Limits to predicting evolution: Insights from a long-term experiment with *Escherichia coli*

Jeffrey E. Barrick[1]

Abstract Our inability to predict how populations of cells will evolve is a fundamental challenge to human health and biological engineering. In medicine, one would like to predict and thwart, or at least have time to adequately prepare for potentially catastrophic events such as the emergence of new pathogens, the spread of drug resistance, and the progression of chronic infections and cancers. In bioengineering, one would like to stop, or at least delay, evolution that inactivates a designed function, in order to make genetic engineering and synthetic biology more reliable and efficient. On a larger scale, one would also like to predict when the presence of recombinant DNA or a certain species might pose a threat to nature or civilization if it has the potential to evolve to become harmful.

Bohr's Hydrogen Atom for Evolution?

Many of these examples of biological systems in which we would like to predict evolution are complex: they involve interactions between heterogeneous populations of cells and our immune system or between cells and entire ecosystems. To make headway on this difficult problem, let's first examine what we can predict in a stripped-down evolving system that includes just a single type of relatively simple cell. Perhaps a good working analogy from chemistry is that we'd like to come up with a system and theory on the order of Bohr's model of the hydrogen atom (Turner 2007). This approach is meant as a first step. We will know from the outset that the study system itself lacks some details that are relevant in real world situations (atoms with more electrons in our analogy). We will also only be able to predict some aspects of evolutionary dynamics and not others (the Rydberg formula but not the Zeeman effect in our analogy). Further development on both fronts (systems and

Department of Molecular Biosciences
The University of Texas at Austin
Austin, Texas 78712, USA
e-mail: jbarrick@cm.utexas.edu

models) will ultimately be needed to achieve completeness and accuracy, but this model is still an instructive waypoint on the path to more complex systems.

A population of *Escherichia coli* bacteria in an Erlenmeyer flask is our evolutionary hydrogen atom. There is little doubt that *E. coli* is the best-characterized free-living organism due to its long history as a model system for molecular biology (Judson 1996). In 1988, Richard Lenski and colleagues began propagating twelve *E. coli* populations in the laboratory under carefully controlled conditions to study evolution (Lenski et al. 1991). Every day, $1/100^{th}$ of each culture is transferred to a new flask with fresh nutrients and the *E. coli* repopulate this flask through ~6.6 generations of binary cell division. These *E. coli* reproduce asexually, with no means for genetic recombination between cells. Evolutionary dynamics in this environment are dominated by competition for a limiting supply of the sugar glucose. These twelve microcosms, each its own (simplified) world in a flask, began from an identical starting point and has now evolved in isolation for more than 60,000 cell generations.

For our discussion of predictability here, we will focus almost entirely on examining the Lenski long-term evolution experiment (LTEE) with *E. coli* and a handful of very similar setups. Before proceeding, it is important to acknowledge that there is a vibrant field of experimental evolution that has developed over the past few decades. Many similar, and equally iconic, experiments have been carried out with viruses, bacteria, yeasts, fruit flies, and mice (Garland and Rose 2009). Most of these other experiments have additional layers of complexity. They purposefully include sex, development, parasites, ecology, social behavior, and more. As a result, they have far richer dynamics than are possible in the "hydrogen atom" of the LTEE.

Levels of Prediction in Biology

What does it mean to predict evolution? There are different levels of detail at which this question can be approached. In the LTEE, the evolutionary process can be described numerically in terms of how well-adapted *E. coli* cells have become to their environment over time. Indeed, ***fitness*** is the only quantity that is directly visible to natural selection; its relevance to evolution is fundamental. One can measure fitness in the LTEE as the relative number of offspring that two different cells contribute to the final population when they compete against one another in the same flask (Lenski et al. 1991). As each population evolves over many growth cycles, more-fit *E. coli* carrying beneficial mutations arise and displace their ancestors and competitors. Thus, the fitness of cells in the population increases over time. Can we predict the future course of this upward fitness trajectory, given historical measurements covering previous generations?

Changes in fitness may reflect a wide array of possibilities in how an *E. coli* cell functions. All of these qualities are summarized as its ***phenotype***. Phenotype encompasses the whole range of observable properties of a cell. Some changes in the LTEE are readily visible (under a microscope), such as a cell's size and shape (Lenski and Travisano 1994, Philippe et al. 2009). Some reflect a cell's simple

behaviors: how quickly it starts growing when nutrients first become available each day, how rapidly it replicates while nutrients remain abundant, and to what extent it is able to survive once nutrients become scarce (Oxman et al. 2008, Rozen et al. 2009). Finally, a cell's properties at a molecular level are also part of its phenotype: for example, how many copies of an enzyme are in a cell or how one of its proteins responds to an environmental signal (Cooper et al, 2003). These various types of phenotypic characteristics are often interdependent. A change in the activity of a protein in a key metabolic or regulatory network may lead to more rapid growth of cells, which in turn may lead to a correlated increase in cell size. Can we predict which growth strategies will dominate and how the physiology of cells will evolve?

At the most basic level, phenotypic evolution is determined by changes in an *E. coli* cell's genome, i.e., its ***genotype***. Many mutations that alter this DNA sequence will change the activity of a gene, leading to differences in cellular physiology, behavior, and ultimately fitness. In some cases, it might take multiple mutational steps to rewire cellular networks to achieve a new phenotype. In the LTEE, all cells started with the same genotype and evolution is driven by natural selection on *de novo* mutations that arise due to errors in copying and repairing DNA as cells replicate. Genetic variation often transiently builds up in the population due to competition between genotypes that are descended from the same cell but have since acquired different mutations. Then, genetic diversity typically declines when one genotype has a fitness that is so superior to others for long enough that it drives them extinct (Barrick and Lenski 2009, Maddamsetti et al. 2015). One outcome is certain: over time, mutations will accumulate in the genomes of the successful lineages of cells. Can we predict which genes will mutate? Can we predict how rapidly mutations will accumulate over time in the *E. coli* genome?

Mutational Stochasticity Limits Predictability

One major challenge in predicting evolution, at any level, is that the appearance of new genotypes due to mutations is random. It turns out that in the LTEE we can mostly ignore this stochasticity when making certain types of predictions (much like Bohr could ignore the probabilistic parts of quantum theory in his model of the hydrogen atom). But, it's important to understand why this is the case for thinking more broadly about limits to predicting evolution.

A mutation anywhere in the entire *E. coli* genome occurs just once in every ~1000 cell divisions (Lee et al. 2012). Among these rare mutations, those that happen to be beneficial in a given environment are even rarer. On the order of 1% may give a fitness benefit in a laboratory flask (Perfeito et al. 2007). Those mutations that are at the leading edge of being the most beneficial of these, the ones that will drive adaptation and have a reasonable chance of fixing in a population, are much rarer still. Fewer than one in a million mutations (<0.0001%) may really matter as far as determining the ultimate winners (Gerrish and Lenski 1998, Hegreness et al. 2006, Woods et al. 2011). Thus, one expects variation in evolutionary outcomes due

simply to the uncertainty as to whether a particular new mutation creating a novel genotype will appear in a population.

Counterbalancing the astronomical odds against an important beneficial mutation appearing in any given cell is the fact that each bacterial culture as a whole has many cells. Having a large population size makes evolution more predictable (Szendro et al. 2013). For example, an LTEE population grows up to approximately half a billion (5×10^8) cells each day, and about five million (5×10^6) of these will be transferred to the next flask (Lenski et al. 1991). According to the estimates here (multiplying cell number times mutation rate), there will be on the order of 5×10^5 mutations generated in each LTEE flask every day. The genome size of *E. coli* is only ~5 million (5×10^6) base pairs and it has ~5000 genes. Many mutations will have similar effects on a gene and a cell's phenotype, so most of the next moves in the evolutionary game will be sampled each day!

Of course, many of these mutations, even highly beneficial ones, are lost each day due to the 1/100 dilution bottleneck. Less obviously, competition between diverged lineages of *E. coli* that are accumulating different sets of beneficial mutations further limits the chances that any one mutation will matter in determining how the population as a whole evolves (Fogle et al. 2008). For example, even a very "good" mutation that is unlucky enough to arise in a "bad" genome—i.e., in the company of a cohort of prior mutations that is already lagging in fitness in the evolutionary race—is unlikely to win (Lang et al. 2013). Still, in the aggregate, many different and similar mutations will appear and have a chance to win in every LTEE population. The influence of rare events on the overall outcome, in terms of fitness evolution, is thus relatively weak compared to what it would be in a smaller population.

Evolutionary unpredictability from mutational stochasticity is not, in and of itself, insurmountable or even unusual for a complex system. The LTEE and similar microbial evolution experiments offer two main ways of dealing with the resulting uncertainties. First, even very large bacterial populations require minimal feeding and upkeep, so many replicate populations can be started from precisely the same initial conditions to survey the array of possible outcomes. Thus, one very important aspect of the LTEE was that it consists of not one, but twelve separate populations that have all evolved in precisely the same environment. We can attribute variation between these cultures in how evolution progresses to chance sampling of initial mutations that may cascade into larger differences over time. Other evolution experiments have used even more populations to define the degree to which mutational stochasticity leads to different evolutionary solutions dominating in different populations (Tenaillon et al. 2012).

As is also common in complex systems, the mutations and phenotypes that are successful in the longer term in a population also sometimes critically depend on the initial conditions (and subsequent events that, in effect, become new initial conditions for yet later dynamics). Here, mutations that appear and dominate at early generations set up a genetic background in each population in which further mutational steps can only appear in genomes with these initial mutations. Interactions between the fitness effects of mutations are common in the LTEE. For example, combining the first few mutational steps taken in one winning lineage in all possible

orders showed that one mutation that was highly beneficial when it occurred would have been neutral if it had happened before a certain earlier mutation was already present (Khan et al. 2011). The second way that microbial evolution experiments can deal with mutational stochasticity is that frozen samples of entire *E. coli* populations can be revived to "replay" the dynamics, with additional replication, at various critical points in this process or from genetically defined starting points to reveal these types of contingency in the evolutionary process (Blount et al 2008, Woods et al 2011). Imagine the implications for weather and earthquake prediction if we could watch for patterns in these phenomena time and time again on different earths that were nearly identical before they were set into motion!

Rates of Fitness and Genome Evolution are Predictable

Once we recognize the inherent stochasticity of evolution and quantify the uncertainties in the exact outcomes by studying replicate and replay populations, we can now put our ability to predict the future trajectory of evolution to the test. We will begin at the two levels of prediction that we discussed: changes in the competitive fitness of *E. coli* cells and how quickly mutations accumulate in surviving genomes over time. The in-between predictions of phenotypic characteristics are harder. We'll revisit them in a later section.

Remarkably, most of the replicate populations of *E. coli* in the LTEE display very similar fitness trajectories over the course of the entire >60,000 generation experiment (Wiser et al. 2013, Lenski et al. 2015). With a couple of exceptions (described in the next section), fitness measurements are surprisingly robust to possible artifacts that could complicate their interpretation, such as non-transitivity and frequency-dependence (Elena and Lenski 1997, Wiser et al. 2013). Precise fitness trajectories were measured at various points in the history of the LTEE as it was in progress. At each point, modeling of the trajectory was done in an attempt to predict how fitness would continue to increase in the future. Originally, it was noted that a rectangular hyperbolic curve fit the data well at both 2,000 generations (Lenski et al. 1991) and through 10,000 generations (Lenski and Travisano 1994). However, a hyperbolic curve assumes an asymptote, a maximal fitness ceiling that can never be broken. The asymptote calculated for the data through 2,000 generations was broken by 10,000 generations. The asymptote predicted with the data through 10,000 generations was also later surpassed, so it became clear that a hyperbolic model has a fatal shortcoming in its functional form and leads to poor long-term predictions.

More recently, the rate of fitness increase through 50,000 generations has been fit to an improved "diminishing returns" power law curve (Wiser et al. 2013). This model reflects an intuitive aspects of evolution toward a fitness optimum: it typically becomes harder and harder to improve fitness over time with each new beneficial mutation. Even though it neglects the details of interactions between the fitness effects of individual mutations that are known to be more complex (Chou et al. 2011, Khan et al. 2011) and the detailed dynamics of competition between mutations in a population (Barrick and Lenski 2009, Maddamsetti et al. 2015), this model

makes remarkably accurate predictions. Fitting the model to the LTEE fitness data from 0 to 5,000 generations can predict quite accurately the fitness trajectory out to at least 50,000 generations. Furthermore, extrapolating the model's predictions, even to exceptionally long time horizons (2.5 billion bacterial generations), still makes physiologically reasonable predictions (an evolved *E. coli* doubling time of ~23 minutes) (Wiser et al. 2013). Thus, the average trajectory of fitness evolution into the future can be predicted surprisingly well for a typical population in the LTEE.

With the revolution in next-generation sequencing (Conrad et al. 2011, Barrick and Lenski 2013), it became possible to comprehensively reconstruct the dynamics with which new mutations accumulate over time in the genomes of *E. coli* sampled at different generations from the LTEE. In an initial study, the rate at which mutations accumulated was found to indistinguishable from a linear model based on data from one population (Barrick et al. 2009). However, it was linear in a discontinuous fashion, with two different rates early and late in the experiment, and the linearity was for two different reasons within each time period. Before 20,000 generations, the near linear rate of increase appears to be due to the fitness of the best new genotypes in the population over second-best "also-ran" genotypes remaining near constant, leading to their takeover and fixation in the population happening more or less regularly, except possibly for an initial burst when the first mutations are mainly competing versus the ancestral genotype.

After 20,000 generations, the linear rate at which mutations accumulated in genomes in this LTEE population steeply increased by a factor of more than 20-fold. This acceleration was due to *E. coli* with a much higher mutation rate—due to a defect in a gene that normally prevents the incorporation of damaged nucleotides into DNA—evolving and taking over this population. Similarly high mutation rates have evolved at some point in five of the other eleven LTEE populations (Sniegowski et al. 1997, Tenaillon et al. 2016). Mutations like these, which lead to hypermutation, can be successful in asexual microbial populations because genomes that contain them have a larger per-capita chance of sampling other beneficial mutations that enable them to be successful (Tenaillon et al. 2001, Wielgoss et al. 2013). The mutational trajectories in each of these hypermutator LTEE populations become constant in way that is typical of the clock-like genetic evolution of neutral models (Kimura 1985, Ohta 1992). That is, the accumulation of neutral or nearly neutral mutations in the hypermutators now so greatly outpaces the dynamics with which beneficial mutations appear and sweep through these populations, that it defines the overall rate.

Here, too, analyzing more data has led to a more refined model of the mutational dynamics in the LTEE (Tenaillon et al. 2016). Specifically, after sequencing a total of 264 genomes from all twelve populations, a model that mixes in a diminishing rate of beneficial mutations over time with the normally low clock-like rate of neutral mutations was found to fit the curve for the non-mutator populations better than the original linear model. The overall effect is a slight decrease in the rate of mutations that accumulate over time, though not as strong a deceleration as was found for the fitness trajectories. The form of this model, which combines two evolutionary processes, is supported by various genetic signatures of neutral versus

adaptive evolution and by an ancillary evolution experiment that observed genome dynamics under conditions of relaxed selection to estimate mutation rates (Barrick and Lenski 2013, Tenaillon et al. 2016). In conclusion, the rates at which new mutations accumulate in genomes over time in the LTEE can also be predicted into the future, except for when hypermutators evolve. Even in these cases, after the switch to a different mutation rate, the trajectories settle on new, at least transiently predictable, rates of genome evolution. In time, however, these new rates may further change as a result of compensatory changes or reversions that readjust the mutation rate to lower values (Wielgoss et al. 2013, Tenaillon et al. 2016).

… Except When Ecology and Innovations Appear

The environment that the *E. coli* cells experience has been kept constant for the duration of the LTEE, although not in the sense that it remains entirely unchanging. Rather, it has seasonal regularity. Every 24 hours the same amount of glucose appears, cells "wake up" and grow until this nutrient is depleted over the course of several hours, and then they become quiescent and "sleep" until the next day. The low concentration of glucose, which limits the cell density to about $1/100^{th}$ of what it would be in a typical microbial culture, leaves few opportunities for complex interactions between cells. In dense populations such interactions are often mediated by excreted metabolic byproducts released by some cells becoming a food source or toxin to other cells, but the LTEE cultures are so sparsely populated that the opportunities for these indirect effects are limited. Each cell is essentially competing for glucose on its own without any other influence from the rest of the cells in the mixture. This is one very important reason that the LTEE behaves so well as a "hydrogen atom" for evolution. Nonetheless, in two of the twelve LTEE populations more complicated ecology has crept back into the experiment. The populations with these deviations were ignored in considering the fitness trajectories in the previous section. These departures from the standard model of evolution in this environment would have been difficult to predict *a priori*.

In the first case, one population diversified into two types of cells by 6,000 generations. Each type accumulated a different set of mutations, and these types continued to co-exist for tens of thousands of generations (Rozen et al. 2005, Plucain et al. 2014). Apparently, one type was a superior competitor during growth on glucose while the other type was better at surviving and scavenging byproducts, and perhaps nutrients from dead cells, after most or all of the glucose had been exhausted (Rozen et al. 2009). These different behaviors led to a situation in which each type had an advantage over the other when it became rare, thus, stabilizing their long-term co-existence. Populations of *E. coli* in another evolution experiment, in an environment that includes a higher concentration of glucose mixed with acetate (which transiently accumulates as a byproduct of glucose metabolism in the LTEE), nearly always diversify into two specialist types: one that grows fastest on glucose and one that switches more rapidly to utilizing acetate in a second growth phase (Spencer et al. 2008, Herron and Doebeli 2013). Thus, unpredictability in this

case in the LTEE likely stems from conditions (low nutrient concentrations) that are on the cusp of a domain in which a more complex ecology is a likely evolutionary outcome. The upshot is that stable coexistence of diverged *E. coli* types with different growth strategies requires a rare sequence of mutations and/or interactions within a population to develop in the LTEE.

In a second population, an even bigger deviation from the typical evolutionary dynamics occurred. This population evolved to utilize citrate, a second potential nutrient that has been present in every flask on every day of the LTEE at a much higher concentration than glucose. *E. coli* cannot normally metabolize citrate under these oxygen-rich conditions, and the citrate innovation was rare—it appeared only after ~30,000 generations of evolution (~15 years) and has remained unique to this one population so far of the twelve (Blount et al. 2008, Blount et al. 2012). Citrate utilization enabled these newly evolved bacteria to colonize a vacant nutrient niche, essentially giving them a private and highly abundant resource. So, it was "big league" beneficial. The citrate innovation is rare in the LTEE because, in part, it is contingent upon a certain set of earlier mutations that alter *E. coli* metabolism in this particular population in a way that pre-adapts them, such that a subsequent mutation that turns on a pump that can exchange citrate into these cells is beneficial, rather than neutral or even deleterious (Quandt et al. 2014, Quandt et al. 2015, Leon et al. 2018).

After efficient citrate utilization arose in this population, a complex ecology also evolved, one related to how the citrate users export other carbon compounds into their environment in exchange for this nutrient. These efflux byproducts accumulate and can be utilized in turn by other genotypes that evolved to specialize on them (Turner et al. 2015). The citrate-eating subpopulation also evolved a high mutation rate shortly after it arose (Blount et al. 2012). The dynamics of adaptation had effectively been reset, such that it was back at the beginning of an increasing fitness trajectory for optimizing growth on citrate instead of deep into the diminishing tail of adaptation to glucose.

Predicting the Genetic Basis of Adaptation is Difficult

Now, let's consider where our predictions start to fail. We can also look not just at how many mutations there are in an *E. coli* genome, but at what genes they affect. In general, there is a lot of parallelism (i.e., convergent molecular evolution) in what genes acquire beneficial mutations among the twelve populations of the LTEE (Barrick et al. 2009, Tenaillon et al. 2016, Good et al. 2017). Though the exact changes to the DNA sequences of those genes are rarely the same, it is likely that the mutations in the same gene from different populations have the same, or at least very similar, effects on molecular and cellular phenotypes. There are even predictable dependencies within some of these genes, such that an earlier mutation in a certain gene can change the probabilities of further mutations accumulating in other genes (Woods et al. 2011, Good et al. 2017). So, we might be able to build up a model of the expected probabilities of different mutational paths impacting certain genes and

cellular pathways in this system, but currently we can only do this on a *post hoc* basis by looking at enough replicate and replay experiments.

What about predicting what mutations in which genes will evolve beforehand? For this aim, we would need a mechanistic model that connects mutations to fitness through our previously neglected level of cellular phenotypes. Advanced models of metabolic and gene regulatory networks exist for bacterial cells (Karr et al. 2012, Monk et al. 2013). In certain cases, one can indeed identify “traffic jams” in metabolic “highways” that are alleviated by “road-widening” mutations in specific enzymes during the course of adaptive laboratory evolution experiments (Ibarra et al. 2002, McCloskey et al. 2018). One can also construct cells with specific changes in gene functions (knockouts, especially) that are predicted by these whole-cell models and show that they are often beneficial to fitness. However, even the complexity of "just" an *E. coli* cell makes it rare that we can predict *a priori* which genes will harbor the best beneficial mutations that will drive adaptation during an evolution experiment.

Why? Often, it is mutations in global regulatory processes instead of single enzymes that are the most impactful (Maharjan et al. 2006, Phillipe et al. 2007, Conrad et al. 2011). The effects of these mutations are difficult to predict because they change many of the links in a cellular network at once. Even though our systems biology knowledge of the *E. coli* strain used in the LTEE continues to improve (Houser et al. 2015, Brown et al. 2017, Cagler et al. 2017), we don't have anywhere near a complete accounting of these subtle effects. Examples of global regulators include RNA polymerase, nucleoid-like DNA-binding proteins, and enzymes that wind and unwind DNA. A change in any one of these proteins may up- or down-regulate the expression of hundreds or thousands of other genes. Some of the changes in the levels of these affected proteins may be beneficial, neutral, or even deleterious on their own. Since the net effect is the sum of many of these weak interactions, regulatory genes may be particularly likely targets for adaptive mutations. It is currently difficult to predict when a specific global regulator will be an effective target for selection. One remaining question is: will mutations in global regulators remain just as common or diminish in importance as LTEE populations reach higher fitness? On the one hand, one might expect evolution to give way to mutations that more precisely adjust the activities of individual genes. On the other hand, mutations in global regulators may still be just as beneficial to fitness overall as those one-gene mutations because they can fine-tune many targets simultaneously.

There are further complications in predicting what mutations will occur in an evolution experiment. There are different mutational target sizes for different types of changes in gene function, and mutations leading to different types of genome variation may arise at vastly different rates (Ryall et al. 2012). For example, many single-base edits to a gene will inactivate it or reduce its activity, whereas very few may result in greater activity or novel functions. Thus, mutations that inactivate genes will appear more often than other types of mutations, and they may have a short-term advantage for this reason in the evolutionary race, even though they may not be optimal in the long run. There are also certain genomic regions and DNA sequences that are especially prone to mutations compared to others. Since they

mutate more rapidly, they may reliably contribute to evolution above their "weight class" (i.e., even when they are not highly beneficial). For example, deletions of the ribose-utilization operon due to a mutational hotspot occur at a high rate in the LTEE and rapidly fix in successful genomes, even though this mutation is only slightly beneficial (Cooper et al. 2001). Though some types of mutational hotspots can be identified computationally (Jack et al. 2015), we are far from being able to comprehensively predict these types of sites accurately enough to take them into account and predict which mutations are likely to contribute to evolution within a gene or genome.

Predicting, Preparing for, and Preventing Evolution

In the end, what can we say about Lenski's flasks as Bohr's hydrogen atom for evolution? First, it truly is possible to predict (or forecast) certain aspects of the evolutionary trajectory fairly well at a non-mechanistic, non-molecular level. Changes in fitness and in the numbers of mutations that accumulate over time follow reliable trajectories. If we have data for a short initial period and fit a few parameters to get the shape of the curve, these predictions hold well into the future. This approach works even though our models of the underlying evolutionary dynamics are relatively coarse grained (e.g., we have one generic term for the diminishing fitness returns of each new mutation and don't need to know any details of genetic diversity in the population).

These types of predictions for the LTEE fail, however, when evolution finds a way to break the rules of the game. In these instances, some of our first-order assumptions are violated, e.g., by the appearance of hypermutators, ecological interactions, or metabolic innovations. The outcome is outside of the realm in which evolution is gradually optimizing fitness on glucose as a limiting nutrient with the ancestral mutation rate. Unfortunately, it is just these types of unanticipated events that we are most concerned about when predicting evolution in the real world. It's not that we are worried that evolution will proceed a bit faster or slower than we planned, it's that an out-of-scope danger will arise: a species will mutate to become invasive in a new environment or an especially virulent pathogen will emerge. It's the risk of these rare events and chance encounters that we struggle to define and mitigate; they are the "hopeful monsters" that keep us up at night.

In closing, let's consider these rare and potentially destructive evolutionary events more directly. In terms of prediction, they are more akin to earthquakes than to electrons, and a shift in analogy at this point is helpful for changing our perspective. It may be near impossible to forecast the exact moment at which an earthquake will occur, but this limitation does not mean that we are helpless against them. Not all of the danger from an earthquake is immediate. In particular, earthquakes may trigger tsunamis that travel thousands of miles across the ocean before wreaking havoc on distant communities. While the propagation of these massive tidal waves also cannot be precisely predicted, providing early warning to at-risk locations by triggering an alarm immediately after an earthquake will save lives.

So, what if we are asking the wrong (or at least a harder than necessary) question of biology in trying to predict outcomes in a complex system of evolving cells? In many cases, such as the evolution of drug resistance, it may be almost as useful as prediction if we could just receive an early warning about what resistance mutations have appeared in a population while they are still very rare. Then we could prepare for these contingencies, by switching the drugs we use before the troublesome variants ever begin to matter. The tsunami warning that we may be after could be deeply sequencing an evolving bacterial or tumor population to profile its rare genetic diversity (Maley et al. 2006). Early warning may be enough in these cases.

Another strategy for mitigating the destruction of earthquakes is disaster preparedness. Although we cannot predict exactly when and where an earthquake will happen, we know that certain areas of the world are far more prone to seismic activity than others. We have mapped out the fault systems of the world based on a long history of recording earthquakes and found they delineate the Ring of Fire. As more and more genomic information becomes available from sequencing tumors and microbial populations involved in chronic infections, such as those in cystic fibrosis (Lieberman et al. 2011, Marvig et al. 2015), we are building up similar genetic maps of how problem cells are likely to evolve over time. The resolution of these maps could potentially be improved by augmenting the outcomes of "natural experiments" or unplanned infections with laboratory experiments like the LTEE, for example with the flu or bacterial pathogens. By repeating and recording many outcomes of evolution in the lab, we would theoretically know what problems to expect and could tailor treatments to undercut those evolutionary paths. This strategy is not unlike enforcing robust building codes in earthquake-prone areas in order to mitigate dangerous forces that are unpredictable and sporadic on a daily basis but essentially certain to occur in the long run.

Finally, what if we invert the prediction problem and seek to purposefully re-engineer an organism's genome to make its evolution more predictable? One could perhaps unravel the tangled network of weak links in cellular networks that has been the product of mindless evolution in a much more complex environment and replace it with a simpler gene expression scheme (Temme et al. 2012). Given that entire microbial genomes are now being constructed or mutated on a large-scale (Esvelt and Wang 2013), this level of re-design is becoming a possibility. Just as mutation rates can evolve to be higher than normal, it is also possible to engineer and evolve "antimutator" organisms that have lower than natural mutation rates (Renda et al. 2014, Deatherage et al. 2018). The limits of this approach have not yet been fully explored, and it might also be possible to block some cellular process that lead to stress-induced mutagenesis with drugs to lower cellular mutation rates (Al Mamun et al. 2012). With these final interventions, evolution of any kind would be expected to be less of a danger, simply because harmful genetic variants would be less likely to ever appear in the first place.

Acknowledgments The inspiration for this chapter came from attending a Santa Fe Institute workshop on "Limits to Prediction". I thank David Krakauer and Joshua Grochow for organizing this event and my fellow participants for stimulating talks about topics ranging from weather and earthquakes to viruses and the evolution of complexity. I would also like to acknowledge many

colleagues and collaborators from the National Science Foundation BEACON Center for the Study of Evolution in Action, particularly Erik Goodman for his leadership in creating this one-of-a-kind environment and Richard Lenski for helpful comments on a draft of this chapter. Research in my lab related to the predictability of evolution has been sponsored by BEACON (DBI-0939454), an NSF CAREER Award (CBET-1554179), the National Institutes of Health (GM087550), the Army Research Office (W911NF-12-1-0390), the Defense Advanced Research Projects Agency (HR0011-15-C0095), and the Welch Foundation (F-1780).

References

Al Mamun AAM et al. (2012) Identity and function of a large gene network underlying mutagenic repair of DNA breaks. Science 338:1344–1348.

Barrick JE, Lenski RE (2009) Genome-wide mutational diversity in an evolving population of *Escherichia coli*. Cold Spring Harb Symp Quant Biol 74:119–129.

Barrick JE et al. (2009) Genome evolution and adaptation in a long-term experiment with *Escherichia coli*. Nature 461:1243–1247.

Barrick JE, Lenski RE (2013) Genome dynamics during experimental evolution. Nat Rev Genet 14:827–839.

Blount ZD, Borland CZ, Lenski RE (2008) Historical contingency and the evolution of a key innovation in an experimental population of *Escherichia coli*. Proc Natl Acad Sci USA 105:7899–7906.

Blount ZD, Barrick JE, Davidson CJ, Lenski RE (2012) Genomic analysis of a key innovation in an experimental *Escherichia coli* population. Nature 489:513–518.

Brown CW, Sridhara V, Boutz DR, Person MD, Marcotte EM, Barrick JE, Wilke CO (2017) Large-scale analysis of post-translational modifications in *E. coli* under glucose-limiting conditions. BMC Genomics 18:301.

Caglar MU, Houser JR, Barnhart CS, Boutz DR, Carroll SM, Dasgupta A, Lenoir WF, Smith BL, Sridhara V, Sydykova DK, Vander Wood D, Marx CJ, Marcotte EM, Barrick JE, Wilke CO (2017) The *E. coli* molecular phenotype under different growth conditions. Sci. Rep. 7:45303.

Chou H-H, Chiu H-C, Delaney NF, Segrè D, Marx CJ (2011) Diminishing returns epistasis among beneficial mutations decelerates adaptation. Science 332:1190–1192.

Conrad TM, Lewis NE, Palsson BO (2011) Microbial laboratory evolution in the era of genome-scale science. Mol Syst Biol 7:509.

Cooper VS, Schneider D, Blot M, Lenski RE (2001) Mechanisms causing rapid and parallel losses of ribose catabolism in evolving populations of *Escherichia coli* B. J Bacteriol 183:2834–2841.

Cooper TF, Rozen DE, Lenski RE (2003) Parallel changes in gene expression after 20,000 generations of evolution in *Escherichia coli*. Proc Natl Acad Sci USA 100:1072–1077.

Deatherage DE, Leon D, Rodriguez ÁE, et al (2018) Directed evolution of *Escherichia coli* with lower-than-natural plasmid mutation rates. Nucleic Acids Res 46:9236–9250.

Elena SF, Lenski RE (1997) Long-term experimental evolution in *Escherichia coli*. VII. Mechanisms maintaining genetic variability within populations. Evolution 51:1058–1067.

Esvelt KM, Wang HH (2013) Genome-scale engineering for systems and synthetic biology. Mol Syst Biol 9:641.

Fogle CA, Nagle JL, Desai MM (2008) Clonal interference, multiple mutations and adaptation in large asexual populations. Genetics 180:2163–2173.

Gerrish PJ, Lenski RE (1998) The fate of competing beneficial mutations in an asexual population. Genetica 102–103:127–144.

Good, BH, McDonald MJ, Barrick JE, Lenski RE, Desai MM. (2017) The dynamics of molecular evolution over 60,000 generations. Nature 551:45–50.

Hegreness M, Shoresh N, Hartl D, Kishony R (2006) An equivalence principle for the incorporation of favorable mutations in asexual populations. Science 311:1615–1617.
Herron MD, Doebeli M (2013) Parallel evolutionary dynamics of adaptive diversification in *Escherichia coli*. PLoS Biol 11:e1001490.
Houser JR et al. (2015) Controlled measurement and comparative analysis of cellular components in *E. coli* reveals broad regulatory changes in response to glucose starvation. PLOS Comput Biol 11:e1004400.
Ibarra RU, Edwards JS, Palsson BO (2002) *Escherichia coli* K-12 undergoes adaptive evolution to achieve in silico predicted optimal growth. Nature 420:186–189.
Jack BR et al. (2015) Predicting the genetic stability of engineered DNA sequences with the EFM calculator. ACS Synth Biol 4:939–943.
Judson HF. 1996. The Eighth Day of Creation (Cold Spring Harbor Press).
Karr JR et al. (2012) A whole-cell computational model predicts phenotype from genotype. Cell 150:389–401.
Khan AI, Dinh DM, Schneider D, Lenski RE, Cooper TF (2011) Negative epistasis between beneficial mutations in an evolving bacterial population. Science 332:1193–1196.
Kimura M (1985) The Neutral Theory of Molecular Evolution (Cambridge University Press).
Lang GI et al. (2013) Pervasive genetic hitchhiking and clonal interference in forty evolving yeast populations. Nature 500:571–574.
Lee H, Popodi E, Tang H, Foster PL (2012) Rate and molecular spectrum of spontaneous mutations in the bacterium *Escherichia coli* as determined by whole-genome sequencing. Proc Natl Acad Sci USA 109:E2774–E2783.
Lenski RE, Rose MR, Simpson SC, Tadler SC (1991) Long-term experimental evolution in *Escherichia coli*. I. Adaptation and divergence during 2,000 generations. Am Nat 138:1315–1341.
Lenski RE, Travisano M (1994) Dynamics of adaptation and diversification: a 10,000-generation experiment with bacterial populations. Proc Natl Acad Sci USA 91:6808–6814.
Lenski RE, et al. (2015) Sustained fitness gains and variability in fitness trajectories in the long-term evolution experiment with *Escherichia coli*. Proc R Soc B 282:20152292.
Leon D, D'Alton S, Quandt EM, Barrick JE. (2018) Innovation in an *E. coli* evolution experiment is contingent on maintaining adaptive potential until competition subsides. PLoS Genet. 14:e1007348.
Lieberman TD et al. (2011) Parallel bacterial evolution within multiple patients identifies candidate pathogenicity genes. Nat Genet 43:1275–1280.
Maddamsetti R, Lenski RE, Barrick JE (2015) Adaptation, clonal interference, and frequency-dependent interactions in a long-term evolution experiment with *Escherichia coli*. Genetics 200:619–631.
Maharjan R, Seeto S, Notley-McRobb L, Ferenci T (2006) Clonal adaptive radiation in a constant environment. Science 313:514–517.
Maley CC et al. (2006) Genetic clonal diversity predicts progression to esophageal adenocarcinoma. Nat Genet 38:468–473.
Marvig RL, Sommer LM, Molin S, Johansen HK (2015) Convergent evolution and adaptation of *Pseudomonas aeruginosa* within patients with cystic fibrosis. Nat Genet 47:57–64.
McCloskey D, Xu S, Sandberg TE, et al (2018) Evolution of gene knockout strains of *E. coli* reveal regulatory architectures governed by metabolism. Nat Commun 9:3796.
Monk JM et al. (2013) Genome-scale metabolic reconstructions of multiple *Escherichia coli* strains highlight strain-specific adaptations to nutritional environments. Proc Natl Acad Sci USA 110:20338–20343.
Ohta T (1992) The nearly neutral theory of molecular evolution. Annu Rev Ecol Syst 23:263–286.
Oxman E, Alon U, Dekel E (2008) Defined order of evolutionary adaptations: experimental evidence. Evolution 62:1547–54.
Perfeito L, Fernandes L, Mota C, Gordo I (2007) Adaptive mutations in bacteria: high rate and small effects. Science 317:813–815.

Philippe N, Crozat E, Lenski RE, Schneider D (2007) Evolution of global regulatory networks during a long-term experiment with *Escherichia coli*. BioEssays 29:846–860.

Philippe N, Pelosi L, Lenski RE, Schneider D (2009) Evolution of penicillin-binding protein 2 concentration and cell shape during a long-term experiment with *Escherichia coli*. J Bacteriol 191:909–921.

Plucain J et al. (2014) Epistasis and allele specificity in the emergence of a stable polymorphism in *Escherichia coli*. Science 343:1366–1369.

Quandt EM, Deatherage DE, Ellington AD, Georgiou G, Barrick JE (2014) Recursive genomewide recombination and sequencing reveals a key refinement step in the evolution of a metabolic innovation in *Escherichia coli*. Proc Natl Acad Sci USA 111:2217–22.

Quandt EM et al. (2015) Fine-tuning citrate synthase flux potentiates and refines metabolic innovation in the Lenski evolution experiment. eLife 4:e09696.

Renda BA, Hammerling MJ, Barrick JE (2014) Engineering reduced evolutionary potential for synthetic biology. Mol Biosyst 10:1668–1678.

Rozen DE, Schneider D, Lenski RE (2005) Long-term experimental evolution in *Escherichia coli*. XIII. Phylogenetic history of a balanced polymorphism. J Mol Evol 61:171–180.

Rozen DE, Philippe N, Arjan De Visser J, Lenski RE, Schneider D (2009) Death and cannibalism in a seasonal environment facilitate bacterial coexistence. Ecol Lett 12:34–44.

Ryall B, Eydallin G, Ferenci T (2012) Culture history and population heterogeneity as determinants of bacterial adaptation: the adaptomics of a single environmental transition. Microbiol Mol Biol Rev 76:597–625.

Sniegowski PD, Gerrish PJ, Lenski RE (1997) Evolution of high mutation rates in experimental populations of *E. coli*. Nature 387:703–705.

Spencer CC, Tyerman J, Bertrand M, Doebeli M (2008) Adaptation increases the likelihood of diversification in an experimental bacterial lineage. Proc Natl Acad Sci USA 105:1585–1589.

Szendro IG, Franke J, de Visser JAGM, Krug J (2012) Predictability of evolution depends nonmonotonically on population size. Proc Natl Acad Sci USA 110:571-576.

Temme K, Zhao D, Voigt CA (2012) Refactoring the nitrogen fixation gene cluster from *Klebsiella oxytoca*. Proc Natl Acad Sci USA 109:7085–7090.

Tenaillon O, Taddei F, Radmian M, Matic I (2001) Second-order selection in bacterial evolution: selection acting on mutation and recombination rates in the course of adaptation. Res Microbiol 152:11–16.

Tenaillon O et al. (2012) The molecular diversity of adaptive convergence. Science 335:457–461.

Tenaillon O et al. (2016) Tempo and mode of genome evolution in a 50,000-generation experiment. Nature 536:165–170.

Turner CB, Blount ZD, Mitchell DH, Lenski RE (2015) Evolution and coexistence in response to a key innovation in a long-term evolution experiment with *Escherichia coli*. bioRxiv:1–36.

Turner JE (2007) Chapter 2: Atomic Structure and Atomic Radiation in Atoms, Radiation, and Radiation Protection, pp 15–53 (Wiley-VCH).

Wielgoss S et al. (2013) Mutation rate dynamics in a bacterial population reflect tension between adaptation and genetic load. Proc Natl Acad Sci USA 110:222–7.

Woods RJ et al. (2011) Second-order selection for evolvability in a large *Escherichia coli* population. Science 331:1433–1436.