Risk Analysis, Vol. 25, No. 3, 2005 DOI: 10.1111/j.1539-6924.2005.00606.x

Empirical Bayes and Adjusted Estimates Approach to Estimating the Relation of Mortality to Exposure of PM_{10}

Alain Le Tertre,[1]* Joel Schwartz,[2] and Giota Touloumi[3]

In the framework of the APHEIS program (Air Pollution and Health: A European Information System), a health impact assessment of air pollution in 26 European cities was performed for particles of an aerodynamic diameter less than or equal to 10 μm (PM_{10}). For short-term effects, it was based on overall estimates from the APHEA-2 project (Air Pollution and Health: A European Approach). These city-specific risk assessments require city-specific concentration-response functions, raising the question of which concentration-response is most appropriate. Estimates from city-specific models are more specific, but have greater uncertainty than those provided from multicity analyses. We compared several estimates derived from the city-specific analyses in cities that were part of the APHEA-2 project, as well as in a city that was not included in APHEA-2 but was part of the APHEIS project. These estimates were: the estimates from a local regression model, the adjusted estimates based on two significant effect modifiers identified through meta-regression models, and the city-specific empirical Bayes (shrunken) estimates and their underlying distribution. The shrunken and adjusted estimates were used to improve the estimation of city-specific concentration-response function. From these different estimates, attributable numbers of deaths per year were calculated. The advantages and limits of the different approaches are discussed through real data and in a simulation study.

KEY WORDS: Empirical Bayes; hierarchical models; mortality; particulate matter; risk assessment

1. INTRODUCTION

In the last decade, a number of epidemiological studies have shown that ambient air pollution adversely affects human health even at levels lower than current national standards.(1) Particulate matter is the pollutant that has most consistently been associated with short-term effects on mortality. Risk analyses based on these results have already been published.(2–5) Typically, an overall estimate (based on a meta-analysis) is used for all cities, assuming that the concentration-response relationship is the same everywhere. Many studies are using an overall estimate derived either from multicity studies, such as APHEA-1(1) or APHEA-2(6) and NMMAPS (National Mortality and Morbidity Air Pollution Study),(7) or from the literature. Overall estimates are computed using a random-effect approach that takes into account the heterogeneity of the effects among cities/studies. However, significant heterogeneity, if present, suggests that the use of a single estimate in all cities is not appropriate. In many cases risk assessment in a particular city, where an original study has been conducted, is based on the city-specific estimate rather than on the overall estimate.

In general, it is naïve to assume that the city-specific effect estimate is better than the overall

1 InVS, Environmental Health Unit, Saint-Maurice Cedex, France.
2 Harvard School of Public Health, Department of Environmental Health, Boston, MA, USA.
3 University of Athens Medical School, Department of Hygiene and Epidemiology, Athens, Greece.
* Address correspondence to Alain Le Tertre, InVS, Environmental Health Unit, 94415 Saint-Maurice Cedex, France; a.letertre@invs.sante.fr.

estimate merely because it is derived from the city of interest. In the case where there is no true heterogeneity, variations in the city-specific effect estimates about the overall mean are purely stochastic. In that case, the overall effect estimate is clearly superior.

In the presence of heterogeneity, city-specific estimates vary around the overall effect estimate for two reasons: (1) due to true heterogeneity in the estimates, and (2) due to additional stochastic error. A city-specific estimate that reflects the first source of variation, but not the second, is preferable. This is obtainable by using a shrunken empirical Bayes estimate.

True variation in the concentration-response estimates among cities presumably reflects differences (e.g., sources of particles, ventilation characteristics of housing, health of population) that are, in principle, identifiable. Meta-regression is an approach that seeks to identify the sources of such heterogeneity. An alternative city-specific estimate is obtained by using the slope predicted for that city through a meta-regression on effect modifiers. This approach has the added advantage that it can be applied to risk assessment in cities that were not part of the original studies, and therefore do not have city-specific estimates. The APHEA-2 study found that annual level of NO_2 and annual temperature mean at the city level both act as modifiers of the effect of PM_{10} on all-cause mortality, explaining a substantial part of the observed between-cities heterogeneity in Europe.[6] A geographical (East-West) difference was also identified in the APHEA-1 study.[1] Although geographical area could act as a surrogate for differences in health or environmental indices, such a finding could imply a bimodal underlying distribution of risk.

In the framework of the APHEIS program (Air Pollution and Health: A European Information System—www.apheis.net)[3,4] a health impact assessment of air pollution in 26 European cities was performed. This raised the question of whether it was preferable to use risk coefficients from city-specific models, from empirical Bayes estimates for each city, or from meta-regression estimates for each city.

In this article we will demonstrate and discuss the impact of implementing different approaches to estimating the short-term air pollution effect on the attributable number of deaths in 21 European cities that were part of the APHEA-2 project and one that was not, but was part of the APHEIS.[4] Robustness of the various estimates to model assumption violations has been evaluated in a simulation study.

2. METHODS

In multicity studies data analysis is generally implemented in two stages. In the first stage, data from each city are analyzed separately (to allow more local covariate control) whereas in the second stage, evidence across cities is combined using meta-regression techniques. Briefly, daily counts of deaths from each city were assumed to come from a nonstationary, overdispersed Poisson process. The concentration-response function is then assumed to be exponential and to follow

$$\ln\left(E(y_t^c)\right) = \beta_p^c P_t^c + \sum_{i=1}^{m} S_i^c\left(Z_{it}^c\right),$$

with $E(y_t^c)$ denoting the mean daily counts of the relevant health outcome in city c on day t, P_t^c daily levels of PM_{10} in city c, β_p^c the corresponding parameter to be estimated, Z_{it}^c the independent covariates other than PM_{10}, and $S_i^c(Z_{it}^c)$ the smooth functions of these covariates. Nonpollutant covariates include long-term time trend, season, calendar events, influenza, and meteorological variables. The smooth functions capture the nonlinear relationships with time-varying covariates and calendar time. Penalized regression splines were used as smoothing functions, as implemented by Wood[8] in the R Development Core Team.[9]

City-specific PM_{10} estimates were combined using the method developed by Berkey *et al.*[10] In summary, a random-effects model regressing the estimates from each city against potential effect modifiers was performed. The model assumptions are

$$\hat{\beta}_p^c \sim N\left(\mu_p^c,\ \sigma_{W,c}^2\right),$$
$$\mu_p^c \sim N\left(X^c\alpha,\ \sigma_B^2\right),$$

and hence

$$\hat{\beta}_p^c \sim N\left(X^c\alpha,\ \sigma_{W,c}^2 + \sigma_B^2\right),$$

where $\hat{\beta}_p^c$ is the estimated PM_{10} effect estimate in city c and $\sigma_{W,c}^2$ and σ_B^2 are the within-city c and the between-cities variances, respectively. It should be noted that $\sigma_{W,c}^2$ is estimated in the first-stage analysis. The between-cities variance, σ_B^2, was iteratively estimated using the maximum likelihood method described in Berkey *et al.*[11] For details about APHEA-2 first-stage and second-stage analysis, refer to Touloumi *et al.*[12]

An alternative to the city-specific estimates and to the overall (pooled random effects) estimate is the use of the city-specific shrunken estimates. These were derived following Longford.[13] Let $\bar{\beta}$ be the overall pooled cities estimate, without regressing on any effect modifier ($X^c\alpha = \bar{\beta}$ in this case). The shrinkage estimator is the conditional expectation of the city mean, μ_p^c, given the observed mean $\hat{\beta}_p^c$. The variance matrix of $(\mu_p^c, \hat{\beta}_p^c)$ is

$$\begin{pmatrix} \sigma_B^2 & \sigma_B^2 \\ \sigma_B^2 & \sigma_B^2 + \sigma_{W,c}^2 \end{pmatrix}.$$

Therefore,

$$\begin{aligned} \hat{\beta}_{p,Shr}^c &= E\big(\mu_p^c \mid \hat{\beta}_p^c; \bar{\beta}, \sigma_B^2, \sigma_{W,c}^2\big) \\ &= \bar{\beta} + \frac{\sigma_B^2}{\sigma_B^2 + \sigma_{W,c}^2}\big(\hat{\beta}_p^c - \bar{\beta}\big) \end{aligned} \quad (1)$$

with variance

$$\begin{aligned} \mathrm{Var}\big(\big(\mu_p^c \mid \hat{\beta}_{\mathrm{p}}^{\mathrm{c}}; \bar{\beta}, \sigma_B^2, \sigma_{W,c}^2\big) &= \sigma_B^2\left(1 - \frac{\sigma_B^2}{\sigma_B^2 + \sigma_{W,c}^2}\right) \\ &= \frac{1}{\sigma_B^{-2} + \sigma_{W,c}^{-2}}. \end{aligned} \quad (2)$$

Shrunken estimates, known also as empirical Bayes estimates, could be considered as posterior probability distributions as they include information from the overall and the city-specific estimates.

Risk analyses are often required in cities where no air pollution study has been carried out. To obtain an estimate in this case, we followed the approach proposed by Post *et al.*,[14] and estimated the underlying distribution of βs as an equally weighted mixture of the shrunken estimates derived in the previous step. That is, $f(\bar{\beta}_{p,\mathrm{mix}}) = \frac{1}{k}\sum_{i=1}^{k} f_c \hat{\beta}_{p,Shr}^c$, where k is the number of cities and f_c denotes the distribution of the empirical Bayes estimate at city c. The mean of this mixture distribution was used as an estimate for the new city.

Finally, we used a meta-regression-based estimate derived from the APHEA-2 study. The two effect modifiers found to explain a substantial part of the observed heterogeneity were annual mean of NO_2 and annual mean temperature. Therefore, for each city we predicted the coefficient for PM_{10} based on the model:

$$\begin{aligned} \hat{\beta}_{p,\mathrm{NO_2-Temp}}^c &= E\big(\hat{\beta}_p^c \mid \mathrm{NO_2}, \mathrm{Temp}\big) \\ &= \bar{\beta}' + \beta_1'\mathrm{NO_2} + \beta_2'\mathrm{Temp}. \end{aligned} \quad (3)$$

Regression coefficients and associated variance-covariance matrix were provided by the APHEA-2 project.[15] This approach can also accommodate a new city.

We used the above-described effect estimates (observed city-specific $\hat{\beta}_p^c$, shrunken city-specific $\hat{\beta}_{p,Shr}^c$, pooled $\bar{\beta}$, mean of estimated mixture distribution $\bar{\beta}_{p,\mathrm{mix}}^c$, and adjusted-for-effect modifiers $\hat{\beta}_{p,\mathrm{NO_2-Temp}}^c$) to calculate the attributable number of deaths in each city. The relative risk (RR) is expressed in our analysis as $RR_{\Delta x} = \exp(\beta \times \Delta x)$ where Δx represents a change by x μg/m^3 in the daily PM_{10} levels. To estimate the attributable number of deaths we need to define a baseline exposure. Let $\bar{Y}$ be the annual mean of daily mortality, which reflects the impact of mean daily PM_{10} levels, $\bar{x}$. The baseline mortality incidence Y_F at the baseline PM_{10} level $x0$ can then be estimated as

$$Y_F = \bar{Y} \times \left(1 - \frac{RR_{\Delta(x0-\bar{x})} - 1}{RR_{\Delta(x0-\bar{x})}}\right).$$

The attributable number of deaths when the PM_{10} levels increase from $x0$ to $x1$ is

$$AR = Y_F \times \big(RR_{\Delta(x1-x0)} - 1\big).$$

We have set the baseline level to 10 μg/m^3 in this study.

3. RESULTS

Table I shows the city-specific estimated regression coefficients and their standard errors for the effect of each 1 μg/m^3 increase in PM_{10} levels on total mortality, as reported in the APHEA-2 project.[15] It also shows the annual mean of daily temperature and NO_2 levels, the two modifiers in the meta-regression. The coefficients ranged from −0.000465 in Erfurt to 0.001554 in Lyon. The annual temperature exhibited a range of around 14°C with the minimum value observed in Helsinki and the maximum in Tel Aviv. Stockholm was the cleanest city in terms of NO_2 levels with an annual mean of 26 μg/m^3 whereas Milan showed the highest levels with annual mean of 94 μg/m^3.

There was significant between-cities heterogeneity in the PM_{10} effect estimates and therefore the pooled effect was calculated based on a random-effects model.[10] The overall estimate was 0.00055 (SE = 0.000098). We then calculated the shrunken estimators in each location following Equations (1) and (2). Fig. 1 shows the estimated density for each of the shrunken estimators (i.e., in each city). The estimated

Table I. Local Estimates for Total Mortality PM_{10} Regression Coefficients with Their Associated Standard Errors and Annual Mean of Daily Temperature and NO_2 Levels in Cities Participating in the APHEA-2 Study

City	PM_{10} Coefficient	SE	Annual Temperature	Annual NO_2
Athens	0.001311	0.0003	17.8	74.0
Barcelona	0.000575	0.0002	16.4	68.6
Basel	0.000462	0.0005	10.7	38.2
Birmingham	0.000305	0.0003	9.6	45.9
Budapest	−0.000248	0.0005	10.5	76.3
Cracow	0.000155	0.0004	8.3	43.5
Erfurt	−0.000465	0.0004	8.8	39.5
Geneva	−0.000059	0.0005	9.5	44.9
Helsinki	0.000389	0.0004	6.1	32.6
London	0.000591	0.0002	11.8	60.7
Lyon	0.001554	0.0005	12.4	63.0
Madrid	0.000372	0.0003	14.5	70.0
Milan	0.000901	0.0002	13.7	93.5
Paris	0.000411	0.0003	12.0	52.8
Prague	0.000097	0.0002	9.9	57.5
Rome	0.001333	0.0003	16.7	87.7
Stockholm	0.000479	0.0009	7.5	25.7
Tel Aviv	0.000522	0.0003	20.4	69.7
Teplice	0.000876	0.0004	8.8	32.4
Torino	0.000938	0.0002	14.3	75.9
Zurich	0.000365	0.0004	10.9	40.1
Toulouse	NA	NA	13.8	30.0

distribution of the pooled estimate is superimposed (i.e., overall), based on the random-effects model, and the estimated mixture distribution of the empirical Bayes estimates across all the cities. Substantial departures from the population mean (overall) estimate can be seen in several cities. The underlying distribution of the empirical Bayes estimates displays the same mean as the pooled estimate, but it is more flat, reflecting the heterogeneity between cities. Consequently, the corresponding 95% credible interval for the relative risk for the total mortality associated with a 10 μg/m^3 increase in PM_{10} (0.994, 1.014) is larger than that derived from the pooled estimate (1.002 to 1.006). We also applied Equation (3) to calculate the city-specific estimate based on the meta-regression.

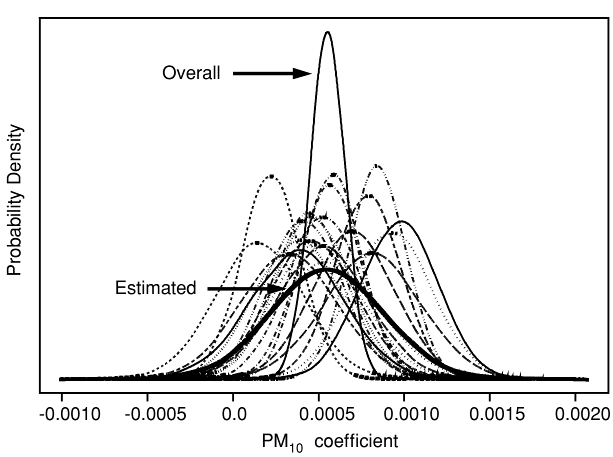

Fig. 1. Probability densities of PM_{10} shrunken coefficients for mortality in each of the 21 cities and the resulting estimated mixture distribution from all cities.The probability density of the pooled overall cities (using random-effects model) coefficient is also shown.

Fig. 2 shows the city-specific relative risks and its 95% confidence interval of mortality for a 10-μg/m^3 increase in daily PM_{10} levels as estimated using the shrunken estimates. The local estimates of the *RR*s as well as the estimates predicted from the meta-regression are also shown. For comparison, the overall estimate provided by the RE model is also shown. As expected, the shrunken estimates are located between the local and the overall estimates so that they can be considered as a weighted mean of these two estimates.

Except in four cities (Budapest, Stockholm, Tel Aviv, and Teplice) the adjusted estimates are close to the shrunken estimates. This concordance is due to the fact that temperature and NO_2 annual means tend to explain the observed heterogeneity between the local estimates.

In Table II, the number of deaths attributable to PM_{10} exposure calculated using the various estimates for the mortality-PM_{10} RR are shown. A reduction in the annual PM_{10} mean from the observed value to 10 μg/m^3 was assumed in each city. There is substantial variability in the attributable numbers of deaths at the city level depending on the choice of the *RR* estimate. The calculated number of deaths using the shrunken estimates and those adjusted for temperature and NO_2 ranged from, respectively, 1/6 and 1/15 in Erfurt and in Stockholm to more than double in Rome the one using the overall *RR* estimate. The effect of the choice-of-risk coefficient on overall risk in all the cities is much less. When we sum the attributable risks across all locations, the expected numbers of deaths using the shrunken estimate or the meta-regression estimate are, respectively, 25% less or 11% larger than those estimated using the overall pooled estimate in each city.

4. SIMULATION STUDY

To evaluate the robustness of the overall and of the shrunken estimates when the underlying distribution of the city-specific estimates is bimodal, we

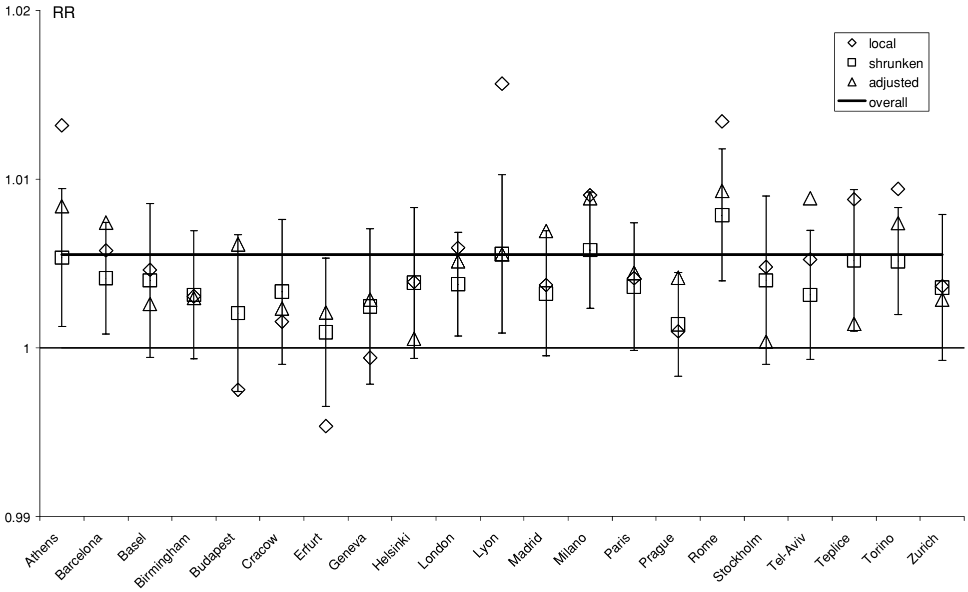

Fig. 2. City-specific shrunken estimates (95% CI) of relative risk for mortality per 10 $\mu g/m^3$ increase in PM_{10} levels. Local and adjusted for effect modifiers (NO_2 and temperature) estimates are also shown. The overall effect as estimated by the random-effects model is also shown.

Table II. Attributable Number (95% CI) of Deaths for a Reduction of the Annual Mean PM_{10} from the Observed Value to 10 $\mu g/m^3$ Calculated Under Various Estimates for Corresponding Relative Risk

	Local Estimates			Shrunken Estimates			Adjusted for Effect Modifiers Estimates			Random-Effects Pooled Estimate		
City	N	CI−	CI+	N	CI−	CI+	N	CI−	CI+	N	CI−	CI+
Athens	1119	559	1691	461	108	819	720	517	924	476	309	643
Barcelona	443	118	774	319	62	579	567	414	723	424	275	574
Basel	32	−32	98	28	−4	59	18	2	35	38	25	51
Birmingham	98	−81	280	101	−21	224	95	34	157	177	115	239
Budapest	−226	−1134	712	185	−233	609	550	301	800	494	322	668
Cracow	50	−182	290	106	−30	245	74	4	145	175	114	237
Erfurt	−48	−128	35	9	−35	55	21	−1	44	55	36	75
Geneva	−3	−49	45	11	−10	33	13	4	22	26	17	34
Helsinki	24	−29	77	23	−4	51	3	−15	22	34	22	45
London	682	236	1132	436	81	793	591	437	745	635	414	858
Lyon	158	49	272	57	9	107	57	44	71	57	37	77
Madrid	229	−95	558	198	−30	428	423	336	510	338	220	457
Milan	403	201	609	260	105	418	395	251	541	248	161	336
Paris	312	−100	727	274	−11	561	336	220	452	417	272	563
Prague	82	−248	419	117	−140	379	348	198	499	458	297	621
Rome	1286	706	1882	766	382	1156	904	682	1129	541	351	732
Stockholm	29	−76	136	24	−6	55	2	−17	21	34	22	45
Tel Aviv	205	−7	423	124	−27	277	345	198	493	216	141	293
Teplice	214	45	387	127	24	231	35	−31	102	135	88	183
Torino	450	272	633	249	94	408	356	278	435	268	174	363
Zurich	51	−55	159	50	−10	111	40	8	72	77	50	104
Total	5592	−29	11338	3926	304	7599	5893	3863	7941	5324	3462	7199

ran a simulation study. City-specific coefficients were generated by a bimodal normal distribution with the means, for each mode respectively, set equal to 0.5 and 1.5 times the overall estimate initially found, and with equal variances, defined as half of the variance of the overall estimate. Twenty coefficients, representing 20 cities, were thus derived from the bimodal distribution. To reflect the random variation for each one coefficient we assigned a standard error, generated by a gamma distribution with a shape equal to 1 and a rate equal to the standard error of the pooled estimate. One thousand replications for each of the 20 coefficients were drawn from a normal distribution with mean equal to the coefficient in city c and standard error derived from the gamma distribution. We label these estimates "original estimates," and they represent the case where we use for each city the results of a regression in that city. For each set of 20 coefficients, the random-effects pooled estimate as well as the city-specific shrunken estimates were calculated, and the percentage of bias was estimated as [(true-estimated)/true] $\times$ 100 for both "original" and shrunken estimates.

Fig. 3 shows the distribution of the true estimates as well as the generated "original" and shrunken estimates. Although there is a slight shift toward the overall mean in the distribution of the shrunken estimates, the overall pattern (i.e., bimodal distribution) is quite well reproduced. The overall estimate (not shown) was, as expected, purely unimodal. The average bias was mostly 0% for the "original" and −3% for the shrunken estimates, but the range of the bias was, respectively, equal to (−510; 571) and (−248; 195). Hence the range of potential bias is reduced when using the shrunken estimates. For the overall estimate, the average bias was −30% and the range was (−125; 42). If we increase the standard deviation in the gamma distribution by a factor of 10, then the "original" and shrunken estimates' distributions are no longer bimodal, and the range of the bias increases (−449; 367) for the shrunken estimates, but explodes (−5146; 5667) for the original estimates. The range for the overall estimates stays quite constant (−194; 68).

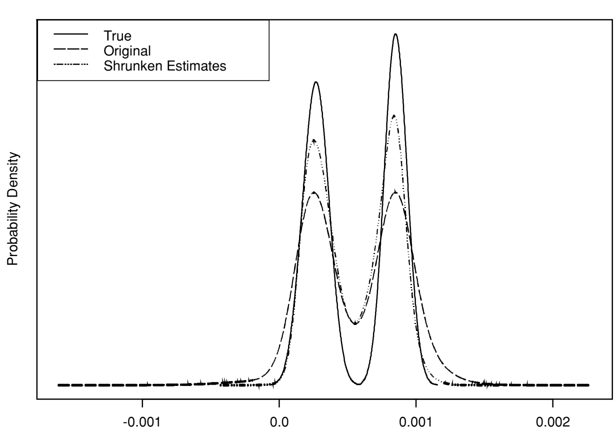

Fig. 3. True, original, and shrunken estimates density from 1,000 simulations.

5. DISCUSSION

In this study we have shown that although the overall sum of the deaths attributable to PM_{10} in 21 European cities is not strongly influenced by the method used to estimate *RR*s, this is not true at the city level. Applying a shrunken estimate in Rome or in Erfurt would lead to almost 80% more deaths or 400% fewer deaths, respectively, than those calculated with the overall estimate. The heterogeneity observed in these cities does not favor applying of a single estimate. Neither does it militate for applying the city-specific estimate, as that estimation would also lead to over- or underestimating the shrunken estimates by 176% and 613%, respectively.

We also applied the overall, the meta-analytic, and the mean of the estimated underlying distribution coefficients to a city, Toulouse in France, not part of the APHEA-2 study, but part of the APHEIS project. The overall or estimated underlying estimates gave 25 deaths per year compared to 13 for the adjusted approach. More interesting was the difference in the 95% confidence interval around these estimates: while the overall estimates show an extremely narrow confidence interval (16, 33), the estimated estimate showed a much larger interval (−26, 62), as did the adjusted one (−5, 31). This indicates that excessive certainty may be suggested by naïve approaches to risk assessment.

The shrunken estimates approach has already been explored and applied in the case of air pollution.[14] The shrunken estimates have the nice property that they derive the estimate at the local level by combining information from the city-specific estimate and the overall estimate. They also reduce the variability of the local estimate by incorporating information from other cities. A key disadvantage of such an estimate is that it can be applied only in cities that are part of the initial analysis.

The adjusted estimate also provides a more local estimate as it takes into account potential effect modifiers. It also reduces uncertainties around the estimate. It is more widely applicable as one just needs to have information on these two effect modifiers to calculate it. The two particular effect modifiers (NO_2 and temperature) that have been identified so far for the PM_{10} mortality relationship should be seen as surrogates for different patterns of air pollution or exposure of the population but they could also be just the best set of covariates from a statistical point of view, explaining the heterogeneity. The effect of temperature could be a surrogate of the ventilation rate between cities. In that case we could not apply annual temperature as an effect modifier on a city with a high rate of air-conditioned houses.[16] This is not common in Europe, but quite common in warmer climates in the United States. This hypothesis, although a plausible one, is only one of several, equally plausible, hypotheses. The use of the adjusted estimate, thus, should be limited to cities presenting similar characteristics to those initially observed. The discrepancy between the shrunken and the adjusted estimates, found in some cities as Stockholm, highlights the limit of such indicators, useful to explore the reasons of the heterogeneity, but not to explain them.

The use of the local estimate is subject to too much noise to be reliable. The two derived city-specific estimates could be alternatively used depending on the data availability. For cities in the initial study, both give similar results. One may still prefer the shrunken estimate, as it does not make any inference on the relation with potential effect modifiers. For cities with available data on effect modifiers, the adjusted estimate has some nice properties but its use requires careful attention.

Applied to a single city, the overall estimate does not adequately reflect the heterogeneity present in the data. We have shown that this could be better taken into account by deriving an estimated underlying distribution that represents the dispersion observed between cities. Both of these techniques are an improvement in reducing the uncertainties surrounding pollutant coefficient, but are still affected by the uncertainties around the initial estimates of these coefficients.

Each of the methods described in this work depends heavily on the assumption that city-specific models from which city-specific estimates were derived have precluded at least the known sources of bias, such as confounding, model misspecification, etc. The city-specific estimates utilized in our work were derived from the APHEA-2 project, which is a well-described European multicity study that used a detailed prespecified protocol for both data collection and data analysis. Potential bias, although as always possible, has been discussed extensively within the project and several actions have been taken to correct it.

Given that unbiased results are obtained in the first-stage analysis, we have presented results from various alternative risk estimates. Whereas shrunken estimates are, in principle, preferable, at least in terms of efficiency, for the local estimates, questions about the effect of departures from unimodal normal distribution of the estimates on the shrunken ones have been raised. Based on our limited simulation study, we have shown that shrunken estimates are preferable to the local estimates even if data are derived from a bimodal distribution. In general, shrunken estimates (and local estimates) can reproduce the underlying distribution, provided that there is enough information in the data. However, when the noise in the city-specific estimates is largely increased, then data do not give enough insight into the true underlying distribution in any case. Even in such situations, shrunken estimates will reduce the margin of error compared to original estimates.

Given the limitations of each of the different estimates outlined above, we recommend the use of the shrunken estimate in cities for which this option is available. For other cities, which estimate is the most appropriate is less straightforward since the adjusted estimate reflects a local situation but requires strong assumptions about the surrogacy of the effect modifiers, whereas the estimated distribution of the estimates reflects a general situation with greater uncertainties. In such cases, applying both coefficients would give a reasonable range of estimates.

REFERENCES

1. Katsouyanni, K., Touloumi, G., Spix, C., Schwartz, J., Balducci, F., Medina, S., Rossi, G., Wojtyniak, B., Sunyer, J., Bacharova, L., Schouten, J. P., Ponka, A., & Anderson, H. R. (1997). Short-term effects of ambient sulphur dioxide and particulate matter on mortality in 12 European cities: Results from time series data from the APHEA project. Air Pollution and Health: A European Approach. *British Medical Journal*, *314*(7095), 1658–1663.
2. Leksell, I., & Rabl, A. (2001). Air pollution and mortality: Quantification and valuation of years of life lost. *Risk Analysis*, *21*(5), 843–857.
3. Medina, S., Plasència, A., Mücke, H.-G., & Contributing Members of the APHEIS Group Health Impact Assessment

of Air Pollution in 26 European Cities—First Results of the APHEIS Study. (2002). *WHO Collaborating Centre for Air Quality Management and Air Pollution Control at the Federal Environmental Agency*. Berlin: Newsletter (30).
4. Medina, S., Plasencia, A., Ballester, F., Mucke, H. G., & Schwartz, J. (2004). Apheis: Public health impact of PM10 in 19 European cities. *Journal of Epidemiology and Community Health*, *58*(10), 831–836.
5. Quenel, P., Zmirou, D., Le Tertre, A., Balducci, F., Medina, S., Barumandzadeh, T., Le Moullec, Y., Ritter, P., Festy, B., & Dab, W. (1995). Impact de la pollution atmosphérique urbaine de type acido-particulaire sur la mortalité quotidienne à Lyon et dans l'agglomération parisienne. *Sante Publique*, *6*(4), 1–14.
6. Katsouyanni, K., Touloumi, G., Samoli, E., Gryparis, A., Le Tertre, A., Monopolis, Y., Rossi, G., Zmirou, D., Ballester, F., Boumghar, A., Anderson, H. R., Wojtyniak, B., Paldy, A., Braunstein, R., Pekkanen, J., Schindler, C., & Schwartz, J. (2001). Confounding and effect modification in the short-term effects of ambient particles on total mortality: Results from 29 European cities within the APHEA2 project. *Epidemiology*, *12*(5), 521–531.
7. Samet, J. M., Dominici, F., Curriero, F. C., Coursac, I., & Zeger, S. L. (2000). Fine particulate air pollution and mortality in 20 U.S. cities, 1987–1994. *New England Journal of Medicine*, *343*(24), 1742–1749.
8. Wood, S. N. (2003). Thin-plate regression splines. *Journal of the Royal Statistical Society [Series B]*, *65*, 95–114.
9. R Development Core Team. (2004). R: *A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation For Statistical Computing.
10. Berkey, C. S., Hoaglin, D. C., Antczak-Bouckoms, A., Mosteller, F., & Colditz, G. A. (1998). Meta-analysis of multiple outcomes by regression with random effects. *Statistics in Medicine*, *17*(22), 2537–2550.
11. Berkey, C. S., Hoaglin, D. C., Mosteller, F., & Colditz, G. A. (1995). A random-effects regression model for meta-analysis. *Statistics in Medicine*, *14*(4), 395–411.
12. Touloumi, G., Atkinson, R., & Le Tertre, A. (2004). Analysis of health outcome time series data in epidemiological studies. *Environmetrics*, *15*, 101–117.
13. Longford, N. T. (1993). *Random Coefficient Models* (p. 18). Oxford: Oxford University Press.
14. Post, E., Hoaglin, D., Deck, L., & Larntz, K. (2001). An empirical Bayes approach to estimating the relation of mortality to exposure to particulate matter. *Risk Analysis*, *21*(5), 837–842.
15. Katsouyanni, K., Touloumi, G., Samoli, E., Petasakis, Y., Analitis, A., & Le Tertre, A. (2003). *Sensitivity Analysis of Various Models of Short-Term Effects of Ambient Particles on Total Mortality in 29 Cities in APHEA2* (pp. 157–164). Boston, MA: Health Effects Institute.
16. Janssen, N. A., Schwartz, J., Zanobetti, A., & Suh, H. H. (2002). Air conditioning and source-specific particles as modifiers of the effect of PM(10) on hospital admissions for heart and lung disease. *Environment and Health Perspectives*, *110*(1), 43–49.