


Day 4: Blocking Transmission (Withstand)

Thursday 29 October · discussion 3:00 to 4:50 pm New York

THE DAY MAP

How do you slow the spread enough that an outbreak doesn't have to become a catastrophe?



Your deep-dive: page 4.3. Respirators for the people who have to keep working. Its readings come first below; the other options follow.

14 readings here, about 145 minutes of reading. Discussion Thursday 29 October, 3:00 to 4:50 pm New York, with Conor Spence.

[4.1-1](#) The Four Pillars: A Hypothesis for Countering Catastrophic Biological Risk
[original paper →](#)

2 + 8 min

<u>4.1-2</u>	Innovate and Stockpile: Respiratory Protection for Essential Workers in a Catastrophic Pandemic original paper →	2 + 10 min
<u>4.1-3</u>	Indoor Air Quality Is the Next Great Public Health Challenge original paper →	3 + 13 min
<u>4.1-4</u>	Far-UVC original paper →	3 + 4 min
<u>4.1-5</u>	Flipping the switch on far-UVC original paper →	3 + 4 min
<u>4.3-1</u>	Blueprint for Pandemic PPE original paper →	3 + 15 min
<u>4.4-1</u>	Glycol vapor pathogen disinfection: a defensive technology for emerging biothreats original paper →	3 + 0 min
<u>4.4-2</u>	On far-UVC and air filtration original paper →	3 + 12 min
<u>4.4-3</u>	Blueprint for Far-UVC original paper →	3 + 20 min
<u>4.4-4</u>	Glycol vapors for pathogen suppression original paper →	3 + 6 min
<u>4.5-1</u>	Executive Summary to the Royal Society report "COVID-19: examining the effectiveness of non-pharmaceutical	3 + 13 min

interventions"

[original paper →](#)

4.5-2	Epidemiological impacts of the NHS COVID-19 app in England and Wales throughout its first year original paper →	3 + 0 min
------------------------------	--	-----------

4.5-3	Effectiveness of international border control measures during the COVID-19 pandemic: a narrative synthesis of published systematic reviews original paper →	3 + 0 min
------------------------------	--	-----------

4.5-4	International air travel-related control measures to contain the Covid-19 pandemic: A companion review to a Cochrane rapid review original paper →	3 + 0 min
------------------------------	---	-----------

WHAT THE COURSE ASKS YOU TO DO

4.2 NPI deep-dive

Pick 1 of the following 2 topics to dive deeper into. We've given you some resources on the following pages as a starting point.

- **Option 1: Pandemic-proof PPE**
- **Option 2: Biohardening indoor spaces**
- **Option 3: Social interventions**

Aim to answer the following questions: - What are the intended goals of this intervention? - What is the current state of this intervention? - What are the technical or policy blockers to having the intervention ready to deploy effectively in a pandemic?

You will bring the results of your research to share in your group discussion.

THE READINGS

4.1-1 · CORE · COURSE TIME 20 MIN

The Four Pillars: A Hypothesis for Countering Catastrophic Biological Risk

Andrew Snyder-Beattie · 2025

[Read the original paper →](#) · the course assigns: whole piece, including footnotes

Instead of listing the pathogens that might come for us and building a countermeasure for each, build four defences that work against anything: respirators, pathogen-free indoor space, pathogen-agnostic detection, and fast-built medicines.

Average summary

Snyder-Beattie starts from a worry: biological risk is worse than most people think. Recent work on **mirror bacteria** describes one organism that could infect

humans, plants and animals, survive in soil or dust, and spread worldwide within several months. Less extreme engineered pathogens could still cause hundreds of millions, or billions, of casualties.

The usual approach is to list pathogen families, rank them by risk, and develop medical countermeasures for the top of the list in advance. He says that works for nature and fails against a determined adversary. No list is ever complete — mirror bacteria would have been left off most of them — and the public history of the Soviet bioweapons programme shows how creatively assumptions get broken: attempts at smallpox-Ebola hybrids, a plague strain resistant to over 15 kinds of antibiotics, pathogens modified to trigger autoimmune reactions.

So he proposes four pillars that are "future-proof" because they exploit constraints every pathogen faces. Any pathogen has to physically enter a body; any catastrophic one has to spread widely and cause harm, which makes it detectable.

1. **PPE.** Of all the ways a pathogen gets in, the airway is the weakest link, so respirators matter most. He prices five options and lands on elastomeric respirators as the sweet spot: protection factor over 100 (they filter 99% of incoming particles), under \$50, reusable for six months on one set of filters, over 10 years in storage.

2. **Biohardening.** People still need pathogen-free places to eat, sleep and work. Far-UVC light, air filtration, and glycol vapours are candidates, plus the speculative option of households building cleanroom-like conditions from materials they already own.
3. **Detection.** PPE only works if people know to wear it. A slow, stealthy pathogen (he uses HIV, where severe symptoms typically come 8 to 10 years after infection) could be everywhere before anyone noticed. Metagenomic sequencing, run constantly whether or not there is an emergency, is the current state of the art.
4. **Medical countermeasures,** created after the threat is identified. He leaves this one mostly outside the post.

The headline claim: widespread protection under the first three pillars "might be feasible by the end of 2027 with a budget of less than \$1 billion." Stockpiling 40 million respirators at \$50 each would cost less than a single B-2 bomber.

Why it's assigned

This is the frame for the whole withstand day. The course's picture is four layers of defence stacked like Swiss cheese, and this reading is the argument for why two of the withstand layers — PPE and biohardening of indoor spaces — are the ones worth

building in peacetime: they are pathogen-agnostic, so they can be bought and installed before anyone knows what is coming. Snyder-Beattie leads Coefficient Giving's Biosecurity and Pandemic Preparedness program, so this is also a funder telling you where the money is pointed.

In depth

The case against lists

The piece opens by raising the ceiling on how bad things could get. Mirror bacteria are the extreme case: one organism able to infect humans, plants and animals, persist in the environment, and go global in several months. In the worst case he calls that an existential risk to humanity, "especially if the responses/countermeasures were inadequate." Below that sit engineered pathogens that could cause hundreds of millions or billions of casualties.

Then the argument. You could enumerate pathogen types, then subtypes (adenoviruses, coronaviruses, paramyxoviruses), score each for risk, and build countermeasures in priority order. Good for natural threats, "ultimately doomed against a determined adversary," for two reasons. First, you can never be confident the list is exhaustive. Second, even for threats on the list, a breakthrough can break your assumptions about whether your countermeasure still works. His evidence is the Soviet programme as

publicly documented by Leitenberg and Zilinskas: hybrid attempts, antibiotic-resistant plague, autoimmune triggers. The conclusion is that defences have to be fully generalised.

The four pillars: PPE, biohardening, detection, and reactive medical countermeasures. The first three are aimed at the earliest stage of a catastrophe, with the goal of saving lives and "preserving industrial and scientific capacity for the rest of the world to respond more fully." That is the leverage: protect the people who keep the lights on, and you buy the time to make everything else.

PPE: the airway is the weak point

He asks which PPE is most needed by listing every route into a body — ingestion, inhalation, skin breaks — and argues the respiratory route is the weakest link on first principles (lung surface area is large, you cannot stop breathing, other routes are easier to sterilise) and empirically (respiratory disease still causes pandemics in wealthy countries; water-borne disease largely does not).

Then the table, which is the heart of the reading. Five levels, with degree of protection, fixed cost per unit, and cost per person per day over six months:

- Surgical mask: under 10, about \$0.05–\$0.10, about \$0.05–\$0.10 a day.

- Disposable N95: about 10–100, about \$0.50–\$1.00, about \$0.25–\$0.50 a day.
- Elastomeric half-face respirator: about 100–1000+, about \$20–\$80, about \$0.10–\$0.50 a day.
- Powered air purifying respirator (PAPR): about 10,000, about \$600–\$2,500, about \$5.00–\$15.00 a day.
- HAZMAT suit with its own air supply: about 1 million, about \$2,000–\$10,000, about \$50.00–\$100.00 a day.

Elastomerics win on cost per unit of protection. A dedicated effort could push the price down; some overseas manufacturers already claim to make them for below \$5. In the footnotes: surgical masks consistently test below 10 and sometimes below 2; N95s show 5th-percentile protection factors above 10 and 95th-percentile usually above 100, with performance depending on aerosol size and face shape; one vendor sells elastomerics at about \$40 with filters, and the EPA's own cost modelling uses \$22 up front.

How many do you need? "Although it's uncertain," different studies suggest protecting between 12 and 45 million people in the US would likely be enough to prevent infrastructure collapse. He proposes 40 million respirators at \$50 — less than one B-2 bomber. A 20-year-old survey found over 2 million US workers already had elastomerics at work, not

counting mining or the military, and HHS and DHS bought about 2.7 million elastomerics in 2023 and 2024.

The most interesting claim is about **source control**. If a mask with a protection factor of 100 filters exhaled as well as inhaled air, and both people wear one, the combined factor is 10,000, a 99.99% reduction. Transmissibility has a physical ceiling set by how much aerosol a human can produce, so a 10,000-fold cut "might be sufficient to prevent transmission even under worst-case theoretical limits." His footnote works it out: measles cannot on average infect more than one person every four minutes in typical workplace conditions, so a 10,000-fold increase in that time allows several weeks of continuous work before transmission becomes probable.

Two assumptions carry it. People have to wear the gear properly — he thinks motivation would be high if infection meant near-certain death. And there have to be pathogen-free places to eat and sleep, which is easy in a human-to-human pandemic and hard in a mirror bacteria world where the environment itself is contaminated.

Victory condition: most of the world's critical workers have access to an elastomeric respirator rated above fit factor 100 within 24 hours. For the US alone: 30 million respirators, under \$500M, end of 2027.

Biohardening

We know how to protect one immunocompromised person. Hospital cleanrooms do it with aggressive filtration; David Vetter, a child with severe immunodeficiency, was kept alive for 12 years in a sterile room and a plastic bubble. Militaries do it too, with filtration in tanks and units trained for "collective protection."

The problem is scale. Finland and Switzerland give every citizen access to shelters built to withstand biological and chemical threats, but that still covers about 0.2% of the world's population.

The promising paths: far-UVC and basic air filtration, neither needing much retrofitting. And emergency measures that could be stood up in two weeks — triethylene glycol vapour, already made for natural gas processing, which "seems effective at killing pathogens when deployed in the air or on surfaces, and also seems extremely safe for humans to breathe." He is careful: more studies should be done on efficacy and safety for glycols. His footnote sizes the supply: the US probably makes about 80 thousand tonnes of TEG a year (global production about 880 kilotonnes), you need about 0.1 to 1 mg/m³ for maximum effect, and US output could cover at least 3 billion m³ of indoor space continuously and potentially over 180 billion m³ — against roughly 3 billion m³ of

US industrial buildings, 27 billion commercial and 60 billion residential.

Then the speculative one: DIY biohardening.

Household insulation converted into improvised HEPA filters (about 70% of US houses have fibreglass insulation), furnace fans strong enough to push air through them and pressurise a house, surfaces sterilised with hypochlorous acid made from salt water and an electric current. He calls an initial investigation "less hopelessly far-fetched than one might initially expect," and says stress-testing it to validate or rule it out "would be an obvious next step." Victory condition: a guidebook or AI assistant good enough that a randomly chosen group could retrofit a building in two weeks to 3 to 5 logs of protection, achievable by end of 2027 at perhaps \$5–10M a year.

He is blunt that biohardening alone is not enough in the mirror bacteria case: someone still has to go out for food, power and water. "Staying inside is not a viable long-term strategy."

Detection

The first pillars assume society takes the threat seriously. The failure mode is a pathogen that spreads without symptoms and turns lethal late. HIV is the example: severe symptoms typically 8 to 10 years after infection, and decades of spread before it was discovered in the 1980s. Imagine that, airborne.

So detection must be pathogen-agnostic and run constantly, with no emergency to prompt it.

Metagenomic sequencing is the state of the art, with protein sequencing, mass spectrometry and electron microscopy as supplements. His structural argument: detection should favour the defender, because verifying that something is dangerous is fundamentally easier than generating the dangerous thing — so if AI gets good enough to redesign pathogens around screening, it is equally good at the screening.

Victory condition: any novel pathogen is widely recognised at least a month before most of the world would have been infected. With a doubling time of 3 days and a simple model, that means detecting at about 0.1% cumulative prevalence. Under ~\$100M a year, end of 2027.

The figure

The reading's argument rests on one table rather than a chart: five rows of respiratory PPE, columns for degree of protection, fixed cost per unit, and cost per person per day over six months. Reading down it, protection climbs by factors of ten while daily cost climbs faster at the top end — surgical masks and N95s are cheap and weak, HAZMAT suits are strong and unaffordable, and the elastomeric row is the one place where a protection factor above 100 costs ten

to fifty cents a day. That single row is the paper's recommendation.

Words to know

- **Protection factor (or fit factor)** — how much cleaner the air inside the mask is than the air outside. A factor of 100 means 99% of incoming particles are blocked. It matters here because it is the one number he uses to rank every option, and because leakage around the seal, not the filter, usually sets it.
- **Elastomeric respirator** — a reusable rubber half-mask with replaceable filter cartridges, the kind used in industry. The whole PPE recommendation rests on it.
- **PAPR (powered air purifying respirator)** — a respirator with a battery-powered fan that pushes filtered air in, so it does not need a tight seal. Very protective, very expensive to run.
- **Source control** — filtering the air a wearer breathes out, not just in. It is what lets two masks multiply into a combined factor of 10,000.
- **Mirror bacteria** — a hypothetical organism built from mirror-image versions of life's normal molecules. It is the worst case the piece is written against: environmental persistence, and no immune recognition to rely on.
- **Biohardening** — building physical barriers and sterilisation into ordinary rooms so indoor space

stays pathogen-free, as opposed to protecting one body at a time.

- **Metagenomic sequencing** — sequencing all the genetic material in a sample at once, rather than testing for one known pathogen. The basis of pathogen-agnostic detection.
- **Log of protection** — a tenfold reduction. "3 to 5 logs" means cutting pathogen levels by 1,000 to 100,000 times.

Numbers in this reading

- Several months: how long mirror bacteria might take to spread worldwide.
- Over 15 kinds of antibiotics: the resistance engineered into a Soviet plague strain, as publicly documented.
- Under 10, 10–100, 100–1000+, ~10,000, ~1 million: the degrees of protection for surgical masks, disposable N95s, elastomerics, PAPRs and HAZMAT SCBA suits.
- \$0.05–\$0.10 up to \$2,000–\$10,000: the matching range of fixed cost per unit.
- \$0.05–\$0.10 up to \$50.00–\$100.00: the matching cost per person per day over six months.
- Over 100 protection factor, equal to filtering 99% of incoming particles: the elastomeric claim.
- Six months on one set of filters, over 10 years in storage, under \$50 each, possibly below \$5 from

some overseas manufacturers: the rest of the elastomeric case.

- 10,000-fold, a 99.99% reduction: the combined effect if both people wear a factor-100 mask that also filters exhaled air.
- One person every ~4 minutes: the fastest measles can transmit in typical workplace conditions, used to set the upper bound.
- 12–45 million people: the range of US workers different studies suggest need protection to prevent infrastructure collapse.
- 40 million respirators at \$50: cheaper than a single B-2 bomber.
- 30 million respirators, under \$500M, end of 2027: the US PPE victory condition.
- Over 2 million US workers with elastomerics at work (20-year-old survey); 2.7 million bought by HHS and DHS in 2023 and 2024.
- 12 years: how long David Vetter was kept alive in a sterile room and plastic bubble.
- ~0.2%: the share of the world's population covered by Finnish and Swiss civil-defence shelters.
- ~80 thousand tonnes a year US TEG production, ~880 kilotonnes global; 0.1 to 1 mg/m³ needed for maximum effect; enough for at least 3 billion m³ and potentially over 180 billion m³ of indoor space.
- 3 to 5 logs of protection in two weeks, \$5–10M a year: the biohardening victory condition.

- 8 to 10 years: HIV's typical delay to severe symptoms.
- ~0.1% cumulative prevalence, assuming a 3-day doubling time: the detection target, under ~\$100M a year.
- Under \$1 billion by end of 2027: the whole three-pillar ask.

What the author is unsure about

He flags his own softness repeatedly. The 12-to-45-million figure is "uncertain." A protection factor of 100 "might" be enough to block human-to-human transmission. He is "less confident" that 100 is enough in a mirror bacteria scenario, where exposure comes from the environment and no mask on a source can help, which argues for higher factors at higher cost. The glycol option needs more studies on efficacy and safety. DIY biohardening needs stress-testing to validate or rule out. Whether medical countermeasures could work against any possible future weapon is "an open question" and his answer is explicitly a hunch. Two assumptions are named as assumptions: that people will wear the equipment properly, and that safe pathogen-free areas exist. And the whole thing is labelled a hypothesis, not a finding — he calls the victory conditions "tentative."

For the discussion

- The strongest claim in the reading is also the cheapest: if a factor-100 mask filters exhaled as well as inhaled air, two of them give 10,000-fold protection, which may be enough to stop transmission at the physical limit of what human lungs can emit — so source control, not filter grade, is the feature to argue for.
- The cost framing is doing a lot of persuasive work: 40 million respirators for less than one B-2 bomber is a political argument, not an epidemiological one, and it is worth asking whether the bottleneck was ever the money rather than who owns the decision.

4.1-2 · CORE · COURSE TIME 20 MIN

Innovate and Stockpile: Respiratory Protection for Essential Workers in a Catastrophic Pandemic

Nadia X. Montazeri and Jonas B. Sandbrink · 2023

[Read the original paper →](#) · the course assigns: whole commentary (Europe PMC full text, article body)

The respirators we stockpiled for COVID-19 are not good enough for a pandemic that kills most of the people it infects, and the two fixes are to pay for better products now and to stockpile reusable ones.

Average summary

This is a short commentary in *Health Security* by Nadia Montazeri (Stanford's Center for International Security and Cooperation) and Jonas Sandbrink (Oxford). Their subject is **respiratory protective equipment (RPE)** in a **global catastrophic biological risk (GCBR)** — the authors' term for a worst-case biological event that could threaten the stability of society.

Their starting move is to make the severity concrete. Almost by definition a GCBR event is more lethal than COVID-19. If COVID-19 had caused severe disease at the rate SARS did, with a case fatality ratio of 11%, the result might have qualified. Case fatality ratios for respiratory viruses can run higher still: Middle East respiratory syndrome at 32.7%, avian influenza at 52.6%. Engineered agents could combine that lethality with high transmissibility.

Then the argument in four parts.

Who needs protection. In a pandemic lethal across all age groups, most people who can stay home will. A small number cannot: not only nurses and doctors but workers in food and water supply, waste management and energy production. They need equipment both to survive and to be willing to keep showing up. The authors point out that existing respiratory protection programmes essentially stop at the hospital door — healthcare workers are the only group with any real chance of being fit-tested and trained.

Why today's products may not be enough.

Assuming a perfect fit, N95-style respirators (including KN95s, KF94s, FFP2s) cut particles at the most penetrating size — on average 0.3 microns — to 5% or 6%. Whether that suffices depends on the pathogen's minimum infectious dose, which for a novel agent you will not know. And the measured performance depends on fit, which depends on a fit test. The CDC estimates one experienced operator with good logistics could fit-test up to 400 people in an 8-hour day. A fast-moving GCBR gives you no time for that.

Trust. In COVID-19 the problem was compliance; in a GCBR, the authors expect the problem to be convincing workers that what they have been handed will actually keep them alive. Surveys show healthcare workers preferred reusable equipment — powered air purifying respirators and elastomerics — in high-risk scenarios even though they rated them less comfortable.

Stockpiles and innovation. Just-in-time manufacturing worked for COVID-19 and will be too slow next time. So: stockpile reusable products, and use prizes, advanced market commitments and changed testing standards to get better ones built.

Why it's assigned

This is the withstand layer applied to the people who cannot shelter. The course's framing for the day is that non-pharmaceutical interventions are mostly pathogen-agnostic and can therefore be prepared in advance, and that lockdowns still leave critical workers exposed; this paper is the specific argument for what those workers should be handed and why the current stockpile answer is wrong. It also sets up the deep-dive on page 4.3 by naming fit, trust and reusability as the design problems.

In depth

The threat model

The paper opens by defining GCBRs as worst-case biological events that could threaten the stability of society, and establishing that such an event would be more lethal than COVID-19. The comparisons are SARS (case fatality ratio 11%), MERS (32.7%) and avian influenza (52.6%). High lethality and high transmissibility "may occur especially for engineered agents." The authors concede they cannot predict the exact properties of such a pathogen, but argue that airborne transmission is the most likely route to exponential spread — which is precisely why physical countermeasures can be prepared in advance: they protect against many different respiratory viruses rather than one.

They then lay out the equipment landscape. The most common RPE is the **filtering facepiece respirator (FFR)**: N95 and N99 in the United States, FFP2 and FFP3 in Europe. During COVID-19, medical masks and barrier face coverings were pressed into service to compensate for respirator shortages. Reusable options — elastomeric half-mask respirators, powered air purifying respirators, and chemical, biological, radiological and nuclear (CBRN) gear — were used in varying amounts but played a minor role next to disposables. They note plainly that "it is unknown how countries will adapt their pandemic preparedness plans."

Protecting essential workers

The behavioural assumption is that in a pandemic lethal across all age groups, anyone who does not urgently need to leave home will stay in. The evidence offered is from shielding behaviour during COVID-19: cystic fibrosis patients in the UK reported adhering to strict government-recommended shielding, and a Spanish survey found people who perceived themselves as vulnerable were more likely to adhere to protective measures.

But society runs on division of labour, so a small number of people are indispensable: healthcare, yes, and also critical infrastructure — food and water supply, waste management, energy production. The authors make the point twice that protective

equipment serves two purposes for these workers, preventing infection and death, and "encourag[ing] them to keep doing their indispensable work."

The gap they identify is institutional. Healthcare professionals are the only group with a chance of getting regularly fit-tested and trained. Most essential sectors have no respiratory protection programme at all, and therefore no route to provide training or equipment.

One argument here is worth holding onto. RPE "enables individual resilience and agency." Unlike limits on gatherings or distancing, a respirator does not require quick government action or anyone else's compliance to protect you. That makes it a defence a careful subset of the population can mount on its own, which the authors suggest could both prevent collapse and reduce societal friction.

Properties of better products

The physical requirement is stated simply: protect the wearer from infection without interfering with their work. The authors acknowledge that medical masks protected well against early SARS-CoV-2 variants and that even cloth masks contributed to slower spread — and then argue that against a GCBR-level pathogen it is "much less likely that homemade masks could offer sufficient protection to make people feel comfortable enough to leave their homes."

On filtration, the number to hold is 5% or 6%: that is what gets through an N95-style FFR at the **most penetrating particle size (MPPS)**, 0.3 microns on average, assuming a perfect fit. Whether that is enough "may or may not" suffice depending on the pathogen's minimum infectious dose and the concentration and size of virus particles in the air. They recommend more research on protection where a small number of particles might be enough to infect.

On fit, the problems stack up:

- Non-powered half-mask respirators, both elastomerics and FFRs, need a seal around chin, cheeks and nose bridge. Matching product to face and learning correct donning is what a fit test is for.
- Just-in-time fit testing is not realistic in a GCBR. The CDC's estimate of up to 400 people per operator per 8-hour day is already ambitious in a well-run hospital and unrealistic in rural places with no equipment or trained operator on site.
- Negative-pressure devices actively "incentivize incorrect fit because leakage eases breathing." That is a design flaw, not a user error.
- KN95-style flatfold respirators with ear loops became popular with the public, and replacing head straps with ear loops decreases fit significantly, as investigated by Niu et al.

On the high end, CBRN gear is reusable and "would likely be effective in protecting against a GCBR-level pathogen." It did not catch on during COVID-19, though a study with self-selected participants in an Australian hospital found staff preferred the CBRN respirator to their usual N95, donned it faster, and needed only short training. Two downsides: it is incompatible with eyeglasses, and it has limited source control, so pathogens from an infected wearer can escape through the exhalation valve.

Their product recommendations follow: create more options that are highly effective, can be worn with glasses, and filter exhaled air. "The ideal product provides universal fit without fit testing or training and incentivizes correct donning, for example, by integrating an automated feedback mechanism." They also name additional transmission routes a future product should handle — contaminated surfaces, aerosols reaching the eyes — pointing to antimicrobial surfaces, easy disinfection, full-face designs, and modular designs that integrate with full-body PPE.

Social factors

Here the authors flip the COVID-19 experience on its head. Compliance "might not be the key challenge" in a GCBR, because the risk to the individual would be obvious. The challenge instead is convincing users

that the devices reliably protect them from death and severe disease.

The evidence: an integrative review of 40 surveys and interviews with healthcare workers found PPE is a relevant factor in willingness to work during biological events. During COVID-19 people tolerated suboptimal PPE with substantial residual risk, such as surgical or cloth masks, because severe disease was unlikely for healthy young individuals. That tolerance should not be expected to hold when exposure means probable death. And a survey by Hines et al. found reusable RPE — PAPRs and elastomerics — was more popular among healthcare workers in high-risk scenarios despite lower comfort ratings.

So the recommendation is two-sided and unusually candid: communicate clearly about efficacy and evidence in government messaging, and "design products that invoke trust in their effectiveness, even if trading off against comfort or style," while collecting the scientific evidence that justifies that trust. The feeling of protection is treated as a design requirement, not a bias to be corrected.

Why stockpiles, not manufacturing

The events of the first weeks determine whether an outbreak becomes a catastrophe, so speed decides everything. Just-in-time manufacturing was adequate for COVID-19 and will not be adequate for a GCBR: a

deliberately released agent might be seeded at multiple sites and optimised for transmissibility. Their reference point for how much worse transmissibility can get is measles, with a basic reproduction number between 12 and 18, against about 9.5 for Omicron.

The COVID-19 response did scale existing production and build ad hoc supply chains for disposables, which they call remarkable — "but it took valuable time." A second problem is informational: during a fast GCBR, authorities will have less capacity to issue iterative updates on which products to use, and "gradually increasing the recommended level of protection would require too many victims." You have to choose the best product before the event.

Then the fragility of manufacturing itself: border closures affect transport, sudden demand exposes bottlenecks such as meltblown polypropylene for FFRs, and military conflict can destroy infrastructure outright. Hence protected stockpiles of reusable, durable products, plus decentralised manufacturing capacity.

Incentivising innovation

Their diagnosis is that existing products were designed for other jobs. FFRs, elastomerics and CBRN gear "have not been designed to be used by untrained workers in the face of a deadly pandemic threat," and

different use cases trade off against each other in design.

Financial levers: there is currently limited monetary incentive to optimise for highly effective RPE. COVID-19 prompted many innovative ideas on style and practicality but none have been widely adopted. Competitions such as the Mask Innovation Challenge can crowdsource ideas; advanced market commitments like COVAX can fund the move from promising prototype to widespread adoption; and agencies such as BARDA in the US and HERA in the EU, or a partnership modelled on CEPI, could fund testing and iteration.

Regulatory levers: current ISO and CEN standards "do not adequately cover the use case of respiratory protection in a GCBR." European testing protocols rely on people experienced in wearing respiratory protection, whereas in a GCBR untrained people must make as few errors as possible. Dropping tests for dust, radioactive fallout and chemical vapour from a pandemic-specific standard would open space for innovation — for instance RPE that inactivates pathogens with ultraviolet irradiation or nanomaterials rather than filtering them.

Testing levers: protocols should model real-world effectiveness better than filtration of sodium chloride particles or latex spheres. They float human challenge trials with safe airborne pathogens that cause no

severe symptoms and do not transmit between humans. And they propose borrowing the **Pearl Index** idea from contraception — publishing a range that captures ideal use against imperfect use, on an ongoing basis.

The stockpiling strategy

Innovation takes time, and improving current stockpiles of disposable gear is urgent, so adding elastomeric respirators to stockpiles "could be a good interim solution." As products expire and are replaced, that itself creates demand for better ones.

To size a stockpile you need to know how many people require protection. The authors ask governmental health and defence agencies to work out how many essential workers need high-end RPE to maintain the absolute minimum critical infrastructure, stockpile enough for all of them, and distribute on a defined trigger. They add that a small fraction of critical infrastructure workers could be encouraged to keep durable equipment at home or in local offices, and that decentralised storage and local supply chains deserve consideration.

The conclusion is three sentences long: lower levels of protection sufficed in COVID-19 and will not defend against pandemics of unprecedented fatality; the essential workforce needs RPE that is reliable, durable and instantly available; governments need to

incentivise innovation and stockpile reusable products. Sandbrink's doctoral research is funded by Open Philanthropy, which the paper discloses.

The figure

The paper's one table is its argument in miniature. Each row pairs a feature of a GCBR pathogen with what follows for respiratory protection and what to do about it: more lethal, so N95-level protection could be insufficient even if fitted perfectly, so develop and stockpile products beyond N95; scarier, so workers need justified trust, so build measures of real-world effectiveness beyond filtration efficiency; spreads faster, so products must ship from day one, so stockpile the right product; disrupted supply chains, so stockpile reusables; and the ideal product does not exist yet, so incentivise innovation through regulation and prizes. Read down the middle column and you have the paper.

Words to know

- **Respiratory protective equipment (RPE)** — the respirator part of PPE, as distinct from gloves, gowns and goggles. The paper is only about this.
- **GCBR (global catastrophic biological risk)** — a worst-case biological event that could threaten the stability of society. The whole argument hangs on designing for this rather than for another COVID-19.

- **Filtering facepiece respirator (FFR)** — the disposable moulded respirator: N95 and N99 in the US, FFP2 and FFP3 in Europe. What most stockpiles contain.
- **Case fatality ratio (CFR)** — the share of diagnosed cases that die. Used here to rank how much worse than COVID-19 a pandemic could be.
- **Most penetrating particle size (MPPS)** — the particle size a filter is worst at catching, on average 0.3 microns. Filter ratings are set at this size, which is why "95%" is a floor rather than a typical performance.
- **Fit testing** — the procedure that checks a particular respirator seals on a particular face. The paper's central bottleneck: up to 400 people per operator per day is not a pandemic-scale process.
- **Source control** — filtering what the wearer breathes out, so an infected wearer does not infect others. CBRN respirators are weak at this because of the exhalation valve.
- **Advanced market commitment** — a funder's binding promise to buy a product if someone builds it to specification, used to pull a product through development. COVAX is their example.

Numbers in this reading

- 11%: SARS's case fatality ratio, the comparison that would have made COVID-19 a GCBR.

- 32.7%: case fatality ratio for Middle East respiratory syndrome.
- 52.6%: case fatality ratio for avian influenza.
- 5% or 6%: the share of particles at the most penetrating size that gets through an N95-style FFR assuming perfect fit.
- 0.3 microns: the most penetrating particle size, on average.
- Up to 400 people in an 8-hour workday: the CDC's estimate for one experienced fit-test operator with appropriate logistics.
- 40: the number of surveys and interviews with healthcare workers in the integrative review linking PPE to willingness to work.
- 12 to 18: the basic reproduction number of measles.
- About 9.5: the basic reproduction number of the Omicron variant.

What the author is unsure about

A lot, and they say so. They cannot predict the exact properties of a GCBR-level pathogen. It is unknown how countries will adapt their pandemic preparedness plans. It is "unclear if commonly used N95-style FFRs offer enough protection for a highly lethal, highly infectious pathogen," because that depends on an unknown minimum infectious dose — hence the call for more research on protection against very small

infectious doses. The claim that CBRN respirators would work against a GCBR pathogen is hedged as "likely." The expectation that compliance will not be the main problem is explicitly a forecast about human behaviour under conditions nobody has observed. And the testing instruments they want — human challenge trials with safe airborne pathogens, a Pearl Index for respirators — do not exist yet; they are proposals.

For the discussion

- The most useful reframe in this paper is that the feeling of protection is a specification, not a soft issue: workers rated reusable respirators less comfortable and still preferred them in high-risk scenarios, so comfort is not the variable to optimise when the pathogen is lethal.
- The fit-testing arithmetic is the sharpest policy point available: at 400 people per operator per day, fit testing cannot be the plan for a whole country's essential workforce, so either the products stop needing a fit test or the plan is fiction. (The course page 4.1 puts a number on that workforce, citing the Federation of American Scientists: about 50 million essential workers in the United States, against a Strategic National Stockpile that as of 2022 held about 540 million disposable N95 respirators, barely enough for 10 days.)

Indoor Air Quality Is the Next Great Public Health Challenge

Juan Cambeiro and Brian Potter · 2023

[Read the original paper](#) → · the course assigns: whole piece, including the appendix on metrics and building codes

We ended the era of city-destroying fires by rewriting building codes, and the same move — standards, enforcement, incentives — could do the same for airborne infection, because the problem is not the virus but the air in our buildings.

Average summary

Cambeiro and Potter, writing for the Institute for Progress, build the whole piece on one analogy. American cities used to burn: the 1871 Great Chicago Fire destroyed more than 17,000 buildings, left about 100,000 people homeless and killed about 300. At the time this was thought to be the unavoidable price of having a city. Then codes restricted flammable materials, set electrical standards, and required fire doors, escapes, sprinklers and alarms. The last major urban fire in the US was the Great Baltimore Fire of 1904, and over the past 100 years the risk of fire death has fallen by roughly 90%.

Their claim is that indoor air is the same kind of problem. Humans spend 90% of their time indoors,

person-to-person pathogens overwhelmingly spread indoors, and the costs are enormous: US excess deaths over 1.3 million as of June 2023 and \$16 trillion for COVID-19, plus more than \$100 billion a year — roughly 0.6% of US GDP — for flu and other respiratory infections in normal years. Poor ventilation also raises carbon dioxide, which has its own effects; one Harvard study found cognitive function scores 61% higher in better-ventilated space.

The regulatory diagnosis is sharp. Almost every air quality law in the US addresses outdoor air. There is no "Indoor Clean Air Act." The only legally enforceable federal indoor limits come from OSHA, cover pollutants rather than pathogens, and are outdated and inadequate by OSHA's own admission. The dominant standards, **ASHRAE 62.1 and 62.2**, define acceptable indoor air as air with no known contaminants at harmful concentrations that at least 80% of occupants do not complain about — and the complaint survey is largely about thermal comfort, not infection. They are also design standards, so they describe the building on the day it opens, not the day you are in it.

Three tools can clean indoor air: **ventilation** (replace it), **filtration** (strain it) and **disinfection** (inactivate what is in it, mainly with UV light). The authors want all three, measured in air changes per hour.

Their policy ask: require 5 air changes per hour or the equivalent in high-occupancy buildings, make the standard apply to existing buildings and not just new ones, require re-inspection of mechanical systems every 3 years, pay jurisdictions to adopt it, and fund far-UVC research and trials.

The payoff they calculate: more than 50% of airborne infections happen in just 5% of buildings by floor area, so fixing the worst high-occupancy buildings could cut indoor infections by roughly 17%, worth nearly \$20 billion a year against five to six times less in cost.

Why it's assigned

This is the withstand layer at the level of the building rather than the body. The course page says air filtration would be valuable in a pandemic and also pays off in normal times, and that far-UVC and glycol vapours are promising but not yet proven; this reading is where those claims get their numbers, their units and their politics. It is also the reading that teaches the vocabulary — air changes per hour, single-pass efficiency, clean air delivery rate — that the rest of the day assumes.

In depth

The fire analogy, and what it is doing

The piece opens with large urban fires as a solved problem: 1,800 buildings destroyed in Portland, Maine in 1866, more than 17,000 and about 300 deaths in Chicago in 1871. Contemporaries treated this as intrinsic to cities, a tinderbox waiting to be set alight.

What changed was not firefighting but construction: codes restricting flammable materials in tall buildings, electrical safety standards, then fire doors, escapes, automatic sprinklers and alarms as design features. Result: the last major US urban fire was Baltimore in 1904, the risk of fire death has fallen roughly 90% in a century, and fires today are almost always contained to the room or building where they start.

The analogy sets up the thesis: indoor air is a public health crisis caused by the built environment, and the same combination of standards and technology could prevent future pandemics and reduce seasonal infections.

What is wrong with the air

Four contaminant classes, each with its own evidence.

Pathogens. People spend 90% of their time indoors — more, the authors note, than some whale species spend underwater. Person-to-person pathogens tend to transmit indoors: moderate evidence for pathogenic bacteria and fungi, extensive evidence for respiratory viruses. Bioaerosols rarely infect people

outdoors, and hunter-gatherer ancestors living outdoors were not troubled by most of the transmissible pathogens we deal with now. The science shifted during COVID-19: before it, most transmission was thought to happen through large droplets that fall out of the air quickly, but newer studies convincingly show that smaller aerosols, which stay aloft for minutes or hours, are an important vector, and inhalation of aerosols is likely the main route for respiratory viruses. Outdoors these aerosols dilute or are inactivated by sunlight; indoors they accumulate.

Carbon dioxide. CO₂ is a by-product of breathing and builds up in poorly ventilated rooms, causing headaches, drowsiness and impaired cognition. A Harvard study found cognitive function scores 61% higher with increased ventilation compared with a conventional building environment. A 2022 study in schools found a one standard deviation increase in the school-term average daily peak CO₂ led to a 0.11 standard deviation decrease in subsequent test scores — and the ventilation improvements cost about \$10 per student per year, less than 0.1% of average annual per-student spending in US public schools.

Particulates and volatile organic compounds, from cooking, smoking, cleaning and outdoor pollution leaking in, linked to respiratory and cardiovascular disease and cancer, and to "Sick Building Syndrome,"

which appeared in the 1970s as buildings were made airtight for energy efficiency without matching changes to ventilation.

The cost. In the US, COVID-19 caused excess deaths over 1.3 million as of June 2023 and cost \$16 trillion — far more than every large urban fire in US history combined. The authors are careful to say COVID-19 was comparatively mild on the scale of possible pandemics: about 24 million dead globally as of June 2023, under 0.3% of the world's population, against an estimated 1 to 5% for the 1918 influenza pandemic. Even outside pandemics, flu and other respiratory infections cost more than \$100 billion a year in medical costs and lost work, roughly 0.6% of US GDP; a UK analysis put the figure at about 1% of GDP a year.

Who decides, and why nothing happens

The Clean Air Act sets outdoor standards. There is no indoor equivalent. The only legally enforceable federal indoor limits are OSHA's workplace exposure limits for pollutants, not pathogens, which OSHA itself admits are outdated and inadequate.

State building codes do address indoor air quality, but not pathogen transmission or CO₂. The real authority sits with **ASHRAE** — the American Society of Heating, Refrigerating and Air-Conditioning Engineers — whose standards 62.1 (commercial) and 62.2

(residential) set minimum ventilation rates. Their definition of "acceptable" indoor air quality is worth reading twice: air with no known contaminants at harmful concentrations, and with which 80% or more of exposed people do not express dissatisfaction. The known contaminants in question are almost entirely bioeffluents, pollutants or chemicals with OSHA limits. The dissatisfaction survey does not ask about infectious disease at all — occupants are not asked whether they get more colds since they started using the space — and is largely measured through thermal comfort. The 2022 edition acknowledges that acceptable air quality takes more than minimum ventilation, but does not emphasise infection control.

At the time of writing there were no infection control standards from ASHRAE or anyone else aimed at public buildings, though ASHRAE had announced the imminent publication of **Standard 241, Control of Infectious Aerosols**: technology-neutral, evidence-based, built around a target equivalent clean airflow rate achieved through any mix of ventilation, filtration and disinfection, with testing and safety requirements for air cleaners, and an "infection risk management mode" that raises controls during periods of elevated risk — explicitly modelled on how fire controls activate during a fire.

The second structural problem is that 62.1 and 62.2 are design standards: they say what performance

should be achieved at construction. In practice buildings run well below design. CO₂ works as a proxy — anything above 1,000 parts per million would usually fail ASHRAE — and a review article finds many classrooms often exceed 2,000 or even 3,000 ppm. Office measurements routinely find ventilation below recommended rates. A UK estimate based on expert interviews concluded 50% of buildings had an inadequate ventilation system. And commercial buildings often shut their HVAC off in the evening with people still inside.

The three tools, and the units that measure them

Ventilation replaces indoor air with outdoor air, and reduces pathogens, CO₂ and everything else. It is measured in **air changes per hour (ACH)**: how many times an hour the full volume of air in the room is replaced. The authors' feel for the scale: an ACH of 10 cuts contaminant concentration by 99.9% in minutes; an ACH of 2 takes hours. Required rates are typically between 1 and 4 ACH. Going from 2 ACH, typical for offices, to 6 ACH brings that 99.9% reduction down from 207 minutes to 69 minutes. You can raise either the outdoor air fraction or the total ventilation rate, and both cost energy, since incoming air has to be heated or cooled; energy recovery systems help but are expensive and hard to retrofit. Two honest caveats: the arithmetic assumes air mixes completely and uniformly, which real rooms do not, with stagnant

zones and uneven emission sources; and mechanical equipment, comfort and the exponential nature of pathogen concentration all put practical ceilings on how far you can push it.

Filtration strains contaminants out and is measured in **effective air changes per hour (ACHe)**: the number of real air changes that would have achieved the same removal. The COVID-19 success story is the **Corsi-Rosenthal box** — four or five MERV 13 filters taped into a cube around a box fan, buildable from hardware store parts for less than \$100, which a 2022 study found cost roughly 10 times less than commercial air filters. The weakness of existing systems is the filter grade: ASHRAE 62.1 and 62.2 typically require a MERV-8 filter, which catches most particles larger than 3 microns — much larger than pathogens, which pass straight through. HEPA filters do capture airborne viruses including influenza and COVID-19, but an HVAC system may not be able to push air through the denser media, which cuts airflow and raises energy use, and HEPA needs more frequent replacement.

Disinfection inactivates pathogens in place, also measured in AChE. Ultraviolet light is the main method. The wavelength map: UV-A 315 to 400 nm, UV-B 280 to 315 nm, UV-C 200 to 280 nm, with UV-C the most germicidal because it damages the DNA and RNA of microorganisms so they cannot replicate or

infect. Because most UV exposure harms people, systems are designed to keep it away from them: in-duct UV, which they call complex and not especially promising, and upper-room UV, mounted near the ceiling behind shields, which has achieved dozens of ACHe and depends on air mixing to bring room air up into the beam. Then **far-UVC**, at 200 to 230 nm rather than the conventional 254 nm, which cannot penetrate past the outer layer of dead skin cells or the tear layer of the eye and so can run in occupied rooms. Laboratory studies show effect against bacteria (*Staphylococcus aureus*), fungi (*Aspergillus niger*) and viruses (influenza A, COVID-19). The headline study they cite: continuous far-UVC in occupied public locations within current regulatory limits, in a room with 3 air changes per hour and five filtered sources, cut the steady-state pathogen load by 98.4% — an additional 184 equivalent air changes per hour.

The size of the prize

Infections cluster. An analysis for the UK Royal Academy of Engineering estimated that more than 50% of airborne infections occur in just 5% of buildings by floor area: airports, schools, theatres, gyms, churches, exhibition halls, which are both crowded and badly ventilated, typically at one to two air changes per hour. A 2017 study of recently constructed US schools found an average of 2 ACH. A

1989 study of US office buildings found every building below 2 ACH and most below 1.

The UK analysis estimated that in poorly ventilated buildings, going to six air changes per hour cuts the spread of airborne-transmitted disease by 30 to 50%. Apply that only to the highest-occupancy buildings and the authors estimate indoor infections overall could plausibly fall by roughly 17%. Against \$100 billion a year in airborne infection costs, that is nearly \$20 billion saved annually — five to six times the cost of implementation, with no technological improvement assumed. Separately, they put the productivity loss from elevated CO₂ at 1 to 10% in poorly ventilated buildings, and cite Macomber and Allen's estimate that if better cognitive performance translated directly into revenue, the benefit of doubling the ASHRAE recommended ventilation rate would be 75 times the extra energy cost.

What they want done

Three groups of recommendations.

Better standards. Require an ACH or ACHe of 5 in all high-occupancy buildings in ASHRAE 62.1 and 62.2 — 5 being a recently specified CDC target achievable with existing equipment. Combine ventilation with filtration and disinfection rather than relying on airflow alone, with minimum requirements for air-cleaning devices. Harmonise Standard 241's safety

procedure A1.7 with 62.1's Indoor Air Quality Procedure, because as written A1.7 sets limits on air-cleaner emissions that are not sourced to recognised authorities and could force a testing authority to switch off air cleaners that are effective and pose no health risk. And require "commissioning" of mechanical systems every 3 years, the way fire sprinklers are inspected, because performance degrades.

Get it into buildings. Most model codes already reference 62.1 and 62.2, so jurisdictions may adopt updates on their normal 3-to-6-year cycle — but provisions seen as adding cost get stripped before adoption, and waiting for the cycle is slow. So the federal government should pay jurisdictions to adopt the updated standards and 241, potentially cost-neutral given the savings. And because existing buildings outnumber new ones roughly 100 to 1, the requirements have to apply to existing buildings too, with a reasonable compliance timeline.

Fund the technology. More human studies outside the lab, especially large-scale epidemiological studies of far-UVC's real-world efficacy. Standards that allow safe use rather than inhibiting devices that pose little risk. Support for cheaper, more easily manufactured far-UVC devices, including LEDs and second harmonic generation devices. Public education on how far-UVC

works and why it is safe. International collaboration and global standards.

The appendix, which is where the measurement problem lives

The appendix is the honest part of the paper and worth the reading time. **Single-pass removal efficiency** is the percentage of contaminants a cleaner removes in one pass, which is what MERV, ePM, HEPA, FPR and MPR ratings describe — higher is better, but the number alone tells you nothing until you know the airflow rate and the room volume, because what matters is how often the room's air goes through. **Clean air delivery rate (CADR)** measures the volume of clean air delivered per unit time, typically 50 to 500 cubic feet per minute, tested under AHAM AC-1 using smoke, dust and pollen particles from 0.09 to 11 microns, with about 800 products listed on AHAM's site; as a rule of thumb CADR is about two-thirds of maximum airflow, and ACH or ACHe equals CADR in cubic feet per minute times 60, divided by room volume in cubic feet. Safety has its own metrics: ozone emissions under UL 867 for low ozone (under 50 ppb) and UL 2998 for near-zero ozone (under 5 ppb), with no standard test method at all for other chemical by-products.

On measurement in practice, their verdict is that no single accessible method covers every intervention. CO2 monitors are the common tool, cost a few

hundred dollars, and can estimate ACH either from steady-state concentration or from the decay rate after a room empties — a normal-sized classroom with 15 teenagers at 4 ACH should sit around 800 ppm. But CO₂ is a poor indicator in sparsely occupied rooms, and it cannot see filtration or UV at all, because those remove microorganisms and not CO₂. ASHRAE has no indoor CO₂ limit. Belgium is the exception they hold up: a national ventilation plan for commercial spaces open to the public, requiring a CO₂ meter visible to the public and an inventory of readings, with Level A at under 900 ppm (at least 40 m³ per hour per person) and Level B at under 1,200 ppm (at least 25 m³ per hour per person), plus air quality labels from 2025.

They also air two ideas they are not sure about. The **path principle** holds that what matters in a mechanically ventilated room is the path between the contaminant source and the exhaust vent, not the total air changes — the worst case being an exhaust far from the source with air streams cutting across the path. That would mean designing to push contaminated air up toward ceiling exhausts, like a fume hood, rather than mixing room air for uniform temperature. And real-time bioaerosol monitoring: viable particle counters that tell live from dead bioaerosols by fluorescence, currently hundreds of thousands of dollars with tens of thousands a year in upkeep and historically not used for viruses;

sequencing of HVAC condensate; an automated on-site PCR device, which does not yet exist at feasible size and cost; and aerosol alarm systems, which may not distinguish viable particles and may drift. What is needed, they say, is to measure air quality directly and in real time, the way we already do for temperature.

The final appendix section tours how US building codes actually get adopted — the Authority Having Jurisdiction, model codes from the ICC, IAPMO and NFPA, hard paths through mandatory code and soft paths through incentives — with one deflating data point: LEED, by far the most successful voluntary standard, is used on approximately 20% of commercial building square footage after 20 years.

Words to know

- **Air changes per hour (ACH)** — how many times an hour the room's whole volume of air is replaced with outdoor air. Higher is better. 1 to 2 is a typical crowded building, 5 is the CDC target the authors want required, 10 clears a room by 99.9% in minutes.
- **Effective (or equivalent) air changes per hour (ACHe)** — the same yardstick applied to filtration or UV, which clean the air in place instead of replacing it: how many real air changes would have done the same job. This is what lets you add a filter and a lamp and a ventilation rate together.

- **Single-pass removal efficiency** — the share of contaminants a filter or cleaner removes each time air goes through it, which is what MERV and HEPA ratings describe. MERV-8, the usual HVAC filter, catches most particles above 3 microns and misses pathogens; HEPA catches them.
- **Clean air delivery rate (CADR)** — the volume of clean air a portable cleaner delivers per minute, typically 50 to 500 cubic feet per minute. Efficiency times airflow; it converts to ACH_e once you know the room's volume.
- **ASHRAE 62.1 and 62.2** — the ventilation standards nearly every US building is designed to, for commercial and residential buildings. They are the lever the authors want pulled, and they currently say nothing about infection.
- **ASHRAE Standard 241** — the new standard for Control of Infectious Aerosols, technology-neutral, with a heightened "infection risk management mode" for periods of elevated risk.
- **Far-UVC** — germicidal ultraviolet at 200 to 230 nm rather than the conventional 254 nm, absorbed by the outer dead layer of skin and the tear film, so it can in principle run in occupied rooms.
- **Commissioning** — testing and tuning a building's mechanical systems to confirm they perform as designed. Done once at construction today; the authors want it every 3 years, like a sprinkler inspection.

Numbers in this reading

- 1,800 buildings destroyed and about 10,000 people homeless: the 1866 Great Fire of Portland, Maine.
- More than 17,000 buildings, about 100,000 homeless, about 300 killed: the 1871 Great Chicago Fire.
- 1904: the Great Baltimore Fire, the last major urban fire in the US; roughly 90% fall in fire death risk over the past 100 years.
- 90%: the share of their time humans spend indoors.
- 61%: how much higher cognitive function scores were with increased ventilation in the Harvard study.
- 0.11 standard deviations lower test scores per one standard deviation rise in school-term average daily peak CO₂; about \$10 per student per year for the relevant ventilation improvements, under 0.1% of average annual per-student spending.
- Over 1.3 million US excess deaths as of June 2023 and \$16 trillion in cost for COVID-19; about 24 million deaths globally, under 0.3% of world population, against 1 to 5% for the 1918 influenza pandemic.
- More than \$100 billion a year, roughly 0.6% of US GDP, for flu and other respiratory infections; about 1% of GDP in the UK analysis.

- 80% or more: the share of occupants who must not express dissatisfaction for air to count as "acceptable" under ASHRAE.
- 1,000 ppm CO₂: roughly the ASHRAE pass mark; classrooms often exceed 2,000 or even 3,000 ppm; 50% of UK buildings judged to have inadequate ventilation.
- ACH of 10 clears 99.9% in minutes; ACH of 2 takes hours; typical requirements 1 to 4 ACH; moving from 2 to 6 ACH cuts the time for a 99.9% reduction from 207 minutes to 69 minutes.
- Under \$100 for a Corsi-Rosenthal box, roughly 10 times cheaper than commercial air filters in a 2022 study.
- MERV-8 filters most particles greater than 3 microns, which is far larger than infectious aerosols.
- UV-A 315–400 nm, UV-B 280–315 nm, UV-C 200–280 nm; conventional germicidal UV at 254 nm; far-UVC at 200–230 nm.
- 98.4% reduction in steady-state pathogen load and an additional 184 equivalent air changes per hour: the far-UVC study, in a room at 3 ACH with five filtered sources.
- More than 50% of airborne infections in just 5% of buildings by floor area.
- Average 2 ACH in recently built US schools (2017 study); every office below 2 ACH and most below 1

(1989 study).

- 30 to 50%: the reduction in airborne disease spread from taking a poorly ventilated building to 6 ACH; roughly 17% fewer indoor infections overall if the highest-occupancy buildings are targeted; nearly \$20 billion a year saved, five to six times the cost.
- 1 to 10%: the productivity loss attributed to elevated CO₂ in poorly ventilated buildings; 75 times the extra energy cost, in Macomber and Allen's estimate of doubling recommended ventilation.
- ACH or ACH_e of 5: the standard they want required, a recently specified CDC target.
- Every 3 years for commissioning; 3 to 6 years for code update cycles; existing buildings outnumber new ones roughly 100 to 1.
- Belgium's targets: under 900 ppm CO₂ with at least 40 m³ per hour per person (Level A), or under 1,200 ppm with at least 25 m³ per hour per person (Level B); labels from 2025.
- 50 to 500 cubic feet per minute: typical CADR; about 800 products tested and listed by AHAM; CADR is about two-thirds of maximum airflow.
- Under 50 ppb ozone for UL 867, under 5 ppb for UL 2998; no standard test method for other by-products.

- Hundreds of thousands of dollars for a real-time viable particle counter, with tens of thousands a year to keep it calibrated.
- About 20%: the share of US commercial building square footage using LEED after 20 years.

What the author is unsure about

The authors are confident about ventilation and filtration and careful about everything newer. Their far-UVC recommendation is explicitly a research recommendation: more human studies outside the laboratory, "especially large-scale epidemiological studies," are needed to evaluate real-world efficacy, and the next generation of devices needs work on lifetime, efficiency, safety and output. They present the path principle as an alternative worth exploring, backed by "some evidence" and simulations, not as settled. Every real-time measurement option they review comes with a list of defects — cost, false positives, sensor drift, inability to tell live particles from dead ones, no track record with viruses — and they conclude that no single accessible method can account for all interventions. The 17% infection reduction and the \$20 billion are built on a chain of other people's estimates, which they label as plausible rather than measured. And they flag that by-product formation from additive air cleaners has no standard test method at all, so manufacturer claims are hard to assess without independent verification.

For the discussion

- The analogy is the argument: nobody today accepts a building that cannot contain a fire, and the question for the day is what it would take for a building that cannot clear its air to become equally unacceptable.
- The sharpest fact for prioritising is that more than 50% of airborne infections happen in 5% of buildings by floor area, which turns a vast retrofit problem into a targeted one — schools, transit hubs, gyms, churches first.

4.1-4 · CORE · COURSE TIME 5 MIN

Far-UVC

Jeff Kaufman · 2025

[Read the original paper →](#) · the course assigns: whole page, including the buying advice and the research list

A short, practical explainer on far-UVC: what it is, why it can run in a room with people in it, and what to look for if you want to buy a lamp that actually works.

Average summary

This is a one-page resource site, written by Jeff Kaufman in his personal capacity, aimed at someone deciding whether to install far-UVC. It is the most

concrete reading on the page and the only one that tells you what to buy.

The setup is the gap in existing tools. We share air, and when someone sick breathes, talks, coughs or sneezes, they release tiny infectious particles that can float for hours. COVID-19 made the airborne route visible, but colds and flu travel the same way.

Ventilation and air purifiers both work, and both have limits: ventilation usually does not remove dirty air fast enough and is expensive to expand, while air purifiers work well when used as designed but get run on low or switched off because of the fan noise. What is wanted is something that cleans the air continuously without an expensive retrofit or a noisy fan.

Far-UVC is a kind of light that continuously inactivates viruses and bacteria in the air and on surfaces. It installs quickly and runs with almost no sound. Kaufman names the benefits he expects in ordinary life rather than in a pandemic: fewer students absent from school, fewer hospital-acquired infections, steadier workplace productivity through cold and flu season.

The mechanism is the interesting part. Ultraviolet light covers a range of wavelengths, and the effect on both microbes and human tissue depends strongly on where in that range you are. Far-UVC's main effect on viruses and bacteria is to damage their proteins so

they can no longer infect us — and that same strong protein absorption is what makes it safer for skin.

Unlike the UVA and UVB in sunlight, which reach deep enough to cause lasting damage, far-UVC is absorbed by proteins in the dead outer layers of skin and never reaches the living cells underneath.

Eyes are different, because they have no dead outer layer. They are protected instead by eyebrows, eyelids, lashes and the tear film, which cut the fraction of light getting in, and by the outermost layers of the cornea, which absorb what does get through and regenerate quickly. Kaufman says plainly that eyes are at higher risk than skin, and that exposure limits and lamp placement rules are driven primarily by what is safe for eyes.

On evidence, he points to Görlitz et al. (2023), a review of risks and benefits that concluded the current evidence supports using far-UVC systems within existing guidelines, and notes that multiple studies have found no evidence of skin or eye damage even at levels well above what air cleaning requires.

Then the buying advice, which is where the page earns its place: the market is full of poor-quality products, including some that emit dangerous longer wavelengths, and a dead bulb that has stopped producing far-UVC will still glow purple as if it were working.

Why it's assigned

The course page says far-UVC is being explored as a promising way to cut indoor transmission while noting we do not yet have conclusive studies on safety and efficacy. This reading is the short version of the optimistic case, written for a practitioner rather than a policymaker, and it is the one that makes far-UVC feel like a thing you could install rather than a research programme. Paired with 4.1 #5, you get the practical view and the field-building view of the same technology.

In depth

The gap far-UVC is meant to fill

Kaufman frames the problem as shared air. Infectious particles released by one person can stay airborne for hours, and anyone breathing that air can pick them up. He treats this as a general fact about respiratory disease rather than a COVID-19 fact, listing colds and flu alongside it.

He then puts the two established tools in their place. Ventilation systems usually do not remove dirty air fast enough, and expanding them is expensive. Air purifiers do work as designed — his complaint is behavioural rather than technical: fans are loud, so people turn them down or off. That observation is doing real work in the argument, because it is why a technology that is almost soundless and needs no

retrofit could matter even if it were not strictly better at cleaning air.

How it works, and why it can be used around people

Ultraviolet light spans a range of wavelengths, and the biological effect at each wavelength is very different. Some kinds of UV are dangerous to humans; far-UVC is the much safer range.

The core mechanism: far-UVC's strongest effect on viruses and bacteria is to damage their proteins, which stops them infecting us. The same protein absorption explains the human safety margin. UVA and UVB in sunlight penetrate deep enough to cause lasting damage. Far-UVC is absorbed so strongly by proteins in the dead outer layers of skin that it does not reach the living cells beneath. He cites Figure 4 from Görlitz et al. (2023), which illustrates how little far-UVC at about 222 nm penetrates those outer dead layers.

Eyes get a more careful treatment, and this is the part worth remembering. They have no equivalent dead outer layer. Their protection comes from physical barriers — eyebrows, eyelids, lashes, tear film — which reduce the fraction of light getting in, plus biological resilience: what does get in is absorbed by the outermost layers of the cornea, which regenerate rapidly. He states the asymmetry rather than

smoothing it over: eyes are at higher risk from far-UVC than skin, and the guidelines for exposure limits and lamp placement are set primarily by what is safe for eyes.

On the safety evidence he gives one review and one claim. The review is Görlitz et al. (2023), which assessed risks and benefits and concluded that the current evidence supports using far-UVC systems within existing guidelines. The claim is that multiple studies have not found evidence of skin or eye damage even at levels much higher than air cleaning needs. The whole page's research list is made up of studies about safety rather than efficacy — mouse eye pathology, DNA damage across skin types, corneal safety, a 36-month clinical observation of ocular safety in a full-room installation, eye irritation in a simulated office, long-term effects in UV-susceptible mice, acute corneal damage comparing 222 nm with 254 nm, and the early Buonanno and Welch papers establishing germicidal efficacy with mammalian skin safety.

What to buy, and what to watch for

Far-UVC is most valuable in crowded indoor spaces where people spend extended time together. Because it works by light, its advantage over air purifiers is largest in bigger rooms, where a beam can travel further.

His warnings about the market are specific:

- There are many poor-quality products, and some emit dangerous longer wavelengths.
- A low-quality bulb that has stopped producing far-UVC will still glow purple, so the lamp looks like it is working when it is not. There is no way to tell by eye.
- Look for lamps evaluated by OSLUV and found to perform as designed, and for KrCl ("222nm") bulbs rated for thousands of hours.

Two examples: the **Aerolamp**, more expensive, built around the Ushio Care222 B1 module, a long-lasting design rated for 10,000 hours — one year of running continuously, five years of working hours, or longer if run less often — and the **Nukit Torch**. Both have published evaluations.

On installation: follow the manufacturer's instructions, including the minimum distance, and do not set lamps up where people will be looking straight into them for extended periods. Well-filtered 222 nm lamps produce minimal amounts of ozone, so manufacturers often recommend running an air purifier with an activated carbon filter as an added precaution.

The closing note is the one caveat that matters most for reading the rest of the literature: all the safety data he presents applies specifically to filtered 222 nm

krypton chloride excimer lamps. Other kinds, such as 233 nm LEDs, may have a lower safety margin — he cites Zamudio Díaz et al. (2025) — and should not be treated as equivalent.

Words to know

- **Far-UVC** — germicidal ultraviolet in a narrow band around 222 nm, strongly absorbed by protein, which is why it inactivates microbes without reaching living human tissue.
- **KrCl (krypton chloride) excimer lamp** — the lamp type that produces 222 nm light. "222nm" on a product listing means this. It is the only source the safety evidence here covers.
- **Filtered** — fitted with an optical filter that strips out the longer, more harmful wavelengths a KrCl lamp also emits. Unfiltered or badly filtered lamps are the hazard he warns about.
- **OSLUV** — the group whose independent evaluations he recommends checking before buying a lamp. Manufacturer claims are not enough, because a failed bulb still glows.
- **Rated hours** — the lamp's expected working life, 10,000 hours in the Aerolamp's case. It converts into real coverage: a year of continuous running, or five years of working hours.
- **Activated carbon filter** — a filter that adsorbs gases rather than particles, recommended here as

a precaution against the small amount of ozone a far-UVC lamp makes.

Numbers in this reading

- 222 nm: the wavelength of the lamps all the safety data applies to.
- 233 nm: the LED wavelength he warns may have a lower safety margin and should not be treated as equivalent.
- 10,000 hours: the Aerolamp's rated lamp life, which he converts to one year of 24/7 use or five years of working hours.
- Thousands of hours: the minimum rated bulb life to look for.
- Hours: how long infectious particles released by a sick person can stay floating in the air.

What the author is unsure about

The page is written to persuade, so its uncertainty is narrow but real. Eyes are explicitly flagged as being at higher risk than skin, with exposure limits set by eye safety rather than skin safety. The safety case rests on a review that endorses use within existing guidelines — not a claim that far-UVC is safe at any dose. The safety evidence is bounded to filtered 222 nm KrCl lamps, with other sources explicitly excluded. Ozone is acknowledged, at minimal levels, with a carbon filter suggested as a precaution. And the benefits he lists — fewer absences, fewer hospital-acquired infections,

steadier productivity — are stated as practical expectations, not as results from a trial. Note what the page does not contain: any real-world effectiveness study showing fewer infections. The research list is about safety.

For the discussion

- The reason this reading matters for a city is the failure mode he names: a dead far-UVC bulb still glows purple, so an institution can run a lamp that stopped working years ago and believe it is protected — which makes independent evaluation and verification part of the intervention, not an optional extra.
- His honest argument for far-UVC over air purifiers is not that it cleans better but that people leave it on, which is worth raising when the group compares technologies on cost per unit of clean air alone.
- Every study on his research list is a safety study, which is a useful thing to notice before the group accepts that far-UVC is ready: nobody on this page is claiming a measured drop in infections.

4.1-5 • CORE • COURSE TIME 5 MIN

Flipping the switch on far-UVC

Richard Williamson • 2025

[Read the original paper →](#) · the course assigns: whole piece

Far-UVC looks extraordinary in the lab and almost nobody uses it, and the reason is not physics: there is no regulator to vouch for it and no company that can afford to prove it works.

Average summary

This is a short post by Richard Williamson, written to announce the release of Blueprint Biosecurity's far-UVC report as a preprint on 10 March 2025. It is the clearest short statement of the field-building problem, and it is structured as a puzzle.

First, the promise. For all our progress against water-borne, food-borne and vector-borne disease, we remain badly exposed to germs that spread through the air. Far-UVC is an emerging type of **germicidal UV (GUV)** that researchers have spent a decade showing can eliminate pathogens while being safe for humans. The landmark 2022 study cut airborne bacteria in a room-sized chamber by **98.4 percent**, all while staying inside safe UV exposure limits.

Converted into the unit building engineers use, that is the equivalent of replacing all the air in the room 184 times an hour — against a CDC recommendation of 5 or more air changes per hour in a workplace, and a US requirement of only 20 for hospital operating rooms. Viruses are not the hard case: airborne coronaviruses are more susceptible to far-UVC than the bacteria used in that study.

Williamson says he called far-UVC the most promising new candidate for a pandemic firewall in these pages two years earlier, and that after thousands of hours in the literature and conversations with hundreds of experts across photobiology, atmospheric chemistry, indoor air quality, building science, environmental engineering, epidemiology and public health, he still stands by it.

Then the puzzle. Given those numbers, you would expect a booming market. Instead most people have never heard of far-UVC.

His answer has three parts. The lab studies are promising but nobody has produced the real-world data needed to support safety and efficacy claims. If far-UVC were a drug, that would simply mean it was not yet approved, and you could not buy one until trials were done. But there is no FDA for putting a lamp in an office, a school or a nightclub — which cuts both ways. Lamps are freely available to anyone, and there is also no body able to confer public trust in the technology's safety and efficacy. Nor is there a company that will fund expensive clinical research when every competitor free-rides on a good result; no for-profit entity can afford the standard of evidence that the risk-averse institutions capable of granting legitimacy will expect. And there is no coordinator between research, industry and those institutions, so

the limited research that happens sits in academic silos.

The report is the proposed fix: it identifies the critical research questions, beginning with work analogous to FDA Phase I and II trials — establishing the doses needed in real applications and completing the picture of biological and chemical effects — and then Phase III, cluster-randomised controlled trials for high-quality evidence of real-world effectiveness.

Why it's assigned

The course page states that we do not yet have conclusive studies on the safety and efficacy of far-UVC. This reading explains why that gap has persisted despite a decade of encouraging results, and it names the mechanism: no regulator, no funder, no coordinator. It is the piece that turns far-UVC from a technology question into an institutions question, which is the level the course teaches at, and it is the natural companion to the full Blueprint report on page 4.4.

In depth

The number, and the unit

The reading leans on one study: the 2022 result in which far-UVC reduced concentrations of airborne bacteria by 98.4 percent in a room-sized chamber while operating inside safe UV exposure limits.

What makes the post useful is that Williamson translates that into the language building engineers already use. A 98.4 percent reduction in steady-state concentration is equivalent, in terms of air cleaning, to changing over all the air in the room 184 times every hour. He then gives two reference points so you can feel the size: the CDC recommends 5 or more air changes per hour in the workplace, and even US hospital operating rooms only require 20. So the claim is not "far-UVC helps" but "far-UVC delivers an order of magnitude more clean air than the most demanding room in a hospital."

He adds that viruses are not the limiting case. Airborne coronaviruses are more susceptible to far-UVC than the bacteria used in that study, so the bacterial result is, if anything, conservative for the pathogens of most concern.

He also establishes his own position carefully. He works at Blueprint Biosecurity. Two years earlier, writing for the same publication, he described far-UVC as the most promising new candidate for building a pandemic firewall. Since then the organisation has spent thousands of hours on the academic literature and spoken with hundreds of experts across photobiology, atmospheric chemistry, indoor air quality, building science, environmental engineering, epidemiology and public health. His verdict after all

that is unchanged — which he states as a judgement, not as a finding.

Why nobody has one

The post's real contribution is the diagnosis, built as four linked failures.

No real-world data. Promising lab studies exist.

What is missing is the kind of compelling real-world evidence needed to support safety and efficacy claims.

No regulator, which is a double-edged sword. If far-UVC were a drug, the situation would be straightforward: unapproved, clinical trials required, nothing on sale. But no approval is needed to install a lamp in an office, a school or a nightclub. So lamps are freely purchasable — and the flip side of having no regulator is that there is no entity to confer public trust in the technology's safety and efficacy. The thing that would slow adoption in a regulated market is the same thing preventing legitimacy in an unregulated one.

No commercial funder. No corporation will make the upfront investment in expensive clinical research when other far-UVC manufacturers can free-ride on a successful outcome. Williamson states the structural conclusion plainly: no for-profit entity can afford to obtain the standard of evidence that risk-averse

institutions — the ones in a position to confer public trust — would expect.

No coordinator. Instead of the limited ongoing far-UVC research being pointed at a shared goal, it happens in academic silos. There is no body sitting between research, industry and the trusted institutions.

He then names what is genuinely novel about the problem: there is no clear precedent for a technology like far-UVC gaining the trust of the public. That is the problem the organisation set out to solve.

What the report proposes, and what was launched

The plan in the report is explicitly borrowed from drug development, which is a useful mental model for the whole of Unit 4. In addition to setting out what is known about how and how well far-UVC works, it identifies a handful of critical research questions that must be answered before trusted institutions can confidently endorse the technology. The sequence:

- First, work analogous to the goals of FDA Phase I and II trials: quantifying the doses of far-UVC needed in real applications, and making sure the understanding of its biological and chemical effects is complete.
- Then Phase III: cluster-randomised controlled trials, to obtain high-quality evidence of real-world

effectiveness.

Alongside the preprint, three things were launched the same day: a month-long call for public feedback on the plan and for additional external reviewers; a job listing for a central role managing the forthcoming far-UVC field-building work; and a Substack called Far-UVC Field Notes to document the progress of the technology and of their own work.

Words to know

- **Germicidal UV (GUV)** — ultraviolet light used deliberately to inactivate microbes. Far-UVC is the newest member of the family; conventional upper-room UV is the older one.
- **Air changes per hour** — how many times an hour a room's whole volume of air is replaced. It is the unit that lets an air cleaner, a lamp and a ventilation system be compared on one scale: 5 is the CDC's workplace recommendation, 20 is a US operating room, 184 is what the 2022 study implies.
- **Safe UV exposure limits** — the published ceilings on how much UV a person may receive. The whole claim depends on the 98.4 percent result being achieved under them, not above them.
- **Cluster-randomised controlled trial** — a trial that randomises whole places (buildings, wards, schools) rather than individuals, which is the only

way to test an intervention that protects a room rather than a person.

- **Free riding** — benefiting from research someone else paid for. It is the reason no single manufacturer will fund the trials.
- **Preprint** — a research document released publicly before formal peer review, which is the form the report took on 10 March 2025.

Numbers in this reading

- 98.4 percent: the reduction in airborne bacteria concentration in the landmark 2022 room-sized chamber study, achieved within safe UV exposure limits.
- 184: the equivalent number of complete air changes per hour that reduction represents.
- 5 or more air changes per hour: the CDC's workplace recommendation.
- 20: the air changes per hour required in US hospital operating rooms.
- Two years: the gap since the author's earlier piece calling far-UVC the most promising candidate for a pandemic firewall.
- Thousands of hours of literature review and hundreds of experts consulted: the basis for the report.
- 10 March 2025: the date the report was released as a preprint.

- One month: the window opened for public feedback on the plan.

What the author is unsure about

He is openly advocating, and the uncertainty he concedes is the central one: the real-world safety and efficacy data does not exist yet. The 98.4 percent figure is a chamber result, not a measured reduction in human infections, and his own report sets out the trials that would be needed to get that. His claim that far-UVC remains the most promising candidate is presented as a standing judgement after review, not as a finding. The comparison with FDA phases is an analogy he is choosing, since no regulator actually requires any of it. And he does not claim the plan will work — the open call for feedback and external reviewers is, in effect, an admission that the plan itself is a draft.

For the discussion

- The sharpest line in the reading is that having no regulator is a double-edged sword: anyone can buy a lamp today, and precisely because nobody has to approve one, there is no institution able to tell the public it is safe.
- The free-rider problem explains the whole stall, and it is the standard argument for public or philanthropic money: a manufacturer that funds the

trial hands the benefit to every competitor, so the trial does not get funded by anyone.

- Worth putting to the group: 184 air changes an hour against an operating room's 20 is the most impressive number in the unit, and it comes from one chamber study with bacteria, which is exactly the kind of single result the rest of the day teaches you to be careful with.

4.3-1 · CORE · COURSE TIME 65 MIN

Blueprint for Pandemic PPE

Blueprint Biosecurity · 2024

[Read the original paper →](#) · the course assigns: executive summary and introduction in full (printed pages 2 to 6), plus every chapter's opening and recommendations for the chapter guide

Against a pandemic as infectious as measles, as fast as Omicron and as deadly as 1918 flu, the disposable N95 is the wrong product: the world needs reusable elastomeric respirators, a 150-day stockpile, and a manufacturing base that can start before the outbreak is called.

Average summary

This is a four-phase study by Blueprint Biosecurity, built from literature, government plans, after-action reports, interviews with manufacturers and designers,

and two stakeholder workshops. The question: what would it take for the world's **vital workers** — the people who cannot work from home and without whom society stops — to have the protective equipment they need in the next pandemic.

Phase 1 catalogued what went wrong. Vulnerability to shortages of key materials and machines, international trade disruption, inadequate stockpiles, and designs that reduce protection or usability. Two shortcomings stand out. Fit: low-cost respiratory protection only works if it seals, and workers struggle to get a good fit, keep it through a shift, and know whether they still have it. And the mindset of making the workforce adapt to the PPE we happen to have, rather than adapting PPE to a diverse workforce — tight-fitting respirators cannot accommodate facial hair, headwear or assistive devices, and body-covering PPE was not built for women's bodies.

Phase 2 modelled a plausible worst case: a virus **as infectious and hardy as measles, spreading as fast as SARS-CoV-2, as deadly as the 1918 influenza**. The authors note this is not science fiction — rinderpest, a deadlier sibling of measles, evolved exactly those properties, and only chance sent it to cattle instead of us. The findings: indoor workers who spend a shift with infected people need **more protection than a disposable N95**; outdoor workers are adequately protected by a well-fitting disposable

N95; and anyone in direct contact with infected people needs barrier protection too. Within 20 days of emergence the world's vital workers would need more than 1.4 billion units of respiratory protection, 1.1 billion gloves and 43 million gowns per day — 10 to 100 times what industry can supply.

Phase 3 wrote requirements. Fit must be achievable, measurable, maintained and visibly lost. Body-covering PPE should be designed to the female worker, which also fits men. Respiratory protection must work for religious, cultural and functional needs.

Phase 4 is the plan. The central finding: **elastomeric half-mask respirators (EHMRs) are vastly superior to disposable N95s** for pandemic preparedness — better protection, better fit, better lifecycle cost — with **powered air purifying respirators (PAPRs)** in the mix for people who cannot get a seal.

Governments should run early detection that tells manufacturers to ramp up before a pandemic is declared and guarantees to buy the surplus if it fizzles; pay to keep standby capacity and trained reserve staff ready; hold a **150-day stockpile**, with respiratory protection centralised and government-managed at 90% EHMRs and 10% PAPRs; split barrier PPE across industry and government; monitor supply chains; on-shore or friend-shore production and precursor materials; and pre-approve alternate raw materials.

Their closing arithmetic: COVID-19 cost the US an estimated \$16 trillion, and everything they propose for the whole world costs under \$50 billion.

Why it's assigned

This is her deep-dive option, and it is the withstand layer applied to the one group that cannot shelter. The course page says properly fitted respirators are the most reliable PPE wherever airborne transmission dominates, and that they keep healthcare systems and the industrial base running during the window before medical countermeasures exist — including the capacity needed to make and distribute those countermeasures. This report is the full version of that claim: technology, fit, manufacturing base, stockpiles and distribution, with the numbers for each. The course suggests reading the executive summary and introduction and then choosing a chapter; the chapter guide below is for that choice.

In depth

The four phases, and how the study was built

The study is organised as a sequence: characterise the gaps that hampered past responses, parametrically analyse how much protection and how much product a next-generation threat demands, use both to write requirements, then identify the funding and policy moves that get there. Phase 1 asks what

the gaps are, Phase 2 and 3 what the goals are, Phase 4 how the goals are achieved.

The evidence base is deliberately mixed: scientific literature, governmental policies and plans, real-world after-action reports, media and other open-source material for the gap analysis; parametric modelling for protection and demand; and for solutions, interviews with industry, government, healthcare and academia plus two workshops. The first workshop was manufacturers and innovators, on the grounds that industry has to operationalise most of the recommendations, so suggestions had to be compatible with real business models. The second was wider — industry, government, think tanks, philanthropy. Some proposed solutions were dropped or rewritten on that feedback, and the rejected ones are listed in an appendix.

Phase 1: what the gaps are

The gap list is grouped into five areas, and it is worth knowing the shape of it even if you choose a different chapter.

Standards. Regulatory standards for the same item differ by country, there is no standardised nomenclature for PPE, not all vital workers have clear PPE standards at all, and there are no standards for the public, including children.

Product design. Respirators are not designed for facial diversity; PPE is not designed for body diversity or biological needs; some PPE cannot meet religious and cultural needs; PPE is not designed for extreme environments; and fit is hard to obtain, hard to confirm, and hard to keep.

Supply chain and manufacturing. Pandemic control policies disrupt production and distribution; geopolitics and regulatory change disrupt international distribution; offshoring makes countries vulnerable and also hollows out the domestic industry that would otherwise compensate; specialised machinery and facilities limit how fast anyone can scale.

Quality control. Counterfeits and failures, with testing and approval processes too slow for a surge.

Culture, communication and training. Guidance that does not reach people in their own language or at their literacy level, inconsistent messaging, misinformation and polarisation, and incorrect use that contaminates the wearer.

The gaps were then put to PPE stakeholders in government, NGOs and industry, who ranked how critical each one was — so the list is not just the authors' view.

Phase 2: the pathogen, and how much protection it demands

The modelled threat combines the worst traits of pathogens that have actually existed: as infectious and hardy as measles virus, spreading globally as fast as SARS-CoV-2, as deadly as the 1918 influenza virus. The authors anchor this in evolution rather than imagination. Measles sheds more and infects more efficiently than SARS-CoV-2 and stays infectious in air and environment for several hours. The 1918 virus was unusually deadly and, unlike COVID-19, hit young healthy workers hardest. And rinderpest — a deadlier sibling of measles with those properties — already evolved; it went to cattle rather than people by chance.

Three exposure scenarios were modelled: two occupants of an enclosed indoor space with no direct contact, over a full shift; a short visit (under an hour) to a room an infected person has occupied for hours; and close contact with the cloud of particles produced by someone speaking, which applies indoors or out. A variant of the third looked at spray landing on fingertips, eyes, nose or mouth. Five viruses were run; measles and SARS-CoV-2 drove the requirements, because of extreme infectiousness and heavy shedding. Crucially, the model did not use an average infected person but also the people who shed far more than most — the report notes that an estimated 10 to 20% of SARS-CoV-2 infections were responsible for 80% of transmission events.

The conclusions are the part to remember:

- Indoors, over a full shift, a disposable N95 is **not sufficient**. Protection equivalent to a properly fitting elastomeric respirator is needed.
- For a short indoor visit, a surgical mask fails within minutes, while an average-quality, well-fitting N95 protects for up to an hour.
- Outdoors in close contact, a well-fitting N95 is sufficient even for long periods.
- Barrier protection — gloves, eye protection, gowns — is required wherever direct contact cannot be avoided, though the authors say existing biomedical data was insufficient to rank barrier items against each other.

An important technical note from this chapter: the difference between a disposable N95 and an EHMR is **leakage around the seal, not the filter**. Total penetration is not much affected by whether a filter is rated N95, N99 or N100, because most of what gets through comes past the face seal. That is the answer to the course's guiding question about what the 95% in N95 refers to: it is the share of particles the filter material catches at a specific particle size in laboratory testing, not the protection the wearer gets.

On demand, the workforce is sized at approximately **1.4 billion vital workers**, which is 40% of the global labour force and 17% of the world's population — 39

million responders, 858 million indoor accompanied workers, 471 million other vital workers. High-income countries have more of their vital workforce in services; low and lower-middle income countries have more in agriculture.

The demand curve rises fastest in the first 20 days and keeps climbing more slowly to 100 days, which means that without stockpiles, production has to multiply in under a month. The executive summary gives the 20-day figure as more than 1.4 billion units of respiratory protection, 1.1 billion gloves and 43 million gowns per day; Chapter 2's own text gives more than 1 billion respirators, 770 million gloves and 31 million gowns by day 20. Either way, cumulative demand reaches **113 billion units of respiratory protection in the first 100 days**, nearly 75 times the number of N95s produced globally in a year.

Phase 3: what next-generation PPE has to do

Phase 3 turns all that into **target product profiles** in the format the WHO uses, split into design features, material performance and use desirability, with a rationale, a performance target, the evidence and the remaining knowledge gaps for each.

The headline requirement is fit, and it is stated as a design obligation rather than a user responsibility: either the respirator does not need a fit, or it must achieve fit on the worker whatever their body type,

make fit easy to measure, not lose fit over a workday or through normal movement, and make the loss obvious so the wearer knows to leave the hazardous area. The authors say bluntly that they know of no disposable respirator that meets this today.

The quantitative target is **inward penetration of 0.5 to 1 micron particles**: under **2%** for workers who may share a room with an infected person all day, under **6%** for workers who are outdoors or alone or making visits under 60 minutes. The under-2% figure is not met by the vast majority of disposable N95s because of leakage even when fit is obtained, and is met by well-fitting EHMRs and by PAPRs. The under-6% figure today's disposable N95s can meet on filtration — but they do not meet the fit requirements.

The second theme is adapting PPE to the workforce rather than the reverse. Respiratory protection should accommodate facial hair, cultural headwear, assistive devices such as cochlear implants, and a range of head shapes; between 30 and 60% of the global male population can grow a beard, and several religions encourage it. Body-covering PPE should take the female worker as the design standard, which also accommodates men: women are 90% of the nursing workforce and 70% of health workers globally, and most PPE has been designed to fit average American and European men. The report is specific about why this is a protection issue and not a comfort issue —

oversized garments snag and trip, modifications made to fix fit break doffing protocol, and one-piece suits that must be fully removed to use a bathroom led female healthcare workers during COVID-19 to adjust birth control to skip periods, wear adult diapers under PPE, or not work during their periods.

The volume requirements follow from the same modelling. A five-month global supply, which is how long manufacturing took to scale last time, would be **128 billion respirators** at under 2% inward penetration, 5.7 billion gowns, 5.7 billion face shields or goggles and 140 billion gloves — with the explicit note that reusable products would reduce those figures enormously, ideally to roughly the daily figures instead.

Phase 4 and the recommendations

Phase 4's organising finding is the comparison between disposable N95s and EHMRs, which runs through everything else: EHMRs are more cost-effective over time, fit better initially and over a shift, protect better, generate far less waste, and are preferred by wearers in high-risk situations. So pandemic respirator stockpiles should be largely EHMRs.

The recommendations summarised in the executive summary:

- **Early detection as a production trigger.**

Governments should identify outbreaks with pandemic potential as early as possible, tell industry to activate surge capacity at the earliest opportunity, and guarantee to purchase the excess if no pandemic materialises. This shifts the production ramp earlier, so inventory builds before demand arrives.

- **A warm base.** Agreements with manufacturers before a pandemic to keep standby equipment, plus training programmes to create reserve staffing.

- **Stockpiles, split by product.** Respiratory PPE stockpiles centralised and government-managed, holding 150 days' worth, 90% elastomeric respirators and 10% PAPRs. Barrier PPE also 150 days, but split between multiple owners: manufacturers, distributors and users holding a cumulative third, government the remaining two thirds. Fill them over five to ten years, which forces manufacturers to expand sustainable production beyond current consumption and thereby raises their surge capacity.

- **Supply chain robustness.** Comprehensive monitoring systems for a real-time view of the domestic market; domestic production or reliance on close allies; stockpiles of precursor materials; and pre-approval of alternate raw materials so manufacturers can substitute during a shortage.

- **Innovation and markets.** Long-term government contracts that create stable demand, support for manufacturing during emergencies, programmes to build PPE systems in regions that have none, open-source design libraries carrying the latest research, and alignment of PPE policy across government agencies to reduce the burden on employers.

The closing claim of the executive summary is a bet on cost: even if another pandemic as severe as COVID-19 took 100 years to arrive, the investments would cost far less than the harms they prevent. And if one as deadly as 1918 arrived, these measures may be necessary — though, they add carefully, not sufficient — to keep vital workers working and avert societal collapse.

Scope, from the introduction

The introduction fences the study. PPE means equipment worn to prevent or minimise exposure to biological hazards: masks, respirators, gloves, face shields, body covers. Engineering controls and collective protective measures are excluded, as are vaccines and treatments. Three of the four product lifecycle stages are in scope — design, manufacturing, service and reuse — while end-of-life disposal costs are not, and technologies for biodegradable PPE were not considered.

The focus population is the vital workforce only. The authors name what that leaves out, and do not soften it: many millions more people worldwide will need PPE to avoid infection, including those living in poverty who cannot practically isolate, and supplying them is a moral imperative and an additional challenge outside the scope of this study.

They also stress that solutions work best at national and regional level with cooperation between countries, because long multinational supply chains add risk from geopolitics, outbreaks, weather and national priorities such as export bans.

The figures

Two diagrams in the assigned pages carry the argument.

Figure 3 lays the COVID-19 gaps along the supply chain as a row of six boxes: material inputs (material shortages, machinery shortage, intellectual property protection), suppliers (long-term contracts), manufacturers (lockdown delays, lengthy approvals, geographic bottlenecks), distributors (transport disruptions, geopolitical barriers), retailers (affordability, predicting PPE needs), and consumers (inadequate stockpiles, hoarding, disposable PPE, expiration dates). The point of reading it left to right is that there was no single failure to fix — every link had its own.

Figure 4 is the one to keep. The horizontal axis is days, the vertical axis units. A flat baseline of demand and stock runs along the bottom, where production tracks consumption and is only slightly higher because of waste. At the marked pandemic start, daily demand turns almost vertical. The production line stays flat for a long stretch — the market lag — then bends upward much later, because obtaining supplies, buying and installing lines, and hiring staff takes up to five months while demand peaks inside 100 days. The gap between the two curves is labelled as the stockpile requirement. Every recommendation in the report is an attempt to change one of those lines: start the ramp earlier, raise the baseline, add flex capacity, or fill the gap with stock. The same figure reappears twice later, redrawn with early detection and with warm basing in place, and a final version in Chapter 5 shows the gap closed.

The chapters

This is the part to choose from. Five chapters plus conclusions and three appendices.

Chapter 1: Gaps in the PPE Enterprise (printed page 7). The backward-looking chapter: standards, design, supply chain, quality control, and culture, communication and training, each with its evidence. It holds the most quotable failures — the US Strategic National Stockpile held only about 1% of the respirators needed for COVID-19, the UK's pandemic

influenza stockpile about two weeks' worth of NHS needs, US price markups from pre-pandemic to April 2020 of about 15x for N95s and surgical masks, 2x for nitrile gloves and 20x for isolation gowns, over 70% of US respiratory PPE sourced from China, NIOSH approval averaging three months, roughly 60% of international respirators NIOSH tested in 2020 performing below their claimed N95 standard, a single recall removing about 89 million face masks, and fit test pass rates of under 60% for French healthcare workers on FFP2s and 45% and 65% in two Chinese studies. Its recommendation is the ranked gap list itself. Choose this chapter for the politics and the failure mechanisms.

Chapter 2: Parametric Analysis of PPE

Performance and Demand (page 15). The modelling chapter: the worst-case pathogen, the three exposure scenarios, the time-to-infection results by PPE type, the sizing of the global vital workforce, and the demand curve against current production capacity (medical gloves 350–400 billion units a year, N95s 1–5 billion, surgical masks 7–35 billion, PAPRs 2–5 million, gowns 1–8 billion). Its conclusion is that capacity is 10 to 100 times short, with estimated unmet first-year demand of 444 billion respirators, 12 billion face shields and 5 billion gowns. Choose this chapter for the quantitative backbone.

Chapter 3: PPE Requirements (page 24). The target product profiles, respiratory and barrier, with the inward penetration targets (under 2% indoors, under 6% outdoors or alone), the fit requirements, the human factors requirements, and the stock volumes needed. Its recommendation is the TPPs themselves, offered as procurement and design specifications anyone can adopt. Choose this chapter if the useful output is a specification she could hand to a purchasing committee or a union negotiator.

Chapter 4: Recommended Solutions (page 36).

The longest and most practical chapter: the full EHMR-versus-N95 comparison on lifecycle cost, fit, storage and waste; early detection with guaranteed purchase; warm basing and reserve staffing; supply chain resilience and precursor stockpiles; the four stockpile models (vendor-, distributor-, user-managed inventories and centralised government stockpiles) and the recommendation to combine them; market and financing strategies including long-term contracts, volume guarantees, pilot programmes and labour union contract negotiation; regulatory barriers; post-market surveillance, an imperfect-use index and anti-counterfeiting; and future research needs. Its headline numbers: well-fitting EHMRs have 60–97% lower total inward leakage than well-fitting disposable N95s and 78–95% lower than ill-fitting ones; 28% of workers lost fit with a disposable N95 after CPR chest compressions against 0% of EHMRs; a 150-day

stockpile of disposable N95s would need 20+ million pallets and 8.4 square miles of warehouse against 0.32 million pallets and 0.13 square miles for the equivalent in EHMRs and filters; filters generate roughly 200 times less waste; and a global 150-day stockpile of respiratory and barrier PPE represents between \$72 and \$430 billion in product. Choose this chapter for the deep-dive: it is where stockpiles, manufacturing and distribution all live.

Chapter 5: Putting It All Together (page 57). The prioritisation chapter: how the solutions combine, and what to do if you cannot fund everything — prioritise shortening the time between outbreak detection and increased production (early detection, warm basing, long-term domestic contracts), because a marginally larger stockpile matters less than a marginally earlier date at which PPE can be made. Choose this chapter if the question is sequencing.

Conclusions (page 62). Short, and carries the cost-benefit claim: \$16 trillion estimated cost of the pandemic to the US alone, against under \$50 billion for everything proposed for the entire world, and at least a 30-fold return even assuming only a 10% chance of another COVID-scale pandemic in the next twenty years.

Appendix 1 tabulates performance standards for respirators at filter efficiency of 97% and above across jurisdictions (N99, N100, FFP3, KN100, DS3,

KF99). **Appendix 2** lists the solutions they rejected and why — mostly cooling technologies judged non-functional or unlikely to be stockpiled — which is useful as an honesty check. **Appendix 3** maps design solutions onto the TPP requirements.

Words to know

- **Vital workers** — the people who keep society running and cannot work from home: healthcare, emergency response, food and water, waste, energy, transport. The report puts them at about 1.4 billion worldwide, and the whole study is scoped to them.
- **EHMR (elastomeric half-mask respirator)** — a reusable rubber half-mask with replaceable filter cartridges. The report's central recommendation.
- **PAPR (powered air purifying respirator)** — a respirator with a battery-powered blower. The only type that does not require a fit, which is why 10% of the stockpile is reserved for it: for beards, headwear, cochlear implants, scars and the highest-risk jobs.
- **FFR (filtering facepiece respirator)** — the disposable moulded respirator. "Disposable N95" in this report means an FFR with any filter efficiency of 95% or above.
- **Filter efficiency versus inward leakage** — filter efficiency is the share of particles the filter material catches at a specified size in laboratory

testing, which is the 95% in N95. Total inward leakage is everything that actually gets in: through the filter, past the face seal, and through any exhalation valve. Most of it comes past the seal, which is why filter grade matters far less than fit.

- **Inward penetration of 0.5 to 1 micron particles** — the report's own performance yardstick. Under 2% is required for workers sharing indoor space all day; under 6% for outdoor or solitary work. Lower is better; a well-fitted EHMR meets 2%, a typical disposable N95 does not.
- **Target product profile (TPP)** — a structured specification of what a product must and ideally should do, in the format the WHO uses. Chapter 3 is a set of them.
- **Warm basing** — paying to keep manufacturing capacity ready but idle so it can be switched on in an emergency, borrowed from the Defense Logistics Agency's Warstopper programme. Its counterpart is a managed inventory: stock held and rotated by a vendor, distributor or user rather than sitting in a government warehouse, so it turns over through normal sales instead of expiring.

Numbers in this reading

From the executive summary and introduction: - As infectious and hardy as measles, as fast-spreading as SARS-CoV-2, as deadly as the 1918 influenza: the modelled worst-case pathogen. - More than 1.4 billion

units of respiratory protection, 1.1 billion gloves and 43 million gowns per day, as soon as 20 days after emergence (executive summary). Chapter 2 states the day-20 figures as more than 1 billion respirators, 770 million gloves and 31 million gowns — the report gives both. - 10 to 100 times: how far that demand exceeds what current industry can provide. - 150 days: the recommended stockpile duration for both respiratory and barrier PPE. - 90% EHMRs and 10% PAPRs: the recommended composition of the respiratory stockpile. - A cumulative third held by manufacturers, distributors and users, two thirds by government: the recommended split for barrier PPE. - Five to ten years: the period over which stockpiles should be filled. - 100 years: the interval at which, even so, the authors argue the investment still pays. - Up to five months for production to ramp up against a demand peak inside 100 days: the gap Figure 4 describes.

From the chapters, for choosing one: - About 1% of the respirators needed: what the US Strategic National Stockpile held for COVID-19; the UK's stockpiles held roughly two weeks of NHS needs. - About 15x for N95s and surgical masks, 2x for nitrile gloves, 20x for isolation gowns: US price markups from pre-pandemic to April 2020. - Over 70%: the share of US respiratory PPE sourced from China before the pandemic. - Roughly 60%: the share of international respirators NIOSH tested in 2020 that performed below their claimed N95 standard; about

89 million masks removed in a single recall; 9.1 million US gowns recalled in January 2020. - Under 60% of French healthcare workers passed FFP2 fit tests; 45% and 65% pass rates in two Chinese studies; 55% of 6,000 healthcare workers passed on the first disposable N95 selection and 93% by the third; 92% achieved fit with an EHMR on first selection; 50% lost fit within an hour of a successful fit test; 28% lost fit after CPR chest compressions against 0% of EHMRs. - 47% of people wearing PPE for more than four hours get skin reactions, rising to 95% at 12 hours or longer. - Women are 90% of the nursing workforce and 70% of health workers globally; 30 to 60% of the global male population can grow a beard. - 1.4 billion vital workers, 40% of the global labour force and 17% of the world's population: 39 million responders, 858 million indoor accompanied, 471 million other. - 113 billion units of respiratory protection in the first 100 days, nearly 75 times annual global N95 production. - Annual global production capacity: 350–400 billion medical gloves, 1–5 billion N95s, 7–35 billion surgical masks, 2–5 million PAPRs, 1–8 billion gowns, 0.65–0.88 billion face shields. Estimated unmet first-year demand: 444 billion respirators, 12 billion face shields, 5 billion gowns, no shortfall in gloves. - Under 2% and under 6% inward penetration of 0.5 to 1 micron particles: the two protection requirements. - A five-month global stock: 128 billion respirators, 5.7 billion gowns, 5.7 billion face shields or goggles, 140 billion gloves.

Peak daily need for indoor accompanied workers: 850 million respirators, 38 million gowns, 38 million goggles or face shields, nearly 1 billion gloves; plus 422 million respirators a day for outdoor or solitary workers. - Disposable N95s cost \$0.25–\$1.57 each, about 21 per employee per quarter in normal times; one year costs \$21–\$132 per person normally and \$62.50–\$392.50 in a pandemic, against \$30–\$180 for an EHMR with filters either way; a sufficient-quality elastomeric can be bought for under \$90 and disinfected repeatedly for up to a year, which is why a disposable would have to cost under \$0.25 to compete. - 60–97% lower total inward leakage for well-fitting EHMRS than well-fitting disposable N95s, and 78–95% lower than ill-fitting ones. - 20+ million pallets and 8.4 square miles of warehouse for a 150-day disposable N95 stockpile, against 0.32 million pallets and 0.13 square miles for EHMRS and filters; three months of filter waste is 4.1 million cubic feet against 960 million cubic feet for disposables, roughly 200 times less. - A false alarm under early detection, estimated at two months of production, would cost the US government \$180 million to \$2.8 billion; the model assumes a trigger about every five years with 1 in 4 becoming a pandemic. - A 10% management fee on the cost of goods for managed inventories; over 20 years, US respiratory protection costs \$6.57–\$39.45 billion as a managed inventory of disposable N95s or \$12.22–\$56.04 billion centralised, against \$4.38–\$22.79 billion managed or \$2.97–

\$15.25 billion centralised for EHMRs. - \$72 to \$430 billion: the value of a global 150-day stockpile of respiratory and barrier PPE. PAPRs for 10% of the workforce account for 40–50% of respiratory stockpile cost and more than 60% of its 20-year cost. - Three months of normal barrier PPE use equals about 18 days of pandemic consumption, or 12% of responder barrier PPE; nine months spread across vendors, distributors and users equals 54 days, or 36% of the 150-day total, leaving government 64% — given as 96 days in Chapter 4 and as 106 days in Chapter 5. - \$5–7 million to build the diverse 3D body and face scan database needed for better fit; current low-cost PAPRs under \$150 up front with hoods at about \$2 a day. - \$16 trillion estimated US cost of the pandemic against under \$50 billion for everything proposed worldwide, with at least a 30-fold return assuming a 10% chance of another COVID-scale pandemic in twenty years.

What the author is unsure about

The report is unusually explicit about its own limits. The worst-case pathogen is a modelling choice, vetted by stakeholders as realistic and evidence-based, not a prediction. Existing biomedical data was insufficient to compare the protective value of different barrier PPE, so the model can only show that barrier protection is necessary in direct contact, not rank the items. The fit and leakage studies that favour

EHMRs are consistently small-sample, which the authors flag and ask to be repeated with larger cohorts. No facepiece respirator, including an EHMR, achieves universal fit, and about 10% of people seeking protection without a fit test would not be adequately protected — enough, they note, that inability to fit even 5% of workers could undermine confidence in a deadly pandemic. Validated decontamination methods for reusable respirators and barrier PPE do not yet exist. Whether filter covers protect EHMRs adequately in high-splash healthcare settings needs more research. Experimental studies of PPE use in cold environments are lacking. The vital worker count is a rough estimate from World Bank sector data, which may wrongly include or exclude particular professions, and the authors invite national governments to redo it with better data. The demand model assumes emergency rationing at one respirator per worker per day, which they call a low-end, already herculean figure. Cost ranges are wide — a tenfold spread on the stockpile value — because reliable quantitative data on PPE manufacturing was not generally available and capacity had to be triangulated. And the two internal inconsistencies worth noticing yourself: the day-20 demand figures differ between the executive summary and Chapter 2, and the government share of the barrier stockpile is given as both 96 and 106 days.

For the discussion

- The single most useful fact from this report is that the gap between a disposable N95 and an elastomeric is leakage around the face, not the filter — which means every procurement argument conducted in terms of 95 versus 99 percent filter efficiency is being conducted in the wrong units.
- Worth raising as a tension: the report scopes itself to vital workers and then states that supplying everyone else who cannot isolate is a moral imperative outside its scope, which is an honest boundary and also the place where a 150-day stockpile becomes a rationing policy nobody has written yet.

4.4-1 • FURTHER • COURSE TIME 15 MIN

Glycol vapor pathogen disinfection: a defensive technology for emerging biothreats

Sarah Siddiqui and Jaspreet Pannu • 2026

[Read the original paper →](#) • the course assigns: whole review (article body)

A review of what is actually known about glycol vapours as air disinfectants: they work in chambers against most classes of pathogen, the mechanism is becoming clear, and the open question is which real indoor conditions they work in and whether breathing them every day is safe.

Average summary

This is a peer-reviewed review in *Frontiers in Microbiology*, and it is the sober companion to Blueprint Biosecurity's advocacy for the same technology. Siddiqui and Pannu gather the historical and modern evidence on glycol vapours inactivating pathogens in air, and say plainly what remains undefined: the mechanism, the effectiveness across different environments, and the safety.

The evidence, by pathogen class, is uneven. It is **strongest for enveloped respiratory viruses** — the ones that cause flu and COVID-19. Mice exposed to aerosolised influenza survived when propylene glycol vapour was present; modern aerosol chamber studies show dose-dependent inactivation of aerosolised SARS-CoV-2 and influenza A; and in vitro work shows broad activity against seasonal coronaviruses, SARS-CoV-2 variants, MERS and Ebola. It is **weaker for non-enveloped viruses**, which have tougher protein shells: triethylene glycol cut airborne MS2 bacteriophage infectivity by about 2 to 3 logs, and rotavirus needed higher concentrations or temperatures. **Bacteria** are the oldest and most extensively documented application, across gram-positive and gram-negative species. **Fungi** have the thinnest evidence — a few species, one modern office study.

The mechanism is the clearest contribution. Glycol vapour partitions or condenses into pathogen-carrying respiratory droplets, producing high local concentrations inside the droplet that permeabilise membranes. On top of that, glycols appear to disrupt protein structure: for enveloped viruses they may push metastable surface proteins into premature shape changes that leave the virus unable to infect, which fits what is known about the fragility of viral fusion proteins. Non-enveloped viruses have more stable capsids, hence the higher doses.

The environmental findings are the problem. Relative humidity is central and the relationship is not simple. Early work found aerosolised *Serratia marcescens* was inactivated at 25% humidity rather than 80%, because higher humidity pulls more water into the droplet and dilutes the glycol. But a recent study under simulated office conditions found efficacy peaked at 55% humidity, declined as temperature rose from 22.0 to 25.0 °C, rose with higher recirculation — and **fell with higher ventilation rates and higher efficiency filters**. In other words the better your ventilation and filtration, the worse your glycol performs, which makes layering these defences a genuine design question rather than a slogan.

On safety, there is no direct evidence. What exists is occupational data on theatrical fog, at concentrations

far above what disinfection would use. At an average personal inhalable concentration of 0.49 mg/m^3 — within the 10 mg/m^3 occupational limit — theatre workers showed acute nasal and throat irritation, lower lung function closest to the fog source, and chronic wheezing and chest tightness associated with cumulative exposure over two years. EPA chamber studies achieved significant reductions at roughly 1.2 to 1.5 mg/m^3 .

The authors also disclose that generative AI assisted with literature search, drafting and figure ideation, with all references independently verified.

Why it's assigned

The course page notes that studies on glycol vapours are more nascent than those on far-UVC. This reading is how you tell the difference between nascent and absent: there is a real and old evidence base in chambers and animals, and almost nothing on occupied real-world buildings or on long-term human exposure. It is also the only reading in the unit that quantifies how an air disinfection tool interacts with the other layers — and finds that it competes with them rather than simply adding to them.

Words to know

- **Enveloped versus non-enveloped virus** — enveloped viruses (flu, coronaviruses) are wrapped in a fatty membrane studded with proteins, which is

relatively easy to disrupt. Non-enveloped viruses (noroviruses, rotavirus, the MS2 bacteriophage used as a surrogate) have a tough protein shell and need much more of everything.

- **Partitioning** — glycol in the air dissolving into the water of a respiratory droplet, as distinct from condensing onto its surface. Either way the glycol concentration inside the droplet ends up far higher than in the room air, which is why the vapour can be dilute and still work.
- **Relative humidity (RH)** — how much water vapour the air holds as a share of what it could hold. It is the single most important environmental variable here, because it determines how much water is in the droplet and therefore how dilute the glycol becomes.
- **MS2 bacteriophage** — a hardy, harmless virus that infects bacteria, used as a stand-in for SARS-CoV-2 in chamber experiments. Results against it are conservative, since it is non-enveloped and tougher than the real target.
- **8-hour time-weighted average (TWA) limit** — the occupational ceiling on average exposure across a working day, 10 mg/m^3 for glycol mists. It permits brief peaks, which is why peak limits exist separately.
- **Carbonyls (formaldehyde, acetaldehyde)** — breakdown products that form when glycols are heated, documented in e-cigarette studies. The

reason the authors want by-products measured under realistic operating conditions.

- **Recirculation rate** — how much of a building's air is returned and re-delivered rather than exhausted. More recirculation keeps glycol in the room, which is why it helps here and ventilation hurts.

Numbers in this reading

- 2 to 3 log: the reduction in airborne MS2 bacteriophage infectivity reported for TEG vapour in controlled chambers.
- 25% versus 80% relative humidity: the early finding that aerosolised *Serratia marcescens* was inactivated at low humidity and not high, with TEG most effective, then PG, then ethylene glycol.
- 55% relative humidity: where TEG efficacy peaked in the recent simulated-office study.
- 22.0 to 25.0 °C: the temperature range across which efficacy declined, with the optimum at 22.0 or 23.5 °C.
- 0.49 mg/m³: the average personal inhalable aerosol concentration measured during theatrical fog use, with short-term peaks occasionally reaching several tens of mg/m³.
- 10 mg/m³: the ACGIH and WCB 8-hour time-weighted average limit for glycol mists, which the measured theatrical exposures did not exceed.

- 1.2 to 1.5 mg/m³: the average TEG concentrations at which EPA chamber studies achieved significant reductions in aerosolised viral viability — comparable to background near a fog machine and well below documented occupational peaks.
- 2 years: the period of cumulative glycol fog exposure significantly associated with chronic work-related wheezing and chest tightness in theatre workers.

What the author is unsure about

Almost everything that matters for deployment, and the review is structured around saying so. The mechanism involves multiple pathways whose relative contributions remain unresolved, especially for resistant pathogens like non-enveloped viruses and spores. Translating chamber findings into real-world occupied spaces remains challenging, because the historical studies were run in unventilated chambers under controlled conditions that do not represent buildings with HVAC systems. The humidity relationship is more complex than the early inverse pattern suggests, with droplet size and aerosol residence time also involved. Evidence for fungi is limited to a small number of species, and whether it extends to clinically relevant fungi is unclear. Safety data is indirect — theatrical fog and e-cigarettes, not continuous low-level disinfection — and the authors name the real question: not whether glycols are

acutely toxic at low concentrations, but whether long-term routine exposure is safe in occupied buildings that contain asthmatic and immunocompromised people. Degradation products need measuring under realistic conditions. The AI and machine learning section is explicitly conceptual: no studies have applied these methods to glycol systems. And their closing statement of the most pressing gap is worth quoting as a standard for the whole unit: identifying the specific indoor conditions under which the antimicrobial effect is robust enough to produce a meaningful public health benefit.

For the discussion

- The finding that should reshape the conversation: in simulated office conditions, glycol efficacy fell with higher ventilation rates and higher-efficiency filters, so the layers of a biohardening strategy can work against each other and "layered defence" needs engineering, not just enthusiasm.
- The safety question is not the one people expect — nobody is worried about acute toxicity at disinfection concentrations; the unanswered question is what continuous exposure does over years to the people least able to leave the building.
- Set this against the Blueprint piece on the same technology: the two agree on the chemistry and diverge on the verdict, and the difference is entirely that this review weighs the absence of real-world

and long-term data rather than the abundance of chamber data.

4.4-2 • CORE • COURSE TIME 30 MIN

On far-UVC and air filtration

Damon Binder • 2026

[Read the original paper →](#) • the course assigns: whole piece, including the three appendices (version updated 29 April 2026)

Far-UVC and air filtration do the same job in different ways, and on today's prices the box fan with a filter taped to it is the cheaper way to clean a room — so the case for far-UVC is about scale and safety, not about it being better.

Average summary

Damon Binder's argument is a corrective. Discussions of pandemic defence in buildings treat far-UVC as the headline technology and air filtration as an afterthought, and he thinks that is backwards. The two are close substitutes: both deliver **effective air changes per hour (eACH)**, and a given eACH cuts airborne pathogen concentration by the same amount whichever way you got it. On current pricing, filtration is cheaper. Far-UVC will get cheaper as krypton chloride emitters scale, but he doubts it will ever be substantially cheaper than filtration, because filter

media and box fans are mature commodities sitting near a cost floor while excimer emitters are pre-scale with room to fall.

He works it through for one room: 30 m² of floor, a 2.5 m ceiling, 75 m³. Over five years, the cost per eACH comes out at **\$53 for an AirFanta 3Pro, \$63 for a Corsi-Rosenthal box, \$79 for a box fan with a MERV-13 filter, \$99 for a commercial Philips HEPA unit** — against **\$74 for an Aerolamp DevKit, \$295 for a commercial UV Medico fixture, \$729 for a consumer Krypton-11 lamp, and \$261 for conventional upper-room 254 nm germicidal UV.**

The one far-UVC row that beats filtration is a projection: Blueprint Biosecurity's possible future fixture at \$110 and 30,000 hours, which comes to \$18 per eACH, conditional on filter costs falling, lifetimes actually reaching 30,000 hours, and the market scaling.

Where far-UVC wins is scaling. Filters get noisy and draughty, and you hit dead spots and short-circuiting. Getting 25 eACH is not hard — aircraft cabins run about 20 ACH at cruise and operating rooms are mandated at 20 or more — but going much past that gets progressively harder. Far-UVC scales by adding fixtures until you run out of ceiling, and its real limit is safety: long-term human exposure data is limited, there are open questions about secondary chemistry in indoor air, and at current exposure limits the ceiling

is probably somewhere in the tens to low hundreds of eACH.

Part 2 is bleaker and more useful. He models what installing UVC in one workplace or school actually buys. An office worker at 5 eACH avoids roughly **0.1 respiratory infections a year**, worth \$69; a school child avoids roughly **0.5**, worth \$208. That makes a four-person office a benefit-cost ratio of about 3.7 at the cheapest fixture pricing, and a 25-student classroom about 35. Clean air everywhere would cut transmission by about a third, which is not enough to eliminate rhinoviruses or coronaviruses.

His conclusion: neither technology will be deployed widely enough to matter in a catastrophe by default. Filtration already should be everywhere on cost-effectiveness grounds and is not, which he reads as a bad sign for far-UVC. For a near-term emergency ramp-up, filtration wins by a lot — fans and filter media are commodities, and people built Corsi-Rosenthal boxes at home during COVID-19 — while far-UVC is a specialty lighting niche with a global install base of roughly 100,000 fixtures and about 20 small vendors.

Why it's assigned

This is the counterweight on the biohardening page. The course page presents filtration and disinfection as two tools with different niches and asks directly

how high single-pass efficiency should be and what the use cases for each are; this reading answers both questions with prices. It is also the piece that makes the day's honest uncertainty concrete: the same lamp's effectiveness varies by one or two orders of magnitude depending on which pathogen susceptibility number you believe.

In depth

The unit that makes everything comparable

He starts by building the yardstick. A typical room in a house has natural ventilation — air leaking through gaps, doors and windows — which replaces the room's whole volume roughly once an hour, so about **1 air change per hour (ACH)**. Opening a window raises it; running an HVAC system raises it further.

Air changes matter because they remove airborne pathogens. An infectious aerosol lingers about an hour in a 1 ACH room, and anyone there during that hour has a chance of breathing it. Double to 2 ACH and you halve the lingering time and roughly halve the infection risk.

Filters and UVC work through the same channel by different means. A filter pulls room air through media that captures aerosols and returns clean air; one that processes the room's full volume once an hour is worth 1 ACH. UVC inactivates the pathogens inside aerosols without removing the aerosols, and a system

that does that across the room's volume once an hour is also worth 1 ACH — except that the rate depends on the pathogen and on aerosol size, since UVC works better on some pathogens than others and less well on larger droplets.

Added together, ventilation plus filtration plus UVC gives **effective air changes per hour (eACH)**. A room with 1 ACH of background ventilation, a filter delivering 2 and UVC delivering 2 has 5 eACH. Each unit adds a fixed amount; two lamps deliver twice one lamp.

The cost comparison

Cost per eACH depends on room size, so he fixes a reference room — 30 m², 2.5 m ceiling, 75 m³, about a bedroom or small office — and prices eight real products plus one projection. Upfront and five-year cost per eACH:

- **HEPA filtration:** AirFanta 3Pro, 8 eACH, \$20 upfront and \$53 over five years. Corsi-Rosenthal box, 5 eACH, \$20 and \$63. Box fan plus MERV-13 filter, 2.5 eACH, \$16 and \$79. Philips 4000i commercial unit, 8 eACH, \$67 and \$99.
- **Conventional 254 nm germicidal UV**, upper-room, 12 eACH, \$208 and \$261 — the only row that includes professional installation, which it requires.
- **Far-UVC at 222 nm:** Aerolamp DevKit, 7 eACH, \$71 and \$74. UV Medico commercial fixture, 7 eACH,

\$286 and \$295. Krypton-11 consumer lamp, 4 eACH, \$625 and \$729.

- **The projection:** Blueprint Biosecurity's possible future KrCl luminaire, 7 eACH, \$16 and \$18.

The underlying numbers are in his first appendix, with sources for every row: measured clean air delivery rates (111 cubic feet per minute for the box fan from Holder et al. 2022; 400 m³/h for the Corsi-Rosenthal box from independent testing in Dal Porto et al. 2022; 600 m³/h for the AirFanta; 610 m³/h for the Philips), prices, power draw, lifetimes and consumables. Operating cost assumes six hours a day at \$0.15 per kilowatt-hour, which is both the occupied hours and the duty cycle.

Two fairness notes he makes himself. Ventilation is left out of the table because where a building supports it, it is essentially free eACH — but availability is heavily constrained by HVAC capacity, climate, and heating and cooling costs that scale with outdoor airflow. And the comparison depends on room size: UVC's cost per eACH falls relative to filtration in larger or taller spaces, because filters multiply with volume while one fixture can illuminate a bigger column of air.

The Blueprint projection is the pivot of the whole piece: \$110 for the device, 30,000 hours, \$0.19 per milliwatt per year, an order of magnitude below today's commercial pricing. At those numbers far-UVC

lands well below commodity HEPA — conditional on dichroic filter costs falling to about \$50 a unit, emitter lifetimes actually reaching 30,000 hours in practice, and the KrCl market scaling to match the model. His verdict: filter media and box fans are mature commodity products sitting near a cost floor, excimer emitters are pre-scale with meaningful headroom, rough parity seems plausible, and he would be surprised to see far-UVC substantially beat filtration.

Where far-UVC genuinely wins, and where it does not

On scaling, far-UVC has the better story if safety allows. Filters get noisy and generate uncomfortable airflow, and because they physically move air you can run into dead spots or short-circuiting, where one unit's clean output is sucked straight into another's intake without mixing through the room. Reaching 25 eACH is not that hard — he anchors it with aircraft cabins at roughly 20 ACH at cruise (WHO 2008) and operating rooms mandated at 20 or more by ASHRAE Standard 170 — but beyond that it gets progressively harder, and the facilities that hit hundreds of eACH, like semiconductor fabs, are purpose-built with laminar flow.

UVC instead scales by adding fixtures until you run out of space, and the binding constraint becomes safety. He is careful here: 222 nm light does not

penetrate far into skin or eye, but long-term human exposure data is limited for both, and there are open questions about secondary chemistry in indoor air. At current ACGIH threshold limit values he puts the safety-limited ceiling somewhere in the tens to low hundreds of eACH for most respiratory viruses, depending on the pathogen's susceptibility, room geometry and how conservatively the limit is applied. He also notes UVC degrades plastics and other surfaces over time, which is not well characterised.

Everything else, he says, is minor and does not clearly net out. Far-UVC's advantages are form factor and sound: ceiling fixtures are soundless and out of the way, where filters are floor-standing units with audible fans. It also disinfects illuminated surfaces, cutting into fomite and short-range transmission that filtration cannot touch, though he expects both effects to be minor for most respiratory viruses. Filtration's advantages run the other way: filter media captures particulate matter, pollen and some volatile organic compounds as a side effect, which is an independent reason to run one in a city or during wildfire season, and filters carry neither the open human-exposure question nor the material-degradation question.

What a real installation buys you

Part 2 is the part that changes how you argue about this, because it prices the benefit rather than the hardware.

Two factors limit what UVC in a workplace can achieve. First, workplaces are a modest source of infection for adults — most respiratory viruses come from family members, public transit and social contacts, so eliminating workplace transmission entirely would avert only a fraction of a year's infections. Second, UVC acts primarily on long-range aerosol transmission, which carries perhaps half of total respiratory transmission, with the rest going through short-range droplets, direct contact and fomites. So even a perfect installation in every room blocks at most half of within-workplace transmission.

Running those through a two-component epidemic model, at 5 eACH an office worker avoids roughly 0.1 respiratory infections a year, worth \$69, and a school child avoids roughly 0.5, worth \$208. At the cheapest current fixture pricing that makes a four-worker office a benefit-cost ratio near 3.7 and a 25-student classroom closer to 35. The classroom case is robust because children spend most of their waking hours at school and have higher attack rates per strain than adults. He flags gyms, restaurants and public transit as other promising crowded settings, and notes the general rule that the case for air cleaning tracks occupant density times time spent in the space — the same intuition behind ASHRAE 241 setting much higher targets for gyms than offices.

For the whole population, his conclusion is that ubiquitous air cleaning would not eliminate colds and flu. Getting 5 eACH everywhere — homes, workplaces, transit, retail — blocks about a third of transmission, which pushes the effective reproduction number below 1 only for a pathogen whose basic reproduction number started below about 1.43. Seasonal flu, at about 1.3, sits just under that boundary, and he does not believe the model's answer: today's measured strains are the ones that win under today's countermeasures, universal UVC is a different selection environment, and new strains may emerge with higher transmissibility or lower susceptibility. He keeps a stated 50% guess for flu throughout rather than accepting the model's "eliminated." Rhinoviruses at about 2 and coronaviruses at about 2.5 are simply too transmissible for elimination, and he notes this ceiling is not specific to UVC: any intervention that cleans air mostly addresses long-range aerosols, and some transmission is just people being close to each other.

The totals: roughly \$640 per adult per year in averted productivity and healthcare costs against roughly \$300 per adult to cover their indoor space at the cheapest pricing, a benefit-cost ratio of 2.1 — which drops to 0.53 at commercial fixture pricing. Scaled to 260 million US adults, that is a recurring programme of about \$80 billion a year, on the order of 0.3% of GDP, against roughly \$170 billion a year in seasonal-

disease benefit. And a warning on his own numbers: that is hardware cost only, and at population scale installation and maintenance labour would plausibly dominate the fixture cost.

Why none of this happens by default

His own interest, he says, is catastrophic pandemic risk rather than seasonal disease. He would pay a lot to have fewer colds, but the seasonal case is marginal at current prices and only modest even as costs fall — not a return that drives a ubiquitous infrastructure rollout.

The most uncomfortable argument in the piece is the one from filtration's own history. Filtration is mature, cheap, and has existed for decades; by any reasonable standard of cost-effectiveness it should already be ubiquitous in schools, workplaces and public buildings. That it is not, despite favourable economics, is a bearish sign for far-UVC. The two have different frictions — filters are loud and bulky, which UVC avoids; UVC has safety and material-degradation questions, which filters avoid — so it is not clear UVC wins on net.

Both could be deployed with a deliberate push: building codes, a consumer-grade quality standard, or a major funder-backed procurement programme could put effective filters in most public-facing buildings, with UVC following as costs fall. Safety

studies and clearer regulation would specifically de-risk UVC. The total bill would be small relative to what is already spent on HVAC.

Absent pre-deployment, catastrophe relevance comes down to emergency ramp-up, and here the gap is wide. Filter media and fans are commodity items, and COVID-19 showed consumers can build effective Corsi-Rosenthal boxes at home from off-the-shelf parts — he calls not pushing that during COVID-19 a massive missed opportunity. UVC is a specialty lighting niche: a global install base of roughly 100,000 fixtures built on roughly 20 small vendors and specialised supply chains — KrCl excimer emitters, narrowband optical filters, driver electronics — that cannot be converted from adjacent industries the way filter media can. Standing up tens of millions of fixtures in a year is not a capacity you can buy your way into; the factories and the trained workforce do not exist and building them takes years.

He still ends constructive. For a near-term ramp-up in building defence, filtration is the bet by a lot and covers most of what UVC would. But the UVC field is small enough that a few million dollars of targeted de-risking — safety studies, regulatory pathways, anchored deployments — is highly leveraged, and further work may yet show UVC performs better than current numbers suggest. The destination is a world

where indoor air cleaning is ubiquitous by filtration, UVC, or more likely a mix.

The appendix that undercuts every headline number

Appendix 1's second half is the most important passage in the reading for judging the rest of the unit. The eACH figures for UVC are not pathogen-neutral. UV inactivation is log-linear in dose, so the eACH a fixed installation delivers scales with the pathogen's susceptibility, written **k** (log inactivation per unit of UV dose at 222 nm). Each UVC row in his cost table inherits whichever pathogen its source used.

Holding one installation fixed and rescaling by **k**, the spread is enormous: coronaviruses have a **k** of 4.1 on the bench and 12 to 16 measured in aerosol, implying 4 or 12 to 16 eACH; influenza A 1.8 bench and 20 aerosol; a murine norovirus surrogate 2.4 bench; adenovirus 0.14 bench; *Bacillus subtilis* spores 0.36 bench, implying 0.35 eACH from the same hardware. Bench and room-scale aerosol measurements disagree, with aerosol **k** running 3 to 10 times above bench **k**, most of which is probably an artefact of how dose is delivered rather than a real difference. And inter-lab **k** values for the same pathogen in nominally similar setups commonly vary by one to two orders of magnitude, so any single headline number carries wide implicit uncertainty. He uses the conservative bench anchors, and notes that Jeff Kaufman's post

reaches about \$53 per eACH for the same Aerolamp by using the aerosol coronavirus numbers instead — the difference between the two posts is this choice, not a disagreement about the hardware.

Droplet size matters too. At 222 nm, light is absorbed within roughly 1 micron of dried respiratory fluid, so sub-micron particles generated by breathing — where most long-range aerosol mass sits — are efficiently inactivated, while cough- and sneeze-sized droplets need roughly 4 to 8 times the dose for the same effect. Experimental data across the droplet size range is sparse. Filtration rows do not shift at all under any of this, because filters capture whole aerosols regardless of what is inside them or how big they are.

Words to know

- **ACH (air changes per hour)** — how many times an hour a room's whole volume of air is replaced. A leaky house is about 1.
- **eACH (effective air changes per hour)** — the same scale extended to filters and lamps, which clean air in place. It is what makes a \$20 box fan and a \$2,500 lamp comparable, and it adds up across technologies.
- **CADR (clean air delivery rate)** — the volume of clean air a device delivers per unit time, in cubic

feet per minute or m^3/h . Divide by room volume to get eACH.

- **MERV-13** — a mid-grade filter rating. His footnote argues consumer HEPA units are over-specified, because what matters is capture rate times airflow: 99.97% efficiency at $500 \text{ m}^3/\text{h}$ is much the same as 90% at $550 \text{ m}^3/\text{h}$, and MERV-13 delivers it at lower pressure drop, less noise and lower cost.
- **Corsi-Rosenthal box** — the DIY air cleaner made from a box fan and four or five filters. Tested at $400 \text{ m}^3/\text{h}$, about \$100 in parts.
- **KrCl excimer emitter and its dichroic filter** — the krypton chloride lamp that produces 222 nm far-UVC, plus the narrowband optical filter that passes 222 nm and blocks the longer, more harmful wavelengths. The filter is the single biggest cost component, and has to fall to about \$50 a unit for the cheap future projection to hold.
- **Susceptibility (k)** — how readily a given pathogen is inactivated per unit of UV dose, in cm^2 per mJ. Everything about far-UVC's value scales with it, and published values for the same pathogen vary by one to two orders of magnitude.
- **Benefit-cost ratio (BCR)** — value delivered divided by money spent. Above 1 means worth doing on the modelled assumptions.

Numbers in this reading

- 1 ACH: a typical room's natural ventilation, which leaves an infectious aerosol lingering about an hour; doubling to 2 ACH roughly halves infection risk.
- 75 m³ (30 m² floor, 2.5 m ceiling): the reference room for every cost figure.
- Five-year cost per eACH: \$53 AirFanta 3Pro, \$63 Corsi-Rosenthal box, \$79 box fan plus MERV-13, \$99 Philips 4000i, \$261 upper-room 254 nm GUV, \$74 Aerolamp DevKit, \$295 UV Medico, \$729 Krypton-11, \$18 for Blueprint's projected future fixture.
- Upfront cost per eACH: \$16 to \$67 for filtration, \$71 to \$625 for current far-UVC fixtures, \$208 for upper-room GUV, \$16 for the projection.
- \$110 device cost, 30,000 hours, \$0.19 per mW per year: Blueprint's possible future KrCl luminaire, an order of magnitude below today's commercial pricing, conditional on dichroic filters falling to about \$50.
- Six hours a day at \$0.15 per kWh: the operating assumption behind every five-year figure.
- 400 m³/h for the Corsi-Rosenthal box, 600 for the AirFanta, 610 for the Philips, 111 cubic feet per minute for the box fan plus filter: the measured clean air delivery rates.
- 25 eACH is achievable with filtration; aircraft cabins run roughly 20 ACH at cruise; ASHRAE 170 mandates 20 or more for operating rooms;

semiconductor fabs reach hundreds with purpose-built laminar flow.

- Tens to low hundreds of eACH: the safety-limited ceiling for far-UVC at current ACGIH limits.
- About half: the share of respiratory transmission carried by long-range aerosols, and therefore the ceiling on what air cleaning can block.
- 0.1 infections a year averted, worth \$69, for an office worker at 5 eACH; 0.5 a year, worth \$208, for a school child.
- Benefit-cost ratio about 3.7 for a four-worker office and about 35 for a 25-student classroom at the cheapest fixture pricing; 0.9 and 9 respectively at commercial pricing.
- α of about 0.3: the transmission cut from 5 eACH everywhere, which only eliminates a pathogen whose basic reproduction number was below about 1.43. Seasonal flu is about 1.3, rhinoviruses about 2, coronaviruses about 2.5.
- 43%, 28%, 33%, 29% and a 50% guess: the per-strain relative risk reductions for rhinovirus, endemic COVID, seasonal coronaviruses, RSV and flu.
- \$510 per cold, \$1,800 per flu, \$1,400 per endemic COVID case, \$1,000 per adult RSV case; an hour of lost work valued at \$48.78.
- \$640 in value averted per adult per year against about \$300 per adult in hardware, a ratio of 2.1 —

or \$1,200 per adult at commercial pricing, a ratio of 0.53.

- About 300 m³: the indoor space a typical adult occupies, 180 at home and 120 across workplace, transit and retail — four reference rooms.
- \$80 billion a year for 260 million US adults, about 0.3% of GDP, against roughly \$170 billion a year in benefit — hardware only.
- Roughly 100,000 far-UVC fixtures installed globally, supplied by roughly 20 small vendors.
- Susceptibility k at 222 nm: 4.1 bench and 12 to 16 aerosol for coronaviruses, 1.8 and 20 for influenza A, 2.4 bench for a murine norovirus surrogate, 0.14 for adenovirus, 0.36 for *B. subtilis* spores — implying anything from 0.35 to 20 eACH from the same hardware.
- 3 to 10 times: the gap between aerosol and bench susceptibility measurements; one to two orders of magnitude: the variation between labs for the same pathogen.
- Roughly 1 micron: the depth 222 nm light penetrates dried respiratory fluid, which is why cough-sized droplets need 4 to 8 times the dose of breath-sized particles.
- A few million dollars: what he thinks targeted de-risking of the far-UVC field would cost.

What the author is unsure about

He labels his own uncertainty throughout, which is what makes the piece usable. The future far-UVC price is a projection with three named conditions. The safety-limited eACH ceiling is "probably" in the tens to low hundreds, and where exactly depends on pathogen susceptibility, room geometry and how conservatively limits are applied. Long-term human exposure data for 222 nm on both skin and eyes is limited, and secondary indoor air chemistry is an open question. Material degradation from UVC is "not very well characterised." The share of transmission carried by long-range aerosols is "poorly quantified but plausibly around half," and he assumes 0.5 for simplicity while saying it varies by disease and setting. The flu elimination result is one his model produces and he rejects, substituting a stated guess. His rhinovirus parameters involve a modelling choice worth a factor of 1.5. The per-case dollar values carry wide defensible ranges, especially for long COVID. He says plainly that at population scale installation and maintenance labour would plausibly dominate hardware cost, which his totals exclude. And on the central question he is explicit: it is not clear that UVC wins on net against filtration once you count safety and material questions, and further work may show it performs better than current numbers suggest.

For the discussion

- The reframe worth bringing is that far-UVC and filtration are substitutes measured in the same unit, so the question is never "does far-UVC work" but "what does a unit of clean air cost here" — and in a 30 m² room today, a box fan with filters wins.
- The strongest argument in the piece is historical rather than technical: filtration is cheap, mature and should already be everywhere, and it is not, so expecting far-UVC to diffuse on its merits is optimistic.
- The hard question for a pandemic, as opposed to a cold season: 100,000 far-UVC fixtures worldwide and 20 vendors means there is no surge capacity, whereas anyone can build a Corsi-Rosenthal box from hardware store parts — which makes filtration the only one of the two that is actually a pandemic tool today.

4.4-3 • CORE • COURSE TIME 60 MIN

Blueprint for Far-UVC

Blueprint Biosecurity · 2025

[Read the original paper →](#) · the course assigns: executive summary and all twelve recommendations in full (printed pages 2 to 11), plus the opening and summary of every chapter for the chapter guide

Far-UVC could clean indoor air far better than anything we have, and the report's whole argument is that it should not be installed

everywhere until four specific research questions are answered, because the public would be exposed without being able to consent.

Average summary

This is the long version of the far-UVC case, built from the literature plus two years of consultation with over 100 experts across photobiology, atmospheric chemistry, indoor air quality, building science, environmental engineering, epidemiology and public health. It opens with the scale of the problem it is aimed at: COVID-19 has claimed an estimated **27 million lives**, tuberculosis kills **1.6 million** people a year, and **one billion** people are infected by influenza every year, producing millions of serious illnesses and hundreds of thousands of deaths.

The technology in one paragraph. Germicidal ultraviolet light has been used in water treatment for over 100 years, and studies in the 1940s showed that treating the air above people's heads — upper-room UV — could control measles in schools; it has been used for decades against drug-resistant tuberculosis. **Far-UVC** is a newer form. Because it is strongly absorbed by proteins in the outer layer of human skin and eyes, it can inactivate a wide range of pathogens while barely penetrating human tissue, which allows higher human exposure limits and therefore continuous disinfection of rooms with people in them.

It is also soundless, energy-efficient, commercially viable at scale, less prone to resistance than drugs, and — the key property for pandemics — can be installed before an outbreak.

The report is organised around **twelve recommendations**, and it deliberately puts them before the evidence. Four are urgent research priorities: establish how installations must be designed to actually suppress transmission; identify the biological effects specific to far-UVC as distinct from sunlight and conventional germicidal UV; understand the unintended air quality effects, principally ozone, and how to mitigate them; and obtain high-quality real-world evidence of effectiveness, which they judge means cluster-randomised trials. Five concern adoption: simple deployment guidance that public health agencies can communicate, better consensus standards, installation alongside ventilation and filtration rather than instead of them, cheaper and better emitters, and cost-benefit frameworks. Three form a long-term agenda: long-term safety studies in diverse populations, the effects of far-UVC on everyday materials, and the mechanism by which it inactivates pathogens.

The moral argument underneath is unusual and worth noticing. Far-UVC works without anyone changing their behaviour, which is exactly why it could matter —

and it means that in a world where it is widespread in public spaces, only limited personal consent can be given to the exposure. The report says the thoroughness of its research agenda reflects the trust that has to be earned before that is acceptable.

On urgency: climate change, factory farming and human encroachment into wild habitats more than double the expected risk of spillover, while AI and synthetic biology increase exposure to engineered pandemics, and experts estimate a **1 in 8 chance of a pandemic killing over 90 million people by 2050** — three times COVID-19's toll. The funding ask is public and philanthropic money on the order of **\$100 million over five years**.

Why it's assigned

The course page says early lab studies indicate far-UVC is effective and safe for humans but real-world data is needed, and that it cannot be scaled rapidly in a pandemic, so it depends on peacetime installation to control ordinary colds and flu. This report is where both halves of that sentence come from. It is also the deep-dive text for the biohardening option, and the course suggests reading the executive summary and recommendations and then picking a chapter — the chapter guide below is for that.

In depth

The problem, and what far-UVC would add

The executive summary frames airborne disease as the unfinished business of public health. We have made enormous progress against water-borne, food-borne and vector-borne disease; airborne infection remains one of the biggest challenges, with 27 million estimated COVID-19 deaths, 1.6 million annual tuberculosis deaths, and a billion influenza infections a year. The authors note that the interventions working in the background of daily life — sanitation, pest control, food hygiene — get less credit than vaccines and drugs, even though they stop us meeting pathogens at all. Far-UVC belongs to that unglamorous category.

The physics in one sentence: strong absorption by proteins in the outermost layer of skin and eyes means far-UVC inactivates pathogens with minimal penetration into human tissue, which permits higher exposure limits, which permits continuous disinfection of occupied spaces at a level of air cleaning significantly beyond current alternatives.

The chapter on why this is worth accelerating sets out what the alternatives cannot do. Medical countermeasures take time — even at record speed, COVID-19 vaccines were not available for a year, longer in many countries, and some people including the immunocompromised, infants and the elderly may not be fully protected even then. Respiratory protection is essential and individually effective, but

faces cost, comfort, communication difficulty and compliance fatigue, and is particularly hard for young children, people with some medical conditions or disabilities, and settings where clear communication matters. Ventilation and filtration are proven and the authors are enthusiastic about them where practical, but many HVAC systems cannot meet current air cleaning standards, especially in crowded, ageing or energy-constrained buildings, and portable cleaners run at the levels needed for substantial protection generate noise and draughts that cause people to turn them down or off. The concrete example: ASHRAE Standard 241 recommends clean air delivery rates that can far exceed both CDC guidelines and typical building capacity — a restaurant would need an additional 60 cubic feet per minute per person of clean airflow beyond standard ventilation, comparable to a modern operating theatre. Conventional upper-room 254 nm UVC can meet those targets and is recommended by CDC/NIOSH and WHO for tuberculosis wards and operating theatres, but its intensity in the upper room must be far above safe human exposure limits, so it needs expert installation and occupants who know about the hazard overhead.

Against that, the promise list: far-UVC inactivates a wide range of pathogens, with strong results against influenza and coronaviruses, bacteria including *Staphylococcus aureus* and *Pseudomonas aeruginosa*, and even pathogenic fungi such as *Candida auris*;

installations can be made safe by design; modelling suggests it can be **up to 450 times more energy-efficient than ventilation and 40 percent more efficient than air purifiers**; it is soundless, compact and simple to install; it shows promise against surface and short-range transmission as well; and it may help against antimicrobial resistance, having inactivated drug-resistant bacteria in lab studies.

Three early cost-benefit results are offered, carefully labelled as preliminary: an analysis of implementing ASHRAE 241 air cleaning targets estimating a 10-to-1 return on investment counting only seasonal illness; a Swiss analysis of far-UVC in indoor public spaces estimating a benefit-cost ratio of 30 to 290 times in a normal winter respiratory season and higher in pandemics; and an analysis of conventional UVC in aircraft cabins finding a 1,000 percent annual return and a cost of \$10,000 per life saved.

The urgency case rests on both the hazard trend — climate change, factory farming and habitat encroachment more than doubling spillover risk, AI and synthetic biology widening exposure to engineered pandemics, and an expert estimate of a 1 in 8 chance of a pandemic killing over 90 million people by 2050 — and on the speed of airborne spread: Omicron doubled cases roughly every 2 to 3 days and infected an estimated 125 million people globally within 10 weeks of being identified, and

measles, the most infectious human pathogen, is airborne too.

Scope, which matters for reading everything else

The report evaluates far-UVC for suppressing **long-range** airborne transmission only — the well-mixed case where pathogens disperse through a room's air and are inhaled minutes or longer after being exhaled. Surface disinfection, prevention and treatment of surgical-site infections, near-field protection such as PPE with far-UVC built in, and short-range transmission are all named as promising and left out, with a short-range review promised in a future update. Other air cleaning technologies are out of scope except where they relate to far-UVC. The primary purpose is coordination: pointing researchers, policymakers, entrepreneurs and funders at the same priorities.

Why the recommendations come first, and the consent argument

The report opens with its recommendations rather than its evidence, and explains why. If far-UVC is to work as pandemic prevention, it has to be widely installed across indoor public spaces before outbreaks happen. That has a consequence the authors follow through: mass installation before an outbreak can only be justified if far-UVC proves highly effective against ordinary endemic respiratory illness,

because otherwise nobody has a reason to install it in peacetime.

Then the part that distinguishes this report from the advocacy around it. Far-UVC does not require most people to change their behaviour, which is a central reason it could be so impactful — and a world where it is widespread in public spaces is one where only limited personal consent can be given to the exposure. The flip side of passive collective protection is that great trust is placed in the scientists who develop it, the professionals who install it, and above all the institutions that set safety and efficacy standards. The thoroughness of the research agenda, they say, reflects the trust that has to be earned through rigorous science. They also list what they did not evaluate and which might matter as much: basic research into the aerobiology and epidemiology of transmission, biosensors that detect infectious aerosols in real time, and better monitoring of infections acquired in high-risk settings such as hospitals.

The four research priorities

1. Establish how installations must be designed to suppress transmission. This is ranked highest because everything else gets harder as the required dose rises. Studies so far show order-of-magnitude differences in how susceptible pathogens are to far-UVC, and the causes of that variability are not

understood and may not be specific to UV at all — standardising research methods for infectious aerosols could benefit every air-cleaning technology. The same protein absorption that makes far-UVC safe will also reduce its penetration into the protein-rich human respiratory aerosols where pathogens actually sit, and nobody knows how big that effect is or what else it does to pathogen stability; representative human respiratory aerosols are also hard to make in a lab. Modelling — computational fluid dynamics plus lighting simulation — is needed to turn experiments into deployment guidance, and because end users cannot run those models, the output has to be practical rules of thumb that public health agencies can endorse. Environmental factors, relative humidity in particular, are known to affect susceptibility to conventional germicidal UV and are likely to matter here. The four sub-tasks: get inactivation data from actual human respiratory aerosols; run controlled bioaerosol chamber studies on environmental variables; find the causes of experimental variability and standardise methods; model dose distribution across room geometries, ventilation regimes, occupancy and source location.

2. Identify the biological effects specific to far-UVC. Protein absorption is the foundation of the safety claim, and its consequences are not fully understood. There is no current evidence of significant harm, but exposure limits are built on

which wavelengths induce which harmful effects, so the list of measurable endpoints has to be complete. Sub-tasks: get a mechanistic understanding of protein absorption in skin and eye; study other pathways by which far-UVC's higher photon energies may differ from longer wavelengths; use that to identify biomarkers and build action spectra that can inform the limits set by expert bodies such as ICNIRP and ACGIH.

3. Understand the air quality impacts. Far-UVC generates some ozone. What we know about ozone harm comes from epidemiology of outdoor ozone, and using it indoors requires assuming that the complex indoor chemistry seen with far-UVC is all downstream of ozone generation — testing that assumption is a high priority. Ozone and its by-products can be mitigated with ventilation, catalysts or activated carbon filters, and reducing indoor ozone may have health benefits independent of far-UVC. Sub-tasks: in-depth field measurements of indoor chemistry in representative spaces, comparing far-UVC against introducing an equivalent quantity of ozone; replication of lab ozone-generation measurements across devices to build a predictive model and a standard test method; epidemiological work quantifying the health effects of indoor ozone and its reaction by-products; and investigation of catalysts and activated carbon for removal.

4. Obtain high-quality evidence of real-world effectiveness. Cluster-randomised trials are the gold standard for infection control, and the report is unusually candid about their risks. Previous upper-room UV trials produced mixed results because of design flaws that are hard to avoid; it is difficult to power a trial well enough to detect a real reduction in transmission, and an underpowered study risks setting the whole field back, because failure to find statistical significance gets misread as evidence that the intervention does not work. Their judgement is that without a shift in what public health agencies accept as evidence, successful cluster-randomised trials are necessary to catalyse adoption — while noting that ventilation became widely accepted without one, and that if they are wrong about this, the other recommendations still stand. Sub-tasks: power trials adequately; use doses likely to be effective while keeping occupants within photobiological safety limits; and collect temperature, humidity, CO₂, ventilation rates, pathogen concentrations and sequencing alongside the infection endpoints, so the result can be interpreted either way.

The five adoption recommendations

5. Simple deployment guidance that trusted institutions can communicate. Public health, risk and radiation safety communication are mature fields with established principles. Guidance must be

implementable by practitioners and comprehensible to a lay audience, must account for different ventilation levels rather than being one-size-fits-all — the impact of ozone-initiated chemistry is especially sensitive to ventilation — and the institutions the public trusts should be brought into the process early.

6. Better consensus standards. Some exist, more are needed, all will need revision. 2023 brought UL 8802 for UV germicidal equipment and ASHRAE 241, the first standard attempting a technology-neutral framework for preventing long-range airborne transmission in public spaces. Others were written without germicidal lamps in mind and need amending — UL 2998 on zero ozone emissions from air cleaners — and IEC 62471 on photobiological safety of lamps has not been updated since 2006. They list six standards to refine, ask for new standards on ozone generation from germicidal applications and on manufacturing and labelling (consistent consumer information on expected lifetime, UV output, emissions spectrum and exposure limits), ask that sector standards such as the Facilities Guidelines Institute's hospital standards be modified to allow germicidal UV in high-risk spaces, and ask for better test methods for air cleaners claiming to remove infectious aerosols, which agencies like the US EPA could use to police marketing claims. There is an explicit invitation: if you are reading this, consider participating in standards development.

7. Install according to standards, as one layer among several. Layered interventions with different mechanisms are more resilient to diverse threats, and far-UVC will be better against some pathogens than others. Practical advice for anyone installing today: buy devices certified to appropriate product standards such as IEC/EN 62471 or ANSI/CAN/UL 8802; when assessing exposure limits, count all the device's emissions and not just the 222 nm peak; make sure ventilation is adequate and working first, because far-UVC is not an alternative to minimum ventilation standards; typically avoid wall mounting and prefer ceiling fixtures facing down for extra eye-safety margin, while recognising this depends on room geometry and occupant behaviour; and weigh the benefits and risks carefully before use in private homes. Risks are higher where occupants dwell longer, ventilation is poorer, and there is no professional facilities management.

8. Improve performance and affordability of emitters. Current emitters are expensive and may not be cost-effective outside high-risk settings such as healthcare. The right metric is cost per milliwatt of far-UVC output that can be safely installed in a space. Scale alone could bring substantial reductions in a very small industry; there is also room in component costs and in wholly new emitter technologies. Industry asks: better diffuser technology, which would cut cost per milliwatt and raise efficiency for low ceilings;

proximity sensors or cameras with dimming and boosting, for dynamic output based on occupancy; cheaper filters that pass far-UVC and block everything else, since filters are a significant share of lamp cost; longer lamp lifetime; next-generation semiconductor emitters; and markets for far-UVC outside indoor air disinfection — water treatment, surface disinfection, healthcare tools, agriculture, food processing, pest control, scientific equipment — which they list as proposals they have not evaluated.

9. Cost-benefit frameworks. Far-UVC is not meant for everywhere, and adoption should be evidence-guided. Institutional decision-makers will weigh upfront capital, operating costs and deployment risk against fewer infections, less absenteeism, more productivity and lower healthcare costs. Where infection control is required and far-UVC substitutes for other methods, the benefit can show up as cost or energy savings. They want frameworks tailored to specific settings, starting with healthcare facilities, public gathering spaces and schools; regularly updated; and adaptable to other air-cleaning technologies so comparisons are fair.

The long-term agenda

10. Long-term safety studies in diverse populations. Post-approval studies matter for pharmaceuticals and the same principle applies here. They say explicitly that requiring long-term studies of

hypothetical side effects before deploying an effective life-saving innovation would be an excessive application of the precautionary principle — and that this does not make the studies unnecessary. Because powering an effectiveness trial constrains which facilities and therefore which people can be studied, long-term safety in populations with different sensitivities will likely have to be studied in different settings. Start as soon as practical, which requires long-term commitment from participating buildings and occupants.

11. Study effects on materials ubiquitous in the built environment. Far-UVC will interact with common materials, and the priority is deciding which to study, weighing mechanism, importance, ubiquity and intended lifespan. Controlled lab exposures on common materials, cosmetics, clothing and plants, looking at appearance, performance and off-gassing of harmful compounds; long-term real-world studies; and mitigation including resistant coatings and guidance for sensitive materials. They note buildings have been re-engineered before — sanitation, fire safety, energy use — and renovations are an opportunity to fix several problems at once.

12. Understand the mechanism of inactivation. Knowing how far-UVC inactivates microbes would let us predict efficacy across threats without testing every pathogen, and it connects to a wider goal: real-

time feedback on whether an intervention is working. Today you cannot reliably tell infectious from inactivated pathogens without slow bioassays, so better mechanistic understanding could give targets for practical biosensors.

The chapters

Ten chapters and two appendices. The course suggests picking one.

1. Germicidal UV: introduction and history (printed page 12). The history from Ritter's discovery of ultraviolet in 1801 through Downes and Blunt showing sunlight inhibits bacterial growth in 1877 and Koch's 1890 demonstration on *Mycobacterium tuberculosis*, up to the present, placing far-UVC in the context of how COVID-19 changed the understanding of airborne transmission. Choose it for the long view of why air disinfection was abandoned and rediscovered.

2. Far-UVC primer (page 20). A glossary chapter, meant to be dipped into: the electromagnetic spectrum, dose and fluence, action spectra, and exposure limits. The key content is that ACGIH and ICNIRP both recommend a spectrally weighted UV exposure limit of 3 mJ/cm^2 over 8 hours in a 24-hour period, where each wavelength is multiplied by a weighting that reflects its biological impact — which is why an unfiltered lamp's small emission at 259 nm matters so much, since 259 nm counts more than 50

times as heavily as 222 nm. Choose it if the units are the obstacle.

3. Efficacy (page 30). The chapter the rest depends on. It explains why clinical trials of place-based interventions are hard to power, introduces the engineering approach (eACH and clean air delivery rate), and lays out the evidence. The famous result is Eadie et al. 2022: airborne bacteria cut by 98.4% in a room-sized chamber, equivalent to 184 eACH. The chapter also prints the rest of the table, which is the honest part — the same five-device setup gives 184, 133 and 122 eACH against *S. aureus* and *P. aeruginosa*, but only 1.24 to 2.06 eACH against MS2 bacteriophage and 1.1 to 1.79 against T1 bacteriophage. And one real-world study in an occupied mouse cage cleaning room found a 99.8% reduction in aerosolised murine norovirus, which implies anywhere from 200 to 14,800 eACH depending on assumptions about mixing and base decay. Its conclusion: if real-world efficacy looks like Eadie, far-UVC is extremely promising; if it looks like the bacteriophage studies, it holds far less promise. Choose this chapter for the central uncertainty of the whole technology.

4. Skin and eye safety (page 55). Overwhelming evidence that sunburn, corneal inflammation and the usual DNA damage markers are all significantly reduced at far-UVC wavelengths, no evidence of

significant harm within recommended limits, and a clear statement that our understanding of far-UVC's other biological effects is incomplete. Notable details: much of the influential recent evidence is from animal models and human studies are needed; protein-rich wound fluid is thought to protect underlying tissue in open wounds; and acute eye overexposure to far-UVC appears to cause immediate discomfort that subsides, rather than the delayed, prolonged corneal inflammation conventional germicidal UV causes. Choose it for the safety argument in detail.

5. Evaluating cancer risk from different UV

wavelengths (page 67). Mouse studies have not shown increased cancer risk from far-UVC, consistent with how little it penetrates the epidermis — but UV as a class is classified as a carcinogen by the US National Toxicology Program and the International Agency for Research on Cancer, and far-UVC currently sits inside that classification, which does not grade magnitude of risk. Its recommendation is more animal work, mechanistic study of which DNA damage types far-UVC causes, and epidemiological studies if deployment becomes widespread. Choose it if the question is how a regulator or a sceptical public will treat the carcinogen label.

6. Ozone and indoor air quality (page 74). Far-UVC generates ozone through photolysis, and the generation rate per unit of fluence is relatively well

quantified: 7 to 10.5 parts per billion per hour per $\mu\text{W}/\text{cm}^2$ for the filtered krypton chloride lamps assessed so far. What occupants are actually exposed to is highly sensitive to outdoor air ventilation rate and to what other chemicals are present. At 1 air change per hour the expected steady-state increase is low to mid single-digit ppb per $\mu\text{W}/\text{cm}^2$ of average fluence — small against outdoor ozone, but relatively large against the central tendency of average indoor ozone at about 4 to 6 ppb. Its recommendation is mitigation through ventilation, catalysts or activated carbon, and it notes that the products of ozone reacting with other substances may matter more for health than ozone itself. Choose it for the strongest objection to far-UVC.

7. Ozone epidemiology (page 90). The health side of the same question. The most cited evidence is Turner et al. 2016, which estimated a hazard ratio for all-cause mortality of 1.02 for every 10 ppb increase in average outdoor ozone across nearly 669,046 participants from 1982 to 2004; Bell et al. 2004 found a 0.52% increase in daily mortality following a 10 ppb rise in the previous week's ozone across 95 US urban communities; and a more recent study estimated a hazard ratio of 1.064 per 10 ppb for under-five mortality in low- and middle-income countries, though only above 50 ppb. Its conclusion: the magnitude of harm from indoor ozone generation is uncertain but potentially large enough to belong

explicitly in cost-benefit calculations, and not so large that it could not be outweighed if far-UVC is highly effective. Choose it to see how a small environmental harm gets weighed against an infection benefit.

8. Emitters and luminaires (page 100). The hardware and the economics. Krypton chloride excimer lamps are used almost exclusively today and will likely lead for 5 to 10 years; current luminaires cost on the order of thousands of dollars per thousand square feet protected, which is prohibitive for many uses. The chapter's claim is that no breakthrough is needed: better diffusers, dynamic output control and longer lifetime (manufacturers claim up to 10,000 hours; industry experts think 30,000 achievable) could transform cost per milliwatt. Its possible-future costing is the number everyone else quotes — \$110 device plus \$200 installation, 110 mW, 15 years, \$0.19 per mW per year, which works out to about \$191 a year for a typical classroom and, at the fluence rates that may prove effective, potentially under a few dollars a year per student. Semiconductor emitters, LEDs and frequency-doubled blue lasers, could offer different form factors and lower cost on a 5-to-10-year horizon. Choose this chapter if the question is whether this ever gets cheap.

9. Guidance, standards, and regulations (page 116). The governance map: who sets guidance, standards and regulations on UV safety, ozone, air

cleaning efficacy, devices and sector practice. Very few US regulations pertain directly to far-UVC, while the EU has incorporated consensus standards on UV safety into product regulation. Useful insight: standards often apply exposure limits more conservatively than the expert guidance they reference, by managing peak irradiance rather than occupants' time-weighted average exposure. Its recommendation is that effective standards and guidance be created in advance of governments and the public wanting legal regulation. Choose it for the policy lane.

10. UVC impacts on materials (page 133). The least-known area. Most UV degradation research covers the UVA and UVB in sunlight, because the atmosphere filters out UVC, so far-UVC's effects on materials are largely unknown and degradation mechanisms are complex and material-specific. Points worth having: wood and paper are vulnerable through the breakdown of lignin; commercial polymers contain additives that complicate everything; metals reflect 222 nm most strongly, which can produce higher fluence rates and exposures than expected unless coated or carefully placed; and most studies only measure colour change when they should also test off-gassing of harmful compounds. Choose it if the practical question is what happens to a real room over ten years.

Appendix A acknowledges the more than 100 contributing experts. **Appendix B** covers impacts on non-human life.

Words to know

- **Germicidal UV (GUV) and far-UVC** — GUV is the family of ultraviolet used deliberately to inactivate microbes: in-duct, upper-room 254 nm, and far-UVC. Upper-room UV is the older, proven technology whose intensity must exceed safe human exposure limits, which is why it stays above your head. Far-UVC is the roughly 200 to 230 nm band, 222 nm in practice, absorbed by protein in the outer layer of skin and eye.
- **KrCl* (krypton chloride excimer) lamp** — the dominant 222 nm source, and likely to stay dominant for 5 to 10 years.
- **Fluence rate** — how much UV energy is arriving per unit area per unit time, in microwatts per square centimetre. Both the ozone generation rate and the disinfection rate scale with it.
- **Spectrally weighted exposure limit and action spectrum** — the action spectrum is the curve of relative biological effect across wavelengths; the exposure limit is the ceiling after each wavelength is multiplied by its weighting. The limit is 3 mJ/cm² over 8 hours, and because 259 nm counts over 50 times more than 222 nm, a lamp's stray emissions can dominate its weighted dose. Recommendation

2 is essentially a request to complete the action spectrum for far-UVC.

- **eACH and CADR** — effective air changes per hour and clean air delivery rate: the engineering metrics that let a lamp be compared with a fan or a window.
- **Susceptibility** — how easily a pathogen is inactivated per unit of UV dose. Order-of-magnitude disagreement here is the report's number one research priority.
- **Photolysis and ozone loss** — photolysis is light breaking a molecule apart: far-UVC photons split oxygen molecules, and each split yields two ozone molecules. Ozone loss is that ozone then reacting with other substances indoors, and the products of those reactions may matter more for health than the ozone itself.
- **Cluster-randomised trial (CRT)** — a trial randomising places rather than people, the only design that can test an intervention protecting a room. Hard to power, and an underpowered one can damage the field. Its policy counterpart is the **consensus standard**: a technical document developed through a formal multi-party process over years, not law, but often adopted into law and used as a de facto purchasing requirement.

Numbers in this reading

- 27 million estimated COVID-19 deaths; 1.6 million annual tuberculosis deaths; one billion influenza infections a year.
- Over 100 years: how long germicidal UV has been used in water treatment; the 1940s for the first upper-room UV work on measles in schools.
- Over 100 experts consulted across two years.
- 1 in 8: the expert-estimated chance of a pandemic killing over 90 million people by 2050, three times COVID-19's toll.
- Every 2 to 3 days: Omicron's case doubling time; 125 million people infected globally within 10 weeks of its identification.
- 60 cubic feet per minute per person: the extra clean airflow a restaurant would need beyond standard ventilation to meet ASHRAE 241, comparable to a modern operating theatre.
- Up to 450 times more energy-efficient than ventilation and 40 percent more efficient than air purifiers, in modelling.
- 10-to-1 return on investment for implementing ASHRAE 241 targets counting only seasonal illness; a benefit-cost ratio of 30 to 290 times for far-UVC in Swiss indoor public spaces in a normal winter; 1,000 percent annual return and \$10,000 per life saved for conventional UVC in aircraft cabins.
- \$100 million over five years: the public and philanthropic funding the authors say is needed to

reach the evidence standard public health agencies will expect.

- 3 mJ/cm² over 8 hours in a 24-hour period: the spectrally weighted UV exposure limit recommended by both ACGIH and ICNIRP; 259 nm counts over 50 times more heavily than 222 nm, so an unfiltered lamp hits the limit 3 times faster than a filtered one at the same irradiance.
- 98.4% reduction and 184 eACH: Eadie et al. 2022. The same chapter's other chamber results: 133 and 122 eACH against *S. aureus* and *P. aeruginosa*, and 1.1 to 2.06 eACH against two bacteriophages.
- 99.8% reduction of aerosolised murine norovirus in an occupied mouse cage cleaning room with a nominal 36 air changes per hour, implying 200 to 14,800 eACH depending on mixing assumptions.
- 7 to 10.5 ppb per hour per $\mu\text{W}/\text{cm}^2$: ozone generation efficiency of the filtered KrCl lamps assessed. Steady-state increase at 1 ACH: low to mid single-digit ppb per $\mu\text{W}/\text{cm}^2$, against an average indoor ozone central tendency of about 4 to 6 ppb.
- Hazard ratio 1.02 per 10 ppb of average outdoor ozone across 669,046 participants, 1982 to 2004 (Turner et al. 2016); 0.52% increase in daily mortality per 10 ppb in the previous week across 95 US urban communities (Bell et al. 2004); hazard ratio 1.064 per 10 ppb for under-five

mortality in low- and middle-income countries, above 50 ppb.

- Thousands of dollars per thousand square feet: current upfront cost of effective far-UVC coverage. Manufacturers claim up to 10,000 lamp hours; 30,000 is thought achievable.
- \$110 device plus \$200 installation, 110 mW output, 15-year lifetime, \$0.19 per mW per year, \$0.04 per mW per year in energy, \$191 a year for a typical classroom at 840 mW: the possible-future costing in Chapter 8.
- 2006: the last update to IEC 62471 on photobiological safety of lamps. 2023: the publication of UL 8802 and ASHRAE 241.
- 5 to 10 years: how long KrCl is expected to lead, and the horizon on which semiconductor emitters might replace it.

What the author is unsure about

This report is largely a catalogue of its own uncertainty, which is the point of it. On efficacy: experimental pathogen susceptibility varies by order of magnitude, and if the lower estimates are right, far-UVC compares less favourably to ventilation and filtration and practical deployment gets harder. Most efficacy research has used pathogens aerosolised in media importantly different from human respiratory aerosols, and they expect protein content to attenuate efficacy by an unknown amount. Assumptions about

air mixing and ventilation affect how lab results translate to rooms. On safety: there is no evidence of significant harm within recommended limits, and our understanding of far-UVC's biological effects is incomplete; much of the influential recent evidence is from animal models; the effects of protein absorption in skin, eye and tear film remain to be worked out; and they expect expert bodies to update exposure guidance in the coming years. On cancer: far-UVC sits inside a carcinogen classification for UV as a class, mouse findings need replication in other models, and second-order mechanisms have barely been considered. On air quality: the potential health impacts of increased indoor ozone are highly uncertain, there may be photochemical effects they have not characterised, and whether a threshold from outdoor epidemiology can be applied indoors at all is an open question. On trials: cluster-randomised trials are risky and an underpowered one could set the field back, and their judgement that such trials are necessary may itself be wrong. On materials: largely unknown. On cost: the attractive numbers are a projection contingent on investment that will only happen if a market appears. And the whole report is explicit that it reflects Blueprint Biosecurity's views, not a consensus of the experts consulted.

For the discussion

- The most interesting argument in this report is not about light at all: because far-UVC protects people without asking anything of them, nobody can meaningfully consent to it, which is why the authors load the burden onto standards bodies and demand trials before mass installation.
- The honest version of the efficacy case is the full table, not the headline: the same installation that delivered 184 effective air changes against one bacterium delivered close to one against two bacteriophages, and nobody yet knows which number the pathogens we fear resemble.
- Worth raising as the practical catch: far-UVC only prevents pandemics if it is already installed, and it only gets installed in peacetime if it clearly reduces ordinary colds and flu — so the pandemic case depends entirely on winning an argument about seasonal illness.

4.4-4 · CORE · COURSE TIME 10 MIN

Glycol vapors for pathogen suppression

Blueprint Biosecurity · 2025

[Read the original paper →](#) · the course assigns: whole page (web archive capture)

The chemical in theatrical fog kills airborne pathogens, was studied and then forgotten in the 1940s, costs pennies a day per room, and

is the only air disinfection tool that could plausibly be surged into millions of buildings in the middle of a pandemic.

Average summary

This is Blueprint Biosecurity's short case for a technology nobody is working on. Their framing is a gap in the toolkit. Air disinfection belongs in the portfolio against biological threats, but every existing option has a weakness: building-scale ventilation is expensive to install and expensive to run because outdoor air has to be heated or cooled; filtration and portable air cleaners are affordable and accessible but need very high airflow rates to deliver real protection, and the noise and air currents can disrupt a workplace; germicidal UV is promising but currently expensive and will take time to reach economies of scale. What is missing is something effective, cheap, barely noticeable, and deployable quickly if a pandemic starts soon.

Their candidate is the vapour of three glycols — **propylene glycol (PG), dipropylene glycol (DPG) and triethylene glycol (TEG)**. These are abundant industrial organic compounds, registered with the US EPA for air sanitisation and deodorisation, and already present in ordinary life: salad dressings, perfume, toothpaste, natural gas refining, e-cigarettes, and decades of theatrical fog and non-hazardous smoke

simulants. They are extremely **hygroscopic**, meaning they pull moisture out of the air.

The efficacy claims are strong. Vaporised into a room, glycols have been shown to rapidly inactivate viruses, bacteria and fungal spores, including SARS-CoV-2, influenza A and *Streptococcus*, achieving **1,000-fold or greater reduction within an hour** — three or more log reductions per hour — against many aerosolised pathogens, depending on the indoor environment. The field evidence is historical: a study across three winters from 1941 to 1944 in a healthcare facility for bedridden children, where direct contact between patients was limited so infections were most likely airborne, found a **90% reduction in infections** in glycol-treated wards compared with controls, including a 96% reduction in colds and 90% reductions in tracheobronchitis, otitis media and acute pharyngitis.

Mechanism, as far as it is understood: because the vapours are hygroscopic they dissolve quickly into any moisture in a particle, and the condensed glycol then appears to dehydrate the pathogen, disrupt its membranes and structural proteins, and disable the machinery it uses to enter cells. This works on surfaces as well as in air, and significant inactivation happens at concentrations low enough to be invisible, odourless and tasteless.

On safety, the evidence to date suggests inhaling these vapours at air disinfection concentrations is safe for healthy adults, with more research possibly needed for sensitive populations. The EPA reviewed in vitro, animal and human exposure data for PG and DPG in 2006 and for TEG in 2003 and found no significant evidence of negative health concerns.

On cost, TEG is likely the most economical, because PG and DPG have lower boiling points and therefore need higher concentrations to saturate the air enough to work. The estimate is roughly **10 to 50 cents a day to protect a 1,000 square foot room**, including the liquid and the energy.

Why it's assigned

The course page says glycol vapour studies are more nascent than far-UVC, notes that propylene glycol vapour is used in theatrical fog and e-cigarettes but also has antimicrobial properties, and makes the key observation: it is unlikely people would want theatre fog in their offices permanently, but it could be rapidly manufactured and deployed in a pandemic. This reading is the source for that argument. In the structure of the day it is the counterpoint to far-UVC — far-UVC has to be installed in peacetime and cannot be surged, while glycols could be surged and probably would not be tolerated in peacetime.

In depth

The gap in the portfolio

The piece opens with the premise that suppressing indoor transmission requires a portfolio of interventions layered together, and that air cleaning and disinfection compensate for the weaknesses of the others. The example given is PPE: necessary, but often burdensome, eventually removed for eating and drinking, and requiring advance stockpiling to avoid shortages.

Then each existing air technology and its weakness, stated briefly and fairly:

- Building-scale ventilation can be expensive, both in installation and in the energy needed to condition outdoor air.
- Filtration and portable air cleaners can be affordable and accessible, but need very high airflow rates to achieve significant protection, and depending on the setting the noise and air currents could significantly disrupt a work environment.
- Germicidal UV is promising, but currently expensive and likely to take time to achieve economies of scale.

The gap, then, is a tool that is efficacious, low-cost and minimally disruptive, and that could be deployed quickly if another pandemic arrives soon. That last criterion is doing most of the work: it is a surge-capacity argument, not a best-technology argument.

What the glycols are

PG, DPG and TEG are abundant organic compounds and extremely hygroscopic — they absorb moisture from the air, which turns out to be the basis of how they work. They are registered for use in air sanitisation and deodorisation by the US Environmental Protection Agency, and are already widely used in food, drug, cosmetic and industrial applications: the piece lists salad dressings, perfumes, toothpaste and natural gas refining. In aerosol and vapour form they have been used for decades in theatrical effects and as non-hazardous smoke simulants.

What the evidence shows

Two tiers of evidence, laboratory and field.

In the laboratory, vaporising glycols into a room has been shown to rapidly inactivate viruses, bacteria and fungal spores, including common pathogens such as SARS-CoV-2, influenza A and *Streptococcus*, with an appendix listing the full set. The quantitative claim: 1,000-fold or greater reduction of pathogens within an hour — three or more log reductions per hour — against many aerosolised pathogens, depending on the indoor environment. That last clause is the caveat the next reading on this page explores in detail.

In the field, the evidence is 80 years old and the only field evidence there is. Across three winters from 1941

to 1944, in a healthcare facility for bedridden children where direct contact between patients was limited so that any infections were most likely airborne, glycol-treated wards saw a 90% reduction in infections compared with controls: 96% fewer colds, and 90% reductions in tracheobronchitis, otitis media and acute pharyngitis. The authors are explicit about why that setting matters — limited contact between patients is what lets you attribute the reduction to the air rather than to hygiene.

How it works, as far as anyone knows

The mechanism section is carefully hedged. Because the vapours are hygroscopic, they quickly dissolve into any moisture in airborne particles. The mechanisms of action are not yet entirely understood, but it appears the condensed glycols then dehydrate pathogens, biophysically disrupt membranes and structural proteins, and inactivate the mechanisms by which pathogens enter cells. These effects have been demonstrated against both airborne and surface-borne pathogens. And crucially for deployment, significant inactivation of airborne pathogens can occur at concentrations low enough to be largely invisible, odourless and tasteless — which is the difference between a disinfection system and a fog machine.

Safety

The evidence to date suggests that inhaling these vapours at the concentrations typically used for air disinfection is safe for healthy adults, although more research may be needed for sensitive populations.

The regulatory basis: after reviewing in vitro studies, animal studies and human exposure reports, the EPA found no significant evidence for negative health concerns — for PG and DPG in 2006, for TEG in 2003.

The known adverse reports are handled directly rather than omitted. There are case reports of health effects in theatre performers exposed to fog and haze effects — throat, skin and eye irritation, laryngitis, allergies — but these involve mixtures of glycols and other compounds at much higher concentrations than would be used for air disinfection, and occur in only a small fraction of people.

Cost

TEG is likely the most economical glycol to deploy, because PG and DPG have lower boiling points and therefore require higher concentrations to saturate the air enough to be efficacious. The headline estimate: roughly 10 to 50 cents per day to protect a 1,000 square foot room, including both the TEG liquid and the energy cost. Set against the far-UVC numbers elsewhere in this unit, that is the entire argument for taking glycols seriously.

Why this was forgotten, which is the most interesting part

US and UK researchers investigated glycol vapours for air disinfection in the 1940s and 1950s, alongside germicidal UV. The account of why both were dropped has four strands, and none of them is that the technology failed:

- A consensus grew — incorrectly — that airborne transmission was not a significant mode of infection.
- Antibiotics and vaccines appeared to be panaceas for infectious disease.
- Post-war funding dwindled, and much of the remaining air disinfection research in the US was classified as part of biological weapons programmes.
- Pharmaceutical approaches shifted the focus from the public to the individual, moving the field from public health professionals such as sanitary engineers into the domain of medical doctors.

That last point is the one worth carrying into the discussion: a whole class of collective, environmental defences was not disproved, it was organisationally displaced.

There was a resurgence during COVID-19. Bleu Garde developed a TEG-based product series for continuous use and achieved emergency approval in six US

states for use in occupied spaces, while Reckitt Benckiser released a DPG-based spray-can product for temporary disinfection of unoccupied spaces.

But the public health agencies involved in pandemic response — the CDC in the US, SAGE in the UK — largely declined to promote glycol vapours for air disinfection. The authors do not treat that as a failure of nerve. Since the safety and efficacy profile had not been fully studied at that point, public health experts were understandably sceptical about adding substances to indoor environments when other tools like ventilation, filtration and PPE had not yet been fully employed.

What they want next

The conclusion is a research agenda rather than a deployment plan: glycol vapours appear promising for quickly and affordably surging air disinfection capacity during future airborne pandemics, and warrant further investigation to understand and realise their potential. The specific studies recommended are safety of degradation products in controlled buildings; real-world efficacy across a range of environments; toxicology; industrial surge production; layerability with filtration and germicidal UV; and scalable dispersion methods.

Words to know

- **Propylene glycol (PG), dipropylene glycol (DPG), triethylene glycol (TEG)** — the three candidate compounds. TEG is the most economical for air disinfection because it does not need as high a concentration to work.
- **Hygroscopic** — able to absorb and hold moisture from the air. This is both why these compounds are used industrially and, apparently, how they reach the pathogen: they dissolve into the water in an airborne droplet.
- **Log reduction** — a tenfold cut. Three or more log reductions per hour means at least a thousandfold drop in viable pathogens in an hour.
- **Air disinfection** — inactivating pathogens in the air where they are, as distinct from removing them (filtration) or replacing the air (ventilation).
- **Smoke simulant** — a harmless substance used to stand in for smoke, for example in fire training. One of the long-standing industrial uses of these glycols, and part of the safety record.
- **Surge capacity** — the ability to produce and deploy something fast, after an emergency has already begun. The central claim for glycols, and the thing far-UVC lacks.

Numbers in this reading

- 1,000-fold or greater reduction within an hour, equal to three or more log reductions per hour,

against many aerosolised pathogens, depending on the indoor environment.

- Three winters, 1941 to 1944: the duration of the field study in a facility for bedridden children.
- 90% reduction in infections in glycol-treated wards compared with controls; 96% reduction in colds; 90% reductions in tracheobronchitis, otitis media and acute pharyngitis.
- 10 to 50 cents per day to protect a 1,000 square foot room, including liquid and energy.
- 2006 and 2003: the EPA review dates for PG and DPG, and for TEG, both finding no significant evidence of negative health concerns.
- Six US states: where a TEG-based continuous-use product received emergency approval for occupied spaces during COVID-19.
- 1940s and 1950s: when this was first investigated, alongside germicidal UV, before both were dropped.

What the author is unsure about

The piece is advocacy for research, so its uncertainty is built into the ask. The mechanisms of action are not yet entirely understood. Efficacy depends on the indoor environment, which is stated but not quantified here. Safety evidence supports use for healthy adults, and the authors say more research may be needed for sensitive populations. The 1,000-fold figure comes with a conditional clause; the 90% field result comes

from one study in one unusual setting 80 years ago. And the recommended research list is itself the admission: degradation products in real buildings, real-world efficacy across environments, toxicology, whether production could be surged industrially, whether it layers safely with filtration and germicidal UV, and how to disperse it at scale are all open. They also note that public health agencies were reasonable to decline to promote this during COVID-19, which is an unusual concession for a piece arguing the technology was neglected.

For the discussion

- The strongest argument here is not efficacy but surge capacity: glycols are abundant industrial chemicals that could be redirected into buildings during a pandemic, which is exactly what far-UVC cannot do.
- The history is the part worth raising: this was not disproved, it was abandoned when the field decided airborne transmission did not matter much, antibiotics and vaccines looked like the answer, and air disinfection moved from sanitary engineers to doctors and then into classified weapons work.
- The honest objection is in the reading itself — the best field evidence is a single 1941 to 1944 study in one children's ward, and the agencies that declined to recommend this during COVID-19 were

right that a full safety and efficacy picture did not exist.

4.5-1 • CORE • COURSE TIME 60 MIN

Executive Summary to the Royal Society report "COVID-19: examining the effectiveness of non-pharmaceutical interventions"

Mark Walport, on behalf of the Royal Society's Expert Working Group • 2023

[Read the original paper →](#) • the course assigns: whole executive summary including all three boxes (Europe PMC full text, article body)

Stringent packages of non-pharmaceutical measures did reduce COVID-19 transmission, individual measures varied in how well they worked, all of them worked better when case numbers were low — and the reason we cannot be more precise than that is that nobody designed a way to find out.

Average summary

This is the executive summary of the Royal Society's report on non-pharmaceutical interventions, written by Mark Walport on behalf of an expert working group, and published as the introduction to a themed issue carrying six commissioned evidence reviews.

Non-pharmaceutical interventions (NPIs) are defined as any measure taken during an outbreak to reduce transmission that is not a vaccine or a drug: behavioural, social, physical or regulatory, encouraged through anything from advice to law. They are the first line of defence before biological interventions exist, and were used alongside vaccines where those failed to stop transmission. The report covers six categories: masks and face coverings; social distancing and lockdowns; test, trace and isolate; travel restrictions and border controls; environmental controls (ventilation, air treatment, screens, occupancy limits, surface cleaning and handwashing); and communications.

Two framings matter before the findings. First, the report deliberately makes no judgement about the timing or manner in which measures were applied, and no assessment of economic or other societal impacts — that is left to the various COVID-19 inquiries. Second, for most of the first year, NPIs were the only measures available to slow a novel virus.

The methodological heart of the piece is an explanation of why the evidence is weaker than anyone would like. Controlled trials played a relatively small role for three reasons: urgent action took precedence over designing trials nobody had pre-written protocols for; measures were applied nationally and in combinations, which made isolating

any single one nearly impossible; and while off-the-shelf protocols existed for drugs and vaccines, equivalents for complex social interventions did not. So the evidence is overwhelmingly observational, and when the working group applied methodology originally built for clinical trials, it flagged the inevitable limitations — a point the summary pauses on, explaining that "bias" here means features of study design that reduce reliability, not partiality on the part of researchers.

What was learnt, in brief: **social distancing measures were the most effective of all the NPIs**, with more stringent versions having greater effects. Masks generally reduced transmission, respirators more than surgical masks, mandates more than voluntary use. Test, trace and isolate reduced deaths where implemented at high levels, with strong evidence for contact tracing apps. At borders, symptomatic and exposure-based screening including temperature checks had no meaningful effect, while **quarantine on entry was the most effective border measure**. Environmental controls did reduce transmission within particular settings, but could not be quantified individually or at population level, and there was insufficient evidence on surface cleaning and barriers.

Three case studies — Hong Kong, New Zealand and South Korea — show it was possible to hold

transmission near zero for over a year with early, stringent border controls plus strict domestic measures, and that all three abandoned that strategy once Delta and Omicron arrived.

The conclusion that travels furthest: NPIs were more effective when case numbers and transmission intensity were lower, because the risk of exposure is proportional to the number of cases in the community.

Why it's assigned

This is the evidence base for the third deep-dive option, social interventions, and it is the course's answer to its own guiding question about why there is so much argument over whether these measures work. The answer is not that the evidence is contradictory but that it is structurally weak: measures were deployed in bundles, at national scale, without comparison groups. For the withstand layer as a whole, this is the reading that explains what buys time and how well, and it ends on the point the whole week turns on — NPIs buy the time in which vaccines and treatments can be developed.

In depth

What NPIs are, and the six categories

The definition is deliberately broad: any measure implemented during an infectious disease outbreak to

reduce transmission that is not a vaccine or a drug. They can be behavioural, social, physical or regulatory, and their uptake can be encouraged along an escalating scale from advice and guidance through to regulation. They are the first line of defence before biological interventions are available, and have also been used alongside vaccines and drugs where those fail to prevent transmission.

The six categories, with the mechanism each relies on:

- **Masks and face coverings** act as barriers to virus particles being inhaled or exhaled. Materials and features determine which particle sizes are filtered, and how well the mask fits the wearer's face is also key. N95 masks, also known as respirators, are highly effective barriers when worn correctly.
- **Social distancing and lockdowns** increase physical distance so less infectious material reaches others, with a commonly recommended minimum of two metres — though the summary notes aerosols typically travel further than droplets. At population level this included closures of schools, workplaces, places of worship and entertainment venues, plus stay-at-home orders.
- **Test, trace and isolate** breaks transmission chains by finding infectious people and their contacts and keeping both away from others until the risk has passed.

- **Travel restrictions and border controls** aim to stop the pathogen moving between countries, ranging from complete or partial bans targeted at higher-risk regions to testing and quarantine requirements at departure or arrival.
- **Environmental controls** include enhanced ventilation to replace air carrying infectious aerosols, filtering or treating indoor air, screens, reduced occupancy limits, surface cleaning and enhanced handwashing.
- **Communications**, because none of the above works unless people understand the reason for it and adopt and maintain it correctly.

The scientific principle underneath is that transmission can be prevented if the transmission pathway is blocked. Early clinical studies suggested the primary routes were direct inhalation or exposure of the mucosal surfaces of nose and mouth to virus in airborne droplets — or, as was realised some months into the pandemic, in aerosols. Early evidence of contaminated surfaces pointed to hand-to-face contact as a possible additional route. This understanding was informed by prior knowledge of influenza, RSV and the coronavirus behind the 2003 SARS outbreak.

Why there is so little good evidence

This section is the most useful part of the reading, and it is about method rather than findings.

There are two ways to generate evidence. Controlled trials randomise closely matched groups to interventions differing in strictly defined ways, which lets you attribute outcomes with high confidence; their weakness is small, tightly selected populations whose results do not extrapolate straightforwardly to large heterogeneous ones. Cluster-randomised designs randomise populations rather than individuals. Observational studies compare outcomes against similar data that is historical (same population before and after), geographical (a population with the intervention against one without), or modelled (actual outcomes against a projection). Their advantage is real-world scale; their weakness is confounding.

For drugs and vaccines, trials happened. The RECOVERY trial enrolled more than 47,000 patients to test whether existing anti-inflammatory and antiviral treatments could be repurposed, and the new vaccines were found highly effective against severe disease and death.

For NPIs, trials mostly did not happen, for three reasons the summary spells out:

1. Urgency. In early 2020 the virus was spreading fast, respiratory spread was already the likely

dominant route, NPIs were the only available measures, and such measures are known to work best when case numbers are still low. Formal prior evaluation was not a dominant consideration for policymakers.

2. Scale and bundling. NPIs were implemented nationally and in combinations, on the reasonable grounds that they complement each other — masks plus handwashing plus distancing plus ventilation — augmented by local or large-scale lockdowns as cases rose. That made trials of individual measures almost impossible to implement.

3. No protocols. Rigorous protocols for drug and vaccine trials were available off the shelf. Equivalents for complex interventions with strong social and behavioural components are harder to design, have historically been run far less often, and would have needed to measure both the intended effect on transmission and the undesirable effects across a wide range of personal and societal variables — from the mental and physical health consequences of isolation to lost education, jobs and businesses.

The consequence: very few studies of adequate scale compared different types of NPI, or compared the presence and absence of mask wearing, or measured different levels of distancing. There were, however, very large numbers of observational studies, and the

summary is clear that a great deal can be learnt from well-conducted observational work at scale.

When the working group assessed those studies using methodology originally developed for clinical trials, three limitations came up repeatedly: interventions were almost invariably implemented in combination, so single effects could not be separated; many studies used routinely collected datasets never designed for later evaluation; and comparison groups were often absent and rarely well matched. The summary takes a paragraph to explain that calling these "biases" does not mean researchers were partial — it means features of the study design could reduce the reliability of the conclusions, in either direction.

What was learnt, measure by measure

Masks. The weight of evidence from all studies suggests that wearing masks, particularly higher quality masks, supported by mandates, generally reduced transmission. Studies reported this consistently though not universally. There is evidence, mainly from healthcare settings, that respirators such as N95s were more effective than surgical-type masks; that masks with greater barrier function outperformed those with less; and that mask wearing under a mandate was more effective than voluntary mask wearing.

Social distancing. The most effective of all the NPIs. Stay-at-home orders, physical distancing and restrictions on gathering size were repeatedly associated with significant reductions in transmission, with more stringent measures having greater effects. Measures aimed specifically at protecting the elderly — restrictions on visitors, and cohorting staff with residents so each group of residents is cared for by a specific group of staff — were frequently associated with reduced transmission and fewer care home outbreaks. School closures and other school-based measures were associated with reduced incidence both within schools and in the community, but effectiveness varied compared with community-wide measures, was time-dependent, and was often contingent on adherence and on the age group of the children.

Test, trace and isolate. Used as a key intervention in many countries, especially those pursuing zero-COVID policies. Studies from several countries implementing high levels of contact tracing with isolation of cases and contacts found reductions in deaths. Strong evidence was also found for contact tracing apps, with the trial of the UK app on the Isle of Wight — alongside communications and manual tracing — associated with a substantial reduction in transmission.

Borders. Observational evidence from national case studies including New Zealand showed comprehensive border control policies could reduce but not eliminate the number of infected travellers or their contacts entering a country. Despite nearly universal adoption of travel restrictions, few studies examined these measures implemented alone. On the available evidence: symptomatic or exposure-based screening, including temperature screening before travel, had no meaningful effect on reducing importation or transmission; targeted restrictions such as early bans on entry from specific countries probably had a moderate effect but quickly became less effective as case numbers rose elsewhere; and quarantine at entry borders had the highest levels of effectiveness.

Environmental controls. Enhanced ventilation, air treatment to remove infectious virus, and reduced room occupancy did reduce transmission within particular settings. But these were typically applied in combination with other NPIs, so quantifying them individually and accurately was not possible; many were retrospective rather than prospectively planned observational studies, so confounders could not be fully controlled; and effectiveness was only judged within the setting where the control was applied, not at population level. There was insufficient evidence to judge enhanced surface cleaning or the use of

barriers — gaps the summary says laboratory studies could help fill.

Communications. Considered only in the UK, because political, social and cultural differences make findings about communication extremely hard to extrapolate between countries. The limited evidence confirmed communication was sufficiently effective to ensure high adherence, while also identifying what led to weaker adherence. Trust and confidence in those communicating mattered, as did clarity and consistency of messaging and the opportunity for personal control. Social media communications appear less likely to be associated with higher adherence than traditional media.

Three countries that held the line

The second approach taken was to examine observational data from three places that identified domestic transmission early in 2020 and then held cases at very low numbers for roughly the first 18 months: Hong Kong, New Zealand and South Korea. All three maintained stringent packages of NPIs until the second half of 2021, when large Delta and Omicron waves arrived — causing little harm to the fully vaccinated — and switched to living with the virus.

Hong Kong had suffered badly in the 2003 SARS outbreak, with almost a quarter of the 8,098 cases

worldwide and 302 deaths, which prompted significant public investment in health infrastructure and testing capacity. During COVID-19, people testing positive isolated for 21 days and their contacts for 14, with similarly strict border quarantine. The telling detail is that only an estimated 27% of all cases that occurred in Hong Kong were confirmed by laboratory test, which means containment cannot be attributed to those policies alone. Other measures included minimum distancing, curfews on restaurant opening times, bans on large events, working from home and school closures, with mask wearing mandated in all public settings and high compliance. About 60% of the population was vaccinated by the end of 2021, with lower uptake among older adults — and when Omicron arrived and spread rapidly, more than 10,000 deaths occurred, largely among vulnerable elderly unvaccinated people.

New Zealand is atypical, being a geographically isolated island group with a small population, but it is a clean example of a national strategy. The package: stringent border controls with tightly restricted entry and pre-departure and post-arrival testing; 14-day quarantine, initially self-isolation and later supervised hotel-managed isolation; strict test, trace and isolate; and local or national lockdowns whenever domestic transmission was detected or at high risk. First recorded case 28 February 2020; by 8 June 2020 all domestic NPIs had been lifted, with a total of 1,504

cases and 22 deaths. The country remained mostly transmission-free until late 2021, despite regular positive tests among quarantined arrivals. Delta was first detected in August 2021; facing rising daily cases and the prospect of an extended lockdown in a highly vaccinated population, the government declared the end of the elimination strategy on 4 October 2021. Border controls stayed; Omicron was detected in arriving travellers in mid-December, community transmission was identified on 23 January 2022, and a large wave followed.

South Korea had experienced a MERS outbreak in 2015 with 186 cases and 38 deaths, which prompted significant reform of pandemic preparedness. Testing infrastructure was established and ready to deploy nationwide in drive-through facilities, giving effective estimates of caseload, and was coupled with GPS data from mobile phones used to monitor movements, alert citizens who had been near a confirmed case, and instruct them to isolate. Arrivals were quarantined for 14 days, and those from Hubei were banned outright. Compliance was demonstrably higher than during MERS, suggesting the population was more conscious of the risks, and early adoption of these packages contained the pandemic effectively enough that an early lockdown was avoided.

The summary draws two lessons and one warning from the three. The lesson: it was possible, in certain

contexts, to control transmission for over a year through early stringent border controls plus strict domestic measures. And the effectiveness of NPIs varied inversely with the transmissibility of the virus — as increasingly transmissible variants evolved, particularly Omicron, control became harder and harder even with the most stringent application. The warning: these results cannot simply be replicated elsewhere. National contexts varied by geography, politics, demography, socio-economics and regulation, and cross-country comparisons are affected by differences in demographics, healthcare systems, economic prosperity, trust between citizens and public authorities, and testing and reporting. Countries were differentially affected depending on older populations, obesity levels, chronic disease, concentrations of lower-income and larger households, and population density. Even death counting differed: Belgium included all deaths where COVID-19 was suspected to contribute, producing higher early reported rates, while others counted only hospital deaths.

The conclusion, and the lessons for next time

There is clear evidence that stringent implementation of packages of NPIs was effective in some countries in reducing transmission, and evidence for individual measures too — though as the pandemic progressed

and the virus became more transmissible, NPIs became less effective.

The common denominator across both the individual-measure reviews and the national case studies: NPIs were generally more effective when case numbers and the associated transmission intensity were lower. The mechanism is stated plainly — the size of the exposure, and therefore the risk of infection, of uninfected non-immune people is proportional to the number of cases in the community. Stringency also mattered, with respirator masks more effective than surgical masks and two weeks of quarantine more effective than shorter periods.

The lessons are addressed to three audiences. Policymakers and their advisers need to learn from national and international experience and understand the differing contexts in which measures were implemented, because national context was an important influence on outcomes. Researchers and funders should recognise that observational studies can be made much more useful if national and international collaborations are established in advance of the next pandemic, with standardised protocols for data collection — randomised trials should not be discounted, but most information in a future pandemic will still be observational, and the natural variation in how measures are implemented within and between countries could be exploited far

better if protocols were prepared beforehand, including protocols designed to disaggregate effects by social group and demographic factors.

And the recommendation that reaches furthest: future assessments should consider the costs as well as the benefits of NPIs — impacts on livelihoods, economies, education, social cohesion, and physical and mental wellbeing. The analogy offered is drug regulation, where recommendations rest on evidence of both effects and side effects. Policymakers will make better decisions on measures that are, in the main, complex social interventions if they have evidence on broader health and societal impacts to weigh alongside transmission effects. That requires pre-planned protocols, in some cases prior research, and an embedded system of expert research advice to help policymakers make extremely difficult decisions during a severe pandemic.

The final point is the one that connects this reading to the rest of the week: effective application of NPIs buys time for the development, evaluation and manufacture of treatments and vaccines at scale. So combinations of NPIs will matter in future pandemics, particularly early on, with novel pathogens, when knowledge is thin and therapeutics do not yet exist.

Words to know

- **Non-pharmaceutical intervention (NPI)** — any outbreak measure to reduce transmission that is not a vaccine or a drug. Behavioural, social, physical or regulatory.
- **Randomised controlled trial (RCT), and the cluster-randomised version** — the design that randomises matched groups to different interventions, allowing confident attribution; the cluster version randomises populations or places rather than individuals, which is how you would trial a mask mandate or a ventilation standard. Barely used for NPIs.
- **Observational study** — evaluation by comparison with historical, geographical or modelled data rather than by randomisation. Real-world scale, but vulnerable to confounding.
- **Bias, in evidence review** — not partiality, but features of a study's design that reduce the reliability of its conclusions, and which can push an estimate up or down.
- **Cohorting** — splitting residents of a care home into groups, each cared for by a specific group of staff, so an infection cannot spread across the whole home.
- **Test, trace and isolate** — the three-step chain-breaking strategy: find infectious people, find their contacts, keep both away from others until the risk has passed.

- **Exposure-based and symptomatic screening** — border checks based on health questionnaires or on symptoms and temperature. Both were found to have no meaningful effect.
- **Stringency** — how demanding and strictly enforced a measure is. One of the few variables that consistently predicted effectiveness.

Numbers in this reading

- Six: the categories of NPI covered, and the number of commissioned evidence reviews.
- Two metres: the commonly recommended minimum separation during the pandemic.
- More than 47,000 patients: enrolment in the RECOVERY trial, which the summary contrasts with the near-absence of NPI trials.
- Three vaccines approved in the UK during December 2020 and January 2021; about half the UK population had two doses by July 2021; half the global population not until January 2022; about 63% by January 2023.
- 8,098 SARS cases worldwide in 2003, of which Hong Kong had almost a quarter, with 302 deaths.
- 21 days of isolation for positive cases and 14 days for contacts in Hong Kong; only an estimated 27% of Hong Kong's cases were confirmed by laboratory test; about 60% of the population vaccinated by the end of 2021; more than 10,000 deaths in the

Omicron wave, largely among vulnerable elderly unvaccinated people.

- New Zealand: first case 28 February 2020; all domestic NPIs lifted by 8 June 2020 with 1,504 cases and 22 deaths; 14-day quarantine; Delta detected August 2021; elimination strategy ended 4 October 2021; Omicron detected in arrivals mid-December 2021; community transmission identified 23 January 2022.
- South Korea: 186 MERS cases and 38 deaths in 2015; 14-day quarantine for arrivals; outright ban on arrivals from Hubei.
- About 18 months: how long the three case-study locations held transmission at very low levels.
- Two weeks of quarantine more effective than shorter periods: one of the two stringency findings, the other being respirators over surgical masks.

What the author is unsure about

The summary is largely a careful account of what cannot be concluded. Controlled trials played only a small role, so most evidence is observational and subject to confounding. Because measures were almost always implemented in combination, it was extremely hard to isolate the effect of any single one. Many studies used routinely collected data never designed for evaluation, and comparison groups were often missing or poorly matched. For environmental controls specifically, individual contributions could not

be accurately quantified, effects were judged only within the setting studied rather than at population level, and there was insufficient evidence to judge surface cleaning or barriers at all. Few studies examined border measures implemented on their own. Communications findings are confined to the UK because they do not transfer. The three country case studies cannot be replicated elsewhere, and cross-country comparisons are confounded by demography, healthcare systems, prosperity, trust, and differences in testing and in how deaths were counted. And two deliberate silences are stated up front: the report makes no judgement on the timing or manner in which measures were applied, and no assessment of their economic or other societal impacts — while recommending that future assessments must do exactly that.

For the discussion

- The sentence to bring is the one that reorders every argument about lockdowns: NPIs were more effective when case numbers were lower, because exposure risk is proportional to community prevalence — which means the same measure can look effective or useless depending only on when it was imposed.
- The strongest practical finding at the border is counterintuitive and specific: symptom and temperature screening had no meaningful effect,

while quarantine on entry had the highest effectiveness of any border measure, so the visible intervention was the useless one.

- The tension worth raising is the report's own boundary: it assesses only transmission and explicitly leaves out economic, educational and mental health costs, then recommends that future work weigh both — which is an admission that the most contested question of the pandemic was never inside the study design.

4.5-2 • FURTHER • COURSE TIME 15 MIN

Epidemiological impacts of the NHS COVID-19 app in England and Wales throughout its first year

Michelle Kendall, Daphne Tsallis, Chris Wymant et al. • 2023

[Read the original paper →](#) • the course assigns: whole paper, including methods

England and Wales's contact tracing app reached roughly a quarter of the population, sent seven million exposure notifications in its first year, and the authors estimate it averted about a million cases and 9,600 deaths.

Average summary

This is the full-year evaluation of the NHS COVID-19 app, launched on 24 September 2020 across England and Wales using Google and Apple's Bluetooth exposure notification platform. The paper matters for two reasons: it is the strongest evidence anywhere that digital contact tracing works at national scale, and it is unusually honest about how wide the uncertainty is.

Uptake. Active users passed 10 million within days. After a December 2020 release fixed the counting, the figure settled around **13.5 million — 23% of the total population, or 29% of the eligible over-16 population.** Of those, between 71 and 88% had Bluetooth contact tracing switched on, a share that broadly declined over the year. When venue check-in became mandatory in some public places during the spring 2021 reopening, uptake jumped: active users reached 18 million in early July 2021, 38% of the eligible population. Then came the wave of notifications in June and July that the media called the "pingdemic," and the numbers fell steadily; by the end of the study period 14.6 million were active and 10.6 million had contact tracing enabled.

Volume. Over the year there were about **2,138,000 app-reported cases and 7,005,000 notifications** — an average of 3.28 notifications per app-recorded case. Between 40 and 55% of app users who

recorded a positive test consented to contact tracing. Notifications usually arrived within four hours.

Accuracy. The paper's most useful analysis asks whether notified people were actually at risk. Notified app users were at least twice as likely to report a positive test as a randomly selected person over 16 in England, at every point in the year, peaking at 26 times more likely in May 2021. Compared with app users who had not recently been notified, the odds of testing positive were usually above 3 and peaked at 77. The authors show the drop in apparent accuracy later in the year came mostly from higher prevalence in the population, not from the app getting worse.

Impact. Their headline estimate is that the contact tracing function alone averted about **1 million cases** (sensitivity analysis 450,000 to 1,400,000), corresponding to **44,000 hospital cases** (20,000 to 60,000) and **9,600 deaths** (4,600 to 13,000).

The comparison with manual contact tracing is the structural insight. The two work in complementary ways: manual tracing reaches more contacts inside households, schools and nurseries, while the app reaches people whose details the case does not have or whose company they do not remember, such as on public transport. And the app scales — it processes exponentially rising case numbers, whereas manual tracing's reach is set by staffing, so its contacts-per-

case figure actually fell during periods of high case load.

Why it's assigned

The course page lists contact tracing as one of the four social interventions, and asks what predicts strong adherence and which countries executed well. This paper is the quantitative case for the digital version, and it sits directly under the Royal Society summary's finding of strong evidence for contact tracing apps. For the withstand layer it is the clearest example of an intervention whose effect depends almost entirely on uptake and adherence rather than on the technology.

Words to know

- **Exposure notification** — the alert sent to a phone that was near a device belonging to someone who later tested positive. The protocol is privacy-preserving: the app does not track location, and the data the researchers used is anonymised and aggregated.
- **Active users** — devices with the app installed that contacted the servers on a given day, which includes people who had paused Bluetooth tracing, so it overstates how many were actually traceable.
- **TPAEN (testing positive after exposure notification)** — the share of notified users who later report a positive test. A lower bound on how

many notified people were infected, since not everyone tests or reports.

- **Venue check-in** — scanning a QR code on entering a place, which became mandatory in some settings and drove a visible jump in app uptake.
- **Case hospitalisation rate and case fatality rate** — the share of known cases that end in hospital or in death. Both fell sharply across waves, which is why averted cases translate into fewer deaths later than earlier.
- **Counterfactual** — the world the estimate is measured against: here, the same pandemic with every other intervention and behaviour unchanged but no app.
- **Risky contact** — their term for a physical interaction where an infectious person could infect a susceptible one. They assume a notification cut these by 60% before Delta and 40% afterwards.

Numbers in this reading

- Launched 24 September 2020; over 10 million active users within days.
- About 13.5 million active users after stabilisation: 23% of the total population, 29% of the eligible over-16 population.
- 71 to 88%: the share of active users with Bluetooth contact tracing enabled, broadly declining over the year.

- 18 million active users in early July 2021 (38% of the eligible population); 13.9 million with contact tracing enabled in late June 2021 (29%); 14.6 million active and 10.6 million traceable at the end of the study on 24 September 2021.
- About 2,138,000 app-reported cases and 7,005,000 notifications; 3.28 notifications per app-recorded case on average.
- 40 to 55%: the share of app users testing positive who consented to contact tracing.
- Within 4 hours: the usual delay before a notification arrives.
- 5 to 7%: the previously estimated share of notified users who went on to test positive, October to December 2020.
- At least 2 times, and up to 26 times (16 to 46) in May 2021: how much more likely a notified app user was to report a positive test than a random person over 16 in England.
- Usually above 3, lowest 2.4, peak 77 (54 to 106): the odds of testing positive for recently notified app users relative to app users not recently notified.
- About 1 million cases averted (450,000 to 1,400,000), 44,000 hospital cases (20,000 to 60,000) and 9,600 deaths (4,600 to 13,000) in the first year.

- 13% (5 to 19%) or 24% (14 to 33%): the earlier estimate of the reduction in total cases over the app's first three months, depending on attribution technique.
- 2.9 times more notifications per index case after the 29 October 2020 release, with the estimated prevention effect rising by a factor of 2.4 to 2.8 — the closest thing in the study to an experiment.
- 60% (38 to 82%) and 40% (20 to 60%): the assumed reduction in risky contacts after a notification, before and after Delta.
- 9.8%, 5% and 2.7%: case hospitalisation rates used for the pre-Alpha, Alpha and Delta waves; 1.5%, 1.9% and 0.2% for case fatality.
- 1.5 to 3 times: how much higher ONS estimated true infections were than reported cases.
- 68%: the share of recent close contacts reached by manual tracing in its final week, before routine contact tracing ended in England on 24 February 2022.

What the author is unsure about

The confidence intervals are wide and the authors explain exactly why: they rely on an estimated share of notified people who are infected, estimated delays, estimated overlap with other ways of discovering infection, and estimated adherence to app guidance. Several assumptions are flagged as poorly informed by data, particularly the timing of a notification

relative to learning you are at risk by other means, and how much behaviour changed as a result. They assume app users and their onward chains of contacts are representative of England, which is unlikely for short transmission chains; they assume transmission chains stay within the same local authority, which they call a limitation; and they neglect saturation of risk for repeatedly exposed people. The Alpha-wave hospitalisation rate is likely an underestimate. Conversely, they argue their total is probably an underestimate, because they credit the app only for behaviour change between a notification and learning the same thing another way, and because indirect effects are not captured: pre-emptive reduction in socialising to avoid being notified, network effects that make app users less likely to be infected in the first place, and the app's non-tracing functions. They could not extend the earlier spatial analysis across the full year because of greater geographical mixing, immunity and vaccination. And they explicitly did not assess the "pingdemic" of June and July 2021, which was followed by a steep decline in cases that others have hypothesised was causal.

For the discussion

- The number to carry is the ratio, not the total: a notified user was at least twice and sometimes 26 times more likely to test positive than a random

adult, which is the measure of whether a notification means anything.

- The complementarity finding is the practical one for a city: manual tracing reaches households and schools and collapses under high case loads, while the app reaches strangers on transit and scales without limit, so they are not substitutes and funding one against the other is a mistake.
- Worth raising as the honest caveat: the authors' own counterfactual assumes everything else about the pandemic stayed the same, including that no lockdown would have been timed differently — which is the same confounding problem the Royal Society review describes for every other social intervention.

4.5-3 • FURTHER • COURSE TIME 35 MIN

Effectiveness of international border control measures during the COVID-19 pandemic: a narrative synthesis of published systematic reviews

Karen Ann Grépin, John Aston and Jacob Burns · 2023

[Read the original paper →](#) · the course assigns: whole paper (Europe PMC full text, article body)

Every country closed its borders in some way and almost nobody studied it properly: of the border measures that can be assessed,

quarantine was the most effective, symptom and temperature screening were not effective at all, and the early travel bans worked for a few weeks and then stopped mattering.

Average summary

This is a synthesis of syntheses, written for the Royal Society's NPI programme by three authors who between them wrote three of the five systematic reviews it covers. Rather than pooling primary studies, it reads the five reviews in publication order to see where consensus emerged.

The scale of what happened first: by early 2020 **every country had adopted some form of international border control measure**, some maintained into 2023, and the range was unprecedented — complete closures, vaccination requirements for arrivals, centralised hotel quarantine. Before COVID-19 the assumption had been that border measures were not worth implementing.

The authors insist on separating measures usually lumped together, because they do different jobs through different mechanisms: **border closure** (complete restriction both ways), **travel restrictions** (partial, aimed at specific places or travellers), **screening** (symptomatic, exposure-based, or diagnostic testing, plus exemptions for vaccinated or recently infected travellers), and **quarantine**

(separating travellers not known to be infected; isolation is for those known to be).

Their findings, in order of confidence:

- **Symptomatic and exposure-based screening**, including temperature checks, was not effective enough to meaningfully reduce importation or transmission: too few cases relative to the number of travellers, and too many asymptomatic infections.
- **Diagnostic screening** (PCR or antigen testing) was better, with studies suggesting detection of 50 to 90% of cases, but at best such testing alone is unlikely to find more than about half of potentially infected travellers. More recent work suggests screening on vaccination or recent infection status performed even better in the short term.
- **Targeted travel restrictions** on travellers from Wuhan and Hubei were effective, with an estimated **70 to 80% of exported cases prevented** in the weeks after 23 and 24 January 2020. Additional bans on other countries added on top of those reduced exported cases by only a few hundred through the end of February 2020, against roughly 80,000 confirmed cases in China. Both became relatively ineffective within weeks as other places became the dominant source.
- **Quarantine had the highest effectiveness of any single measure**, but the evidence comes from

long periods, up to 14 days. Compliance mattered, and was lower when self-monitored than in mandatory facility-based quarantine. Combined with testing, quarantine was more effective and could be shortened: four observational studies found 68 to 92% detection, and the forthcoming update found 14-day and 7- or 8-day quarantines both detected around 80 to 95% of cases with multiple tests.

Two structural points close the paper. The comprehensive border regimes that let Hong Kong, mainland China, Singapore and New Zealand hold zero-COVID status for months were largely excluded from the reviews, which were built to evaluate individual measures — so the cases that worked best are the ones the evidence base says least about. And border measures are among the most economically costly interventions: cross-border and informal trade accounts for 40% of GDP and 55.7% of employment in sub-Saharan Africa. Showing that a measure works may not justify it when cheaper interventions are at least as effective.

Why it's assigned

Border closures are one of the four social interventions on the course page, and this is the best available answer to whether they worked. It also answers the page's question about why there is so much argument: 135 studies exist, only 43 are

observational, almost all of the observational ones cover the first few months of 2020, and hardly any controlled for the other measures running at the same time. The paper is a lesson in reading an evidence base rather than a verdict.

Words to know

- **Narrative synthesis** — a method for integrating evidence from multiple sources in text when the studies are too heterogeneous to pool statistically. The authors follow the SWiM guidelines, written for exactly this situation.
- **Modelling versus observational study** — a model predicts what a measure would do under stated assumptions; an observational study measures what happened. Of the 135 studies here, 92 were modelled, which is the central weakness of the field.
- **Quarantine versus isolation** — quarantine separates travellers not known to be infected; isolation separates those known to be. The terms are used interchangeably in practice, which muddies the literature.
- **Diagnostic-based screening** — testing travellers with PCR or rapid antigen tests, as opposed to taking temperatures or asking about symptoms.
- **Risk of bias tools** — ROBINS-I for observational intervention studies, QUADAS-2 for studies of screening performance, GRADE for rating overall

certainty. All five reviews used some of these, and nearly every body of evidence came out low or very low certainty.

- **Phylogenetic evidence** — using the genetic sequences of viral lineages to see which were imported and when, an approach outside the reviews' inclusion criteria that the authors argue is informative.
- **PHEIC (public health emergency of international concern)** — the WHO declaration made under the International Health Regulations. When it was declared for COVID-19, the WHO recommended no international border measures, citing limited evidence.

Numbers in this reading

- Five published systematic reviews synthesised, with a sixth in progress; 135 studies pooled across them, of which 88 (65%) appeared in only one review and only 43 (32%) were observational.
- 65% and 85% of the observational studies had study periods complete by the end of April and of June 2020 respectively; 93% evaluated national-level measures.
- 70 to 80%: the estimated share of exported cases prevented in the weeks following the Wuhan (23 January 2020) and Hubei (24 January 2020) outbound travel restrictions.

- A few hundred exported cases: the estimated effect of other countries' additional measures against Chinese travellers through the end of February 2020, against approximately 80,000 confirmed cases in China by then.
- 26 to 90% reduction in imported cases and up to 26 days of delay to an outbreak: the early review's estimate for international travel restrictions, with low confidence in the precision.
- 1.8% to 97.8%: the range of estimated reductions in community cases across 13 studies of travel restrictions, 10 of which found a positive effect.
- 50 to 90%: the share of cases diagnostic screening could detect, described as highly context-specific; at best, testing alone is unlikely to detect more than about half of potentially infected travellers.
- 68 to 92%: detection rates in four observational studies combining testing with quarantine; around 80 to 95% for both 14-day and 7- or 8-day quarantine with multiple testing in the forthcoming update.
- 2.3 to 7.5 days: the range of estimates for SARS-CoV-2's mean generation time at different points in the pandemic, shortening with new variants — the reason a shorter quarantine with testing can match a longer one.
- Up to 95%: the share of symptomatic Omicron cases that developed symptoms within six days, in

a Korean study outside the reviews.

- 10 to 72%: the reduction in transmission risk during travel from pre-departure PCR testing, depending on the interval between test and travel.
- 86 to 98%: the reduction in imported viral lineages attributed to Switzerland's strictest border measures in phylogenetic work.
- 40% of GDP and 55.7% of employment: the share accounted for by cross-border and informal trade in sub-Saharan Africa, cited to show how costly border measures can be.

What the author is unsure about

The paper is mostly an inventory of what the evidence cannot support. Certainty was rated low or very low for almost every body of evidence across all five reviews. Border measures were rarely the only measures in place and few studies controlled for the others, which likely biased effectiveness estimates upward; only a handful used quasi-experimental designs. Most of the literature is modelling rather than observational, and the quality of a model depends on assumptions that are hard to assess. Many observational studies used ecological comparisons across countries with very different tracking and enforcement, or composite policy indices that lump different measures together. There is very limited evidence on border measures implemented after community transmission was established, despite countries continuing them for

years, and almost none on measures aimed at stopping new variants of concern. The global collapse in travel means even countries without measures saw fewer imported cases, and measures in one country displaced travellers to others — effects few studies accounted for. Geographical coverage is thin outside Australia, Hong Kong, Singapore and Japan, and the authors note that over-representing places with robust measures and good data may bias the overall findings toward higher effectiveness. The number of studies was too small to draw firm conclusions on the optimal quarantine length or testing schedule. And the comprehensive regimes that actually held the line were excluded by the reviews' own inclusion criteria, which the authors say may miss important lessons.

For the discussion

- The cleanest finding to bring is that the most visible border measure was the least useful: symptom and temperature screening had no meaningful effect, while quarantine — the costly, unglamorous one — was the most effective single intervention.
- The generation-time argument is the policy-relevant one: because SARS-CoV-2's mean generation time ran from 2.3 to 7.5 days and shortened with new variants, a 7-day quarantine with repeated testing detected roughly as many cases as 14 days, which converts a political fight

about duration into a question about testing schedules.

- The tension worth raising: the places that succeeded did so with whole regimes, not single measures, and those are precisely the cases a systematic review of individual measures has to exclude — so the strongest practical evidence sits outside the strongest methodological evidence.

4.5-4 • FURTHER • COURSE TIME 20 MIN

International air travel-related control measures to contain the Covid-19 pandemic: A companion review to a Cochrane rapid review

Ameer S-J Hohlfeld, Leila Abdullahi, Ahmed M Abou-Setta and Mark E Engel • 2022

[Read the original paper →](#) • the course assigns: whole review, including the three summary-of-findings tables

A WHO-commissioned update to the Cochrane review on air travel measures: quarantine, testing, and especially the two together look effective, longer quarantine and more tests detect more cases, and every single finding is rated very low certainty.

Average summary

This review updates a Cochrane rapid review, following the same methods, and covers the literature

from November 2020 to 6 June 2022. From 6,202 citations screened, 22 studies were included — **nine modelling and 13 observational** — alongside the studies in the original Cochrane review. The interventions are narrow and specific: quarantine, microbiological detection (testing), or a combination of the two, applied to international air travel. The outcomes are cases avoided, a shift in epidemic development, and cases detected.

The findings line up in a clear order.

Quarantine length is the dominant variable.

Reductions in the proportion of imported cases ranged from 55% for 7 days to 99% for 21 days, with one study putting a 21-day quarantine at **98.6% effectiveness against imported cases, against 91.8% for 14 days and 70.4% for 7 days.**

Quarantining arrivals from high-risk countries was estimated to cut imported cases by 89 to 99%. The probability that an infectious person leaves a 14-day quarantine with no testing was put at 4% under moderate transmission and 28% under high transmission.

Testing works best stacked and timed. Pre-departure testing reduced the risk of transmission during travel by **10 to 72%**, rising as the gap between test and travel shrinks. Adding a test on arrival improved detection of imported cases from 57.2% (with a 2-day pre-departure test) to 84.8%. A one-day

interval between pre-departure and arrival tests risks false negatives, with only 50% of infected travellers detected.

The combination is the strongest option.

Quarantine with a day-7 PCR test gave a median 94% reduction in infectious travellers released into the community compared with no quarantine and no test; a 14-day quarantine alone exceeded 99%; a 5-day quarantine with a negative PCR gave 89%. The best overall performance came from a pre-departure test plus PCR on arrival plus a rapid antigen test on leaving quarantine: detection of 91.2% with three days of quarantine, rising to 99.3% with 21 days.

The modelled outbreak-delay figures are the most striking and the shakiest. Pre-flight saliva testing alone was estimated to delay an outbreak by 0.3 years, with the annual risk of an outbreak at 88.4%. Adding PCR tests on days 1, 3 and 12 pushed the delay to 0.7 years; adding mask use until the last test, symptom reporting and contact tracing pushed it to 1.5 years and cut the annual outbreak risk to 36.6%. With a 21-day quarantine the modelled risk approaches zero.

The real-world detection numbers are small.

Across ten observational studies, the share of all travellers found positive ran from 0.01% to 1.5%.

The authors' bottom line: weak evidence supports using these measures, and more real-world studies are needed.

Why it's assigned

This is the detail behind the Royal Society summary's finding that quarantine at the border was the most effective border measure, and it is the only reading in the unit that prices quarantine duration against testing schedule. It also demonstrates what the course page means when it asks why there is so much argument about social interventions: the summary-of-findings tables here are a column of effect sizes next to a column of "very low certainty" ratings, every one downgraded for a stated reason.

Words to know

- **Microbiological detection** — the review's term for testing travellers, by PCR or rapid antigen test, as opposed to symptom or temperature screening.
- **Cochrane rapid review** — a systematic review conducted on a compressed timeline using Cochrane methods. This paper is a companion, meaning it applies the same methods to later literature so the two can be read together.
- **GRADE** — the system for rating how much confidence to place in a body of evidence, from high to very low. Every outcome in this review came out very low.

- **ROBINS-I** — the tool used to assess risk of bias in non-randomised intervention studies. Five of the included observational studies were rated critical risk and three moderate.
- **Shift in epidemic development** — the outcome category covering time to diagnosis, time to outbreak and risk of outbreak, as distinct from cases avoided or detected.
- **False negative** — a test that misses an infection, typically because it was taken too soon after exposure. The reason a one-day gap between tests detected only half of infected travellers.
- **Sensitivity** — the share of true infections a test detects. The saliva test in the outbreak-delay model was assumed to have 62% sensitivity.

Numbers in this reading

- 6,202 citations screened; 22 studies included (nine modelling, 13 observational), added to those in the original Cochrane review.
- 55% (7 days) to 99% (21 days): the range of reductions in the proportion of imported cases attributed to quarantine.
- 98.6%, 91.8% and 70.4%: one study's efficiency estimates for 21-day, 14-day and 7-day quarantine at averting imported cases.
- 89 to 99%: the reduction in imported cases from quarantining passengers arriving from high-risk countries.

- 4% and 28%: the probability of an infectious person leaving a 14-day quarantine with no scheduled testing, under moderate and high transmission.
- 10 to 72%: the reduction in transmission risk during travel from pre-departure testing, greatest when the test is closest to the day of travel.
- 57.2% rising to 84.8%: the improvement in detection of imported cases when entry testing is added to a 2-day pre-departure test; 83.1% for rapid antigen pre-departure plus on arrival.
- 50%: the share of infected travellers detected when only one day separates the pre-departure and arrival tests.
- 94% median reduction (89 to 98%) in infectious arrivals released, for quarantine plus a day-7 PCR with a one-day result delay; over 99% (98 to 100%) for 14-day quarantine with no testing; 89% (83 to 95%) for a 5-day quarantine with a negative PCR.
- 91.2% to 99.3%: detection with pre-departure test plus arrival PCR plus rapid antigen on exit, for 3 days and 21 days of quarantine respectively.
- 0.3 years delay to an outbreak from pre-flight saliva testing alone (62% assumed sensitivity; 4 days to 1.7 years; mean 170 flights); 0.4 years with day-1 testing (225 flights); 0.7 years with days 1, 3 and 12 (359 flights); 1.5 years with the full combination including masks, symptom reporting and contact tracing (802 flights).

- 88.4% annual risk of an outbreak with pre-flight testing alone, falling to 36.6% with the full combination; 67.8% with pre-flight testing plus 7-day quarantine, 13.1% at 14 days, virtually zero at 21 days.
- 0.01% to 1.5%: the share of all travellers detected as positive across ten observational studies of testing.
- 0.87% to 1.2% positivity on arrival in the combination studies; 0.2% and 0.6% among more than 25,000 PCR-negative arrivals to Bahrain after 10 to 14 days of quarantine; 26.1% positivity among 337 symptomatic quarantined individuals, with 63.6% of those detected in the first week.
- 5.1% and 23.3%: the reduction in local cases in China and Singapore from quarantining only symptomatic travellers; 30 to 35% from immediate isolation at symptom onset before or during travel; 64 to 95% at 7 days and 96 to 100% at 14 days for quarantine after exposure with no monitoring or testing.
- 1.7 days: the maximum reduction in the mean interval from arrival to laboratory confirmation.

What the author is unsure about

The certainty rating for every outcome in all three tables is very low, and each is downgraded for named reasons: lack of external validation, concerns about model structure or uncertainty assessment,

incomplete technical documentation, imprecision from a single contributing study or a wide range of plausible effects, and indirectness because the travellers studied were not representative of usual travellers. Overall risk of bias was critical in five observational studies and moderate in three. The authors say plainly that the actual effects may differ considerably from what is reported, in either direction, and that it is questionable whether screening detects enough cases to prevent regional outbreaks. Three limitations are structural. No included study reported any data on COVID-19 vaccines, which limits what the findings say about effectiveness once populations had immunity. None reported the financial or human resources needed to run these measures, which the authors note matters most for low- and middle-income countries with already overburdened systems. And none reported the bio-psycho-social impacts on the people quarantined — mental wellbeing, the financial effect of being unable to work, the social treatment of those with a positive diagnosis. They also note that each study was conducted in a different place at a different point in the pandemic, so waves and variants differed between them. The paper also carries an internal inconsistency worth noticing: the results section describes the Cochrane review as containing 62 studies while the conclusion refers to 31.

For the discussion

- The useful pattern across the tables is that layering beats lengthening: a short quarantine with a well-timed test sequence reaches detection rates close to a long one, which is the difference between a policy people comply with and one they do not.
- The timing finding is the kind of operational detail that decides whether a measure works at all — testing one day before departure and one day after arrival caught only half of infected travellers, because both tests fall inside the window where an infection is invisible.
- The honest framing for the group: this review reports impressive-looking percentages and then rates every one of them very low certainty, which is a good test of whether the room can hold a number and its uncertainty at the same time.