Explaining Metabolic Innovation: Neo-Darwinism *versus* Design

Douglas D. Axe and Ann K. Gauger

Biologic Institute, 16310 NE 80th Street, Redmond, WA 98052, USA.
daxe@biologicinstitute.org, agauger@biologicinstitute.org

Abstract

Like all life, bacterial life depends on a complex, integrated network of precise metabolic processes. These processes are carried out by more than a thousand enzymes — genetically encoded proteins with information-rich three-dimensional structures that catalyze specific chemical reactions. Can neo-Darwinian theory explain the origin of this network of enzymes that orchestrates metabolic complexity? Building on previous experimental and theoretical work, we argue here that it cannot. But instead of merely listing the theory's shortcomings, we attempt to construct a full and coherent picture of *how* it has failed to explain metabolic innovation, from the level of single enzymes all the way up to the network of enzymatic pathways that composes metabolism as a whole. Then, from this critical synthesis we identify six key principles of a new theory of biological innovation. Although these principles only hint at the substance of the new theory, they show clearly that it will be strikingly unlike neo-Darwinism. Whereas the old theory focuses on the simple material processes of mutation and selection in the hope that these can drive innovation, the new one focuses on innovation itself — on the concepts that guide effective designs. Consequently, the new theory will look more like the systematic concepts of an engineering discipline than a set of causal laws.

Key words: metabolic complexity, innovation, pathway evolution, complex adaptation, enzyme recruitment, cost of gene expression, causal circularity, design principles, critique of neo-Darwinism

Introduction

Life exhibits extraordinary functional complexity on many scales, from the molecular to the organismal and on up to whole ecosystems. Near the bottom of this arrangement is *metabolic* complexity, which refers to the intricate networks of coordinated chemical reactions that undergird all biological phenomena. Even the very simplest organisms, bacteria, are highly complex in this respect, which makes metabolic complexity a universal hallmark of life. Its universality also makes it a *benchmark* for assessing theories of biological origins. That is, any theory claiming to explain the origin of biological complexity in general must tackle the particular challenge of explaining metabolic complexity.

How well has the dominant theory, neo-Darwinism, met this challenge? The structure of metabolism itself suggests that this should be assessed in a hierarchical way. At the lower level the question is how well the theory explains the origin of new functions for single enzymes, while at the higher level it is how well it explains the origin of the more complex metabolic functions that emerge when enzyme functions are combined to form metabolic pathways, and the integrated networks of pathways that constitute metabolism as a whole. Notice that natural selection relates more directly to the higher level, in that this is where phenotypic traits are manifested, whereas mutation relates more directly to the lower level, in that individual mutations typically alter single genes, and therefore single enzymes. The perennial challenge for neo-Darwinism has been to explain how mutation and selection, two disparate phenomena operating at different levels, can combine to produce such spectacular functional innovations at both levels.

The hope has always been that explaining evolutionary innovation at the level of single genes would eventually simplify the task of explaining innovation at the level of complete pathways. That reductionistic hope seems to be fading. Even at the level of single genes, explaining innovation is growing harder, not easier, as more and more distinct protein structures are discovered. The count of fundamentally distinct structures, or *folds*¹ as they are known, now stands at about 2,000, with more being added every year.

The extraordinary difficulty that neo-Darwinism encounters with single-gene innovations requiring a new protein fold has recently been described in detail [1]. That raises an obvious question. If the Darwinian mechanism cannot reliably explain innovation at the level of a single protein fold, what *can* it explain? This prompted us to investigate the more modest case of enzymatic innovation within a fold family, which we regard as metabolic innovation on the smallest scale possible.² With that aim, we attempted to modify one particular bacterial enzyme so as to make it perform the function of another that closely resembles it [2]. Although we were ultimately unable to achieve this functional conversion, extensive testing of the kinds of amino-acid substitutions that ought to promote it

¹Proteins have three-dimensional folded structures that determine their function. Those with secondary structural elements (alpha helices and beta strands) in the same order and similar spatial arrangement are said to have a *common* fold, or in other words, to be members of the same fold family. Proteins with fundamentally *distinct* folds differ in the arrangements of secondary structural elements and/or in their order.

²Although adaptations can certainly occur on a smaller scale, 'innovation' refers to the first-time appearance of a genuinely new function, not the adjustment of an existing function.

demonstrated that success would, for our test case, require many more specific changes than the Darwinian mechanism can accomplish, even over billions of years.

It would be tempting to disregard that result if there were a body of contrary evidence. Instead, as we have discussed [2], our result is just one contribution to a consistent picture based on numerous studies (see below). No one denies the possibility of converting enzymes to new functions, but it seems that anyone attempting it with the assumption that it can be done with just a few nucleotide changes is in for a surprise.

Where to go from here is a matter of perspective. Darwin's theory certainly will not benefit from ignoring or denying the severity of the problems that have beset it. Once that is conceded, the most important question is whether the theory needs to be remedied or replaced. Among the things that will be needed to answer that question is a full picture of what has gone wrong with the standard evolutionary account. In other words, it will be increasingly helpful to go beyond a mere catalog of inexplicable facts to something more like a *synthesis* of the whole problem. We use the word 'helpful' here because a synthesis of this kind should, we think, be the start of something much more positive than the dismantling of an old theory. It should instead be seen as an opportunity to gain key insights for constructing a new theory by building a clear understanding of how the old theory went wrong.

With that in mind, we here take a step toward such a synthesis by describing briefly the general aspects of metabolic innovation that most profoundly challenge the current neo-Darwinian model. The aspects are logically separable, which allows them to be examined as distinct topics, but their effects are highly interconnected. We will show this by developing a synthesis of the whole problem in a progressive way as each aspect is considered. Based on this critical synthesis we then offer the beginnings of a *positive* synthesis — a set of principles that hint at a new theory of innovation. The ultimate aim, of course, is to develop a theoretical framework from which to understand all biological innovation. Metabolic innovation will admittedly be only a small part of that big picture, but its relative simplicity makes it a promising *first* part for getting the whole project underway.

As should now be obvious, this paper is written primarily for readers who are willing to at least consider the possibility that Darwin's theory might be fundamentally deficient as an explanation for innovation in the history of life. We recognize that a great many talented biologists may not place themselves in that category, but we think the time is right for the evidential case against the standard Darwinian model to be presented in order to begin a serious discussion of the alternatives.

Problem 1: Offsetting the cost of gene expression

The most widely accepted explanation for the origin of new enzymes is gene duplication and recruitment [3, 4]. This process involves duplication of an existing gene, followed by divergent evolution of one of the copies to a new function. For this process to work, though, the diverging duplicate must continue to be transcribed and translated. But these processes of gene expression carry a resource cost [5–8]. Consequently, a duplicate gene undergoing divergent evolution will only confer a net benefit if that cost is more than offset by its positive biological contribution. In many cases this makes cost reduction by deletion or inactivation of the duplicate gene much more likely than innovation as an adaptive response. Several recent papers have demonstrated this by finding that cells reduce expression of nonessential or duplicate genes, or completely inactivate them, in competitive environments [8–12]. When under continuous selection for metabolic efficiency, such as when growing under nutrient-limiting conditions, cells that reduce the total cost of gene expression by inactivating or deleting unneeded genes have a significant fitness advantage and can quickly overtake the population [8, 10].

In judging the degree to which the cost of gene expression impedes metabolic innovation, it is particularly important to distinguish natural selection from laboratory selection. Reported experimental conversions of two enzymes to *o*-succinylbenzoate synthase (OSBS) activity illustrate this point. Working with an *Escherichia coli* (*E. coli*) strain in which the chromosomal gene encoding OSBS was deleted, Schmidt and coworkers identified single mutations that enable two other genes to replace this missing function well enough for selection *in vivo* under specified laboratory conditions. Among those conditions, though, was high-level expression of the replacing gene,³ which was needed in order to compensate for the very low activity of the converted function (0.0004% or 0.06% of wild-type activity based on k_{cat}/K_m , depending on the source gene [13]). Even with the boosted expression, though, the converted genes fell well short of fully restoring growth [13]. So while the enzyme conversions reported in that study provide useful information, it should not be assumed that they would succeed in nature.

Considering that newly evolved functions are likely to be extremely weak, it should be expected that they would need amplified expression in order to be of any use. But if so, the expression cost might easily outweigh any functional benefit. Natural genes, of course, escape this dilemma by having extremely high catalytic proficiencies and by minimizing expression costs through regulated expression (turning expression off when it is not needed).

³Achieved with an induced *tac* promoter on a multi-copy plasmid [13]. For vector details, see <http://www.ncbi.nlm.nih.gov/entrez/viewer.fcgi?db=nucore&id=45614>.

First obstacle: *Because gene expression is costly, it cannot be assumed that weakly converted enzyme functions isolated by laboratory selection would provide net selective benefit in wild populations.*

Problem 2: Winning the fixation lottery

Bacteria reproduce rapidly enough to exhaust any pool of nutrients, no matter how large, in a short time frame. This means that local extinction (by starvation) figures much more prominently in the dynamics of bacterial populations than it does for higher organisms. Many bacterial cells alive now, for example, will manage to have billions of descendants alive a year from now. But for each of these cells destined for success, billions in the current population are destined to expire in that time frame, leaving no descendants. Thus, losers in the game of bacterial procreation vastly outnumber winners.

The overall effect of these frequent local extinctions or near-extinctions is a dramatic reduction in genetic variability, which means a dramatic increase in the time required for rare genotypic variants to become *fixed* (i.e., to become the new wild-type). In population genetics, the parameter that characterizes this phenomenon is the *effective population size*, N_e . Roughly speaking, N_e is the size of the subpopulation in each generation that will influence the genetic makeup of future generations. So the smaller N_e is relative to the true population size, N , the more rare winners are in the propagation lottery.

The estimated value of N_e for wild bacterial populations is 10^9 [14, 15], roughly eleven orders of magnitude lower than estimates of N [16]. Consequently, particular beneficial mutations have to appear on the order of 10^{11} times before they have any reasonable likelihood of being fixed. And because that likelihood scales with the coefficient of selection, s [17], which is commonly assumed to have a small fractional value, something like 10^{12} or more appearances may be needed in order for fixation to become probable. In a population of 10^{20} organisms that passes through 10^3 generations per year [18], this does not prevent fixation of common mutations. A beneficial mutation that occurs once in 10^9 cells, for example, will appear 10^{11} times per generation, which means that a cell line destined to carry this mutation to fixation will probably be present within roughly 10 generations. But the situation changes for rare mutations or rare combinations of mutations. At an incidence rate of one new carrier in the population per generation, some 10^{12} generations ($\sim 10^9$ years) would be required for fixation to become likely, even though the genotype in question exists somewhere in the population most of the time.

Second obstacle: *Beneficial mutations appearing less than about once per generation in a global bacterial population may remain unfixed for a billion years or more.*

Problem 3: Complex adaptation — Combining rare genetic events

From here on it will become increasingly apparent that each of the problems we describe is compounded by the others. If new enzyme functions can evolve by consecutive adaptive mutations,⁴ each known to occur spontaneously with reasonable frequency, then Problem 2 would be of no consequence. The difficulty arises from the fact that they typically appear *not* to be achievable in this way.

As mentioned in the introduction, when we attempted to convert an enzyme to perform a new function, we found it to be surprisingly difficult [2]. The starting point was an enzyme we designated Kbl₂ (2-amino-3-ketobutyrate CoA ligase), and the target function was that of BioF₂ (8-amino-7-oxononanoate synthase). The structures of Kbl₂ and BioF₂ are so similar (Fig. 1) that the enzymes are commonly assumed to be close evolutionary relatives. However, after extensive testing of mutations that were carefully chosen for their potential to achieve the desired conversion, we found success to be elusive. Still, we were able to deduce from our results that the shortest path to conversion would involve seven or more mutations. That is, at least seven mutations would be required before any level of the new function would be achieved. The true number is probably much higher, considering that we introduced many more than seven substitutions without success. But seven is high enough to cause a severe problem. Mathematical analysis shows that even this seemingly modest number of mutations places the conversion *well* beyond what neo-Darwinian evolution can explain (Fig. 2) [2, 21].

There is an understandable tendency for defenders of a theory, when faced with challenging evidence like this, to marshal as much opposing evidence as can be found. Indeed, if there were a solid body of evidence showing that genuine conversions of enzyme function usually *are* achievable with one or two nucleotide substitutions, we would conclude that the case we examined happened to be exceptionally problematic. But the result of our study is actually quite consistent with the whole body of work on functional conversions in enzymes, even as others have summarized it. For example, two well-known contributors to the field, John Gerlt and Patricia Babbitt, recently gave this sobering assessment of the field:

Interchanging reactions catalyzed by members of mechanistically diverse superfamilies might be envisioned as “easy” exercises in (re)design: if Nature did it, why can’t we? ...Anecdotally, many attempts at interchanging activities in mechanistically diverse superfamilies have since been attempted, but few successes have been realized [22].

⁴Adaptive mutations are those that increase the fitness of the organism that carries them, meaning that the organism can grow and reproduce faster than its neighbors. Most mutations are neutral or deleterious.

Fig. 1. Structural similarity of BioF and Kbl. a) Dimeric enzymes BioF₂ (left; 1DJ9 [17]) and Kbl₂ (right; 1FC4 [18]) viewed along axes of symmetry. Active sites are at the monomer interfaces. b) Aligned backbones of BioF and Kbl monomers. c) Identical side chains in the BioF₂ (dark) and Kbl₂ (light) active sites, labeled according to BioF positions. PLP-external aldimines are shown in the center of the active sites. This figure was originally published as Fig. 5 in reference 2.

Similarly, Philip Romero and Frances Arnold drew the conclusion that many researchers (including us) have reached:

Some functions, however, simply cannot be reached through a series of small uphill steps and instead require longer jumps that include mutations that would be neutral or even deleterious when made individually. Examples of functions that might require multiple simultaneous mutations include the appearance of a new catalytic activity... [23]

Apart from neo-Darwinian expectations, perhaps the difficulty of enzyme conversion should not have been a surprise. The information content of an enzyme is quite large. Its one-dimensional protein sequence bears a complex causal relationship to

Fig. 2. Expected waiting times for an enzyme conversion requiring from seven to twelve specific base changes. The assumed starting point is a population lacking a duplicate version of the gene to be converted. As discussed (Problem 1), cells in which a duplicate appears are disadvantaged by the cost of expressing a raw duplicate. Shown are the predicted times for a 1% fitness cost ($\bar{s} = -0.01$; top line), a 0.1% fitness cost ($\bar{s} = -0.001$; second from top), and a 0.01% fitness cost ($\bar{s} = -0.0001$; third from top), and no cost (bottom line). Other parameter values are as listed in Table 1 of [19]. The dashed line marks the boundary between feasible waiting times (below) and waiting times that exceed the age of life on earth (above), assuming 10^3 generations per year. This figure was originally published as Fig. 11 of reference 2.

its three-dimensional folded structure *and* to its dynamic behavior as an enzyme. Its activity depends upon many distinct and context-dependent interactions that enable it to form a stable folded structure and to carry out its chemistry. Converting an enzyme to a new catalytic activity is therefore likely to require the simultaneous reconfiguration of many amino-acid interactions, so any step-wise process of enzymatic conversion almost inevitably will involve non-functional intermediates.

In the end, two things seem inescapable. One is that enzymatic innovations requiring more than two specific mutations in a spare gene (provided by a duplication event) are implausible in neo-Darwinian terms [21]. The other is that once this limitation is taken into account, most reported experimental conversions of enzyme function are beyond the reach of neo-Darwinian processes under natural conditions.

Third obstacle: Adaptations requiring duplication and modification of an existing gene should not be presumed feasible if they require more than two specific base substitutions, which seems to exclude most functional conversions.

Problem 4: The complexity of metabolic pathways

The severe challenge to the Darwinian model posed by the first three problems becomes exponentially more severe when we recognize that the relevant scale of genetic innovation is not a single new enzyme function, but rather the coordinated sequence of enzymatic steps needed to produce a new phenotypic trait. Our reported attempt to change Kbl₂ into a BioF₂-like enzyme in *E. coli* illustrates this point [2]. To make selection of successful mutants possible, one of us (AG) engineered a strain that lacks the gene encoding BioF₂. Without that gene the engineered strain is unable to make biotin, an essential cofactor for fatty acid biosynthesis and other carboxylation reactions [24–26]. This makes growth impossible unless either functional conversion is achieved (which never happened) or biotin is supplied as a nutrient (which is how we maintained the strain). This suited our experimental objectives well, but it is important to recognize that our engineered strain is *wholly* unrealistic as a natural evolutionary context for the origin of BioF₂.

The complete metabolic pathway for biotin synthesis (Fig. 3) shows why this is so. BioF₂ is just one of four enzymes that are exclusively dedicated to biotin production. This means that any proposed explanation of the origin of biotin production as a phenotypic trait must account for innovation on a considerably larger scale than the already problematic scale of a single functional conversion. The full impact of this becomes evident when we realize that quadrupling the scale of a complex adaptation increases the evolutionary difficulty not merely by a *factor* of four, but rather by a *power* of four [21].⁵

The biotin example illustrates the problem of pathway complexity nicely, but is it typical or exceptional for metabolic pathways to depend on four dedicated enzymes? To answer this we need to examine the whole metabolic picture. When we do this, we see that the biotin pathway is unexceptional in its complexity. According to EcoCyc, a comprehensive database of metabolic information on *E. coli*, this common bacterium uses 1,467 enzymes to carry out the functions of 281 metabolic pathways.⁶ That amounts to just over five enzymes per pathway, on

⁵More precisely, it increases the required probabilistic resources (opportunities for success) by a power of four, which would increase the waiting time in generations by *more* than a power of four (assuming each generation provides multiple opportunities for success).

⁶See <http://ecocyc.org/ECOLI/organism-summary?object=ECOLI>.

Fig. 3. The dedicated pathway for microbial biotin biosynthesis.

average. Similarly, in 2001 Teichmann *et al.* reported 581 proteins used in 106 small-molecule metabolic pathways in *E. coli* [27]. Although the definition of “pathway” is somewhat imprecise, these figures give us at least a rough picture of the complexity of metabolic processes in terms of enzymatic steps, and from that picture we deduce that most of the innovations that brought new metabolic traits did indeed involve multiple enzymatic innovations.

This poses a severe challenge for neo-Darwinism. Mechanisms that have been proposed in attempts to meet this challenge, such as retrograde evolution [28], or serial duplication and recruitment [29] do not match the actual distribution of protein domains across and within pathways [30]. Rather, most pathways employ several different protein folds, which, as we discuss next, raises another problem.

Fourth obstacle: Accounts of metabolic innovation must recognize that beneficial metabolic traits typically depend on multiple dedicated genes.

Problem 5: Radical innovation — the need for new protein folds

The previous problem makes it clear that a realistic treatment of metabolic innovation has to explain more than a single new enzyme function. Explaining how a new enzyme function might appear is a key part of the problem, but it is not the *whole* problem for several reasons. The first, as just discussed, is that new metabolic

traits typically require multiple new enzyme functions, not just one. The second is that these new functions often call for new protein folds, which adds the problem of *structural* innovation to the already mushrooming problem of functional innovation. The problem of converting an existing fold to a new function is very modest compared to the problem of generating a stable new fold with enzymatic activity from scratch.

The basis for thinking that such structural innovation is typically beyond the reach of Darwinian evolution has been described [1]. The next question is how prevalent structural innovation appears to have been in the early history of life. More specifically, how often did metabolic innovation involve structural innovation? We can get at least a rough answer to this in a couple of different ways. One is to estimate the number of distinct fold types used by a typical bacterial species and divide that by the number of metabolic pathways that these folds serve. This avoids the need to reconstruct history by giving us an average value — the average number of new folds that have to be explained per pathway explanation. A previous analysis found this average to be about four (991 distinct folds serving 263 pathways [1]), which means that the vast majority of early metabolic pathways required new folds.⁷

A complementary approach is to get a rough lower-bound estimate of the total number of distinct protein folds used in bacterial life. Analysis of the bacterial genomes that have been sequenced so far indicates that a substantial majority (>80%) of the 1,962 known protein folds are used in at least one bacterial species.⁸ Although there is no reliable way to estimate the actual total number of folds, that result suggests that bacterial life uses most of them. Currently, about 40% of the proteins known to exist are known only by the sequence of their encoding gene (i.e., nothing is known of their structure or function [31]). As more genomes are sequenced, the list of these uncharacterized proteins continues to grow, and again a substantial fraction of them (about 50%) are of bacterial origin [31]. A concerted effort has been made in recent years to target these proteins for structural analysis, with interesting results. Of 248 newly determined structures described by Jaroszewski *et al.* [31], 44 are completely new folds, and another 23 have only partial similarity to known folds. Thus, the folds that have been identified so far may be only the tip of a very large ‘iceberg.’

Fifth obstacle: *Accounts of metabolic innovations must recognize that they often depend on new protein folds.*

⁷Using the Poisson distribution with an expectation of $991/263 = 3.8$ new folds per pathway gives a 98% likelihood of at least one new fold having been needed for a randomly chosen pathway.

⁸Based on analysis of *Superfamily* assignments for 1,392 bacterial genomes (version 1.75; see <http://supfam.cs.bris.ac.uk/SUPERFAMILY/>).

Problem 6: Causal circularity

Kun, Papp, and Szathmáry have described the problem of “kick-starting metabolic networks” [32]. Their abstract begins, “If chemical *A* is necessary for the synthesis of more chemical *A*, then *A* has the power of replication.” Accordingly, they apply the term “autocatalytic” to *A*. To avoid confusion, we suggest that this term ought to be reserved for cases where *A* is *sufficient* for the production of itself (with no extraordinary preconditions). By contrast, *A* being *necessary* for making *A* does not mean that supplies of *A* are self-renewing. Rather, it means that the absence of *A* assures its continued absence. We will use the term *causal circularity* to describe this case.

Whenever a biosynthetic process exhibits causal circularity (requiring its product, *A*), selection-based accounts of the origin of this process encounter complications. In the first place, since the biosynthesis of *A* as we now see it requires not just the genes encoding the enzymes that produce *A* but also *A* itself, a satisfactory account has to go beyond gene origins. The current biosynthetic apparatus for making *A* must, in such a case, not only come into existence but also be primed with pre-existing *A* in order to begin working. But this presents another complication. If *A* was pre-existing, how would acquiring a way of making *A* provide a selective advantage? Although it is possible to construct answers to this, they all suppose circumstances beyond the simple fact that *A* is useful, which makes the final explanation only as compelling as those suppositions are.

How common is causal circularity, though? By analyzing metabolic network models for various microbial species, Kun and coworkers showed that ATP production involves causal circularity in all organisms, with other metabolites showing circularity in some organisms but not in others [32]. However, because their analysis focused on net reactions rather than on the actual physical requirements for them to occur, they may have underestimated the generality of this phenomenon.

A few examples will illustrate this. One is the biosynthesis of cysteine in bacteria. The reactant that provides the sulfur atom for incorporation into cysteine is hydrogen sulfide (H_2S),⁹ which itself must be produced from sulfate (SO_4^{2-}) in a multi-step enzymatic process.¹⁰ The final step of this process is catalyzed by sulfite reductase, an enzyme that depends upon a prosthetic group consisting of four iron atoms bridged by four inorganic sulfur atoms (an Fe_4S_4 iron-sulfur cluster [33]) and coordinated to the protein by means of four cysteine side chains (Fig. 4).

⁹<http://BioCyc.org/ECOLI/NEW-IMAGE?type=PATHWAY&object=CYSTSYP-WY>.

¹⁰<http://BioCyc.org/ECOLI/NEW-IMAGE?type=PATHWAY&object=SO4ASSIM-PWY>.

Fig. 4. The enzyme sulfite reductase. As shown in the expanded view, the active site of sulfite reductase uses two prosthetic groups. The larger of these is siroheme (dark honeycomb structure). Coupled down below the iron center of siroheme is the cube-like iron-sulfur cluster, which is held in place by four cysteine side chains.

Consequently, without those coordinating cysteine residues, sulfite reductase cannot produce H_2S , and without H_2S , cysteine synthase cannot produce cysteine. So cysteine biosynthesis is a striking example of causal circularity. Other amino acid pathways provide additional examples. The biosynthesis of arginine depends on ornithine carbamoyltransferase (ArgF),¹¹ which has an essential arginine residue in its active site [34], and the biosynthesis of lysine depends on diaminopimelate decarboxylase (LysA),¹² which requires a lysine residue to form a Schiff-base linkage to its PLP prosthetic group.¹³

In fact, there is a simple way to generalize the principle of causal circularity. Since life is a prerequisite for *all* biosynthesis, any biosynthetic product that is necessary for life in its present form is also necessary for its own biosynthesis in modern life. So causal circularity exists for *all* essential biosynthetic products. In some cases the loop is extremely tight. LysA, for example, embodies a causal loop in itself by both producing and requiring lysine directly. More often the causal loop involves multiple activities. Biotin production is a good example of this,

¹¹<http://biocyc.org/ECOLI/NEW-IMAGE?type=PATHWAY&object=ARGSYN-PWY&detail-level=2>.

¹²<http://biocyc.org/ECOLI/NEW-IMAGE?type=PATHWAY&object=DAPLYSINESYN-PWY>.

¹³See PDB entry 1KNW and reference 35.

biotin being necessary for fatty acid biosynthesis, which is necessary for building the cell membrane, which is necessary for life, which is necessary for the biosynthesis of everything, including biotin.

So, in order to conceive of an evolutionary origin of biotin biosynthesis, we must suppose that prior to this origin either A) cells were making their membranes without biotin, or B) cells had an abiotic source of biotin. Either way, the question of how the ability to make biotin would have been beneficial is raised. To answer it, we have to contrive a selective scenario that goes well beyond plain facts, which means we end up having to justify both a contrived selection story and a seemingly unlikely supposition (either A or B) about the state of life prior to biotin biosynthesis. Of course it is *possible* to suppose any number of additional things in an attempt to do this, but each of these suppositions adds to the complication of an already complicated story.

Sixth obstacle: *The fact that life depends on numerous components jointly means that no simple relationship exists between the functions of these components and the selective story that would be needed for them to have arisen as simple adaptations.*

Discussion

When the key shortcomings of neo-Darwinism are examined in any detail, it is hard to escape the impression that the theory is unraveling. All theories encounter unsolved problems, but for a solid theory these are challenges in the positive sense of the word — opportunities to prove itself further. With neo-Darwinism, on the other hand, things appear to be moving in the other direction. As we learn more about biological systems, we encounter apparently insoluble problems at every level. To make matters worse, as we have seen here the interdependence of these individual failures compounds them greatly, making repair of the theory seem very unlikely.

As negative as this may sound, it has a positive side: the insights we gain from identifying the obstacles facing neo-Darwinism can and should inform the construction of a new theory to take its place. That is, in pinpointing the key problems with the old theory we are identifying crucial respects in which its replacement must differ from it. We ourselves have become convinced that intelligent causation is essential as a starting point for any successful theory of biological innovation. If this is so, what is needed now is an elaboration of the general principles by which living things have been designed. Accordingly, we have attempted to

identify design principles from each of the problems described above. The six principles, paired with the obstacles they address, are as follows:

First obstacle: *Because gene expression is costly, it should not be assumed that weakly converted enzyme functions isolated by laboratory selection would provide net selective benefit in wild populations.*

First principle: *Innovations are more like investments than quick cash. They must be well implemented to offset their cost, and even then the benefits tend to accrue over a long period.*

Second obstacle: *Beneficial mutations appearing less than about once per generation in a global bacterial population may remain unfixed for a billion years or more.*

Second principle: *For innovations to be established reliably they need to be carried past a critical 'tipping point' in numerical representation, beyond which they become self-establishing.*

Third obstacle: *Adaptations requiring duplication and modification of an existing gene should not be presumed feasible if they require more than two specific base substitutions, which seems to exclude most functional conversions.*

Third principle: *The substantial reworking of a homologous structure needed to give it a genuinely new function is more suggestive of reapplication of a concept than adjustment of a physical thing.*

Fourth obstacle: *Accounts of metabolic innovation must recognize that beneficial metabolic traits typically depend on multiple dedicated genes.*

Fourth principle: *Useful innovations tend to require the simultaneous solution of multiple new problems, which means they tend to be compound innovations.*

Fifth obstacle: *Accounts of metabolic innovations must recognize that they often depend on new protein folds.*

Fifth principle: *Useful innovations often involve both the reapplication of proven design concepts and the de novo invention of new ones.*

Sixth obstacle: *The fact that life depends on numerous components jointly means that no simple relationship exists between the functions of these components and the selective story that would be needed for them to have arisen as simple adaptations.*

Sixth principle: *The implementation of innovation is nearly the opposite of ordinary physical causation. It is the top-down arrangement of matter in such a way*

that the resulting bottom-up behavior of that matter serves the intended purpose of the innovator.

Even in this rough form these principles suggest some interesting things. One is that biological innovation seems similar in essence to human innovation, though certainly beyond it in degree. This realization is attracting an increasing number of engineers to biology with the aim of reapplying biological innovations in human technology [36]. Although that field of study, known as *biomimetics*, has practical ambitions, the fact that it exists (and is thriving) also implies an essential similarity between intelligent design in engineering and intelligent design in life.

Another interesting aspect of the above set of principles is that while they were drawn from observations at the molecular scale of metabolic innovation, they do not appear to be restricted to that scale. Indeed, they have the appearance of general rules that make sense irrespective of the particulars of the innovation, including its physical scale. Since universality of that kind is precisely what we expect of a useful theory, this suggests that the principles may be a starting point for framing the first successful theory of biological innovation.

Next, and perhaps most significantly, it is clear that this new theory will be of an entirely different *kind* than the one it hopes to replace. Darwinism is purely mechanistic in its approach, in that it offers a bottom-up causal explanation for the origin of all biological forms and phenomena. In this respect it is also intrinsically reductionistic — it takes physical causation to be the *fundamental* explanation of all origins events. The design-based theory hinted at in this paper will differ radically in both respects. The new theory will be fundamentally top-down in its approach and therefore fundamentally non-reductionistic. It will focus mainly on design principles rather than on mechanisms. Just as students of engineering and design focus mainly on high-level principles that leave a great deal of freedom as to their physical implementation, so too students of the new theory will focus mainly on the principles that inform biological designs [37] rather than on the processes by which these designs may be implemented.

Might this new theory transform biology beyond the topic of origins? Most who reflect on the current state of biology sense a need for understanding to catch up with the enormous flow of new data. Sydney Brenner, one of the pioneers of modern molecular biology, has concluded that “biology urgently needs a theoretical basis to unify it and it is only theory that will allow us to convert data to knowledge [38].” He continues by pointing out that the trend toward performing measurements on whole systems instead of their isolated parts (one of the emphases of *systems biology*) brings us no closer to the needed theory,

but his suggestion that we should return to hard-core reductionism also misses the mark:

Our approach directly reflects the structure of biological systems and, as we reduce each level to the level below — organisms to cells and cells to molecules — we can then confidently complete the reductionist programme because the properties of molecules can be reduced to physics [38].

The problem with this approach is that reducing a living thing to its simplest material causes does not lead to an understanding of it. By way of analogy, those who want to understand software should have some exposure to the zeros and ones of machine language, but they would do well to spend *most* of their time studying principles of software design that are nowhere to be found among the bits. More generally, one can acquire a great deal of knowledge of the *operation* of a complex innovative system without having the slightest grasp of the genius behind it. To grasp that, we need to consider how it was designed.

In the end Brenner's search for a new theory seems to be hamstrung by the old theory. He thinks "we need to remember that whereas mathematics is the art of the perfect and physics is the art of the optimal, biology, because of evolution, is only the art of the satisfactory [38]." We think it may actually be much more than that.

References

1. Axe DD (2010) The case against a Darwinian origin of protein folds. *BIO-Complexity* 2010(1):1–12. doi:10.5048/BIO-C.2010.1
2. Gauger AK, Axe DD (2011) The evolutionary accessibility of new enzyme functions: a case study from the biotin pathway. *BIO-Complexity* 2011(1): 1–17. doi:10.5048/BIO-C.2011.1
3. Ohno S (1970) *Evolution by gene duplication*. Springer-Verlag, New York.
4. Jensen RA (1976) Enzyme recruitment in evolution of new function. *Annu Rev Microbiol* 30: 409–425.
5. Dykhuizen D (1978) Selection for tryptophan auxotrophs of *Escherichia coli* in glucose-limited chemostats as a test of the energy conservation hypothesis of evolution. *Evolution* 32: 125–150. doi:10.2307/2407415.
6. Stoebel DM, Dean AM, Dykhuizen DF (2008) The cost of expression of *Escherichia coli* lac operon proteins is in the process, not in the products. *Genetics* 178: 1653–1660. doi:10.1534/genetics.107.085399.

7. Wagner A (2005) Energy constraints on the evolution of gene expression. *Mol Biol Evol* 22: 1365–1374. doi:10.1093/molbev/msi126.
8. Gauger AK, Ebnet S, Fahey PF, Seelke R (2010) Reductive evolution can prevent populations from taking simple adaptive paths to high fitness. *BIO-Complexity* 2010(2):1–9. doi:10.5048/BIO-C.2010.2
9. Mira A, Ochman H, Moran NA (2001) Deletional bias and the evolution of bacterial genomes. *Trends Genet* 17: 589–596. doi:10.1016/S0168-9525(01)02447-7
10. Lang GI, Murray AW, Botstein D (2009) The cost of gene expression underlies a fitness trade-off in yeast. *Proc Natl Acad Sci USA* 106: 5755–5760. doi:10.1073/pnas.0901620106.
11. Kuo C-H, Ochman H (2010) The extinction dynamics of bacterial pseudogenes. *PLoS Genet* 6(8): e1001050. doi:10.1371/journal.pgen.1001050.
12. Behe M (2010) Experimental evolution, loss-of-function mutations, and “the first rule of adaptive evolution.” *Q Rev Biol* 85: 419–445.
13. Schmidt DM, Mundorff EC, Dojka M, Bermudez E, Ness JE, *et al.* (2003) Evolutionary potential of (β/α)₈-barrels: functional promiscuity produced by single substitutions in the enolase superfamily. *Biochemistry* 42: 8387–8393. doi:10.1021/bi034769a.
14. Lynch M (2005) Simple evolutionary pathways to complex proteins. *Protein Sci* 14: 2217–2225. doi:10.1110/ps.041171805.
15. Lynch M, Conery JS (2003) The origins of genome complexity. *Science* 302: 1401–1404. doi:10.1126/science.1089270.
16. Milkman R, Stoltzfus A (1988) Molecular evolution of the *Escherichia coli* chromosome. II. Clonal segments. *Genetics* 120: 359–366.
17. Lenski RE, Rose MR, Simpson SC, and Tadler SC (1991) Long-term experimental evolution in *Escherichia coli*. I. Adaptation and divergence during 2,000 generations. *The American Naturalist* 138: 1315–1341. doi:10.1086/285289.
18. Levin BR (1981) Periodic selection, infectious gene exchange and the genetic structure of *E. coli* populations. *Genetics* 99: 1–23.
19. Alexeev D, Alexeeva M, Baxter RL, Campopiano DJ, Webster SP *et al.* (1998) The crystal structure of 8-amino-7-oxononanoate synthase: a bacterial PLP-dependent, acyl-CoA-condensing enzyme. *J Mol Biol* 284: 401–419. doi:10.1006/jmbi.1998.2086.
20. Schmidt A, Sivaraman J, Li Y, Larocque R, Barbosa J *et al.* (2001) Three dimensional structure of 2-amino-3-ketobutyrate CoA ligase from *Escherichia coli* complexed with a PLP-substrate intermediate: Inferred reaction mechanism. *Biochem* 40: 5151–5160. doi:10.1021/bi002204y.
21. Axe DD (2010) The limits of complex adaptation: an analysis based on a simple model of structured bacterial populations. *BIO-Complexity* 2010(4):1–10. doi:10.5048/BIO-C.2010.4.
22. Gerlt JA, Babbitt PC (2009) Enzyme (re)design: Lessons from natural evolution and computation. *Curr Opin Chem Biol* 13(1):10–18. doi:10.1016/j.cbpa.2009.01.014.

23. Romero and Arnold (2009) Exploring Protein Fitness Landscapes by Directed Evolution. *Nature Reviews: Molecular Cell Biology* 10:866–875. doi:10.1038/nrm2805.
24. Del Campillo-Campbell A, Kayajianian G, Campbell A, Adhya, S (1967) Biotin requiring mutants of *Escherichia coli* K-12. *J Bacteriol* 94: 2065–2066.
25. Rolfe B, Eisenberg MA (1968) Genetic and biochemical analysis of the biotin loci of *Escherichia coli* K-12. *J Bacteriol* 96: 515–524.
26. Streit WR, Entcheva P (2003) Biotin in microbes, the genes involved in its biosynthesis, its biochemical role and perspectives for biotechnological production. *Appl Microbiol Biotechnol* 61: 21–31.
27. Teichmann SA, Rison SCG, Thornton JM, Riley M, Gough J (2001) Small-molecule metabolism: an enzyme mosaic. *Trends Biotech* 19: 482–486. doi:10.1016/S0167-7799(01)01813-3.
28. Horowitz NH (1945) On the evolution of biochemical syntheses. *Proc Natl Acad Sci USA* 31:152–157. doi:10.1073/pnas.31.6.153.
29. Horowitz NH (1965). The evolution of biochemical syntheses - retrospect and prospect. In: Bryson V, Vogel HJ (eds) *Evolving genes and proteins*, pp 15–23, Academic Press, New York.
30. Rison SCG, Thornton JM, Riley M, Gough J, Chothia C (2001) The evolution and structural anatomy of the small molecule metabolic pathways in *Escherichia coli*. *J Mol Biol* 311: 693–708. doi:10.1006/jmbi.2001.4912.
31. Jaroszewski L, Li Z, Krishna SS, Bakolitsa C, Wooley J (2009) Exploration of uncharted regions of the protein universe. *PLoS Biol* 7: e1000205. doi:10.1371/journal.pbio.1000205.
32. Kun A, Papp B, Szathmáry E (2008) Computational identification of obligatorily autocatalytic replicators embedded in metabolic networks. *Genome Biol* 9: R51.
33. Crane BR, Siegel LM, Getzoff ED (1995) Sulfite reductase structure at 1.6 Å: Evolution and catalysis for reduction of inorganic anions. *Science* 270: 59–67. doi:10.1126/science.270.5233.59.
34. Kuo LC, Miller AW, Lee S, Kozuma C (1988) Site-directed mutagenesis of *Escherichia coli* ornithine transcarbamoylase: Role of arginine-57 in substrate binding and catalysis. *Biochem* 27: 8823–8832. doi:10.1021/bi00424a021.
35. Momany C, Levnikov V, Blagova L, Crews K (2002) Crystallization of diaminopimelate decarboxylase from *Escherichia coli*, a stereospecific D-amino-acid decarboxylase. *Acta Crystallogr D Biol Crystallogr* 58: 549–552. doi:10.1107/S0907444902000148
36. Bar-Cohen Y (2006) *Biomimetics: biologically inspired technologies*. CRC Press, Taylor and Francis Group, Boca Raton.
37. Csete ME, Doyle JC (2002) Reverse engineering of biological complexity. *Science* 295:1664–295. doi: 10.1126/science.1069981.
38. Brenner S (2010) Sequences and consequences. *Phil Trans R Soc B* 365: 207–212. doi:10.1098/rstb.2009.0221.