. There is a certain conventional code, therefore, for indicating deliberate idealizations: where possible, set the parameter in question, or some appropriately related parameter, to zero or infinity. Zero air resistance, zero irrationality, zero drift, zero intermolecular force, zero wavelength: scientists know how to read these round numbers as asserting that air resistance, irrationality, drift, forces, and wavelike behavior are predictively and explanatorily irrelevant. (The code is not infallible, of course; as with any system of representation, messages may be misinterpreted or seen where there are none.)

To conform to the convention is not always possible; it is therefore not obligatory. But it is in the interests of scientific communication to make compliance as easy as possible—thus, to find a way, in as many cases of idealization as possible, to set a parameter to an extreme value. That is the purpose, where such ascriptions cannot be accomplished by simple idealization, of what I am calling asymptotic extrapolation. It gives scientists additional avenues for saying “Look; I have set something to zero or infinity” (or “I have set something to zero by setting something else to infinity”, as with drift and

population)—and so additional avenues for saying, loudly and clearly, that something is not relevant.

In the absence of this extreme-value protocol, scientists might perhaps follow Norton in renouncing talk of infinite idealizations in cases where infinities cannot be inserted directly into models by way of simple idealization. But the protocol is real, and so there is real value in enlarging the scope of infinity talk as broadly as possible. That is the rationale for understanding deterministic models in population genetics not merely as good approximations of the behavior of actual populations, but as good approximations that are also models in which populations are infinitely large.

6. Asymptotic Extrapolation Generalized

Let me characterize asymptotic extrapolation in more general terms, providing a widely applicable method for understanding scientists who say that some idealized model represents the limiting behavior of the real system as a certain parameter goes to infinity. Such claims, I hold, can be rationally reconstructed as asserting that the model in question is derived from a more realistic model by way of asymptotic extrapolation.

Asymptotic extrapolation from a realistic model requires:

1. An extrapolation parameter, that is, a parameter in the model which in the intended idealization goes to zero or infinity or some other extreme value,
2. A template for an extrapolation model, which is derived from the realistic model by assigning a fixed value to the extrapolation parameter and removing all representation of that parameter from the model while retaining the model's ability to represent behavior relevant to the explanatory task, and
3. An extrapolation space, which provides the mathematical structure for

finding the limiting form of the extrapolation models as the extrapolation parameter tends to the extreme value.

The idealized model is the limiting form of the extrapolation models. It is said to model the behavior of the system with respect to the target phenomenon when the extrapolation parameter takes on its extreme value (when population size is infinite, when wavelength is zero, and so on)—even if it is impossible to set the parameter to that value in the realistic model, and even though the idealized model does not represent the infinitude explicitly.

Asymptotic extrapolation looks, then, at a certain behavior of a feature of the realistic model as some parameter goes to an extreme value. In those numerous cases across the sciences where the model itself falls apart at the limit, the behavior or feature derived from the model often takes on a distinctive limiting form, and this form can be used as the basis for an idealized model. Such a model will be useful for explaining a phenomenon whenever the features of the system represented by or dependent on the extrapolation parameter are not relevant to explaining the phenomenon.

Appropriate choices of an extrapolation model and of a metric for the extrapolation space may vary significantly with the phenomenon to be explained. Thus extrapolation typically does not produce general theories, but rather models that are targeted at a restricted range of phenomena.

Population genetics provides some further simple examples of asymptotic extrapolation: the “infinite alleles” and “infinite sites” idealizations (Gillespie 2004, §§2.3, 2.4). They work in similar ways; let me consider the former.

Occasionally a mutation will modify an allele. Typically, the modification—a change in one or a few nucleotides in a long string of DNA—will create an entirely new allele. There is a small chance, however, that a mutation could convert an existing allele into another allele that is already represented in the population, or that has been represented in the population in the past. Population geneticists normally “idealize away” this probability, setting it to zero. They call their idealization the “infinite alleles” model, in effect

supposing that there are infinitely many possible alleles at the relevant locus (rather than a very large finite number), and so that the chance of mutation creating any one of the finite number currently or previously represented in the gene pool is zero.

That all seems straightforward, but even here the move to infinitude creates a problem. The fundamental stochastic assumption of the (relatively) realistic model is that a mutation at a genetic locus is equally likely to create any of that locus's possible alleles.¹² Assume there are infinitely many alleles, then, and you are positing a uniform probability distribution over a countably infinite set of outcomes. There are considerable mathematical difficulties in constructing such a distribution. No such thing is possible, some would say. At the very least, the axiom of countable additivity must be abandoned (an axiom that is necessary, by the way, to derive the strong law of large numbers).

Yet population geneticists no more wrestle with this difficulty than they worry about the order-relativity of limiting frequencies. They have some way of making sense of the “infinity” talk that requires no high-level mathematical maneuvers.

I suggest that the infinite alleles model is another elementary example of asymptotic extrapolation. In this case, an extrapolation model is created by substituting for the probability distribution over alleles as mutational endpoints a simpler probability distribution, namely, the distribution giving the probability that a new mutation will produce an allele that currently exists or has previously existed in the population—call it the probability of mutational replication. The rest of the dynamic structure in the realistic model is left unchanged, and so is incorporated in its entirety into the extrapolation model.

The extrapolation parameter is the number of possible alleles. An extrapo-

12. In this regard, the realistic model is already idealized. More realistic models take into account the molecular realization of the genetic material, and replace the infinite alleles idealization with the infinite sites idealization.

lation model does not represent this number, but contains only the replication probability that it entails. Otherwise, the model is identical to the realistic model.

As long as the number of possible alleles n at the relevant locus is finite, the mutational replication probability is non-zero. Using the realistic model's assumption of uniform probability, a sequence of replication probabilities for increasing values of n can be constructed; these probabilities can then be incorporated into a sequence of extrapolation models containing the rest of the dynamic structure of the realistic models for each value of n . The limiting value of the sequence of replication probabilities is of course zero; thus, the extrapolation models' limit is a model in which the probability of mutational replication is zero. And that is why population geneticists say that when there are infinite alleles, the replication probability is zero.

The mathematics of the extrapolation is elementary, but that is as it should be: it is supposed to reflect roughly what is going on in the minds of the builders of the infinite alleles model, and all the evidence suggests that what is going on is mathematically quite straightforward.

7. The Ray Theory of Light

This section is omitted from the published version of this article.

Let me now return to the place where it all began: the relationship between the wave and ray theories of light. It is often said that the ray theory approximates the wave theory in cases where the wavelength is much smaller than the objects (prisms, raindrops, occluders) that affect the waves, or more strongly and more asymptotically, that the wave theory reduces to the ray theory as the wavelength goes to zero, or perhaps more perspicuously, as the ratio of the wavelength to the size of the relevant objects goes to zero. But as Berry (1994) and Batterman argue, the situation is not so straightforward, because wavelength in the wave theory cannot be set to zero, either simply or asymptotically, without creating a breakdown in the wave-theoretical structure.

Ray-theoretic models of specific optical phenomena can nevertheless be derived using asymptotic extrapolation. The structure of a rainbow is one example, important in both Berry's and Batterman's thinking.

Consider a monochromatic rainbow, that is, a rainbow produced by a single frequency of light. It lacks the characteristic colorfulness of a rainbow, but it has the same structure as each of a normal rainbow's bands: the top of the bow is sharply defined, bright below and dark above; the intensity of the light then falls off below the bow, eventually fading away entirely.

These features can be explained by the ray model of light, from which can be derived a formula for the intensity of light in the rainbow having the form shown in Figure 3. In the figure, the vertical axis marks the top of the rainbow's arc. To the right of axis the intensity is zero, representing complete darkness above the arc. The intensity at the axis is infinite, representing a caustic, or line of infinite brightness, at the top of the arc. To the left of the axis the intensity then diminishes, representing a slow fading away of the rainbow below the caustic.

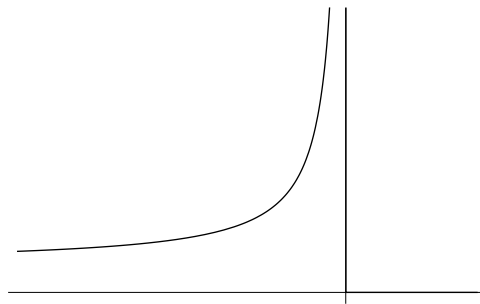


Figure 3: Light intensity in a monochromatic rainbow according to the ray theory

A real rainbow, monochromatic or otherwise, does not have these features: the dark band (Alexander's band) above the rainbow is not completely lightless; the upper boundary of the rainbow's arc is not infinitely bright; the brightness of the rainbow falls off rather suddenly a certain distance below the top;

and “supernumerary bands” exist after this drop-off, that is, additional bands appear below the rainbow (typically pastel in appearance and difficult to see when conditions are less than perfect). The ray theory does not explain features such as supernumerary bands; nor, of course, does it explain the features that it predicts but that are not there, such as complete darkness above the bow or the smooth fading away below. What it explains are the features that it predicts and that are there: a sharp transition in brightness at the top of the rainbow and a fading away below. Call these the *ray features*. (There are other even more important ray features not represented in Figure 3, notably the shape of the rainbow’s arc and its perceived position in the sky relative to the viewer and the sun.)

In the first half of the nineteenth century, the additional features of the rainbow not predicted by the ray theory were explained by the new wave theory of light. Particularly important was the work of George Airy, who derived a formula for the intensity of the rainbow’s light that conformed much more closely than the ray theory to what is observed (Figure 4).

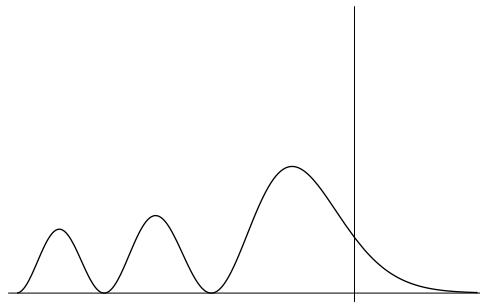


Figure 4: Light intensity in a monochromatic rainbow according to Airy’s wave-theoretic model. The vertical axis marks, as in Figure 3, the line of maximum intensity predicted by the ray theory—though not, as you can see, by Airy’s theory.

Airy’s theory predicts a number of features that the ray theory does not: that there is some illumination above the rainbow; that at the top of the rainbow (now predicted to occur at a slightly different position) the light is

not infinitely intense; that at a certain distance below the top the intensity falls off quickly; but that there are then multiple bands of light that show up below this point, themselves decreasing in intensity with distance from the top of the rainbow. (Two such bands are shown in the figure.)

What happens to Airy's intensity curve as the wavelength decreases? Figure 5 shows the curve from the previous figure superimposed on a curve for light of one-tenth of the wavelength. As you can see, the rate of oscillation

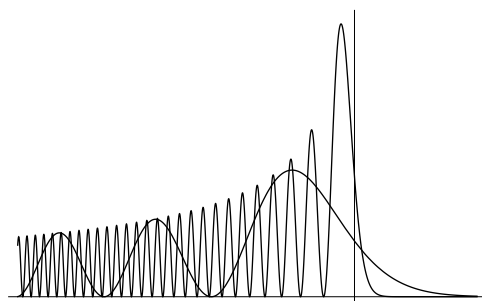


Figure 5: The Airy intensity for two wavelengths, one ten times shorter than the other

of the intensity curve increases by a factor of ten, and the intensity near the top of the rainbow increases by a factor of just over two (to be precise, the cube root of ten). That gives you a sense of what happens in the limit as wavelength goes to zero: the intensity at the top of the rainbow goes to infinity and the intensity curve oscillates infinitely many times in any finite interval.

On the one hand, this limiting behavior is quite unlike the ray-theoretic intensity shown in Figure 3: the ray intensity falls off smoothly while the wave intensity oscillates more and more madly as the wavelength gets closer to zero. On the other hand, the envelope of this mad oscillation traces a curve of the same shape as the ray-theoretic intensity curve. (It is not exactly the same curve; here it should be remembered that the Airy intensity is itself only an approximation).

It is for this reason that optical physicists are (some of them) inclined to say that both the ray model, on the one hand, and the wave model with the

wavelength in some sense asymptotically set to zero, on the other, can in some sense be used to give the same explanation of the rainbow's ray features.

It remains to spell out what the two occurrences of “in some sense” amount to. I suggest that what is happening hinges on an asymptotic extrapolation. It is not the case that the ray model is derived by extrapolation from the wave model. Rather, what is derived from the wave model—by letting the wavelength go to zero—is a much simpler model containing only an intensity curve. This intensity curve is identical to the ray intensity curve. Thus it is said that the wave model offers the same explanation as the ray model as the wavelength goes to zero.

The extrapolation in question proceeds as follows. The extrapolation parameter is of course wavelength. The extrapolation model is, as I have said, the intensity function. (Or perhaps some more complex model containing the intensity function; it does not, however, contain rays or waves themselves.) And, crucially, the extrapolation space uses a metric for taking the limit that cares only about the overall shape traced by a function but not about the way that the shape is filled in—thus, only about the “envelope” of the small-wavelength Airy function. To put it another way, the metric “blurs” functions before it compares them, counting two intensity functions as close if they have almost identical average intensities over small intervals. Relative to such a metric, the extrapolation model will converge, as the wavelength goes to zero, on an intensity function having the qualitative structure of the ray-theoretic intensity function.¹³ Thus, such a function will qualify as an asymptotic extrapolation of the wave model, giving a literal sense to the claim that the ray theory and the asymptotic idealization of the wave model give the same explanation of the ray features of rainbows.

* * *

13. It will converge on many other functions as well, all agreeing on the small-interval moving average. Every one of these functions qualifies as an idealization of the wave model as wavelength tends to zero.

What of the ray model itself? I do not think that it is best understood as an asymptotic idealization of the wave model. It is an idealization, but not one that can be perspicuously regarded as obtained from the wave theory in the limit as wavelength goes to zero.

The point of the ray idealization, on my approach to these matters, is to communicate clearly and effectively that wave effects, especially interference and diffusion, do not make a difference to the target phenomenon. But saying “the wavelength is zero” does not communicate very clearly at all that interference is a non-difference-maker, for reasons described by Berry and Batterman and made clear in the rainbow case: interference does not decrease as wavelength decreases. This is not a problem with the ray idealization itself. Clearly rays do not interfere, and so the use of a ray model constitutes an ostentatious putting aside of interference. The obfuscation is in asserting that interference is put aside by reducing wavelength to nothing. (With respect to diffusion, the zero-wavelength extrapolation is, I think, far more effective.)

Some scientists, however, apparently do think of the ray idealization in terms of zero wavelength. Asymptotic extrapolation gives us a way to understand what they are trying to say—even if we then go on to criticize their message.

8. Conclusion

It would be noble and fitting at this point to return to Batterman’s more complex cases of asymptotic idealization in physics, to see whether they can be regarded as asymptotic extrapolations. That is, however, more than can be comfortably fitted into a single paper. So I leave you instead with the following thoughts.

First, asymptotic idealization is pervasive in certain parts of biology as well as in physics. Second, it can be casual and mathematically straightforward, but at the same time distinct from simple idealization. Third, even in cases where the realistic model disintegrates in the limit, an asymptotic extrapolation may

produce a rich, dynamic, and idealized model of the target system's behavior. Fourth (and at this point I introduce my own ideas about the function of idealization), to regard such models as describing infinite systems has an important communicative role to play in the practice of identifying predictive and explanatory irrelevancies by way of idealization. So we can, I hope, find some illumination in Batterman's insights while respecting Norton's strictures and doing justice to the asymptotic idiom as it is found in actual scientific idealization.

References

- Abrams, M. (2006). Infinite populations and counterfactual frequencies in evolutionary theory. *Studies in History and Philosophy of Biological and Biomedical Sciences* 37:256–268.
- Batterman, R. W. (2002). *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*. Oxford University Press, Oxford.
- . (2010). On the explanatory role of mathematics in empirical science. *British Journal for the Philosophy of Science* 61:1–25.
- Belot, G. (2005). Whose devil? Which details? *Philosophy of Science* 72:128–153.
- Berry, M. (1994). Asymptotics, singularities, and the reduction of theories. In D. Prawitz, B. Skyrms, and D. Westerståhl (eds.), *Logic, Methodology and Philosophy of Science IX: Proceedings of the Ninth International Congress for Logic, Methodology and Philosophy of Science, Uppsala, 1991*. Elsevier, Amsterdam.
- Edwards, A. W. F. (1977). *Foundations of Mathematical Genetics*. Cambridge University Press, Cambridge.
- Elgin, C. (2007). Understanding and the facts. *Philosophical Studies* 132:33–42.
- Gillespie, J. H. (2004). *Population Genetics: A Concise Guide*. Second edition. Johns Hopkins University Press, Baltimore, MD.
- McMullin, E. (1985). Galilean idealization. *Studies in History and Philosophy of Science* 16:247–273.
- Norton, J. D. (2012). Approximation and idealization: Why the difference matters. *Philosophy of Science* 79:207–232.

- Pincock, C. (2012). *Mathematics and Scientific Representation*. Oxford University Press, Oxford.
- Potochnik, A. (2017). *Idealization and the Aims of Science*. University of Chicago Press, Chicago.
- Ramsey, G. (2013). Driftability. *Synthese* 190:3909–3928.
- Strevens, M. (2008). *Depth: An Account of Scientific Explanation*. Harvard University Press, Cambridge, MA.
- . (2017). How idealizations provide understanding. In S. R. Grimm, C. Baumberger, and S. Ammon (eds.), *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science*. Routledge, New York.
- Wayne, A. (2012). Emergence and singular limits. *Synthese* 184:341–356.
- Weinberg, S. (1992). *Dreams of a Final Theory*. Pantheon, New York.
- Weisberg, M. (2007). Three kinds of idealization. *Journal of Philosophy* 104:639–659.
- . (2013). *Simulation and Similarity: Using Models to Understand the World*. Oxford University Press, Oxford.