

Appendix A: Model description

The airway is assumed to be a cylinder with constant radius r and constant length l (one or both may be implicitly defined by specifying the airway volume $\pi r^2 l$ and/or surface area $2\pi r l$). The model (Eq. 1) results from imposing conservation of (NO) mass throughout this domain. That is, for any time interval $(t, t + \Delta t)$, we assume any net change in mass inside an arbitrary region $(z, z + \Delta z)$ equals the net mass passing through the regions's boundary over the same time interval.

[Figure 1 about here.]

Advection of NO axially through the airway

Over the time interval $(t, t + \Delta t)$, a cross-section of air moving at $v(t)$ cm/s will travel a distance $\Delta z \approx v(t)\Delta t$ cm. Therefore, over a time interval of length Δt the mass that flows into the region through the lower boundary is approximately $v(t)\Delta t\pi r^2 c(z, t)$. Similarly, the mass that flows out of the region through the upper boundary is approximately $v(t)\Delta t\pi r^2 c(z + \Delta z, t)$, hence the net change in mass due to advection (flow) is approximately $-v(t)[c(z + \Delta z, t) - c(z, t)]\pi r^2 \Delta t$

The airway wall as a source of NO

Depending on the concentration, the tissue lining the airway wall can act as either a source or sink for NO. Typically, we assume the concentration of NO in the airway wall c_w exceeds the alveolar concentration $c(z_{\text{alv}}, t)$, implying the airway tissue serves as a net source of NO. The net transfer of NO from tissue to lumen is assumed to occur at a rate proportional to the concentration difference. Denoting the proportionality constant by p , the flux per unit area is $p[c_w - c(z, t)]$. Therefore, over a time interval of length Δt the net mass diffusing into the region from the airway wall is approximately equal to the product $p[c_w - c(z, t)]2\pi r\Delta z\Delta t$.

Diffusion of NO axially through the airway

Assuming $c_w > c(z_{\text{alv}}, t)$ implies that during exhalation the concentration will increase as air moves through the airway, i.e. $c(z + \Delta z, t) > c(z, t)$. According to Fick's first law, the diffusive flux per unit area will be proportional to the concentration gradient $c_z(z, t) := (\partial/\partial z)c(z, t)$. Denoting by d the proportionality constant, over a time interval of length Δt the mass diffusing into the region will be approximately $dc_z(z + \Delta z)\pi r^2 \Delta t$. Similarly, the mass diffusing out of the region will be approximately $dc_z(z, t)\pi r^2 \Delta t$, hence the net change in mass due to (axial) diffusion is approximately $d[c_z(z + \Delta z, t) - c_z(z, t)]\pi r^2 \Delta t$.

Imposing conservation of mass requires that over the time interval $(t, t + \Delta t)$, the net change inside the region must equal the net mass passing through the boundary. This implies the approximate equality $[c(z, t + \Delta t) - c(z, t)]\pi r^2 \Delta z \approx -v(t)[c(z + \Delta z, t) - c(z, t)]\pi r^2 \Delta t + d[c_z(z + \Delta z, t) - c_z(z, t)]\pi r^2 \Delta t + p[c_w - c(z, t)]2\pi r \Delta z \Delta t$. Dividing both sides by $\pi r^2 \Delta t \Delta z$ then taking $\Delta t, \Delta z \rightarrow 0$ results in the model (Eq. 1).

Eq. 2 as a special case of Eq. 1

To transform Eq. 1 into Eq. 2, we begin by assuming the airway NO concentration has achieved a steady-state, and thus does not vary through time. Mathematically, this is equivalent to assuming $(\partial/\partial t)c(z, t) = 0$, which also implies $c(z, t) = c(z)$. Because the concentration is flow dependent, a necessary condition for this to occur is for the velocity to be constant through time, that is $v(t) = v$. As a further simplification, the diffusivity of NO in air is assumed to be negligible, i.e. $d = 0$, reducing the problem further to a first order ordinary differential equation (ODE).

By substituting $(\partial/\partial t)c(z, t) = 0$, $v(t) = v$, and $d = 0$ into equation Eq. 1 it simplifies into $0 = -vc'(z) + (2p/r)c_w - (2p/r)c(z)$. Multiplying through by the airway volume $\pi r^2 l$, dividing through by $\dot{V}l$, and rearranging terms yields $c'(z) + [\text{Daw}_{\text{NO}}/(\dot{V}l)]c(z) = J'_{\text{awNO}}/(\dot{V}l)$. Generically, this is a linear ODE with constant coefficients of the form $c'(z) + [a/d]c(z) = b/d$. Using standard techniques for ODEs, such as separation of variables, the general solution can be shown to be $c(z) = b/a + k \exp(-z[a/d])$, where k is the constant of integration. Therefore, by mak-

ing the above simplifying assumptions, the general solution of (Eq. 1) can be written as $c(z) = J'_{\text{awNO}}/D_{\text{awNO}} + k \exp(-z[D_{\text{awNO}}/(\dot{V}I)])$.

The constant k can be found by imposing a boundary condition, and for consistency with (32) the origin (z_0) is shifted from the sensor to the alveolar boundary. That is, by definition the concentration at the origin is equal to the alveolar concentration, or $c(0) = C_{\text{ANO}}$. Combining this condition with the general solution above allows us to solve for the constant, resulting in $k = C_{\text{ANO}} - J'_{\text{awNO}}/D_{\text{awNO}}$, and hence the particular solution $c(z) = J'_{\text{awNO}}/D_{\text{awNO}} + (C_{\text{ANO}} - J'_{\text{awNO}}/D_{\text{awNO}}) \exp(-z[D_{\text{awNO}}/(\dot{V}I)])$.

This solution is applicable along the length of the airway from the alveolar boundary to the mouth, or $0 \leq z \leq l$. Once air enters the sampling device dead space, both the airway wall concentration and permeability are assumed to be zero. This is equivalent to assuming $c'(z) = 0$ in this region, i.e. there is no spatial variation in the concentration. This implies the concentration measured at the sensor equals the concentration measured at the mouth. Mathematically, this means $c(l) = F_{\text{ENO}}$, and inserting $z = l$ and $c(l) = F_{\text{ENO}}$ into the particular solution yields the equality (Eq. 2).

Appendix B: Numerical integration

Spatial discretizations

To solve the governing PDE (Eq. 1) numerically, the spatial (z) derivatives are replaced with finite difference approximations based on Taylor series expansions. For the diffusive term, a centered three term Taylor series approximation is employed, $(\partial^2/\partial z^2)c(z, t) \approx [c(z - \Delta z, t) - 2c(z, t) + c(z + \Delta z, t)]/[\Delta z]^2$. For the advective term, a biased 4 term Taylor series approximation is employed. The direction of the bias is determined by the direction of flow, as dictated by the sign of $v(t)$. Specifically, the approximation is oriented with an ‘‘upwind’’ bias; two of the terms in the approximation are chosen on the side from which the flow originates, and only one is chosen from the opposite side.

For example, when $v(t) > 0$ the approximation is $(\partial/\partial z)c(z, t) \approx [c(z - 2\Delta z, t) - 6c(z - \Delta z, t) + 3c(z, t) + 2c(z + \Delta z, t)]/[6\Delta z]$. An upwind discretization is employed because centered discretizations

for advection can lead to spurious oscillations in the numerical approximation. A completely one-sided discretization can prevent oscillations; however, despite the formal order of the Taylor series, such an approximation will always have first order accuracy (16). By employing a two-sided, but biased, discretization, higher order accuracy can be achieved while minimizing the potential for oscillatory solutions.

Boundary and initial conditions

For the PDE (Eq. 1) to have a unique solution, initial (time) and inflow boundary (space) conditions must be specified. Moreover, the incoming concentration depends on the direction of flow. During exhalation it corresponds to the alveolar concentration, a parameter to be estimated. During inhalation this concentration is typically the ambient NO level; however, during testing subjects may be provided air that has been “scrubbed” of NO.

By definition, modeling F_{ENO} involves modeling exhalation. However, because respiration is cyclic, the terminal condition in one direction of flow becomes the initial condition for the reverse flow. This relationship means that NO measured during exhalation is determined, in part, by the terminal state of the previous inhalation. In principle, the previous inhalation depends, in turn, on the preceding exhalation, which depends on the inhalation before that, continuing ad nauseam.

In practice, higher flow rates diminish this dependence, and at relatively high rates (300+ ml/s), the terminal airway concentration is effectively independent of the initial. Although 300 ml/s is a relatively rapid rate for exhalation, it is a relatively slow rate for inhalation. For example, in all 900 maneuvers used for simulation this threshold was cleared every time, typically by factors of at least 2-3x.

The implication of this phenomena is that calculating accurate estimates of the airway concentration immediately after inhalation does not require knowing the initial airway concentration when inhalation began. The solution can be calculated based on a simple initial condition (i.e. zero everywhere), and the end result will be essentially unchanged. The terminal condition will also depend on the inflow concentration; however, in the case of “scrubbed” air this can be assumed to be zero.

Appendix C: MCMC parameter estimation

We adopt a Bayesian approach to inference (8), and our goal is to characterize the posterior distribution (generically denoted $f(\theta|x)$ for parameters θ and data x). Using Bayes rule, the posterior can be expressed in terms of the likelihood $f(x|\theta)$ and a prior distribution $f(\theta)$. For many purposes, including ours, the unnormalized posterior is sufficient, simplifying the relationship between the posterior, likelihood, and prior to the proportionality $f(\theta|x) \propto f(x|\theta)f(\theta)$.

Multiple parameterizations of the model (Eq. 1) are possible; for consistency with Eq. 2, we work with the vector of parameters $\theta = (C_{\text{ANO}}, J'_{\text{awNO}}, \text{Daw}_{\text{NO}})$. The data for each subject consists of $x = \{\tilde{c}_{ij}\}$, where \tilde{c}_{ij} is the measured concentration at time t_i during maneuver $j \in 1, 2, \dots, 9$.

To formulate a likelihood, we assume that by fixing θ and solving the corresponding model equation, the model solution can be used to calculate the density of the observed data. If we further assume the observed values \tilde{c}_{ij} arise from a shared parametric conditional distribution with density function f , and that conditional on the model solutions c_{ij} the \tilde{c}_{ij} are independent, then the likelihood can be written as $f(x|\theta) = \prod_i \prod_j f(\tilde{c}_{ij}|c_{ij})$, where θ appears implicitly on the right via the model solution c_{ij} .

When the likelihood $f(x|\theta)$ is combined with a prior $f(\theta)$, the (unnormalized) posterior can be easily calculated for any particular set of parameters θ . To efficiently explore the posterior distribution we employ a Metropolis-Hastings style MCMC algorithm (24), which generically proceeds as follows:

0. Select an initial value θ and calculate the likelihood $f(\theta|x)$.
1. Propose a new value θ' using a transition kernel $q(\theta \rightarrow \theta')$, and calculate the likelihood $f(\theta'|x)$.
2. Accept the proposed value with probability $\min \left[1, \frac{f(x|\theta')f(\theta')q(\theta' \rightarrow \theta)}{f(x|\theta)f(\theta)q(\theta \rightarrow \theta')} \right]$.
3. If the proposal is accepted set $\theta = \theta'$, $f(\theta|x) = f(\theta'|x)$ then return to 1; otherwise, return to 1.

The choice of transition kernel q can have a significant impact on the efficiency of this type of algorithm. Finding an optimal q can be difficult; however, there are a number of more recent MCMC algorithms which incorporate an “adaptive” transition distribution (25). To better account for variability in the posterior across individuals, the adaptive Metropolis algorithm of (12) is employed to automatically calibrate the proposal distribution against the target. This has the dual benefit of both increasing the

efficiency of our sampler, while also simplifying the user experience by largely automating the choice of transition kernel.




Fig. 1: For an arbitrary airway segment of length Δz , and an arbitrary time interval of length Δt , we assume any net change in mass inside the segment must equal the net mass passing through the boundary. As illustrated here, air moves vertically through the airway during exhalation; therefore, the net change due to advection (flow) represents the difference between the mass entering through the lower boundary and exiting through the upper boundary. The airway wall is also assumed to be permeable to NO. Depending on the relative concentrations of the airway wall c_w and the alveolar compartment $c(z_{\text{alv}}, t)$, it will serve as either a net source or sink of NO. Typically, $c_w > c(z, t)$, yielding a net diffusion of NO from the airway wall to lumen (as pictured here). This also implies that the concentration in the lumen will increase as air moves vertically through the airway, i.e. $c(z + \Delta z, t) > c(z, t)$. This concentration gradient results in net NO diffusion opposite to the direction of flow; therefore, the net change in mass due to diffusion will be the difference between the mass diffusing in through the upper boundary and the mass diffusing out through the lower boundary.