

8/28/03

Worker Productivity and Ventilation Rate in a Call Center: Analyses of Time-Series Data for a Group of Registered Nurses

William J Fisk*, Phillip Price, David Faulkner, Douglas Sullivan, Dennis Dibartolomeo
Indoor Environment Department
Lawrence Berkeley National Laboratory
1 Cyclotron Rd, MS 90R3058
Berkeley, CA 94720

Cliff Federspiel, Gang Liu, and Maureen Lahiff
University of California, Berkeley
Berkeley, CA 94720

* Fax: (510) 486-6658, Email: wjfisk@lbl.gov

Worker Productivity and Ventilation Rate in a Call Center: Analyses of Time-Series Data for a Group of Registered Nurses

Abstract

We investigated the relationship of ventilation rates with the performance of advice nurses working in a call center. Ventilation rates were manipulated; temperatures, humidities, and CO₂ concentrations were monitored; and worker performance data, with 30-minute resolution, were collected. Multivariate linear regression was used to investigate the association of worker performance with indoor minus outdoor CO₂ concentration (which increases with decreasing ventilation rate per worker) and with building ventilation rate. Results suggest that the effect of ventilation rate on worker performance in this call center was very small (probably less than 1% or nil, over most of the range of ventilation rate (roughly 12 L s⁻¹ to 48 L s⁻¹ per person). However, there is some evidence of worker performance improvements of 2% or more when the indoor CO₂ concentration exceeded the outdoor concentration by less than 75 ppm.

Keywords

carbon dioxide, field study, office building, performance, productivity, ventilation rate

Practical implications

The results of this study suggest that practical increases in ventilation rates above approximately 12 L s⁻¹ per person are unlikely to affect the speed of work of call center workers; however, further research in different buildings is needed before general conclusions can be drawn.

Introduction

Approximately 60% of the U.S. workforce performs non-industrial and non-agricultural work inside buildings (U.S. Department of Commerce 1997), and about half of the workforce can be characterized as office workers. Much of the work performed by office workers is cognitive, e.g., information processing. Prior research suggests that performance (e.g., speed or accuracy) of cognitive work can be affected by indoor thermal conditions (Wyon 1993, 1996a, 1996b, Seppanen et al 2003). In previous studies, increased ventilation rates and reduced indoor carbon dioxide concentrations have been associated with improvements in health at work (Seppanen et al., 1999). Only a few studies have assessed the relationship of ventilation rates with worker performance. In a study of 35 Norwegian classrooms, higher concentrations of CO₂, which indicate lower rates of outside air ventilation per person, were associated with poorer performance ($p < 0.01$) in computerized tests of reaction time (Myhrvold et al., 1996); however, the percentage change in performance was not specified. In a study by Nunes et al. (1993), workers who reported building-related health symptoms, known to be associated with lower ventilation rates (Seppanen et. al., 1999), took 7% longer to respond in a computerized neurobehavioral test of sustained visual attention ($p < 0.001$) and had 30% higher error rates in a symbol-digit substitution test of speed and coding ability. In laboratory experiments by Wargocki et al. (2000), increasing the ventilation rate in a room with a carpet removed from a

complaint building was associated with improvements of a few percent in speed or accuracy of several simulated work tasks such as text typing, addition, proof reading, and creative thinking ($p < 0.05$).

With few exceptions, prior studies have assessed the influence of ventilation or indoor environmental quality (IEQ) parameters on worker's performance in special work-related tests, such as tests of reaction time and multiplication speed. Few investigations have documented the relationships of IEQ with actual cognitive work performance in real work environments. The difficulty of defining and measuring the cognitive performance of workers has been one of the barriers to this research. However, for a few types of cognitive work, worker performance has been clearly defined and routinely measured by the employer. Workers in call centers are an example. In call centers, large pools of workers interact with clients via the telephone and enter data or process information associated with the telephone calls. The types of work include sales, appointment scheduling, trouble shooting, and providing advice. To track worker performance, call centers frequently have automatic systems that record data on worker speed along with the type or purpose of the calls. Consequently, call centers are an appropriate setting for studies of the dependence of work speed, but not work quality, on IEQ. This paper describes such a study in a call center operated by a health maintenance organization. The primary objectives of this study were to determine if the rate of outside air ventilation affected the performance of workers and to quantify the magnitude of the effect.

Based on the studies discussed above, we would expect practical changes of ventilation rate or IEQ to alter work performance by only a few percent. However, opportunities to improve worker performance by even one or two percent through improvements in IEQ may be very attractive because the cost of workers' salary and benefits is very large relative to the costs of improving IEQ.

Methods

Overall approach

The approach employed in this study was to manipulate outside air ventilation rates and monitor indoor air temperatures (which vary naturally), while collecting telephone call data quantifying worker performance. Data were collected between July 28 and October 24, 2000 from a call center located in the San Francisco Bay Area. Data were analyzed with multivariate statistical models. The workforce was blinded regarding all aspects of the study, except they were aware that indoor air temperatures were being monitored. The study protocol was approved by the University of California's Institutional Review Board.

Building and study population

The study building is a call-center operated by a health maintenance organization located in northern California. The building, constructed in 1998, has two floors, a total floor area of 4,600 m² (50,000 ft²), sealed windows, carpeted floors, concrete ceilings, and walls of glass and concrete. Workstations are predominately located within cubicles that house one to four workers. Each call center worker has a computer and telephone with a headset. The appearance

of the workspace is pleasant and the maximum worker density in the building of 6.3 persons per 100 m² (1076 ft²) is typical of offices.

The call center was heated, cooled, and ventilated by four air handling units (AHUs) located on the rooftop. Each AHU served a portion (air handling zone) of the building's interior, although there was considerable mixing of air among the air handling zones. One AHU served spaces not occupied by the study population (entryway, lounge, and a small number of private offices for management staff and meetings). The variable air volume (VAV) AHUs modulate the rates of supply of cool or warm air to maintain indoor air temperatures in the desired range. Each AHU has an "economizer" control system that modulates the rate of outside air supply, above a minimum rate established by the building code, with the goal of minimizing costs for heating and cooling; however, to prevent unplanned changes in outside air supply the economizer controls were deactivated during most experimental periods.

The workers were predominantly registered nurses (RNs) who provide medical advice and tele-service representatives (TSRs) who schedule appointments. For reasons discussed subsequently, the analyses discussed in this paper included only the RNs. Workers were present in the building at all times (7 days per week and 24 hours per day), although the number of workers was highly variable, with the largest workforce on weekday mornings. During the study, the maximum number of RNs and TSRs within the building at any time were 119 and 173 respectively.

An incoming telephone call from a patient was normally routed to a TSR who scheduled an appointment or transferred the call to the RNs who may have asked the patient questions, provided medical advice, and, (when needed) scheduled an appointment. Both RNs and TSRs used the computer system to obtain and enter information.

Work shifts varied from 0.5 to 12.4 hours (Mean 6.8 hr, median 6.4 hr). For some workers, the days of the week worked and time of day worked were variable.

Manipulation and Measurement of Ventilation Rates

Prior to data collection, we added equipment to each AHU enabling manipulation and approximate measurement of outside air ventilation rates. The outside air flow rate in each AHU was computed as the product of the supply airstream flow rate and the fraction of outside air in this airstream.

For each AHU, we used a carbon dioxide monitor calibrated weekly with five calibration gas standards, to measure concentrations of CO₂ every 7.2 minutes in the outside air, return airstream, and supply airstream. With these measurements and the simple mass balance calculation in equation 1 (Drees et al. 1992), we computed the fraction of the supply airstream that was outside air (fraction of outside air), with the remainder of the supply airstream being recirculated indoor air,

$$FOA = (C_R - C_S) / (C_R - C_{OA}) \quad (1)$$

where FOA is the fraction of outside air, C_R , C_S , and C_{OA} are CO₂ concentrations in the return, supply, and outside airstreams, respectively. In periods of low occupancy, the differences between CO₂ concentrations in the numerator or denominator of equation 1 were too low for

accurate determinations of FOA. Therefore, outside air ventilation rates were based on the average of measured rates when C_R exceeded C_S by at least 5 ppm.

To measure the flow rate of each AHU's supply airstream, we measured an average supply airstream velocity using one to three Pitot-static tubes and electronic pressure transducers calibrated prior to the study with an electronic micromanometer. Data were recorded for each sensor once per minute. The optimal location(s) of the Pitot-static tubes within the supply airstreams were determined during the study preparation phase using an array of eight Pitot static tubes to characterize the velocity profiles across a cross-section of the supply ducts for a range of supply flow rates.

The AHUs had dampers for modulation of the FOA. Fixed damper positions for low ventilation rate periods were selected to match the code-minimum outside air supply rate of 12.0 L s^{-1} per occupant at maximum occupancy (0.76 L s^{-1} per square meter of floor area and 292 persons). The fixed damper positions for medium and high ventilation rates were selected to provide approximately twice and four-times the code minimum. In a fourth ventilation setting, the normal control systems for the AHU's outside air supply, including the outside air economizers, were activated. We anticipated that this mode of operation (called economizer mode) would typically provide a ventilation rate greater than eight times the code minimum. In practice, ventilation rates in economizer mode varied considerably.

Using these methods, we scheduled periods of ventilation in each of the four control modes: low, medium, high, and economizer mode. The intent to have randomized ventilation rates that changed daily during weeks 3 – 6 and 8 – 10 was met reasonably well. During weeks 1, 2, 7, 11 and 13, we intended to fix the ventilation rates at the low, medium, high, or economizer setting for one-week periods; however, the control system failed during some periods. The resulting schedule of ventilation control modes is provided in Table 1. The control system resulted in a wide range of ventilation rates; however, these ventilation rates were not sufficiently repeated to use the control mode as a categorical surrogate for the ventilation rate. Thus, measured ventilation rates and carbon dioxide concentration were used in analyses of the worker performance data.

Measurements of carbon dioxide concentration, temperature, and humidity

CO_2 concentrations were measured, as described above, in the return airstreams of all AHUs. Concentrations in return airstreams represent an approximate average of the concentration of CO_2 in the associated occupied spaces of the building. The difference between indoor CO_2 concentrations and outdoor concentrations (ΔCO_2) is not easily related to accurate estimates of rates of outside air supply; however, these ΔCO_2 concentrations are measures of the degree of control of occupant-generated air pollutants via outside air ventilation. The average of the return-air CO_2 concentrations was used in the analyses described subsequently.

Air temperature and humidity were measured approximately one meter above floor level throughout the spaces occupied by the study population. Temperatures were logged every one

minute at 25 indoor locations and relative humidities were logged every five or 15 minutes at 11 indoor locations. More information on these measurements is provided in Fisk et al. (2001).

Collection of worker performance data and associated data

The call center's automated call distribution (ACD) system monitors several performance-related parameters. Worker performance for each half-hour period is summarized with the "average handle time" (AHT) of all of the calls that ended during that period. The AHT is the average time (averaged over all workers, called agents, and over the entire half hour) taken for each call from beginning to end, starting when the call was answered and ending when the agent completely finished all tasks associated with the call and was available to answer another call.

For each half-hour period, the call center's computer calculated a number, called "nets", that estimated (based on prior experience and current number of incoming calls) how many extra RNs were on hand, compared to the number needed to have the average wait experienced by callers equal to a target wait time. When the target wait times were exceeded, nets was negative. Nets was used as a variable in the data analyses, as a measure of the work demand on RNs.

Statistical analyses

Our primary interest was the relationship between ventilation rate and worker performance -- measured by AHT. Ventilation rate was expected to influence AHT by at most a few percent, which is less than the variation due to several other sources. In principle, in a very long study using randomized levels of ventilation rate, the other sources of variation would cancel out and any effects of changes in ventilation could be examined with a simple categorical analysis. In practice, with only twelve weeks of data, randomization of ventilation would not necessarily be sufficient to ensure that other sources of variation cancel out completely. Therefore, to estimate effects of ventilation-related variables with useful precision we excluded some data from the analyses and controlled for other sources of variation in AHT using multivariate regression models.

The number of calls varied greatly throughout the day, as did the number of workers agents scheduled to handle them. Carbon dioxide concentrations were never high (> 300 ppm above the outdoor concentration) outside the normal workday (7:30 a.m. – 6:00 p.m. Monday through Friday). Although excluding data collected from nights and weekends significantly reduced the amount of data analyzed, the alternative was to rely on regression modeling alone to relate performance during the night-time and weekend, when CO₂ concentrations were always low, to performance during the weekday periods with a wide range of carbon dioxide levels. We did not trust such modeling to be free of small sources of bias that could overwhelm the small expected effects of ventilation rate on worker performance; consequently, we elected to discard data outside the normal work week (Monday – Friday, 7:30 a.m. – 6:00 p.m.).

The tasks performed by the appointment schedulers (the TSRs) were simpler than those performed by the registered nurses (RNs). The RNs had some "wrap-up" time associated with most calls, during which they create a computer record of the medical problem and their response to it. For both RNs and TSRs, the time required to handle a particular call was

substantially dependent on the caller rather than on the agent, but RNs can have a larger influence on their own work rate. Therefore, we concentrated on analyzing the RN data and all of the results provided below are for RNs.

Data from Labor Day (a Monday) and the following day were clearly anomalous in the number of calls, although this is not reflected in unusual values of AHT: many fewer people called on Labor Day than on a typical Monday, and more people than usual called on the following day. We excluded data from Labor Day and the following day from the analyses.

Data from two additional days were excluded. Substantial changes to the computer software used by the agents were made just before the start of the study, so we excluded what would otherwise have been the first day of the study (a Friday) to try to avoid data affected by the learning period. Further software changes were made on day 58 of the study (a Saturday), and these changes apparently slightly affected AHT (by increasing it) until workers became accustomed to the changes. For this reason, we excluded day 60 (a Monday). Except where specified, all results below are based only on RN data from 7:30 a.m. - 6:00 p.m. Monday through Friday, excluding Labor Day and the following day, and excluding day 1 and day 60.

Linear regression was our main tool for analyzing the data: we regressed AHT or $\log(\text{AHT})$ on explanatory variables that are expected to be relevant. Although we performed analyses using both AHT and $\log(\text{AHT})$, we favor the log transformation for two reasons: the resulting residuals have a distribution that is closer to normal; and we anticipate that changes in environmental factors (or in other confounding factors) that reduce or increase work rate should increase or decrease AHT by a relative rather than an absolute amount. We report results only for analyses based on $\log(\text{AHT})$.

We have no direct measurements of noise, call content, caller cooperation, worker motivation, and other factors that are expected to directly influence agent performance, but we do have explanatory variables that serve as proxies for these parameters:

1. Noise and level of activity are related to the number of agents working at a given time, so we included the number of agents as one of our explanatory variables.
2. Callers' cooperation, workers' motivation, and call content were all expected to be related to queue length, so we included the "nets" variable.
3. Call content was expected to vary with time of week, so we included time-of-week indicator variables in the regression.
4. Indoor air pollutant levels were not directly measured with useful frequency, but indoor and outdoor carbon dioxide were measured every 7.2 minutes. CO_2 never reached levels at which it would directly affect performance. However, the difference between the indoor and outdoor concentration (ΔCO_2) concentration is a proxy for ventilation per occupant. The difference between indoor and outdoor CO_2 concentrations, ΔCO_2 , should be correlated with the indoor concentration of any pollutant emitted indoors at a rate that is approximately proportional to the number of people in the building. Examples of such pollutants are body odors, perfumes, dust stirred up by activity, and emissions from equipment used by occupants such as computers and copy machines.

We cannot expect to perfectly adjust for all of the actual sources of variation in AHT. Instead, the goal was to adjust for systematic variation that could interfere with obtaining accurate estimates of the effects of ventilation-related variables.

AHT varied systematically throughout the day and throughout the week. To control for this systematic variation, we created time-of-week indicator variables for each half-hour period. These time-of-week indicator variables alone accounted for about 35% of the variation in $\log(\text{AHT})$.

We performed a variety of regressions on $\log(\text{AHT})$, using different combinations of explanatory variables chosen from the following:

1. time-of-week indicator variables;
2. building-average temperature (Celsius) – 23 °C;
3. $(\text{building-average temperature} - 23\text{ °C})^2$ was included as a variable to account for the possibility of a non-linear relationship between temperature and performance, e.g., a temperature associated with maximum performance;
4. building-average relative humidity;
5. number of agents on duty---we have data on RNs and TSRs separately and using the RN occupancy number alone provided a slightly better fit to the RN data.;
6. normalized “nets” (piecewise linear), normalized to the number of workers on duty; for example, if 30 agents were working, and nets equaled 3, normalized nets was 0.1, since there was a 10% surplus of agents;
7. indoor-outdoor carbon dioxide concentration difference (linear or categorical, or piecewise linear).

The piecewise-linear model for “nets” divided the parameter into bins, and assumed that within each bin the parameter had a linear influence on $\log(\text{AHT})$, but we allowed the different bins to have different slopes and intercepts.

For each model, we first performed an ordinary regression of $\log(\text{AHT})$ on the explanatory variables, weighting each data point by the number of calls received during the half hour. Eighty percent of the half-hour periods in our analysis included between 130 and 260 calls, so the weights were not highly variable, nor were they very influential. After fitting the model, we calculated the serial correlation of the residuals, as a function of time lag. In every model the lagged serial correlation was moderately high (greater than 0.2) for lags of a few hours, but dropped to near zero over several hours. The temporal correlation reduced the statistical power of the analysis: effectively, if observations were highly correlated, then each observation added less independent information. We used the observed serial correlation (or, rather, the observed lagged covariance) to estimate the variance-covariance matrix of the residuals, assuming the covariance to be identically zero for time lags exceeding six hours. Following a standard approach for regression in which the residuals have an off-diagonal variance-covariance matrix (e.g. see Box et al., p.363 or Gelman et al., p. 257), we then performed a linear regression that used the same explanatory variables but adjusted for the temporal correlation of the residuals. Compared to the conventional regression, in which observations are assumed to be independent, the coefficient estimates changed by a small amount and width of the error bars increased slightly.

Results

Ventilation rates and ΔCO_2 concentrations

Figure 1 shows the ΔCO_2 concentration versus time, separately for each day of the week, with all weeks superimposed. Only data from 7:30 a.m. to 6:00 p.m. are shown. The lower right plot in the figure shows histograms of the carbon dioxide concentration difference for all of the days (upper line in each bin), and also for just Monday through Friday (lower line in each bin). As the histogram shows, ΔCO_2 on weekends was never above 200 ppm, and ΔCO_2 concentrations on weekdays were rarely below 100 ppm.

Figure 2 shows ΔCO_2 versus ventilation rate (flow rate of outside air supply), displaying all of the data for the entire study (not just the normal work week). Superimposed is a line showing the results of a local robust regression using the lowess method (Cleveland, 1979; with a smoother span of 1/5, and 3 iterations). ΔCO_2 values tend to cluster into three wide clumps, corresponding to low, medium, and high damper settings, with “economizer” settings also tending to lead to fairly low or very low values of ΔCO_2 .

ΔCO_2 varies with both ventilation rate and the number of people in the building; in steady state, the concentration would be nearly proportional to the number of people, and inversely proportional to the ventilation rate. Thus a given ventilation rate can correspond to a wide range of carbon dioxide concentrations (or other pollutants associated with the presence of people).

Average Handle Times

As illustrated in Fisk et al (2001), AHT varies throughout the day, and from day to day throughout the week. The variation may ultimately be due to several causes, including variation in call content, and in worker alertness and fatigue. During the normal workday, 7:30 a.m.- 6:00 p.m. Monday through Friday the AHT ranges from about 500 seconds in the morning to about 540 seconds in the evening, as well as varying slightly from day to day.

Temperatures and humidities

The building’s climate control held the building-average temperature within a very narrow range during daytime working hours, although there was a larger variation in temperatures in specific zones (Federspiel et al 2002). Fifty percent of the half-hourly work-day temperatures were between 23.1 °C and 23.3 °C, and ninety percent were between 22.9 °C and 23.5 °C. Building average relative humidity almost never straying outside the range 46% to 47%. The very narrow observed ranges of these parameters yielded low statistical power to estimate coefficients of these variables. Unsurprisingly, then, the regression coefficients associated with temperature, (temperature – 23 °C)², and relative humidity, are all very small compared to their uncertainties.

For modeling worker performance and environmental factors, discussed below, we fit models both with and without temperature and humidity data; including these variables had very little

effect on the estimates for the main parameters of interest. The results presented are from models with temperature and humidity included as variables.

Association of worker performance with environmental factors

We fit several dozen regression models, using different definitions for the bin boundaries for number of calls, normalized nets, and ΔCO_2 , and using different subsets of the data. Measures of model fit were very similar for all models that included the full set of explanatory variables, whatever the details of the model.

In every model that we tried, including Models A-C discussed below, the “nets” variables were found to be highly influential. For example, when nets was very negative---that is, when there were many more calls than the agents were able to handle---log(AHT) was elevated by more than 7% compared to the same day and time of day in different weeks. This effect is substantially larger than any expected effect of ventilation, so failure to accurately model the variation of log(AHT) with “nets” could potentially overshadow a ventilation effect. The five-category piecewise-linear relationship between log(AHT) and “nets” that we used in the models allows a more complicated relationship between nets and AHT than would an assumed linear or quadratic relationship across the whole range of “nets” values.

We also included half-hour lagged “nets” in the regressions. Lagged “nets” is the value of “nets” in the previous half-hour. This variable, which was found to be influential, may be important because some calls that terminated in a given half-hour were started in the previous one, and because lagged nets may help predict agent fatigue. Additionally, the mix of calls in a given half-hour period is affected by the wait times from the previous periods: some callers who abandoned calls in the previous period will call back. Although we also examined non-linear relationships between lagged “nets” and log(AHT), these models provided no advantage over linear models for this variable. In the models discussed below, we assumed log(AHT) to vary linearly with lagged “nets.”

We now discuss three specific models for log(AHT) in some detail. Each of these models includes: the time-of-week indicator variables; temperature – 23 °C and (temperature – 23 °C)²; number of agents on duty; five piecewise-linear “normalized nets” categories; and lagged “nets”, allowing linear variation of log(AHT) with the “nets” value of one half hour previous. The three models differ only in their handling of ΔCO_2 . Model A includes no ΔCO_2 variable. Model B includes three ΔCO_2 categorical variables, indicating whether ΔCO_2 for each half hour was: 0-150 ppm, 150-300 ppm, or over 300 ppm. In Model C, the two lower ΔCO_2 categories within Model B have been split, thus, Model C has five ΔCO_2 categorical variables: 0-75 ppm, 75-150 ppm, 150-225 ppm, 225-300 ppm, or over 300 ppm. Table 2 identifies the variables used in each model and provides some of the regression coefficients and associated uncertainties.

Figure 3 shows the residuals from Model A (which did not include ΔCO_2), plotted versus ΔCO_2 . A lowess local regression fit is shown as a solid line (we used a smoother span of 1/5, and 3 iterations). Only for low ΔCO_2 values is there any evidence that the residuals might vary with

ΔCO_2 ; the model tends to predict longer handle times than were actually observed for very low ΔCO_2 concentrations. The right-hand portion of Fig. 3 shows a histogram of the residuals, with a normal distribution (mean 0, standard deviation 0.0448) superimposed. The distribution of residuals is very close to normal, as we assume when we perform least-squares regressions.

Figure 4 shows the estimated model coefficients associated with each ΔCO_2 bin, for Models B (lower plot) and C (upper plot). For each bin, the horizontal bar shows the range of ΔCO_2 spanned by the bin, and the vertical error bar covers plus or minus one standard error. In each case, the lowest bin is defined to have no effect, a coefficient of 0.00.

In the model with only three ΔCO_2 bins (Model B) there is no evidence that lower ΔCO_2 is associated with lower (faster) AHT---indeed, the relationship points the other direction: the estimate for the high- ΔCO_2 bin is about 1% faster than that for the lowest bin (an effect of -0.009 on $\log(\text{AHT})$ corresponds to a factor of $\exp(-0.009)=0.991$ on AHT, which is very close to a 1% speed-up). However, this estimate is not very precise, with an uncertainty (one standard error) of approximately ± 0.6 percentage points.

In contrast, the results from Model C suggest that very low ΔCO_2 concentrations are associated with lower AHT (faster work) than are higher concentrations. All of the estimated coefficients for ΔCO_2 concentrations over 75 ppm are around 0.025 to 0.035, corresponding to handle times that are 2.5% to 3.5% slower than at the lowest ΔCO_2 . Moreover, these effects are all highly statistically significant ($p < 0.05$ for all bin coefficients). However, as we discuss below, we think the statistical uncertainties are understated and that the relationship between AHT and ΔCO_2 is far from conclusive.

Overall, neither the Model B nor Model C results show evidence that AHT increases with ΔCO_2 over most of its range. A dependence of $\log(\text{AHT})$ on ΔCO_2 is apparent only for ΔCO_2 concentrations below about 150 ppm: $\log(\text{AHT})$ is somewhat lower for ΔCO_2 concentrations in the 0-75 ppm range than in the 75-150 ppm range, after adjusting for all of the other explanatory variables. When the 0-150 ppm ΔCO_2 category is split into two categories, as in Model C, the 0-75 ΔCO_2 category has the lowest (fastest) values of $\log(\text{AHT})$, after adjusting for the other explanatory variables. But when the ΔCO_2 data from 0-150 ppm are combined into a single bin, as in Model B, the overall average AHT in this bin is about the same as in the other bins.

In another model, we treated ΔCO_2 as a continuous variable throughout the entire concentration range and found no statistically significant or strong relationship between ΔCO_2 and AHT. These findings were essentially unchanged, when data collected after day 57 were excluded from the analyses in order to eliminate any possible effects of the change in software on day 58.

The analyses based on ΔCO_2 , which is a proxy for ventilation per person, assume that effects on AHT would be caused by pollutants with indoor concentrations approximately proportional to the number of people. However, the building itself can also be a source of pollutants, independent of the number of people in it: walls or carpeting may emit volatile organic compounds, for example. If the building is the source of pollutants that affect performance, then it is total ventilation rate, not ventilation per person, that should predict variation in AHT. We

investigated this possibility by fitting Models like A-C, but using ventilation rate categories rather than ΔCO_2 categories. There is no evidence for a dependency of AHT on ventilation rate; in fact, even for the highest values of ventilation rate there is no evidence for reduced handle time compared to lower ventilation rate values. To the extent that there is an apparent ventilation-related effect in this study, it is due to ventilation rate per person (as indicated by ΔCO_2) rather than ventilation per unit indoor air volume.

Discussion

Considering that all of the workers in the study received at least the code-minimum ventilation, any performance differences associated with ventilation would be expected to be only a few percent at most. It is very hard to design an experiment to quantify such small performance differences in the real world. The analysis of the call center experiment that is described in this paper has inadequate statistical power to find effects smaller than about 2%.

Several characteristics of the call center (or, indeed, of almost any potential study environment) differ from the environments of controlled small-scale experiments. For example, new workers are hired and become part of the work place, while others quit and drop out. Presumably, the new workers are slower than experienced workers, because they haven't learned all the tricks of the trade; or perhaps they are faster, because they are under increased scrutiny or because the novelty keeps them more alert. Changes to software may have effects that are virtually impossible to estimate; for instance, changes may first slow the workers as they learn the new software, but then increase their speed in the long run if the software allows them to work more efficiently. The mix of call types, and thus the average handle time, may change subtly over long periods, perhaps, for this medically-related call center, depending on the types of illnesses that become more or less common over the study period. Over a very long study period these effects will cancel out, if the ventilation parameters are randomized, but over a short period any chance correlation between ventilation-related parameters and the other effects can eliminate the ability to determine the effects related to ventilation.

In the present study, in the M-F, standard workday data there were only forty half-hour periods (out of 1051) in which the ΔCO_2 concentration was below 75 ppm. Nineteen of these periods occurred on a single day (the 78th day of the study, a Friday), and all of the rest occurred during the following week, very close to the end of the 88-day study. The entire apparent speed-up in AHT indicated by Model C, for the below-75 ppm bin relative to the others, is based on data from only six different days, and 65% of those data are from two consecutive Fridays.

In our analysis, when calculating standard errors (and p-values) we took into account the intra-day correlation in residuals, but not the correlation over longer periods. Throughout the study, day-to-day and week-to-week correlation in residuals was usually very low; however, there are some periods in which correlation is evident across several days. One example is days 58-60, discussed previously, when new software led to increased AHT (compared to what the models predict) for several consecutive days. Although investigation of the data and residuals reveals nothing suspicious about the days during which the ΔCO_2 concentrations were very low, the possibility of an unknown effect that led to decreased AHT during that period cannot be ruled out, and is not accounted for in the model. Consequently; in spite of the low p-values, we do not

consider the results of Model C to conclusively indicate a faster work-rate when ΔCO_2 was very low and ventilation per worker was very high. If the very high values of ventilation had occurred on 6 days spread throughout the study period, and the same analyses resulted, we would have confidence that the observed effect is really due to ventilation; in the present case, though, we are simply not sure.

One limitation in these analyses is a consequence of the incomplete separation of time of day from ΔCO_2 . As time of day increases, ΔCO_2 generally increases (and then decreases late in the workday when occupancy is reduced); therefore, controlling for time of day via the regression models could have partially obscured a relationship of ΔCO_2 with worker speed. Including nets in the regression models may also have partially obscured a relationship of ΔCO_2 with worker speed, because a decrease in work speed caused by higher ΔCO_2 would result in a decrease in nets. However, the relationship AHT with ΔCO_2 was similar in limited analyses without the nets variable; i.e., AHT was decreased for the lowest ΔCO_2 concentrations, but there was no other significant relationship of AHT with ΔCO_2 .

We have used ΔCO_2 as a surrogate of ventilation rate per person, recognizing that previous studies have demonstrated the difficulty of accurately computing the ventilation rate per person from CO_2 measurements. The errors result primarily from variable occupancy, variable outside air supply rates, failure to measure outdoor CO_2 concentrations, and inaccurate CO_2 measurements. In this study, outdoor CO_2 concentrations were measured and all measurements were made with research-grade, frequently-calibrated instruments. In addition, outside air supply rates were relatively stable, except in the economizer control mode. Despite these efforts to reduce errors, we do not claim that ΔCO_2 data can be precisely converted to ventilation rates per person. Nevertheless, it is clear that the categories of ΔCO_2 used in our analyses represent different average ventilation rates per person.

Conclusions

If we exclude periods of very high ventilation rates per occupant (indicated by very low ΔCO_2) experienced during the study, we can rule out effects on AHT that are greater than about 2%. There is no evidence of any effect at all, although the present analysis does not have sufficient statistical power to eliminate the chance of effects in the 1% range; such effects, if they occur, would still be of practical importance.

There is some evidence that very high ventilation rates per occupant (very low ΔCO_2) may lead to lower AHT (faster work rates), but the possibility of an unknown confounding variable makes this result less than conclusive in spite of high statistical significance ($p < 0.05$). The statistical model from which the p-values were derived assumes that no temporal correlation in the residuals lasts more than one workday. A “one-time-only” event that lowered AHT slightly (on average) for a week cannot be ruled out, and could be the real cause of the apparent ventilation effect.

The results of the present analysis may not apply to other buildings, or even to other call centers. For instance, it may be that in this building there are no strong indoor sources of pollutants that

influence performance, but that in other buildings such sources exist. If that is the case, then ventilation may have larger effects in those other buildings.

Acknowledgments

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technology Program of the U.S. Department of Energy (DOE) under contract No. DE-AC03-76SF00098 and by the Center for the Built Environment at U.C. Berkeley. The authors thank Michael Sohn and Michael Apte for their reviews of a draft of this paper.

References

- Box GEP, Jenkins GM, Reinsel G, 1994. *Time Series Analysis*, New Jersey: Prentice Hall.
- Cleveland, WS. 1979. Robust locally weighted regression and smoothing scatterplots. *J. Amer. Statist. Assoc.* 74:829-836.
- Drees, K., J. Wenger, and G. Janu, 1992, "Ventilation air flow measurement for ASHRAE Standard 62-1989," ASHRAE Journal, 34(10), 40-45.
- Federspiel, CC, Liu, G, Lahiff, M, Faulkner, D, Dibartolomeo, DL, Fisk, WJ, Price, PN, Sullivan, DP (2002) Worker performance and ventilation: analyses of individual data for call-center workers. *Proceedings of Indoor Air 2002*, vol. 1, pp. 796-801, Indoor Air 2002, Inc., Santa Cruz, CA. LBNL-50124
- Fisk WJ, Price PN, Faulkner D, Sullivan D, Dibartolomeo D, Federspiel C, Liu G, Lahiff M (2001). *Worker Productivity and Ventilation Rate in a Call Center: Analyses of Time-Series Data for a Group of Workers*, Lawrence Berkeley National Laboratory Report, LBNL-49356, Berkeley, CA
- Gelman A, Carlin JB, Stern HS, and Rubin DB, 1995. *Bayesian Data Analysis*. New York: Chapman and Hall.
- Myhrvold AN, Olsen E, Lauridsen O. 1996. Indoor environment in schools – pupils health and performance in regard to CO₂ concentrations. *Proc. Indoor Air 1996, The 7th International Conference on Indoor Air Quality and Climate*. 4: 369-374. SEEC Ishibashi Inc., Japan.
- Nunes F, Menzies R, Tamblyn RM, Boehm E, Letz R. 1993. The effect of varying level of outside air supply on neurobehavioral performance function during a study of sick building syndrome. *Proc. Indoor Air 1993, The 6th International Conference on Indoor Air Quality and Climate*. 1: 53-58. Indoor Air 1993, Helsinki.
- Seppanen OA, Fisk WJ, Mendell MJ. 1999. Association of ventilation rates and CO₂-concentrations with health and other responses in commercial and institutional buildings. *Indoor Air* 9: 226-252.
- Seppanen O, Fisk WJ (2003) A conceptual model to estimate cost effectiveness of the indoor environment improvements. To be published in proceeding of the Healthy Buildings 2003 Conference, December 7-11, 2003, Singapore.

U.S. Department of Commerce, Bureau of the Census. Table 645, *Statistical Abstract of the United States* (117th ed.), 1997.

Wargocki P, Wyon DP, Sundell J, Clausen G, and Fanger PO (2000a) The effects of outdoor air supply rate in an office on perceived air quality, sick building syndrome (SBS) symptoms, and productivity. *Indoor Air* 10(4): 222-236

Wyon DP. 1993. Healthy buildings and their impact on productivity. *Proc. Indoor Air'93, the 6th International Conference on Indoor Air Quality and Climate*. 6: 3-13. Helsinki. Indoor Air 93.

Wyon DP. 1996a. Indoor environmental effects on productivity. *Proc. IAQ'96 "Paths to Better Building Environments"*. pp 5-15, ASHRAE, Atlanta.

Wyon DP. 1996b. Individual microclimate control: required range, probable benefits, and current feasibility. *Proc. of Indoor Air 96*. 1:1067-1072. Tokyo. Institute of Public Health.

Table 1. Ventilation control schedule. L, M, and H refer to fixed damper positions for low, medium, and high ventilation rates. E refers to control of ventilation rates by the economizer control system.

Day	Week												
	1	2	3	4	5	6	7	8	9	10	11	12	13
F	L	H	E	H	L	E	M	M	H	M	E	E	E
Sa	L	H	L	E	H	M	M	M	M	L	E	E	E
Su	L	H	H	L	M	L	M	M	E	M	E	E	E
M	L	H	E	H	L	M	M	M	M	H	E	E	L
Tu	L	H	M	E	H	L	M	L	H	E	H	H	E
W	L	H	H	L	E	H	M	M	L	E	H	E	--
Th	L	H	M	H	M	E	M	L	E	E	H	E	--

Table 2: Regression coefficient estimates and standard errors for three statistical models of log(AHT) on the listed set of explanatory variables. Coefficient estimates for the 105 time-of-week effects, five “nets” categories, and five “nets” slopes are not shown, but are similar for all three models.

Coefficient estimate [units]	Model A	Model B	Model C
<Time of week> [log(AHT)]	Included in model	Included in model	Included in model
< five “nets” categories> [log(AHT)]	Included in model	Included in model	Included in model
< five “nets” slopes> [log(AHT) / nets]	Included in model	Included in model	Included in model
Lagged “nets” [log(AHT) / nets]	0.048 +/- 0.015	0.053 +/- 0.015	0.054 +/- 0.015
(temp – 23 °C) [log(AHT) / degree C]	-0.24 +/- 0.45	-0.23 +/- 0.44	-0.13 +/- 0.43
(temp – 23 °C)^2 [log(AHT) / (degree C)^2]	0.007 +/- 0.027	0.006 +/- 0.03	0.004 +/- 0.029
Relative Humidity [log(AHT) / %]	0.11 +/- 0.22	0.11 +/- 0.22	0.06 +/- 0.22
Agents on duty [log(AHT) / agent]	0.0016 +/- 0.0004	0.0017 +/- 0.0004	0.0017 +/- 0.0004
0 < ΔCO ₂ < 75 [log(AHT)]	Not in model	0.00 +/- 0.00	0.00 +/- 0.00
75 < ΔCO ₂ < 150 [log(AHT)]	Not in model	Not in model	0.036 +/- 0.010
150 < ΔCO ₂ < 225 [log(AHT)]	Not in model	-0.002 +/- 0.004	0.031 +/- 0.010
225 < ΔCO ₂ < 300 [log(AHT)]	Not in model	Not in model	0.027 +/- 0.010
300 < ΔCO ₂ [log(AHT)]	Not in model	-0.009 +/- 0.006	0.022 +/- 0.010
[log(AHT)]Residual standard error	0.0446	0.0445	0.0442

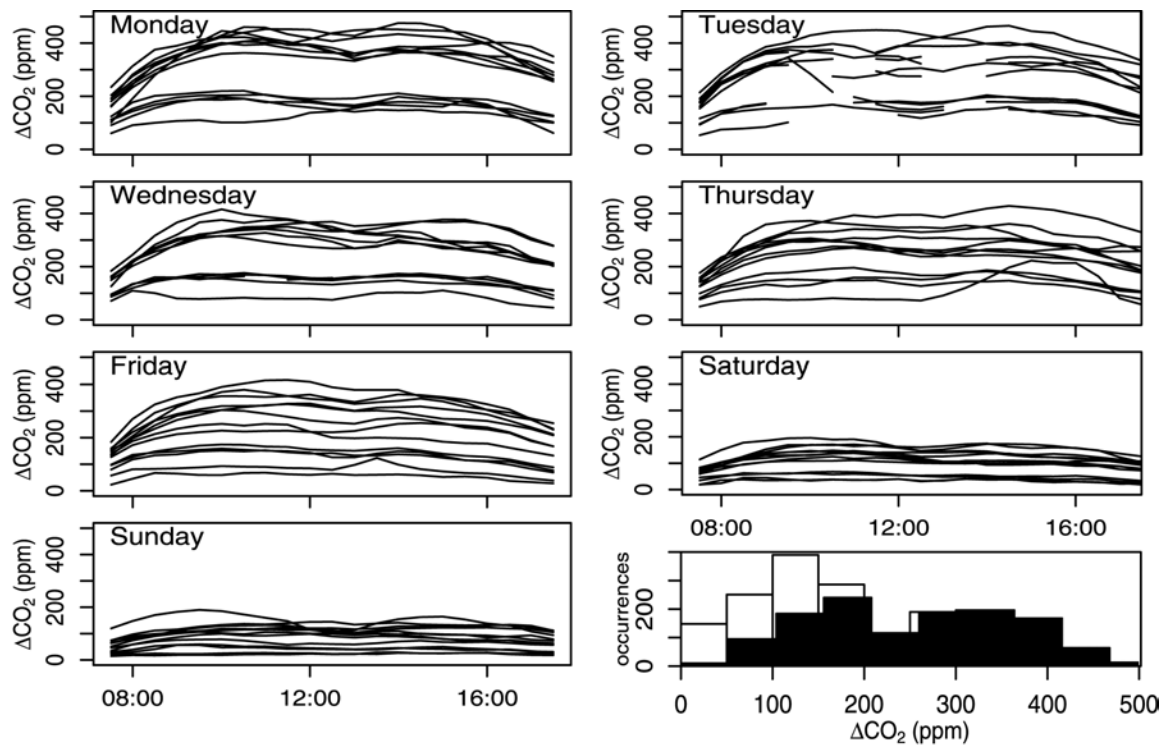


Figure 1. Time trends and histogram for the differences between indoor and outdoor carbon dioxide concentrations. In time trends, each line represents data from a workweek. In the histogram, the upper line in each bin reflects from 7:30 a.m. to 6:00 p.m. of all days while the lower line (shaded section) represents data from the same periods of Monday through Friday.




Figure 2. Indoor minus outdoor carbon dioxide concentrations plotted versus ventilation rates.




Figure 3. Residuals from regression model A, which does not include a carbon dioxide variable, versus indoor minus outdoor carbon dioxide concentration.




Figure 4. Model coefficients for bins of ΔCO_2 concentration, indicating the effect of ΔCO_2 on $\log(\text{AHT})$ with the lowest ΔCO_2 bin used as the reference. The lower and upper plots are results of Model B and Model C, respectively. Horizontal bars indicate ΔCO_2 bin boundaries and vertical error bars represent \pm standard deviation.