Duplicate genes and robustness to transient gene knock-downs in *Caenorhabditis elegans*

Gavin C. Conant and Andreas Wagner[*]

Department of Biology, 167 Castetter Hall, The University of New Mexico, Albuquerque, NM 87131, USA

We examine robustness to mutations in the nematode worm *Caenorhabditis elegans* and the role of single-copy and duplicate genes in it. We do so by integrating complete genome sequence and microarray gene expression data with results from a genome-scale study using RNA interference (RNAi) to temporarily eliminate the functions of more than 16 000 worm genes. We found that 89% of single-copy and 96% of duplicate genes show no detectable phenotypic effect in an RNAi knock-down experiment. We find that mutational robustness is greatest for closely related gene duplicates, large gene families and similarly expressed genes. We discuss the different causes of mutational robustness in single-copy and duplicate genes, as well as its evolutionary origin.

Keywords: *Caenorhabditis elegans*; gene duplication; mutational robustness

1. INTRODUCTION

Genes whose loss of function has no detectable effect number in the thousands in a typical eukaryotic genome (Winzeler *et al.* 1999; Steinmetz *et al.* 2002; Kamath *et al.* 2003). Duplicate genes comprise at least one-third of eukaryotic genomes (Rubin *et al.* 2000; Li *et al.* 2001), a fact that might explain this observation, because duplicate genes often share similar functions. Losing one duplicate gene can thus be tolerated because others can buffer the organism against this loss. This candidate explanation for many genes without phenotypic effects is appealing but also inadequate. A systematic analysis of the effects of knock-out mutations in the yeast *Saccharomyces cerevisiae*, a single-celled eukaryote, showed that much robustness against null mutations is caused by single-copy genes (Wagner 2000*b*). This analysis, based on over 250 synthetic null (gene-knockout) mutations, found that more than 40% of mutations with no phenotypic effect occurred in single-copy genes. It also showed little support for the role of gene duplications in robustness, a result caused by the limited amount of gene-knockout data available at the time. A more recent study (Gu *et al.* 2003), based on more than 5700 synthetic null mutations in yeast, showed that gene duplications have an important role in mutational robustness. However, this later study also underscored the importance of single-copy genes in conferring robustness. Between 41% and 77% of non-detectable mutational effects were caused by single-copy genes, a number higher even than that found in the more limited study.

Whether single-copy or duplicate genes are primarily responsible for mutational robustness has implications for the mechanisms providing robustness. The question itself, however, has thus far been asked only in the unicellular eukaryote yeast *S. cerevisiae*. Multicellular organisms might yield different answers, both because they contain more duplicate genes, which form larger families (Rubin *et al.* 2000; Qian *et al.* 2001; Conant & Wagner 2002), and because developmental processes that arose with

multicellular life may rely on different mechanisms to buffer the effect of null mutations. A recent genome-wide analysis that transiently eliminated the function of more than 16 000 *Caenorhabditis elegans* genes through RNA interference (RNAi) (Fire *et al.* 1998; Kamath *et al.* 2003) allowed us to ask this question for the first time in a higher organism.

Any such analysis has caveats. First, RNAi only temporarily deactivates genes and may not reveal all effects of a synthetic null mutation. This fact, in addition to errors in genome annotation such as the accidental inclusion of pseudo-genes, may contribute to the low proportion of genes with phenotypic effects identified in the RNAi analysis of the worm genome (*ca.* 9% of the genes studied here). However, the effects that the RNAi approach detects are representative of those found with other approaches: 64% of *C. elegans* genes with known knock-out phenotypes can also be detected with RNAi, and, out of those, over 92% give RNAi phenotypes similar to those observed previously (Kamath *et al.* 2003). Second, because RNAi relies on base complementarity between a (denatured) double-stranded RNA and its cognate mRNA, the method may not distinguish between closely related gene duplicates. We alleviate this problem by using only the 13 565 genes for which an RNAi clone specific to the gene—and not affecting multiple targets—was available (Kamath *et al.* 2003) and for which an unambiguous identification in release 73 of Wormpep (all protein-coding genes in the *C. elegans* genome; Stein *et al.* 2001) could be made. Third, unlike microbes, where growth-rate differences can be measured with great accuracy (Steinmetz *et al.* 2002), indicators of fitness cannot be as reliably estimated for multicellular organisms. Fourth—and this is a limitation shared by all laboratory studies—phenotypic effects of mutations are usually assessed in only a small number of environments. That is, they do not necessarily reflect fitness differences. Despite these caveats, the resolution of such experiments is sufficient for our purpose: to distinguish the role of duplicate and single-copy genes in the buffering of mutations.

[*]Author for correspondence (wagnera@unm.edu).

(*a*)

(*b*)

(*c*)

Figure 1. Proportions of genes with (*a*) no RNAi knock-down phenotype, (*b*) a detectable and viable phenotype and (*c*) a lethal phenotype, categorized by gene-family size, that is, the number of paralogues per gene. Adjacent categories with the same shading correspond to cases where the phenotype proportions did not differ significantly (only the categories with one duplicate and two or three duplicates show no differences here; see § 3b). Note the trends in each panel. The proportion of genes with no RNAi phenotype increases with gene-family size, whereas the proportion of genes with either a viable or a lethal phenotype decreases. Absolute numbers of genes in each family-size category are 8861 (single-copy genes), 316 (genes with one duplicate), 1624 (genes with two or three duplicates), 1209 (genes with four to seven duplicates) and 1555 (genes with eight or more duplicates).

$K_a < 2.0$ identified only 20 more duplicate genes, suggesting that our results are not strongly biased by this cut-off. Genes not identified as duplicates under this criterion were treated as single-copy genes. We used release 73 of Wormpep for this analysis (Stein *et al.* 2001), and only genes present in this release of the genome were analysed.

(b) *RNA interference (knock-down) data*

Data on gene knock-down effects were obtained from the RNAi (Fire *et al.* 1998) experiments of Kamath *et al.* (2003). Because interfering RNAs may not distinguish between closely related gene duplicates, we excluded clones annotated as affecting multiple targets (Kamath *et al.* 2003).

We grouped phenotypic knock-down effects into three categories: no phenotype, viable but detectable phenotype and lethal phenotype, and assigned numerical scores to the categories in order of increasing defect: 0 for no phenotype, 1 for moderate (viable) phenotype and 2 for lethal phenotype.

(c) *Effect of gene-family size and evolutionary distance on knock-down phenotype*

We first asked whether the distribution of genes among the three phenotypic categories was affected by the number of paralogues a gene has. We grouped genes into five classes (genes with 0, 1, 2/3, 4–7 or 8 or more paralogues, see § 3b). We then asked whether the proportions of genes with each of the three phenotypes differed (i) between single-copy genes and genes that have one or more duplicates, and (ii) between the gene families of the various sizes.

To address question (i), we calculated the expected number of genes with each of the three knock-down effects among the genes with at least one duplicate, using the phenotypic proportions seen in the single-copy genes. By comparing these three expected values with the observed number of genes of each phenotype among duplicated genes, we were able to use a χ^2 goodness-of-fit test with two degrees of freedom to ascertain statistical significance.

To address question (ii), we used the same approach, limiting our comparisons to adjacent duplication classes. For example, we asked whether the phenotype distribution is the same for genes with one duplicate as for genes with two or three duplicates.

To determine whether phenotypic effects were correlated with the evolutionary distance between duplicates, we compared the proportion of genes in the three phenotypic categories with both the amino acid distance (the fraction, K_a, of substitutions per non-synonymous site) between closest duplicates and the

2. METHODS

(a) *Identification of gene duplicates*

We identified duplicates in the *C. elegans* genome (The *C. elegans* Sequencing Consortium 1998) using our previously described whole-genome analysis tool (Conant & Wagner 2002). For this analysis, we used only duplicate pairs separated by a non-synonymous distance (K_a) of 1.0 or less (calculated by the maximum-likelihood methods of Goldman & Yang (1994) and Muse & Gaut (1994)). Use of a more liberal threshold of

Figure 2. Relationship between non-synonymous distance to nearest gene duplicate (K_a) and proportion of genes with no RNAi knock-down phenotype (Pearson's $r = -0.92$, $p = 0.002$). The dashed line indicates the proportion of single-copy genes with no knock-down effect. In all, 4639 gene pairs were analysed. Inset shows proportions of viable (dashed line) and lethal (solid line) knock-down phenotypes.

synonymous distance (the fraction, K_s, of substitutions per synonymous site) between closest duplicates. We calculated the Pearson product-moment correlation, r, between the distance (K_s or K_a) and the proportion of genes in each of the three phenotypic categories (see § 3c, and figure 5 in electronic Appendix A, available on The Royal Society's Publications Web site). In the case of K_s, we included only duplicate pairs where $K_s < 2.0$ and in which both genes showed an effective number of codons (ENC) (Wright 1990) of greater than 43. This choice of ENC cut-off excludes the *ca*. 10% of genes in the *C. elegans* genome with the lowest values of ENC. Although failing to exclude any genes with high codon bias yields a correlation between knock-down effect and K_s (presumably owing to the association between expression level and knock-down effect seen in § 3d), varying the ENC cut-off so as to exclude between 4% and 30% of genes yields the same result (no significant association) as that reported in § 3c. To test the statistical significance of r, we randomly reshuffled the phenotypic effects with respect to the distances 1000 times and recalculated r for each reshuffled dataset.

To examine whether duplicate genes show similar phenotypic effects, we counted the number of duplicate-gene pairs within a given window of K_a where one member showed no knock-down phenotype and the other showed either a lethal or a moderate phenotype. We tested for significance using the same randomization test.

(d) Association between knock-down effect and gene expression

To identify a statistical relationship between knock-down effect and gene expression, we used a large microarray expression dataset comprising 553 experiments and most *C. elegans* genes (Kim *et al.* 2001). The data consist of logarithmically (log$_2$) transformed expression changes relative to a reference condition that depends


Figure 3. Relationship between non-synonymous distance to nearest duplicate (K_a, *x*-axis) and proportion of genes with asymmetric knock-down effects (*y*-axis; see § 3c for details). Pearson's $r = 0.96$, $p = 0.001$.

on the particular experiment (Kim *et al.* 2001). We identified pairs of duplicate genes (see § 2a) for which RNAi data were present and that were separated by a pairwise K_s of 0.2 or more. For each duplicate pair, we assembled all microarray experiments for which data were available for both genes and calculated the Pearson correlation coefficient (r) between the two genes' expression changes. We then calculated the correlation between the expression similarity of the pairs and their average knock-down effect (calculated using the numerical scheme given in § 2b). A randomization analysis was used for significance testing. We repeated this analysis, substituting the proportion of gene pairs where exactly one gene had a lethal effect (see § 3c) for the average knock-down effect. Using the duplicate pairs identified for the expression analysis above, we next calculated the statistical association between the pairwise correlations in duplicate expression from the experiments by Kim *et al.* (2001) and the pairwise K_a between the duplicates (figure 6 in electronic Appendix A).

To assess whether highly expressed genes show strong knock-down effects, we used results of an experiment (Hill *et al.* 2000) that had determined the expression levels of 18 791 *C. elegans* open-reading frames at eight time points during the worm's life cycle. Using Affymetrix gene chips, these authors estimated the concentration of transcripts (in parts per million (p.p.m.)) at each time point. We considered only the 2624 genes where RNAi knock-down data were available, where a transcript was detected by all hybridization replicates and where that transcript showed an expression level above 20 p.p.m. We compared the log$_{10}$ transform of each gene's highest concentration across the eight time points with the RNAi knock-down effect. We again evaluated significance using a randomization test as outlined in § 2c.

Our final analysis compared the level of gene expression (using the same concentration values as above) with amino acid distance (K_a). We again used the base-ten logarithm of the maximum concentration, comparing it with K_a and determining significance with a permutation test.
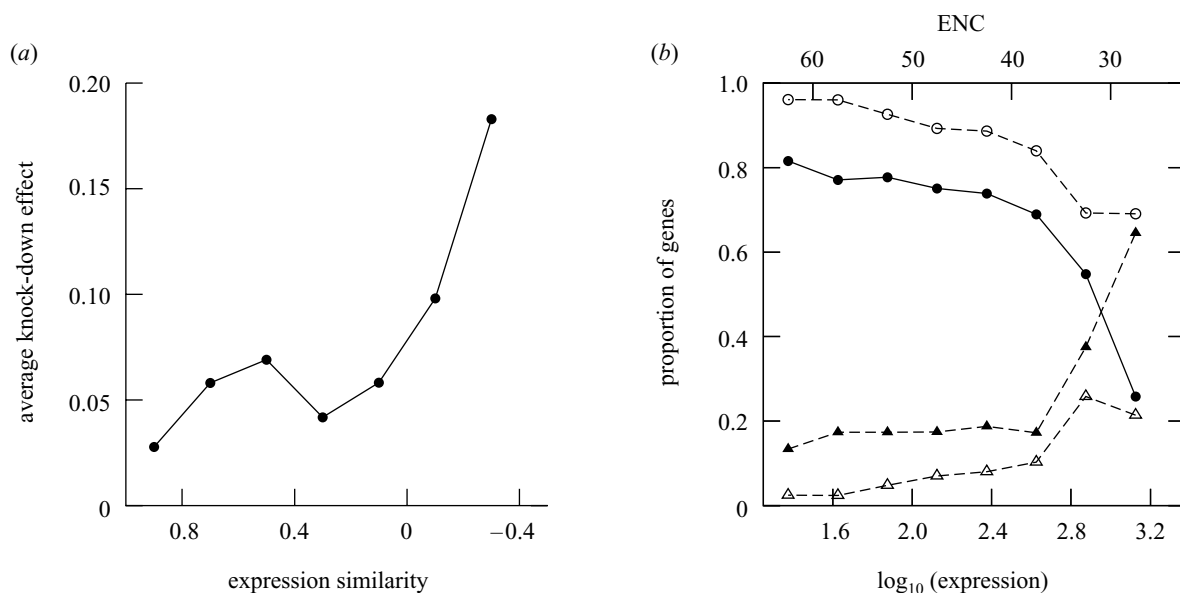
Figure 4. (*a*) Association of duplicate expression similarity and average knock-down effect. The *x*-axis shows the Pearson's *r* for the correlation of the expression levels of the two duplicates, while the *y*-axis shows the average RNAi knock-down effect (see § 2d). Pearson's $r = 0.80$, $p = 0.04$. (*b*) Distribution of knock-down effects by expression level. The *x*-axes show two measures of gene expression: the log of the p.p.m. counts for each gene (the relative expression level, taken at the maximal expression timepoint (closed symbols)—see § 2d) or the ENC (open symbols—see § 2d). The *y*-axis shows the proportions of genes with no knock-down phenotype (circles) and with a lethal knock-down phenotype (triangles).

3. RESULTS

(a) *Many weak phenotypic effects are caused by single-copy genes*

We grouped the phenotypic effects of RNAi (gene knock-down) into three categories: no phenotype (12 387 genes), viable but detectable (moderate) phenotype (395 genes) and lethal phenotype (783 genes). The detectable category groups phenotypes with slow or arrested post-embryonic growth and post-embryonic phenotypes without such growth defects (Kamath *et al.* 2003) together.

Among the 13 565 genes analysed, 8861 are single copy, out of which 88.8% (7872) show no detectable phenotype in an RNAi knock-down experiment. The proportion of genes occurring in gene families of size two or greater that have no RNAi phenotypes is somewhat larger: 96.0% (4515 out of 4704) of such genes have no knock-down phenotype. For the other two classes of phenotype, the relationship is reversed: more lethal phenotypes are caused by single-copy genes (7.5% or 668 genes) than by duplicated genes (2.4% or 115 genes), as are more moderate phenotypes (single-copy genes: 3.6% or 321 genes; duplicate genes: 1.6% or 74 genes). The large numbers of genes involved makes even these small differences statistically highly significant ($\chi^2 = 243$; d.f. = 2; $p < 10^{-10}$). We now examine in greater detail the relationship between gene-family size and RNAi phenotype.

(b) *Gene-family size is correlated with RNA interference phenotype*

Figure 1 demonstrates a correlation between the size of a gene family and the frequencies of the different RNAi knock-down effects. Specifically, the larger a gene family, the more likely it is that its members have no RNAi phenotype (figure 1*a*), and the less likely it is that they have either a detectable (figure 1*b*) or a lethal (figure 1*c*) pheno-

type. Absolute differences in proportions are again small: 88.8% of single-copy genes but 94.0% of genes with one duplicate have no detectable RNAi phenotype. We asked whether any two adjacent size categories in the panels of figure 1 contain equal proportions of genes (see § 2c). Because we are making four comparisons in this analysis, we used a Bonferroni correction (Sokal & Rohlf 1995), performing individual tests at a significance level of 0.0125 to yield a family error rate of 0.05. Adjacent categories of gene-family sizes with the same shading in figure 1 indicate cases where we cannot reject the hypothesis of equal proportions across the three phenotypes. Only the categories with one duplicate and with two or three duplicates show such equal proportions: all others contain different proportions of genes ($p \leqslant 0.0125$). In summary, there is strong evidence that the phenotypic effect detected in knock-down experiments changes with increasing gene-family size.

(c) *The more similar two duplicates are, the less severe is their knock-down effect*

We next examined the proportion of genes with a given phenotypic effect as a function of similarity between duplicates, using the amino acid distance K_a (number of non-synonymous substitutions per non-synonymous site; Li 1997) to measure similarity. The proportion of genes with no phenotypic effect decreases with amino acid distance to the nearest paralogue (figure 2; Pearson's $r = -0.92$, $n = 4639$, $p = 0.002$, correlation calculated using a permutation test on the binned data, see § 2c). Likewise, the proportion of genes with moderate and lethal effects increases with increasing amino acid distance (figure 2 inset; $r = 0.77$, $p = 0.04$ and $r = 0.95$, $p = 0.001$, respectively, $n = 4639$ for both). We also asked whether two duplicates

generally have similar knock-down effects, and found that K_a and the proportion of duplicate pairs with different knock-down effects have a strong positive correlation (Pearson's $r = 0.96$, $n = 3314$, $p < 0.001$). That is, the more distant two duplicates are, the more likely it is that one of them has a more severe knock-down effect than the other (figure 3). Previous genome-scale analyses in various organisms have shown that many duplicate genes have asymmetric sequence or functional divergence, as indicated by protein interactions, sequence divergence and gene-expression patterns (Wagner 2002; Conant & Wagner 2003). For example, for some 30% of worm duplicate genes, one duplicate diverges faster than the other on the amino acid level (Conant & Wagner 2003). Asymmetric divergence, which may increase with amino acid distance and divergence time, could explain why distantly related duplicates often show different mutational effects.

As noted in § 2b, we removed from our analysis all genes with possible cross-reactivity according to Kamath *et al.* (2003). In addition, we assessed whether there were any remaining cross-reactivity biases in the above two analyses by repeating these analyses excluding gene pairs with $K_a < 0.1$. Doing so changed neither the association of knock-down effect and amino acid sequence similarity nor the association of asymmetry of knock-down effect and sequence similarity (data not shown).

We also assessed whether time since duplication affects knock-down phenotypes by comparing knock-down effect with K_s (Li 1997). K_s is a better indicator of divergence time than is K_a because it is subject to fewer evolutionary constraints and thus may change at an approximately constant (neutral) rate (Li 1997). Interpretation of K_s values is confounded by codon usage bias, a feature of very highly expressed genes that can lead to slower rates of synonymous evolution in such genes (Bernardi & Bernardi 1986; Comeron & Aguade 1998). In *C. elegans*, a measure of codon usage bias is the ENC (Wright 1990). It shows a significant correlation (Pearson's $r = -0.57$, Spearman's $s = -0.45$, $n = 3160$, $p < 0.0001$ for both) with a gene's maximum expression level during *C. elegans* development, as measured by oligonucleotide microarrays (Hill *et al.* 2000). We thus eliminated genes with a high codon usage bias (low ENC) before analysis. The remaining genes showed no significant association between K_s and the propensity to have no, a viable or a lethal phenotypic defect ($n = 1791$; no phenotype: $r = -0.13$, $p = 0.39$; viable phenotype: $r = 0.53$, $p = 0.10$; lethal phenotype: $r = -0.59$, $p = 0.14$; see figure 5 in electronic appendix A). To be certain that this lack of association is not an artefact of our permutation test, we have also applied a χ^2 goodness-of-fit test to these data, testing the null hypothesis that the different ranges of K_s all show the same proportions of null, moderate and lethal phenotypes. This test is conservative in the sense that it can reject the null hypothesis even if there is no linear trend in the data. However, the χ^2-test reinforces our conclusions of no association ($\chi^2 = 6.7$, d.f. = 17, $p = 0.99$).

(d) *Expression level and knock-down effect*

Similarity in amino acid sequence is only one indicator of functional similarity among gene duplicates. Studies of individual gene duplicates have shown that functional divergence sometimes occurs through diverging expression patterns rather than diverging sequences (Li & Noll 1994; Hanks *et al.* 1995; Wang *et al.* 1996). This raises the question of whether expression divergence among gene duplicates, which is generally rapid (Wagner 2000a; Gu *et al.* 2002), is also associated with phenotypic effect. To address this question, we compared similarity in expression levels (see § 2d) between duplicate genes with the average RNAi knock-down effect. To avoid artefacts from cross-reactivity in microarray experiments, we excluded duplicate pairs where $K_s < 0.2$. There is a significant correlation between similarity of expression pattern and the average knock-down effect ($r = -0.80$, $n = 3028$, $p = 0.04$; figure 4a). We observe a similar association if we replace the average knock-down effect with the proportion of gene duplicates where one gene shows a lethal knock-down effect while the other does not ($r = -0.88$, $n = 3028$, $p = 0.02$). Excluding genes with high codon usage bias (low ENC) does not change this pattern ($r = -0.81$, $n = 2535$, $p = 0.05$).

It is possible that sequence similarity and expression similarity covary, and hence that the association of each with knock-down effect therefore reflects the same underlying phenomenon. However, the magnitude of the correlation between expression similarity and amino acid sequence distance is small (Pearson's $r = -0.29$, Spearman's $s = -0.27$, $n = 3032$, $p < 0.001$ for both, figure 6 in electronic appendix A, see § 2d). Moreover, considering only restricted ranges of K_a in the above analysis should eliminate the observed correlation in figure 4a if it is truly a result of the covariance of K_a and expression. We determined the association between gene expression level and knock-down effect separately for gene duplicates within five ranges of K_a (0.1–0.2, 0.2–0.3, 0.3–0.4, 0.4–0.5 and 0.5–0.6). Despite small sample sizes (several bins had fewer than 30 elements), four out of the five bins showed a negative association, just as did the complete data.

We also found a statistically significant relationship between maximal expression level and knock-down effect, consistent with the results of others in yeast (Gu *et al.* 2003; Pál *et al.* 2003). Using the highest expression level of each gene measured during eight time points in the worm's life cycle (Hill *et al.* 2000), we found that highly expressed genes are more likely to show a lethal effect (Pearson's $r = 0.77$, $n = 2624$, $p = 0.02$) and less likely to show no effect from knock-down ($r = -0.83$, $n = 2624$, $p = 0.005$; figure 4b). A similar statistical association holds if high codon usage bias (low ENC) is used as an indicator of high expression. (No-effect knock-downs: Pearson's $r = 0.93$, $p = 0.001$; lethal knock-downs: Pearson's $r = -0.89$, $p = 0.005$; $n = 13529$ for both; figure 4b). This result is unsurprising, given the negative correlation of ENC and microarray gene expression levels seen in § 3c.

Finally, it has been noted in yeast (Pál *et al.* 2001) that highly expressed genes are under stronger evolutionary constraints and thus evolve more slowly. Data from duplicate genes in *C. elegans* are consistent with this finding: amino acid distance (K_a) and expression level (data as described for figure 4b) show significant negative correlations (Pearson's $r = -0.80$, $p = 0.01$, $n = 1552$).

4. DISCUSSION

Although the absolute number (7872) of single-copy genes with no knock-down effect is higher than the number of duplicate genes with no knock-down effect, proportionally more duplicate genes have no knock-down effect than do single-copy genes. Kamath *et al.* (2003) noted a similar pattern using a different method of identifying duplicates. In addition, mutational robustness is greatest for closely related and similarly expressed gene duplicates, as well as for duplicates in large gene families. These findings show the important role of both single-copy genes and duplicate genes in robustness against mutation. Weak knock-down phenotypes for duplicate genes can be explained by gene redundancy and overlapping gene functions. Much less clear is how single-copy genes can be eliminated without detectable effect, even though this phenomenon has now been established in two organisms. One possibility is that for many single-copy genes the worm genome harbours at least one other gene with a convergent function, but no sequence similarity. Consistent with this possibility is the observation that sequence-similarity search algorithms miss many genes with dissimilar sequences but convergent tertiary structures (Hubbard *et al.* 1998). Whether such convergent evolution could explain most cases of single-copy genes with no phenotypic effect is unknown. However, the massive scale—more than 7000 genes—at which such convergence would have to occur makes this seem unlikely. A second possibility is that much mutational robustness is the result of interactions of unrelated genes in genetic networks. Mechanistically, this kind of buffering is best understood in metabolic networks. Such networks can compensate for loss-of-function mutations in many (non-redundant) genes by re-routing the flux of metabolites through alternative pathways (Edwards & Palsson 2000).

Is gene redundancy more important in the multicellular worm than in the unicellular yeast? In yeast, 39.5% of single-copy genes versus 64.3% of duplicate genes cause synthetic null mutations with weak or no effect on growth (Gu *et al.* 2003), a ratio of 1 : 1.63. The proportions we found in the worm indicate a ratio of 1 : 1.08, less strongly skewed towards gene duplicates. Conversely, in yeast 29.0% of single-copy genes versus 12.4% of duplicate genes cause synthetic null mutations with lethal effects, a ratio of 1 : 0.43. In the worm, the corresponding percentages are 7.5% and 2.4%, yielding a ratio of 1 : 0.32. From this perspective, gene duplication in the worm is less important than in yeast for causing weak phenotypic effects. However, gene duplication is slightly more important in the worm for preventing lethal phenotypic effects.

A complementary analysis follows that of Gu *et al.* (2003), who estimated lower and upper bounds on the proportion of weak gene knock-out effects that can be attributed to duplicate genes. Their lower bound derives from the assumption that the difference in proportions of mutations with no effect between single-copy genes and duplicate genes is the result of gene duplication. From our worm data, 89% of single-copy genes and 96% of duplicate genes had no knock-down effect. This difference of 7% indicates that at least 323 duplicate genes show no knock-down effect because they are duplicates. The lower

bound in the worm is thus *ca.* 3% (323 out of 12 387), compared with 23% in yeast. The main caveat to this lower bound is that RNAi detects fewer phenotypic effects than does gene knock-out in yeast, biasing the estimate. To obtain an upper bound on the contribution of gene duplicates, Gu *et al.* (2003) assumed that all weak knock-out effects in duplicate genes are caused by redundancy among duplicates. In the worm, this implies that all 4515 duplicate genes with no phenotypic effect showed this phenomenon because of functional redundancy, and hence that roughly 36% of robustness is the result of buffering from duplication. In summary, the available yeast data suggest that the contribution of duplicate genes to weak phenotypic effects is between 23% and 59%, whereas the corresponding range for the worm is 3–36%. An important caveat to this comparison is that synthetic null mutations in yeast and RNAi represent fundamentally different approaches to generating phenotypic effects. Moreover, the patterns of duplication in these two organisms have resulted in different functional distributions of duplicate genes (Conant & Wagner 2002).

Despite uncertainties in estimating the relative contribution of gene duplicates to the buffering of null mutations, it is clear that much gene redundancy exists in eukaryotes. Why is this so? At least three possibilities exist. First, gene redundancy may be an accidental by-product of gene duplication, serving no adaptive role. If so, redundancy is just a transient state after gene duplication. Because multiple lines of evidence indicate that sequence and functional divergences after gene duplication are rapid (Lynch & Conery 2000; Gu *et al.* 2002; Wagner 2002), redundancy should then be observed only in recent gene duplicates. This prediction is contradicted by at least two lines of evidence. First, many genomes contain ancient gene duplicates with very similar functions. Examples include the yeast TPK gene family (catalytic subunits of cyclic AMP-dependent protein kinase; Toda *et al.* 1987) and the yeast CLN gene family (cyclins required for the G1-S transition in the cell cycle; Nasmyth 1993). Although synthetic null mutations in member genes of both (well-characterized) families show only subtle fitness defects (Benton *et al.* 1993; Smith *et al.* 1996), even the youngest duplicate pair within each family is ancient (more than 100 million years old; Wagner 2001). A second line of evidence is our figure 2, which shows that even highly diverged duplicate genes are more likely to show no phenotypic effect in RNAi than are single-copy genes. The age of duplicates cannot be reliably estimated from amino acid divergence. However, for a third of the duplicates shown, synonymous sites on DNA have completely diverged (results not shown), demonstrating that these duplicates are ancient. Mutational robustness through gene redundancy is not just a transient phenomenon.

The second possibility is that redundancy is maintained whenever it is advantageous for an organism to produce copious amounts of gene product (Seoighe & Wolfe 1999). Clearly, for duplicate genes to fulfil such a role, they must maintain a high degree of functional similarity. Consistent with this notion is our observation that highly expressed genes are more likely to have a duplicate with high sequence similarity than are other genes. This pattern has been previously described for duplicate genes in yeast

although it may have other causes (Pál *et al.* 2001). The major difficulty with this argument is that if most redundant gene duplicates are maintained because the genes must be highly expressed, then gene duplication cannot be responsible for many weak gene knock-out effects, because eliminating one of two duplicates would then have deleterious effects. Indeed, our results show that the loss of highly expressed genes in the worm tends to result in severe phenotypic effects.

The last remaining possibility concerns an adaptive role for redundant gene functions. Population genetic modelling (Cooke *et al.* 1997; Nowak *et al.* 1997; Wagner 1999, 2000*c*) has shown that gene redundancy can be maintained by natural selection of genotypes robust against mutations. Such robustness is maintained indirectly, as organisms with redundant genes do not have higher fitness but rather accumulate in populations because they are less susceptible to deleterious mutations. The problem is that the selection pressure is very weak, of the order of the genic mutation rate μ (Wagner 1999, 2000*c*). Redundancy can thus be indefinitely maintained only if mutation rates are very high or populations are very large (effective size $N_e > 1/\mu$; Hartl & Clark 1997). However, even in small populations, this evolutionary mechanism can substantially delay the functional divergence of duplicates and the concomitant loss of redundancy (Wagner 2000*c*). In addition, multifunctional gene duplicates with many pleiotropic interactions can also diverge very slowly in function, even in small populations (Wagner 2000*c*). *Caenorhabditis elegans*, whose populations consist largely of self-fertilizing hermaphrodites, is likely to have a small effective population size. Nevertheless, it shows considerable redundancy in ancient gene duplicates (figure 2), consistent with a slowing of duplicated-gene divergence owing to an adaptive role of redundancy.

In summary, the worm genome contains thousands of single-copy genes with absent phenotypic effects. This phenomenon is most probably the result of complex interactions in genetic networks that are still incompletely understood. Whether such robustness is an evolved or an intrinsic feature of genetic networks is an open question for future research. Conversely, gene duplications also contribute to numerous cases of genes with absent phenotypic effects. Many of these duplicates are ancient, raising the possibility that the functional divergence of genes may be slowed by selection for mutational robustness.

REFERENCES

Benton, B. K., Tinkelenberg, A. H., Jean, D., Plump, S. D. & Cross, F. R. 1993 Genetic analysis of Cln/Cdc28 regulation of cell morphogenesis in budding yeast. *EMBO J.* **12**, 5267–5275.

Bernardi, G. & Bernardi, G. 1986 Compositional constraints and genome evolution. *J. Mol. Evol.* **24**, 1–11.

Comeron, J. M. & Aguade, M. 1998 An evaluation of measures of synonymous codon usage bias. *J. Mol. Evol.* **47**, 268–274.

Conant, G. C. & Wagner, A. 2002 GENOMEHISTORY: a software tool and its application to fully sequenced genomes. *Nucleic Acids Res.* **30**, 3378–3386.

Conant, G. C. & Wagner, A. 2003 Asymmetric sequence divergence of duplicate genes. *Genome Res.* **13**, 2052–2058.

Cooke, J., Nowak, M. A., Boerlijst, M. & Maynard-Smith, J. 1997 Evolutionary origins and maintenance of redundant gene expression during metazoan development. *Trends Genet.* **13**, 360–364.

Edwards, J. S. & Palsson, B. O. 2000 The *Escherichia coli* MG1655 *in silico* metabolic genotype: its definition, characteristics, and capabilities. *Proc. Natl Acad. Sci. USA* **97**, 5528–5533.

Fire, A., Xu, S., Montgomery, M. K., Kostas, S. A., Driver, S. E. & Mello, C. C. 1998 Potent and specific genetic interference by double-stranded RNA in *Caenorhabditis elegans*. *Nature* **391**, 806–811.

Goldman, N. & Yang, Z. 1994 A codon-based model of nucleotide substitution for protein-coding DNA sequences. *Mol. Biol. Evol.* **11**, 725–736.

Gu, Z., Nicolae, D., Lu, H. H.-S. & Li, W.-S. 2002 Rapid divergence in expression between duplicate genes inferred from microarray data. *Trends Genet.* **18**, 609–613.

Gu, Z., Steinmetz, L. M., Gu, X., Scharfe, C., Davis, R. W. & Li, W.-H. 2003 Role of duplicate genes in genetic robustness against null mutations. *Nature* **421**, 63–66.

Hanks, M., Wurst, W., Ansoncartwright, L., Auerbach, A. B. & Joyner, A. L. 1995 Rescue of the En-1 mutant phenotype by replacement of En-1 with En-2. *Science* **269**, 679–682.

Hartl, D. L. & Clark, A. G. 1997 *Principles of population genetics*, 3rd edn. Sunderland, MA: Sinauer.

Hill, A. A., Hunter, C. P., Tsung, B. T., Tucker-Kellogg, G. & Brown, E. L. 2000 Genomic analysis of gene expression in *C. elegans*. *Science* **290**, 809–812.

Hubbard, T. J. P., Ailey, B., Brenner, S. E., Murzin, A. G. & Chothia, C. 1998 SCOP, structural classification of proteins database: applications to evaluation of the effectiveness of sequence alignment methods and statistics of protein structural data. *Acta Crystallographica Section D, Biol. Crystallogr.* **54**, 1147–1154.

Kamath, R. S. (and 12 others) 2003 Systematic functional analysis of the *Caenorhabditis elegans* genome using RNAi. *Nature* **421**, 231–237.

Kim, S. K. (and 12 others) 2001 A gene expression map for *Caenorhabditis elegans*. *Science* **293**, 2087–2092.

Li, W.-H. 1997 *Molecular evolution*. Sunderland, MA: Sinauer.

Li, W.-H., Gu, Z., Wang, H. & Nekrutenko, A. 2001 Evolutionary analyses of the human genome. *Nature* **409**, 847–849.

Li, X. L. & Noll, M. 1994 Evolution of distinct developmental functions of 3 *Drosophila* genes by acquisition of different *cis*-regulatory regions. *Nature* **367**, 83–87.

Lynch, M. & Conery, J. S. 2000 The evolutionary fate and consequences of duplicate genes. *Science* **290**, 1151–1155.

Muse, S. V. & Gaut, B. S. 1994 A likelihood approach for comparing synonymous and non-synonymous nucleotide substitution rates, with application to the chloroplast genome. *Mol. Biol. Evol.* **11**, 715–724.

Nasmyth, K. 1993 Control of the yeast cell cycle by the Cdc28 protein kinase. *Curr. Opin. Cell Biol.* **5**, 166–179.

Nowak, M. A., Boerlijst, M. C., Cooke, J. & Maynard-Smith, J. 1997 Evolution of genetic redundancy. *Nature* **388**, 167–171.

Pál, C., Papp, B. & Hurst, L. D. 2001 Highly expressed genes in yeast evolve slowly. *Genetics* **158**, 927–931.

Pál, C., Papp, B. & Hurst, L. D. 2003 Rate of evolution and gene dispensability. *Nature* **421**, 496–497.

Qian, J., Luscombe, N. M. & Gerstein, M. 2001 Protein family and fold occurence in genomes: power-law behavior and evolutionary model. *J. Mol. Biol.* **313**, 673–681.

Rubin, G. M. (and 49 others) 2000 Comparative genomics of the eukaryotes. *Science* **287**, 2204–2215.

Seoighe, C. & Wolfe, K. H. 1999 Yeast genome evolution in the post-genome era. *Curr. Opin. Microbiol.* **2**, 548–554.

Smith, V., Chou, K. N., Lashkari, D., Botstein, D. & Brown, P. O. 1996 Functional analysis of the genes of yeast chromosome V by genetic footprinting. *Science* **274**, 2069–2074.

Sokal, R. R. & Rohlf, F. J. 1995 *Biometry*, 3rd edn. New York: Freeman.

Stein, L., Sternberg, P., Durbin, R., Thierry-Mieg, J. & Spieth, J. 2001 Wormbase: network access to the genome and biology of *Caenorhabditis elegans*. *Nucleic Acids Res.* **29**, 82–86.

Steinmetz, L. M. (and 10 others) 2002 Systematic screen for human disease genes in yeast. *Nature Genet.* **31**, 400–404.

The *C. elegans* Sequencing Consortium 1998 Genome sequence of the nematode *C. elegans*: a platform for investigating biology. *Science* **282**, 2012–2018.

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.

Wagner, A. 1999 Redundant gene functions and natural selection. *J. Evol. Biol.* **12**, 1–16.

Wagner, A. 2000*a* Decoupled evolution of coding region and mRNA expression patterns after gene duplication: implications for the neutralist–selectionist debate. *Proc. Natl Acad. Sci. USA* **97**, 6579–6584.

Wagner, A. 2000*b* Robustness against mutations in genetic networks of yeast. *Nature Genet.* **24**, 355–361.

Wagner, A. 2000*c* The role of population size, pleiotropy, and fitness effects of mutations in the evolution of overlapping gene function. *Genetics* **154**, 1389–1401.

Wagner, A. 2001 The yeast protein interaction network evolves rapidly and contains few redundant duplicate genes. *Mol. Biol. Evol.* **18**, 1283–1292.

Wagner, A. 2002 Asymmetric functional divergence of duplicate genes in yeast. *Mol. Biol. Evol.* **19**, 1760–1768.

Wang, Y. K., Schnegelsberg, P. N. J., Dausman, J. & Jaenisch, R. 1996 Functional redundancy of the muscle-specific transcription factors Myf5 and myogenin. *Nature* **379**, 823–825.

Winzeler, E. A. (and 50 others) 1999 Functional characterization of the *S. cerevisiae* genome by gene deletion and parallel analysis. *Science* **285**, 901–906.

Wright, F. 1990 The 'effective number of codons' used in a gene. *Gene* **87**, 23–29.