20 October 2004

Molecular dates for the "Cambrian explosion": the influence of prior assumptions

John J. Welch, Eric Fontanillas, and Lindell Bromham

Centre for the Study of Evolution, School of Life Sciences,
University of Sussex, Falmer, Brighton. BN1 9QG

Fossil evidence and DNA sequence data provide complementary sources of information for dating the divergence time of lineages. However, it is well known that estimates using the two sources differ dramatically in some cases. This applies particularly to dating the origins of animal phyla (for a review see Smith and Peterson 2002). The first uncontroversial members of at least half the modern animal phyla appear suddenly in the fossil record of the early to mid-Cambrian period (542-501 Million years ago, Mya). This has led to hypotheses of an explosive origin of metazoans, in which novel body plans were generated in a very short period. However, a wide range of molecular dating studies have suggested that the major lineages of animals arose long before the Cambrian, at over 630 Mya (e.g., Wray et al. 1996; Ayala et al. 1998; Bromham et al. 1998; Lynch 1999). This raises the possibility that there was a long cryptic period of animal evolution preceding the explosion of fossils in the Cambrian. Resolving this controversy is important for understanding the tempo and mode of evolution: were there extraordinary episodes in the evolutionary past that shaped the present, or does evolution work gradually and continually?

A recent study by Aris-Brosou and Yang (2002) introduced a novel molecular dating technique to address this problem. A second study (Aris-Brosou and Yang 2003) used a variant of this method to date two evolutionary events: the divergence of the deuterostomes from the other bilaterian clades, and the divergence of chordates and echinoderms. Their estimate for the former date was 582 +/- 112 Mya, which is the median obtained over a dataset of 22 genes. In contrast to most previous molecular estimates, this date was considered compatible with an origin of the metazoan phyla during or close to the Cambrian period. In addition, the results suggested that rates of molecular evolution were dramatically higher in the late Precambrian.

Aris-Brosou and Yang further showed that the relatively young date estimates were attributable, in part, to the fact that rates were allowed to vary. When the rate of evolution was held fixed across the tree, much older date estimates were obtained (Aris-Brosou and Yang 2002, 2003). This is a point of crucial importance, because many previous molecular dating studies have assumed that a constant-rate molecular clock applies (see Smith and Peterson 2002). There is now a great deal of evidence suggesting that rates of molecular evolution may vary substantially between lineages (e.g., Bromham and Penny 2003). For their data, Aris-Brosou and Yang (2002) used model comparison statistics to show that models allowing rate variation were always preferred (see also Langley and Fitch 1974; Cutler 2000). In this way, Aris-Brosou and Yang (2002) showed how a questionable *a priori* assumption (the assumption of a fixed substitution rate) can lead directly to misleading molecular date estimates. Here, however, we suggest that this conclusion also applies to their own date estimates (Aris-Brosou and Yang 2003), and that these dates, no less than fixed-rate estimates, are strongly influenced by questionable *a priori* assumptions.

BAYESIAN METHODS

Aris-Brosou and Yang's (2002, 2003) studies use a Bayesian statistical approach, in which prior knowledge about the parameters to be estimated is exploited in their estimation (see e.g., Thorne et al. 1998; Huelsenbeck et al. 2001; Holder and Lewis 2003;

Felsenstein 2004). Specifically, the parameters of interest (the divergence dates and molecular rates) are assigned a prior probability density, representing the researcher's beliefs about their values in the absence of the molecular data under consideration. This prior distribution is then multiplied by the likelihood of the branch lengths; suitably normalised, this yields the posterior distribution – that is the probability density of the dates and rates, given the sequence data.

It is shown here that Aris-Brosou and Yang's choice of prior distribution is the major determinant of their young date estimates (or, equivalently, that these estimates are rather insensitive to the molecular data). In Bayesian terms, we say that their prior is highly informative. In principle, there is nothing wrong with assigning informative priors, if we have good reasons to believe the assumptions that they embody, or if we have enough data to override our false assumptions. We suggest, however, that Aris-Brosou and Yang's prior does not accurately reflect the biological situation being modelled, and was not "overridden" by the data. We now examine three components of the prior distribution: the assumptions about evolutionary rates, the distribution of node ages, and the use of fossil information.

Prior distribution of evolutionary rates

Aris-Brosou and Yang's dating method allows for each branch of the tree to be characterised by a distinct rate. To avoid degeneracy, this flexibility must be balanced by some constraints on the possible rate values (Sanderson 1997). In common with most

previous authors, Aris-Brosou and Yang rely on the reasonable assumption that rates change gradually over the tree (see also Sanderson 1997, 2002; Thorne et al. 1998; Huelsenbeck et al. 2000). This basic assumption may be embodied in a large variety of stochastic models of rate change, each of which may be expressed as a prior distribution of rates. Aris-Brosou and Yang (2003) use two such models to obtain their date estimates.

Their first model assumes that molecular rates change according to an Ornstein-Uhlenbeck process (OUP). To understand this model, denote the rate associated with a particular branch by r, and the rate associated with its ancestral branch by r_A. Under the OUP, the probability density function of r, conditional on r_A, is normal (Gaussian) with a mean value that depends on r_A. This is not a legitimate probability density for molecular rates, as it has non-zero probability for negative values of r while, by definition, the rate of molecular evolution cannot be negative. This, however, matters little if negative values take up a very small proportion of the density. More importantly, the expected value of r is given by $E[r \mid r_A] = r_A e^{-\beta \Delta t}$, where Δt is the time duration of the branch, and β is a non-negative parameter (Aris-Brosou and Yang 2002). From this expression, it is clear that each branch is expected have a lower rate than its ancestral branch, and so the OUP predicts that the rate of evolution will decline over time. (Both of these problems with the OUP stem from the fact that it was developed to model the velocity of a particle undergoing Brownian motion; this is because velocities, unlike molecular rates, may be negative, and because friction systematically retards such motion over time.)

Aris-Brosou and Yang's second model of stochastic rate change is the exponential model (EXP). This model simply assumes that the probability density of the rate, r, is an exponential distribution with a mean equal to r_A, the rate of the ancestral branch. In this case, the expected rate does not change between branches, $E[r \mid r_A] = r_A$, so the exponential model is non-directional in this important sense. However, because the exponential distribution is a monotonically decreasing function of r, it follows that any given rate decrease will be assigned a higher probability than any given rate increase, and that a decrease of rate to zero will be assigned the highest probability of all. Furthermore, when the overall probability of a rate decrease is considered, we have, $\Pr[r < r_A] = 1 - e^{-1} \approx 63\%$. For these two reasons, it is fair to say that the model favours rate decreases, even though the expected rate does not change. In addition, as Aris-Brosou and Yang point out, the exponential model has the corollary that rapidly evolving lineages are more likely to give rise to daughter lineages with highly variable rates (this is because the variance of an exponential distribution is equal to its squared mean).

These considerations show that both models of rate change embody assumptions that cannot be easily justified on biological grounds. In particular, both predict that molecular rates change in a directional fashion, favouring rate decreases over time. This is likely to have influenced Aris-Brosou and Yang's (2003) finding that estimated rates were faster in the Precambrian, towards the root of their phylogenies. However, Aris-Brosou and Yang (2002) compared a wide range of variable-rate models, and showed that the date estimates obtained were quite insensitive to the particular model chosen (although all differed substantially from results when a fixed rate was assumed). As such, it is unlikely

that the young date estimates were an artefact of the assumptions of these models alone.

Prior distribution of relative node ages

The second component of Aris-Brosou and Yang's prior is the probability density of a set of internal node times. These times are scaled relative to the age of the ingroup root so they lie on the range zero (the present day) to one (the basal node). To obtain the relevant distribution, Aris-Brosou and Yang use a mechanical model of speciation and extinction. Specifically, it is assumed that diversification has taken place via a linear birth-death process with constant speciation rate λ, and extinction rate μ. The sequence data are then assumed to contain a random sample of the extant species, comprising a fraction ρ of the total number (Nee et al. 1994; Yang and Rannala 1997). As with the rate priors, the date prior contains strong assumptions about the processes that gave rise to the sequence data. In particular, the model assumes i) that speciation and extinction rates have remained constant throughout the relevant period, ii) that neither rate shows any correlation with the rate of molecular evolution, and iii) that the dataset contains a random sample of extant metazoans. There is empirical evidence that all three of these assumptions do not hold. For example, contrary to assumption i), Orme et al. (2002) demonstrate highly uneven diversification rates between metazoan lineages; while contrary to assumption ii), Barraclough et al. (1998) suggest a possible correlation between diversification rate and the rate of molecular evolution. Regarding assumption iii), the taxa chosen do not appear to be a random sample in any case (for example, the phylogenies contain proportionally

too many chordates and too few arthropods given relative species richness – see Orme et al. 2002). Of course, any theoretical model is an idealisation that is susceptible to nit-picking criticisms; but, unlike the rate priors, it is clear that, as recognised by Aris-Brosou and Yang (2003), the functional form of the date prior does have a substantial effect on the date estimates obtained. Three lines of evidence show that this is the case.

Firstly, dates estimated using *only* the prior distribution are very similar to those obtained from the posterior. Note that the prior distribution depends only on the topology of the tree, which was fixed at the outset of the analysis, and so makes no use at all of the molecular data. Aris-Brosou and Yang (2003) reported a date estimate obtained from their prior alone, using the topology assigned to their mitochondrial dataset. We repeated this procedure with their 18S rRNA topology (since this is their largest dataset, and was used in both studies). The results are shown in Table 1 in the row labelled "no data." Also shown, for comparison, are the results with the data reported by Aris-Brosou and Yang (2003); these are labelled "original data". The table contains date estimates for the divergence of the deuterostomes from other bilaterian clades (referred to as the protostome-deuterostome split), and for the divergence of echinoderms and chordates. In addition, estimates obtained under the fixed-rate molecular clock assumption (labelled CLOCK), and estimates obtained under the two stochastic models of rate change (labelled EXP and OUP) are compared. Clearly, for the variable rate models (though not for the fixed rate molecular clock), date estimates from the prior alone are very close indeed to those obtained with the sequence data. For example, including the data alters the estimate of the protostome-deuterostome split by less than 30My.

A second indication that the date estimates were predominately due to the prior specified comes from the insensitivity of these estimates to serious errors in the data. As has been mentioned, the topology of each tree was assumed known and fixed before the date estimation. For the 18S gene, Aris-Brosou and Yang used the topology of Nielsen (1995). Due to subsequent developments, this topology is not undisputed – it has been suggested, for example, that the hemichordates group with the echinoderms rather than the chordates (see Bromham and Degnan 1999) – but evidence from the dating study of Yoder and Yang (2000) suggests that most reasonable topologies will yield similar date estimates. However, the supplementary material of Aris-Brosou and Yang (2003) shows that some of the additional genes used were assigned topologies with errors that are serious enough to affect the results. For example, for the 11 mitochondrial genes the Collembola (springtails) were grouped with the brachiopods, and both placed basal to the deuterostomes. Although the brachiopod placement follows Nielsen (1995), they are more commonly placed within the Lophotrochozoa, while the Collembola placement conflicts with the 18S topology (where they are grouped with the arthropods). These phylogenetic relationships directly affect the placement of the nodes to be dated. Errors in the topologies assigned to EF-1 and α-tubulin genes were more widespread (e.g., insect and crustacean taxa were placed in multiple clades containing nematodes, mammals, and mollusks; while *Xenopus* is contained within a clade of bony fish, which is nested within mammals). In addition to these topological errors, other trees showed errors in the placement of the fossil calibration points, and of the nodes to be dated (e.g., the node denoting the divergence of the deuterostomes from other bilaterians was misidentified).

Despite this, the results reported by Aris-Brosou and Yang (2003) show that none of the genes with serious errors yielded anomalous results. To demonstrate this further, consider a corrected dataset, excluding the 13 genes with serious topological errors. H1-Histone was also excluded, as none of the eight fossil calibration points used by Aris-Brosou and Yang (2002, 2003) appears on this tree. In addition, the date estimates for 18S and actin were recalculated using the supplementary material (although no estimate of the echinoderm-chordate split could be obtained for actin, as the tree contains no echinoderms). Table 1 shows the median date obtained from this corrected dataset. Despite the significant changes in the data, the estimated dates for the variable rate models remain very close indeed to the estimates obtained from the original dataset.

The third indication that the divergence-time prior has a major influence on the results is the sensitivity of the date estimates to the parameters used in this prior. Rather than estimate the parameters of the birth-death process, λ, μ, and ρ, Aris-Brosou and Yang have them integrated out; so, for example, ρ is assigned a "hyperprior" – a uniform probability density on the range (ρ_{min}, ρ_{max}) – and then the divergence time prior is averaged over all possible values of ρ weighted by their probability of occurrence. This is a standard Bayesian practice for dealing with "nuisance parameters" that are not of direct interest. In a sense, however, this procedure defers the problem because hyperpriors will themselves require parameterisation (in this case, we must specify ρ_{min} and ρ_{max}) and the possibility remains that the date estimates will be sensitive to *these* parameters. Because ρ is a sampling proportion, its absolute limits are $\rho_{min}=0$ and $\rho_{max}=1$. However, for all of their main results, Aris-Brosou and Yang set $\rho_{max}=0.001$ (the value used here for Table

1). In addition, two results are reported for ρ_{max}=0.0001 and ρ_{max}=0.5. In the former case (ρ_{max}=0.0001), date estimates were close to those obtained with ρ_{max}=0.001, but when ρ_{max}=0.5, the estimate of the protostome-deuterostome split increases by over 200My, to 791 +/- 246 Mya. Figure 1 suggests why this is so. Yang and Rannala (1997) showed that the multivariate probability density of a set of node times under the birth-death process can be expressed as the order statistics of a kernel distribution (to sample a complete set of node ages, one can make independent samples from this kernel and then order them). Figure 1 plots this kernel distribution for the three values of ρ_{max} discussed (with ρ_{min}=0 in all cases). Two sets of curves are shown. The dotted lines show the distribution with the three parameters (λ, μ, and ρ) integrated out over uniform hyperpriors, and these hyperpriors are those specified by Aris-Brosou and Yang (2003). The solid lines, in contrast, show the distribution for fixed values of each of the parameters, and these fixed values are the hyperprior means. Formally, denoting the kernel distribution $K(\bullet)$, the dotted lines show

$$\int_0^{15} d\lambda \int_0^{5} d\mu \int_0^{\rho_{max}} d\rho \;\; K(\text{age};\, \lambda,\mu,\rho) \;\; (15 \times 5 \times \rho_{max})^{-1},$$

while the solid lines show $K(\text{age};\, \bar{\lambda},\bar{\mu},\bar{\rho})$, where $\bar{\lambda} = 15/2$, $\bar{\mu} = 5/2$ and $\bar{\rho} = \rho_{max}/2$. From Figure 1, it is clear that in both cases (solid and dotted lines) ρ_{max}=0.001 and ρ_{max}=0.0001 give qualitatively similar curves. This explains the similar date estimates that they yield. In addition, it is clear that these small values of ρ_{max} will assign higher probabilities to internal nodes that are closer to the root. In other words, choosing small values of ρ_{max} favours trees with short internal branches and clustered internal nodes. Crucially, this tendency of the prior to place the internal nodes in close proximity, "pushes" the fixed calibrated nodes close to the (basal) nodes whose dates are being

estimated. This is tantamount to assuming that an explosive radiation occurred in the Cambrian, and explains why young date estimates are obtained (see also Perez-Losada et al. 2004).

So far, it has been suggested that the birth-death model with random sampling is a questionable model of speciation and extinction, and that it leads to a highly informative prior. However, use of the model might still be justified if it were treated as a phenomenological model. As such, ρ would be viewed not as a sampling fraction, but merely as a parameter altering the shape of the curve plotted in Figure 1, and so controlling the relative lengths of internal branches, and those leading to the tips of the phylogeny. Taking this view, the use of $\rho_{max}=0.001$ could be justified by showing that a model with short internal branches best fits the data. This might involve, for example, the use of Empirical Bayes methods (e.g., Carlin and Louis 2000). Aris-Brosou and Yang (2003) compare the fit of the model with the three different values of ρ_{max} by calculating the Posterior Bayes Factor (PBF). (The PBF is the likelihood, weighted by the posterior, and integrated over all possible values of the rates and dates – the use of the posterior rather than the prior contrasting with the conventional Bayes Factor.) Using this test, Aris-Brosou and Yang showed that the smallest value, $\rho_{max}=0.0001$, was preferred to $\rho_{max}=0.001$ (which was the value used for the main body of results), while the value favouring long internal branches, $\rho_{max}=0.5$, was strongly rejected. Although the use of model testing is a potentially powerful approach, there are reasons to view even a thorough test with caution in this case. Firstly, as Aris-Brosou and Yang (2003) recognise, the use of the PBF is controversial. Particular problems stem from the multiple

uses of the rather limited data – twice in the PBF (since it appears in the likelihood and the posterior) and then again to estimate the dates (see commentaries to Aitken 1991). Secondly, a rate prior favouring rate decreases was specified, and this will lead to a date prior with short internal branches being preferred, even if no such trend is present in the raw branch-length estimates. Thirdly, of course, all tests rely on accurate assignment of topology.

Incorporation of fossil information

Although we are unavoidably ignorant about the events that have given rise to the metazoan tree, there is some evidence of which we can be relatively confident – namely fossil evidence. Aris-Brosou and Yang (2003) make use of palaeontological dates, but only after the Bayesian estimation procedure has been completed. As has been mentioned, this procedure estimates internode ages scaled relative to that of the ingroup root. Afterwards, point estimates of particular divergence dates are used to convert these relative dates into real time. When multiple calibration points were available (and each mitochondrial tree had four, for example), calibration was undertaken separately for each fossil date, and the median value of the resulting times was reported (this was the procedure used to produce Table 1). This procedure differs from that of previous studies, in which fossil evidence was incorporated directly into the date prior (e.g., Sanderson 1997, 2002; Thorne et al. 1998; Kishino et al. 2001; Perez-Losada et al. 2004).

Aris-Brosou and Yang (2003) remark that their approach is inferior to the simultaneous use of multiple genes, and calibration points. A major reason for this is that the use of multiple calibrations, one at a time, can lead to internally inconsistent results (Douzery et al. 2003; Perez-Losada et al. 2004). For example, the upper bound of the estimated date of the protostome-deuterostome split given by Aris-Brosou and Yang (2003) is younger than four of their eight fossil calibration points, all of which are its daughter nodes. Although, this inconsistency stems in part from the separate use of multiple genes, similar problems arise within a single tree. To demonstrate this, Table 2 compares the date estimates obtained for the 18S tree using each of the eight available calibration points. (These were taken from Bromham et al. 1998, and are listed below Table 2). Each row shows the estimated ages of a set of internodes using a single calibration point. Each column gives the estimates obtained for a single internode using each of these calibrations. The eight nodes chosen are those corresponding to the eight fossil calibration points (the number in square brackets is the node label from the supplementary material of Aris-Brosou and Yang 2003). Due to this arrangement, the diagonal (shown in bold) reproduces the calibration dates. If Aris-Brosou and Yang's procedure led to reasonably consistent results, then we would expect each column to contain very similar dates. Clearly, this is not the case. The fixed-rate molecular clock (Table 2a) leads to results that are wildly inconsistent, while even the variable-rate OUP model (Table 2b), whose results are far more consistent, contains a wide variation in each column. Consider, for example, calibration point 6 (which is the divergence of jawless fish from the rest of the vertebrates). Column 6 of Table 2b shows that, in this case, all seven other calibrations produce date estimates younger than the fossil evidence.

One result of this inconsistency is that it becomes difficult to provide meaningful confidence intervals for the estimated dates. This negates one of the great attractions of Bayesian methods: the ability to obtain a measure of confidence directly from the posterior distribution (Holder and Lewis 2003). Difficulties with interpreting the confidence intervals are exacerbated by the separate use of multiple genes, and the fact that errors associated with the fossil estimates are not taken into account (e.g., Weiss and Marshall 1999).

CONCLUSIONS

Aris-Brosou and Yang's (2002, 2003) work shows that the results of molecular dating studies may be very sensitive to their underlying assumptions. They placed particular emphasis on the assumption of the fixed-rate molecular clock – stressing that this restrictive assumption is both biologically implausible, and can lead to consistently inaccurate date estimates (a conclusion supported here by Tables 1 and 2). We have shown that similar conclusions apply to the variable-rate date estimates of Aris-Brosou and Yang (2003). Their Bayesian analysis requires the assignment of a prior distribution of molecular dates and divergence dates, and this prior has been shown to embody strong assumptions about directional changes in molecular rates, and about patterns of diversification. The latter assumptions, in particular, bias the estimation procedure in favour of young date estimates. As such, the young estimates obtained are more plausibly attributed to these prior assumptions than to any signal from the data. Aris-Brosou and

Yang (2003) justify their prior assumptions about diversification by employing model-testing methods. Although we have suggested that the particular tests used are not convincing, the use of model testing in this way sets a valuable precedent that should be followed by future studies (Arbogast et al. 2002).

It is important to note that caveats similar to those mentioned above, may also apply to the conclusions of other studies applying molecular dating techniques to the origin of the metazoan phyla. For example, Bromham and Hendy (2000) suggested that applying universally rapid rates across early branches of the phylogeny could not reconcile molecular and fossil dates. However, this conclusion is contingent on the assumption that rates in the Cambrian fell within the range of rates empirically estimated from their data set. Contrasting conclusions were reached by Peterson et al. (2004), whose date estimates are close to those of Aris-Brosou and Yang (2003). However, amongst other things, the study of Peterson et al. (2004) demonstrates the sensitivity of molecular dating to assumptions about variation in evolutionary rates across sites. When all sites were assumed to evolve at the same rate, Peterson et al. dated the protostome-deuterostome split at 573Mya, but when between-site rate variation was incorporated, this estimate increased by over 80 million years, to 656Mya. No formal tests for between-site rate variation were reported, but given previous results for individual proteins (e.g., Zhang and Gu 1998), it is most likely that extensive between-site rate heterogeneity would be detected for Peterson et al.'s concatenated sequence.

As a result of the above considerations, it would be premature to conclude that the conflict between palaeontological and molecular dates for the origin of animals has been resolved. Instead, attention should continue to focus on the susceptibility of molecular dating to methodological artifacts resulting from inaccurate prior assumptions.

ACKNOWLEDGEMENTS

Stephane Aris-Brosou and Ziheng Yang were generous with their time and source code, and provided many helpful corrections to our manuscript. Thanks are also due to David Waxman, Meg Woolfit and Mark Broom.

REFERENCES

Aitken, M. 1991. Posterior Bayes factor. J. R. Stat. Soc. B Met. 53:111-142.

Arbogast, B. S., S. V. Edwards, J. Wakeley, P. Beerli and J. B. Slowinski. 2002. Estimating divergence times from molecular data on phylogenetic and population genetic timescales. Ann. Rev. Ecol. Syst. 33:707-740.

Aris-Brosou, S. and Z. Yang. 2002. Effects of models of rate evolution on estimation of divergence dates with special reference to the metazoan 18S ribosomal RNA phylogeny. Syst. Biol. 51(5):703-714.

Aris-Brosou, S. and Z. Yang. 2003. Bayesian models of episodic evolution support a late precambrian explosive diversification of the metazoa. Mol. Biol. Evol. 20:1947-1954.

Ayala, F. J., A. Rzhetsky, and F. J. Ayala. 1998. Origin of the metazoan phyla: Molecular clocks confirm paleontological estimates. Proc. Natl. Acad. Sci. USA 95(2):606-611.

Barraclough, T. G., A. P. Vogler, and P. H. Harvey. 1998. Revealing the factors that promote speciation. Philos. T. Roy. Soc. Lond. B 353:241-249.

Bromham, L. D., and B. M. Degnan. 1999. Hemichordates and deuterostome evolution: robust molecular phylogenetic support for a hemichordate+echinoderm clade. Evol. Dev. 1:166-171.

Bromham, L. D. and M. D. Hendy. 2000. Can fast early rates reconcile molecular dates with the Cambrian explosion? Proc. R. Soc. Lond. B. 267:1041-1047.

Bromham, L., and D. Penny. 2003. The modern molecular clock. Nat. Rev. Genet. 4(3):216-224.

Bromham, L., A. Rambaut, R. Fortey, A. Cooper, and D. Penny. 1998. Testing the Cambrian explosion hypothesis by using a molecular dating technique. Proc. Natl. Acad. Sci. USA 95:12386-12389.

Carlin, B. P., T. A. Louis. 2000. Bayes and empirical Bayes methods for data analysis: 2nd Ed. Chapman and Hall.

Cutler, D. J. 2000. Estimating divergence times in the presence of an overdispersed molecular clock. Mol. Biol. Evol. 17(11):1647-1660.

Douzery, E. J. P., F. Delsuc, M. J. Stanhope, and D. Huchon. 2003. Local molecular clocks in three nuclear genes: divergence times for rodents and other mammals and incompatibility among fossil calibrations. J. Mol. Evol. 57:S201-S213.

Felsenstein, J. 2004. Inferring Phylogenies. Sinauer Associates, Sunderland MA.

Holder, M. and P. O. Lewis. 2003. Phylogenetic estimation: traditional and Bayesian approaches. Nature Rev. Genet. 4:275-284.

Huelsenbeck, J. P., B. Larget and D. Swofford. 2000. A compound Poisson process for relaxing the molecular clock. Genetics 154:1879-1892.

Huelsenbeck, J. P., F. Ronquist, R. Nielsen and J. P. Bollback. 2001. Bayesian inference of phylogeny and its impact on evolutionary biology. Science 294:2310-2314.

Kishino, H., J. L. Thorne and W. J. Bruno. 2001. Performance of a divergence time estimation method under a probabilistic model of rate evolution. Mol. Biol. Evol. 18(3):352-361.

Langley, C. H., and W. M. Fitch. 1974. An estimation of the constancy of the rate of molecular evolution. J. Mol. Evol. 3:161-177.

Lynch, M. 1999. The age and relationships of the major animal phyla. Evolution 53(2): 319-325.

Nee, S., R. M. May and P. H. Harvey. 1994. The reconstructed evolutionary process. Philos. T. Roy. Soc. B. 344:305-311.
Nielsen, C. 1995. Animal evolution: interrelationships of the living phyla. Oxford University Press, Oxford.
Orme, C. D. L., D. L. J. Quicke, J. M. Cook, and A. Purvis. 2002. Body size does not predict species richness among the metazoan phyla. J. Evol. Biol 15:235–247.
Perez-Losada, M., J. T. Hoeg, and K. A. Crandall. 2004. Unravelling the evolutionary radiation of the Thoracian barnacles using molecular and morphological evidence: a comparison of several divergence time estimation procedures. Syst. Biol. 53:244-264.
Peterson, K. J., J. B. Lyons, K. S. Nowak, C. M. Takacs, M. J. Wargo, and M. A. McPeek. 2004. Estimating metazoan divergence times with a molecular clock. Proc. Natl. Acad. Sci. USA 101(17):6536-6541.
Sanderson, M. J. 1997. A nonparametric approach to estimating divergence times in the absence of rate constancy. Mol. Biol. Evol. 14(12): 218-1231.
Sanderson, M. J. 2002. Estimating absolute rates of molecular evolution and divergence times: A penalized likelihood approach. Mol. Biol. Evol. 19(1):101-109.
Smith, A. B., and K. J. Peterson. 2002. Dating the time of origin of major clades: molecular clocks and the fossil record. Annu. Rev. Earth and Pl. Sc. 30:65-88.
Thorne, J. L., H. Kishino and I. S. Painter. 1998. Estimating the rate of evolution of the rate of molecular evolution. Mol. Biol. Evol. 15(12):1647-1657.
Weiss, R. E., and C. R. Marshall. 1999. The uncertainty in the true end point of a fossil's stratigraphic range when stratigraphic sections are sampled discretely. Math. Geol. 31(4):435-453.
Wray, G. A., J. S. Levinton, and L. H. Shapiro. 1996. Molecular evidence for deep precambrian divergences among metazoan phyla. Science 274(5287):568-573.
Yang, Z., and B. Rannala. 1997. Bayesian phylogenetic inference using DNA sequences: a Markov chain Monte Carlo method. Mol. Biol. Evol. 14(7):717-724.
Yoder, A. D., and Z. Yang. 2000. Estimation of primate speciation dates using local molecular clocks. Mol. Biol. Evol. 17(7):1081-1090.
Zhang, J. and X. Gu. 1998. Correlation between the substitution rate and rate variation among sites in protein evolution. Genetics 149:1615-1625.

TABLE 1: estimated divergence dates (Mya) under three models of rate change

	protostome-deuterostome			echinoderm-chordate		
	CLOCK	EXP	OUP	CLOCK	EXP	OUP
No data:	556	555	556	545	543	546
Original data:	1090	579	582	736	515	533
Corrected data:	789	597	576	707	585	533

Estimated ages of divergence (millions of years before present) obtained with the method of Aris-Brosou and Yang (2003) are shown. Results obtained under the assumption of a constant-rate molecular clock (CLOCK), are compared to those obtained under two models of rate change: the Ornstein-Uhlenbeck Process (OUP) and the Exponential model (EXP) – see text. Three data-sets are compared. "No data" gives the estimates obtained solely from the prior distribution, assuming Aris-Brosou and Yang's (2002, 2003) 18S topology; "Original data" gives results directly from Aris-Brosou and Yang (2003); and "Corrected data" shows a subset of these results from which genes with seriously erroneous topology have been excluded, and with corrected placement of fossil calibration points (see text). Each date estimate is the median of the mean estimate obtained for each relevant gene and calibration point.

TABLE 2: estimated dates (Mya) using individual calibrations and 18S rRNA data

2a) Fixed rate molecular clock (CLOCK)

		node dated							
		1 [57]	2 [58]	3 [46]	4 [48]	5 [69]	6 [55]	7 [66]	8 [64]
calibration used	1	**390**	121	113	128	142	284	170	362
	2	1310	**405**	381	429	476	953	572	1215
	3	1437	444	**418**	470	523	1045	627	1332
	4	1308	404	380	**428**	476	951	571	1213
	5	1375	425	400	450	**500**	1000	600	1275
	6	701	217	204	230	255	**510**	306	650
	7	1192	368	347	390	433	867	**520**	1105
	8	572	177	166	187	208	416	249	**530**

2b) Ornstein-Uhlenbeck Process model of rate change (OUP)

		node dated							
		1 [57]	2 [58]	3 [46]	4 [48]	5 [69]	6 [55]	7 [66]	8 [64]
calibration used	1	**390**	359	353	365	421	390	409	427
	2	440	**405**	398	412	475	440	461	482
	3	462	425	**418**	433	499	462	484	506
	4	457	421	413	**428**	493	457	479	501
	5	463	426	419	434	**500**	463	485	507
	6	510	470	461	478	550	**510**	534	559
	7	496	457	449	465	536	496	**520**	544
	8	484	446	438	453	522	484	507	**530**

Estimated ages of divergence (millions of years before present) obtained with the method of Aris-Brosou and Yang (2003) are shown. Table 2a shows results obtained under the assumption of a constant-rate molecular clock, and Table 2b gives the equivalent results under the Ornstein-Uhlenbeck Process model of rate change. Each row contains date estimates obtained with a different fossil-calibration point, while each column contains date estimates obtained for a single node using each of these calibrations. On the diagonal of each table, the node dated corresponds to the calibration point, and here the calibration dates are given in bold-type. Retaining the terminology of Bromham et al. (1998) and Aris-Brosou and Yang (2002, 2003), the calibration points are (1) Collembola-Pterygota, (2) Aranaea-Scorpionida, (3) Coelacanth-Dipnoi/Tetrapoda, (4) Osteichthyes-Dipnoi/Tetrapoda, (5) Asteroidea-Echinoidea, (6) Agnata-Gnathostoma, (7) Arachnida-Merostomata, (8) Cephalochordata-Chordata. The numbers in square brackets correspond to the arbitrary labels given to each node in the supplementary material of Aris-Brosou and Yang (2003). Each estimate is obtained from the mean of the marginal posterior distribution for the relevant node.

Figure 1 Caption:

The kernel distribution associated with Aris-Brosou and Yang's (2002, 2003) node age prior is shown. This distribution is derived from the birth-death process with random sampling of extant species, and is the fairly lengthy expression given by Eqs. (1), (2), (4) and (14) of Yang and Rannala (1997). The kernel distribution is a function of node age scaled between 0 (the present day), and 1 (the basal node). It takes three parameters, λ (speciation rate), μ (extinction rate) and ρ (proportion of species sampled). Two different versions of the function are plotted. The dotted lines show the kernel with its three parameters integrated out over uniform hyperpriors; these hyperpriors are those used by Aris-Brosou and Yang 2003. The solid lines show the kernel with a fixed value of each parameter; these values are the means of the hyperpriors. In each case, curves for three values of ρ_{max} (the upper bound of the sampling proportion hyperprior) are shown. These are the three values compared by Aris-Brosou and Yang (2003).

FIGURE 1. node age prior, kernel distribution