Computational Engineering and Technology Innovations

www.ceti.reapress.com

Comput. Eng. Technol. Innov. Vol. 3, No. 2 (2026) 68–88.

Paper Type: Original Article

Exact Predictive Ventilation Control with Causal Occupancy Forecasts and Concentration Guarantees

Omar Mar Cornelio[1], Mingyue Wang[2], Rania A. Mahmoud[3]

[1] Centro de Estudio de Matematica Computacional, Universidad de las Ciencias Inform aticas, La Habana, Cuba; omarmar@uci.cu.
[2] School of computer and information, Lanzhou University of Technology, China; wang.mingyue9811@gmail.com.
[3] Department of Computer Science, Alexandria University, Alexandria, Egypt; rhs00@fayoum.edu.eg.

Citation:

Received: 02 June 2025 *Revised: 08 August 2025* *Accepted: 15 October 2025*	Cornelio, O. M., Wang, M., & Mahmoud, R. A. (2026). Exact predictive ventilation control with causal occupancy forecasts and concentration guarantees. *Computational Engineering and Technology Innovations, 3*(2), 68-88.

Abstract

Predictive ventilation depends on both the mathematical plant and the information available when an airflow command is selected. We study these dependencies in a fully specified, single-zone apartment benchmark. For a well-mixed carbon dioxide balance with nonnegative generation and piecewise constant airflow, we derive an exact propagation map and prove that predicted concentration is jointly convex in the planned airflow sequence after the state variables are eliminated. The resulting finite-horizon problem retains the cubic fan law without successive linearization. A monotone one-step filter then enforces a selected concentration bound under declared occupancy and plant-parameter envelopes. We evaluate seven controllers on twelve independent, paired week-long simulations, using eight separate training weeks for occupancy-model estimation and reactive-controller tuning. Filtered Model Predictive Control (MPC) consumes 26.386 kWh/week of modeled fan-plus-heating energy, compared with 27.726 for the tuned rule and 28.965 for the original 1000/850 ppm rule. Its paired difference from the tuned rule is −1.340 kWh/week, with a pointwise 95% interval of [−1.465, −1.214]. A causal conditional-mean forecast with a five-minute occupancy-report delay yields only 0.045 occupants of one-step Mean Absolute Error (MAE), yet its unfiltered controller spends 44.175 minutes per occupied person-hour above 1000 ppm, predominantly through small exceedances. The filter removes these exceedances in the matched model at an added 0.037 kWh/week. Four paired mismatch tests show why the guarantee must use valid parameter envelopes. The results establish a reproducible separation between forecast quality, energy performance, and conditional concentration invariance; they do not constitute residential field validation or a general indoor-air-quality guarantee.

Keywords: Demand-controlled ventilation, Exact discretization, Convex optimization, Causal occupancy forecasting, Model predictive control, Concentration invariance.

1 | Introduction

Demand-controlled ventilation adjusts outdoor-air supply to a varying indoor load. Its potential energy benefit has long motivated sensor-based control [1], but a defensible comparison must specify what is

Corresponding Author: omarmar@uci.cu
https://doi.org/10.48314/ceti.v3i2.83
Licensee System Analytics. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (http://creativecommons.org/licenses/by/4.0).

controlled, what is measured, and what future information is available. A controller supplied with the realized sequence of future occupancy states answers a different question from one operating on delayed observations. Likewise, a small occupancy prediction error need not imply a small duration above a concentration target.

Carbon dioxide is useful for tracking occupant-generated load within a defined mass balance. It does not summarize all indoor pollutants or determine a universal acceptable indoor environment. The ASHRAE position document distinguished these uses and cautioned against interpreting a single CO_2 concentration as a general air-quality criterion [2]. Occupant generation rates also depended on physiology and activity [3]. We therefore use 1000 ppm as an explicit mathematical design target for the present benchmark. Every concentration guarantee below refers to that target and the stated plant assumptions.

Model Predictive Control (MPC) selects an airflow sequence using predictions, applies its first command, and repeats the calculation after receiving new information. General constrained-MPC theory established the importance of feasibility and the relation between an optimization problem and its closed-loop behavior [4]. Building applications combined predictions of weather and internal loads with energy objectives [5], while ventilation-specific studies addressed robust control [6], probabilistic occupancy forecasting [7], and uncertainty in occupancy-driven operation [8]. These developments make the information contract part of the method rather than a minor implementation detail.

The plant considered here contains a product of airflow and concentration. This bilinearity rules out treating its ordinary state-input formulation as a linear program, but it does not settle the convexity of every equivalent formulation. For nonnegative generation, the exact solution is a nonnegative sum and integral of exponentials of affine functions of planned airflow. We use this representation to establish convexity after eliminating the predicted states. The model thus permits a more precise statement than the generic observation that ventilation dynamics are nonlinear.

We connect this structural result to a complete numerical experiment. A training-only Markov predictor supplies conditional-mean occupancies from observations delayed by one control interval. An oracle supplies realized future occupancy solely as a privileged comparator. A one-step filter uses an occupancy cap rather than future transition events. Reactive and constant-flow comparators expose the influence of baseline tuning and analytically sufficient ventilation rates. The resulting argument moves from exact dynamics to a constrained optimization problem, then to an implemented control comparison and a parameter-mismatch test.

The study establishes a convexity result for this specific scalar balance, an explicit invariant-set filter, and reproducible comparative evidence. It does not propose a universal building model or claim priority for convex building control or safety filtering. Its practical value is to make the assumptions supporting a concentration bound visible and testable. All reported numerical results come from the complete benchmark specified here; no measured dwelling dataset is used.

2|Literature and Research Position

The development of predictive building control involved several distinct questions. Drgoňa et al. [9] organized modeling, optimization, implementation, and performance assessment within a unified building-MPC framework. Killian and Kozek [10] examined the design choices that governed successful energy-efficient applications. Their perspectives supported comparing a complete control architecture, including disturbances, constraints, numerical solution, and baseline design, rather than attributing performance to the label MPC alone. Data-driven predictive-control research further connected model construction with control integration and operational flexibility [11].

Occupancy information introduced another layer. Oldewurtel et al. [12] quantified its role in integrated building climate control. Candanedo and Feldheim [13] estimated occupancy from environmental measurements using statistical learning. Detection of a current state, prediction of a future state, and advance knowledge of realized transitions are different information products. Caballero-Peña et al. [14] reviewed the

estimation and prediction process through its data-acquisition, modeling, evaluation, and testing stages. This process view motivated the explicit temporal boundary used in Section 5: A forecast at decision time may depend on prior occupancy reports and current concentration, but not on later realized transitions.

The ventilation literature also distinguished the interpretation of CO_2 from the choice of controller. Lu et al. [15] examined the relation between indoor concentration and ventilation demand, and subsequently reviewed control principles, sensing, implementation, and performance evaluation across CO_2-based strategies [16]. Persily developed a space-specific CO_2 metric that accounted for occupants, generation, ventilation, and elapsed time [17]. These works showed why the same absolute concentration cannot serve every occupancy and ventilation context. The present model fixes these quantities so that a concentration constraint can be derived exactly; it does not use that simplification to infer a health threshold.

Research on uncertainty adopted richer information models than perturbing a known future trajectory. Li et al. [6] developed tube-based robust MPC for multi-zone demand-controlled ventilation. Zhuang et al. [7] combined probabilistic occupancy forecasts with risk-aware ventilation. Jiang et al. [8] studied occupancy-driven predictive ventilation with energy and infection-related objectives. Their different objectives and uncertainty treatments cannot be reduced to a single error percentage. In particular, a probability model should be assessed by its forecast horizon and calibration, while a deterministic envelope requires a justified range of admissible loads. Our causal forecast supplies a mean, whereas the filter supplies an envelope; these objects perform different mathematical jobs.

Convex control formulations were already used in building operation. Naqvi et al. [18] derived convex steady-state expressions for joint air-quality and temperature management in multi-zone buildings. Andersen [19] and Jradi compared differently weighted predictive ventilation strategies within a component-based framework for a university building. The result developed in Section 4 concerns the entire finite-horizon transient of a scalar, fixed-outdoor-concentration balance with held airflow. It relies on positive source terms and eliminates the concentration states. It is therefore neither a claim that all bilinear building models are convex nor a replacement for the coupled thermal and multi-zone formulations addressed by those studies.

The mathematical distinction between performance optimization and constraint enforcement also had established foundations. Convex optimization supplied the relevant closure rules and optimality interpretation [20]. Wabersich and Zeilinger [21] formulated a predictive safety filter that modified a proposed input when required by constraints. Ames et al. [22] expressed safety through forward invariance and control-barrier conditions. Our filter is a scalar specialization: monotonicity reduces its correction to a bounded one-dimensional root search. It certifies only the modeled concentration set, under the specified occupancy, parameter, and state-information bounds. The word guarantee has that conditional meaning throughout the paper.

Measurement quality remained consequential even when the plant equation was simple. Fisk et al. [23] found substantial errors in deployed CO_2 sensors and documented their relevance to ventilation control. Peng and Jimenez [24] derived CO_2-related infection-risk proxies whose interpretation varied with environment and activity. We do not equate these risk proxies with the present concentration target. Recent synthesis by Rizo-Maestre et al. [25] considered intelligent ventilation through the integration of sensing, control, and indoor-environment assessment. These studies motivate keeping sensor uncertainty and other pollutants outside a guarantee until they have been included explicitly.

Simulation benchmarking developed its own evaluation requirements. Page et al. [26] modeled occupancy as an inhomogeneous Markov process with periods of absence, while Richardson et al. [27] developed a high-resolution domestic occupancy model for energy-demand simulation. These studies justified stochastic occupancy as a useful modeling approach; they did not calibrate the numerical matrices selected here. BOPTEST provided repeatable building emulators, baseline controls, and common performance indicators for comparing algorithms [28]. Field implementation by Blum et al. [29] demonstrated the additional work

required to evaluate MPC in an operating office HVAC system. The present benchmark favors transparent analytical structure over emulator fidelity and remains several validation steps short of a field deployment.

Finally, simulation evidence depends on experimental organization. Morris et al. [30] emphasized explicit aims, data-generating mechanisms, estimands, methods, and performance measures. We separate the weeks used to estimate the forecast and tune the reactive baseline from those used for final comparison. Paired week-level differences use common disturbances and retain the week as the independent unit. Their intervals are pointwise; a family of confirmatory hypothesis tests would require a declared multiplicity procedure, such as Holm's [31] sequential method. No collection of five-minute observations is treated as thousands of independent replications.

Table 1 compares the relevant modeling choices, while *Table 2* organizes selected developments into overlapping phases. The phases are an interpretive reading of the literature, not an exhaustive historical classification. The gap addressed here is a model-specific connection between exact transient convexity, causal information, and a directly verifiable concentration filter. Section 8 identifies the completed result associated with each part of that gap.

Table 1. Representative research strands, mathematical questions, and scope boundaries.

Studies	Model or Method	Information or Uncertainty	Relation to the Present Result
Building MPC [5], [9–11]	Predictive climate and energy control	Weather, internal loads, model choice	Establish the broader design and implementation context
Occupancy estimation and prediction [12–14]	Measurement and statistical prediction	Available sensor data and prediction horizon	Motivate an explicit decision-time information set
CO_2-based ventilation [1], [15–17]	Concentration-based demand response	Generation, ventilation, measurement time	Define the physical meaning of the controlled variable
Robust and probabilistic ventilation [6–8]	Tube-based or risk-aware control	Bounded uncertainty or probability forecasts	Distinguish an envelope from a conditional mean
Convex and multi-objective control [18–20]	Convex expressions and energy–air-quality objectives	Model-specific constraints and weights	Precede the finite-horizon scalar convexity proof
Constraint filtering [21], [22]	Predictive safety and invariant sets	Explicit state and model bounds	Support separating proposed and admissible commands
Occupancy and control benchmarks [26–29]	Stochastic profiles, emulators, field operation	Repeated disturbances and physical validation	Locate this transparent simulation within a larger validation path
Measurement and evidence [2], [3], [17], [23–25], [30], [31]	CO_2 interpretation, sensing, simulation assessment	Sensor error and inferential scope	Limit the conclusions that energy and threshold results support

Table 2. Key developments relevant to the benchmark across overlapping modeling, sensing, and control timelines.

Period	Representative Development	Evaluation Consequence
1998–2008 [1], [4], [20], [26], [27]	Sensor-based ventilation, constrained control, stochastic presence	Specify plant dynamics and exogenous load generation
2010–2017 [3], [5], [11], [13], [22], [23]	Occupancy-informed control and measurement interpretation	Distinguish observations, forecasts, and physical uncertainty
2019–2023 [6–9], [11], [15–18], [21], [24], [28–30]	Robust ventilation, probabilistic forecasts, repeatable evaluation	Connect control guarantees with a declared uncertainty model
2024–2026 [14], [19], [25]	Integrated prediction pipelines and intelligent ventilation	Evaluate the complete information and control architecture

3 | Physical Model and Concentration Envelopes

Consider a well-mixed zone of volume $V = 90$ m³. Indoor and outdoor concentrations are $C(t)$ and $C_{out} = 420$ ppm. Define excess concentration $x = C - C_{out}$, infiltration $a = 0.3$ h⁻¹, mechanical rate $0 \le u \le 3$ h⁻¹, expressed in Air Changes per Hour (ACH), and occupancy $0 \le N \le 2$. With fixed generation $G = 192$ ppm h⁻¹ per person, the balance is

$$\dot{x}(t) = GN(t) - [a + u(t)]x(t), \quad x(0) = 280 \text{ ppm}. \tag{1}$$

The generation coefficient corresponds to 0.0048 L s⁻¹ per person in the stated volume. It is a concentration-rate coefficient, not a volumetric flow. All rates in *Eq. (1)* use hours, and all control intervals have length $\Delta = 1/12$ h. Occupancy and airflow are held constant within each interval. Outdoor concentration is fixed; window opening, inter-zone transport, deposition of other pollutants, and thermal-state dynamics are not represented.

For $q = GN$ and $r = a + u > 0$, define the exact held-input map

$$\Phi(x, u; q, a) = e^{-r\Delta}x + q\,\phi(r), \quad \phi(r) = \frac{1 - e^{-r\Delta}}{r}. \tag{12}$$

Thus $x_{k+1} = \Phi(x_k, u_k; GN_k, a)$. The implementation evaluates $1 - e^{-r\Delta}$ through the numerically stable expression $-\mathrm{expm1}(-r\Delta)$. Unlike forward Euler, *Eq. (2)* introduces no integration approximation when the inputs satisfy the stated hold convention. The within-interval solution moves monotonically between its initial value and its equilibrium, so its maximum is attained at an interval endpoint.

Proposition 1. Let $a > 0$, $G \ge 0$, $0 \le N(t) \le N_{max}$, and $x(0) \ge 0$. For a constant mechanical rate $u_c \ge 0$, the first line below bounds the state. The second line

$$0 \le x(t) \le \max\left\{x(0), \frac{GN_{max}}{a + u_c}\right\}, \qquad u_c \ge \left[\frac{GN_{max}}{L} - a\right]_+ \tag{3}$$

is sufficient to preserve $x(t) \le L$ whenever $x(0) \le L$, and is necessary for that property over all constant occupancies up to N_{max}.

Proof: Variation of constants gives a nonnegative solution. Replacing $GN(t)$ by its upper bound yields an exponential interpolation between $x(0)$ and $GN_{max}/(a + u_c)$, proving the envelope. If the second inequality holds, the vector field at $x = L$ is nonpositive for every admissible occupancy. Conversely, if it fails, constant occupancy N_{max} has equilibrium above L and a trajectory initialized at L immediately increases.

For $L = 1000 - 420 = 580$ ppm, the smallest constant rate satisfying this uniform condition is

$$u_{rob} = \frac{384}{580} - 0.3 = 0.3620689655 \text{ h}^{-1}. \tag{4}$$

A 0.5 ACH schedule gives the stronger ceiling $C \le 900$ ppm from the stated initial condition. It therefore supplies more mechanical ventilation than is needed for the selected 1000 ppm target under this occupancy cap. This analytical difference matters when comparing its energy with a controller that operates close to 1000 ppm. It does not imply that 0.362 ACH meets any external ventilation requirement.

Forward Euler would give $x_{k+1} = [1 - \Delta(a + u_k)]x_k + \Delta GN_k$. Nonnegative interpolation requires $0 \le \Delta(a + u_k) \le 1$; the benchmark's upper value is 0.275. Euler therefore preserves the corresponding constant-rate envelope here, but it still changes finite-time trajectories. We use the exact map throughout the new calculations to keep integration error separate from forecast and control effects.

Modeled fan and sensible-heating power are

$$P_{fan}(u) = 75(u/3)^3, \qquad P_{heat}(u, T) = 30.15[0.3 + 0.3u]\,[21 - T]_+ \tag{5}$$

Both expressions are in watts. The coefficient $30.15 = 1206 \times 90/3600$ converts hourly air changes to a sensible-heat flow, and a fixed heat-recovery efficiency of 0.7 leaves a mechanical-heating multiplier of 0.3. The infiltration term is independent of the controller. The energy measure is fan plus this modeled heating, not whole-home consumption or a coupled building heating simulation. Cubic fan power and constant recovery efficiency are declared equipment assumptions.

Table 3. Plant and computational parameters. Concentration is in ppm and time in hours unless stated otherwise.

Quantity	Value	Role
Zone and outdoor level	$V = 90$ m³; $C_{out} = 420$	Fixed well-mixed volume and background
Initial state	$C_0 = 700$	Shared initialization in every week
Generation and infiltration	$G = 192$; $a = 0.3$ h⁻¹	Matched-model plant
Occupancy and actuation	$N \in \{0,1,2\}$; $u \in [0,3]$ h⁻¹	Load cap and actuator limit
Decision interval	$\Delta = 1/12$ h	Five-minute held inputs
Predictive horizon	$H = 12$ steps	One-hour optimization
Concentration bounds	Evaluation target 1000; optimizer target 999.999	Separate mathematical target and numerical margin
Energy law	Fan 75 W at 3 ACH; recovery 0.7; setpoint 21 °C	Fixed fan and sensible-heating model
Independent weeks	8 training; 12 testing	Training-only estimation and tuning
Parameter stress	Actual $a = 0.2$; actual $G = 230.4$	Four paired test weeks

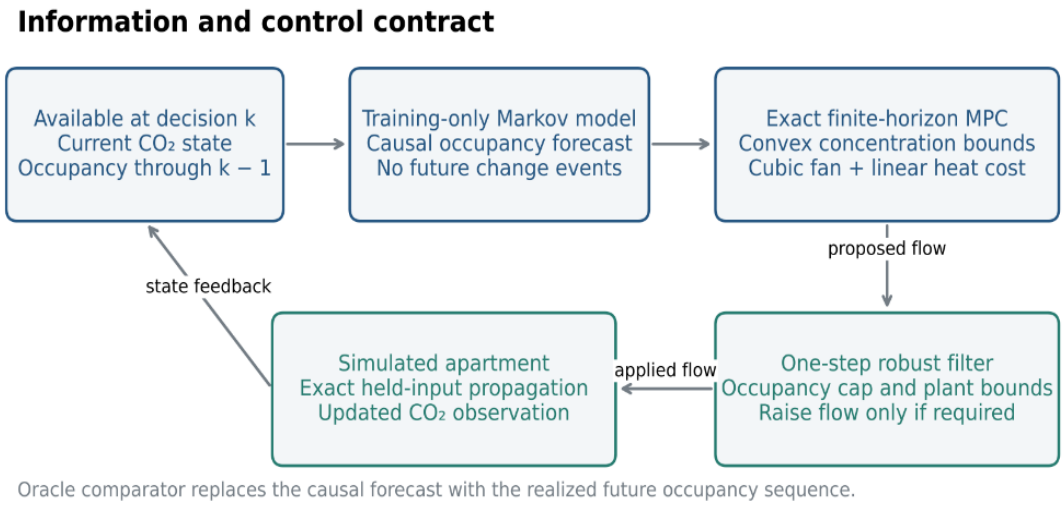
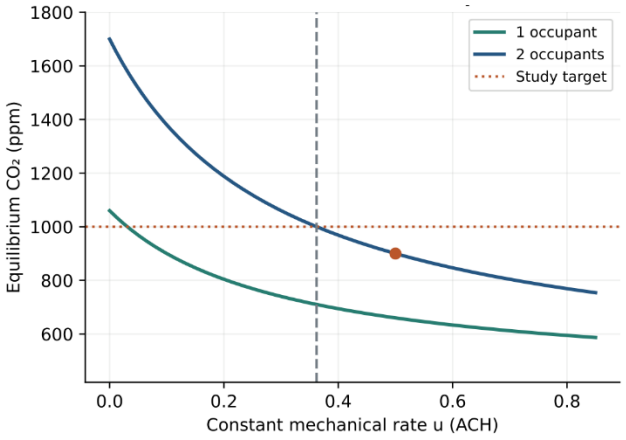
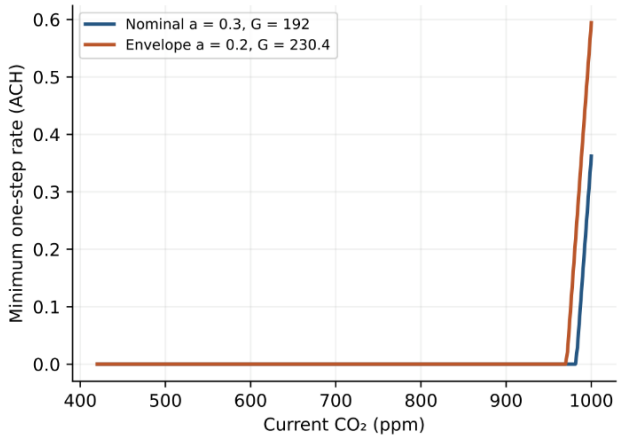


Fig. 1. Implemented information and control framework: Causal occupancy forecast, concentration-based filtering with load envelope, and oracle comparator with access to future occupancy.



a.



b.

Fig. 2. Analytical implications of the exact balance; a. equilibrium concentrations for one and two occupants, with nominal robust airflow rate, b. minimum one-step airflow rate versus current concentration and plant-envelope conditions.

4 | Exact Predictive Optimization and Constraint Filtering

At decision time k, let $\widehat{N}_{k+j|k} \geq 0$ be a fixed occupancy forecast over H intervals. Consider a planned airflow vector $u = (u_0, \ldots, u_{H-1})$. Once the forecast and initial state are fixed, the exact concentrations are functions of this vector. The following property permits optimization without a linearized flow–concentration product.

Proposition 2. Under *Eq. (1)*, with nonnegative initial excess concentration and nonnegative prescribed generation, each finite-horizon predicted excess concentration is jointly convex and nonincreasing in u.

Proof: Let $\ell_j(s,t)$ be the length of the intersection of $[s,t]$ with control interval j. Variation of constants at a prediction node gives

$$K(s,t;u) = \exp\left[-a(t-s) - \sum_j u_j \ell_j(s,t)\right].$$
$$x(t;u) = x_0 K(0,t;u) + \int_0^t q(s) K(s,t;u)\, ds. \qquad (6)$$

Every exponent is affine in u. Its exponential is convex; multiplication by a nonnegative coefficient, integration, and addition preserve convexity [20]. Each ℓ_j is nonnegative, so increasing any airflow component cannot increase *Eq. (6)*. This proves both statements.

The proof concerns the state-eliminated exact prediction. It does not assert joint convexity of a formulation that introduces independent state variables and equality constraints involving products. It also depends on the positivity of the initial excess and generation terms. Coupled thermal dynamics, variable outdoor concentration, and additional pollutant mechanisms require separate analysis.

The implemented finite-horizon objective is

$$\min_{u} \ \frac{\Delta}{1000}\sum_{j=0}^{H-1}\left[75\left(u_j/3\right)^3+\beta_{k+j}u_j\right]. \tag{7}$$

$$\text{Subject to} \quad 0 \le u_j \le 3, \qquad x_{j+1}(u) \le L_{opt},$$

where $\beta_{k+j} = 9.045\left[21 - T_{k+j}\right]_+$ and $L_{opt} = 579.999$ ppm. The constant infiltration-heating term is omitted from optimization and restored when reporting total modeled energy. The fan objective is convex on the nonnegative domain, and Proposition 2 makes each concentration sublevel constraint convex. Consequently, *Eq. (7)* is a convex optimization problem. It is not a linear program, and no terminal penalty, slack variable, or terminal-set constraint is used in this benchmark.

The numerical solver is Sequential Least-Squares Programming (SLSQP), using exact analytical first derivatives. For clarity, the required one-step derivatives are

$$\frac{\partial \Phi}{\partial x} = e^{-r\Delta}, \qquad \frac{\partial \Phi}{\partial u} = -\Delta e^{-r\Delta}x + q\phi'(r), \qquad \phi'(r) = \frac{\Delta r e^{-r\Delta} - \left(1 - e^{-r\Delta}\right)}{r^2}. \tag{8}$$

The horizon Jacobian follows by the chain rule. The implementation shifts the previous optimal sequence to initialize the next solve, raises a predicted-infeasible initial sequence componentwise to at least 0.7 ACH, and uses tolerance 10^{-9} with at most 60 iterations. The objective is scaled by a positive constant to improve numerical conditioning; scaling does not alter its minimizer. A failed solve or an accepted-sequence concentration excess greater than 10^{-5} ppm triggers a full-flow fallback and is counted explicitly. Convexity characterizes the mathematical problem; solver status and a feasibility check alone are not a numerical dual certificate of optimality.

A conditional-mean occupancy forecast constrains the model's mean-load trajectory. It does not restrict every realization. A separate command filter can enforce a deterministic load envelope. With current state x and upper generation $q_{max} = G_{max}N_{max}$, define

$$u_{min}(x) = \min\{u \in [0,3]: \Phi(x, u; q_{max}, a_{min}) \le L\}, \qquad u^F = \max\{u^P, u_{min}(x)\}. \tag{9}$$

Here u^P is the first predictive command, $L = 580$, and u^F is applied to the plant. If zero airflow meets the inequality, $u_{min} = 0$. Otherwise, monotonicity permits bisection; the implementation uses 35 iterations. This filter specializes the separation of desired and constraint-compatible inputs used in safety-filter and barrier-function methods [21, 22]. It does not optimize the full horizon again after changing the command.

Proposition 3. Suppose $0 \le x_k \le L$, actual generation satisfies $0 \le q_k \le q_{max}$, actual infiltration satisfies $a \ge a_{min} > 0$, and

$$q_{max} \le (a_{min} + u_{max})L, \quad u_{max} = 3. \tag{10}$$

Then the set in *Eq. (9)* is nonempty. With the exact current state and held inputs, the filtered command preserves $0 \le x(t) \le L$ during the next interval, for any proposed command in $[0,3]$.

Proof. By Proposition 1, maximum flow preserves the interval under the worst admissible source and infiltration, so it is feasible in *Eq. (9)*. Increasing generation increases the exact solution. Increasing removal

decreases it when the state and source are nonnegative. The actual next state is therefore no greater than the worst-envelope prediction, which is at most L. Within the interval the scalar solution is monotone toward its equilibrium, so neither endpoint below L can conceal an interior overshoot. Positivity follows from *Eq. (2)*. Repetition proves invariance for subsequent intervals whenever the same assumptions remain valid.

For the nominal envelope, *Eq. (10)* holds with a wide actuator margin. For the stress envelope, $a_{min} = 0.2$ and $G_{max} = 230.4$, its smallest robust constant rate becomes

$$u_{rob,env} = \frac{2(230.4)}{580} - 0.2 = 0.5944827586\ h^{-1}. \quad (11)$$

A state observation with a certified upper error bound can replace x by its upper estimate in *Eq. (9)*. The guarantee then remains one-step conditional on the upper estimate being valid and the resulting filter problem being feasible. An unbounded Gaussian sensor model cannot provide a deterministic error envelope merely by choosing a multiple of its standard deviation. The numerical campaign below deliberately uses exact concentration observations to isolate the occupancy-information question; the bounded-error extension is analytical, not a sensor-validation experiment.

Corollary 1. If current-interval occupancy is known exactly, the plant is matched, and the first hard predicted constraint is satisfied, receding-horizon operation preserves the concentration bound independently of errors at later forecast horizons, provided a feasible command remains available.

The current predicted and actual one-step balances then coincide. Proposition 1 supplies a feasible maximum-flow command for bounded occupancy, and induction applies. This corollary explains why the timing of an occupancy report matters: the implemented causal experiment withholds the current interval's occupancy until the next decision. A test that reveals it immediately would have a different constraint contract.

5 | Reproducible Benchmark Design

The numerical campaign uses eight training weeks and twelve independent test weeks. Each week contains 2016 five-minute intervals, begins at 700 ppm, and uses the same occupancy and outdoor-temperature realization across all seven controllers. Exact concentration measurement is available at each decision. Occupancy reports are delayed by one interval for the causal predictor. There is no sensor noise, actuator lag, open-window event, or weather-forecast error in this campaign. These idealizations are shared and declared rather than hidden inside a forecast-error parameter.

Occupancy is a three-state, time-inhomogeneous Markov chain. At interval k, the generator draws N_k from the row indexed by N_{k-1} in the matrix for the current clock-time regime. The preceding state is initialized as $N_{-1} = 2$. Table 4 provides all transition probabilities. Their values create an illustrative household-like daily pattern; they were not estimated from a dwelling survey and are not claimed to reproduce a population distribution. Markov occupancy models have an established role in building simulation [26], [27], but that methodological precedent does not validate these particular probabilities.

Table 4. Complete five-minute occupancy transition matrices.

Regime	Row for State 0	Row for State 1	Row for State 2
00:00–06:00	(0.960, 0.030, 0.010)	(0.010, 0.965, 0.025)	(0.001, 0.003, 0.996)
06:00–09:00	(0.985, 0.010, 0.005)	(0.030, 0.960, 0.010)	(0.025, 0.020, 0.955)
09:00–17:00	(0.992, 0.006, 0.002)	(0.035, 0.955, 0.010)	(0.025, 0.020, 0.955)
17:00–24:00	(0.950, 0.030, 0.020)	(0.005, 0.965, 0.030)	(0.001, 0.004, 0.995)

*Rows and columns use state order 0, 1, 2. Every row sums to one. Clock-time regimes repeat each day.

The predictor estimates four matrices from training transitions only. A pseudocount of one is added to every cell before row normalization. At time k, its initial probability vector is the point mass at the last reported state, N_{k-1}. For forecast interval $k + j$, the propagation and conditional mean are

$$p_{k+j|k} = e_{N_{k-1}} \prod_{\ell=0}^{j} \widehat{P}_{r(k+\ell)}, \quad \widehat{N}_{k+j|k} = p_{k+j|k}(0,1,2)^{\top}. \tag{12}$$

The product follows increasing time from left to right, using row probability vectors. The initial decision uses the known preceding state 2. A forecast can therefore assign positive mass to an arrival without knowing whether that arrival actually occurs. The oracle arm replaces these means with the realized $N_k, \ldots, N_{k+H-1}$. It is an information comparator, not a deployable estimator or a certified lower bound on weekly energy, since receding-horizon operation need not solve the full-week optimum.

The outdoor-temperature sequence is generated hourly and linearly interpolated to decision nodes:

$$T_h = 6 + 3\sin\left(\frac{2\pi(h-8)}{24}\right) + 1.5\sin\left(\frac{2\pi h}{60}\right) + \zeta_h, \qquad \zeta_h \sim \mathcal{N}(0, 0.8^2). \tag{13}$$

Hourly indices run from 0 to 169, providing the terminal forecast extension. Temperature is treated as known over each predictive horizon. Each week uses a NumPy SeedSequence with two spawned streams for occupancy and temperature. Training seeds are 60601 through 60608; test seeds are 60701 through 60712. Occupancy generation includes twelve extra intervals for terminal oracle forecasts. The code and all week-level outputs accompany the manuscript in the computational package.

The original reactive comparator commands 3 ACH when current concentration is above 1000 ppm, commands zero below 850 ppm, and otherwise retains its previous state. The tuned reactive comparator is selected from 30 combinations: active flow in {0.75, 1, 1.5, 2, 3} ACH and upper switching threshold in {875, 900, 925, 950, 975, 1000} ppm, with a 100 ppm hysteresis width. Among configurations with no target crossing in any training week, the one with the smallest mean modeled energy is locked before test evaluation. This procedure selects 1 ACH with thresholds 975/875 ppm. Its tuning is finite-grid selection, not a claim of optimality over every reactive policy.

The fixed 0.5 ACH and nominal robust-constant arms provide two different reference margins. The three predictive arms share *Eq. (7)*, solver settings, exact plant propagation, and current concentration. They differ in occupancy information or in whether the one-step filter is applied. *Table 5* records the contracts. The filter's operation changes subsequent states, so its MPC optimization is rerun on its own trajectory; filtered performance is not calculated by modifying a saved command sequence after the experiment.

Table 5. Controller definitions and information contracts.

Controller	Decision Rule or Occupancy Input	Target Mechanism	Interpretation
Rule 1000	3 ACH; thresholds 1000/850 ppm	Reactive switching after upper crossing	Untuned high-flow reference
Tuned rule	1 ACH; thresholds 975/875 ppm	Training-only threshold selection	Feasible in training; no general invariance proof
Fixed 0.5	Constant 0.5 ACH	Analytical 900 ppm ceiling	Conservative constant comparator
Fixed robust	Constant 0.3620689655 ACH	Proposition 1 at 1000 ppm	Minimal uniform constant-flow condition
Oracle MPC	Realized future occupancy	Hard predicted constraints	Privileged information comparator
Causal MPC	Training-only conditional means	Hard mean-load constraints	No realization-wise occupancy guarantee
Filtered MPC	Same causal forecast; filtered first command	Occupancy-cap one-step condition	Proposition 3 under the declared envelope

*The mathematical guarantees refer to the matched model and exact concentration observation.

Energy is accumulated from held power over each interval. Mechanical-ventilation energy contains fan power plus recovery-adjusted mechanical heating; total modeled energy additionally includes infiltration heating. The latter term is identical across nominal controllers within a week. Its subtraction therefore leaves every paired energy difference unchanged, although it changes percentage savings and explains why total-energy percentages are smaller than mechanical-energy percentages.

Threshold exceedance is evaluated in continuous time within each interval, using the exact crossing of the exponential solution. The person-weighted outcome is

$$E_C = 60\,\frac{\sum_k N_k\,\lambda_k(C > C_{max})}{\sum_k N_k\,\Delta}, \tag{14}$$

where λ_k is the duration in hours above the threshold during interval k. The units are minutes per occupied person-hour, and the range is 0 to 60. A two-person interval contributes twice the exposure of a one-person interval.

This definition is distinct from dividing unweighted room exceedance minutes by person-hours. Maximum concentration is the maximum of all interval endpoints, which is also the continuous maximum under the hold assumptions.

For a paired energy contrast between filtered MPC and comparator b, define $d_r = E_r^F - E_r^b$. The interval reported for its mean is

$$\bar{d} \pm t_{0.975,11}\frac{s_d}{\sqrt{12}}. \tag{15}$$

These are pointwise t intervals across independent simulated weeks. The primary comparison is filtered MPC versus the training-tuned rule; other contrasts are descriptive. The intervals describe variability under the declared generator and do not quantify uncertainty across real dwellings, weather regions, or physical models. Five-minute observations are not substituted for the twelve independent weeks.

6 | Nominal Results and Forecast Interpretation

The implementation uses Python 3.12.14, NumPy 2.3.5, and SciPy 1.17.0 with double-precision arrays. The nominal campaign executes 72,576 horizon optimizations across the three predictive arms.

All return successful solver status and satisfy the declared accepted-sequence feasibility check; no full-flow fallback is invoked. An independent central-difference check of the analytical concentration Jacobian gives a maximum discrepancy of 4.08×10^{-9}. Two hundred random convex-combination checks also satisfy the predicted convexity inequality. These checks support implementation correctness; *Proposition 2* supplies the mathematical proof.

Fig. 3 shows the first 48 hours of test week 60701, selected by its position in the locked test list. It illustrates why controller labels alone are insufficient. The tuned rule cycles at a lower active flow than the original 3 ACH rule. The oracle anticipates future occupancy, whereas the causal forecast responds to the preceding interval's report and the learned clock-dependent transition probabilities. The filtered causal controller raises a proposed flow when the occupancy-cap condition requires it.

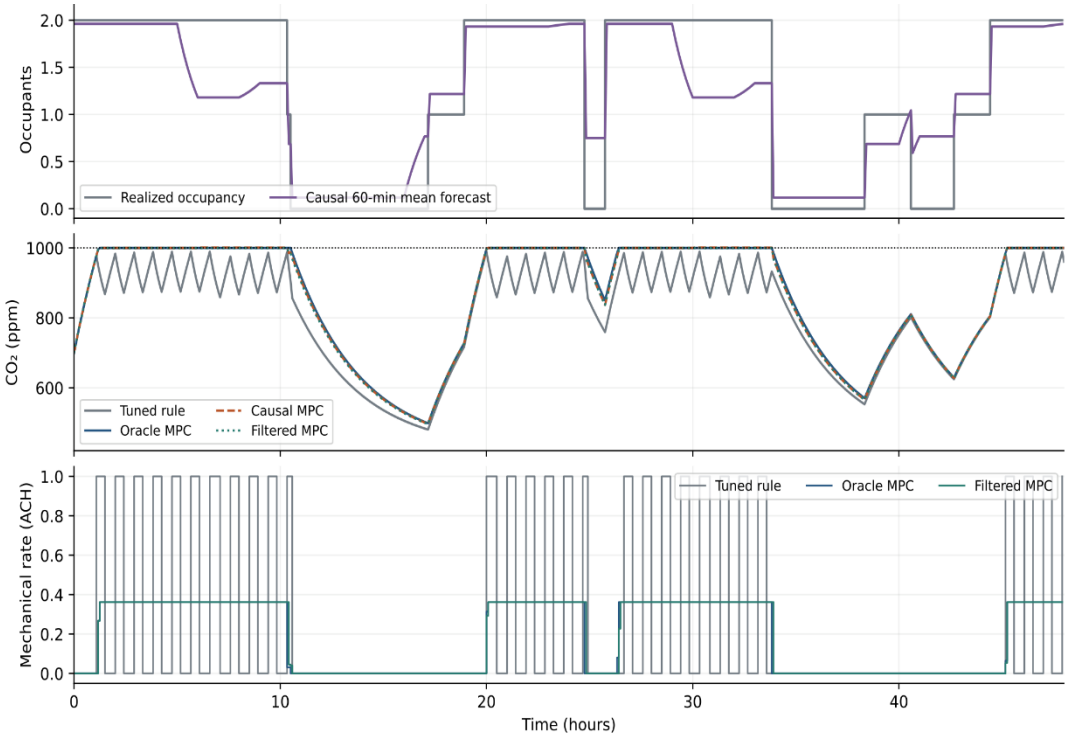


Fig. 3. First 48 hours of test week 60701 (Occupancy and causal 60-minute mean forecast, concentration, and airflow. Exact current concentration and 5-minute occupancy-report delay. Illustrative trajectory; 12-week results in Tables 6–7).

Table 6 reports energy outcomes. Filtered MPC uses 26.3863 kWh/week, compared with 27.7260 for the tuned rule and 28.9649 for the original rule. Its paired mean difference from the tuned rule is −1.3397 kWh/week, with interval [−1.4655, −1.2140]. The ratio of the mean reduction to the tuned-rule mean is 4.832%. Against the original rule that ratio is 8.903%. Baseline tuning therefore reduces the apparent predictive-control advantage substantially, while leaving a positive modeled benefit in this testbed.

Table 6. Mean energy and airflow over twelve independent test weeks.

Controller	Total energy (kWh/week)	Mechanical energy (kWh/week)	Fan energy (kWh/week)	Mean rate (ACH)
Rule 1000	28.9649	6.2343	0.89635	0.21342
Tuned rule	27.7260	4.9954	0.09145	0.19597
Fixed 0.5	34.1543	11.4236	0.05833	0.50000
Fixed robust	30.9828	8.2522	0.02215	0.36207
Oracle MPC	26.3357	3.6051	0.00853	0.14292
Causal MPC	26.3490	3.6184	0.00871	0.14350
Filtered MPC	26.3863	3.6557	0.00860	0.14516

* Total energy includes fan plus modeled mechanical and infiltration heating. Mechanical energy excludes infiltration heating.

The common mean infiltration-heating contribution is 22.7306 kWh/week. The fixed robust rate reduces total modeled energy by about 9.29% relative to the 0.5 ACH schedule while preserving the same 1000 ppm target. Thus part of the comparison with the 0.5 ACH schedule is a comparison of different concentration margins. Filtered MPC uses 14.836% less total modeled energy than the robust constant comparator, reflecting occupancy adaptation beyond constant robust sizing. It uses 22.744% less than the 0.5 ACH schedule, but that larger percentage should be read alongside the latter's 900 ppm ceiling.

The cubic fan law also explains why a lower average flow does not alone determine fan energy. The original rule uses a relatively small mean rate through full-flow bursts and incurs much greater fan energy than the

smoother predictive commands. Jensen's inequality implies that a constant rate minimizes cubic fan power among trajectories with the same arithmetic mean, but the compared controllers do not share that mean or the same concentration trajectory. The fan component therefore supports a mechanism-level explanation, not a universal ranking of scheduled and predictive ventilation.

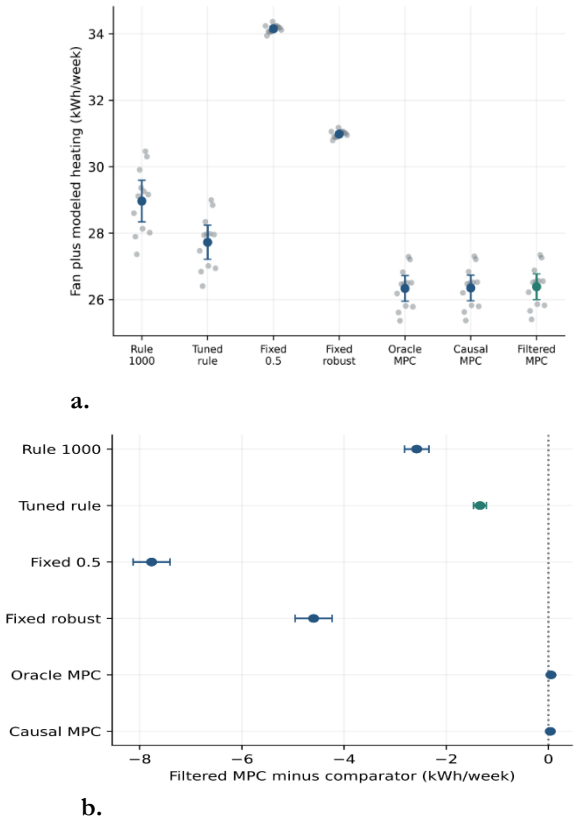


Fig. 4. Nominal energy performance; a. weekly totals, means, and 95% confidence intervals, b. paired mean energy differences (filtered MPC minus comparator), with 95% confidence intervals. Negative values favor filtered MPC; tuned rule as primary comparison.

Table 7 shows a different ordering for threshold exceedance. The fixed schedules, oracle MPC, and filtered MPC have zero exceedance in every tested week. The tuned rule also remains below 1000 ppm, with a maximum across test weeks of 992.866 ppm. Its empirical success is broader than its training feasibility but remains a finite test rather than a uniform proof. The original rule records 1.295 minutes per occupied person-hour above 1000 ppm, with a maximum of 1017.042 ppm across weeks.

Table 7. Person-weighted exceedance and peak concentration across 12 test weeks: Pointwise t-intervals for mean exceedance and maximum campaign concentration.

Controller	Mean Exceedance (Min/Person-Hour)	Pointwise 95% Interval	Largest CO_2 (ppm)
Rule 1000	1.295	[1.223, 1.367]	1017.042
Tuned rule	0	[0,0]	992.866
Fixed 0.5	0	[0,0]	899.999
Fixed robust	0	[0,0]	999.992
Oracle MPC	0	[0,0]	999.999
Causal MPC	44.175	[42.489, 45.862]	1016.316
Filtered MPC	0	[0,0]	1000.000

The unfiltered causal controller combines a low one-step occupancy error with frequent small threshold exceedances. Its one-step Mean Absolute Error (MAE) is 0.0453 occupants and its Root-Mean-Square Error (RMSE) is 0.1690. At 30 and 60 minutes, the MAEs increase to 0.2130 and 0.3274. These values describe predictions based on delayed reports and training-only transition estimates. They do not establish a probability guarantee for occupancy or concentration.

The duration outcome is particularly sensitive to operating close to a hard boundary. Causal MPC records 44.175 minutes per occupied person-hour above 1000 ppm, but only 0.190 above 1005 ppm and 0.093 above 1010 ppm; it never reaches 1020 ppm. Its maximum excess over the target is therefore small even though the duration above the exact target is large. Figure 5 displays both forecast error and threshold sensitivity. Reporting only duration would conceal the small amplitude, while reporting only peak would conceal persistence near the target.

The filter removes these nominal exceedances with a mean energy increment of 0.03725 kWh/week relative to causal MPC. The paired interval is [0.03085, 0.04365], and the increment is 0.1414% of causal MPC's mean total modeled energy. Relative to the oracle comparator, filtered MPC uses an additional 0.05055 kWh/week, or 0.1919%. The filter modifies the first command in 49.67% of intervals, on average. Frequent intervention need not be energetically large when the required corrections are small near the concentration boundary.

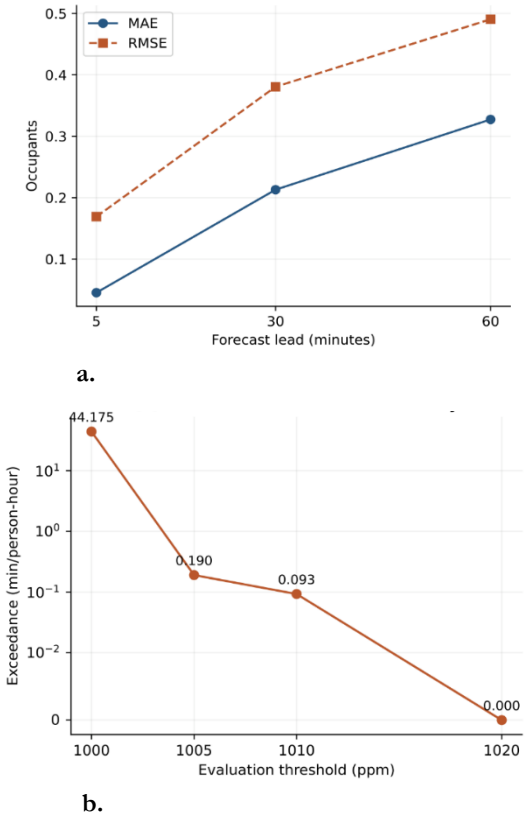


Fig. 5. Causal forecast diagnostics and exceedance analysis; a. Occupancy MAE and RMSE at three lead times, b. Exceedance under varying evaluation thresholds using a symmetric logarithmic scale (Fixed causal-MPC trajectories).

An oracle-informed trajectory is therefore informative about the value of future information, but it cannot validate an online forecasting pipeline. Conversely, a conditional-mean predictor can perform well on a conventional forecast metric while failing a realization-wise concentration target. In this model, Corollary 1 shows how immediate current-occupancy observation would change that conclusion. The implemented five-minute reporting delay and the filter's independent occupancy cap make the information boundary observable in the numerical results.

7 | Parameter Mismatch and Conditional Guarantees

A concentration guarantee is only as broad as its plant envelope. We test this dependence by lowering actual infiltration from 0.3 to 0.2 ACH and increasing actual per-person generation by 20%, from 192 to 230.4 ppm h^{-1}. The first four test weeks supply the same occupancy and temperature realizations to each stress-test arm. The predictive optimizer remains nominal. Only the envelope-aware filter uses the changed lower-infiltration and upper-generation bounds. The comparison therefore isolates the protection supplied by valid bounds rather than refitting the predictive model.

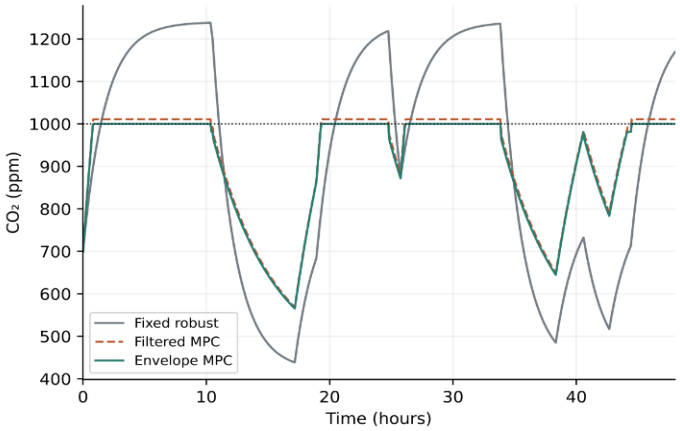
Four paired policies are evaluated: the nominal robust constant rate, the envelope robust constant rate from (11), nominal filtered MPC, and envelope-filtered MPC. The stress campaign executes a further 16,128 predictive optimizations. All pass the same solver and predicted-feasibility checks without fallback. The four weeks are a controlled mismatch diagnostic; they are not used to estimate an unrestricted distribution of building uncertainty.

Table 8. Four-week mismatch diagnostics: actual plant (a = 0.2 ACH, G = 230.4), nominal predictive model, mean values, and maximum concentration.

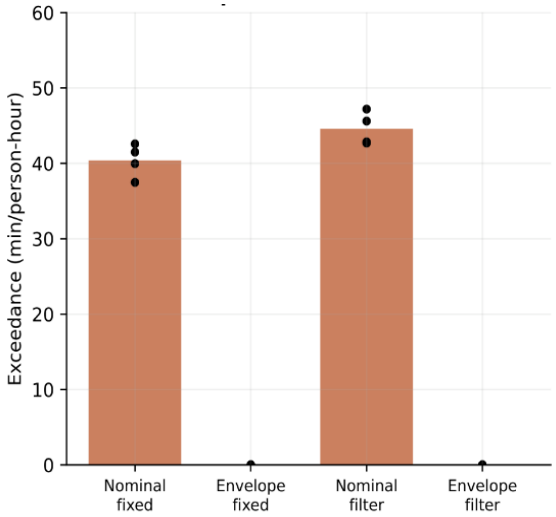
Applied Policy	Modeled Total Energy (kWh/Week)	Exceedance (Min/Person-Hour)	Largest CO_2 (ppm)
Nominal robust constant	23.3531	40.383	1239.661
Envelope robust constant	28.7000	0	999.996
Nominal filtered MPC	21.1881	44.587	1011.063
Envelope-filtered MPC	21.3724	0	1000.000

The nominal constant rate no longer satisfies the invariant-set condition and permits concentrations near 1240 ppm. Nominal filtering also loses its guarantee: its worst-case prediction uses the wrong source and infiltration bounds. The envelope-aware filter restores zero exceedance in all four weeks, consistent with Proposition 3. Relative to nominal filtering within the same stress regime, it adds 0.18425 kWh/week on average. It also consumes substantially less modeled energy than the envelope constant-flow comparator because it adapts its commands to the current state.

The stress energy totals should not be compared directly with the nominal totals as an efficiency gain. Lower actual infiltration reduces the uncontrolled heating term, while greater generation raises the ventilation requirement. Energy comparisons in *Table 8* are meaningful within its common stress regime. *Fig. 6* pairs the concentration trajectories with the four-week exposure outcomes so that the change in guarantee is visible independently of the common heating offset.



a.



b.

Fig. 6. Mismatch diagnostics under reduced infiltration and increased generation; a. first 48 hours of week 60701; b. mean and individual outcomes across four paired weeks.

This experiment does not establish tolerance to every model discrepancy. A bounded measurement error, variable outdoor CO_2, imperfect mixing, airflow saturation, or an unmodeled arrival above the occupancy cap can invalidate the simple condition. The filter also assumes that the commanded mechanical rate is realized. A practical implementation would require calibrated concentration and flow measurements, a justified load envelope, explicit handling of infeasibility, and validation of the hold model. Those are substantive modeling tasks rather than a change in optimizer tolerance.

8 | Discussion and Scope

The principal analytical result is that exact transient prediction can reveal a favorable optimization structure that is hidden by the state-input product in a differential equation. For the scalar balance studied here, positivity turns every predicted concentration into a convex function of the planned airflow sequence. This result permits retaining the cubic fan law and exact held-input propagation in the horizon problem. It does not extend automatically to a multi-zone thermal model, and it should not be used to reclassify every bilinear HVAC problem as convex.

The information result is equally specific. Forecast error at later horizons and lack of knowledge about the current interval have different consequences for a hard one-step target. With immediate current occupancy, a matched plant and feasible first constraint suffice for concentration invariance. With delayed reporting, a point forecast replaces the current unknown load and can permit violations. The model-based filter closes that gap through an occupancy cap, at the cost of conservatism whenever the actual load is below the envelope. Its nominal energy increment is small in the declared winter benchmark; that magnitude is an empirical result, not a property proved for all weather and occupancy processes.

Baseline design changes the interpretation of energy percentages. The 0.5 ACH schedule imposes a 900 ppm ceiling, whereas the predictive controllers target 1000 ppm. The robust constant comparator removes part of this margin difference. Training-only rule tuning further reduces the comparison's dependence on a rule that starts ventilation after the target has already been crossed. The resulting 4.832% saving against the tuned rule is more informative than the larger saving against the original rule, but neither is a universal estimate of residential savings.

Table 9 connects the literature to the contribution actually completed. The framework combines established optimization and invariant-set ideas with a direct scalar proof and a complete comparative campaign. It is intended to support more discriminating evaluation, not to replace high-fidelity emulators or field tests. BOPTEST-style benchmarking [28] and operating-system demonstrations [29] remain relevant next validation stages because they introduce actuator, thermal, communication, and modeling effects that the current balance intentionally isolates.

Table 9. Model-specific gaps and corresponding contributions: completed derivations and computational results.

Question	Relevant Research	Completed Response
Does a bilinear differential balance require a linearized horizon model?	MPC and convex building operation [4], [9], [18], [20]	Exact state elimination and convexity in Proposition 2
Does forecast accuracy certify a realized concentration target?	Occupancy and risk-aware ventilation [6–8], [12–14]	Delayed causal forecast, oracle comparator, and threshold sensitivity
Can a proposed flow be corrected using a declared load envelope?	Safety filters and invariant sets [21], [22]	Monotone filter and conditional guarantee in Proposition 3
How much does baseline selection affect the reported benefit?	Ventilation evaluation and benchmarking [16], [28], [30]	Training-tuned rule, two constant comparators, and twelve paired weeks
What happens when the plant leaves the nominal envelope?	Generation and sensing limitations [3], [17], [23], [25]	Four paired mismatch tests with nominal and envelope-aware policies

Several limitations remain. The occupancy matrices create one declared stochastic regime and contain no weekday–weekend distinction, household heterogeneity, or long absence state. The temperature series is generated from simple periodic components, and its forecast is exact. The concentration observations are exact, so the benchmark does not measure sensitivity to sensor noise or drift. Both generation and infiltration are constant within each nominal or stress campaign, and the recovery and fan laws are uncalibrated equipment assumptions.

The statistical scope is also limited. Twelve test weeks support an assessment of repeatability under the selected generator; they do not estimate variability across a housing stock. Pointwise t intervals assume that week-level paired differences are adequately described by that approximation. The six energy contrasts are not simultaneous confidence statements, and no nonsignificant result is used to assert equivalence. The deterministic proofs carry the concentration guarantee only where their assumptions apply. The empirical comparisons quantify energy and exposure under one implementation and one information contract.

A physical extension should first validate the well-mixed balance and the realized airflow. It should then estimate generation and infiltration envelopes, evaluate measurement error, and compare forecasts using information available at their actual timestamps. Coupling temperature, humidity, particulate matter, or other pollutants would change both the objective and the admissible set. These extensions require new evidence; lowering CO_2 alone cannot establish overall indoor-air-quality protection [2], [17], [24], [25].

9 | Conclusion

Exact held-input dynamics give the single-zone ventilation problem a useful structure: with nonnegative initial excess concentration and generation, predicted concentration is convex in the planned airflow sequence. This permits a convex finite-horizon formulation that retains the cubic fan law. A monotone one-step filter then separates energy-oriented prediction from an explicit occupancy and plant envelope.

The complete twelve-week comparison shows why these distinctions matter. Filtered MPC reduces modeled energy by 4.832% relative to a training-tuned rule and maintains the selected concentration target in the matched model. An accurate causal mean forecast with delayed occupancy reports still permits frequent small exceedances; enforcing the load envelope removes them at a small measured energy increment. The mismatch experiment shows that nominal filtering does not survive invalid parameter bounds, while envelope-aware filtering restores the conditional guarantee. The resulting contribution is a mathematically explicit and reproducible control benchmark whose claims remain tied to its information, plant, and evaluation assumptions.

Data and computational availability

All occupancy and temperature sequences are model-generated; no human-participant or measured apartment dataset is used. The accompanying computational package contains the complete generator, controller and filter implementation, figure code, parameters, software versions, seed-level metrics, and saved trajectories. No external repository or permanent archive is claimed.

Funding

The research received no specific grant.

Competing interests

The authors declare no competing interests.

AI assistance

AI assistance was used under the Interdisciplinary Research Portfolio Program.

References

[1] Fisk, W. J., & De Almeida, A. T. (1998). Sensor-based demand-controlled ventilation: A review. *Energy and Buildings*, *29*(1), 35–45. https://doi.org/10.1016/s0378-7788(98)00029-2.

[2] Persily, A., Bahnfleth, W. P., Kipen, H., Lau, J., Mandin, C., Sekhar, C., ..., & Weekes, L. C. N. (2022). ASHRAE position document on indoor carbon dioxide. *ASHRAE*, *64*, 50-52. https://www.ashrae.org/file%20library/about/position%20documents/pd-on-indoor-carbon-dioxide-english.pdf

[3] Persily, A., & de Jonge, L. (2017). Carbon dioxide generation rates for building occupants. *Indoor Air*, *27*(5), 868–879. https://doi.org/10.1111/ina.12383.

[4] Mayne, D. Q., Rawlings, J. B., Rao, C. V, & Scokaert, P. O. M. (2000). Constrained model predictive control: Stability and optimality. *Automatica*, *36*(6), 789–814. https://doi.org/10.1016/s0005-1098(99)00214-9

[5] Oldewurtel, F., Parisio, A., Jones, C. N., Gyalistras, D., Gwerder, M., Stauch, V., …, & Morari, M. (2012). Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings*, *45*, 15–27. https://doi.org/10.1016/j.enbuild.2011.09.022

[6] Li, B., Wu, B., Peng, Y., & Cai, W. (2022). Tube-based robust model predictive control of multi-zone demand-controlled ventilation systems for energy saving and indoor air quality. *Applied Energy*, *307*, 118297. https://doi.org/10.1016/j.apenergy.2021.118297

[7] Zhuang, C., Choudhary, R., & Mavrogianni, A. (2022). Probabilistic occupancy forecasting for risk-aware optimal ventilation through autoencoder Bayesian deep neural networks. *Building and Environment*, *219*, 109207. https://doi.org/10.1016/j.buildenv.2022.109207

[8] Jiang, Z., Deng, Z., Wang, X., & Dong, B. (2023). PANDEMIC: Occupancy driven predictive ventilation control to minimize energy consumption and infection risk. *Applied Energy*, *334*, 120676. https://doi.org/10.1016/j.apenergy.2023.120676

[9] Drgoňa, J., Arroyo, J., Figueroa, I. C., Blum, D., Arendt, K., Kim, D., ..., & Helsen, L. (2020). All you need to know about model predictive control for buildings. *Annual Reviews in Control*, *50*, 190-232. https://doi.org/10.1016/j.arcontrol.2020.09.001

[10] Killian, M., & Kozek, M. (2016). Ten questions concerning model predictive control for energy efficient buildings. *Building and Environment*, *105*, 403–412. https://doi.org/10.1016/j.buildenv.2016.05.034

[11] Kathirgamanathan, A., De Rosa, M., Mangina, E., & Finn, D. P. (2021). Data-driven predictive control for unlocking building energy flexibility: A review. *Renewable and Sustainable Energy Reviews*, *135*, 110120. https://doi.org/10.1016/j.rser.2020.110120

[12] Oldewurtel, F., Sturzenegger, D., & Morari, M. (2013). Importance of occupancy information for building climate control. *Applied Energy*, *101*, 521–532. https://doi.org/10.1016/j.apenergy.2012.06.014

[13] Candanedo, L. M., & Feldheim, V. (2016). Accurate occupancy detection of an office room from light, temperature, humidity and CO_2 measurements using statistical learning models. *Energy and Buildings*, *112*, 28–39. https://doi.org/10.1016/j.enbuild.2015.11.071

[14] Caballero-Peña, J., Osma-Pinto, G., Rey, J. M., Nagarsheth, S., Henao, N., & Agbossou, K. (2024). Analysis of the building occupancy estimation and prediction process: A systematic review. *Energy and Buildings*, *313*, 114230. https://doi.org/10.1016/j.enbuild.2024.114230

[15] Lu, X., Pang, Z., Fu, Y., & O'Neill, Z. (2022). The nexus of the indoor CO2 concentration and ventilation demands underlying CO2-based demand-controlled ventilation in commercial buildings: A critical review. *Building and Environment*, *218*, 109116. https://doi.org/10.1016/j.buildenv.2022.109116

[16] Lu, X., Pang, Z., Fu, Y., & O'Neill, Z. (2022). Advances in research and applications of CO_2-based demand-controlled ventilation in commercial buildings: A critical review of control strategies and performance evaluation. *Building and Environment*, *223*, 109455. https://doi.org/10.1016/j.buildenv.2022.109455

[17] Persily, A. (2022). Development and application of an indoor carbon dioxide metric. *Indoor Air*, *32*(7), e13059. https://doi.org/10.1111/ina.13059

[18] Naqvi, S., Kar, K., Bhattacharya, S., Chandan, V., Mishra, S., & Salsbury, T. (2023). Air quality and comfort constrained energy efficient operation of multi-zone buildings. *Building and Environment*, *244*, 110716. https://doi.org/10.1016/j.buildenv.2023.110716

[19] Andersen, A. H., & Jradi, M. (2025). Multi-objective optimization of building ventilation systems using model predictive control: integrating air quality, energy cost, and environmental impact. *Applied Sciences*, *15*(1), 451. https://doi.org/10.3390/app15010451

[20] Boyd, S. (1977). Convex optimization-boyd & vandenberghe. *Journal on Numerical Analysis*, *8*, 6. https://web.stanford.edu/~boyd/cvxbook

[21] Wabersich, K. P., & Zeilinger, M. N. (2021). A predictive safety filter for learning-based control of constrained nonlinear dynamical systems. *Automatica*, *129*, 109597. https://doi.org/10.1016/j.automatica.2021.109597

[22] Ames, A. D., Xu, X., Grizzle, J. W., & Tabuada, P. (2016). Control barrier function based quadratic programs for safety critical systems. *IEEE Transactions on Automatic Control*, *62*(8), 3861–3876. https://doi.org/10.1109/tac.2016.2638961

[23] Fisk, W. J. (2010). *CO_2 monitoring for demand controlled ventilation in commercial buildings*. https://indoor.lbl.gov/publications/co2-monitoring-demand-controlled

[24] Peng, Z., & Jimenez, J. L. (2021). Exhaled CO2 as a COVID-19 infection risk proxy for different indoor environments and activities. *Environmental Science & Technology Letters*, *8*(5), 392–397. https://doi.org/10.1021/acs.estlett.1c00183

[25] Rizo-Maestre, C., Flores-Moreno, J. M., Nebot Sanz, A., & Echarri-Iribarren, V. (2025). Intelligent ventilation and indoor air quality: state of the art review (2017-2025). *Buildings*, *16*(1), 65. https://doi.org/10.3390/buildings16010065

[26] Page, J., Robinson, D., Morel, N., & Scartezzini, J. L. (2008). A generalised stochastic model for the simulation of occupant presence. *Energy and Buildings*, *40*(2), 83–98. https://doi.org/10.1016/j.enbuild.2007.01.018

[27] Richardson, I., Thomson, M., & Infield, D. (2008). A high-resolution domestic building occupancy model for energy demand simulations. *Energy and Buildings*, *40*(8), 1560–1566. https://doi.org/10.1016/j.enbuild.2008.02.006

[28] Blum, D., Arroyo, J., Huang, S., Drgoňa, J., Jorissen, F., Walnum, H. T., ..., & Helsen, L. (2021). Building optimization testing framework (BOPTEST) for simulation-based benchmarking of control strategies in buildings. *Journal of Building Performance Simulation*, *14*(5), 586-610. https://doi.org/10.1080/19401493.2021.1986574

[29] Blum, D., Wang, Z., Weyandt, C., Kim, D., Wetter, M., Hong, T., & Piette, M. A. (2022). Field demonstration and implementation analysis of model predictive control in an office HVAC system. *Applied Energy*, *318*, 119104. https://doi.org/10.1016/j.apenergy.2022.119104

[30] Morris, T. P., White, I. R., & Crowther, M. J. (2019). Using simulation studies to evaluate statistical methods. *Statistics in Medicine*, *38*(11), 2074–2102. https://doi.org/10.1002/sim.8086

[31] Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 65–70. https://www.ime.usp.br/~abe/lista/pdf4R8xPVzCnX.pdf

Appendix

Appendix A. Derivatives and Verification Details

For interval j, let X_j denote the predicted excess concentration and $J_{j\ell} = \partial X_j / \partial u_\ell$. The exact horizon Jacobian used by the solver obeys

$$\begin{aligned} J_{j+1,\ell} &= e^{-(a+u_j)\Delta} J_{j\ell}, \qquad \ell < j. \\ J_{j+1,j} &= -\Delta e^{-(a+u_j)\Delta} X_j + G\widehat{N}_j \phi'(a+u_j), \end{aligned} \tag{A.1}$$

with zero derivatives for future controls and $J_{0\ell} = 0$. The objective derivative before constant positive scaling is $(\Delta/1000)\left[(75/9)u_j^2 + \beta_j\right]$. The code checks *Eq. (A.1)* against central differences using a step of 10^{-5} ACH at a fixed random interior point. The maximum absolute discrepancy is 4.08×10^{-9} in the computed units. Random convex-combination checks use 200 independent pairs of control vectors, nonnegative forecast occupancies, and mixing weights in $[0,1]$; the largest signed convexity residual is -5.82×10^{-7} ppm. These finite checks diagnose implementation errors and do not replace *Proposition 2*.

For an interval with constant inputs, its equilibrium concentration is $C_{eq} = C_{out} + GN/(a+u)$. When a threshold lies strictly between the two endpoint concentrations, its crossing time relative to the interval start is

$$t_\times = -\frac{1}{a+u} \log\left(\frac{C_{max} - C_{eq}}{C_k - C_{eq}}\right). \tag{A.2}$$

An increasing trajectory contributes $\Delta - t_\times$ to exceedance duration; a decreasing one contributes $t_\times$. Intervals with both endpoints above or below the threshold contribute Δ or zero. This avoids rounding every crossing to a five-minute duration and makes the person-weighted outcome in *Eq. (A.2)* consistent with the continuous balance.

The complete campaign uses 88,704 predictive solves when nominal and stress tests are combined. No fallback occurs. Exact equilibrium identities, the nominal and envelope constant-rate limits, finite-difference derivatives, and the predicted invariant bounds are checked by executable assertions. The saved outputs include all twelve nominal concentration and airflow trajectories and the first stress-test trajectory for each policy, together with all sixteen stress-test metric rows. The package's README states the commands needed to regenerate the figures and repeat the numerical campaign.