Appendix

Some intuition on the structural model itself

As presented in the paper, the structural model is:

$$T_{\bar{0}}(\psi) = \int_0^T \exp[\psi \times A(t)]\, dt$$

In this equation, T is the observed survival time, $A(t)$ denotes the observed exposure (mg/m^3) in year t, and ψ represents the unknown constant parameter of the model; when we use its true value, $T_{\bar{0}}(\psi)$ represents the counterfactual survival time under no exposure.

Note that if $\psi = 0$, then regardless of the exposure, $\exp[0 \times A(t)] = 1$, so the equation reduces to

$$T_{\bar{0}}(\psi) = \int_0^T 1dt = T - 0 = T$$

That is, if the true value of ψ is zero, then the counterfactual survival time if unexposed is the same as the observed survival time, regardless of exposure. In other words, exposure has no effect.

Likewise, for workers who were never exposed, $A(t) = 0$, and the calculation will work out the same way. No matter what the true value of ψ is, workers who were never exposed have a counterfactual survival time if unexposed that is equal to their observed survival time.

The equation holds true if we replace the observed quantities by what would happen under a counterfactual exposure scenario, too.(1) Counterfactual exposure history is denoted $\bar{a}(t)$, and the counterfactual survival time is $T_{\bar{a}}$:

$$T_{\bar{0}}(\psi) = \int_0^{T_{\bar{a}}} \exp[\psi \times a(t)]\, dt$$

If workers were exposed every year at an intensity of one unit, then the counterfactual exposure $a(t) = 1$ for all t and we denote the counterfactual always-exposed survival time by $T_{\bar{1}}$. In that case,

$$T_{\bar{0}}(\psi) = \int_0^{T_{\bar{a}}} \exp[\psi \times a(t)]\, dt$$

$$= \int_0^{T_{\bar{1}}} \exp[\psi \times 1]\, dt$$

$$= T_{\bar{1}} \exp[\psi].$$

Thus, solving for the parameter, we find that $\psi = -\ln\left[\frac{T_{\bar{1}}}{T_{\bar{0}}}\right]$: the parameter equals the negative log of the ratio of (median) survival times comparing everyone always exposed at an intensity of one unit to everyone never exposed.

In our application, time is measured in years (other than the possible fractions of years at the very beginning and very end of a worker's time on follow-up), so this integral is actually a discrete sum. In each year, a worker could be unexposed (in which case that year contributes $\exp[0] = 1$ to the sum) or exposed (in which case that year contributes $\exp[\psi \times A(t)]$ to the sum). If exposure is harmful, then being unexposed results in longer survival time (*i.e.*, $T_{\bar{0}}(\psi) > T$ for workers who were ever exposed), which is equivalent to $\psi > 0$. If being exposed is beneficial (as in the case of medications), then $T_{\bar{0}}(\psi) < T$ for workers who were ever exposed (i.e. they were observed to live longer than they would have lived if they had not been exposed), which implies that $\exp[\psi \times A(t)] \leq 1$, so $\psi < 0$.

Handling of administrative censoring

In practice, not all workers died during follow-up; those who survived to the administrative end of follow-up do not have an observed survival time, so we could not calculate a candidate counterfactual unexposed survival time. Furthermore, such survival may depend on exposure, so simply excluding those workers still alive from the analysis while including all those who died during follow-up would cause bias.(1) In order to adjust for this, we did not use $T_{\bar{0}}(\psi)$ itself directly in the exposure model. Instead, we used the following function of $T_{\bar{0}}(\psi)$:

$$\Delta(T_{\bar{0}}(\psi), C) = \begin{cases} \min[T_{\bar{0}}(\psi), C]\text{, if } \psi \geq 0 \\ \min[T_{\bar{0}}(\psi), C \exp(\psi \times A)]\text{, if } \psi < 0 \end{cases}$$

(where C represents the time from that worker's entry into the cohort to the administrative end of follow-up [*i.e.*, their maximum observable survival time] and A denotes the maximum observed exposure among all of the person-years included).(2) That is, if the candidate value for ψ was greater than 0, corresponding to exposure being harmful, we used $T_{\bar{0}}(\psi)$ only if a worker's outcome under no exposure would have occurred before the administrative end of follow-up, and otherwise used the total time the worker would have been observed during follow-up if unexposed. For further details, see Robins et al.(2)

This method of adjustment for administrative censoring results in a locally smooth estimating function, and the local slope is used to estimate confidence intervals.(3, 4)

Sensitivity analyses

Sensitivity analyses were conducted using binary exposure metrics corresponding to indicators for exposures exceeding certain cutoffs: 0.117mg/m^3 (the median observed annual average daily exposure), 0.05mg/m^3 (the OSHA standard), 0.025mg/m^3 (the American Conference of Governmental and Industrial Hygienists recommended limit), and 0mg/m^3 (a ban). The analysis

using a binary exposure metric indicating whether a worker was ever exposed at all in that year (*i.e.*, a cutoff of 0mg/m^3) theoretically asks the same question asked in the main analysis: what would have happened if no one were ever exposed to silica? However, the assumptions made in specifying the models in the two analyses are different. In the main analysis, exposure intensity is assumed (if held constant over time) to have a log-linear effect on survival time; in the sensitivity analysis, exposure intensity is assumed to have no importance whatsoever, with exposure duration having a log-linear effect on survival time.

For NMRD, we adjusted for censoring by lung cancer death, since lung cancer death appeared to be an important censoring event in the main analyses.

The estimating function for the g-estimation procedure used with the binary exposure metrics was not very well-behaved, so that variance estimation could only be done by running bootstraps. Due to the comparatively small numbers of cases for cause-specific mortality, many of the bootstrap samples did not yield estimates, so we were unable to estimate confidence intervals for lung cancer and NMRD. However, we were mainly interested in the point estimates from the sensitivity analyses.

Assumptions required for causal inference

Consistency requires that exposure correspond to a well-defined intervention.(5) For example, silica exposure could hypothetically be eliminated by requiring the full-time use of a perfect respirator (imagining that such a thing exists). We must assume that the effects of this intervention and of any other intervention achieving the exposure ban would be identical.

Causal inference from any estimate based on observational data depends on the assumption that we have measured all confounders (*conditional exchangeability*); g-estimation leverages

this assumption. The idea is that within strata of confounders, the exposure can be thought of as "randomized"—so differences in outcomes can be attributed to the exposure rather than to unmeasured differences between exposure groups.

As in most analyses, we also assume that we *specified the correct model forms* for both the exposure and the structural model. Our choice of exposure model in the main analysis means that we assume the counterfactual survival times are independent of the observed category of exposure (in quartiles). We used quartiles to avoid issues with sparse data; however, note that the exposure variable used in the main structural model for survival time was a continuous measure of annual average daily exposure. In the sensitivity analyses, we used a logistic model to predict the binary exposure.

The accelerated failure time model we used also makes several assumptions. One is that there is no effect modification. Another is that the people who died of the cause of interest would have died of that cause regardless of their exposure; only the timing might change. In addition, the model assumes that exposure has a log-linear effect on survival time, rather than on probability of event. (In the sensitivity analyses, the model assumes that intensity of exposure within the categories of above and below the cutoff does not affect the outcome but that duration of exposure above the cutoff has a log-linear effect on survival times.) If these assumptions are violated, then causal interpretation will be incorrect.

Causal interpretation of results from this method does not depend on experimental treatment assignment (also referred to as positivity). We therefore do not require all strata of covariates to contain all possible exposures. This is one advantage of g-estimation in occupational mortality studies, since workers who are no longer actively employed cannot be exposed. The

model predicting exposure is conducted on the “actively employed” subset of the person-time. We implicitly extrapolate that exposure after employment termination would have the same etiologic effect it has during active employment.

References

1. Hernan MA, Cole SR, Margolick J, Cohen M, Robins JM. Structural accelerated failure time models for survival analysis in studies with time-varying treatments. Pharmacoepidemiol Drug Saf. 2005;14(7):477-91.
2. Robins JM, Blevins D, Ritter G, Wulfsohn M. G-estimation of the effect of prophylaxis therapy for Pneumocystis carinii pneumonia on the survival of AIDS patients. Epidemiology. 1992;3(4):319-36.
3. Mark SD, Robins JM. Estimating the causal effect of smoking cessation in the presence of confounding factors using a rank preserving structural failure time model. Stat Med. 1993;12(17):1605-28.
4. Naimi AI, Cole SR, Hudgens MG, Richardson DB. Estimating the effect of cumulative occupational asbestos exposure on time to lung cancer mortality: using structural nested failure-time models to account for healthy-worker survivor bias. Epidemiology. 2014;25(2):246-54.
5. Hernan MA, Taubman SL. Does obesity shorten life? The importance of well-defined interventions to answer causal questions. Int J Obes (Lond). 2008;32 Suppl 3:S8-14.