this error is very small, but in an asbestos exposed cohort it may have a substantial effect. Leaving the miscoded mesotheliomas in the lung cancer count will overstate the true lung cancer SMR. Excluding them will in theory understate it, but only to the small extent that this error affects the population as a whole. The best available approximation to a true estimate of the risk is therefore to exclude the miscoded mesotheliomas, and this has been done for this review.

Derivation of cohort mean exposure estimates

Mean exposure for cohorts was calculated in different ways, depending on the available information. When data was given for separate exposure groups, the cohort mean was calculated by weighting the individual group means by the expected deaths from lung cancer in the group. On the assumption that excess risk is proportional to cumulative exposure, this weighting preserves the same proportionality when the results from subgroups with different exposures are aggregated, it is therefore the optimal statistical measure of aggregate exposure.

Where mean exposure values for individual dose categories were not given, the midpoints were used. The top exposure category was usually given as an open interval (e.g. exposures>100 f/ml.yr): in these cases a value was chosen based on a view of the highest likely exposure and the distribution of individuals across all exposure categories. It was assumed that where the highest category contains a relatively small proportion of the population, the category mean will be a smaller multiple of the lower band than otherwise.

For cohorts where results for exposure specific subgroups were not given, the cohort mean was either given directly (cohorts 4, 13 and 15); derived from information given on the distribution of individual doses (cohorts 1 and 7), or on the exposure of internal controls (cohort 17), or by multiplying a mean exposure level by mean exposure duration (cohorts 8 and 14).

Exposure estimates given in particle counts were converted to counts of 'regulated fibres' (fibres with an aspect ratio greater than 3:1, and length>=5 microns), using conversion factors calculated by the report authors where possible. The most commonly used conversion was 1 mppcf (million particles per cubic foot)=3 f/ml (fibres per millilitre), and this was the value adopted for the Johns Manville cohort, where a conversion was not given. For the Massachusetts cohort, where the fibre involved was crocidolite (rather than chrysotile as in the other cohorts with particle counts), an independent expert hygienist was asked for an assessment (see Appendix B).

The exposure estimates for Wittenoom have been questioned by Rogers (1990) who has suggested—having re-examined some of the original samples using modern light and electron microscopy—that the levels may have been underestimated by up to a factor of 10. Details of this reassessed data were to be published, but these have not so far appeared in print. It is therefore difficult to know whether to make an adjustment to the published estimates, and if so by how much. Similar comments may of course apply to other cohorts and introducing a correction might then distort rather than correct the overall picture. de Klerk and colleagues, developing estimates of environmental risk at Wittenoom (1992) use a factor of 4 without detailed discussion. The effect of using this adjusted exposure level is examined as a variant of the main analyses.

Exposure-specific risk estimates

It is generally assumed that the most reliable guide to dose-specific risk is provided by exposure analyses using estimates of individual exposure. This is clearly the case when these individual exposure values can be accurately determined. However this assumption is very much not the case in the studies in this review. Not only are there the inevitable problems of extrapolating earlier exposures on the basis of more recent measurements; there are also problems of converting the most usual historic measurements (in terms of particle counts) to the more relevant measure of fibre counts. Direct fibre counting only became generally used in the 1970s.

In these circumstances it is at least arguable that global assessments of average exposure, set against overall mortality outcomes, should be preferred. Exposure–response regressions with inaccurate individual exposure assignments will produce a slope estimate biased downwards. Use of an overall assessment will also minimise the error introduced by conversion from particle counts to fibres, since these average conversion factors will represent a more accurate conversion for the totality of exposure than for a particular individual.

However, the arguments are not all one way. Overall mortality outcomes can only be assessed against some outside reference—usually the regional or national population—and this may not represent a true baseline level for the exposed population in question. Assessment of an internal exposure response gives some check on this issue. A complete absence of exposure response must cast some doubt on any overall excess being counted as a measure of risk (the Albin and Connecticut cohorts are examples of this).

Cohort-level risk measures were chosen for this review both because these allow a wider range of data to be assessed than if attention is restricted to internal exposure response analyses and since (as argued above) cohort-level exposure estimates are likely to be more accurate than individual exposures.

Smoking

The evidence on the joint effect of smoking and asbestos exposure on lung cancer has been reviewed