Figure 3.3
Two styles of self-service restaurant, (A) the sushi conveyor and (B) the buffet (photo credits: see Acknowledgments).

begin with a practically inexhaustible abundance of static choices in full view, and fill our plate with the desired amount of each dish, often choosing many different items; in the latter, we iteratively make a yes-or-no choice of the chef's latest creation as it passes by, typically accepting only a few choices and rejecting the vast majority. In either case, we exercise choice, and we may end up with a satisfying meal.

The thinking of molecular evolutionists frequently corresponds to the sushi conveyor model, which illustrates a proposal-acceptance process like the one described in "Climbing Mount Probable." We choose (we select), but we don't control what is offered or when: instead, we accept or reject each dish that passes by our table. Though we decide whether each dish is right for us, initiative and creativity belong largely to the chef.

By contrast, the architects of the Modern Synthesis were committed to the buffet view, in the form of a totipotent "gene pool" with sufficient variation to respond to any challenge (addressed in more detail below). Just as the staff who tend the buffet will keep it stocked with a variety of choices sufficient to satisfy every customer, the gene pool is said to *maintain* abundant variation, so that selection never has to wait for a new mutation. Adaptation happens when the customer gets hungry and proceeds to select a platter of food from the abundance of choices, each one ready at hand, choosing just the right amount of each ingredient to make a well-balanced meal.

A bias in the choices offered to customers (analogous to a bias in variation) will have different effects in the two regimes. Let us suppose that the buffet has five apple pies and one cherry pie. This quantitative bias makes no difference. A rational customer who prefers cherry pie is unaffected by relative amounts and will choose a slice of cherry pie every time. The only kinds of biases in supply that are relevant at the buffet are absolute constraints—the complete lack of some possible dish—that is, the customer who prefers cherry pie will end up with apple pie only if there are no cherry pies.

But at the sushi conveyor, the effect of a bias will be different. Let us suppose that the dishes of sushi on the conveyor show a 5 to 1 ratio of salmon to tuna. Even a customer who would prefer tuna in a side-by-side comparison may eat salmon more often, because a side-by-side comparison simply is not part of the process.

How does this difference in analogies map to a difference in regimes of evolutionary genetics? We have already seen sushi conveyor dynamics (figure 3.2). The question is whether the buffet regime exists, and if so, what are its defining characteristics? The key condition in a hypothetical buffet regime would be that the variants relevant to the outcome of evolution are abundantly present in the initial gene pool.

Therefore, let us consider what is the effect of initial variation in the model presented earlier. The result, shown in figure 3.4, is simple and compelling. Here we have simulated the model exactly as before. The upper line repeats the results (from figure 3.2) of evolution from a pure starting population, and the lower lines represent cases in which both alternative types are present in the initial population at non-zero frequencies (Yampolsky and

behavior is readily understood with origin-fixation dynamics, yielding Eqn 2. In the limiting case as $N \to \infty$ with u fixed, i.e., a deterministic case of shifting frequencies with all types present immediately, we may consider a set of coupled differential equations whose behavior dictates that the fittest type asymptotically approaches a frequency near 1. This behavior corresponds to the buffet regime. In the intermediate regime where concurrent mutation is possible (as per Weissman et al. 2009; Desai and Fisher 2007), the problem has not been solved, yet the simulation results in the preceding section indicate clearly that biases in mutation are important.

Distinctive Implications

The results above reveal a kind of cause–effect relationship that seems very fundamental but is absent from standard treatments of evolutionary causation. Stated more precisely, it is the cause–effect relationship by which biases in the introduction of variation have a difference-making power characterized by the following statements:

- **It applies when fixations are selective.** Hypotheses in which mutation bias shapes features, though not allowed in the original Modern Synthesis (OMS), quickly became accepted by molecular evolutionists as implications of neutrality, under the assumption that mutation bias can be effective only when selection is absent. By contrast, biases in the introduction of variation do not require neutrality, leading to the novel prediction of mutation-biased adaptation.

- **It poses a directional, quantitative dependence.** In conventional thinking, variation is only a material cause, a source of substance but not form, and so the obvious questions to ask about variation concern the amount of raw material available. By contrast, the cause–effect relationship proposed here identifies quantifiable *directions* of variation that may influence directions of evolution and proposes idealized conditions under which the effect of a mutational bias of magnitude B is a B-fold bias on evolution.

- **It is regime-dependent.** The maximal effect just mentioned is fully realized in the sushi conveyor regime but is negligible in the buffet regime, where the variants relevant to the outcome of evolution are abundantly present. By contrast, conventional thinking assigns *fixed roles* to mutation and selection, independent of population-genetic regime.

- **It depends on the rareness of mutations.** In this kind of causation, the influence of biases in the generation of variation emerges as $uN \to 0$. By contrast, conventional arguments based on the force of mutation assume that mutation rates must be large in magnitude for mutational effects to be important in evolution.

- **It establishes a condition of parity with selection.** Under ideal conditions, the effect of a bias in the introduction of new alleles is the same in magnitude as that of a bias in fixation probabilities of equal magnitude. By contrast, conventional thinking treats variation as a different *kind* of cause that is passive and subordinate to selection.

set of A:T → C:G transversions, while resistant cultures from the mutH parent tended to adapt by another small set of G:C → A:T and A:T → G:C transitions.

Meyer et al. (2012) reported changes in the J gene of bacteriophage Lambda in 48 replicate cultures of *E. coli*, half of which had acquired the ability to utilize the OmpF receptor as an attachment site. The complete set of 241 differences (all non-synonymous) from the parental J gene in 48 replicates is shown in figure 3.5. Among these non-synonymous mutations, 22 are found at least twice (asterisks), including 16 transitions observed in adapted strains 181 times, and 6 transversions observed 42 times. Thus the transition:transversion ratio is 16/6 = 2.7 when we count paths (i.e., columns in figure 3.5), and 181/42 = 4.3 when we count events (cells in figure 3.5). This is far above the null expectation of 0.5 under an absence of mutation bias.

To explore the role of mutational biases in parallel adaptation more systematically, Stoltzfus and McCandlish (2017) gathered data for a set of experimental cases such as Meyer et al. (2012), and for a comparable set of natural cases of parallel adaptation. They used these data to test for an effect of transition:transversion bias, a widespread kind of mutation bias. In the dataset of 389 parallel events along 63 paths from experimental studies of evolution, they find a highly significant tendency—from 4-fold to 7-fold in excess of null expectations—for adaptive changes to occur by transition mutations rather than transversion mutations. For the dataset of natural cases of parallel adaptation consisting of 231 parallel events along 55 paths, they found a bias of 2-fold to 3-fold over null expectations, which was statistically significant for both paths and events. They conclude that parallel adaptation takes place by nucleotide substitutions that are favored by mutation, noting that the size of this effect is not a small shift, but a substantial effect of 2-fold or more.

This analysis was made possible by prior work that used empirical data on the effects of mutations to show that transitions and transversions do not differ importantly in their distributions of fitness effects (Stoltzfus and Norris 2016). This was important to establish because the lore in molecular evolution held that amino acid changes tend to occur by transitions because they are more conservative in their effects (Keller, Bensasson, and Nichols 2007; Rosenberg, Subramanian, and Kumar 2003; Wakeley 1996). Thus, taking the work of Stoltzfus and McCandlish (2017) at face value, a several-fold effect previously attributed to selection is now best understood as being entirely or almost entirely due to mutation.

More generally, I take it that the way to establish that a kind of cause is important in evolution is to show that it has quantitatively large effect-sizes in regard to features that evolutionary biologists consider important or interesting. An effect-size of 2-fold is large. In a data-rich field such as molecular evolution, enormous numbers of papers are published every year based on far smaller effect-sizes. The subjectivity of "important or interesting" cannot be avoided, but may be exploited, as in the above example, by stealing explanatory

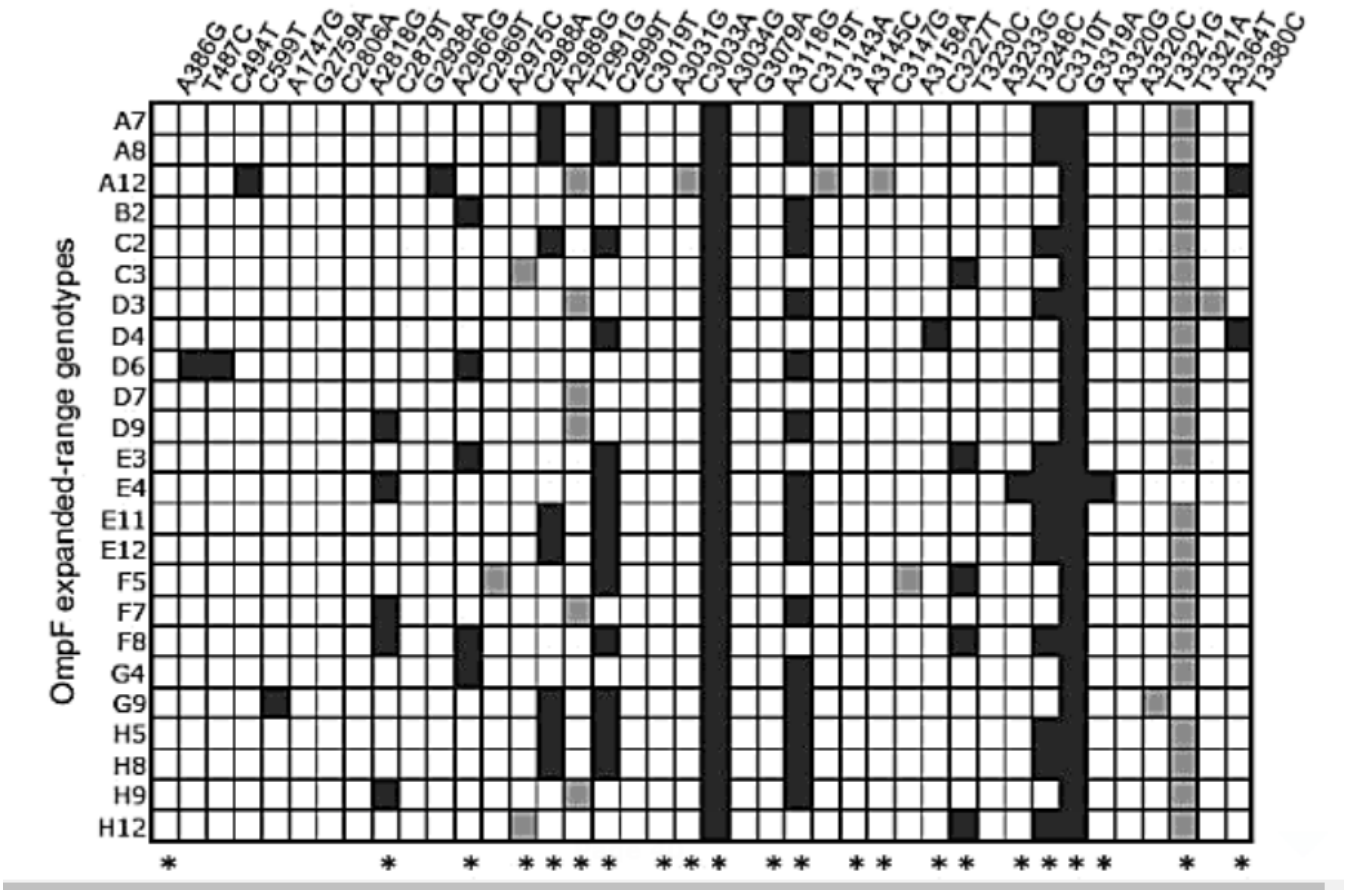
A386G
T487C
C494T
C599T
A1747G
G2759A
C2806A
A2818G
C2879T
G2938A
A2966G
C2969T
A2975C
C2988A
A2989G
T2991G
C2999T
C3019T
A3031G
C3033A
A3034G
G3079A
A3118G
C3119T
T3143A
A3145C
C3147G
A3158A
C3227T
T3230C
A3233G
T3248C
C3310T
G3319A
A3320G
A3320C
T3321G
T3321A
A3364T
T3380C
OmpF expanded-range genotypes
A7
A8
A12
B2
C2
C3
D3
D4
D6
D7
D9
E3
E4
E11
E12
F5
F7
F8
G4
G9
H5
H8
H9
H12

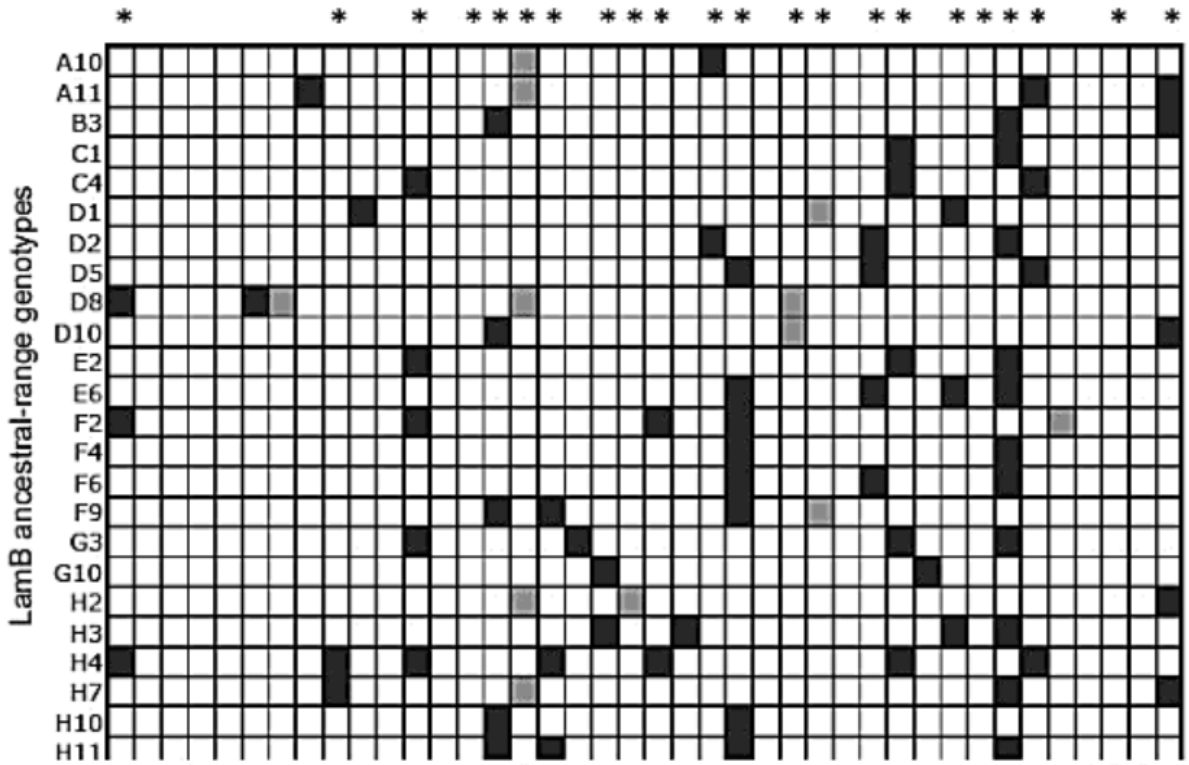

Figure 3.5
Parallel evolutionary changes (asterisks) in the Lambda tail tip gene (after figure 1 of Meyer et al. 2012). Black (transition) and gray (transversion) boxes indicate specific nucleotide mutations (columns) found among 48 replicates (rows).

power from selection: if some evolutionary phenomenon (such as transition:transversion bias) was important enough to elicit a selective explanation, then an alternative mutational explanation is also important.

Note that this evidence does not address all of the distinctive implications outlined above. Specifically, it addresses the claim that an effect of mutation bias is possible under non-neutral conditions. With more data, it will be possible to address a graduated relationship between cause and effect (given that the magnitude of mutation bias varies with taxonomic context), the composition of causes, and the role of population-genetic regime.

Three Larger Implications

The Potential for Directional Trends

The model above addresses what theoreticians would call a 1-step adaptive walk. Is it possible for a sustained bias in the introduction of new variation to cause a sustained trend over many steps? To address this question, one must consider a multistep adaptive walk on some larger landscape.

Precisely this issue is addressed by Stoltzfus (2006), using a model of a protein-coding gene, based on an abstract NK fitness model to represent interactivity of amino acids in a protein. The question addressed in the model is whether a mutational bias toward G and C nucleotides could result in a biased protein composition, even when all of the changes are beneficial. This question is relevant to interpreting observed differences in protein composition that correlate with GC content (Singer and Hickey 2000): organisms with high-GC synonymous sites and inter-genic regions also have proteins enriched for amino acids encoded by high-GC codons (Gly, Ala, Arg, Pro) and depleted for those encoded by low-GC codons (Phe, Tyr, Met, Ile, Asn, Lys). Typically such differences are attributed to mutational biases operating in the context of neutral evolution. However, we can now understand the potential for mutation biases to operate in non-neutral evolution.

The general result of simulating evolution of a protein-coding gene via beneficial changes, subject to biases in the introduction of variants, is to confirm the intuitions of the metaphor of Climbing Mount Probable, showing that it is not misleading in its implications. On a perfectly smooth landscape, the trajectory of evolution is deflected initially by mutation bias, but the ultimate destination—the optimal sequence—is unaffected. On rough landscapes, a mutation-biased composition evolves along a trajectory to a local peak. The rougher the landscape, the shorter the trajectory and the greater the per-step effect of mutation bias. In the model used by Stoltzfus (2006), a modest mutational bias toward GC (or AT), with a magnitude consistent with mutational biases inferred from models of genome composition evolution, can cause a progressive shift in protein composition of a magnitude consistent with observed biases in protein composition.

or slightly deleterious mutations (McCandlish and Stoltzfus 2014). Outside of molecular evolution, theoretical population genetics still relied on the gene-pool assumption. In the 1990s, theoreticians began to notice this restriction; for example, Yedid and Bell (2002) write:

> In the short term, natural selection merely sorts the variation already present in a population, whereas in the longer term genotypes quite different from any that were initially present evolve through the cumulation of new mutations. The first process is described by the mathematical theory of population genetics. However, this theory begins by defining a fixed set of genotypes and cannot provide a satisfactory analysis of the second process because it does not permit any genuinely new type to arise. (p. 810)

Likewise, Hartl and Taubes (1998) write:

> Almost every theoretical model in population genetics can be classified into one of two major types. In one type of model, mutations with stipulated selective effects are assumed to be present in the population as an initial condition. ... The second major type of models does allow mutations to occur at random intervals of time, but the mutations are assumed to be selectively neutral or nearly neutral. (p. 525)

Eshel and Feldman (2001) make a similar distinction:

> We call short-term evolution the process by which natural selection, combined with reproduction (including recombination in the multilocus context), changes the relative frequencies among a fixed set of genotypes, resulting in a stable equilibrium, a cycle, or even chaotic behavior. Long-term evolution is the process of trial and error whereby the mutations that occur are tested, and if successful, invade the population, renewing the process of short-term evolution toward a new stable equilibrium, cycle, or state of chaos. (p. 182)

They argue that,

> Since the time of Fisher, an implicit working assumption in the quantitative study of evolutionary dynamics is that qualitative laws governing long-term evolution can be extrapolated from results obtained for the short-term process. We maintain that this extrapolation is not accurate. The two processes are qualitatively different from each other. (p. 163)

Thus, due to the way that prevailing views shaped approaches to modeling, theoreticians were not prepared to recognize a role for biases in the introduction of variation. A half-century after theoretical population genetics emerged as a discipline, the introduction process began to appear as an unnamed technical feature of certain types of models, though not as a feature of contending theories (which were focused instead on selection, drift, recombination, sex, and so on). It took several more decades for the argument to emerge that recognizing the introduction process is not a minor technical detail, but a major theoretical innovation that challenges previous thinking about how evolution works, and opens new avenues for theoretical and empirical research.

Structure and Implications of the OMS

In the early years of the twentieth century, Johannsen's experiments with beans showed that selection can be effective at sorting out existing varieties, but it does not create new types from masses of environmental fluctuations. That experiment spelled the end of Darwinism among those who embraced genetics (Gayon 1998). In the mutationist view that emerged among early geneticists such as Morgan (1916), "Evolution has taken place by the incorporation into the race of those mutations that are beneficial to the life and reproduction of the organism" (p. 194).

Given this view, one might imagine that the introduction of a new mutation would be a key event that provides initiative for evolutionary change, and thus influences dynamics as well as direction. Following Shull (1936), one might suppose that "If mutations are the material of evolution, as geneticists are convinced they are, it is obvious that evolution may be directed in two general ways: (1) by the occurrence of mutations of certain types, not of others, and (2) by the differential survival of these mutations or their differential spread through the population" (p. 122), and we might suppose that a new allele "produced twice by mutation has twice as good a chance to survive as if produced only once" (p. 140).

We might think this way today, but the architects of the OMS emphatically did not. In their view, Johannsen's experiments showed nothing. Instead, the true nature of evolution was revealed in Castle's experiments with the hooded rat, because unlike Johannsen, "Castle had been able to produce new types by selection" (Provine 1971, 114). Castle was able to shift coat-color from mottled to nearly all black, or nearly all white, in less than 20 generations of selection, not enough time for new mutations to play any important role—that is, *selection can create new types without mutation*.

The genetic interpretation of this result was that selection simultaneously shifts the frequencies of available alleles at many loci, leveraging recombination to combine many small effects in one direction (Provine 1971). Thus, recombination (not mutation) is the proximate source of new genetic variation every generation. Evolution, rather than being a process of the mutational introduction and reproductive sorting of variation, is envisioned as a process of shifting allele frequencies in the gene pool. Even though this mode of change requires abundant standing variation, it prevails in nature (they argued), because natural populations have a "gene pool" that "maintains" variation. Thus, the maintenance of variation in the gene pool, logically necessary to prop up the Castle experiment as the paradigm of evolution, became a major theme, what Gillespie (1998) called "The Great Obsession" of population genetics.

Thus, in the OMS, the term "gene pool" is not merely descriptive but evokes the theory that natural populations *maintain* abundant genetic variation; for example, Stebbins (1966) writes that "a large 'gene pool' of genetically controlled variation exists at all times" (p. 12). This theory, proposed by Chetverikov and popularized by Dobzhansky, holds that various features of genetics and population genetics—including recessivity, chromosome

assortment, crossing over, sexual mixis, frequency-dependent selection, and heterosis—come together to create a dynamic in which variation is soaked up like a "sponge" and "maintained." In the OMS, this gene-pool dynamic ensures that evolution is always a multifactorial process in which selection never waits for a new mutation but can shift the population to a completely new state based on readily available variation.

This gene-pool theory was used to argue against the early geneticists in regard to (1) the source of initiative in evolution (mutation or selection), and (2) rates of evolution, which do not depend on mutation rates. For example, Stebbins (1959) writes:

> Second, mutation neither directs evolution, as the early mutationists believed, nor even serves as the immediate source of variability on which selection may act. It is, rather, a reserve or potential source of variability which serves to replenish the gene pool as it becomes depleted through the action of selection.... The factual evidence in support of these postulates, drawn from a wide variety of animals and plants, is now so extensive and firmly based upon observation and experiment that we who are familiar with it cannot imagine the appearance of new facts which will either overthrow any of them or seriously limit their validity. (p. 305)

Stebbins shows an irreversible commitment to denying any direct role of mutation in evolution. An episode of evolution begins when the environment changes, bringing on selection of an abundance of small-effect variation in the gene pool so that, in the words of Mayr (1963, 613), "mutation merely supplies the gene pool with genetic variation; it is selection that induces evolutionary change." As Mayr (1963, 101) explains at greater length,

> In the early days of genetics it was believed that evolutionary trends are directed by mutation, or, as Dobzhansky (1959) recently phrased this view, "that evolution is due to occasional lucky mutants which happen to be useful rather than harmful." In contrast, it is held by contemporary geneticists that mutation pressure as such is of small immediate evolutionary consequence in sexual organisms, in view of the relatively far greater contribution of recombination and gene flow to the production of new genotypes and of the overwhelming role of selection in determining the change in the genetic composition of populations from generation to generation.

The gene-pool theory has a particular implication about rates of evolution. If an event of mutation is never the token cause that initiates evolutionary change, then the rate of evolution will not depend in any strong way on the rate of mutation; for example, Stebbins (1966, 29) writes:

Mutations are rarely if ever the direct source of variation upon which evolutionary change is based. Instead, they replenish the supply of variability in the gene pool which is constantly being reduced by selective elimination of unfavorable variants. Because in any one generation the amount of variation contributed to a population by mutation is tiny compared to that brought about by recombination of pre-existing genetic differences, even a doubling or trebling of the mutation rate will have very little effect upon the amount of genetic variability available to the action of natural selection. *Consequently, we should not expect to find any relationship between rate of mutation and rate of evolution. There is no evidence that such a relationship exists.* [my emphasis]

A

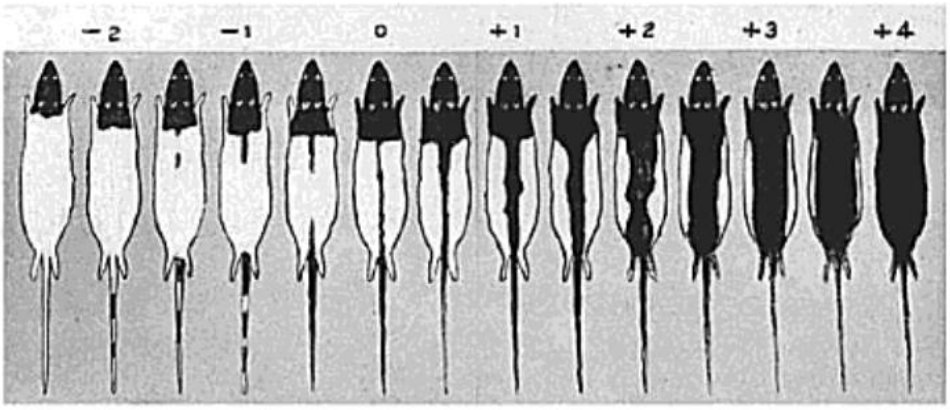

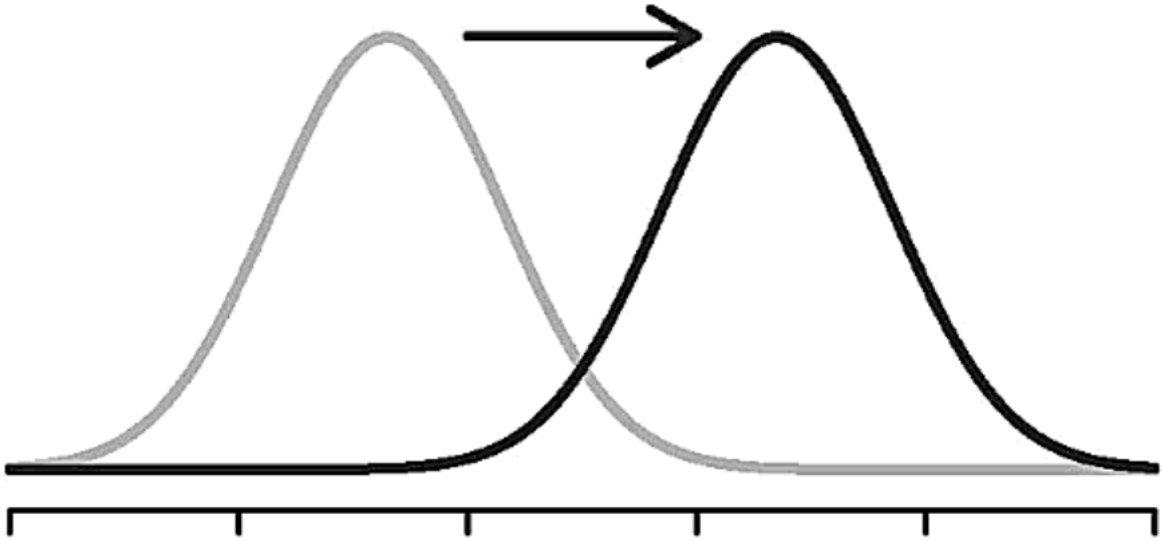

Figure 3.7
The forces view of the OMS. In the OMS, evolution happens when the environment changes, and (A) selection shifts the phenotypic value, sometimes shifting it so far as to result in a new type. (B) At the genetic level, this shift is a multifactorial process, dependent on small-effect variation at many loci. Change is not based on mutation-fixation events but quantitative shifts in frequency of alleles maintained in the gene pool. Therefore, (C) we can understand the process as taking place in the topological interior of an allele-frequency space (i.e., the interior of the cube, but not its edges or surfaces). In this interior space, the trajectory of an evolving system can be understood with the theory of forces. However, the forces theory becomes inadequate when movement occurs outside of this interior space.

Understanding Bias in the Introduction of Variation as an Evolutionary Cause

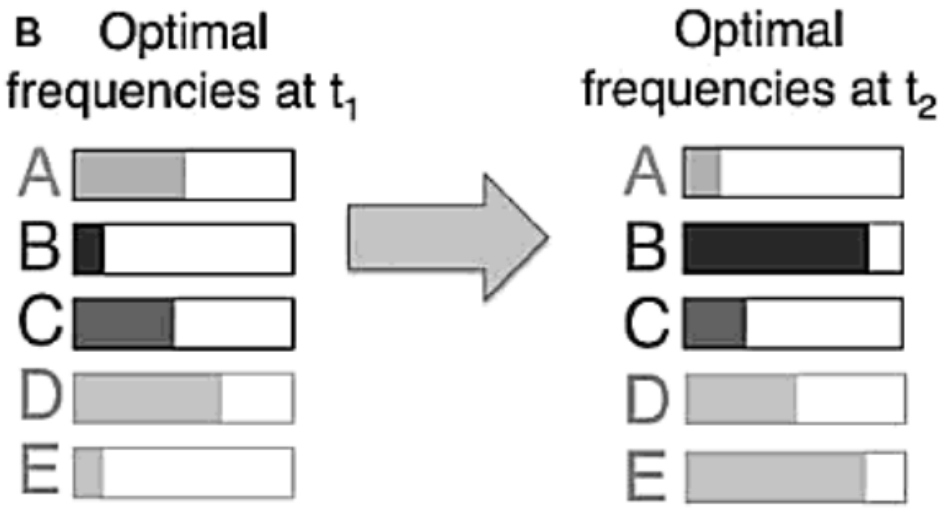

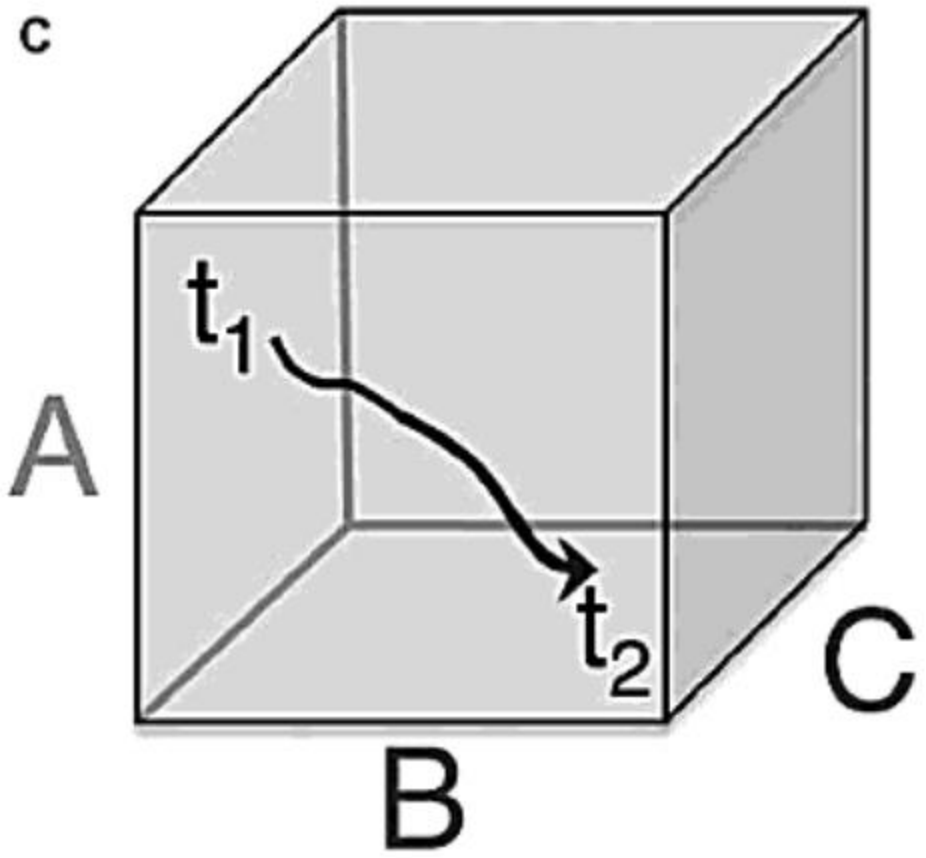

Figure 3.7
(continued)

Likewise, in their 1970s textbook, Dobzhansky et al. (1977, 72) write:

> The large number of variants arising in each generation by mutation represents only a small fraction of the total amount of genetic variability present in natural populations. . . . It follows that rates of evolution are not likely to be closely correlated with rates of mutation. . . . Even if mutation rates would increase by a factor of 10, newly induced mutations would represent only a very small fraction of the variation present at any one time in populations of outcrossing, sexually reproducing organisms.

The Forces Theory

The structure of the OMS encourages a particular conception of evolutionary causes as frequency-shifting forces. Evolution at the phenotypic level is envisioned as a smooth shift to a new adaptive optimum, as in figure 3.7A, with a rate and direction that depend entirely

regimes that are close to it. We described these failures earlier: neither increasing the magnitude of selection coefficients (which increases the force of selection), nor decreasing the magnitude of mutation rates (which decreases the force of mutation), causes selective preferences to overpower mutational ones.

The limitations of the forces view are also revealed in how scientists construct hypotheses and explanations. In molecular evolution, we are regularly faced with comparative evidence suggesting that molecular features, particularly aggregate features like codon usage or proteome composition, are systematically shaped by mutation biases. By the logic of opposing forces, this must mean that the force of selection is effectively absent. Accordingly, for decades, molecular evolutionists habitually treated mutational explanations as references to neutral models (e.g., Sueoka 1988; Gillespie 1991; Gu, Hewett-Emmett, and Li 1998; Knight, Freeland, and Landweber 2001; Lafay et al. 1999; Wolfe 1991).

This interpretation literally suggests that the pressure of recurrent mutation is driving alleles to fixation in the absence of opposing selection. Although this is the correct way to apply the forces theory, the resulting model is not very reasonable, because fixation by recurrent mutation is too slow. The more likely basis for mutation-biased neutral evolution would be that fixation happens by drift, while the mutational bias resides in the introduction process. That is, once we recognize two kinds of causes with two different currencies—introduction and fixation—this immediately suggests a different interpretation in which the mutation bias is applied only to the introduction process. From this point, it is a small step to consider that the fixation process might be either drift or selection, leading to the novel prediction of mutation-biased adaptation.

Conclusion

The notion that a fundamental cause–effect relationship could emerge 80 years after the origins of theoretical population genetics must be considered surprising, as suggested by the epigraph at the head of this chapter (Charlesworth 1996). One eminent theoretician argued that explaining the nature of the results of Yampolsky and Stoltzfus (2001) was not a worthwhile activity because we have "matured past the phase when discussions of basic principles are useful for professionals. We all believe in mutation, selection, and drift, I hope" (see Stoltzfus 2012). Lynch (2007) cites Yampolsky and Stoltzfus (2001) along with a handful of other sources going back to Charles Darwin, explaining that "the notion that mutation pressure can be a driving force in evolution is not new."

A theory is new if scientists have not considered it before, either because it was never introduced, or because it was introduced and then ignored. The question of whether some set of past statements about mutation or variation constitutes a prior description of the kind of cause–effect relationship described here is a matter of applying the duck test: if it looks like a duck, swims like a duck, and quacks like a duck, it's a duck. Either the distinctive implications of biases in the introduction process have been previously shown to arise from a consideration of basic principles of population genetics, or they have not.

One's disbelief at the belated discovery of this kind of cause–effect relationship must begin to fade rapidly when the historical development of evolutionary thought is considered. In fact, (1) the founders of theoretical population genetics considered the possible role of biases in variation as an influence on the course of evolution, (2) they reached the conclusion that such an effect was impossible, and (3) this conclusion became a foundation of the OMS, cited until contemporary times. A role for internal tendencies of variation was rejected by leading thinkers *precisely on the grounds of lacking a population-genetic mechanism*. Indeed, the idea was ridiculed, for example, as when Simpson (1967) argued that selection is the only plausible cause of directional changes in the fossil record, dismissing "the vagueness of inherent tendencies, vital urges, or cosmic goals, without known mechanism" (159).

Decades of results from molecular studies of evolution have conditioned us to accept the mutationist idea that the timing and character of evolutionary change may depend on the timing and character of mutations. Because of this, scientists today may find the evolutionary role of biases in the introduction process to be intuitively obvious. However, the foundations of contemporary evolutionary thought were laid by scientists who explicitly rejected mutationist thinking as an apostasy, and who were committed to the doctrine that variation plays only the role of a material cause, a source of substance only, never a source of form or direction.

From this context, we can begin to see that evolutionary reasoning about causes and explanations is not merely a matter of mathematics, and is not merely a matter of following obvious implications of basic principles that everyone accepts.

Other factors are at work to make the recognition and appreciation of causes non-obvious. We have addressed two of them already. One factor is the role of high-level conjectural theories asserting that evolution works in some ways and not in others, and in particular, the role of neo-Darwinism in shaping the OMS theory of evolutionary genetics. To understand evolutionary reasoning requires not merely a knowledge of mathematics but an awareness of the substantive conjectural theories that have shaped prevailing mathematical approaches to evolution—for example, the gene-pool theory. One cannot understand the strengths and weaknesses of contemporary evolutionary thinking if one does not understand how it relies on the OMS conjecture that evolution operates in the buffet regime of abundant preexisting variation (Stoltzfus 2017).

A second factor, which requires more explication, is the role of verbal theories of causation and explanation. Even if we attempt to set aside the OMS and other broad conjectures, we must recognize that the conceptual tools we use to negotiate issues of causation and explanation include verbal theories that draw on metaphors and analogies (e.g., pressures, raw materials). Mathematics itself is not a language of causation but a notation for relating quantities, along with a set of rules for manipulating these relations. The "equals" sign has no direction, thus it cannot possibly represent an arrow of causation. There is no necessary causal translation of $a = bc$, or the mathematically equivalent statements $b = a/c$ or $a/b - c$.

constructed. In practice, it is constructed in a context that includes not only equations and facts but low-level folk theories of causation (e.g., pressures, raw materials) and high-level conjectures about how causes work together to account for larger phenomena. Recognizing the introduction process as a causal process in evolution, and then recognizing the consequences of biases in the introduction process, makes variation-biased evolution intelligible in a way that led to the prediction of mutation-biased adaptation (now verified), and which yields new insights into the potential role of development and self-organization in evolution.

To summarize, the novelty of scientific theories is not judged by whether they are considered intuitively obvious with benefit of hindsight, or whether they require a complex mathematical derivation. To recognize a causal role of biases in the introduction process is novel because this kind of causation has distinctive behavior and implications. Generic references to mutation pressure, contingency, chance, constraints, and so on, simply are not the same thing as references to biases in the introduction process, because *the former concepts were never previously understood to exhibit the implications that the latter concept can be demonstrated to exhibit*. Generations of theoreticians who invoked the term "mutation" failed to comprehend the implications of biases in the introduction process: again, such comprehension does not emerge magically by uttering the word "mutation," but requires a modeling framework and a conceptualization of causation.

Mutational and developmental biases in the introduction of variation appear to be important determinants of the outcome of evolution, representing a novel kind of causation absent from the Modern Synthesis. The extent of their importance, and particularly their importance as predictable factors acting at the level of phenotypic evolution, is a major unresolved issue in evolutionary biology.