
Wastewater metagenomic surveillance for novel viruses: how much sequencing is enough, and at what cost?¹

Harry Gray

With
Apart Research

Abstract

Wastewater metagenomics has been pitched as a universal early-warning layer for pandemics, including AI-enabled engineered viruses, but its practical limits are rarely quantified. Using estimates of virus abundance from recent published work, I map sequencing depth to detection probability for DNA, RNA enteric, RNA respiratory and a broad “unknown pathogen” at 1% incidence. Results show extreme heterogeneity: some plausible pathogens are easy to detect with tens of millions of reads, others remain effectively invisible even at $\sim 10^{12}$ reads. Robust designs that cover most unknowns are orders of magnitude more expensive than median-case designs, and still struggle against “stealth” scenarios. Wastewater metagenomics is valuable, but not a silver bullet.

Keywords: biosecurity; wastewater-based epidemiology; metagenomic sequencing; early warning systems; infectious disease surveillance; artificial intelligence; pandemic preparedness; defensive biosurveillance

1. Introduction

Pandemic-scale infectious diseases remain a central global risk, capable of causing mass mortality, economic disruption and long-term societal damage. Recent history has shown that even a naturally emerging virus, such as SARS-CoV-2, can exploit global connectivity to spread exponentially, with the outcome highly sensitive to the speed of detection and response (Nucleic Acid Observatory Consortium; 2021).

¹ Research conducted remotely for the Apart Research def/acc Hackathon, 2025

Large language models and related AI systems are increasingly being explored for tasks such as literature synthesis, sequence design and protocol optimisation, and early red-team studies indicate that they can already assist in planning and conceptualising large-scale biological attacks, potentially closing knowledge gaps for less expert actors (Mouton CA, Dunham JL, Baker GP, et al.; 2024). As AI tools become more capable and accessible, they may lower the barrier to engineering novel viral agents, accelerate research and development cycles, and enable strategic planning of outbreaks for maximal harm. In such an environment, defensive systems must be robust to rapid shifts in the threat landscape, including deliberately engineered pathogen properties.

Early-warning surveillance is one of the few interventions that can, in principle, change the game in our favour. The NAO proposal argues that sufficiently deep, routine sequencing of wastewater and waterways, combined with anomaly detection against local historical baselines, could flag any rapidly expanding biological signal, including uncharacterised or engineered agents that would evade existing diagnostic panels. This “universal early warning” framing has motivated increasing attention to environmental sequencing as a cornerstone of biosecurity.

Within this space, wastewater-based epidemiology (WBE) has emerged as a practical and relatively mature tool for population-level surveillance. Historically used to monitor poliovirus, WBE has expanded to a broad range of gastrointestinal and respiratory pathogens and rose to prominence during the COVID-19 pandemic as a cost-effective, scalable complement to clinical testing (Parkins MD, Lee BE, Acosta N, Bautista M, Hubert CRJ, Hruday SE, et al.; 2024). By aggregating excreted viral material from both symptomatic and asymptomatic individuals, wastewater surveillance can provide an anonymised, composite readout of community infection trends and has repeatedly detected SARS-CoV-2 signals days to weeks before corresponding increases in reported clinical cases (Parkins MD, Lee BE, Acosta N, Bautista M, Hubert CRJ, Hruday SE, et al.; 2024). National-scale systems such as the French SUM’EAU network and related programmes elsewhere now routinely integrate wastewater viral loads into public-health situational awareness and decision-making (Jourdain F, Bouzid D, Blanchon T, Roque-Afonso A-M, Bonifait L, Schvoerer E, et al.; 2025).

Current wastewater systems typically combine targeted molecular assays and, increasingly, metagenomic sequencing. Targeted RT-qPCR or digital PCR panels offer high sensitivity and relatively low cost for a predefined list of pathogens (for example SARS-CoV-2, influenza viruses, and enteroviruses), but they require continual maintenance as pathogens evolve around primer sites and, by design, cannot detect truly novel agents. Shotgun metagenomic sequencing, by contrast, can capture a wide spectrum of human, animal and plant viruses from a single composite sample, enabling detection of both expected and unexpected pathogens and characterising their seasonality. Recent ultra-deep sequencing work, generating on the order of 10^9 reads per weekly sample over 18 months, has shown that such unbiased approaches can simultaneously track endemic viruses (e.g. noroviruses, rotaviruses, SARS-CoV-2) and pick up unexpected signals such as hepatitis A surges and highly pathogenic avian influenza H5N1. This aligns with the NAO argument that sufficiently deep environmental metagenomics can, in

principle, detect any exponentially growing biological threat without needing to anticipate its exact genome in advance.

However, turning wastewater signals into decision-relevant estimates of risk is still difficult. Reviews of WBE emphasise unresolved challenges in linking measured viral concentrations to true infection incidence, handling transport- and dilution-related variability, and standardising methods across sites (Singh S, Goyal R, Chander Y, Gupta N.; 2024). Experimental work with virus-like DNA tracer particles underscores how advection, decay and measurement noise in sewer systems can distort the relationship between an upstream shedding event and a downstream measurement, and shows that we currently lack well-validated controls to characterise these effects in operational networks (Gopal A, Greenwald HD, Crits-Christoph A, Kantor RS, Krammer F, Marras SAE, et al.; 2023). More broadly, environmental surveillance data can be noisy, with dilution by rainfall, variable water use and complex microbial backgrounds complicating analysis and interpretation. At the same time, digital-twin and machine-learning approaches are beginning to demonstrate that it is possible to move from descriptive to predictive models of wastewater microbiome dynamics, suggesting a path toward more mechanistic use of these data in public health. (Gurung BDS, Gweon HS, Jabbari S, Shadoun S, Sails AD, Singer AC, et al.; 2025)

Recent statistical modelling has begun to quantify the sensitivity and cost of metagenomic wastewater surveillance (Grimm SL, Rushford C, Bley D, Kamm J, Crits-Christoph A, Kantor RS, et al.; 2025). Grimm et al. fit a hierarchical model to several untargeted wastewater metagenomic datasets and associated epidemiological indicators, using it to predict viral relative abundance at specified incidence levels and to infer the sequencing depth and financial cost needed to hit chosen detection thresholds. Even after adjusting for local infection levels, they estimate that viral relative abundance differs by several orders of magnitude between studies, consistent with large effects of sewershed characteristics and protocol choices on W-MGS sensitivity and cost. Targeted enrichment panels can massively boost apparent abundance and cut sequencing cost for the viruses they target, but this comes at the price of reduced agnosticism and weaker performance on genuinely novel agents (Grimm SL, Rushford C, Bley D, Kamm J, Crits-Christoph A, Kantor RS, et al.; 2025). Related work on global aircraft wastewater surveillance networks shows that relatively small numbers of strategically placed sentinel sites can, in principle, provide early detection of emerging respiratory pathogens, with diminishing returns beyond a critical network size (St-Onge C, Delarocque-Astagneau E, Schofield H, Zeller H, Van der Werf S, Obadia T, et al.; 2025). Together, these studies suggest both the promise and the design constraints of metagenomic wastewater surveillance as a global early-warning tool.

From a biosecurity and AI-safety perspective, this omission matters. An attacker with access to AI-assisted design tools could, in principle, seek to engineer viruses with atypical genomic signatures, altered shedding into wastewater, or clinical presentations that reduce the chance of timely clinical detection. The NAO proposal explicitly notes that sophisticated biological weapons would likely be designed to evade pre-existing diagnostics, and argues that metagenomic surveillance is one of the few approaches robust to such adaptation (Nucleic Acid Observatory

Consortium; 2021). Yet even metagenomic systems are constrained by budgets for sampling, sequencing and analysis. Decision-makers must choose how many sites to monitor, how frequently to sample, and how deeply to sequence, under uncertainty about both natural and engineered pathogens.

Building on previous modelling of viral relative abundance in wastewater (Grimm SL, Rushford C, Bley D, Kamm J, Crits-Christoph A, Kantor RS, et al.; 2025), this project develops a sequencing design and planning framework that links pathogen incidence, sequencing depth and detection probability to explicit cost assumptions. I extend RA-based models to handle both known viral classes (for example RNA respiratory, RNA enteric and DNA viruses) and an “unknown” class represented as a mixture over these priors, to approximate the detection problem for novel or AI-enabled agents whose shedding behaviour is uncertain. For each class and incidence target, the framework uses Monte Carlo methods to compute detection probability curves as a function of sequencing depth, identify the minimum depth required to achieve a specified detection goal, and characterise the distribution of required depths across “good”, “typical” and “bad” relative abundance cases. This is coupled to a UK-specific cost model that decomposes per-sample processing and sequencing costs, incorporates empirical or forecast trajectories for sequencing price declines, and scales from per-sample costs to annual per-site and UK-wide budgets for a given sampling frequency and network size.

In doing so, the project provides a decision-focused tool for evaluating how different wastewater metagenomic surveillance designs perform under uncertainty about pathogen properties, including stylised AI-enabled threat scenarios, and what they are likely to cost over the next decade. The framework produces concrete outputs that are directly interpretable by policymakers and technical planners, such as recommended sequencing depths for given detection targets, credible intervals for required depths under relative abundance uncertainty, and projected annual costs for sustaining those depths across a UK-wide sentinel network. It is not intended as a complete biosurveillance blueprint: important dimensions such as detailed sewer-network hydrology, operational logistics and integration with clinical or air-based surveillance are treated only implicitly or left for future work. However, by quantifying the trade-offs between coverage, depth, detection probability and cost in a transparent way, the project offers a practical starting point for designing wastewater metagenomic systems that remain robust in the face of AI-accelerated biological threats.

2. Methods

2.1 Overview

The aim of this work was to quantify how well wastewater metagenomic surveillance could detect an emerging respiratory virus, including “unknown” or deliberately engineered agents, and what sequencing depths and budgets would be required for robust performance.

The analysis proceeds in three steps:

1. **Viral relative abundance priors.** We construct probability distributions for the relative abundance (RA) of human-infecting viruses in wastewater at a fixed incidence of 1% in the upstream population, using published untargeted wastewater metagenomic datasets and the hierarchical model of Grimm et al. (2025).
2. **Detection and design.** Using a binomial detection model with logit-normal overdispersion, we translate RA priors into detection probabilities as a function of sequencing depth and derive “minimum depth” requirements for achieving specified detection targets under both typical and conservative (robust) assumptions.
3. **Costs and AI-enabled scenarios.** We link sequencing depth to financial cost using a simple cost model calibrated to current UK prices and published sequencing cost trajectories, and examine both baseline and “stealth” scenarios representing AI-assisted low-detectability pathogens.

All analyses were performed in Python using reproducible scripts; code and pre-computed outputs are designed so that heavy Monte Carlo calculations only need to be run once and can then be explored interactively.

2.2 Data sources and viral relative abundance priors

We based our viral relative abundance priors on the combined untargeted wastewater metagenomic (W-MGS) datasets used by Grimm et al. (2025), which link observed read counts for human-infecting viruses to independent estimates of infection incidence in the upstream population. In brief, the Grimm model specifies a hierarchical relationship between:

- Population incidence for each virus and study,
- Wastewater relative abundance (fraction of metagenomic reads mapping to that virus), and
- Study- and virus-specific random effects capturing differences in sewershed hydrology, sampling, concentration, extraction and sequencing protocols.

Using the posterior samples provided by Grimm et al. (2025), we constructed empirical priors for the relative abundance at 1% incidence (RA_{1pct}) for three broad virus classes:

- DNA viruses
- RNA enteric viruses
- RNA respiratory viruses

For each virus v in a class, we extracted posterior draws of RA_{1pct} from the Grimm model and pooled across viruses to obtain a class-specific distribution. These distributions capture both biological differences in shedding and environmental behaviour and methodological variability between studies.

To represent an unknown novel pathogen, we built a mixture prior over classes. Specifically, we defined an “unknown” RA_{1pct} prior as a convex combination of the three class-specific priors, with mixture weights chosen to roughly reflect the frequency with which human-infecting viruses from each class appear in the underlying datasets. This yields a broad distribution whose lower tail resembles low-abundance DNA or low-end RNA viruses and whose upper tail resembles relatively abundant enteric viruses.

All RA values are defined as the fraction of total metagenomic reads in a sample that would map to the virus of interest when 1% of the upstream population is infected, before any targeted enrichment.

2.3 Detection model

We modelled detection from a single wastewater sample as a two-stage process:

1. Overdispersed sampling of the viral read fraction.

Even at fixed incidence, the true fraction of reads belonging to a virus varies between samples because of variability in shedding, transport, concentration and library preparation. Following Grimm et al. (2025), we represent this by taking the logit of RA_{1pct} and adding a Normal random effect with class-specific variance. This yields a logit-normal distribution for the sample-specific viral read fraction p .

2. Binomial sampling of reads.

Given a sequencing depth n (total reads per sample) and a realisation of p , the observed virus-specific read count Y is modelled as a *Binomial*(n, p) draw.

A sample is declared detected if the observed count Y is at least a threshold T . In the main analyses we fix:

- Detection threshold $T = 100$ reads (following Grimm (2025))
- Population incidence $I^* = 1\%$ infected
- One wastewater sample per site per week

For a given virus class or RA draw and sequencing depth n , we estimate the single-week detection probability $P_{det, week}(n)$ by Monte Carlo simulation over both the logit-normal dispersion and the binomial sampling. For cumulative surveillance over a horizon of H weeks, with independent samples and constant incidence, we approximate the cumulative detection probability as:

$$P_{det, cum}(H, n) \approx 1 - (1 - P_{det, week}(n))^H.$$

We considered $H \in \{4, 8, 12\}$ weeks, with eight weeks used as the primary horizon.

2.4 Performance metrics and sequencing depth targets

We evaluate surveillance performance using three main metrics:

1. Single-week detection probability $P_{det, week}(n)$ at 1% incidence and depth n .
2. Cumulative detection probability $P_{det, cum}(H, n)$ over an H -week horizon.
3. Minimum sequencing depth required to meet a detection target (“design depth”).

Because the combination of realistic relative abundances and a stringent threshold $T = 100$ makes extremely high weekly detection probabilities difficult to attain, we define our primary design target as 70% single-week detection at $I^* = 1\%$ incidence and $T = 100$ reads.

For a given virus class and depth grid $n \in [10^6, 10^{12}]$ reads per sample, we estimate:

- For known classes (DNA, RNA enteric, RNA respiratory), a single minimum depth n_{min} at which the class-averaged $P_{det, week}(n)$ first reaches or exceeds 0.7; if the curve never reaches 0.7 even at 10^{12} reads, the target is recorded as unattainable on the explored grid.
- For the unknown mixture, a distribution of minimum depths across RA draws. For each sampled RA_{1pct} value we compute the smallest depth $n_{min, i}$ at which the corresponding per-RA detection curve reaches 0.7; if the target is never reached, that draw is marked as unattainable. We then summarise the distribution of attainable $n_{min, i}$ values by its median and 95th percentile, and record the fraction of RA draws for which the target is attainable on the grid.

We interpret the median unknown depth as a “typical” design for an unknown pathogen and the 95th percentile depth as a robust design intended to perform adequately even for low-abundance pathogens in the tail of the RA prior.

2.5 Cost model and UK network scenarios

To relate sequencing depth to financial cost, we use a simple per-sample cost model:

$$\text{cost per sample}(n) = \text{processing cost} + \text{sequencing cost per read} * n,$$

where the processing cost includes collection, concentration, extraction and library preparation and is independent of depth, and the sequencing cost per read reflects platform-specific pricing. We calibrate these quantities using current UK price estimates for short-read sequencing, and we incorporate forecasts of sequencing

cost decline using data from 2015-2022 from Our World In Data (Wetterstrand KA.) to project costs from 2025 to 2035. (Note that this forecasting model was not implemented in the results presented in this paper, see Appendix B6 for more information.)

Assuming one sample per site per week, we convert per-sample costs into:

- Annual cost per site as $52 \times$ cost per sample, and
- Annual network cost by multiplying per-site costs by the number of sentinel sites N_{sites} .

We consider a notional UK network of several hundred wastewater treatment plants (wastewater network citation) and compute annual network costs for both the median and robust design depths for the unknown pathogen class.

To visualise trade-offs, we also compute a cost–detection frontier by evaluating, for a range of depths, the pair:

- single-week detection probability for the unknown mixture, and
- annual cost per site at a reference year.

This frontier shows how much additional cost is required to move from lower to higher levels of detection probability.

2.6 AI-enabled “stealth” pathogen scenario

To approximate an adversarial, AI-enabled pathogen design focused on reducing detectability, we consider a stealth scenario in which the novel virus has a deliberately low relative abundance in wastewater.

Operationally, we define the stealth RA_{1pct} as the 10th percentile of the unknown mixture prior. This corresponds to a pathogen that lies near the low-abundance tail of what we consider plausible for an unknown virus.

We then:

1. Compute single-week and cumulative detection curves for the stealth pathogen using the same binomial + logit-normal model as for other pathogens.
2. Determine the minimum depth required to achieve 70% weekly detection at 1% incidence in the stealth scenario, if attainable on the extended depth grid.

3. Translate this depth into per-site and network costs and compare with the baseline unknown-mixture design depths.

This allows us to quantify how much additional depth and cost would be required to maintain a given detection target if an adversary could push shedding behaviour toward the low-detectability tail of the RA distribution.

2.7 Implementation and reproducibility

All analyses were implemented in Python and are available at: <https://github.com/HGray384/apart-hackathon-2025>. Monte Carlo simulations were run using vectorised code and fixed random seeds for reproducibility.

The workflow is structured in two layers:

1. Heavy pre-computation. Offline scripts load the Grimm posterior samples, construct RA priors, run the detection simulations across large depth grids for each class and RA draw, compute minimum-depth distributions, and evaluate cost trajectories. The outputs are saved as compressed NumPy arrays and CSV tables.
2. Lightweight exploration. A separate interactive demo (Streamlit application) and plotting scripts read the pre-computed arrays and generate figures and tables without additional Monte Carlo sampling.


This separation ensures that the main results and figures can be regenerated deterministically from the raw data and configuration files, while interactive use is fast and does not require access to high-performance computing.

3. Results

3. Results

3.1 Viral relative abundance priors for known and unknown pathogen classes

Posterior samples of viral RA_{1pct} show substantial heterogeneity both within and across viral classes (Figure 1, Table 1). For each class (DNA, RNA enteric, RNA respiratory), RA_{1pct} spans roughly 4.5–6 orders of magnitude on the $\log_{10}$ scale between the 5th and 95th percentiles. Specifically, DNA viruses have RA_{1pct} in the range 4.91×10^{-14} – 1.49×10^{-9} (5th–95th percentile), RNA enteric viruses 4.54×10^{-9} – 3.15×10^{-4} , and RNA respiratory viruses 5.69×10^{-12} – 6.50×10^{-6} (Table 1). Median RA_{1pct} values differ by about five orders of magnitude across classes: 5.93×10^{-12} for DNA, 1.22×10^{-6} for RNA enteric, and 1.49×10^{-9} for RNA respiratory viruses.




Class	Median RA_{1pct} ($\log_{10}$)	P5 RA_{1pct} ($\log_{10}$)	P95 RA_{1pct} ($\log_{10}$)
DNA	5.93e-12 (-11.23)	4.91e-14 (-13.31)	1.49e-09 (-8.83)
RNA enteric	1.22e-06 (-5.91)	4.54e-09 (-8.34)	0.000315 (-3.50)
RNA respiratory	1.49e-09 (-8.83)	5.69e-12 (-11.24)	6.5e-06 (-5.19)
Unknown	1.54e-09 (-8.81)	2.2e-13 (-12.66)	0.0002 (-3.70)

RNA enteric viruses are typically most abundant in wastewater, with a median RA_{1pct} of 1.22×10^{-6} (5th–95th percentile 4.54×10^{-9} – 3.15×10^{-4}), substantially higher than RNA respiratory viruses (median 1.49×10^{-9} ; 5th–95th percentile 5.69×10^{-12} – 6.50×10^{-6}) and DNA viruses (median 5.93×10^{-12} ; 5th–95th percentile 4.91×10^{-14} – 1.49×10^{-9}). DNA viruses also exhibit a particularly heavy lower tail, with 5th percentile RA_{1pct} values near 5×10^{-14} , indicating that some DNA pathogens are systematically harder to detect at fixed sequencing depth than typical RNA enteric or respiratory viruses.

To approximate the detection problem for novel agents, an “unknown pathogen” prior was constructed as an equal-weight mixture over the three class-specific RA_{1pct} priors (DNA, RNA enteric, RNA respiratory). The resulting RA_{1pct} distribution (Figure 2, Table 1) is broader than any single class and strongly skewed toward lower values. The median RA_{1pct} for the unknown mixture is 1.54×10^{-9} , with a 5th–95th percentile range of 2.20×10^{-13} – 2.00×10^{-4} , spanning roughly nine orders of magnitude. The lower tail of the mixture is dominated by low-abundance DNA-like and low-end RNA realisations, while the upper tail reflects more favourable enteric-like shedding. This mixture prior captures both uncertainty


about which viral class a novel agent resembles and the wide variation in environmental abundance.



3.2 Detection probability as a function of sequencing depth

3.2.1 Single-week detection probabilities for known classes


Weekly detection probability for a single composite wastewater sample was evaluated as a function of total sequencing depth for each viral class, assuming 1% incidence in the catchment and $T = 100$ (Figure 3). Depths range from 10^6 to 10^{12} reads per sample. For each depth, we compute the single-week detection probability $P_{det, week}(n)$ by averaging over class-specific RA_{1pct} draws and the logit-normal dispersion.



For known classes, single-week detection probability increases with depth but saturates at very different scales. RNA enteric viruses achieve high detection probabilities at comparatively modest depths: the transition from near-zero to near-certain detection occurs over roughly two orders of magnitude in depth, and probabilities above 70% are attainable at $n \approx 10^9$ reads per sample. RNA respiratory viruses require much higher depths, with 70% detection only appearing close to the upper end of the grid (around 10^{12} reads). DNA viruses are substantially more challenging: even at 10^{12} reads per sample, the weekly detection probability remains below the 70% target under the sampled RA_{1pct} distribution, so no finite minimum depth is found on the grid for DNA alone.

3.2.2 Cumulative detection over multiple weeks for known classes


Because sampling is weekly, cumulative detection over multiple weeks is also relevant. For each class and depth, we compute the single-week detection probability $P_{det, week}(n)$, then derive cumulative detection over H weeks for horizons $H \in \{4, 8, 12\}$ (Figure 4).



For RNA enteric viruses, once single-week detection exceeds about 30–40% (at depths near 10^8 – 10^9 reads), cumulative detection over 8–12 weeks rapidly approaches 1. RNA respiratory viruses benefit from multi-week compounding but still require depths around 10^{11} – 10^{12} reads for cumulative detection to reliably exceed 90%. DNA viruses show negligible cumulative detection across the considered horizon for depths up to 10^{12} reads, reflecting the very low typical RA_{1pct} values. These results highlight that extended operation can compensate partially for sub-maximal weekly sensitivity, but cannot rescue designs that sit too far into the low-abundance regime for a given class.

3.2.3 Unknown pathogen detection under RA uncertainty


For the unknown-pathogen mixture, we propagate RA uncertainty. For each depth in the 10^6 – 10^{12} grid and each RA draw from the mixture prior, we compute a per-RA weekly detection probability, then aggregate across RA draws to obtain the median and 5th–95th percentile range at each depth (Figure 5).



The median single-week detection curve for the unknown class remains essentially at zero up to around 10^{10} reads, and then transitions sharply to high detection between roughly 10^{10} and a few $\times 10^{11}$ reads. Under the mixture prior, 70% weekly detection at 1% incidence is only achieved at depths of order 10^{11} reads if one designs for the “median RA” behaviour.

RA uncertainty translates into very wide credible bands. For much of the depth range, the 5th percentile of $P_{det, week}(n)$ is close to zero while the 95th percentile is close to one, so the 5–95% interval covers most of the $[0,1]$ range even though the median remains near zero. This reflects a strongly bimodal situation: for a subset of RA realisations (high-abundance, enteric-like), detection becomes almost trivial at moderate depths, whereas for others (low-abundance, DNA-like) detection remains extremely unlikely even near 10^{12} reads.

Cumulative detection over $H = 8$ weeks for the unknown mixture is computed separately for each RA realisation and then aggregating across RA draws (Figure 6). The median cumulative detection curve shows a similar sharp transition, with negligible detection up to $\sim 10^{10}$ reads and rapid convergence to near-certain detection once weekly probabilities cross into the 30–50% range. Again, the 5–95% band is wide over much of the depth range, indicating that even with multi-week sampling, some plausible RA realisations remain essentially invisible while others are detected almost immediately.



3.3 Minimum sequencing depth required for specified detection targets

3.3.1 Class-specific minimum depths (median-case)

To make the design problem more concrete, we invert the detection curves to estimate the minimum sequencing depth required to meet a 70% single-week detection target at 1% incidence with $T = 100$. For each class, we search over a refined depth grid up to 10^{12} reads and return the smallest depth where the target is reached.


The resulting minimum depths (Table 3) are highly class-dependent. For RNA enteric viruses, the median minimum depth is 1.70×10^9 reads per sample ($\log_{10} n \approx 9.23$). For RNA respiratory viruses, the corresponding value is 7.02×10^{11} reads ($\log_{10} n \approx 11.85$). For DNA viruses, the 70% weekly detection target is not reached anywhere on the grid up to 10^{12} reads, consistent with the extremely low RA_{1pct} values for many DNA pathogens. These class-specific results confirm that a single global sequencing depth cannot guarantee uniform performance across viral families.

Class	$n_{min70\%}$
DNA	N/A (not reached up to 10^{12} reads)
RNA enteric	1.70×10^9
RNA respiratory	7.02×10^{11}

3.3.2 Robust depth under RA uncertainty for the unknown class

For the unknown mixture, we examine minimum depths under both median and robust criteria. Using the same refined grid, we compute for each RA draw the minimum depth at which the 70% weekly detection target is achieved, then summarise the resulting distribution.

The median design depth for the unknown class, defined as the 50th percentile of the $n_{min70\%}$ distribution, is 1.93×10^9 reads ($\log_{10} n \approx 9.29$). The robust design depth, defined as the 95th percentile, is 6.15×10^{11} reads ($\log_{10} n \approx 11.79$) (Figure 7). Overall, 63.0% of RA draws reach the 70% detection target somewhere within the 10^6 – 10^{12} grid, with $n_{min70\%}$ for those draws spanning roughly from 10^6 to 10^{12} reads. The ratio between robust and median depths is $\approx 3.2 \times 10^2$, quantifying the additional sequencing required to protect against most (95%) RA scenarios rather than a typical (median) one.



3.4 Cost of achieving robust detection in a UK wastewater system

(Note: there were some discrepancies in the cost model which were noticed prior to submission. The actual results in this section may be inaccurate.)

3.4.1 Per-sample and per-site costs at different depths


To link sequencing depth to financial constraints, we use a cost model calibrated to mid-range UK prices in 2025, including both library preparation and sequencing. For each candidate depth, we compute:

- a per-sample cost, and
an annual per-site cost for one sample per week.

Using the median and robust design depths for the unknown class, the per-sample cost at the median design depth $n = 1.93 \times 10^9$ reads is about £1,020, corresponding to an annual per-site cost of £53,053. At the robust design depth $n = 6.15 \times 10^{11}$ reads, the per-sample cost rises to £276,944, and the annual per-site cost to £14.4 million. Moving from the median to the robust depth therefore increases per-sample and per-site costs by a factor of ≈ 270 , reflecting both the higher depth and the cost structure of ultra-deep sequencing.

3.4.2 Scaling to a UK-wide network and future cost trajectories


We then scale these per-site costs to a notional UK-wide network of 200 sentinel wastewater treatment plants, each sampled weekly. For each year from 2025 to 2035, we adjust sequencing costs using an external forecast and recompute annual network costs for both the median and robust depths (Figure 8).



At 2025 prices, the total annual network cost is about £10.6 million for the median design and £2.88 billion for the robust design. By 2030, these values fall to approximately £5.6 million and £1.28 billion, respectively; by 2035, to about £4.6 million and £0.96 billion. The additional annual expenditure required to move from the median to robust depth decreases from roughly £2.87 billion in 2025 to £0.96 billion in 2035. Although still substantial, the absolute premium for robustness shrinks markedly over time, making conservative designs progressively more affordable if sequencing costs continue to fall.

3.4.3 Cost–detection trade-offs

Finally, we trace the cost–detection frontier for the unknown class by varying depth across the range that connects the median and robust designs (Figure 9). For each depth, we compute the median weekly detection probability for the unknown mixture and the corresponding annual per-site cost at 2025 prices.



The resulting curve shows modest increases in depth from the low end of the range yield large gains in detection probability at relatively low cost, but beyond about 10^{11} reads further increases in depth produce sharply diminishing returns. The median design depth (1.93×10^9 reads; $\approx \text{£}53\text{k}$ per site per year) sits near the “knee” of this curve, while the robust design depth (6.15×10^{11} reads; $\approx \text{£}14.4\text{m}$ per site per year) lies far along the frontier, representing a deliberate choice to buy protection against low-RA scenarios at dramatically higher cost.


3.5 Stealth pathogen scenario

3.5.1 Definition and detection curves

To approximate an AI-enabled “stealth” pathogen deliberately engineered to be hard to detect in wastewater, the script defines a stealth scenario by taking the 10th percentile of the unknown-mixture RA_{1pct} distribution, placing the pathogen near the low-abundance tail of the prior. We then compute single-week detection curves for:


1. the unknown-mixture median curve, and
2. the stealth scenario,

over the 10^6 – 10^{12} depth grid (Figure 10).



The unknown-median curve again shows a steep transition from near-zero to near-certain detection over roughly one order of magnitude in depth, but shifted to higher depths than the mixture-based design analysis (Section 3.3.2), since it conditions on a single, moderately low RA. By contrast, the stealth curve remains essentially flat at zero across the entire range up to 10^{12} reads: even ultra-deep sequencing fails to reach meaningful weekly detection probabilities for such a low-abundance pathogen.

Cumulative detection over 8 weeks ($H = 8$) is then computed from these weekly probabilities (Figure 11). At depths where the unknown-median scenario achieves high cumulative detection (close to 1), the stealth scenario still remains near zero. This illustrates that even multi-week sampling struggles when RA is deliberately pushed into the extreme low tail.



3.5.2 Impact on required depth and cost

We re-run the minimum-depth logic for the stealth scenario using the same 70% weekly detection target. In this case, the target is never reached anywhere on the grid up to 10^{12} reads (Table 7).

Scenario	$n_{min70\%}$	Cost per sample (2025)	Cost per site (2025)
Unknown median	1.93×10^9	£1,020	£53k
Unknown robust	6.15×10^{11}	£276,944	£14.4m
Stealth	N/A (not reached up to 10^{12} reads)	N/A	N/A

The median and robust designs for the unknown mixture require 1.93×10^9 and 6.15×10^{11} reads respectively, with 2025 per-sample costs of approximately £1,020 and £276,944, and annual per-site costs of about £53k and £14.4m (Table 7). Designing to the 95th percentile of the n_{min} distribution therefore already pushes sequencing very close to the practical limits considered here; pushing further into the stealth regime would demand depths (and costs) beyond 10^{12} reads per sample and is likely infeasible with current technology.

3.6 Summary of key findings

Overall, the results indicate that:

- Wastewater viral abundance at 1% incidence is highly heterogeneous across and within viral classes, spanning 4–6 orders of magnitude per class and roughly 9 orders of magnitude for the unknown mixture. An “unknown pathogen” is therefore best modelled by a broad mixture prior rather than a single typical value.
- Achieving robust early detection for unknown agents requires sequencing depths substantially above those suggested by median-case analyses, especially when targeting high weekly detection probabilities. For the unknown mixture, moving from median-case to 95th-percentile robustness multiplies required depth by $\approx 3 \times 10^2$.
- The additional cost of robustness is very large at current sequencing prices (billions of pounds annually for a 200-site UK network) but decreases materially under plausible future cost trajectories, making conservative designs progressively more affordable.
- Simple stealth-like modifications to viral behaviour (approximated by choosing a low-percentile RA from the mixture prior) can render metagenomic wastewater surveillance effectively blind within realistic depth ranges. Designs tuned only to median abundance can therefore fail badly in adversarial scenarios, whereas robust-depth designs substantially reduce but do not fully eliminate this risk.

Taken together, these findings suggest that metagenomic wastewater surveillance can meaningfully contribute to defence against AI-enabled viral threats, provided that sequencing depths and budgets are chosen with explicit attention to RA uncertainty. However, highly adversarial scenarios do not appear to be appropriately covered by this protocol and so require further analysis.

4. Discussion and Conclusion

This work quantifies how far metagenomic wastewater surveillance can be pushed as a defence against emerging and potentially AI-enabled respiratory viruses. Using posterior draws from Grimm et al. (2025) and a simple epidemic/surveillance model, we show that viral relative abundance at 1 % incidence spans many orders of magnitude even within a single class. As a result, the same sequencing depth can reliably detect high-abundance enteric viruses while providing only weak protection against typical respiratory viruses and almost none against low-abundance DNA-like or deliberately “stealthy” threats. The unknown-pathogen mixture prior makes this heterogeneity explicit: at any fixed depth the weekly detection probability varies from essentially zero to near one across plausible RA_{1pct} draws.

The detection and cost curves imply that there is no cheap route to robust wastewater early warning. For a UK-scale network, achieving 90 % weekly detection against a median unknown pathogen requires tens of millions of reads per sample but remains vulnerable to low-abundance designs; pushing to depths that cover 90–95 % of the RA prior multiplies sequencing costs by an order of

magnitude while still leaving residual blind spots. Multi-week sampling substantially improves cumulative detection once weekly sensitivity passes ~20–30 %, but cannot compensate for pathogens whose wastewater signal remains below the detection threshold even at very high depths.

These results should not be interpreted as a verdict against wastewater surveillance. Rather, they underscore that metagenomic WBE is a powerful but limited component of a broader defensive stack. Its strengths are breadth of coverage and the ability to detect “unknown unknowns” without target-specific assays; its weaknesses are the dependence on poorly understood shedding and transport processes, the cost of very deep sequencing, and a structural vulnerability to low-abundance or “off-sewer” attack designs. The modelling here abstracts away from many real-world complexities, including detailed sewer network structure, joint optimisation with clinical and environmental sensors, and adaptive adversaries using AI systems to search for subtle evasion strategies. These are important directions for future work, but they are unlikely to reverse the basic cost–robustness trade-offs we identify.

This analysis suggests that metagenomic wastewater surveillance can materially reduce the time to detection for many high-abundance respiratory threats, including some plausible AI-enabled pathogens, provided that sequencing depths and network scales are chosen with explicit attention to variability in wastewater viral abundance. At the same time, the large gap between the depths needed for “typical” and “robust” detection, and the persistence of blind spots for low-abundance or non-sewered releases, mean that wastewater data alone cannot be relied upon as a universal early-warning system. Effective biosecurity policy will require combining such surveillance with targeted clinical testing, environmental sensing, and analytics tuned to adversarial threat models, alongside governance that constrains the most dangerous uses of AI in biological design. This work offers a transparent, reproducible framework for stress-testing candidate wastewater networks against those emerging risks and for making the associated trade-offs explicit.

5. References

- Nucleic Acid Observatory Consortium. A global nucleic acid observatory for biodefense and planetary health. arXiv preprint. 2021. arXiv:2108.02678.
- Mouton CA, Dunham JL, Baker GP, et al. The operational risks of AI in large-scale biological attacks. Santa Monica (CA): RAND Corporation; 2024. (Report RRA2977-1).
- Parkins MD, Lee BE, Acosta N, Bautista M, Hubert CRJ, Hruddy SE, et al. Wastewater-based surveillance as a tool for public health action: SARS-CoV-2 and beyond. *Clin Microbiol Rev.* 2024;37(2):e00103-22. doi:10.1128/cmr.00103-22.
- Jourdain F, Bouzid D, Blanchon T, Roque-Afonso A-M, Bonifait L, Schvoerer E, et

- al. Wastewater-based epidemiological surveillance in France: the SUM'EAU network. *Microorganisms*. 2025;13(2):281. doi:10.3390/microorganisms13020281.
- Singh S, Goyal R, Chander Y, Gupta N. Wastewater surveillance for SARS-CoV-2 and other pathogens: current practice, challenges and future opportunities. *Front Public Health*. 2024;12:1445961. doi:10.3389/fpubh.2024.1445961.
- Gopal A, Greenwald HD, Crits-Christoph A, Kantor RS, Krammer F, Marras SAE, et al. Virus-like DNA tracers for characterizing wastewater-based pathogen monitoring systems. *bioRxiv*. 2023. doi:10.1101/2023.12.24.573197.
- Gurung BDS, Gweon HS, Jabbari S, Shadouh S, Sails AD, Singer AC, et al. Forecasting urban wastewater microbiome dynamics using a digital twin framework. *medRxiv*. 2025. doi:10.1101/2025.07.21.666059.
- Grimm SL, Rushford C, Bley D, Kamm J, Crits-Christoph A, Kantor RS, et al. Inferring the sensitivity of wastewater metagenomic sequencing for early detection of viruses: a statistical modelling study. *Lancet Microbe*. 2025;6(11):eXXX–eXXX. (In press; preprint version: 2025.10.27.25338874).
- St-Onge C, Delarocque-Astagneau E, Schofield H, Zeller H, Van der Werf S, Obadia T, et al. Global aircraft wastewater genomic surveillance for early detection of respiratory pathogens. *EBioMedicine*. 2025;103:105040. doi:10.1016/j.ebiom.2024.105040.
- Wetterstrand KA. DNA Sequencing Costs: Data from the NHGRI Genome Sequencing Program (GSP) Available at: www.genome.gov/sequencingcostsdata. Accessed [2025-11-23].

6. Appendix

Appendix A: Security considerations

A1. Scope and intended use

This project is intended as a **defensive** contribution to biosecurity.

It asks:

- How sensitive could a realistically funded wastewater metagenomic network be to emerging respiratory viruses, including unknown or engineered agents?

- What sequencing depths and budgets would be required to reach specified detection targets?
- How does performance degrade for low-abundance “stealth” pathogens that might plausibly be designed with the help of AI tools?

The work focuses on surveillance system design, not on biological agent design.

All modelling operates at the level of relative abundance distributions and detection probabilities, not viral genomes, constructs, or protocols.

A2. Biosecurity benefits

From a security standpoint, the main value of the work is to:

1. Quantify detection blind spots.
 - The analysis identifies regions of parameter space (e.g. low wastewater relative abundance, low shedding into sewage) where even very deep sequencing performs poorly.
 - This helps planners understand which threat models are realistically covered by wastewater metagenomics and which are not.
2. Clarify trade-offs between robustness and cost.
 - The results show how rapidly cost escalates as one moves from “typical unknown virus” performance to “robust against low-abundance variants.”
 - This supports more realistic budgeting and prioritisation of surveillance investments, rather than assuming “more sequencing solves everything.”
3. Stress-test AI-enabled threat scenarios.
 - By embedding an explicit “stealth” scenario, the project gives a concrete way to discuss how AI-assisted design might erode detection and how much additional resource would be needed to claw back performance.
4. Support high-level policy questions.
 - The framework helps answer questions like:
 - *Is metagenomic wastewater surveillance alone a sufficient backbone for early warning?*

- *Under what conditions is it cost-effective compared with other modalities (clinical, air sampling, etc.)?*

The intent is to inform defensive strategy, not to prescribe specific operational parameters for any particular national system.

A3. Limitations and uncertainties

The work has several important limitations that matter for security interpretation:

1. Coarse threat model.
 - Pathogens are represented via relative abundance at 1% incidence and simple shedding into wastewater; within-host dynamics, environmental decay pathways, and behavioural heterogeneity are not explicitly modelled.
 - Sophisticated adversaries might exploit features not captured here (e.g. clustered shedding, time-varying shedding, targeted seeding patterns).
2. Simplified network structure.
 - The model treats wastewater sites as independent sentinel catchments, without detailed sewer network topology, mixing, or travel patterns.
 - Spatial seeding strategies, targeted attacks on unmonitored regions, or exploitation of atypical sewersheds (e.g. small, industrial, or highly transient populations) are therefore not fully addressed.
3. Uncertain RA priors.
 - RA_{1pct} distributions are derived from a limited set of historical datasets and laboratory protocols.
 - They may not generalise to novel technologies, different geographies, or radically different virus families.
 - In particular, the “unknown” mixture prior is an informed but subjective construct.
4. Simplified detection rule.
 - A fixed threshold of 100 reads is used as a proxy for “detectable and credible to analysts.”

- Real systems use more complex pipelines (alignment, filtering, background models, anomaly detection).
 - The estimates here are therefore conservative and approximate, not a substitute for end-to-end pipeline evaluation.
5. Cost model uncertainty.
- Costs are based on current UK-style workflows and forecasting assumptions about sequencing price declines.
 - Operational costs, staffing, QC, and information-system expenses are simplified or omitted; real programmes may be substantially more expensive.
6. No adversarial game dynamics.
- The work evaluates static surveillance designs against static attacker assumptions.
 - It does not model iterated attacker–defender adaptation, nor coordinated multi-modal detection (wastewater + clinical + environmental).

These limitations should be kept in mind when interpreting the results: the model is best seen as an intuitive proof of concept for a more complex design.

A4. Dual-use and misuse potential

Although this work is defensive in intent, elements of it are technically dual-use:

1. Quantifying detection thresholds.
 - The model clarifies what combinations of relative abundance and sequencing depth lead to low detection probabilities.
 - In principle, an adversary could use this to reason about how far they might need to push a pathogen’s wastewater shedding downward to evade detection.
2. Stealth scenario framing.
 - The explicit “stealth” scenario (10th percentile RA of the unknown prior) highlights that low-abundance designs are more likely to slip through.

- If combined with separate, more detailed design tools, this could inform adversarial optimisation.
3. Cost–robustness trade-offs.
 - Showing how expensive robust coverage is could be misused rhetorically to argue against investment in wastewater surveillance entirely, rather than to motivate complementary defences.
 4. AI linkage.
 - The project contextually discusses AI tools lowering barriers to designing novel pathogens or planning release strategies.
 - While no concrete design or protocol information is provided, any discussion of AI-enabled threat models risks being cited in offensive framing.

At the same time, the actual model details are quite far from something a plausible attacker would find practically useful:

- There is no genome-level information, no optimisation over specific constructs, and no lab protocol content.
- The RA priors and detection curves are highly aggregated and uncertain; an attacker attempting to fine-tune a real design from these numbers would likely be misled.
- Truly capable adversaries would almost certainly prefer their own experimental measurements or bespoke simulations.

Although the dual-use angle is real, the findings here provide information which support unsurprising intuitions: rare viruses are hard to detect, and more robust detection incurs more costs.

A5. Future improvements and research directions

Several future directions could improve the defensive usefulness of this work while staying within safe bounds.

1. Richer multi-modal surveillance models.
 - Extend the framework to jointly consider indoor air sampling, clinical diagnostics, and possibly syndromic data alongside wastewater.
 - This would better reflect realistic detection architectures and reduce the temptation to over-interpret wastewater as a single point of

failure.

2. More realistic sewershed and network structure.
 - Incorporate explicit sewer network topology, travel patterns, and realistic catchment sizes to identify which nodes are most critical as sentinels.
3. Adversarially robust design explorations.
 - Explore how much robustness can be gained from combining metagenomics with algorithmic anomaly detection, targeted enrichment, or adaptive sampling strategies.
4. More granular uncertainty treatment.
 - Decompose uncertainty by source (biological, environmental, methodological) to see which components most undermine robustness.
 - This can guide where additional measurement or standardisation would have the biggest defensive payoff.
5. Governance and policy integration.
 - Link design outputs to concrete governance questions:
 - At what cost thresholds should international or national funding mechanisms trigger?

Overall, this project should be viewed as a defensive decision-support tool: it quantifies what wastewater metagenomic surveillance can and cannot do in the face of emerging and potentially AI-enabled viral threats. Used appropriately, it can help steer investments toward more robust early-warning architectures. Misused in isolation, it offers at most a very rough sketch of surveillance blind spots, and it is intentionally designed not to provide actionable offensive guidance.

Appendix B: Mathematical details

The main text deliberately avoids heavy mathematical notation. Here we summarise the formal models used.

B.1 Hierarchical model for wastewater relative abundance

We use posterior samples from the hierarchical model of Grimm et al. (2025) to obtain draws from the distribution of RA_{1pct} for each virus. Here we give a simplified description of the model structure.

For a given study s , virus v and sample j , Grimm et al. (2025) model the $\log_{10}$ relative abundance as

$$\log_{10}(RA_{s,v,j}) = \alpha_s + \beta_s * \log_{10}(Inc_{s,v,j}) + u_v + e_{s,v,j}$$

where:

- $Inc_{s,v,j}$ is an external estimate of the incidence (proportion of the upstream population infected with virus v in study s at time j).
- α_s is a study-specific intercept capturing baseline differences in protocols and sewershed characteristics.
- β_s is a study-specific slope linking incidence to relative abundance.
- u_v is a virus-specific random effect capturing systematic differences in shedding and environmental behaviour between viruses.
- $e_{s,v,j}$ is a residual term.

Priors are specified on α_s , β_s , u_v , and the model is fit jointly across all studies and viruses using Bayesian methods. The exact prior choices and fitting details follow Grimm et al. (2025) and are not modified here.

For each virus and study, the model implies a posterior predictive distribution for the relative abundance at a fixed reference incidence $I^* = 1\%$. Setting $Inc_{s,v,j} = I^*$ and sampling from the posterior gives us draws of RA_{1pct} for each virus, study, and sample.

To construct a virus-level distribution of RA_{1pct} , we aggregate over studies and samples for each virus, treating all posterior predictive samples for that virus at 1% incidence as draws from the virus's RA_{1pct} distribution.

To obtain class-level priors, we pool the virus-level draws across viruses belonging to the same class c .

From these pooled draws we can estimate empirical densities, medians, and quantiles for each class, which are used both in the RA plots and in the detection simulations.

B.2 Construction of the unknown-pathogen mixture prior

We represent an “unknown” novel pathogen as a mixture over the three virus classes. Let:

- $p_c(RA_{1pct})$ denote the empirical distribution of RA_{1pct} for class c , derived from the pooled posterior samples above;
- w_c be non-negative weights for each class that sum to 1.

The unknown-pathogen prior is then defined as

$$p_{unknown}(RA_{1pct}) = \text{sum over classes } c \text{ of } [w_c * p_c(RA_{1pct})].$$

Operationally, we implement this by:

1. Drawing a class label C from the categorical distribution with probabilities w_c .
2. Drawing RA_{1pct} from the empirical class distribution corresponding to C .

In code this is done by concatenating the class-specific samples with appropriate resampling, but mathematically it is equivalent to the mixture above.

The weights w_c are chosen to reflect the relative representation of DNA, RNA enteric, and RNA respiratory viruses in the underlying datasets, and can be specified explicitly for different weighting schemes. Alternative weighting schemes are straightforward to implement and would shift the unknown prior toward the corresponding class distributions. Our main conclusions are driven more by the overall span of the mixture (from very low to relatively high RA_{1pct}) than by small changes in w_c .

B.3 Binomial + logit-normal detection model

Given a virus with relative abundance RA_{1pct} , we model detection in a single wastewater sample at incidence $I^* = 1\%$ as a two-stage process.

1. Overdispersed sampling of the viral fraction p .

Even at fixed incidence, the fraction of reads belonging to a virus varies between samples because of variability in shedding, transport, concentration and library preparation. We model the logit of the sample-specific viral fraction p as

$$\text{logit}(p) = \text{logit}(RA_{1pct}) + \epsilon$$

$$\epsilon \sim \text{Normal}(0, \sigma_c^2)$$

Where $\text{logit}(x) = \log(x / (1 - x))$, σ_c is a dispersion parameter that depends on virus class, and epsilon is a random effect capturing sample-to-sample variation.

2. Binomial sampling of reads.

Conditional on the sample-specific fraction p and sequencing depth n , we model the observed read count Y as

$$Y \mid p, n \sim \text{Binomial}(n, p).$$

We declare the virus detected in that sample if the observed count meets or exceeds a threshold T . In this work, we fix $T = 100$ reads.

3. Detection probability as an integral.

The single-sample detection probability at depth n for a virus with RA_{1pct} is

$$\begin{aligned} P_{det, week}(n \mid RA_{1pct}) &= P(n \geq T) \\ &= E_{\epsilon} [P(n \geq T \mid \epsilon)]. \end{aligned}$$

Expanding the conditional probability:

$$P(n \geq T \mid \epsilon) = \sum_{k=T}^n \text{BinomialPMF}(k; n, p(\epsilon))$$

where

$$\begin{aligned} p(\epsilon) &= \text{logistic}(\text{logit}(RA_{1pct}) + \epsilon) \\ &= 1 / (1 + \exp(-[\text{logit}(RA_{1pct}) + \epsilon])). \end{aligned}$$

So,

$$P_{det, week}(n \mid RA_{1pct}) = \text{integral over } \epsilon [P(Y \geq T \mid \epsilon) * \text{NormalPDF}(\epsilon; 0, \sigma_c^2) d\epsilon].$$

There is no convenient closed-form expression for this integral, so we approximate it using Monte Carlo simulation.

B.4 Estimation of detection probabilities and cumulative detection

B.4.1 Monte Carlo estimation for a given RA and class

For a fixed RA_{1pct} , virus class c , and depth n , we approximate $P_{det, week}(n \mid RA_{1pct})$ by Monte Carlo over the logit-normal random effect:

1. Choose a number of dispersion draws B (e.g. $B = 1,000$).
2. For each $b = 1, \dots, B$:

$$\circ \text{ Sample } \epsilon_b \sim \text{Normal}(0, \sigma_c^2).$$

- Compute $p_b = \text{logistic}(\text{logit}(RA_{1pct}) + \epsilon_b)$.
- Compute the binomial tail probability analytically
- Estimate the weekly detection probability by averaging over B

This estimator is unbiased and its Monte Carlo standard error can be made small by choosing B sufficiently large. In practice, we verify convergence by doubling B in pilot runs and checking that results change negligibly.

We repeat this procedure for each depth n on a logarithmic grid (e.g. 10^6 to 10^{12} reads).

B.4.2 Aggregation across RA draws for the unknown class

For the unknown mixture prior, we start with a set of M draws:

$$RA_{1pct}^{(unknown, i)} \text{ for } i = 1, \dots, M.$$

For each draw i and each depth n_j on the grid, we compute the Monte Carlo estimate:

$$P_{det, week}^{\sim}(i, j) = P_{det, week}^{\sim}(n_j | RA_{1pct}^{(unknown, i)})$$

This yields a matrix of size $M \times K$ where K is the number of depths. From this matrix we compute, for each n_j :

$$median_p(j) = \text{median over } i \text{ of } P_{det, week}^{\sim}(i, j)$$

$$p5(j) = \text{5th percentile over } i \text{ of } P_{det, week}^{\sim}(i, j)$$

$$p95(j) = \text{95th percentile over } i \text{ of } P_{det, week}^{\sim}(i, j)$$

These curves give the median and 5–95% quantile band of weekly detection probability across the unknown RA prior.

B.4.3 Cumulative detection across H weeks

Assuming:

- incidence remains at $I^* = 1\%$ for the horizon of interest, and
- weekly samples are independent conditional on the parameters,

the cumulative detection probability over H weeks at depth n is:

$$P_{det, cum}(H, n) = 1 - (1 - P_{det, week}(n))^H$$

This holds because the probability of not detecting the virus in any of H independent weeks is $(1 - P_{det, week}(n))^H$.

For the unknown mixture, we apply this transformation to each of the summary curves:

$$median_{cum}(j) = 1 - (1 - median_p(j))^H$$

$$p5_{cum}(j) = 1 - (1 - p5(j))^H$$

$$p95_{cum}(j) = 1 - (1 - p95(j))^H$$

We mainly use $H = 8$ weeks.

B.5 Minimum depth distributions and robust design criteria

We define the weekly detection target as

$$\gamma_{week} = 0.7$$

corresponding to 70% single-week detection at $I^* = 1\%$ and threshold $T = 100$.

B.5.1 Minimum depth for known classes

For each known class c (DNA, RNA_enteric, RNA_respiratory), we have its class-averaged weekly detection curve:

$$P_{det, week, c}(n_j)$$

over the depth grid n_j . We define the minimum depth for that class as:

$$n_{min, c} = \text{smallest } n_j \text{ on the grid such that } P_{det, week, c}(n_j) \geq \gamma_{week}$$

If the curve never reaches γ_{week} even at the largest depth (e.g. DNA viruses), we record $n_{min, c}$ as “not reached on grid”.

B.5.2 Minimum depth distribution for the unknown class

For the unknown mixture prior, we compute a per-RA minimum depth for each draw i :

$$n_{min, i} = \text{smallest } n_j \text{ on the grid such that } P_{det, week}^{\sim}(i, j) \geq \gamma_{week}$$

If $P_{det, week}^{\sim}(i, j) < \gamma_{week}$ for all depths up to the maximum (10^{12} reads), we set:

$n_{min,i} = \text{NaN}$ (target not reachable for this RA draw).

From the set of $n_{min,i}$ values we then compute:

- frac_reached = fraction of RA draws for which $n_{min,i}$ is finite (target reachable on the grid).
- $n_{min, median}$ = median of $n_{min,i}$ over finite values.
- $n_{min, 95th}$ = 95th percentile of $n_{min,i}$ over finite values.

We interpret:

- $n_{min, median}$ as the typical design depth for an unknown pathogen, and
- $n_{min, 95th}$ as a robust design depth intended to cover most (95%) of the RA prior.

These depths are the inputs to the cost calculations and the “median vs robust” design comparisons in the main text.

B.6 Cost model and forecasting

NOTE: THIS APPENDIX CONTAINS THE CORRECT VALUES USED IN THIS PAPER. VALUES FOR 2025 ARE AS INTENDED, BUT FORECASTED VALUES ARE INACCURATE. A BETTER FORECASTING MODEL WAS DEVELOPED BUT WAS NOT IMPLEMENTED IN THIS PAPER BY MISTAKE. THE REPO CODE CONTAINS THE CORRECT MODEL IN A DEMO SCRIPT.

We use a simple linear cost model to relate sequencing depth to financial cost.

B.6.1 Per-sample, per-site, and network costs

For a given year y and depth n , the per-sample cost is:

$$\text{cost_per_sample_y}(n) = \text{processing_cost} + c_read_y * n$$

where:

- processing_cost is a fixed amount covering collection, concentration, extraction, and library preparation, taken as £150 per sample
- c_read_y is the sequencing cost per read in year y (GBP per read).

The annual cost per site with weekly sampling is:

$$\text{cost_annual_per_site_y}(n) = 52 * \text{cost_per_sample_y}(n)$$

Given a network of N_sites sentinel sites, the annual network cost is:

$$\text{cost_annual_network_y}(n) = N_{\text{sites}} * \text{cost_annual_per_site_y}(n)$$

In all network-level figures we use $N_{\text{sites}} = 200$ as a notional UK-wide sentinel network.

B.6.2 Sequencing cost trajectories

We assume that processing costs remain constant over time at £150 per sample, and that the per-read sequencing cost declines over time following approximate forecasts from sequencing cost data.

Specifically, we set:

$$c_{\text{read_2025}} = 4.5 \times 10^{-7} \text{ GBP per read}$$

$$c_{\text{read_2030}} = 2.0 \times 10^{-7} \text{ GBP per read}$$

$$c_{\text{read_2035}} = 1.5 \times 10^{-7} \text{ GBP per read}$$

For intermediate years we interpolate smoothly (e.g. using a log-linear interpolation in time).

These values imply, for example, that at $n = 1.417 \times 10^7$ reads (the median design depth for the unknown class):

$$\text{cost_per_sample_2025} \approx 150 + 4.5\text{e-}7 * 1.417\text{e}7 \approx \text{£}156.4$$

and at $n = 7.055 \times 10^8$ reads (the robust design depth):

$$\text{cost_per_sample_2025} \approx 150 + 4.5\text{e-}7 * 7.055\text{e}8 \approx \text{£}467.5$$

which match the values in the main tables.

Using the formulas above, we map each design depth ($n_{\text{min, median}}$, $n_{\text{min, 95th}}$, or any chosen n) into annual per-site and network costs across years.

B.7 Stealth scenario calculations

To approximate an adversarially designed “stealth” pathogen, we define a low-abundance scenario based on the lower tail of the unknown prior.

B.7.1 Defining the stealth relative abundance

Let F_{unknown} denote the cumulative distribution function (CDF) of the unknown $RA_{1\text{pct}}$ mixture. The stealth RA is defined as the 10th percentile of this distribution:

$$RA_{\text{steath}} = F_{\text{unknown}}^{-1}(0.10)$$

Operationally, we compute this by taking the empirical 10th percentile of the RA_{1pct} draws used for the unknown mixture.

B.7.2 Detection and minimum depth for the stealth scenario

Treating $RA_{stealth}$ as a fixed value, we apply exactly the same detection model as before.

We translate the depth into annual per-site and network costs using the cost model in B.6:

$$\text{cost_per_sample_stealth_y} = \text{processing_cost} + c_read_y * n_{min, stealth}$$

$$\text{cost_annual_per_site_stealth_y} = 52 * \text{cost_per_sample_stealth_y}$$

$$\text{cost_annual_network_stealth_y} = N_sites * \text{cost_annual_per_site_stealth_y}$$

Comparing $n_{min, stealth}$ and its costs to $n_{min, median}$ and $n_{min, 95th}$ for the unknown mixture quantifies how much additional sequencing effort (and budget) would be required to maintain a 70% weekly detection target if an adversary deliberately pushes wastewater shedding into the lower tail of the plausible RA_{1pct} range.