

Distributed robustness versus redundancy as causes of mutational robustness

Andreas Wagner

Summary

A biological system is robust to mutations if it continues to function after genetic changes in its parts. Such robustness is pervasive on different levels of biological organization, from macromolecules to genetic networks and whole organisms. I here ask which of two possible causes of such robustness are more important on a genome-wide scale, for systems whose parts are genes, such as metabolic and genetic networks. The first of the two causes is redundancy of a system's parts: A gene may be dispensable if the genome contains redundant, back-up copies of the gene. The second cause, distributed robustness, is more poorly understood. It emerges from the distributed nature of many biological systems, where many (and different) parts contribute to system functions. I will here discuss evidence suggesting that distributed robustness is equally or more important for mutational robustness than gene redundancy. This evidence comes from the functional divergence of redundant genes, as well as from large-scale gene deletion studies. I also ask whether one can quantify the extent to which redundancy or distributed robustness contribute to mutational robustness. *BioEssays* 27:176–188, 2005. © 2005 Wiley Periodicals, Inc.

Introduction

Living things are unimaginably complex, yet also highly robust to genetic change on all levels of organization. Proteins can tolerate thousands of amino acid changes, metabolic networks can continue to sustain life even after removal of important chemical reactions, gene regulation networks continue to function after alteration of key gene interactions, and radical genetic change in embryonic development can lead to an essentially unchanged adult organism.^(1–4) The mutational robustness of organisms is an intriguing phenomenon, not only because complexity suggests fragility. (If you have ever built a house of cards, you will know what I mean.) Mutational robustness also affects an organism's ability for future evolution, because it may increase the rate of neutral

mutations, and thus the amount of genetic variation invisible to natural selection.

Such cryptic genetic variation is ubiquitous and it can affect phenomena as different as variation in genetic disease phenotypes and reproductive isolation.⁽⁵⁾ It is thus not surprising that biologists have a long-standing interest in robustness, going back to Fisher's work on dominance,⁽⁶⁾ as well as to Waddington's work on developmental canalization.^(7,8) However, an understanding of the root causes of this phenomenon had to wait for the mechanistic understanding of living systems afforded by molecular biology. Here, I will explore a very basic question about mutational robustness: what are the mechanistic causes of mutational robustness, and which of these causes is most important? This question is one of the most fundamental questions one can ask about the organization of genetic systems. I will begin by illustrating the two principal causes of mutational robustness, redundancy and distributed robustness. Most of my examples regard systems above the gene level. However, much of the following is equally relevant for systems on lower levels of biological organization, especially macromolecules like protein and RNA.

Redundancy and gene redundancy

Biologists use the term redundancy in more than one way. One usage invokes redundancy if a gene's activity can be changed or a system's part can be removed without affecting key system properties.⁽⁹⁾ Another usage refers to redundancy only if two parts of a system perform the same or similar tasks. The two notions are not synonymous. A system may be unaffected by removing a part, yet no two of its parts may have identical functions. Examples include proteins, whose tertiary structure is the result of cooperative folding involving many amino acids. No two amino acids may play identical roles in the folding process, yet protein structure can be highly robust to changes in individual amino acids. I will use redundancy here strictly in the second sense—redundancy of parts. This usage is not only consistent with one of the main dictionary definitions of the word, namely “unnecessary repetition”.⁽¹⁰⁾ It is also consistent with the usage of engineers concerned with designing reliable systems.⁽¹¹⁾ In engineering, redundancy is the main pillar of system reliability, partly because it is conceptually straightforward to build systems with redundant parts. But how important

Department of Biology, University of New Mexico, 167A Castetter Hall, Albuquerque, NM 87131-1091. E-mail: wagnera@unm.edu
Funding agency: NIH grant GM63882.
DOI 10.1002/bies.20170
Published online in Wiley InterScience (www.interscience.wiley.com).

is redundancy for robustness in biological systems? Less important than one might think, as I will argue here.

Genes encode the most important parts of biological systems above the gene level, proteins. The notion that redundancy of genes might be important for robustness emerged with the ability to generate gene deletion, synthetic null, or gene knockout strains of various organisms, that is, with the ability to eliminate specific genes from a genome. Many such gene deletion experiments yielded surprises. Specifically, they revealed that genes thought to be involved in key biological processes could often be removed without affecting the organism. This phenomenon exists in all organisms and all kinds of genes, be they enzymatic, structural or regulatory.^(12–22) The phenomenon is too pervasive to survey exhaustively, but I will discuss a few examples. The *CLN1* gene of budding yeast regulates the activity of the yeast cyclin-dependent kinase Cdc28p.⁽²³⁾ Proper regulation of Cdc28p is required for the transition from the G₁ phase to the S phase of the cell cycle. However, a null mutant in *CLN1* grows and divides normally on a minimal growth medium.⁽²⁴⁾ The *BarH2* gene from the fruit fly *Drosophila melanogaster* encodes a transcription factor with a homeobox DNA-binding domain. It is involved in sensory organ development, yet its deletion does not cause morphological defects in sensory organs.⁽²⁵⁾ The extracellular matrix protein tenascin of mice is thought to play an important role in morphogenesis during embryonic development. Yet the deletion of a gene encoding this protein leads to no detectable developmental abnormalities, to no change in the distribution of other extracellular matrix proteins, and to no detectable reduction of fertility.⁽²⁶⁾ The yeast *HMG2* gene encodes a key enzyme in the sterol biosynthesis pathway, 3-hydroxy-3-methylglutaryl-coenzyme A reductase. The end products of this pathway take part in processes ranging from electron transport to DNA replication. Yet the *HMG2* gene can be eliminated without affecting cell growth on either rich or minimal growth medium.⁽²⁷⁾ The *Drosophila* gene *knirps* encoding a key transcriptional regulator in embryonic development can be eliminated without affecting head development, during which *knirps* is expressed.⁽²²⁾ And finally, there are the three *TPK* genes in yeast, which encode catalytic subunits of the yeast cyclic AMP-dependent protein kinase, a key molecule in cell signaling. Any two of these three genes are dispensable for cell growth.⁽²⁸⁾

All of these examples are cases of gene redundancy resulting from gene duplication. The three *TPK* genes form a family of genes that originated from two duplications of a single ancestral gene.⁽²⁸⁾ Similarly, the *CLN1* gene has two duplicates, *CLN2* and *CLN3*, with similar functions, the *BarH2* gene has a duplicate, *BarH1*, the *HMG2* gene has a duplicate *HMG1*, and *knirps* has a duplicate, *knirps-related*.^(22,24,25,27) These and numerous other examples have drawn early attention to redundancy as a source of robustness.^(29–33) Eliminating *all* duplicates of a gene often reveals the severe phenotypic

effects that mutations in the gene lack, as expected if redundancy is the cause for weak gene knockout effects. Redundant duplicate genes with similar functions can thus help explain why knock-out mutations in important genes have no phenotypic effects.

Gene duplications are abundant by-products of recombination and DNA-repair processes in most genomes. Typically, between 25 and 50 percent of eukaryotic genes have at least one duplicate within the same genome.^(34,35) Immediately after a gene duplication, many pairs of duplicate genes produce the same two polypeptides with the same function. Subsequently, mutations that accumulate in either gene can lead to elimination of one duplicate, or to divergence in the duplicate's functions.⁽²⁸⁾ Although such divergence is often very limited, as the examples above show, most gene duplicates do not have completely identical functions. That is, they often encode multifunctional proteins that share some functions while they differ in others. For example, the *knirps* gene is dispensable for head development, but it is necessary for abdominal development.⁽²²⁾ Mig1p and Mig2p, two yeast transcriptional regulators, act redundantly in repressing the yeast gene *SUC2*, needed for metabolizing sucrose, but Mig1p is uniquely required for repressing the *GAL* genes necessary for galactose metabolism.⁽²¹⁾ Similarly, the three *TPK* genes, any two of which are normally dispensable for cell division, have different functions under conditions where yeast forms filamentous cell aggregates called pseudohyphae. *TPK2* is an essential gene under these conditions.⁽³⁶⁾ Such observations have led to the notion that duplicate genes can have partially redundant or overlapping functions, functions that reinforce each other in some conditions, but that are different in others.^(29,37) Such partial redundancy can also be an important reason why gene duplicates are maintained in a genome.^(38,39)

Distributed robustness

In distributed robustness, many parts of a system contribute to system function, but all of these parts have different roles. When one part fails or is changed through mutations, the system can compensate for this failure, but not because a “back-up” redundant part takes over the failed part's role. Distributed robustness is a fairly poorly understood cause of mutational robustness, because it requires a detailed, quantitative understanding of the inner workings of a genetic network. I will illustrate it with two examples, one from metabolic networks, the other one from a gene regulation network important for embryonic development.

Metabolic networks include hundreds of chemical reactions that are required to generate energy and biosynthetic building blocks for cell growth and proliferation. A variety of computational methods such as metabolic flux balance analysis can predict the maximally possible cell growth rate that a metabolic network can sustain in a given nutritional environment, given

stoichiometric information about all the chemical reactions that occur in the network.^(40–42) The predictions of these approaches are mostly in good agreement with experimental results. Where they are not, laboratory evolution can create organismal strains that behave as predicted. A case in point is the *E. coli* strain MG1655, which grows on glycerol as sole carbon source at a rate lower than predicted as optimal by flux balance analysis. However, within 40 days or 700 generations of evolution in the laboratory, the strain increases growth substantially, and the growth rate approaches that predicted by flux balance analysis as maximal.⁽⁴³⁾

Metabolic flux balance analysis and related approaches can be used to ask how robust a metabolic network's output (e.g. cell growth) is to the elimination of individual chemical reactions in the network. Such elimination corresponds to the elimination of enzyme-coding genes from a genome, and its effects are thus a good indicator of mutational robustness in metabolic genes. Several studies have systematically examined how elimination of individual reactions would affect a metabolic network.^(2,44–47) To give but one example, Edwards and Palsson assembled a reaction network comprising 436 metabolites and 736 reactions from the biochemical literature, genome sequence information and metabolic databases⁽⁴⁵⁾ of the bacterium *Escherichia coli*. They assessed the effects on cell growth of eliminating each of 48 central chemical reactions in glycolysis, the pentose phosphate pathway and the tricarboxylic acid cycle. Only seven of the 48 eliminated reactions are essential to cell growth. Of the remaining 41 nonessential reactions, 32 reduced cell growth by less than 5% and only 9 reduced growth by more than 5%. The experimental data and computational predictions on growth versus no-growth coincided in 86% of the cases. Where discrepancies occur, they may be due to the fact that wild-type *E. coli* cells have not been allowed to evolve under these inflicted genetic conditions of a missing enzyme-coding gene, and thus that the performance of a metabolic network is not optimal in this situation.⁽⁴⁶⁾

It is worth emphasizing that there is no redundancy in such metabolic networks, in the sense that different enzymes generally catalyze different chemical reactions. How then, can such networks be robust to elimination of individual reactions? The answer is that they shuttle the flow of matter through parts of a metabolic network that are not affected by blocking a reaction. I will illustrate this with a concrete example, that of the enzyme glucose 6-phosphate dehydrogenase—part of the pentose-phosphate shunt—encoded by the *zwf* gene of *E. coli*. Experimental work shows that loss-of-function mutants in the *E. coli* gene *zwf* encoding this enzyme grow at near wild-type levels,⁽⁴⁸⁾ and flux balance analysis predicts a similarly small growth rate reduction.⁽⁴⁴⁾ The elimination of this reaction, however, has profound systemic consequences on the metabolic network. For example, before the deletion, in the wild-type state, about two-thirds of a cell's NADPH is produced by

the pentose phosphate shunt. Most of the NADH is produced by the tricarboxylic acid cycle. One of the major systemic reorganizations of metabolic flow after the mutation regards NADPH production. To compensate for the blocked oxidative branch of the pentose phosphate shunt, most of the NADPH is now produced through an increased flux in the tricarboxylic acid cycle, which generates NADH. This NADH is then transformed into NADPH via a massively increased flux through the transhydrogenase reaction. (The nonoxidative branch of the pentose phosphate shunt can still be entered through other metabolites of glycolysis, and thus still serves its other important function, to produce various biosynthetic precursors.) In sum, the example illustrates a hallmark of distributed robustness: a systemic adjustment of network properties to compensate for a perturbed network part.

Gene regulation networks provide completely different examples of distributed robustness.^(3,49–52) I will just discuss one such example, that of the network sustaining the expression of the segment polarity genes *engrailed* and *wingless* in *Drosophila melanogaster*.⁽³⁾ Proper expression of the genes *wingless* and *engrailed* in 14 stripes is necessary for segmentation of the fly. The expression pattern of these genes is sustained by a network of multiple genes and their products, including *hedgehog*, *patched*, and *cubitus interruptus*. The properties of this network are determined by multiple biochemical parameters. These parameters include the binding constants of transcriptional regulators such as *engrailed* to their target gene's regulatory regions; the rate at which a ligand such as Hedgehog associates with its receptor; the stability of mRNAs and proteins expressed from network genes, and so on. Von Dassow and collaborators assembled a quantitative model of this network, a model that is grounded in empirical observations, and that correctly predicts many aspects of network function. They asked how sensitive the network is to changes in the biochemical parameters (48 in total) needed to describe its properties. The answer is that most of these parameters can be changed by more than one order of magnitude and the network still produces the *engrailed* and *wingless* gene expression patterns required for proper development. Moreover, even if one chooses random values (within a biologically sensible range) of all of the 48 biochemical parameters, there is a high probability that the network will function properly. Finally, the network is also highly robust to gene duplications, which can drastically change the expression levels of network proteins. A caveat to these conclusions is that they derive from a quantitative model based on limited empirical information, and that they thus await experimental confirmation. However, other studies hint at similar phenomena in different gene regulation networks.^(49–51)

Again, there is no redundancy involved in this representation of the segment polarity network. Every gene has a different role in the network, yet the network is highly robust to

changes of its regulatory interactions. The source of such robustness lies again in the distributed architecture of the network, and cannot be pinned down to any one gene. However, in many gene regulation networks, gene expression and activity of gene products are cooperatively regulated by one or more proteins,^(52–57) an observation that may go a long way towards explaining this phenomenon.⁽⁵⁸⁾

These examples illustrate that a quantitative representation of a whole metabolic or gene regulation network—not local information on only a small group of genes—is necessary to study distributed robustness. We will thus understand this phenomenon only to the extent that high-quality quantitative models of such systems based on experimental data will become available.

Candidate examples for distributed robustness also exist on different levels of biological organization. They include the developmental pathways responsible for the embryonic development of many phenotypic characters. The genes acting in such pathways harbor much variation that may affect gene products but does not affect the phenotype. Such cryptic variation is an indicator of mutational robustness. Are redundant gene functions responsible for it? Probably not. If they were, then mutations that reduce this robustness would preferentially affect genes with redundant duplicates. However, the best-explored mutations that reduce robustness are completely different in nature. In fruit flies, they affect the non-redundant gene encoding the heat-shock protein Hsp90, for which heterozygous loss-of-function mutations abolish robustness in a wide spectrum of developmental pathways.⁽⁵⁹⁾ The reason is that Hsp90 interacts with more than 80 different proteins important in development, and helps stabilize their conformation. As a second example, distributed robustness is a well-established phenomenon in artificial and natural neural networks. For instance, the computational performance of neural networks can be highly robust to removal of synaptic interactions or neurons, despite the fact that no two neurons may have the same role in a network.⁽⁶⁰⁾ On lower levels of biological organization, the paradigmatic examples of distributed robustness are biological macromolecules, such as proteins and RNA. Their folding in space involves the cooperation of hundreds of building blocks—nucleotides or amino acids—each of which forms unique (non-redundant) interactions with other building blocks. If a change in one amino acid leaves a protein's three-dimensional conformation unchanged, is this because the protein contained a redundant 'back-up' amino acid? Clearly not: The thermodynamic stability of any macromolecular structure is the sum of all contributions that such interactions make to the structure's free energy.

Large-scale gene perturbation studies

Examples, however many, of individual redundant genes or genetic networks with distributed robustness, are mere

anecdotal evidence for either cause of robustness. They do not answer the question how prevalent gene redundancy or distributed robustness is on a genome-wide scale. I will now turn to this question. To answer it, systematic, genome-wide gene perturbation studies that eliminate many genes in a genome and systematically assess the phenotypic effect of these deletions are very helpful. The first such large-scale study created some 300 different strains of budding yeast, each of which lacked one of the genes on chromosome V.⁽²⁴⁾ This study was followed by larger scale studies in *Saccharomyces cerevisiae*, the worm *Caenorhabditis elegans*, and the plant *Arabidopsis thaliana*,^(61–63) studies that mutated the DNA or reduced the expression levels of thousands of genes. A caveat to comparing the results of such studies is that their perturbation approaches differ. For instance, although deletion of specific genes through homologous recombination is highly effective in yeast, it is not feasible on a large scale in higher organisms such as *C. elegans*. In such organisms, temporary gene activation through interference with the expression of messenger RNA is currently the method of choice.⁽⁶²⁾

The first surprise of systematic gene knockout studies is the sheer number of genes with no detectable effect. For instance, the systematic knockout of most chromosome V genes in yeast shows that almost 40% of genes have no detectable effect on key fitness indicators such as cell division rate and sporulation efficiency in five different nutritional environments.⁽²⁴⁾ Later studies that eliminated more than 90 percent of all yeast genes yielded similar results.^(61,64) Such striking numbers raise the question whether genes with no apparent effect are truly dispensable under all circumstances. Although no laboratory experiment could answer this question definitively, the likely answer is no. First, some genes with no apparent effect in standard assays may affect fitness in subtle ways.⁽⁶⁵⁾ Second, it is easy to conceive of genes with specific functions, such as metabolic genes, where a mutation may have no effect in one environment but may be fatal in another environment. For metabolic genes, such environmental specificity may be the rule rather than the exception.⁽⁶⁶⁾ Even very rare environmental conditions, conditions that may occur only once every few thousand generations, can cause such genes to be maintained in a population.⁽⁶⁷⁾ Also, if genes whose elimination has no effects in one or a few environments were truly dispensable, they should evolve much more rapidly than indispensable genes, because they would be able to tolerate fewer mutations. However, the statistical association between a gene's dispensability and its rate of evolution is weak at best.^(68–72) Unfortunately, except for the limited number of metabolic genes (fewer than 20% of genes in a genome), it can be very difficult to ascertain whether a gene is not used in a particular environment or physiological condition, especially in higher eukaryotes, where 'gratuitous' expression of genes may be frequent.

Genome-scale gene-deletion studies and redundancy

Even though their neutrality may only be temporary, the sheer number of genes—thousands—with no deletion effect still needs a mechanistic explanation. Does gene redundancy provide this explanation? The results of genome-scale gene-deletion studies reveal the next surprise in this regard: Thousands of genes whose deletion has no detectable effect are single-copy genes. They have no duplicates in the genome.

The results of an early pertinent study are summarized in Fig. 1.⁽⁷³⁾ In this figure, I grouped gene-deletion mutations of yeast chromosome V into two categories, those with a non-detectable phenotypic defect, and those with a severe effect leading to a growth-rate reduction of more than 25% relative to the wild type. The figure's vertical axis shows the proportion of genes in each defect category that have the number of duplicates indicated on the horizontal axis. Significantly, more than 40% of duplicate genes with weak or no fitness effect are single-copy genes. A more recent study by Gu and collaborators that used more than 4,700 growth rate measurements of yeast gene-deletion strains reached a qualitatively similar result. 39.6% of yeast genes with a weak or no fitness defect have no duplicate.⁽⁷⁴⁾ Additionally, even if a gene has a duplicate, the duplicate may have diverged in function and may not be responsible for a weak gene-deletion effect. The situation in many-celled eukaryotes might be fundamentally different, because they contain more duplicates which occur in larger families.^(34,35,75) However, a recent study using the results of RNA interference experiments that temporarily silenced more than 16,000 worm genes suggests otherwise. In this study, more than 7,500 single-copy genes show no detectable phenotypic effect when temporarily silenced, versus 4515 of genes with one or more duplicate.⁽⁷⁶⁾

Does all this mean that gene redundancy is unimportant in determining the effect of gene deletions? No. For genes with duplicates, the likelihood that a gene has a lethal effect is lower than for single-copy genes.^(74,76) Moreover, the proportion of genes with weak or no effects is higher for duplicate genes than for single-copy genes.^(74,76) Third, the more similar two duplicates are, the less severe is the effect of deleting or silencing one of them.^(74,76) Fourth, the larger the number of duplicates a yeast gene has, the faster the gene evolves, implying that it can tolerate more mutations.⁽³⁴⁾ Finally, pairs of duplicate genes have more similar phenotypic effects upon deletion than randomly chosen pairs of unrelated genes.⁽⁷⁴⁾ Thus, gene duplications clearly play a role in determining gene-deletion effects.

Rapid functional divergence of duplicate genes

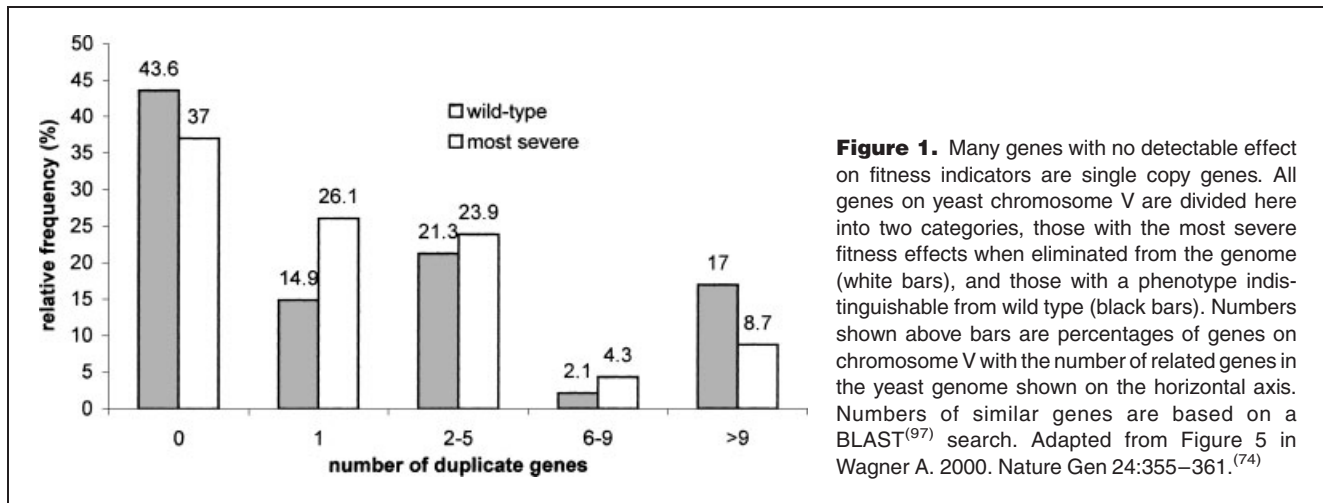
The above results suggest that both distributed robustness and gene redundancy play a role in mutational robustness.

However, they leave unanswered the key quantitative question: on a genome-wide scale (and not just for one functional class of genes) what proportion of weak gene perturbation effects are due to gene duplication? An estimate of this proportion based on large-scale gene-deletion experiments suggests that, in yeast, between 23% and 59% of weak gene perturbation effects are due to gene duplication.⁽⁷⁴⁾ A more recent computational analysis based on flux balance analysis in yeast indicates a greater role for gene duplicates.⁽⁶⁶⁾ However, this study is restricted to metabolic genes—only some 16% of all yeast genes—and it takes environmental specificity in a limited number of environments into account, which makes its results difficult to compare with these figures. In a higher organism, the worm *C. elegans*, the proportion of weak gene perturbation effects due to gene duplication may be lower than in yeast, between 3% and 36%.⁽⁷⁶⁾

Such results make it seem unlikely that gene redundancy is responsible for the vast majority of weak gene perturbation effects, although the numbers in the worm have to be taken with a grain of salt, because they are based on results of RNA interference assays. The true percentage of weak gene-deletion effects would be close to the upper bound of the above intervals (59% for yeast) if most duplicates retained similar functions long after their duplication. It would be close to the lower bound (23%) if most gene duplicates diverged rapidly in function. In the following section, I will argue that most duplicate genes diverge rapidly in important indicators of gene function. This suggests that the true proportion of weak gene deletion effects due to gene duplication is nowhere near the upper bound.

Estimating the rate at which gene duplicates diverge in function faces two technical problems. The first of them is to measure the time that elapsed since a gene duplication. The crudest indicator of this time is the number of nucleotide changes that lead to amino acid changes in either of the duplicates' gene product. However, because natural selection can strongly affect amino acid changes, their rate of accumulation can vary greatly over time. The problem can be alleviated by estimating instead the number of synonymous (silent) nucleotide substitutions that occurred in two genes since the duplication. Silent substitutions are substitutions that do *not* lead to amino acid changes in an encoded protein. They are subject to weaker selection than amino acid replacement substitutions, and thus do not pose the same problem as amino acid substitutions in estimating time since divergence.⁽⁷⁷⁾ In practice, synonymous (silent) divergence of two duplicate genes is measured as the fraction K_s of silent substitutions that occurred since the duplication at synonymous nucleotide positions.⁽⁷⁷⁾

The second, much more serious, problem is to identify whether two gene functions are identical, similar, or completely different. To identify one gene's function is already difficult, because gene function can mean many different things. For instance, gene function can be defined in biochemical terms,



such as through a gene product's DNA binding and catalytic activity, or via biological processes, such as through a gene's role in liver metabolism or brain development.⁽⁷⁸⁾ Another reason is that even presumably very well-characterized genes, such as those encoding central metabolic enzymes or aminoacyl-tRNA transferases, can perform utterly surprising additional functions in processes like cellular signaling or gene expression regulation.⁽⁷⁹⁾ To make matters worse, assessing whether gene duplicates *generally* retain similar functions can only be answered by characterizing hundreds or thousands of gene duplicates. The only current solution to this problem is to use crude indicators of gene function, indicators that are measurable for many genes. Functional genomics provides a wealth of such indicators, including where and when a gene is expressed, which proteins regulate its transcription, and which other molecules its protein products interact with. Each such indicator captures only one dimension of gene function, but it has the advantage that it can be measured for thousands of genes at once. I will now discuss evidence for the divergence of various gene function indicators after gene duplication.

One important indicator of similarity in gene function is similarity in gene expression. Some well-studied gene duplicates, for example, have preserved their biochemical function but differ in their spatial expression pattern, and can thus not be deleted without deleterious effects.⁽⁸⁰⁾ Genome-scale analyses of expression divergence after gene duplication show that even recently duplicated genes can show highly different mRNA gene expression patterns, so much so that there is only a moderate correlation between the silent divergence of two duplicate genes and the similarity in their expression patterns.^(81,82) Fig. 2 shows an example from the nematode worm *C. elegans*.⁽⁷⁶⁾ The figure is based on an analysis of more than 500 different microarray gene expression measurements and sequence information from the whole worm genome. Many gene duplicates with only

moderate silent divergence (and thus very low amino acid divergence) show very different mRNA expression patterns, almost as different as genes that have duplicated a long time ago.

A similar pattern holds for the transcriptional regulators that bind at regulatory regions of duplicate genes.^(83,84) After a gene duplication, the regulatory regions of two duplicate genes diverge. Much of the initial divergence occurs through the disappearance of binding sites for different transcription factors, because random mutations are more likely to eliminate such binding sites than to create new ones. Binding of transcriptional regulators to a large fraction of yeast genes has been studied using the technique of chromatin immunoprecipitation.⁽⁸⁵⁾ Fig. 3 shows how the similarity in the spectrum of transcriptional regulators that bind to the regulatory region of two duplicate genes depends on the silent (synonymous) divergence between the duplicates.⁽⁸³⁾ For yeast, a synonymous divergence of $K_s = 0.1$ among duplicated genes corresponds to a duplication age of approximately 10 million years.^(86,87) The solid horizontal line indicates the average fraction of shared transcriptional regulators between pairs of genes chosen at random from the yeast genome. The figure demonstrates that the regulatory regions of many closely related duplicate genes are bound by different transcriptional regulators. This is further indication of a rapid functional divergence among duplicate genes.

Thirdly, genes whose products interact with other proteins show a rapid divergence of common protein interactions.^(84,86,87) An example of such rapid divergence is shown for duplicate yeast genes in Fig. 4. Even among the most-recent gene duplicates, with a silent divergence of $K_s < 0.2$, only some 60% of gene products share a protein interaction partner. Among older duplicates, fewer than 20% have any common protein interaction partners. Qualitatively, these findings are independent of which among various genome-

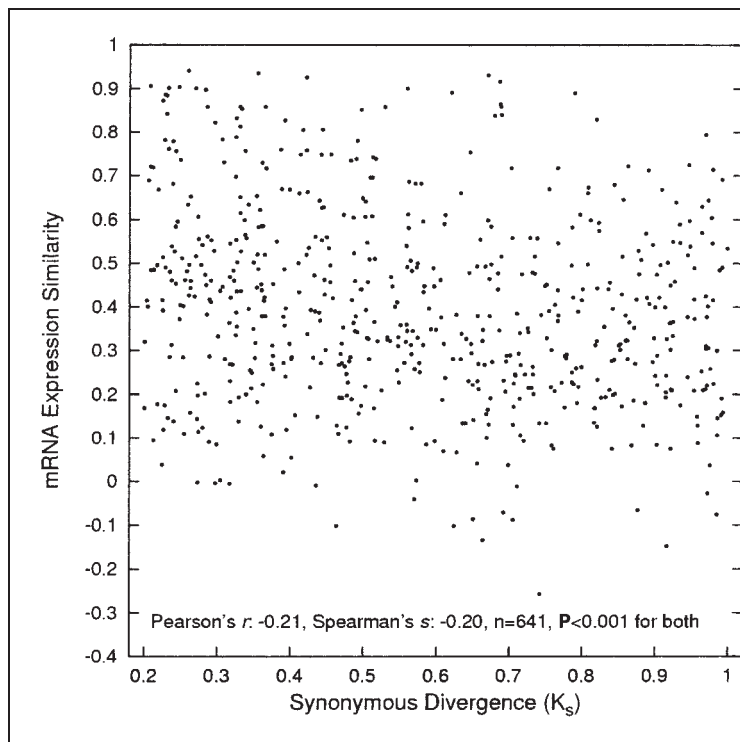


Figure 2. Even closely related duplicate genes can have very different mRNA expression patterns. mRNA gene expression similarity (vertical axis) of duplicate genes in the worm *C. elegans*. The horizontal axis shows synonymous divergence K_s of each duplicate gene pair. K_s is the fraction of synonymous substitutions that occurred since the duplication at synonymous nucleotide positions that do not lead to amino acid replacements. mRNA gene expression similarity is measured by the Pearson product–moment correlation coefficient of logarithmically transformed gene expression ratios from 522 microarray experiments.⁽⁹⁸⁾ Gene duplicates with $K_s < 0.2$ are excluded, because microarray experiments cannot distinguish differences in expression among such closely related duplicates. Most duplicate gene pairs in the silent divergence interval shown would have an amino acid divergence of less than 20%. Adapted from Conant GC, Wagner A. 2003. Proc R Soc Lond Series B 271:89–96.

scale datasets on protein interactions^(34,88–90) one uses in this analysis.

Finally, the only moderately conserved indicator of functional similarity among duplicate genes is the localization of their protein products. Protein localization data based on fluorescence microscopy of proteins tagged with green fluorescent protein have recently become available for over 4,000 yeast genes.⁽⁹⁹⁾ The available data classify a protein's localization according to one or more cellular compartments (nucleus, endoplasmic reticulum etc.) in which it occurs. Fig. 5 shows the proportion of duplicate yeast genes with a given synonymous divergence that occur in identical cellular compartments (Fig. 5a) or completely different cellular compartments (Fig. 5b). Clearly, the fraction of genes with identical localization is somewhat conserved: a substantial fraction of duplicate gene products at all levels of divergence localize to the same cellular compartments.

Each genome-scale dataset on indicators of gene function has its own source of error, but taken together, the above evidence suggests that duplicate genes do not, as a rule, retain indications of similar function long after a duplication. It provides further evidence that gene redundancy may not be the dominant factor determining robustness to mutations. This does, however, not mean that genes never retain similar functions long after duplication. For instance, some of the gene families listed above, families whose members retained similar functions, result from very ancient gene duplications. Examples include the yeast CLN family of cell-cycle regulators, and

the TPK gene family of protein kinase catalytic subunits, in which individual members are dispensable for growth under some conditions. The most-closely-related genes in these families show a silent divergence of $K_s > 1$, indicating that they may be more than 100 million years old.⁽⁸⁷⁾ The observation that some gene duplicates retain similar functions long after others have completely diverged represents an intriguing evolutionary puzzle waiting to be solved.

Quantifying redundancy and distributed robustness?

Can we quantify the extent of redundancy in a genetic system whose inner workings are perfectly well understood? On a crude level, this seems quite feasible. For example, metabolic reactions in living cells are driven by enzymes, some of which may be the products of duplicate genes. One could simply determine the fraction of enzymes in a metabolic pathway that are encoded by more than one gene, and use this fraction as an indicator of redundancy in the network. However, can one go beyond that? Specifically, for any specific indicator of a system's performance, can we say to what extent its robustness is caused by redundancy or by distributed robustness? The answer to this question must depend on how one defines robustness. For instance, a flux-balance representation of metabolic networks, and a differential equation representation of a regulatory gene network are fundamentally different, and will therefore also employ different measures of robustness. But, more disturbingly, redundancy and distributed robustness

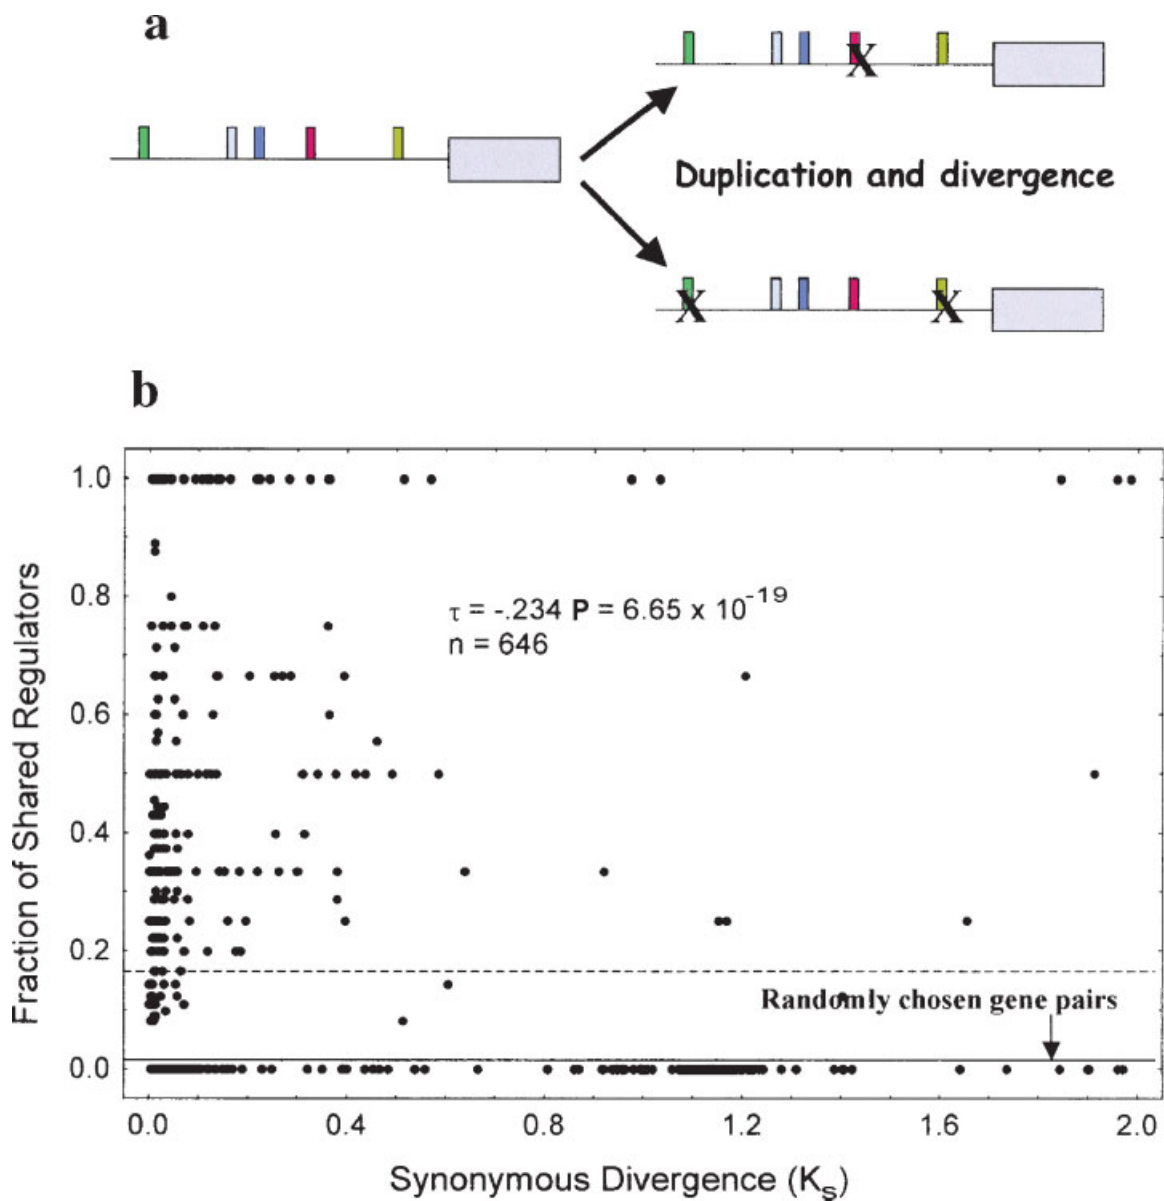


Figure 3. Duplicate genes and divergence in the binding of transcriptional regulators to their regulatory regions. **a:** After a gene duplication, the regulatory regions of two duplicate genes diverge. Much of this divergence occurs through the disappearance of binding by individual transcription factors (indicated by the small rectangular boxes), because random mutations are more likely to eliminate transcription factor binding sites than to create new ones. **b:** The vertical axis shows the fraction f of transcriptional regulators that bind to the regulatory region of *both* genes of a pair of yeast gene duplicates whose synonymous divergence is shown on the horizontal axis. If d_1 transcriptional regulators bind to the regulatory region of gene 1, d_2 transcriptional regulators bind to the regulatory region of gene 2, and if d_{12} transcriptional regulators bind to the regulatory regions of both genes, then this fraction f calculates as $f = d_{12}/(d_1 + d_2 - d_{12})$. A value of $f = 1$ indicates that the same transcriptional regulators bind the two genes' regulatory regions. A value of $f = 0$ indicates that completely different regulators bind the two genes' regulatory regions. The horizontal lines indicate the average (solid) plus one standard deviation (dashed) fraction f of shared transcriptional regulators between pairs of genes chosen at random from the yeast genome. The synonymous (silent) divergence of two duplicate genes is measured by the fraction K_S of synonymous substitutions that occurred since the duplication at synonymous nucleotide positions.^(86,87) From Evangelisti, A, Wagner, A. 2004 J Exp Zool Molec Dev Evol 302B: 392–411, which uses transcription factor binding data based on chromatin immunoprecipitation,⁽⁸⁵⁾ reproduced with permission.

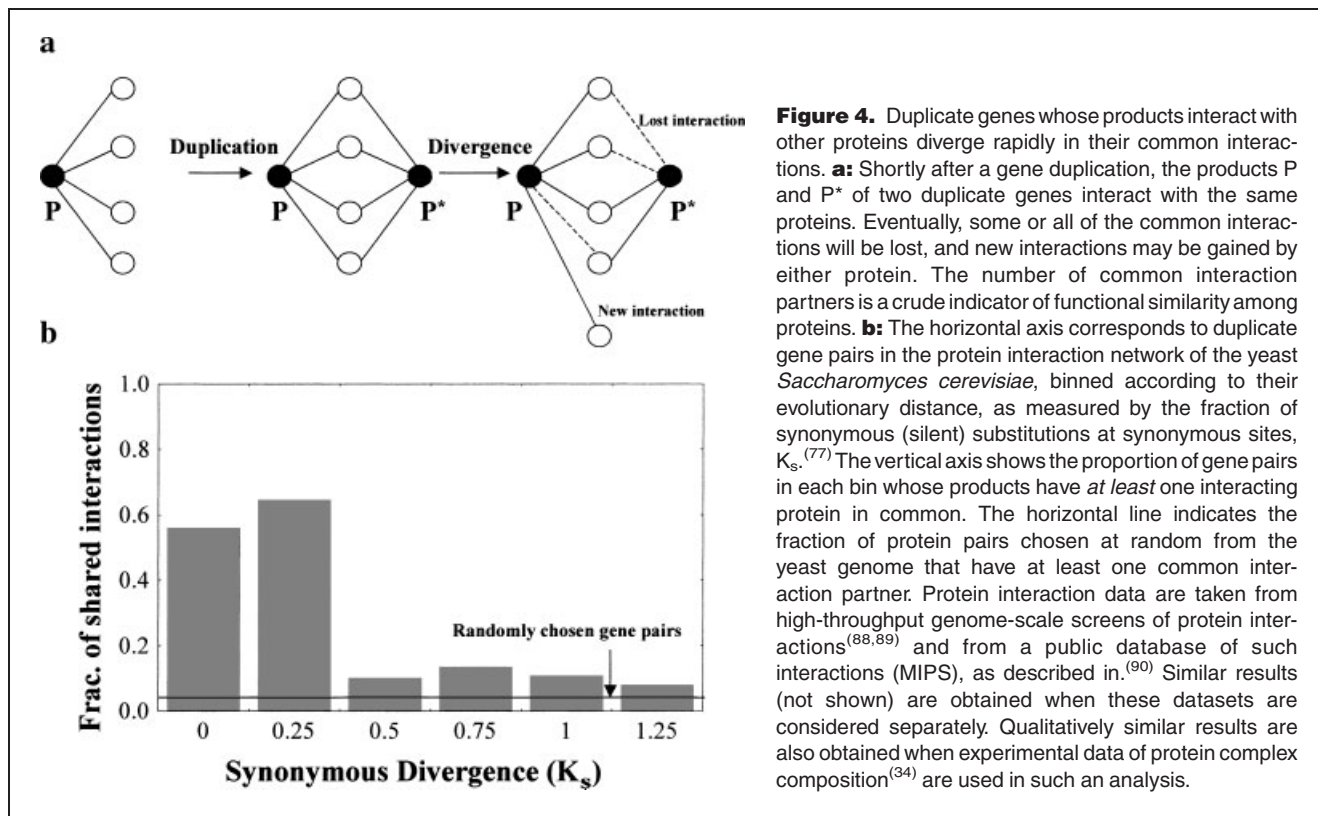


Figure 4. Duplicate genes whose products interact with other proteins diverge rapidly in their common interactions. **a:** Shortly after a gene duplication, the products P and P* of two duplicate genes interact with the same proteins. Eventually, some or all of the common interactions will be lost, and new interactions may be gained by either protein. The number of common interaction partners is a crude indicator of functional similarity among proteins. **b:** The horizontal axis corresponds to duplicate gene pairs in the protein interaction network of the yeast *Saccharomyces cerevisiae*, binned according to their evolutionary distance, as measured by the fraction of synonymous (silent) substitutions at synonymous sites, K_s .⁽⁷⁷⁾ The vertical axis shows the proportion of gene pairs in each bin whose products have *at least* one interacting protein in common. The horizontal line indicates the fraction of protein pairs chosen at random from the yeast genome that have at least one common interaction partner. Protein interaction data are taken from high-throughput genome-scale screens of protein interactions^(88,89) and from a public database of such interactions (MIPS), as described in.⁽⁹⁰⁾ Similar results (not shown) are obtained when these datasets are considered separately. Qualitatively similar results are also obtained when experimental data of protein complex composition⁽³⁴⁾ are used in such an analysis.

can have indistinguishable effects on a system. I will use an example from a simple metabolic system to make this point.

Metabolic control analysis is concerned with the behavior of metabolic pathways—chains of chemical reactions catalyzed by enzymes—when enzyme and metabolite concentrations change. An important quantity in this regard is an enzyme’s flux control coefficient.^(91,p.4) The metabolic flux through an enzymatic reaction is the rate at which an enzyme converts substrate into product per unit time. Flux is a key variable that links the activity of a metabolic enzyme to the phenotype. The reason is that the flux through some chemical reactions is strongly associated with fitness components like cellular growth or biomass production. An enzyme’s flux control coefficient indicates the sensitivity of flux through an enzymatic reaction to changes in the activity of some enzyme in the pathway. Consider the linear metabolic pathway of Fig. 6a. In steady-state, the metabolic flux through each reaction of the pathway must be identical to some value F . The flux control coefficient C_j of each enzyme E_j in the pathway is usually described as a differential quotient (Ref. 91, page 4) which measures the response of flux F to a very small change in enzyme activity E_j .

$$C_j = \frac{\partial F}{\partial E_j} \frac{E_j}{F} \quad (1)$$

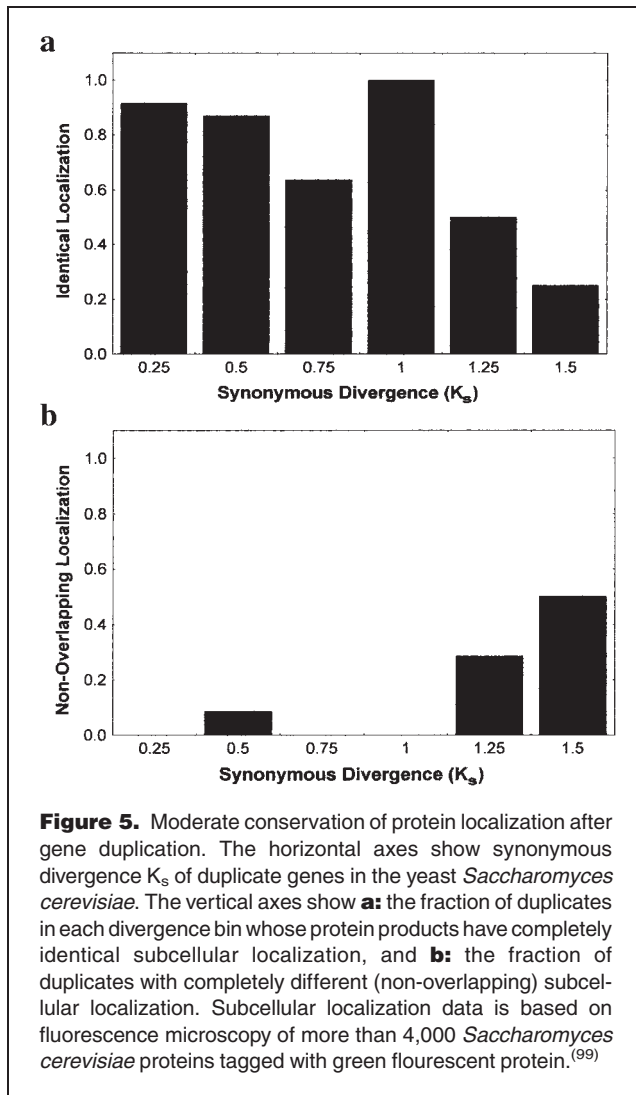
If all enzymes in the pathway are unsaturated and show Michaelis–Menten kinetics, the flux coefficient takes on a

particularly simple form.⁽⁹²⁾

$$C_j = \frac{1/E_j}{1/E_1 + 1/E_2 + 1/E_3 + \dots + 1/E_n} \quad 1 \leq j \leq n \quad (2)$$

For my purpose, it will be sufficient to consider an even simpler case, where every enzyme E_j ($1 \leq j \leq n$) has the same activity e . In this case, the flux coefficient of each enzyme simply becomes $C_j = (1/n)$ ($1 \leq j \leq n$). A pathway like this is robust, in the sense that a small (mutationally caused) reduction in the activity of any one enzyme activity changes the steady-state pathway flux by an amount that is a factor n smaller. Notice that redundancy, as defined above, is not involved here, because each enzyme may catalyze a different chemical reaction. The observation that each and every enzyme E_j has the same effect on flux in this scenario raises the question of whether such even distribution of control is a reasonable way to make the concept of distributed robustness exact, and to distinguish it from redundancy.

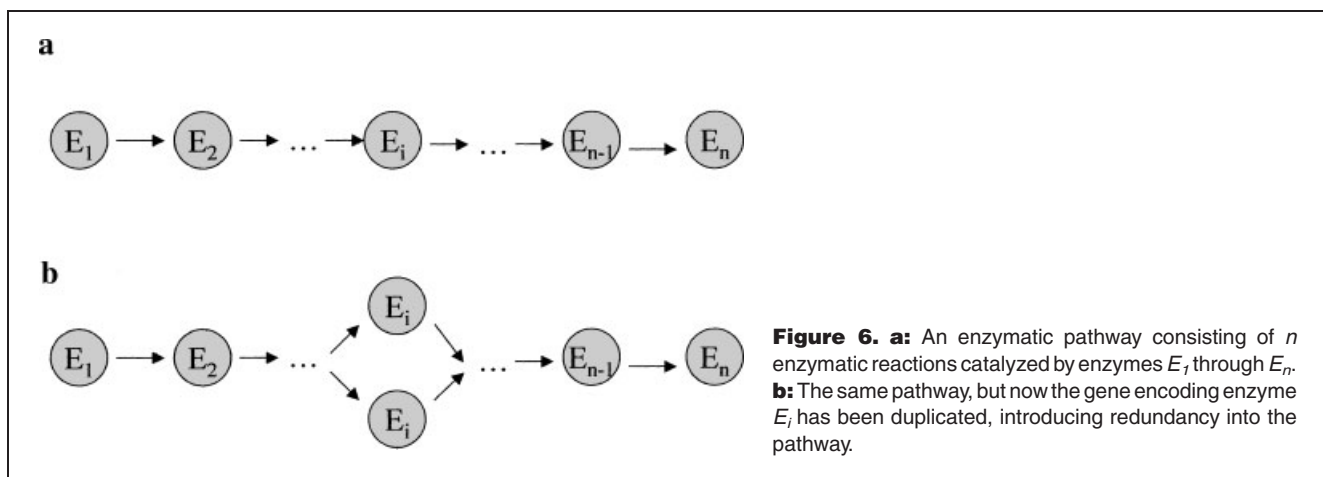
To see whether this is the case, imagine that we introduce redundancy into this pathway, as shown in Fig. 6b. Specifically, we duplicate the gene encoding one of the enzymes, say enzyme E_i . Two copies of the gene produce twice the amount of enzyme and, as a result, this enzyme’s activity doubles to $2e$. (In engineering, this kind of redundancy would be called parallel and active redundancy, because both copies of the system part are active—expressed—and their products act in



parallel; (ref. 11, Chapter 6) How does this duplication affect robustness of the pathway? For enzyme E_j , the control coefficient (2) now calculates as $1/(2n-1)$. If n is large, then the control coefficient of this enzyme has been effectively cut in half. In other words, the pathway is now nearly twice as robust to changes in activities of this enzyme. What about the remaining enzymes, whose genes have not been duplicated? Their control coefficient (2) increases to $2/(2n-1)$, a very small increase if the number of enzymes n is sufficiently large. If we were to introduce not one but $k > 1$ additional copies of the gene for E_i , then the pathway would become increasingly more robust to changes in the activity of E_j . In sum, the net effect of introducing gene redundancy into this metabolic pathway is to render the previously uniform distribution of control coefficients (2) non-uniform. From this perspective, one might propose distinguishing redundancy from distributed robustness through an uneven distribution of control in the pathway. Redundancy causes uneven distribution of flux control across parts of the pathway; the lower an enzyme's control coefficients, the higher its redundancy.

The fundamental flaw in this argument is that uneven distribution of control can be achieved by means that have nothing to do with redundancy. For example, even if each enzyme is encoded by only one gene, and even if each enzyme-coding gene expresses the same numbers of enzyme molecules, each enzyme molecule could have different catalytic activity, which would cause an uneven distribution of flux control. Similarly, if we duplicated all genes in the model pathway of Fig. 6a, every enzyme in the pathway would have the same flux control coefficient—flux control would be perfectly evenly distributed—but now all enzymes are encoded by two redundant genes. In sum, a pathway could have a uniform distribution of control coefficients with much gene redundancy, or a non-uniform distribution without any gene redundancy.

All this implies that our ability to determine how much redundancy a system contains depends on our perspective. If



we view a metabolic pathway from the point of view of the genes that encode its enzymes, it may be simple to determine how much redundancy it contains—just count the redundant genes. The same holds if we view the pathway from the perspective of the individual protein molecules that do the chemical work: it is in principle easy to estimate how many of the molecules catalyzing the same reaction a cell contains, and all of these can be viewed as redundant. However, as I just showed, assessing the extent of redundancy becomes impossible on the level of enzyme activities, the aggregate variables best suited to understand a pathway's robustness to small changes in enzyme activity.

This problem is not a peculiarity of metabolic pathways. For instance, quantities analogous to metabolic control coefficients exist for bacterial and eukaryotic gene regulation systems.^(93,94) More generally, one can characterize the robustness of any property of a regulatory network—the expression level of its genes, their products' activities or phosphorylation state, or their affinity to other molecules—by such control coefficients. And for any network design devoid of redundant genes, there may be a design including redundant genes that shows exactly the same distribution of control coefficients. Even more generally, pertinent and analogous insights exist in control theory, a discipline aiming at understanding and controlling the behavior of any system, whether natural or engineered. An important concept in control theory is that of the transfer function, which specifies how a system's output behaves as a function of its input. General theorems of control theory imply that for any system with redundant parts, one can design a system without redundancy but the same transfer function, i.e. the same behavior (ref. 95, Chapter 4).

All this implies that the qualitative approach to estimate the extent of redundancy in a biological system—enumerating redundant parts—may be the only one. Qualitative, graph-based representations of metabolic and gene regulation networks lend themselves ideally to such enumeration. Beyond such enumeration, however, it may be impossible to distinguish distributed robustness and redundancy, because the two have indistinguishable signatures on the variables necessary to understand a system's function and robustness.

In summary, redundancy as a source of robustness is especially prominent for systems whose parts are genes and gene products. Its importance results from the abundance of gene duplications in most genomes. In such systems, gene redundancy is partly responsible for increased robustness to mutations and an increased incidence of neutral mutations. However, redundancy may not be the most prominent cause of robustness on this level of biological organization. This is because many genes whose elimination has no effect on the organism are genes without duplicates. Also, typically fewer than half of a genome's genes are duplicates, and many of these duplicates diverge rapidly in function after duplication. Is distributed robustness more important in microbes than in

higher organisms, which have more gene duplicates but also more stable internal environments? Is distributed robustness more important in some biological processes like signal transduction than in others, such as metabolism? Ultimately, only empirically grounded quantitative models of biological systems will be able to answer these and many other remaining questions about mutational robustness in living systems.

Acknowledgments

This article is to appear in similar form in ref. 96, whose copyright is held by Princeton University Press. Permission to reprint it is hereby acknowledged.

References

1. Rennell D, Bouvier S, Hardy L, Poteete A. 1991. Systematic mutation of bacteriophage T4 lysozyme. *J Molec Biol* 222:67–87.
2. Edwards JS, Palsson BO. 2000. Robustness analysis of the *Escherichia coli* metabolic network. *Biotech Progress* 16:927–939.
3. von Dassow G, Meir E, Munro E, Odell G. 2000. The segment polarity network is a robust development module. *Nature* 406:188–192.
4. Raff RA. 1996. *The Shape Of Life. Genes, Development, And The Evolution Of Animal Form*. Chicago, IL: The University of Chicago Press.
5. Gibson G, Dworkin I. 2004. Uncovering cryptic genetic variation. *Nature Reviews Genetics* (in press).
6. Fisher RA. 1931. The evolution of dominance. *Biological reviews* 6:345–368.
7. Waddington CH. 1953. The genetic assimilation of an acquired character. *Evolution* 7:118–126.
8. Waddington CH. 1959. *The Strategy of the Genes*. New York: Macmillan.
9. Krakauer D, Plotkin J. 2002. Redundancy, antiredundancy, and the robustness of genomes. *Proc Natl Acad Sci USA* 99:1405–1409.
10. *The American Heritage College Dictionary*. 1993. Boston, MA: Houghton Mifflin.
11. Fuqua NB. 1987. *Reliability Engineering for Electronic Design*. New York, NY: Marcel Dekker.
12. Cadigan KM, Grossniklaus U, Gehring WJ. 1994. Functional redundancy: the respective roles of the 2 sloppy paired genes in *Drosophila* segmentation. *Proc Natl Acad Sci USA* 91:6324–6328.
13. Goldstein LSB. 1993. Functional redundancy in mitotic force generation. *J Cell Biol* 120:1–3.
14. Hoffmann FM. 1991. *Drosophila* *abl* and genetic redundancy in signal transduction. *Trends Gen* 7:351–356.
15. Goodson HV, Spudich JA. 1995. Identification and molecular characterization of a yeast myosin-I. *Cell Mot Cyto* 30:73–84.
16. Joyner AL, Herrup K, Auerbach BA, Davis CA, Rossant J. 1991. Subtle cerebellar phenotype in mice homozygous for a targeted deletion of the *En-2* homeobox. *Science* 251:1239–1243.
17. Wang YK, Jaenisch R. 1997. Myogenin can substitute for *Myf5* in promoting myogenesis but less efficiently. *Development* 124:2507–2513.
18. Wang YK, Schnegelsberg PNJ, Dausman J, Jaenisch R. 1996. Functional redundancy of the muscle-specific transcription factors *Myf5* and *Myogenin*. *Nature* 379:823–825.
19. Lohnes D, Kastner P, Dierich A, Mark M, Lemeur M, et al. 1993. Function of retinoic acid receptor gamma in the Mouse. *Cell* 73:643–658.
20. Lufkin T, Lohnes D, Mark M, Dierich A, Gorry P, et al. 1993. High postnatal lethality and testis degeneration in retinoic acid receptor-alpha mutant mice. *Proc Natl Acad Sci USA* 90:7225–7229.
21. Lutfiyya LL, Iyer VR, Derisi J, Devit MJ, Brown PO, et al. 1998. Characterization of 3 related glucose repressors and genes they regulate in *Saccharomyces cerevisiae*. *Genetics* 150:1377–1391.
22. Gonzalez-Gaitan M, Rothe M, Wimmer EA, Taubert H, Jackle H. 1994. Redundant functions of the genes *knirps* and *knirps-related* for the establishment of anterior *Drosophila* head structures. *Proc Natl Acad Sci USA* 91:8567–8571.

23. Nasmyth K. 1993. Control of the yeast cell cycle by the Cdc28 protein kinase. *Current Opin Cell Biol* 5:166–179.
24. Smith V, Chou KN, Lashkari D, Botstein D, Brown PO. 1996. Functional analysis of the genes of yeast chromosome V by genetic footprinting. *Science* 274:2069–2074.
25. Higashijima S, Michiue T, Emori Y, Saigo K. 1992. Subtype determination of *Drosophila* embryonic external sensory organs by redundant homeobox genes *BarH1* and *BarH2*. *Genes Dev* 6:1005–1018.
26. Saga Y, Yagi T, Ikawa Y, Sakakura T, Aizawa S. 1992. Mice develop normally without tenascin. *Genes Dev* 6:1821–1831.
27. Basson ME, Thorsness M, Rine J. 1986. *Saccharomyces cerevisiae* contains 2 functional genes encoding 3-hydroxy-3-methylglutaryl-coenzyme-A reductase. *Proc Natl Acad Sci USA* 83:5563–5567.
28. Toda T, Cameron S, Sass P, Zoller M, Wigler M. 1987. Three different genes in *S. cerevisiae* encode the catalytic subunits of the cAMP-dependent protein kinase. *Cell* 50:277–287.
29. Wilkins AS. 1986. *Genetic Analysis of Animal Development*. New York, NY: Wiley.
30. Thomas JH. 1993. Thinking about genetic redundancy. *Trends in Gen* 9:395–399.
31. Tautz D. 1992. Redundancies, development and the flow of information. *Bioessays* 14:263–266.
32. Wilkins AS. 1997. Canalization: A molecular genetic perspective. *Bioessays* 19:257–262.
33. Cooke J, Nowak MA, Boerlijst M, MaynardSmith J. 1997. Evolutionary origins and maintenance of redundant gene expression during metazoan development. *Trends Gen* 13:360–364.
34. Conant GC, Wagner A. 2002. Genome History: a software tool and its applications to fully sequenced genomes. *Nucleic Acids Res* 30:1–10.
35. Rubin GM, Yandell MD, Wortman JR, Miklos GLG, Nelson CR, et al. 2000. Comparative genomics of the eukaryotes. *Science* 287:2204–2215.
36. Robertson LS, Fink GR. 1998. The three yeast A kinases have specific signaling functions in pseudohyphal growth. *Proc Natl Acad Sci USA* 95:13783–13787.
37. Edelman G, Gally J. 2002. Degeneracy and complexity in biological systems. *Proc Natl Acad Sci USA* 98:13763–13768.
38. Wagner A. 1999. Redundant gene functions and natural selection. *J Evol Biol* 12:1–16.
39. Force A, Lynch M, Pickett FB, Amores A, Yan YL, et al. 1999. Preservation of duplicate genes by complementary degenerative mutations. *Genetics* 151:1531–1545.
40. Edwards JS, Ramakrishna R, Schilling C, Palsson BO. 1999. Metabolic flux balance analysis. In: Lee S, Papoutsakis E, editors. *Metabolic Engineering*. New York: Marcel Dekker. p 13–57.
41. Schuster S, Hilgetag C, Woods J, Fell D. 1996. Elementary modes of functioning in biochemical networks. In: Cuthbertson R, Holcombe M, Paton R, editors. *Computation in Cellular and Molecular Biological Systems*. Singapore: World Scientific. p 151–164.
42. Schuster S, Dandekar T, Fell DA. 1999. Detection of elementary flux modes in biochemical networks: a promising tool for pathway analysis and metabolic engineering. *Trends Biotech* 17:53–60.
43. 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.
44. Edwards JS, Palsson BO. 1999. Systems properties of the *Haemophilus influenzae* Rd metabolic genotype. *J Biol Chem* 274:17410–17416.
45. Edwards JS, Palsson BO. 2000. The *Escherichia coli* MG1655 in silico metabolic genotype: Its definition, characteristics, and capabilities. *Proc Natl Acad Sci USA* 97:5528–5533.
46. Segre D, Vitkup D, Church G. 2002. Analysis of optimality in natural and perturbed metabolic networks. *Proc Natl Acad Sci USA* 99:15112–15117.
47. Stelling J, Klamt S, Bettenbrock K, Schuster S, Gilles ED. 2002. Metabolic network structure determines key aspects of functionality and regulation. *Nature* 420:190–193.
48. Fraenkel DG. 1996. Glycolysis. In: Neidhardt FC, editor. *Escherichia coli and Salmonella: Cellular and Molecular Biology*. Volume 1. Washington, D.C.: American Society of Microbiology Press. p 189–198.
49. Raya A, Kawakami Y, Rodriguez-Esteban C, Ibanes M, Rasskin-Gutman D, et al. 2004. Notch activity acts as a sensor for extracellular calcium during vertebrate left-right determination. *Nature* 427:121–128.
50. Alon U, Surette M, Barkai N, Leibler S. 1999. Robustness in bacterial chemotaxis. *Nature* 397:168–171.
51. Barkai N, Leibler S. 2000. Biological rhythms—Circadian clocks limited by noise. *Nature* 403:267–268.
52. Jiang G, Sladek F. 1997. The DNA binding domain of hepatocyte nuclear factor 4 mediates cooperative, specific binding to DNA and heterodimerization with the retinoid X receptor alpha. *J Biol Chem* 272:1218–1225.
53. Pinckard J, Sheehan K, Schreiber R. 1996. Ligand-induced formation of p55 and p75 TNF receptor heterocomplexes: A mechanism underlying TNF receptor cooperativity. *FASEB J* 10:1794.
54. Brandt M, Vickery L. 1997. Cooperativity and dimerization of recombinant human estrogen receptor hormone-binding domain. *J Biol Chem* 272:4843–4849.
55. Stoecklin E, Wissler M, Moriggl R, Groner B. 1997. Specific DNA binding of Stat5, but not of glucocorticoid receptor, is required for their functional cooperation in the regulation of gene transcription. *Molec Cell Biol* 17:6708–6716.
56. Klinge C. 1999. Estrogen receptor binding to estrogen response elements slows ligand dissociation and synergistically activates reporter gene expression. *Molec Cell Endocrin* 150:99–111.
57. Wanant S, Quon M. 2000. Insulin receptor binding kinetics: Modeling and simulation studies. *J Theoret Biol* 205:355–364.
58. Wagner A. 1996. Does Evolutionary Plasticity Evolve? *Evolution* 50:1008–1023.
59. Rutherford S, Lindquist S. 1998. Hsp90 as a capacitor for morphological evolution. *Nature* 396:336–342.
60. Amit DJ. 1989. *Modeling Brain Function. The World of Attractor Neural Networks*. Cambridge, UK: Cambridge University Press.
61. Giaever G, Chu AM, Ni L, Connelly C, Riles L, et al. 2002. Functional profiling of the *Saccharomyces cerevisiae* genome. *Nature* 418:387–391.
62. Kamath R, Fraser A, Dong Y, Poulin G, Durbin R, et al. 2003. Systematic functional analysis of the *Caenorhabditis elegans* genome using RNAi. *Nature* 421:231–237.
63. Alonso J, Stepanova A, Leisse T, Kim C, Chen H, et al. 2003. Genome-wide insertional mutagenesis of *Arabidopsis thaliana*. *Science* 301:653–657.
64. Winzeler EA, Shoemaker DD, Astromoff A, Liang H, Anderson K, et al. 1999. Functional characterization of the *S. cerevisiae* genome by gene deletion and parallel analysis. *Science* 285:901–906.
65. Thatcher JW, Shaw JM, Dickinson WJ. 1998. Marginal fitness contributions of nonessential genes in yeast. *Proc Natl Acad Sci USA* 95:253–257.
66. Papp B, Pal C, Hurst LD. 2004. Metabolic network analysis of the causes and evolution of enzyme dispensability in yeast. *Nature* 429:661–664.
67. Wagner A. 2003. Risk management in biological evolution. *J Theoret Biol* 225:45–57.
68. Jordan I, Rogozin I, Wolf Y, Koonin E. 2002. Essential genes are more evolutionarily conserved than are nonessential genes in bacteria. *Genome Res* 12:962–968.
69. Hirsh A, Fraser H. 2003. Rate of evolution and gene dispensability—Reply. *Nature* 421:497–498.
70. Hirsh A, Fraser H. 2001. Protein dispensability and rate of evolution. *Nature* 411:1046–1049.
71. Pal C, Papp B, Hurst L. 2003. Rate of evolution and gene dispensability. *Nature* 421:496–497.
72. Yang J, Gu Z, Li W. 2003. Rate of protein evolution versus fitness effect of gene deletion. *Molec Biol Evol* 20:772–774.
73. Wagner A. 2000. Mutational robustness in genetic networks of yeast. *Nature Gen* 24:355–361.
74. Gu Z, Steinmetz L, Gu X, Scharfe C, Davis R, et al. 2003. Role of duplicate genes in genetic robustness against null mutations. *Nature* 421:63–66.
75. Qian J, Luscombe NM, Gerstein M. 2001. Protein family and fold occurrence in genomes: Power-law behaviour and evolutionary model. *J Molec Biol* 313:673–681.

76. Conant GC, Wagner A. 2003. Duplicate genes and robustness to transient gene knockouts in *Caenorhabditis elegans*. *Proc R Soc Lond Series B* 271:89–96.
77. Li W-H. 1997. *Molecular Evolution*. Massachusetts: Sinauer.
78. Ashburner M, Ball C, Blake J, Botstein D, Butler H, et al. 2000. Gene Ontology: tool for the unification of biology. *Nature Genetics* 25:25–29.
79. Jeffery C. 1999. Moonlighting proteins. *Trends in Biochem Sci* 24:8–11.
80. Li XL, Noll M. 1994. Evolution of distinct developmental functions of 3 *Drosophila* genes by acquisition of different cis-regulatory regions. *Nature* 367:83–87.
81. Gu Z, Nicolae D, Lu H, Li W. 2002. Rapid divergence in expression between duplicate genes inferred from microarray data. *Trends in Genetics* 18:609–613.
82. Wagner A. 2000. Decoupled evolution of coding region and mRNA expression patterns after gene duplication: Implications for the neutralist-selectionist debate. *Proc Nat Acad Sci USA* 97:6579–6584.
83. Evangelisti A, Wagner A. 2004. Molecular evolution in the transcriptional regulation network of yeast. *J Exp Zool Molec Dev Evol* 302B:392–411.
84. Maslov S, Sneppen K, Eriksen K, Yan K-K. 2004. Upstream plasticity and downstream robustness in evolution of molecular networks. *BMC Evolutionary Biology* 4:9.
85. Lee T, Rinaldi N, Robert F, Odom D, Bar-Joseph Z, et al. 2002. Transcriptional regulatory networks in *Saccharomyces cerevisiae*. *Science* 298:799–804.
86. Wagner A. 2003. How large protein interaction networks evolve. *Proc R Soc London Series B* 270:457–466.
87. Wagner A. 2001. The yeast protein interaction network evolves rapidly and contains few redundant duplicate genes. *Molec Biol Evol* 18:1283–1292.
88. Ito T, Chiba T, Ozawa R, Yoshida M, Hattori M, et al. 2001. A comprehensive two-hybrid analysis to explore the yeast protein interactome. *Proc Nat Acad Sci USA* 98:4569–4574.
89. Uetz P, Giot L, Cagney G, Mansfield TA, Judson RS, et al. 2000. A comprehensive analysis of protein–protein interactions in *Saccharomyces cerevisiae*. *Nature* 403:623–627.
90. Mewes HW, Heumann K, Kaps A, Mayer K, Pfeiffer F, et al. 1999. MIPS: a database for genomes and protein sequences. *Nucleic Acids Res* 27:44–48.
91. Fell D. 1997. *Understanding the Control of Metabolism*. Miami: Portland Press.
92. Kacser H, Burns JA. 1981. The molecular basis of dominance. *Genetics* 97:639–666.
93. Savageau MA. 1971. Parameter sensitivity as a criterion for evaluating and comparing the performance of biochemical systems. *Nature* 229:542–544.
94. Savageau MA. 1976. *Biochemical Systems Analysis: A Study Of Function And Design In Molecular Biology*. Reading, MA: Addison-Wesley.
95. Leigh JR. 1992. *Control Theory: A Guided Tour*. London: IEE Publishing.
96. Wagner A. 2005. *Robustness and Evolvability in Living Systems*. Princeton, NJ: Princeton University Press.
97. Altschul SF, Madden TL, Schaffer AA, Zhang JH, Zhang Z, et al. 1997. Gapped Blast and Psi-Blast: a new generation of protein database search programs. *Nucleic Acids Res* 25:3389–3402.
98. Kim S, Lund J, Kiraly M, Duke K, Jiang M, et al. 2001. A gene expression map for *Caenorhabditis elegans*. *Science* 293:2087–2092.
99. Huh W-K, Falvo J, Gerke L, Carroll A, Howson R, et al. 2003. Global analysis of protein localization in budding yeast. *Nature* 425:686–691.