THOUGHT LEADERSHIP | EQUITABLE ACCESS

Equitable Access to the Air We Share

Immunocompromise, disability law, and the economic case for evaluating cleaner indoor air as accessible infrastructure

PJ Piper | CEO, Far UV Technologies | September 2026

Disclosure: The author is CEO of a company that manufactures germicidal ultraviolet air-treatment technology. The legal, economic, and engineering arguments below are offered for evaluation on their merits.

The Policy Question

The Americans with Disabilities Act changed the built environment after the nation recognized that a doorway, stair, or restroom designed for only some bodies could exclude people from work, education, commerce, and civic life. For people whose immune systems make respiratory infection unusually consequential, the barrier may be less visible but no less practical: they can enter a building and still be unable to participate safely in the activity inside it. The policy question is therefore narrow but consequential: **when should cleaner indoor air be evaluated as a reasonable accommodation or modification, and when should it be treated as ordinary accessible infrastructure?**

Existing federal law does not supply a universal clean-air mandate, and this paper does not assume one. It asks whether the ADA's governing idea—remove avoidable environmental barriers so qualified people can participate—offers a useful framework for individualized accommodations, targeted public investment, and future building standards.

The ADA's Governing Idea: Change the Environment

The legal origin story begins with Section 504 of the Rehabilitation Act of 1973, the first federal civil-rights protection against disability discrimination in federally funded programs. The ADA extended that principle across employment, state and local government, public accommodations, transportation, and communications. Its support was unusually broad: the Senate initially passed the bill 76–8; the final conference measure passed the House 377–28 and the Senate 91–6 before President George H.W. Bush signed it on July 26, 1990. Congressional action and vote history confirms that bipartisan record.

Wheelchair access became the ADA's most visible architectural expression, even though the statute protects a much wider range of disabilities. DOJ's demographic analysis estimated that about 1.4% of Americans age 15 or older used a wheelchair or similar mobility device. The country nonetheless made accessible routes, entrances, door clearances, parking, and restrooms routine features of new construction and alterations under the ADA Standards for Accessible Design. The underlying judgment was not that ramps eliminated disability; it was that design should not impose a needless barrier.

A useful economic precedent exists, although not a ramp-only ROI. In 2010, DOJ commissioned a benefit-cost analysis of the revised ADA design standards. At a 7% discount rate, it estimated $22.0 billion in benefits and $12.8 billion in costs—a 1.7-to-1 benefit-cost ratio and $9.3 billion in net present value. At 3%, estimated benefits were $66.2 billion against $25.8 billion in costs, a 2.6-to-1 ratio. DOJ attributed 93% of quantified benefits to time savings for users and noted additional benefits it could not monetize. These figures evaluate the 2010 revisions, not the entire 1990 ADA and not wheelchair ramps alone; they nonetheless show that accessibility rules can create measurable economic value as well as civil-rights value. The mechanism is not identical to the earnings-based estimates used later in this paper—these benefits were driven by physical-navigation time savings, not preserved wages—so the parallel should be read at the level of category, not magnitude: rigorous accessibility analysis can find net-positive economic value, not that a ramp's return rate predicts an air-quality return rate.

Extending the Principle to the Immunocompromised

The constituency is both larger and more varied than the public discussion often suggests. CDC has historically estimated that people who are moderately to severely immunocompromised comprise about 3% of U.S. adults. A nationally representative 2021 survey published in JAMA estimated broader self-reported immunosuppression at 6.6%—roughly one in fifteen adults, or about 17 million people when applied to the adult population. The estimates should be presented as a range, not merged: CDC's clinical category is narrower, while the JAMA measure includes a condition, prescription, treatment, or recent hematologic cancer reported to weaken the immune system.

This paper uses both estimates deliberately, for different purposes. The broader 17 million figure is the right lens for describing who has a stake in this question, including many people whose immunosuppression is real but comparatively mild. The exclusion-cost and ROI estimates later in this paper instead use CDC's narrower estimate—about 3% of U.S. adults, or roughly 7.7 million people—because the leave-duration evidence anchoring those estimates (cancer treatment, transplantation, graft-versus-host disease) describes people whose immune status is moderate to severe, not the full range the JAMA estimate captures. Sizing a severity-anchored cost estimate off the broader population would overstate its plausibility.

Common pathways include active treatment for solid tumors or blood cancers; organ or stem-cell transplantation and anti-rejection therapy; moderate or severe primary immunodeficiency present from birth; advanced or untreated HIV; and high-dose corticosteroids, biologics, chemotherapy, and other immunosuppressive medicines. CDC's clinical examples illustrate why this is not a single disease constituency. Long COVID is best treated as a recent catalyst and an overlapping disability population: HHS and DOJ guidance states that it can qualify as a disability when it substantially limits a major life activity, but qualification remains individualized and Long COVID should not be described as universally immunocompromising.

What Current Disability Law Establishes

The ADA does not currently require every covered building to meet a single infectious-aerosol standard. It does require covered employers, public entities, and public accommodations to consider reasonable accommodations or modifications for qualifying disabilities, subject to setting-specific limits such as undue hardship, direct threat, or fundamental alteration. The analysis is individualized: a diagnosis alone does not dictate the result, and an organization may provide an effective reasonable option even if it is not the requester's preferred option. The statute and implementing framework therefore support evaluation, not a predetermined engineering prescription.

Depending on the setting and documented need, options can include remote or hybrid participation, schedule or room changes, masking in defined circumstances, increased outdoor air, improved HVAC filtration, portable HEPA filtration, or appropriately designed germicidal ultraviolet air treatment. Each option carries its own cost, maintenance, and safety profile—germicidal UV, for instance, requires attention to wavelength selection and exposure limits—that a case-specific reasonableness analysis should weigh rather than assume away. The legal question is whether the measure is effective and reasonable in context—not whether all infection risk can be eliminated.

The Cases Have Begun to Define the Boundary

Disability status and leave. In **Brown v. Roanoke Rehabilitation & Healthcare Center**, a federal district court allowed an ADA claim based on severe, continuing COVID symptoms and a request for temporary leave to survive a motion to dismiss. The court stressed that disability status was fact-driven; it did not hold that every infection or every case of Long COVID is disabling.

Meaningful access and individualized review. In **L.E. v. Superintendent of Cobb County School District**, the Eleventh Circuit held that the relevant access question concerned in-person schooling, not education in the abstract. Its August 2026 opinion held that four medically vulnerable students had standing and noted that the requested masking accommodation was not facially unreasonable. It nevertheless affirmed denial of preliminary relief for two students because other individualized measures had been found largely effective, while remanding the other two claims.

Air filtration as an on-site accommodation. The Third Circuit's nonprecedential August 2026 decision in **Gardner v. Kutztown University** is especially relevant. A professor taking immunosuppressive medication sought remote teaching. The court held that a Spring 2023 package—limited classroom capacity, a separate entrance, specialized air filters, and plexiglass—was reasonable as a matter of law—reversing a lower-court judgment that had favored

the professor and declining to treat full-time remote work as required—while remanding disputes about earlier semesters. That outcome favored the institution's substitute measure over the professor's preferred accommodation, which is precisely why it matters here: it is not binding precedent, but it is a concrete example of a federal appellate court treating engineered air-risk mitigation, on its own, as a legally sufficient accommodation.

The mask cases also expose an implementation distinction. Mask requirements regulate visible human behavior and depend on repeated individual compliance; they became subjects of statutes, litigation, and political conflict. By contrast, NIOSH ranks engineering controls above administrative controls and personal protective equipment and notes that effective engineering controls need minimal user action. It is reasonable to infer—not yet proven by comparative political research—that well-designed ventilation, filtration, or air cleaning may encounter different and potentially lower compliance friction because the protection operates in the background.

The Economic Cost of Exclusion

A single national participation-loss rate for immunocompromised Americans does not exist, and applying a small arbitrary percentage to the full 17 million would obscure the real pattern: burden varies sharply, while severe subgroups can lose months or leave circulation entirely. More specific evidence provides a better anchor. In a nationally representative U.S. study, 41.3% of cancer survivors who had worked at or since diagnosis reported cancer-related employment changes; among those, 75.4% took extended paid time off and 46.1% made other changes such as moving to part-time or less demanding work. NCI separately reports that employed people receiving cancer treatment missed about 22 more workdays per year than workers not receiving treatment.

At the more severe end, a 2025 U.S. claims study of employed hematopoietic stem-cell transplant recipients with graft-versus-host disease found mean short-term disability leave of 103 days and long-term leave of 120 days. For patients with both claim types, mean leave reached 127 of 261 workdays—nearly half a work year. These subgroup findings cannot be generalized to every immunocompromised person, but they support duration-based scenarios measured in months rather than fractions of a percentage point.

Using the 2024 BLS median full-time wage of $1,192 per week ($61,984 annually), the direct-earnings scale is substantial. If 5% of the narrower, 7.7 million-person clinically-severe cohort is unavailable for three months, the annualized loss is about $6.0 billion. If 10% is unavailable for six months, it is about $23.9 billion; if 20% is unavailable for six months, about $47.7 billion. These are transparent scenarios, not prevalence estimates or causal forecasts. They exclude nonworkers, part-time variation, employer output, tax receipts, replacement costs, health spending, unpaid care, and missed education. They also assume, generously, that cleaner air could recover the full loss; in practice it addresses only the infection-risk share of absence—not the fatigue, treatment toxicity, or recovery time that drive much of the leave data cited above—so realized recovery would likely be a fraction of these totals.

The Social Multiplier

Exclusion does not stop with the individual, though what follows is a ceiling for illustration, not a precise count. Census QuickFacts reports 2.53 persons per U.S. household. Applied mechanically to the broader 17 million-person constituency, that arithmetic yields roughly 43 million person-household connections—before correcting for affected people who share a household, and while still excluding relatives living elsewhere, friends, caregivers, classmates, customers, and colleagues. The number should be read only as evidence that accessible air is a household- and community-level question, not an individual one.

Cleaner Air as Measurable Infrastructure—and an ROI Test

The engineering premise is administrable. ASHRAE Standard 241 establishes an equivalent-clean-air framework for reducing exposure to infectious aerosols in new and existing buildings. CDC/NIOSH ventilation guidance identifies outdoor air, filtration, portable air cleaners, and germicidal ultraviolet systems as tools that can reduce airborne pathogen concentrations when properly selected, designed, and commissioned. A policy can therefore specify performance, verification, and maintenance rather than mandate one device.

No credible national estimate yet establishes the annual cost or causal labor-market benefit of a defined clean-air accessibility program, so a claimed national ROI would be premature. A break-even test is more honest. At $61,984 per full-time-equivalent year, every $1 billion in annualized program cost would break even on direct earnings alone if it preserved about 16,100 FTE-years—about 0.209% of the narrower, 7.7 million-person clinically-severe cohort used in this section. The cost side of this test is illustrative spacing, not a vetted budget estimate: no published

source yet prices a defined national clean-air accessibility program by building type, square footage, or equipment class, and that costing is a necessary next step this paper does not attempt. The thresholds scale linearly:

Annualized policy cost	FTE-years needed to break even	Share of 7.7M cohort
$1 billion	16,100	0.209%
$5 billion	80,700	1.048%
$10 billion	161,300	2.095%

Illustrative ROI. If cleaner-air measures preserved 154,000 FTE-years—2% of the narrower, 7.7 million-person cohort, bounded to the infection-risk-attributable share of the leave data cited above—the direct earnings benefit would be about $9.5 billion. Against annualized costs of $1 billion, $5 billion, or $10 billion, the gross benefit-cost ratios would be approximately 9.5:1, 1.9:1, and 0.95:1, respectively. Net ROI would be about 850%, 90%, and -5%, respectively—each derived directly from the rounded $9.5 billion benefit figure (benefit ÷ cost, minus 1), so every number in this box is reproducible by hand from the figures printed here. The last result lands just below the $10 billion break-even threshold identified in the table above (2.095% of the cohort), rather than assuring a positive return at every scale. This remains a sensitivity analysis, not a forecast; it excludes health-care savings and spillover benefits, and it assumes the preserved work is attributable specifically to reduced infection exposure rather than to the other clinical factors that also drive leave.

A Policymaker's Decision Frame

The most defensible next step is to evaluate cleaner air with the same disciplines applied to other accessibility investments: define the covered setting and population; measure baseline air performance; select an equivalent-clean-air target; compare capital, operating, and maintenance costs with attendance, leave, health-utilization, and participation outcomes; and preserve individualized ADA review. Priority settings could be tested first where exclusion carries the highest consequences—health care, schools, workplaces, transportation, and essential public services. If the return in some settings is genuinely this favorable, the more interesting question is why it has not already been captured voluntarily; split incentives between building owners and occupants, and the same coordination and information failures that justified the ADA's original mandates, are the likely answers—and both point toward a policy role rather than away from one.

The ADA's historical lesson is not that every barrier warrants the same remedy. It is that design choices can determine who gets to participate—and that the economic value of access can far exceed the cost of providing it. For millions of people whose immune defenses are compromised by treatment, transplantation, medication, congenital disease, HIV, or other conditions, access to the air we share is now a concrete legal, engineering, and economic question.

Access to the building. Access to the air inside. Engineered air treatment—filtration, ventilation, or germicidal UV among them—can extend an accessible-design mindset from the entrance to the room itself. Illustrated conceptually.