APPENDIX AVAILABLE ON REQUEST

Research Report 155

The Impact of the Congestion Charging Scheme on Air Quality in London

Part 1. Emissions Modeling and Analysis of Air Pollution Measurements

Frank Kelly et al.

Appendix C. Investigation of the Cumulative Sum Technique to Analyze Air Quality Changes

Note: Appendices Available on the Web may appear in a different order than in the original Investigators' Report, and some remnants of their original names may appear in Table and Figure numbers. HEI has not changed the content of these documents, only the letter identifier.

Appendix C was originally Appendix H.

Correspondence may be addressed to Dr. Frank Kelly, Professor of Environmental Health, Environmental Research Group, MRC-HPA Centre for Environment & Health, School of Biomedical Sciences, King's College London, 150 Stamford Street, London SE1 9NH, U.K.
Tel ++44 20 7848 4004 Fax ++44 20 7848 3891 frank.kelly@kcl.ac.uk.

Although this document was produced with partial funding by the United States Environmental Protection Agency under Assistance Award CR–83234701 to the Health Effects Institute, it has not been subjected to the Agency's peer and administrative review and therefore may not necessarily reflect the views of the Agency, and no official endorsement by it should be inferred. The contents of this document also have not been reviewed by private party institutions, including those that support the Health Effects Institute; therefore, it may not reflect the views or policies of these parties, and no endorsement by them should be inferred.

This document was reviewed by the HEI Health Review Committee but did not undergo the HEI scientific editing and production process.

APPENDIX C.
INVESTIGATION OF THE CUMULATIVE SUM TECHNIQUE TO ANALYZE AIR QUALITY CHANGES

INTRODUCTION

The CUSUM statistical techniques, first proposed by Page (1954), were developed for use in industrial process control to detect deviations in production parameters from predetermined values. The statistical methods used in process quality control, including CUSUM techniques, were reviewed by Ryan (2000). This approach has been used to assess the performance of medical professionals, learning curves for trainees, and the efficacy of medical treatments, among others. In addition, Rossi and colleagues (1999) applied the CUSUM technique in an environmental epidemiologic setting where they compared mortality rates in two distinct time periods in an area of Tuscany, Italy, characterized by the presence of chemical plants.

Such a technique would appear to offer a useful approach to identifying changes in mean pollutant concentrations after a traffic management intervention, especially one aiming to detect a sustained shift in mean pollutant concentrations from the levels before implementation. As a first step in assessing its suitability for the CCS study, we used CUSUM to analyze changes in ambient mean air pollution levels after the introduction of a relatively simple traffic management scheme at Westminster—Marylebone Road in central London. After this, the CUSUM method was applied to measurements from indicator sites that showed exceptional changes in mean concentrations over the 4-year analysis period. The primary aim of the analysis was to identify whether the exceptional change could be related to a step change in concentrations on a specific date and whether this date could be related to implementation of the CCS. A secondary aim was to further test the applicability of the CUSUM technique in complex air quality scenarios.

METHODS

The CUSUM Procedure

CUSUM methods apply to observations recorded over time (daily, weekly, monthly). The observations may be physical measurements, counts, or rates and may be grouped (e.g., in production batches) or individual observations. The CUSUM methods of Lucas (1982) and Lucas and Crosier (1982), as applied to individual observations, first compute the standardized deviations of observations from the desired or reference mean for the process or time series:

$$z_i = \frac{x_i - x}{\hat{\sigma}_x} \qquad (1)$$

where $\boldsymbol{x_i}$ is the observed value at time $\boldsymbol{i}$, x is the desired process mean (or reference mean), and $\hat{\sigma}_x$ is an estimate of the standard deviation of the observed values. These are accumulated over time to compute the cumulative sum, S, at each time point $\boldsymbol{i}$ as follows:

$$S_i = S_{i-1} + \boldsymbol{z_i} \qquad (2)$$

where $S_0 = 0$.

Thus, if there is a shift in the process mean away from the target, then $\boldsymbol{z_i}$ will tend to be larger or smaller than the target average and the cumulative sum will steadily increase or decrease. Depending upon the magnitude of the shift in the mean, the CUSUM may not detect the change immediately and may require a number of observations at the new level before such a change is detectable.

If we take an example of detecting a shift in daily mean pollutant concentrations at a particular monitoring site, $\boldsymbol{x_i}$ would be the daily mean concentration on day $\boldsymbol{i}$, $\bar{x}$ would be the mean concentration over the period before any shift (the reference mean), and S_i would represent the cumulative deviation in standardized concentration on day $\boldsymbol{i}$ away from the reference mean concentration.

Lucas (1982) proposed to split S into two to help differentiate between mean increases (S_{Hi}) and mean decreases (S_{Li}) defined as:

$$S_{Hi} = \max[0, (\boldsymbol{z_i} - \boldsymbol{k}) + S_{Hi-1}] \text{ and}$$

$$S_{Li} = \min[0, (\boldsymbol{z_i} + \boldsymbol{k}) + S_{Li-1}] \,. \qquad (3)$$

The parameter $\boldsymbol{k}$ is the allowable 'slack' in the process and is usually set to be one half (in $\boldsymbol{z}$ units) of the mean shift one wishes to detect; in other words

$$k = \frac{1}{2}\Delta \text{ where: } \Delta = \frac{x - \bar{x}}{\hat{\sigma}_x} \,. \qquad (4)$$

The usual choice of $\boldsymbol{k} = 0.5$ is therefore the appropriate choice for detecting a 1-σ shift in the reference mean. Deviations from the reference mean less than this factor will be ignored as process noise. However, a balance has to be achieved between filtering out acceptable noise and reducing the sensitivity of the CUSUM output.

Although time-series data may be irregular or interrupted, the CUSUM technique was developed for applications in which input data are normally distributed with no serial correlation between data points; air quality data rarely conform to these assumptions because they are typically both skewed and correlated.

INVESTIGATIVE APPLICATION OF CUSUM TO AIR POLLUTION DATA

To test the application of the CUSUM technique to assess air quality changes, we applied it to detect a change in monitored concentrations of CO after the introduction of a bus lane on Westminster—Marylebone Road (at the CCZ boundary) on August 18, 2001. This area is a congested six-lane east–west trunk route in central London in which the nearside lane in each direction is a designated bus lane (only buses and taxis are allowed to enter the lane, which is strictly enforced using a system of cameras and automatic fines). The bus lane is permanently in operation 24 hours a day, 7 days a week.

We selected CO as an indicator of traffic-related pollution because it has relatively straightforward atmospheric chemistry, low background concentrations, and has a strong traffic-related signal at the curbside. Measurements of CO concentrations were available from a continuous monitoring site (Westminster—Marylebone Road) located on the southern side of the road with the CO sampling inlet approximately 1 m from the curb. In addition, traffic count data were obtained from an induction loop system fixed within the road surface adjacent to the monitoring site. The system monitors vehicle number, type (by axle length), and speed for each of the six lanes.

Quality Assurance/Quality Control

The monitoring site is part of the U.K. Government's automatic network and is operated to defined QA/QC standards. All data used in this study were fully ratified. The CO valid capture rate for the period of study was 96%, equating to 1,056 valid days (> 75% capture rate for 15-minute means per day) out of a possible 1,096. The measure used in this analysis was the daily mean concentration in milligrams per cubic meter.

Results

Figures C.1 and C.2 show time-series data between January 1, 2000, and December 31, 2002, for the daily total number of vehicles traveling past the vehicle monitors and the daily mean CO concentration measured at the curbside monitor. Figure C.1 clearly shows a fall in the number of vehicles using the nearside lane after the bus lane was introduced. Figure C.2 illustrates the decreasing daily concentrations of CO throughout the 3-year period as well as the considerable seasonality; concentrations increased in the winter months, which was a feature particularly evident in later months of 2000.

For the CUSUM calculations (Figure C.3), the reference mean ($\boldsymbol{x}$ = 2.13 mg/m^3) and standard deviation ($\hat{\sigma}_x$ = 0.92 mg/m^3) were calculated from CO data during the preimplementation period from January 1, 2000, to August 17, 2001.

To differentiate between any decreases in the mean CO concentrations that might be attributable to introducing the CCS from those occurring generally across London, we plotted time-centered running annual mean CO concentrations at monthly increments at five Inner London roadside sites (not including Westminster—Marylebone Road) that were not subject to local traffic management schemes during the period 1999 to 2003 (Figure C.4). The plot showed a decrease over the time period and seasonal variation contained within a linear trend. A simple regression model calculated this linear decrease as −0.368 μg/m^3 per day. As there was no reason to assume that background changes in CO levels did not also apply to Westminster—Marylebone Road, we adjusted the reference mean ($\boldsymbol{x}$) used in the computation of the CUSUMs to take account of this trend. Although the resulting CUSUM chart suggests that a decrease in concentrations occurred over the period, the strong seasonal nature of the CO data disrupts the analysis to a point where no clear change point can be identified (Figure C.5). Any stepped decrease in concentrations caused by the intervention was insufficiently large to override this natural seasonal variation.

The CUSUM procedure assumes the data are independent and normally distributed, whereas air pollution measurements tend to have a skewed distribution and a high degree of autocorrelation between daily mean values. The CO data did show evidence of skewness and strong autocorrelation at a lag time of 1 day. To assess the sensitivity of the results to this modest departure from normality, the analysis was repeated using log-transformed daily CO concentrations. There was little change in the CUSUM chart and therefore no change in the conclusions from the original analyses (data not shown). In order to assess the impact of autocorrelation, the data were condensed to weekly averages and the CUSUM analysis repeated. Although this batch analysis had the effect of producing a lower standard deviation and a smoother CUSUM trace, which clarified the results of the daily mean analysis, the disruption caused by seasonal variation remained (data not shown).

Discussion

In this initial case study we attempted to apply a statistical technique normally used in quality control processes to identify sustained change in ambient pollution levels caused by a traffic management intervention. In its basic form, it appears to be a simple method to identify subtle but sustained step changes in pollution levels, but the range of confounding influences on concentrations, most notably underlying trends and seasonality, complicate its interpretation and act to obscure the identification of change points. In the basic form presented here, the CUSUM procedure was unable to differentiate between these confounding influences and the target intervention.

The detection of small changes in pollutants therefore requires modification of the technique to allow adjustment for time-varying influences on pollution levels, most notably the temporal correlation inherent in air pollutant data. A number of established filtering techniques have been designed to remove such correlation in long-term time-series pollutant data (Rao and Zurbenko 1994; Porter et al. 2001). These spectral de-trending techniques rely on low, medium, and high frequency variations within the time series to isolate short, seasonal, and long-term trends in time-series data. However, none consider the effects of step changes within a time series of the type hypothesized in the study.

The purpose of the CUSUM technique is to identify a date when a change point occurs. Smoothing filters such as the Kolmolgorov-Zurbenko filter could not be applied across step changes because the regression would be disrupted and the change point would be blurred across the filter time span. Kuebler and colleagues (2001) used the concept of an average meteorologic year to remove seasonal and weekly variations in a pollution time series. However, this approach would not be suitable in a time series containing a step change. Filtering could be applied separately to the periods just before and after step changes but this would assume some prior knowledge of the date of the step change, thereby defeating the purpose of the CUSUM method.

In this application, batching input data into weekly mean concentrations was able to remove the effect of short-term correlation. A simple method for addressing independent long-term trends was used, but it relied on the assumption that the intervention would not affect concentrations recorded by surrounding monitoring sites, which may not be justified in larger-scale interventions such as the CCS. Methods for accounting for seasonal fluctuations and long-term trends while retaining the CUSUM technique’s major benefit of speed and simplicity are currently being investigated. Until this issue can be addressed, the use of the technique is limited to step changes in concentrations sufficiently large to override confounders.

We therefore decided to screen measurements from those CCS indicator sites that reported a large change in concentrations over the 2 years before and after implementation of the CCS to establish whether the change occurred gradually or at a certain point in time.

CUSUM SCREENING OF CCS INDICATOR SITES

Methods

Although the predicted concentrations due to the CCS was not large in comparison with seasonal and other independent changes, initial statistical analysis of results (presented in Appendix L, which is available on the HEI Web site) showed that, at certain key indicator sites, the 2-year means before and after CCS implementation did display a larger-than-expected change. The CUSUM method was therefore applied as a straightforward screening exercise to these sites and pollutant species where such change may be sufficiently large to allow robust identification of a change point.

Identification of Data That Show Exceptional Change Geometric mean NO_x, NO_2, PM_{10}, and CO concentrations during CCH over the 2 years before and after CCS implementation were calculated for all monitoring sites within Greater London that met the completeness criteria of 75% capture rate over the 4-year period. Due to the low number of available sites within the CCZ, this threshold was lessened to 50% for Bloomsbury—Russell Square, City of London—Senator House, and Westminster—Horseferry Road. Results from these sites should be considered in this context. There were insufficient data from sites that monitor $PM_{2.5}$ and black smoke to analyze these pollutants.

The Bland–Altman test of repeatability (Bland and Altman 1986) was applied to each pollutant to identify key indicator sites, within or surrounding the CCZ, that recorded a change in mean concentration more than 2 standard deviations from the Greater London network mean change. The Bland–Altman plots show overall 4-year mean concentrations at each monitoring site arrayed against the change in the geometric mean concentration between the 2 years before and the 2 years after CCS implementation. The ***x***-axis was offset to the network mean change with upper and lower confidence limits marked with dashed lines. This analysis identified one key indicator site that measured exceptional changes in PM_{10} (Bloomsbury—Russell Square), in NO_x and NO (Tower Hamlets—Mile End Road), and in NO_2 (Westminster—Marylebone Road). No key indicator sites were identified for CO. The Bland–

Altman plot for NO_2 is shown in Figure C.6. A small number of sites in Outer London were also identified as outliers (one for PM_{10}, two for NO_2 and NO, four for NO_x, and three for CO). The assumption was made that these sites were too far from the CCZ to have experienced such a large magnitude of change directly relating to the implementation of the CCS.

Screening of Selected Indicator Sites Input data for the CUSUM analysis were log-transformed to decrease skewness; only measurements for CCH were included and they were averaged over Monday–through-Friday weekly periods to minimize short-term serial correlation. Suburban pollutant measurements were subtracted from hourly mean concentrations before analysis as a background control in an attempt to isolate the local component of measurements. For all analyses, the parameter ***k*** (the allowable slack in the process)was set at 0.3 due to the high standard deviation of the input data.

Results

The most striking result of the CUSUM analysis was that for NO_2 at Westminster—Marylebone Road, a curbside site on the CCZ boundary. A sudden and sustained increase in NO_2 concentrations was evident and the CUSUM signal was sufficiently strong despite seasonal variation to identify a change-point date between February and April 2003 (Figure C.7). It is worth noting that mean NO_x concentrations at Westminster—Marylebone Road did not increase by significantly more than the network mean, which indicates that the proportion of NO_x existing as NO_2 must have increased. The reasons for this stepped increase close to the timing of the introduction of the CCS at this particular site are discussed more fully elsewhere.

The CUSUM analysis of NO_x at Tower Hamlets—Mile End Road, the only indicator site that showed a large change in mean concentrations, was inconclusive. A sustained negative deviation in the lower CUSUM trace was evident but variable. In this case, the signal was insufficiently strong to override seasonal and long-term effects and no clear change point could be identified (Figure C.8). The results for the analysis of NO at Tower Hamlets—Mile End Road were similar (data not shown).

The only key indicator site to show exceptional change in PM_{10} was Bloomsbury—Russell Square, a background site within the CCZ. The resulting CUSUM output did show a steady gradient in the lower CUSUM trace, indicating a sustained stepped decrease in PM_{10} concentrations in early 2003 (Figure C.9). However, PM_{10} measurements at this site between June 2002 and May 2003 were excluded due to equipment failure, and concentrations before this period were affected by roadwork close to the site that produced an artificially high reference mean (described in the main report in the section on establishing a CCS Study Database). The stepped decrease in 2003 indicates a return to concentrations before the road work started rather than a decrease as a result of the CCS.

Discussion

The major strength of the CUSUM method over other methods used in the CCS study is that it can identify the approximate timing of changes that may have been caused by an intervention. However, this ability is weakened by the effects of serial correlation within the data caused by seasonality and long-term trends. The application of the CUSUM analysis was therefore limited to datasets that were demonstrated as measuring exceptional changes, defined as more than 2 standard deviations from the Greater London network mean change.

The CUSUM analysis did reveal step-change signals at two sites sufficiently strong to overcome the confounders: a stepped increase in NO_2 levels at Westminster—Marylebone Road (IRR–CCZ Boundary) and a stepped decrease in PM_{10} concentrations at Bloomsbury—Russell Square (Within CCZ–Background). The increase in NO_2 at Westminster—Marylebone Road could be dated within a few months of the introduction of the CCS. However, the identified decrease in PM_{10} was related not to the CCS but to unusually high concentrations due to road work started before the intervention. An inconclusive result was found in respect to a decrease in NO_x concentrations at one roadside site, Tower Hamlets—Mile End Road (Inner London–Roadside).

We therefore conclude that CUSUM can provide useful evidence about changes in concentrations from preexisting levels. The interpretation of CUSUM results is most secure when detecting a change in pollution that is large by comparison with seasonal fluctuations and the shifts of long-term trends. In the more usual circumstance of studying comparatively subtle changes in pollution, the secure interpretation of CUSUM requires adaptation of the technique to use estimates of the standard deviation and means of pollutant concentrations that take proper account of the underlying correlation among measurements. If these limitations can be overcome, the CUSUM technique should prove to be a valuable addition to the range of air pollution time-series analysis methods.

Although the CUSUM screening study was not able to provide a quantitative estimate of changes in pollution levels arising from the introduction of the CCS, the strong signals that were identified should be considered in the context of other analytic results in the CCS study. CUSUM results were the only results presented that provided information on the timing of a significant change in pollution levels.

REFERENCES

Bland JM, Altman DG. 1986. Statistical methods for assessing agreement between two methods of clinical measurement. Lancet 1(8476):307–310.

Kuebler J, van den Bergh H, Russell AG. 2001. Long-term trends of primary and secondary pollutant concentrations in Switzerland and their response to emission controls and economic changes. Atmos Environ 35: 1351–1363.

Lucas JM. 1982. Combined Shewhart-CUSUM quality control schemes. J Qual Tech 14:51-59.

Lucas JM, Crosier RB. 1982. Fast initial response for CUSUM quality control schemes: Give your CUSUM a head start. Technometrics 24:199-205.

Page ES. 1954. Continuous inspection scheme. Biometrika 41:100-115.

Porter PS, Rao ST, Zurbenko IG, Dunker AM, and Wolff GT. 2001. Ozone air quality over North America: Part II. An analysis of trend detection and attribution techniques. J Air Waste Manag Assoc 151:283–306.

Rao ST, Zurbenko IG. 1994. Detecting and tracking changes in ozone air quality. J Air Waste Manag Assoc 44:1089–1092.

Rossi G, Lampugnani L, Marchi M. 1999. An approximate CUSUM procedure for surveillance of health events. Stat Med 18:2111-2122.

Ryan TP. 2000. Statistical Methods for Quality Improvement, 2nd edition, John Wiley & Sons.

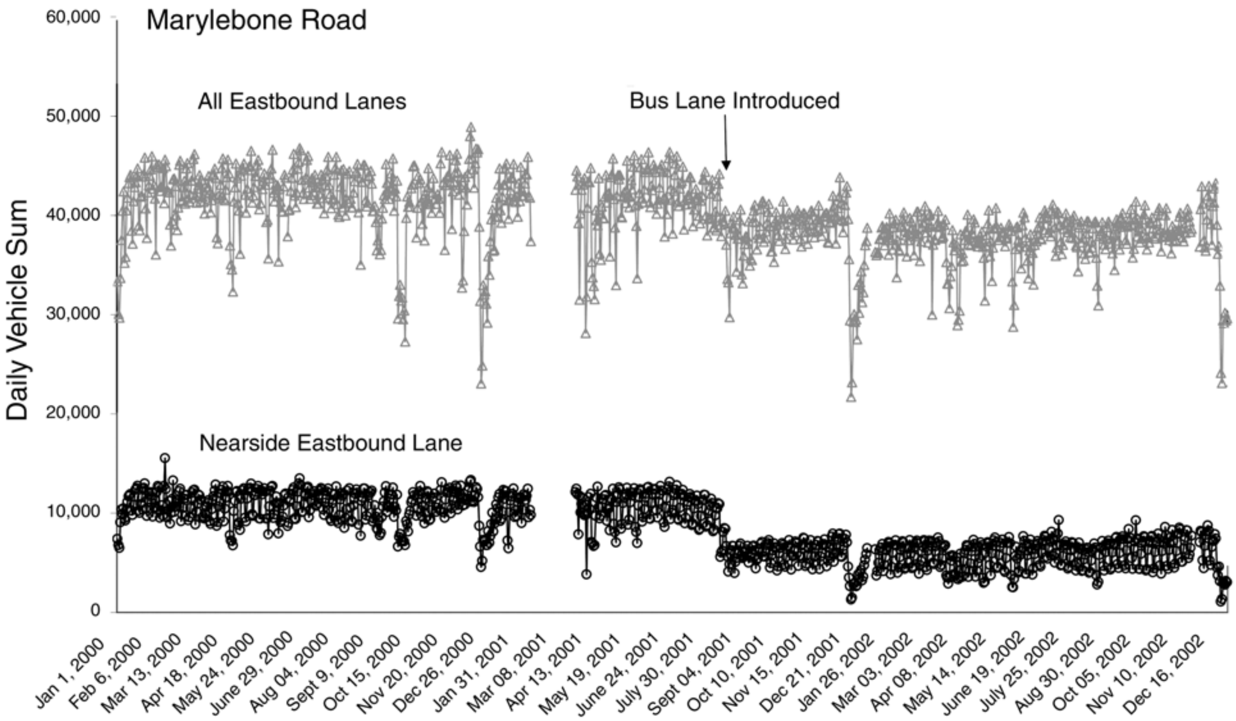

Figure C.1. Time-series plot of daily total vehicle count measured at Westminster—Marylebone Road between January 1, 2000, and December 31, 2002. March and April 2001 data are missing due to equipment failure.

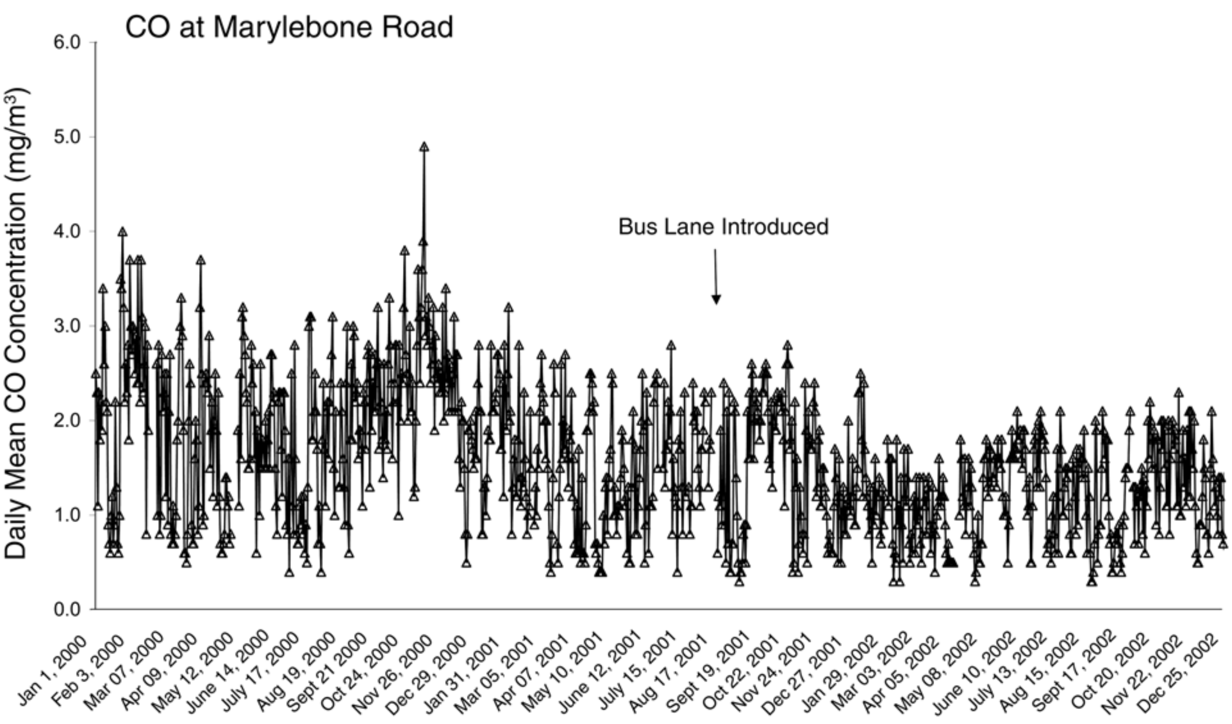

Figure C.2. Time-series plot of daily mean CO concentrations measured at Westminster—Marylebone Road between January 1, 2000, and December 31 2002.

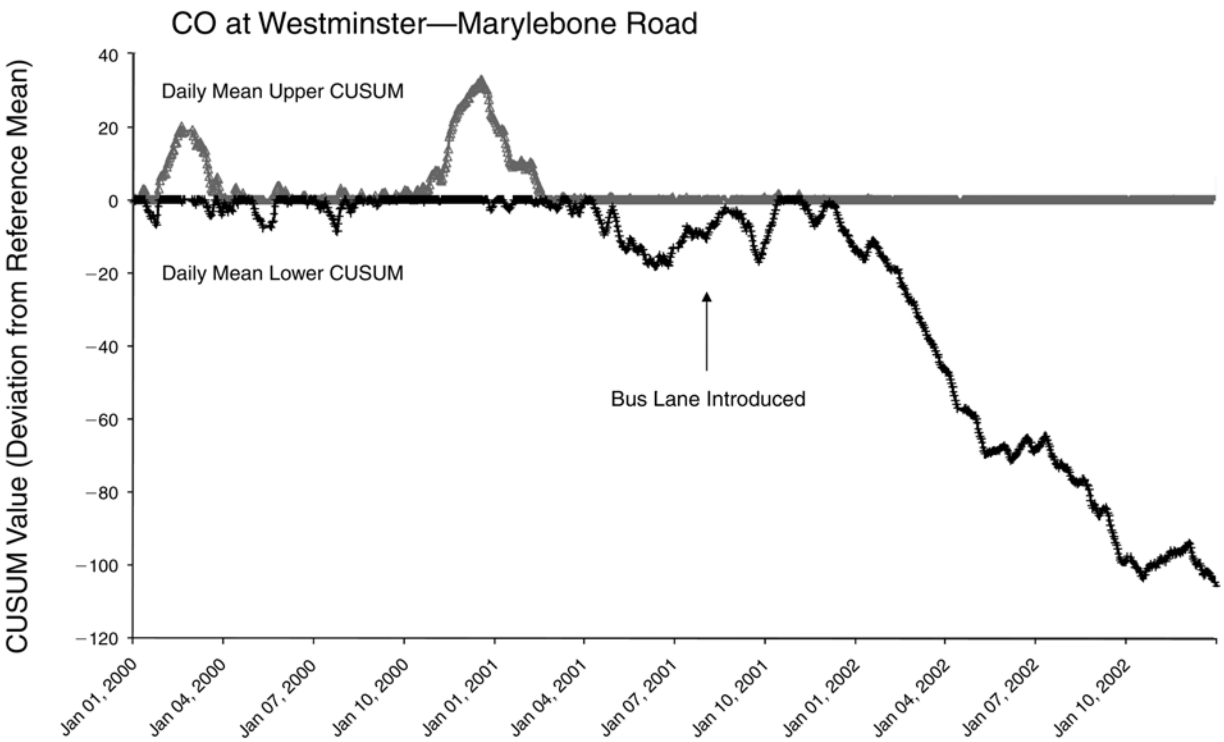

Figure C.3. CUSUM chart for S_{Hi} and S_{Li} CO data at Westminster—Marylebone Road.

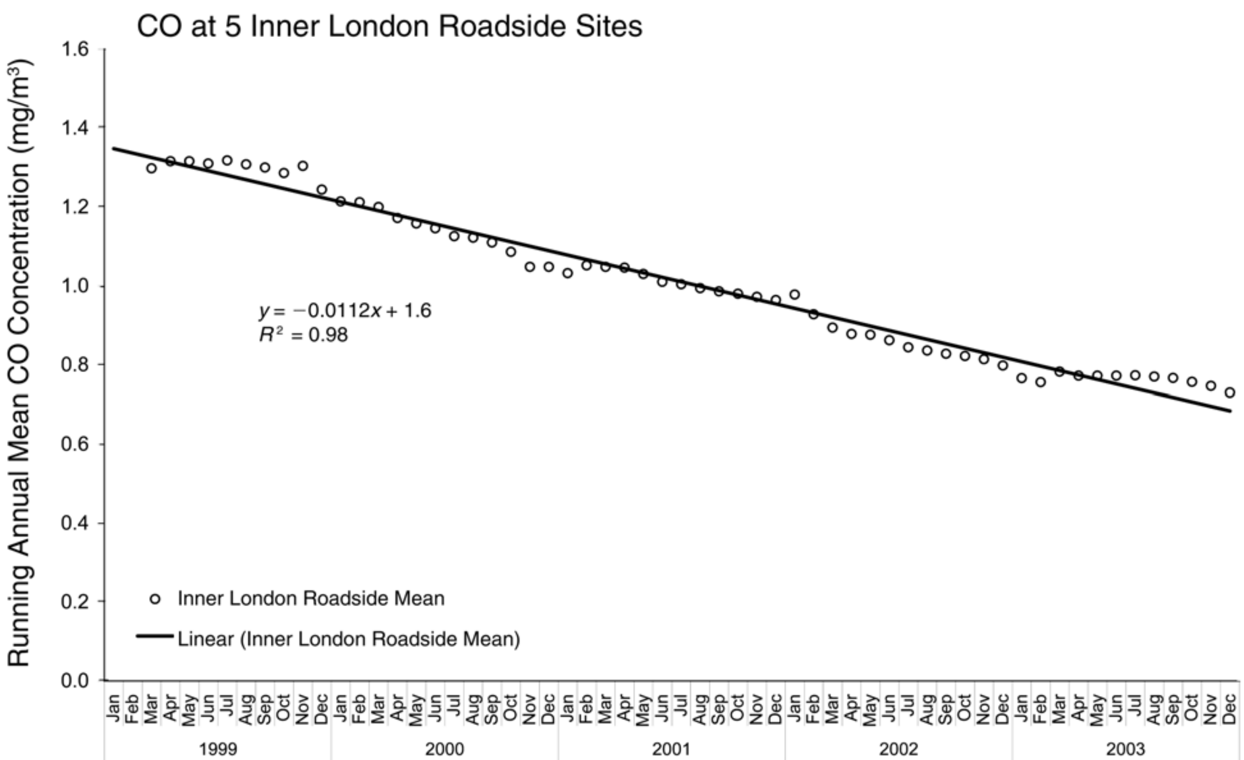

Figure C.4. Monthly increment of running annual mean CO concentrations at five Inner London roadside sites (not including Westminster—Marylebone Road).

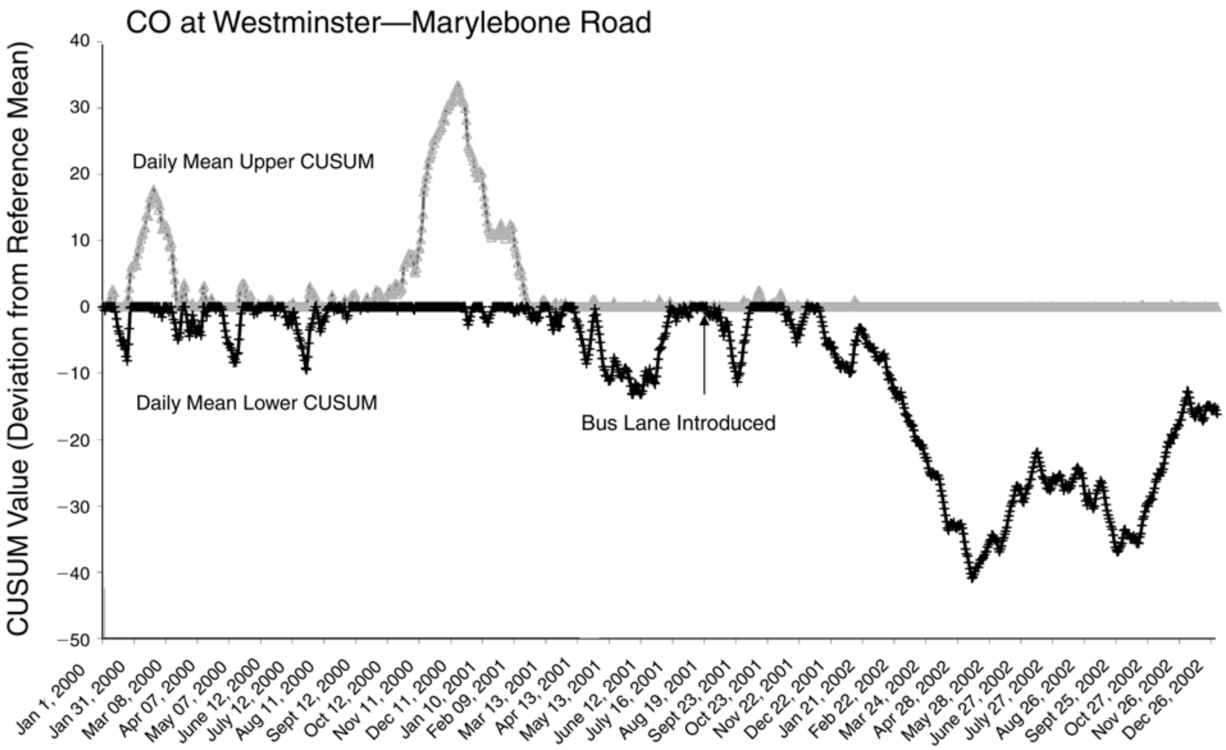

Figure C.5. CUSUM chart for S_{Hi} and S_{Li} CO data at Westminster—Marylebone Road adjusted for background trends.

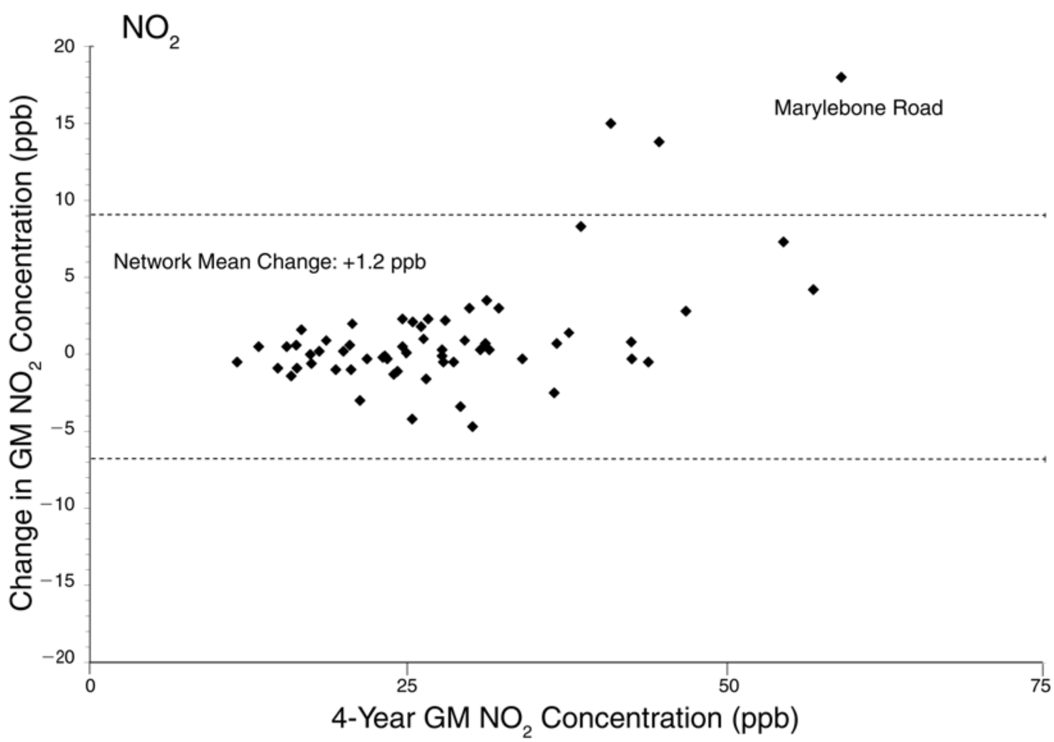

Figure C.6. Bland–Altman plot for NO_2 highlighting those sites that measured exceptional change between the pre- and post-CCS analysis periods in comparison with other sites in Greater London. Upper and lower confidence limits are shown with dotted lines.

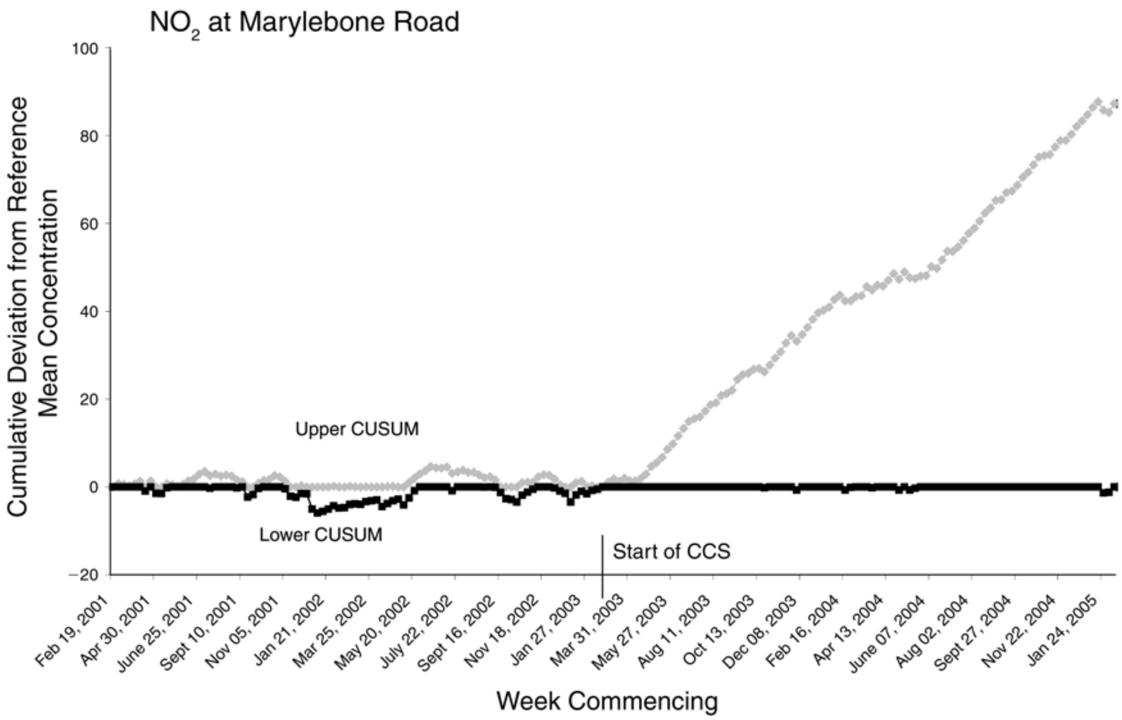

Figure C.7. CUSUM chart of NO_2 concentrations at the Westminster—Marylebone Road curbside site on the CCZ boundary, showing a signal in the upper CUSUM sufficiently strong to overcome the effects of seasonality. Data are for CCH only.

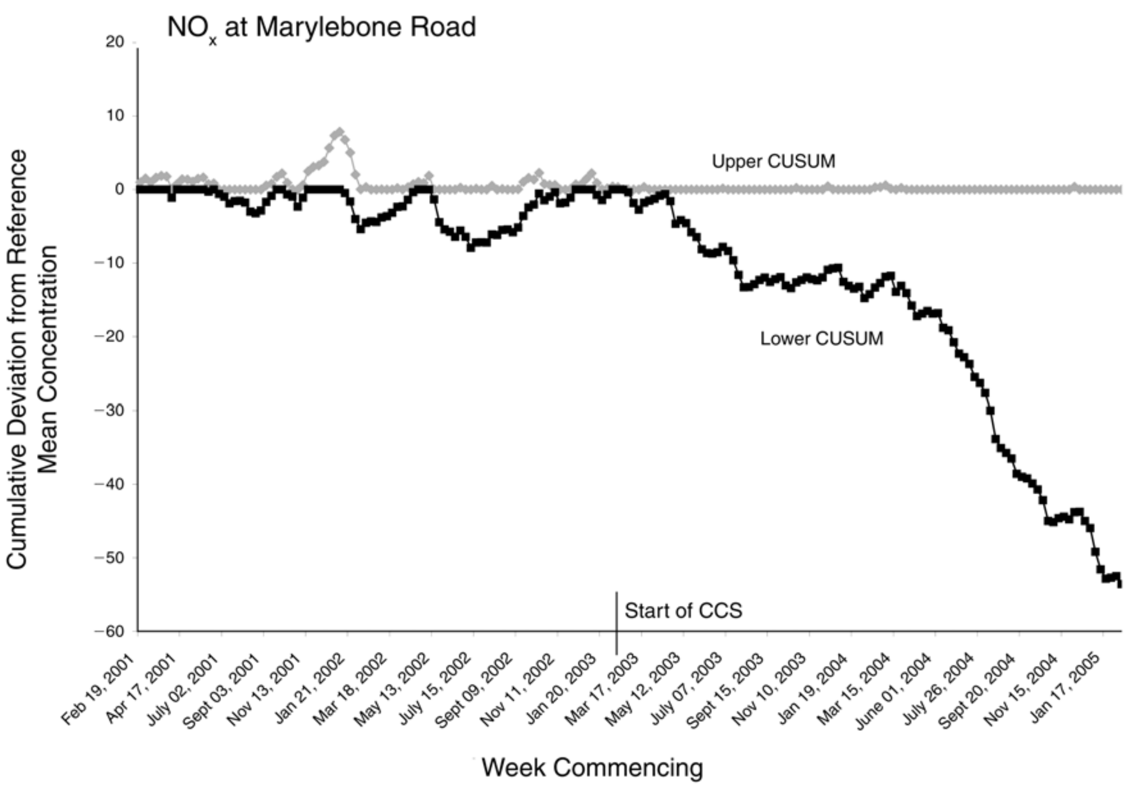

Figure C.8. CUSUM chart for NO_x at Tower Hamlets—Mile End Road showing a lower CUSUM signal insufficiently strong to allow clear identification of a change point. Data are for CCH only.

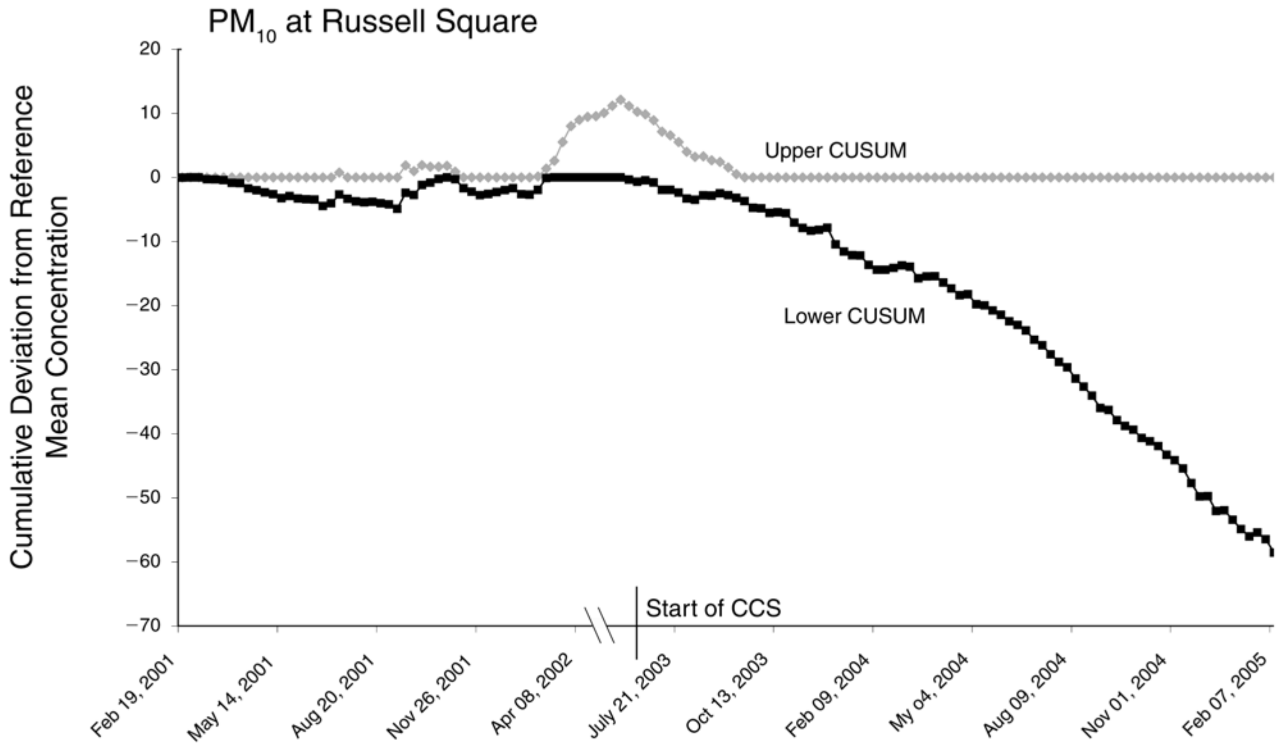

Figure C.9. CUSUM chart for PM_{10} at the Bloomsbury—Russell Square background site within the zone. Measurements between June 2002 and May 2003 were excluded due to equipment failure.