



The landscape of innovation in bacteria, battleships, and beyond

Terence C. Burnham^{a,b,1} and Michael Travisano^{c,d,1,2}

Edited by Simon A. Levin, Princeton University, Princeton, NJ, and approved November 18, 2020 (received for review August 12, 2020)

We draw lessons from microbial experimental evolution and naval warfare to improve the understanding of innovation in financial markets. Major financial innovations often arise without explicit societal planning because novel approaches can be favored by markets, in a manner strikingly parallel to natural selection. We utilize the concept of an adaptive landscape to characterize environments that increase the speed and magnitude of innovation. We apply this adaptive landscape framework to innovation in portfolio management. We create a general taxonomy for understanding and nurturing innovation.

adaptation | experimental evolution | finance | innovation

What environmental characteristics favor innovation? We begin with an examination of innovation evident in naval warfare, primarily the design and function of battleships. Improvements in battleships suggest that fierce competition fosters innovation.

While the battleship heuristic appears straightforward, the study of any historical set of innovations has the challenge of disentangling causation from correlation. Without a time machine, we cannot be sure of the true cause of battleship innovation, nor its message for innovation more broadly.

Experiments, using the venerable scientific method, allow for strong inference. The ability to identify the cause of innovation was one of the primary motivations for Richard Lenski and his coauthors to begin the long-term evolution experiment (LTEE). The LTEE uses *Escherichia coli* and because of their short lifespan relative to humans, provides the opportunity to observe evolution in a forward-looking basis over tens of thousands of generations. Innovation in the early years of the LTEE produces a message consistent with that drawn from battleship design.

An adaptive landscape is a visual representation that elucidates the payoffs to variants. In biological systems, diversity is driven importantly by genetic variation. In human-constructed systems, competing types are produced by design selections.

Both biological and human-constructed systems can utilize an adaptive landscape to represent the

relationship between design and performance. Consequently, both battleship and bacterial innovation can be represented in a similar adaptive landscape.

We draw general lessons regarding approaches to innovation, aided by the adaptive landscape representation. The goal is to understand innovation in finance and other domains. From this perspective, we analyze one area of finance: the management of investment portfolios. We conclude with perspectives on how to best foster innovation.

Dreadnoughts

The HMS *Dreadnought*, commissioned in 1906, featured a steam turbine for power with an unprecedented number of large guns protected by thick armor (Fig. 1). The *Dreadnought* was faster than other ships, so thickly armored that enemy shells simply bounced off, and her own guns were literally able to blow other ships out of the water at superior range.

The word "dreadnought" became an adjective, and naval power was measured by the number of such battleships. Competing global powers commenced rapid and expensive programs to build large fleets filled with dreadnoughts. It was believed (correctly so as it turned out) that the fleet with more dreadnoughts would win subsequent battles and control the oceans.

Britain committed to spending whatever sum of money it took to win this battleship arms race. When

^aArgyros School of Business & Economics, Chapman University, Orange, CA 92866; ^bLaboratory of Financial Engineering, Massachusetts Institute of Technology, Cambridge, MA 02142; ^cDepartment of Ecology, Evolution and Behavior, University of Minnesota, Saint Paul, MN 55108; and ^dBioTechnology Institute, University of Minnesota, Saint Paul, MN 55108

Author contributions: T.C.B. and M.T. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Published under the [PNAS license](#).

¹T.C.B. and M.T. contributed equally to this work.

²To whom correspondence may be addressed. Email: travisano@umn.edu.

Published June 25, 2021.



Fig. 1. HMS *Dreadnought*. Image credit: Wikimedia Commons/Symonds & Co.

World War I (WWI) broke out, Britain formed a blockade around all ports controlled by Germany and her allies. The blockade gave free movement of supplies and troops to Britain while restricting Germany to movement over land.

Recognizing the important strategic advantage provided by the blockade, the German fleet steamed out to confront the British. The two fleets fought the Battle of Jutland in 1916. Measured by the number of ships sunk and casualties inflicted, the German fleet won the battle. Strategically, however, the British blockade held and was never challenged again during WWI.

Britain's blockade held because its Battle of Jutland fleet outnumbered Germany's 28 to 16 in dreadnought-class battleships. The British dreadnoughts directly shaped the outcome of WWI by maintaining the blockade. Additionally, because Germany was unable to compete with surface ships, she turned to submarine warfare.

The United States entered the war on Britain's side after Americans died on the *Lusitania*, sunk by a German U-boat. One million US "doughboys" helped break the European trench warfare stalemate and contributed to Germany's defeat. Thus, the Battle of Jutland helped decide the outcome of WWI. The fight was won because of British naval innovation and economic commitment to dreadnoughts.

What is the dreadnought message for innovation? Innovation is fueled by what we label "hard competition"—competition with winner takes all outcomes. Folk wisdom captures the idea that hard competition produces innovation in a number of well-known adages: necessity is the mother of invention; failure is not an option; and where there is a will, there is a way.

How do we spur innovation? The dreadnought message is that innovation is spurred by increasing the payoff to winning. Put pressure on individuals and groups to win at all costs. Reward first place lavishly and punish even the slightest loss with the most draconian punishment.

Now, we shift to biological innovation in the experimental evolution of bacteria.

Long-Term Experimental Evolution

Innovations structure biology. The evolution of flight, flowers, and photosynthesis, among many other innovations, has transformed living systems and facilitated the evolution of additional innovations. Because life has existed on Earth for roughly 4 billion y, most

investigations of evolutionary innovations are necessarily retrospective (1). Unfortunately, over this history of life, the vast majority of species (>99%) have gone extinct, and there are no living descendants from most lineages.

Hence, the basis for innovation in biology is a topic of both great interest and uncertainty. No one was present to observe the evolution of innovations occurring from tens, hundreds, and thousands of millions of years ago. Much of our understanding of evolutionary innovation comes from deciphering partial information from an incomplete molecular and fossil record.

In theory, one appealing alternative is to look forward in time instead of backward to carry out experimental evolution. Instead of attempting to disentangle a complex and partial history, would it not be better to observe innovation as it evolves in real time?

Experimentation would provide avenues to directly investigate the basis for evolutionary innovation, much in the same way that experimental testing has advanced other topics in biology and science. For many species, however, observing evolution in real time is an exercise in frustration. This is because many organisms require years to reproduce and because adaptive evolution occurs more readily in large, unwieldy populations. Tractability has been the central hurdle in observing evolution as it happens (2).

Many microbes, however, reproduce rapidly, and millions of individuals can be sustained in small volumes of liquid media. Recognizing this, Lenski and colleagues (3) initiated 12 independent replicate populations of the bacterium *E. coli* in an LTEE.

The LTEE populations are propagated daily, going through 6.64 generations every day and each containing 500 million cells by the end of a day's growth. In 1 y, the populations go through 2,400 generations, and each population produces almost 200,000 million cells. With this system, observations of adaptation and evolutionary innovation can be made directly, in one of the most well-studied and genetically tractable model biological systems. The project has continued for over 70,000 generations, with surprisingly parallel observations to the improvement seen in dreadnoughts.

"Survival of the fittest" is a common description of how natural selection operates and provides one way of viewing the LTEE. Throughout the experiment, the populations significantly improved in competitive ability compared with their common ancestor. This is because more fit genotypes increase in frequency during selection, replacing those that are less fit. In the LTEE,




Fig. 2. Trajectories of mean fitness relative to the ancestor in 12 replicate populations of *E. coli* during 10,000 generations.

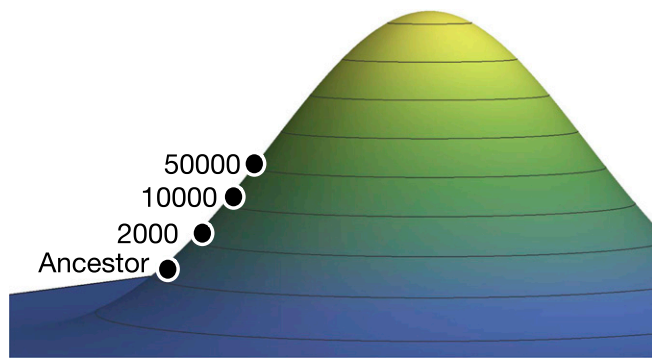


Fig. 3. Innovation in the LTEE visualized in a Fisherian adaptive landscape.

competitive ability is measured by head-to-head assays, in which both ancestral and selected competitors are grown together in the test environment.

The ancestor population can be directly competed with its descendants because bacteria can be preserved indefinitely in suspended animation in ultracold freezers (-112 °F). In direct competition, strains after 2,000 generations of selection are 30% more fit than their shared ancestor (3). By 10,000 generations, the selected populations had increased their selective advantage by 50% (4), which increased to 70% after an additional 10,000 generations (20,000 in total) (Fig. 2).

From these competitive results, we can see improvement was initially quite rapid, and while substantial improvements continued through 20,000 generations, the rate of increase declined. Continuing improvement, but at slower and slower rates, progresses over at least 50,000 generations and fits a power law (5).

Dreadnoughts, LTEE, and the Adaptive Peaks of Ronald Fisher

An adaptive landscape is a visual and conceptual representation of the multidimensional relationship between alternative designs and outcomes. In biological systems, adaptive landscapes are three-dimensional representations of the fitness of all possible genotypes in a particular environment. The x and y axes are representative of genetic diversity, with the z axis as fitness (or sometimes a trait tightly associated with fitness).

While some argue that fitness landscapes should be interpreted literally (6), we employ them here as heuristic models (7, 8). On an adaptive landscape, innovation is represented by moving on the x and y axes to reach a higher level on the z axis. A higher level in the z dimension is interpreted, in the LTEE and other biological systems, as higher fitness.

Fisher (9) argued that most populations are typically in the general vicinity of a global fitness optimum. The "Fisherian fitness landscape" can be visualized as having a single unimodal fitness peak. Because most populations are not too far from optimal, evolution generally proceeds slowly for Fisher (9).

A single unimodal Fisherian fitness landscape can be used to represent the fitness increase in the LTEE (5). Under the force of selection, the bacterial populations move up toward the Fisherian peak. The fitness of the populations increases, and the rate of innovation declines over time (Fig. 3) (second derivative <0).

Two aspects of Fig. 3 should be noted. First, the environment of the LTEE is quite different from the environment experienced by the original population of *E. coli* before the start of the LTEE.

Thus, Fig. 3 shows the original population far from the optimal peak.

Second, Fig. 3 and most of the other adaptive landscape representations of the LTEE represent all 12 replicate populations with a single point based on the average fitness of those populations. There are variations between populations both in average fitness as well as genotype that we discuss more fully below.

One of the central questions motivating the LTEE is determining the likelihood of parallel improvements: would the replicate populations improve similarly during selection (3)? Is the outcome of the experiment largely deterministic, based on shared genetic ancestry and environment (10)?

All of the populations were initiated with the same single ancestral genotype, and all populations experience the same experimental conditions (10 mL glucose supplemented medium at 37 °C). Do the populations inexorably approach an increasingly smaller set of adaptive solutions? Diversity could temporarily persist across the replicates, due to the chance appearance of different beneficial mutations in each replicate. However, if evolution is largely deterministic, the differences would eventually be lost as the experiment proceeds.

That is largely what is observed in the LTEE through at least 10,000 generations. At 2,000 generations of selection, the variance in fitness among populations is quite modest, less than 10% of the improvement (3, 11). Even after 10,000 generations, while there is some hint of divergence among the populations, the differences are dwarfed by the similarities in adaptation to the selected environment (4).

When examined in novel environments, there are significant fitness differences among the replicates (11, 12). However, even this is supportive of a Fisherian unimodal model. The differences among the populations are hidden under normal conditions, not exposed to natural selection, suggesting a single functional solution.

When the differences are exposed to selection, by experimental evolution in a novel environment, the functional differences are largely eliminated (11). Adaptation appears largely deterministic, with ever-declining increases in fitness in any one




Fig. 4. Battleship innovation represented in a Fisherian adaptive landscape.



Fig. 5. The first battle of ironclads. Image credit: Wikimedia Commons/Louis Prang & Co.

environment. Intense selection and competition find additional adaptive improvements, even though there are diminishing returns.

Thus far, the LTEE experiment broadly seems to support a Fisherian perspective, but a caveat is that the populations diverge to some degree in the first 10,000 generations and subsequently. This is to be expected because the LTEE is like the natural world, and the outcome of evolution is always under revision. Some amount of diversity will persist in both the short and long run. Thus, while there is evidence of divergence between the LTEE populations (3, 11–13), we argue that the single-peak Fisherian model is a useful representation of some aspects of innovation in the LTEE.

Innovation Heuristic Drawn from Dreadnoughts and LTEE

An adaptive landscape can be utilized to represent innovation in nonbiologic systems. The z axis is interpreted as a measure of performance with higher being better. Rather than genetic variants, the x and y axes represent different design decisions. We can represent battleships in such an adaptive landscape. In such a battleship landscape, the HMS *Dreadnought* is superior to the USS *Monitor* (the first pure metal ship to fight a battle), which is better than predecessor wooden ships (Fig. 4).

We define hard competition as occurring in environments that place an extreme premium on success—winner takes all outcomes. In contrast, “soft competition” describes environments with relatively more permissive outcomes. “Winning” by having the best relative performance is favored under soft competition, but designs, or genotypes, that have below peak outcomes can persist for some period.

The essential difference between hard and soft competition is the speed with which less adaptive designs decrease in frequency. Hard competition results in faster change and a consequent reduction of variation. Soft competition allows variation in design to persist for longer periods. Our definitions of hard and soft competition are similar but distinct from those of hard and soft selection (14–16). Both of our modes, hard and soft, involve competitive differences among genotypes. However, in hard competition, less fit genotypes are rapidly culled from the population as if they had essentially lethal genes akin to hard selection (14). Conditions favoring soft competition include less steep fitness landscapes, smaller populations, and the absence of recombination (17, 18).

In the adaptive landscape representation, holding population size constant, the steeper the gradient, the harder the competition. One interpretation of the LTEE is that a steeper adaptive landscape fosters innovation. Recall that the ancestral bacterial population, at the beginning of the LTEE, has been under selection outside the laboratory for literally billions of generations. In a mere 10,000 generations, the blink of an eye in evolutionary time, the populations increase their fitness by 30%. Additionally, populations of millions of cells are far from small. Less fit genotypes appear to be rapidly swept from the replicate populations, in as little as 300 generations (3, 4). Therefore, the LTEE message appears to be that competition fosters innovation.

Therefore, in both the LTEE and naval warfare, there is evidence that competition spurs innovation. The idea that innovation is spurred by hard competition in warfare precedes the dreadnought era. Perhaps the most significant innovation in naval warfare, before the dreadnought, was the introduction of metal ships in the form of ironclads.

The “first battle of ironclads” famously took place in 1862 between the *Merrimack* and the *Monitor* (Fig. 5). The American South constructed its first ironclad with the same goal as the German fleet at Jutland—to break a strangling blockade. The *Merrimack* enjoyed 1 d of unfettered success as it destroyed the wooden ships of the Union blockade. On the second day of the *Merrimack*’s metal existence, however, the *Monitor* arrived, and the two ships fought to a draw. The Union blockade held, and the factories of the Union overwhelmed the undersupplied armies of the Confederacy.

While the 1862 *Merrimack–Monitor* battle is famous, almost unknown is that the US Navy approved an ironclad project in 1842 called the *Stevens Battery*. The *Stevens Battery* ironclad project fizzled, and no ship was ever launched. Why did the first funded US ironclad never fight, while the *Monitor* helped win the US Civil War?

The *Stevens Battery* languished, in part, because during peacetime, there was little impetus to complete a novel armored ship design. However, the payoff to an ironclad changed in the intense wartime pressure. In fact, the *Monitor* fought the *Merrimack* less than 5 mo after the start of construction. Furthermore, President Lincoln bypassed all normal channels of military

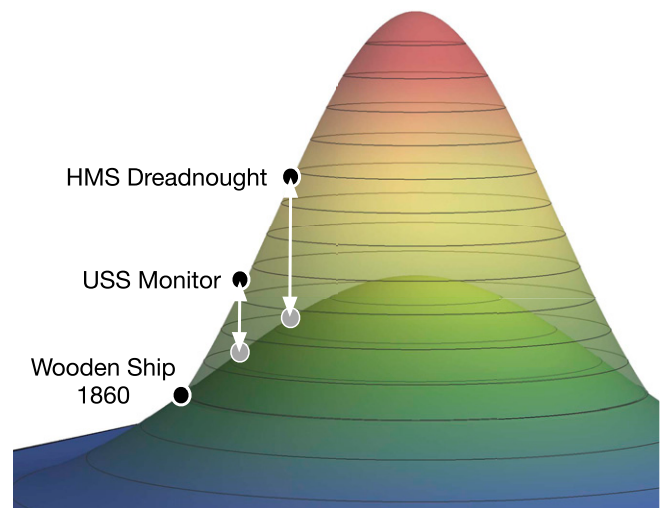


Fig. 6. The rewards to innovation are greater under hard competition.

Table 1. The first and last of the dreadnoughts

Trait	HMS <i>Dreadnought</i> (United Kingdom) 1906	<i>Yamato</i> (Japan) 1941
Displacement (tons)	18,120	71,659
Length (feet)	527	862
Largest guns (inch)	12	18.1
Power (kW)	17,000	110,000
Top speed (knot)	21	27
Armor thickness (inch)	12 maximum	25 maximum
Crew	700	2,500

procurement by personally authorizing the *Monitor*. Had the Union delayed for even a few days, the Confederacy might have won the US Civil War.

Ironclad technology was ready more than a decade before the *Monitor*. Combat-ready ships arose, not because of technological change but rather, because of a change in the payoff to innovation. The war produced a working battleship at a much faster rate than the prior period of peace. In an adaptive landscape representation, war creates a steeper, hard competition environment, whereas the prior peacetime was a more forgiving, soft competition environment (Fig. 6).

The message from naval warfare and the LTEE is that competition fosters innovation. All entities are capable of better performance; they simply need to be forced by the appropriate harsh environment. Do not coddle your possible innovators; eliminate designs that do not win. Innovation will be produced by a caricature of natural selection: “red in tooth and claw.”

Furthermore, it is possible to characterize the nature of innovation. Movement up to the Fisherian peak is a local process where small, incremental change is the key to innovation. Large changes in design are very likely to end in failure. Finally, the outcome of this process is a grinding progression where each generation is likely to be marginally better than the prior design.

A Second Type of Innovation

Superdreadnoughts. The British dreadnoughts won the Battle of Jutland and helped win WWI. As the British fleet sailed home, the future of naval power was obvious: bigger, faster, more powerful battleships under the name superdreadnoughts.

The dreadnought-class ships had dominated all prior ships, instantly making obsolete every earlier design. Superdreadnoughts would feature more powerful engines, higher speeds, bigger guns, more armor, and better command and control systems. No pre-dreadnought ship could stand up to a dreadnought. Similarly, no mere dreadnought could compete with the massive superdreadnoughts that were built between WWI and World War II (WWII).

The Japanese *Yamato*, commissioned in 1941, remains the largest battleship ever built. The *Yamato* was better in every important aspect than the HMS *Dreadnought* (Table 1). The *Yamato* had four times the displacement of the *Dreadnought*, generated six times the power resulting in increased mobility, and bristled with far superior guns and armor.

The dreadnoughts played a central role in winning WWI for Britain and the allies. What was the impact of the *Yamato* on WWII? The answer is that battleships were worse than irrelevant in WWII.

The *Yamato* saw only one minor action and consequently, had no effect on any important aspect of the war (other than diverting materiel and manpower away from productive uses). Rather than

fighting the US forces, the *Yamato* spent her short existence hiding from planes. In a desperation move, she was ordered to beach herself on the shore of Okinawa to become a fixed artillery position. En route to her intended fate as an expensive, inefficient artillery unit, the *Yamato* was sunk in 1945 by planes launched from aircraft carriers. The dreadnought era was officially over.

The *Yamato* was produced by innovation within a fixed architecture. This approach to naval combat extends far back before the *Monitor*. In 480 BCE, for example, the Greek and Persian navies fought the Battle of Salamis (Fig. 7). In this ancient fight, ships from each side attempted to destroy the enemy by hurling projectiles from their ships or by ramming the enemy. More than 2,000 y later, the ships in the Battle of Jutland attempted to destroy the enemy by hurling projectiles and ramming. In fact, the HMS *Dreadnought* itself destroyed exactly one vessel in her career; she rammed and sank a submarine.

While battleships were undergoing rapid innovation between the 1862 *Monitor*, the 1906 *Dreadnought*, and the 1941 *Yamato*, a different approach to naval power was emerging. Naval aviation began by launched balloons off ships, and in 1910, the first fixed wing plane was launched from a ship. These early planes posed no significant threat to battleships. However, the progress in naval air power over the subsequent few decades was dramatic, and by the beginning of WWII, battleships were obsolete. Two decades of radical innovation in air power produced planes that surpassed the product of 2,000 y of warship innovation.

We can represent naval innovation in a series of adaptive landscapes. Between the Battle of Salamis and the launching of the *Yamato*, battleships moved relentlessly upward with new, more capable ships able to dominate predecessors. In an adaptive landscape, each generation of battleship moves closer to a single Fisherian peak (left peak in Fig. 8).

Aircraft as a means to project naval power is innovation in the form of multiple simultaneous design changes. The creation of naval planes utilizes some of the same technologies shared with battleships such as power generation and metal fabrication. However, the manifestation of the physical plane is extremely different from a battleship. Similarly, the aircraft carrier uses a hull and power system that are similar to those of battleships. However, because the goal of an aircraft carrier is to launch planes, not directly engage with other ships, myriad design aspects from layout of the deck to size of guns are very different from battleships.

When planes first made their appearance in the oceans, they were small, with short range and limited offensive power. In these early forms, planes were no threat to battleships. However, over a



Fig. 7. Battle of Salamis—480 BCE. Image credit: Wikimedia Commons/Bayerischer Landtag.




Fig. 8. Planes moved rapidly up a different adaptive peak.

relatively short period (years vs. decades or centuries), plane capabilities increased dramatically.

Fig. 8 shows an adaptive landscape representation of naval power of battleships and aircraft. The two approaches to projecting naval power are located in different areas of the landscape. Over centuries, battleships moved up toward the local peak by continuing to improve in the key attributes for ship vs. ship battle: bigger, faster ships with better defensive protection and increased offensive range and power.

Between WWI and WWII, aircraft increased rapidly in efficacy. By 1942, planes had moved up their adaptive peak to a level superior to that of battleships. The *Yamato*, the largest battleship in history, was destroyed by relatively inexpensive and still rapidly improving aircraft.

An adaptive landscape is a summary, metaphorical representation of an infinitely complex reality. As such, each depiction involves qualitative decisions on what aspects to illustrate and which to abstract. One could put airplanes and battleships into the same adaptive peak under a more general “naval power” representation. Alternatively, a more detailed representation could include submarines and cruise missile. For the purpose of this paper, we prefer the representation of Fig. 8 focusing on the divergence in approaches between battleships and airplanes.

LTEE—Utilization of Citrate. In nature, long periods of evolutionary stasis are punctuated by rapid evolutionary change, for reasons that are intensely debated (19, 20). Specifically, are periods of rapid evolutionary change due to changes in environmental conditions and therefore, consistent with a Fisherian perspective? Or is rapid change the consequences of finding a new evolutionary solution, the appearance of an innovation?

After 31,000 generations of selection, dramatic changes were observed in 1 of the 12 replicate populations in the LTEE (21). The population density increased severalfold, due to the ability of a new genotype to use citrate, a previously unusable carbon and energy resource. Citrate had always been a media component, to facilitate the ability of *E. coli* to acquire iron from the medium (3). However, *E. coli* cannot generally metabolize citrate aerobically, even when it is present at 20 times greater than their typical resource (glucose), as it is in the LTEE. The ability to use citrate is clearly adaptive in the LTEE, as citrate-using mutants dominate the population in which they appear. Eventually, noncitrate-utilizing genotypes go extinct within this one population.

Up to roughly 30,000 generations, this population had been gradually increasing in fitness, albeit increasingly slowly, and then

a revolutionary innovation appeared that changed everything. The bacteria evolved the ability to use a new resource, dramatically increasing in fitness and exploding the population size.

One key difference appears to be a mutation (*gltA1*) in the population that eventually evolves efficient citrate use. The mutation “critically affected the fitness consequences of the pivotal evolutionary step toward innovation” (22), but it is only one of several ways of improving growth on glucose metabolites. Without the *gltA1* mutation, subsequent mutations that could confer efficient citrate growth are deleterious. However, the genotype carrying this mutation was at low frequency and could easily have been lost before being rescued by subsequent citrate-specific mutations (23). These and additional molecular complexities, including gene duplication and activation of an unexpressed nutrient transporter, indicate a fitness landscape with a complex topography, rather than a unimodal fitness landscape. Indeed, substantial cell death is associated with the ability to grow on citrate (24).

Is innovation, in the form of the citrate utilization in the LTEE, consistent with a Fisherian unimodal fitness landscape? Not obviously. The tempo of evolutionary change is not consistent with the eventual appearance of the innovation in the other populations. Furthermore, the same innovation has not appeared in any of the other 12 replicate populations, despite over 70,000 generations of selection, more than double of the length of the time required for initial appearance (24, 25).

In contrast to Fisher, Sewall Wright suggested that evolutionary adaptation was more strongly affected by complex ecological and genetic interactions (17). Such interactions generate complex adaptive landscapes with multiple fitness peaks, local optima, that can potentially trap populations and constrain innovation. Populations are trapped because moving to alternative peaks requires traversing fitness valleys, with losses in the competitive ability of individuals within the population, before moving to the vicinity of the new adaptive peak.

The results from the LTEE provide at least two insights into how to find innovations, even when there are multiple optima. The most obvious is to have multiple replicate (populations). In the LTEE, only 1 population of 12 finds the dramatic evolutionary innovation. If that lineage had somehow not been started or had been lost within the initial 30,000 generations, none of the subsequent observations would have been made. Also, more importantly, there would have been little impetus to explore the possibility of the innovation because it would not have been observed.

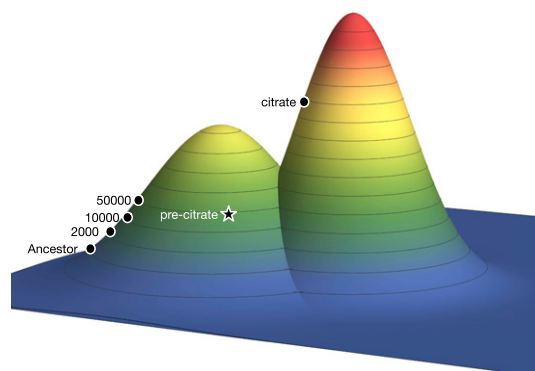


Fig. 9. LTEE moving up a glucose peak and then drifting over to a higher citrate peak.

Table 2. Conditions for incremental and radical innovation

	Incremental	Radical
Competition	Hard	Soft
Changes	Small	Large
Rate	Relentless	Periods of stasis and drift
Peaks	Single	Multiple

The second insight involves nuanced differences between soft and hard competition. In the LTEE, evolutionary success is measured by representation in the population at the time of transfer. Earlier, we discussed how this selection creates increases in fitness in a relatively small number of generations. While this process can be viewed as hard competition, that representation leaves out the possibility of multiple winners occurring roughly contemporaneously. With multiple beneficial mutations, the preexisting less fit genotypes are swept out, as in hard competition, but multiple more fit genotypes increase in frequency. As the preexisting lower fit genotypes are lost, the differences in fitness among the more fit genotypes become increasingly relevant. However, at the same time, even newer beneficial mutations arise, potentially maintaining multiple adaptive lineages in the population. This is known as clonal interference (18).

Such relaxed conditions permit some exploration of evolutionary possibilities because sufficiently good genetic variants can persist for longer in a population before elimination, even if they are not the best variants. Some of these variants can subsequently lead to evolutionary innovations, and this is largely how citrate use evolved in the LTEE. In contrast, under hard competition, populations are pushed to the edge and will go extinct unless the necessary genetic variants appear that can survive in the conditions. Hard competition conditions are thereby less likely to generate innovations because far less genetic exploration is tolerated (Fig. 9). Note, the “precitrate” star in Fig. 9 represents the one population that developed the necessary precursor mutations for subsequent citrate utilization. The remaining 11 populations (or 12 in early generations) are represented by the circles in Fig. 9 and other LTEE adaptive landscapes.

In this representation, all glucose-utilizing *E. coli* populations are drawn on a landscape with a single Fisherian peak. This representation does not feature all of the details of the populations. The LTEE populations have now been maintained for 70,000 generations, with samples frozen every 500 generations. Molecular sequencing of the current and historical populations reveals some interesting variation within some of the glucose-utilizing populations. Several of the glucose-utilizing populations reach a persistent state with multiple staple genotypes. This is labeled as “quasi-stable coexistence” by the authors, possibly caused by frequency-dependent selection (13).

What are the implications for an adaptive landscape representation of multiple genotypes within a population? Taken literally, multiple genotypes imply that the population cannot be represented by a single location in the landscape. While some genes are completely fixed with no functional variation between individuals, there are always loci with variation. Thus, the choice to locate a population on an adaptive landscape is always a qualitative judgment meant, in our approach, as a metaphor.

In the LTEE, all of the glucose-utilizing populations are closer to each other than they are to the citrate-utilizing population. Furthermore, multiple mutations have been documented that allow the utilization of citrate. Thus, the representation of citrate utilization as a separate peak is consistent with our notion of

radical innovation involving multiple relevant genetic changes. No population is literally exactly clonal, so representing a population as a single point on a landscape should be interpreted as a summary abstracting away from some of details (1).

The LTEE citrate utilization message is that radical biological innovation involves drift across the adaptive landscape for long periods of time before the ascent of a different adaptive peak. This message is seen in a variety of other species beyond *E. coli* including yeast (26, 27) and *Drosophila* (28).

Two Different Approaches to Innovation

Innovation can proceed in two different manners. We label these two types of innovation as “incremental” and “radical.” Incremental innovation is the grinding, relentless improvement within a generally fixed design. Incremental innovations tend to be small, local changes within the existing architecture. Radical innovation involves changing multiple design aspects simultaneously.

In the LTEE, improvements in the utilization of glucose are incremental innovations. The ability to utilize citrate is a radical innovation. Incremental innovation in battleships consisted of making bigger, faster ships for directly destroying the enemy. Radical innovation came in the form of airplanes. In both cases, the radical innovations led to the extinction of the previous approaches: battleships became obsolete before WWII, and citrate-utilizing bacteria completely wiped out conspecifics that relied upon glucose.

In the adaptive landscape representation, incremental innovation is viewed as moving up toward the local peak, while radical innovation is a move to a disparate peak in a different area of the landscape. The value of visualizing innovation in an adaptive landscape comes in clarifying both the means of innovation and the methods for fostering innovation.

Necessity is indeed the mother of incremental innovation. Incremental innovation is fostered by a harsh environment with immense rewards for relative performance. Under fierce competitive pressure (hard competition), even small improvements are rapidly adopted. Furthermore, incremental innovation involves the steady accumulation of gradual changes. In the adaptive landscape, hard competition requires sufficient population size and appears as a single fitness peak, and the steeper the landscape, the faster the incremental innovation.

Under hard competition, changes that are nonincremental or not compatible with the current architecture are unlikely to persist. For example, adding a plane in 1906 to the HMS *Dreadnought* would have decreased effectiveness. The planes of 1906 were weak, and any change in ship design to accommodate aircraft would have degraded the ship’s fighting ability. A hybrid dreadnought with a 1906 plane would have been defeated by a pure dreadnought.

In contrast, radical innovation is nurtured by soft competition. Rather than harsh necessity, daydreaming might be a better description of the conditions to foster nonincremental ideas. Albert

Table 3. Active and passive investing

	Active	Passive
Approach	Invest in what you know	Invest blindly
Stock selection	One by one	Buy all stocks
Research	Extensive	None
Portfolio adjustments	Continual	None
Fees	Significant	Low

Einstein wrote, “When I examine myself and my methods of thought, I come to the conclusion that the gift of fantasy has meant more to me than any talent for abstract, positive thinking.”

Radical innovation often requires avoiding winner takes all outcomes. Soft competition helps find radical innovations because of smaller competitive differences among potential winners. The smaller competitive differences allow diversity to appear and persist. This diversity provides the raw material for the radical innovations that can become apparent over time (Table 2). The propensity for these different modes of innovation depend upon the possibilities for divergent adaptive outcomes (the complexity of the landscape) and the potential to explore the landscape.

A final historical detail illustrates the role of soft competition in fostering radical innovation. After WWI, the major powers engaged in an economically devastating superdreadnought arms race. In 1922, fearing economic collapse from the massive naval expenditures, the five most powerful countries signed the Washington Naval Treaty. The treaty set limits on the number and size of ships.

After the Washington Naval Treaty, the pressure to devote all of a country's finances toward battleships softened. Furthermore, the United States had spent significant sums producing the hulls of battleships that could not be completed under the constraints of the treaty. What was the United States do with these costly, partially completed ships?

The aircraft carriers USS *Lexington* and USS *Saratoga* were built on the converted hulls of the halted battleships. Thus, the naval experiment in aircraft was able to innovate at much lower cost. The relaxed competition for battleship construction literally provided the raw material for the radical innovation that ended the battleship era.

Case Example: Radical Innovation in Finance

The incremental and radical innovation framework can be applied to any area. Incremental innovation consists of grinding improvement via small changes. Radical innovation includes multiple nonlocal changes. The incipient phases of radical innovation are marked by failures. In human technological innovation, the early versions are often underestimated due to a combination of novelty and actual shortcomings in the design. The relatively recent revolution in money management is an example of radical innovation that was misunderstood and slow to gain acceptance.

In 1990, Peter Lynch was one of the most famous investors in the world. With over \$13 billion in his Fidelity Magellan fund, Lynch had racked up an average 29% annual return between 1977 and 1990 (29).

Lynch explained his process for investing in Dunkin' Donuts, one of his big winners. After drinking a coffee at one of the Boston area stores, Lynch followed up with detailed analysis of the company's financial statements, calls with Wall Street analysts who were experts in the firm, and discussions with the management of Dunkin' Donuts. After detailed analysis, Lynch made a significant investment in the stock of Dunkin' Donuts and subsequently, watched it soar in price (30).

Peter Lynch was a stock picker. His approach to investing is labeled “active” as he invested in a relatively small percentage of all stocks and changed his investments periodically. Lynch selected his stocks one by one using a research-intensive process similar to his Dunkin' Donuts success. His mantra, preached in books and public talks, was “invest in what you know.”

While Peter Lynch was enjoying success, however, a revolutionary way to invest was gaining traction. The second method is

now called “passive” or index investing, and it differs in several important ways from active management.

Active investment relies upon the expertise of highly trained teams of individuals knowledgeable in every nuance of each company's situation. Passive management, upon first glance, seems like satire. The passive investor does no company research, buys all stocks, never changes investments, and (currently) charges nothing (Table 3).

Perhaps not surprisingly, passive investing had a slow and inauspicious start. The first academic paper suggesting index funds was published in 1960 (31); 16 y later, on 31 August 1976, Jack Bogle and Vanguard launched the first stock index fund, the Vanguard 500.

Investors were not excited by the Vanguard index fund; in fact, the Wall Street underwriters, hired by Vanguard, suggested canceling the offering and not pursuing this silly idea. Vanguard persisted in what became known as “Bogle's folly” (32). The Vanguard 500 index fund started with \$11 million; 8 y passed before Wells Fargo created the world's second index fund. Vanguard itself did not introduce its second passive fund for a decade.

Passive investing now dominates the investment world. There are thousands of passive funds all over the world, investing in every sort of asset from stocks to bonds to commodities and more (33). The total money invested in passive funds surpassed active funds in 2019 (34), and passive funds continue to outgrow active funds.

Fifty-nine years passed from the publishing of the concept in 1960 to the 2019 milestone of passive management exceeding active management in market share. In February 2020, Vanguard had \$6.2 trillion in assets under management, second only to Blackrock and more than twice that of Peter Lynch's former employer, Fidelity (which now offers many passive funds).

In 2018, Jack Bogle wrote, “There no longer can be any doubt that the creation of the first index mutual fund was the most successful innovation—especially for investors—in modern financial history” (35). How did passive investing go from Bogle's folly to (arguably) the “most successful innovation” in modern finance? The answer is superior performance in the form of higher returns and lower risk.

Investors earn higher returns from passive funds than from active funds. Because of lower fees, the average return of passive funds must exceed the average return of active funds (36). (Of course, some active managers do have individual years or even streaks of multiple years of outperformance.)

Investors also achieve lower risk with passive investing. The two most common measures of financial risk are volatility (SD of returns) and beta (covariance of portfolio returns with market returns, scaled by variance in market returns). Passive portfolios, holding the entire market, provide the lowest volatility for any level of expected return.

Thus, theory predicts that passive management will produce both higher, after fee, returns and lower volatility than active management. Live performance data are consistent with the theory. Among US stock funds that invested in large capitalization companies, for example, only 12% of active funds outperformed passive funds over the 15 y ending in 2020 (37). Year after year, the smartest money managers in the world, doing the most detailed analysis possible, lose to the passive approach that buys all stocks without any company-specific research.

In 1990, Peter Lynch retired from investing. Financial organizations tried two different approaches to money management.

The majority of investors tried more sophisticated and expensive efforts to gain specific knowledge about firms: in short, incremental improvements in stock picking in the active style of Peter Lynch. Other investors embraced the passive approach of not picking individual stocks at all.

Passive money management is an example of radical innovation. Its creation required simultaneous change in multiple attributes, which is represented as a movement to a different region in an adaptive landscape. Rather than attempt to move up to a higher level on the current adaptive peak, Vanguard and other passive investors moved across the landscape to a new region, in a manner that has some similarities to the precitrate changes that were on the path to citrate utilization. Just as the aircraft defeated battleships and citrate-utilizing bacteria won in the LTEE, passive investing has outperformed active management.

The message from passive investing is that radical innovation is often underappreciated in its early phases. Indeed, both Wall Street and Main Street thought Vanguard was crazy to buy all stocks without any company-specific research. This once radical approach has become the conventional wisdom.

Summary and Prescription: Nurture Radical Innovation via Soft Competition

Incremental innovation consists of gradual and continuous improvement. This is a grinding process where selection favors small, local improvements. In contrast, radical innovation consists of nonincremental change on multiple design aspects. Radical innovation can produce long periods of stasis punctuated by bouts of rapid improvement.

An adaptive landscape allows visualization and understanding of innovation. Improvements are represented in the landscape as movements to a higher elevation. Incremental innovation is conceptualized as a movement up an adaptive landscape to a local peak. Radical innovation is the movement to a nonlocal area in the landscape and toward a different peak.

Beyond simply representing innovation, an adaptive landscape can provide lessons for how to produce innovation. Hard competition increases the pace of incremental innovation but

decreases the likelihood of radical innovation. Conversely, soft competition slows down incremental convergence to a single winner while increasing the probability of revolutionary change.

To the extent that a person or a group of people has the ability to shape the adaptive landscape, we argue that, perhaps paradoxically, the better approach is to soften competition. The reasons for favoring soft competition are twofold.

First, incremental innovation is likely to proceed without specific efforts for its promotion. The improvement, for example, from an 8" gun to a 12" gun is relatively straightforward and is likely to occur without any effort to influence or nurture the change. Radical innovation, however, requires multiple changes and the opportunity for multiple changes to occur before culling. Furthermore, new radically different designs are likely to have problems. Thus, if these new designs are forced to compete fiercely with incumbent designs, the novel approach is likely to lose and perhaps go extinct.

Second, the landscape is in flux. By this, we mean both the biological landscape on the planet as well as the technological landscape in human society. The rapid increase in human population and the advance of computer and digital technologies are likely to be varying the landscape at a current pace that exceeds the rate of flux in many previous periods. Because of an increased rate of landscape flux, the designs that work today are less likely to persist for as long as they would have in previous eras. Thus, the payoffs to exploring novel parts of the adaptive landscape are likely to be greater now than previously.

Data Availability.

There are no data underlying this work.

Acknowledgments

We thank two anonymous reviewers and thank Joshua Plotkin for excellent discussant comments on a prior version. Andrew Lo and Simon A. Levin created this effort to illuminate financial innovation with insights from the natural sciences, and they expertly managed all aspects of the process. We thank participants in the Massachusetts Institute of Technology Evolutionary Models of Financial Markets meeting, members of the Micropop Laboratory group, Paul Greenberg, and Jay Phelan for helpful comments and discussion.

- 1 D. H. Erwin, D. C. Krakauer, Evolution: Insights into innovation. *Science* **304**, 1117–1119 (2004).
- 2 A. C. Love, M. Travisano, Microbes modeling ontogeny. *Biol. Philos.* **28**, 161–188 (2013).
- 3 R. E. Lenski, M. R. Rose, S. C. Simpson, S. C. Tadler, Long-term experimental evolution in *Escherichia coli*. I. Adaptation and divergence during 2,000 generations. *Am. Nat.* **138**, 1315–1341 (1991).
- 4 R. E. Lenski, M. Travisano, Dynamics of adaptation and diversification: A 10,000-generation experiment with bacterial populations. *Proc. Natl. Acad. Sci. U.S.A.* **91**, 6808–6814 (1994).
- 5 M. J. Wiser, N. Ribeck, R. E. Lenski, Long-term dynamics of adaptation in asexual populations. *Science* **342**, 1364–1367 (2013).
- 6 X. Yi, A. M. Dean, Adaptive landscapes in the age of synthetic biology. *Mol. Biol. Evol.* **36**, 890–907 (2019).
- 7 M. H. Crawford, *Sewall Wright and Evolutionary Biology*. By William B. Provine. xvi and 545 pp. Chicago: University of Chicago Press. 1989, \$18.95 (paper). *Am. J. Hum. Biol.* **2**, 592–593 (1990).
- 8 M. Zuk, M. Travisano, Models on the runway: How do we make replicas of the world? *Am. Nat.* **192**, 1–9 (2018).
- 9 R. A. Fisher, *The Genetical Theory of Natural Selection* (Clarendon Press, 1930).
- 10 Z. D. Blount, R. E. Lenski, J. B. Losos, Contingency and determinism in evolution: Replaying life's tape. *Science* **362**, eaam5979 (2018).
- 11 M. Travisano, J. A. Mongold, A. F. Bennett, R. E. Lenski, Experimental tests of the roles of adaptation, chance and history in evolution. *Science* **267**, 87–90 (1995).
- 12 M. Travisano, R. E. Lenski, Long-term experimental evolution in *Escherichia coli*. IV. Targets of selection and the specificity of adaptation. *Genetics* **143**, 15–26 (1996).
- 13 B. H. Good, M. J. McDonald, J. E. Barrick, R. E. Lenski, M. M. Desai, The dynamics of molecular evolution over 60,000 generations. *Nature* **551**, 45–50 (2017).
- 14 B. Wallace, Hard and soft selection revisited. *Evolution* **29**, 465–473 (1975).
- 15 D. Reznick, Hard and soft selection revisited: How evolution by natural selection works in the real world. *J. Hered.* **107**, 3–14 (2016).
- 16 P. Chen, R. Kassen, The evolution and fate of diversity under hard and soft selection. *Proc. Biol. Sci.* **287**, 20201111 (2020).
- 17 S. Wright, "The roles of mutation, inbreeding, crossbreeding, and selection in evolution" in *Proceedings of the 6th International Congress of Genetics*, D. F. Jones, Ed. (Brooklyn Botanic Garden, Brooklyn, NY, 1932), pp. 356–366.
- 18 S.-C. Park, J. Krug, Clonal interference in large populations. *Proc. Natl. Acad. Sci. U.S.A.* **104**, 18135–18140 (2007).
- 19 S. J. Gould, N. Eldredge, Punctuated equilibria: The tempo and mode of evolution reconsidered. *Paleobiology* **3**, 115–151 (1977).
- 20 S. F. Elena, V. S. Cooper, R. E. Lenski, Punctuated evolution caused by selection of rare beneficial mutations. *Science* **272**, 1802–1804 (1996).

- 21 Z. D. Blount, J. E. Barrick, C. J. Davidson, R. E. Lenski, Genomic analysis of a key innovation in an experimental *Escherichia coli* population. *Nature* **489**, 513–518 (2012).
- 22 E. M. Quandt *et al.*, Fine-tuning citrate synthase flux potentiates and refines metabolic innovation in the Lenski evolution experiment. *Elife* **4**, e09696 (2015).
- 23 D. Leon, S. D'Alton, E. M. Quandt, J. E. Barrick, Innovation in an *E. coli* evolution experiment is contingent on maintaining adaptive potential until competition subsides. *PLoS Genet.* **14**, e1007348 (2018).
- 24 Z. D. Blount *et al.*, Genomic and phenotypic evolution of *Escherichia coli* in a novel citrate-only resource environment. *Elife* **9**, e55414 (2020).
- 25 R. E. Lenski, Convergence and divergence in a long-term experiment with bacteria. *Am. Nat.* **190**, S57–S68 (2017).
- 26 W. C. Ratcliff, R. F. Denison, M. Borrello, M. Travisano, Experimental evolution of multicellularity. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 1595–1600 (2012).
- 27 W. W. Driscoll, M. Travisano, Synergistic cooperation promotes multicellular performance and unicellular free-rider persistence. *Nat. Commun.* **8**, 15707 (2017).
- 28 P. Simões *et al.*, Predictable phenotypic, but not karyotypic, evolution of populations with contrasting initial history. *Sci. Rep.* **7**, 1–12 (2017).
- 29 C. J. Chipello, M. Siconolfi, J. Clements, Both fidelity investors and firm are at sea as Magellan boss goes. *Wall St. J.*, 29 March 1990, Section A, p. 1.
- 30 P. S. Lynch, P. Lynch, J. Rothchild, *One up on Wall Street: How to Use What You Already Know to Make Money in the Market* (Simon & Schuster, 2000).
- 31 E. F. Renshaw, P. J. Feldstein, The case for an unmanaged investment company. *Financ. Anal. J.* **16**, 43–46 (1960).
- 32 J. Zweig, Birth of the index mutual fund: "Bogle's folly" turns 40. *Wall St. J.*, 31 August 2016. <https://blogs.wsj.com/moneybeat/2016/08/31/birth-of-the-index-mutual-fund-bogles-folly-turns-40/>. Accessed 12 January 2021.
- 33 J. C. Bogle, The index mutual fund: 40 years of growth, change, and challenge. *Financ. Anal. J.* **72**, 9–13 (2016).
- 34 D. Lim, Index funds are the new kings of Wall Street. *Wall St. J.*, 18 September 2019. <https://www.wsj.com/articles/index-funds-are-the-new-kings-of-wall-street-11568799004>. Accessed 12 January 2021.
- 35 J. C. Bogle, Bogle sounds a warning on index funds. *Wall St. J.*, 29 November 2018. <https://www.wsj.com/articles/bogle-sounds-a-warning-on-index-funds-1543504551>. Accessed 12 January 2021.
- 36 W. F. Sharpe, The arithmetic of active management. *Financ. Anal. J.* **47**, 7–9 (1991).
- 37 K. Langley, Stock pickers underperformed during coronavirus market turmoil. *Wall St. J.*, 10 June 2020. <https://www.wsj.com/articles/stock-pickers-underperformed-during-coronavirus-market-turmoil-11591786801>. Accessed 12 January 2021.