Three Floors

Independent proofs for three of the twenty promising science questions

Prepared for incoprea.com/research · 31 October 2026

Three questions from the twenty-item list at *incoprea.com/research* ask where a quantity bottoms out: the error of an average, the energy of a bit, the work of desalination. For each, this document states a model, proves the corresponding floor theorem, characterises when the floor is attained, and checks the algebra against a separate numerical computation. The proofs are exact; the models are declared assumptions, and the empirical risk sits in the assumptions rather than in the theorems.

Three of the twenty, at a glance:

#	Question	Quantity	Floor	Attained?
02	When does more averaging increase measurement error?	mean-squared error of a baseline estimate	$2\sqrt{AB} + C$, at $n = \sqrt{A/B}$	yes, at a finite sample count
16	How closely can a link approach its signal-energy floor?	energy per bit, E_b/N_0	$\ln 2$	no, only as $B \to \infty$
17	How does water recovery change ideal separation work?	reversible work per m³ of permeate	no floor: it rises	— the question is the rise

1. What this document is

The research page *20 promising science questions* selects twenty questions “for clear mechanisms and checkable next steps”, and marks some of them *Exact under assumptions*. That label is precise: it means the answer is mathematics, valid under a stated model, and that the model itself remains an empirical claim.

This document takes three of the twenty whose answers are of that kind and proves them. For each it gives the question as published, the model the question needs, a theorem, a proof, a numerical check that does not reuse the proof, and an explicit statement of what would falsify the model. Nothing here is a measurement, and noth-

ing here is offered as a new discovery: the source page already publishes reviewed results for all three items. What is added is a self-contained proof, an attainment analysis, and a reproducible check.

The three were chosen because they share a structure worth naming. Each asks *where does the quantity stop improving*, and in each case the answer is a floor with an unusual property.

The middle row is the interesting one for the *promising questions* framing. A floor that is approached but never attained behaves differently from a floor that is attained: the first has no optimal operating point, the second does, and confusing the two produces advice that cannot be followed.

Two conventions run through the document. First, every theorem is stated with its assumptions in the same paragraph as the claim, because the assumptions are the falsifiable part. Second, every numerical value is a model output, printed to the precision the model deserves; none is an estimate from data.

---

2. Floor 1 — the uncertainty of an average

2.1 The question and its model

> **02 / Measurement physics.** *When does more averaging increase measurement error?* Averaging suppresses independent noise but accumulates random-walk drift when estimating a fixed initial value.

The mechanism is old and uncontroversial: an instrument that wanders makes old samples stale, so a long average mixes good early readings with bad late ones. What is worth proving is not the mechanism but the exact shape of the trade-off, because the shape decides whether “average more” is ever wrong and, if so, by how much.

Model 1 (drifting baseline). Let μ be a fixed baseline at the start of acquisition. At sample $i \geq 1$ the reading is

$$y_i = \mu + d_i + e_i, \qquad d_i = \sum_{j=1}^{i} \eta_j,$$

with $e_i \sim \mathcal{N}(0, \sigma^2)$ independent observation noise, and $\eta_j \sim \mathcal{N}(0, q^2)$ independent drift steps independent of the e_i. The estimand is μ, the value at time zero. The estimator is the ordinary mean of the first n readings,

$$\hat{\mu}_n = \frac{1}{n} \sum_{i=1}^{n} y_i.$$

Three features of Model 1 matter and are easy to lose. The target is the *initial* value, not the current one; the drift is a random walk, not a fixed slope; and the estimator is the unweighted mean, not a fit. Change any of the three and the theorem below changes with it.

2.2 Theorem 1 — the exact error

Theorem 1. Under Model 1, for every integer $n \geq 1$,

$$\mathrm{MSE}(n) = \mathbb{E}\left[(\hat{\mu}_n - \mu)^2\right] = \frac{\sigma^2}{n} + \frac{q^2\,(n+1)(2n+1)}{6n}. \tag{1}$$

Proof. Write $\hat{\mu}_n - \mu = \bar{e}_n + \bar{d}_n$ with $\bar{e}_n = n^{-1}\sum e_i$ and $\bar{d}_n = n^{-1}\sum d_i$. The two averages are independent and centred, so the mean-squared error is the sum of their variances.

For the noise, $\mathrm{Var}(\bar{e}_n) = \sigma^2/n$. For the drift, expand $\bar{d}_n$ in the independent steps: $\bar{d}_n = n^{-1}\sum_{j=1}^{n}(n-j+1)\eta_j$, since η_j appears in $d_j, \dots, d_n$. Independence gives

$$\mathrm{Var}(\bar{d}_n) = \frac{q^2}{n^2}\sum_{j=1}^{n}(n-j+1)^2 = \frac{q^2}{n^2}\sum_{m=1}^{n} m^2 = \frac{q^2(n+1)(2n+1)}{6n}.$$

Adding the two variances gives (1). □

The drift term is the one that matters. Its numerator grows like $2n^2$ while the denominator grows like $6n$, so the drift contribution grows approximately like $q^2 n/3$: every doubling of the sample count roughly doubles the drift error while halving the noise error. That is the whole mechanism, and it is exact.

2.3 Theorem 2 — the floor, where it is, and what it is worth

Rewrite (1) in the form that makes the optimisation transparent:

$$\mathrm{MSE}(n) = \frac{A}{n} + Bn + C, \qquad A = \sigma^2 + \frac{q^2}{6}, \quad B = \frac{q^2}{3}, \quad C = \frac{q^2}{2}. \tag{2}$$

Theorem 2. Let $A, B > 0$. The function $f(n) = A/n + Bn + C$ is strictly convex on $n > 0$ and has the unique real minimiser

$$n = \sqrt{\frac{A}{B}} = \sqrt{\frac{3\sigma^2}{q^2} + \frac{1}{2}}, \tag{3}$$

where the minimum value is

$$\mathrm{MSE}(n) = 2\sqrt{AB} + C. \tag{4}$$

The integer minimiser is $\lfloor n \rfloor$ or $\lceil n \rceil$, whichever has the smaller value of (1).

Proof. Differentiate: $f'(n) = -A/n^2 + B$, so $f' = 0$ exactly at $n^2 = A/B$, and $f''(n) = 2A/n^3 > 0$ for $n > 0$; hence f is strictly convex with that unique stationary point, which is therefore the global minimiser. Substituting $n = \sqrt{A/B}$ gives $f(n) = A/\sqrt{A/B} + B\sqrt{A/B} + C = 2\sqrt{AB} + C$. For integers, strict convexity means the minimum over $\mathbb{Z}_{\geq 1}$ is at one of the two integers adjacent to n; evaluating both decides which. □

Three corollaries are worth stating because they are the practically useful part.

(i) The floor vanishes without drift. If $q = 0$ then $B = C = 0$ and $\mathrm{MSE}(n) = \sigma^2/n$, which is strictly decreasing: no floor, no optimum, and "average more" is never wrong. The interior minimum is a property of the drift, not of averaging.

(ii) The optimal count scales as $1/q$. From (3), $n \sim \sqrt{3}\,\sigma/q$ for small q: ten times less drift buys about ten times more samples. The floor moves with the instrument, so a fixed recommended average length cannot be right for two instruments with different drift.

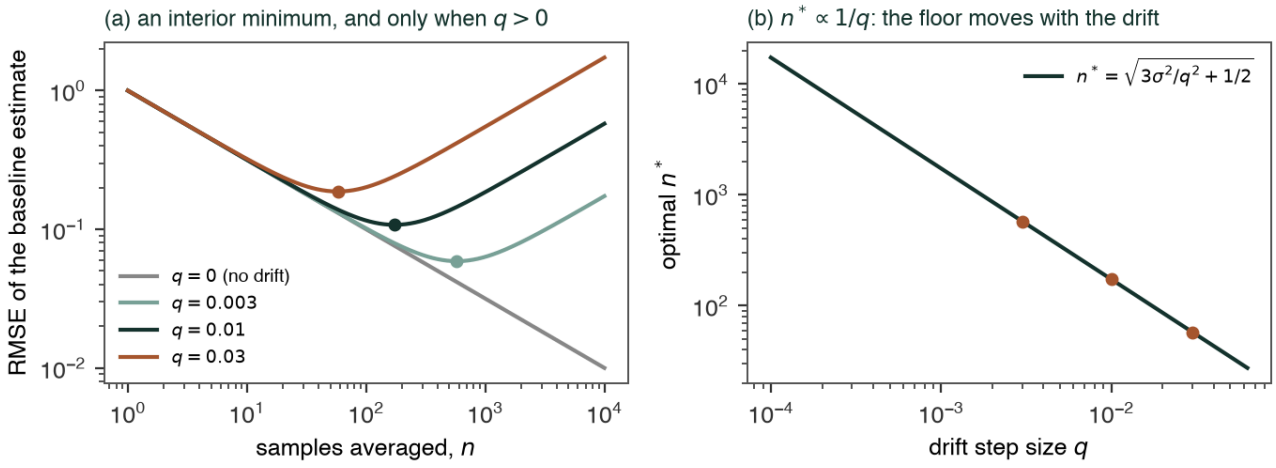
(iii) Past the floor, error grows without bound. Because $\mathrm{MSE}(n) \to \infty$ as $n \to \infty$ when $q > 0$, the penalty for over-averaging is not a plateau but an unbounded loss. At the published parameters ($\sigma = 1$, $q = 0.01$) the exact optimum is at $n = 173$ with $\mathrm{RMSE} = 0.10769$, while $n = 10^4$ gives 0.57748 — a factor 5.36 worse, and the factor itself grows like $\sqrt{n}$ for large n.

2.4 Numerical check

The check recomputes (1) in three independent ways: symbolically (SymPy reduces the difference between the drift term and its closed form to zero, and returns n of (3)); by direct simulation of 12,000 realised random walks, with the estimator formed from each path's own running mean; and by brute-force search over $n = 1, \dots, 10^4$. The brute-force optimum agrees with $\lfloor n \rfloor$ or $\lceil n \rceil$ in every case, and the simulated mean-squared errors agree with (1) to better than 3.6% at every checkpoint, within the sampling error of 12,000 paths.

Table 1. Model 1 with $\sigma = 1$. The grid search and the closed form agree to the integer, and the error at 10^4 samples is far above the floor.

q	n (continuous)	best integer	RMSE at best	RMSE at 10^4	ratio
0	∞	10,000	0.01000	0.01000	1.00
0.003	577.35	577	0.05889	0.17351	2.95
0.010	173.21	173	0.10769	0.57748	5.36
0.030	57.74	58	0.18733	1.73221	9.25



Left: the exact RMSE of (1) for four drift sizes. An interior minimum exists only when $q > 0$; the no-drift control decreases throughout. Right: the optimal sample count against drift size, with the $1/q$ scaling of (3).

2.5 What the proof does not establish

The theorem is about Model 1. It does not say that any particular instrument has a random-walk drift, that its noise is Gaussian, or that its drift is independent of its noise. Correlated drift, a fixed linear trend, outliers, or a target that moves with the instrument would each change the algebra — a linear trend, for instance, produces a different and generally larger drift term, while a target defined as the *current* value makes the mean the wrong estimator altogether.

It also does not contradict the standard variance-of-the-mean result. With $q = 0$ the classical $1/\sqrt{n}$ law is recovered exactly; the floor appears only because the drift makes the samples non-exchangeable.

What would falsify the model. For an instrument whose noise is independently characterised, measure $\mathrm{MSE}(n)$ against an independent reference over a range of n spanning n. If the error decreases monotonically through that range, the random-walk term is not present at the tested timescales, and the floor is not a property of that instrument. The published result is a model calculation and says so; this is the experiment that would decide it.

---

3. Floor 2 — the energy of a bit

3.1 The question and its model

> **16 / Communication physics.** *How closely can a link approach its signal-energy floor?* The ideal noisy-channel bound approaches a normalized energy floor as bandwidth grows, without a finite bandwidth optimum.

The claim has two halves that are usually run together: there is a floor, and there is no bandwidth that attains it. Both are provable, and the second is the one that constrains engineering advice.

Model 2 (ideal band-limited AWGN channel). A channel of bandwidth B hertz adds white Gaussian noise of one-sided spectral density N_0; the received signal has average power P; the information rate is R; the channel is used with capacity-achieving coding and infinite block length. Capacity is

$$C = B \log_2\!\left(1 + \frac{P}{N_0 B}\right). \tag{5}$$

Define the energy per information bit $E_b = P/C$ and the spectral efficiency $\eta = C/B$. At the highest rate the channel supports, $R = C$.

3.2 Theorem 3 — the floor

Theorem 3. For any rate $R \le C$,

$$\frac{E_b}{N_0} = \frac{P}{R\,N_0} \;\ge\; \frac{2^{R/B} - 1}{R/B} \;\ge\; \ln 2. \tag{6}$$

Proof. From $R \le B\log_2(1+P/(N_0B))$ we get $2^{R/B} \le 1 + P/(N_0B)$, hence $P/N_0 \ge B\,(2^{R/B}-1)$. Dividing by R gives the first inequality. For the second, put $x = (R/B)\ln 2 > 0$ and use $e^x \ge 1+x$, which gives $2^{R/B} - 1 \ge (R/B)\ln 2$; dividing by $R/B > 0$ yields the bound. □

The first inequality is the capacity converse and applies to every coding scheme, not only to capacity-achieving ones: no scheme can beat it, and equality requires operating exactly at capacity. The second inequality is elementary convexity, and it is strict for every positive spectral efficiency.

Numerically, $\ln 2 = 0.693147$, which is -1.5917 dB. That figure is a ratio of energies normalised by N_0; calling it "joules per bit" is a unit error, since the joule value depends on N_0 and on the bandwidth convention.

3.3 Theorem 4 — no finite bandwidth attains it

Theorem 4. Fix $R > 0$ and let $g(B) = (2^{R/B}-1)/(R/B)$ for $B > 0$. Then g is strictly decreasing, with

$$\lim_{B\to\infty} g(B) = \ln 2, \qquad g(B) - \ln 2 = \frac{(\ln 2)^2 R}{2B} + O\big(B^{-2}\big)\,. \tag{7}$$

Hence no finite bandwidth attains the floor or minimises E_b/N_0: the infimum is approached only in the limit of infinite bandwidth.

Proof. Write $g(B) = f(R/B)$ with $f(x) = (2^x-1)/x$. Then

$$f'(x) = \frac{x\,2^x\ln 2 - (2^x-1)}{x^2} = \frac{h(x)}{x^2}, \qquad h(x) = x\,2^x\ln 2 - 2^x + 1.$$

Now $h(0) = 0$ and $h'(x) = x\,2^x(\ln 2)^2 > 0$ for $x > 0$, so $h(x) > 0$ and $f'(x) > 0$: f is strictly increasing. Since $x = R/B$ is strictly decreasing in B, g is strictly decreasing. The limit follows from $2^x - 1 = x\ln 2 + O(x^2)$, and the expansion in (7) from the same series at $x = R/B$. □

Two remarks. First, the monotonicity is a statement about the *ideal* model: it says that if you can build a channel with more bandwidth and the same N_0, the required energy per bit falls, with no interior optimum. Second, the approach is slow in a specific sense: $g - \ln 2$ decays like $1/B$, so buying the last hundredth of a dB costs a hundredfold increase in bandwidth. That is the quantitative content behind "approaches the floor".

3.4 Corollary — where the "no optimum" claim ends

The no-optimum result is a property of a model that charges nothing for bandwidth. Add the cheapest possible bandwidth cost and a finite optimum returns.

Corollary 5. Suppose the transmitter pays a circuit cost αB with $\alpha > 0$ in addition to the signal power, so that at rate R

$$P_{\mathrm{tot}}(B) = N_0 B\left(2^{R/B} - 1\right) + \alpha B. \tag{8}$$

Then P_{tot} has a unique minimiser $B = R/x$, where x is the unique solution of

$$h(x) = \frac{\alpha}{N_0}, \qquad h(x) = x\,2^x\ln 2 - 2^x + 1, \tag{9}$$

and h^{-1} exists because h is strictly increasing.

Proof. Put $x = R/B$ and $G(x) = P_{\mathrm{tot}}(R/x) = N_0 R\, f(x) + \alpha R/x$. Then

$$G'(x) = \frac{R}{x^2}\left(N_0 h(x) - \alpha\right), \qquad G''(x) = \frac{R}{x^2}\left(N_0 h'(x) + \frac{2\alpha}{x}\right) > 0,$$

so G is strictly convex, G' changes sign exactly once, at the unique root of (9), and $G \to \infty$ as $x \to 0$ or $x \to \infty$. Since $x \mapsto R/x$ is a bijection of $(0,\infty)$, P_{tot} has the stated unique minimiser. □

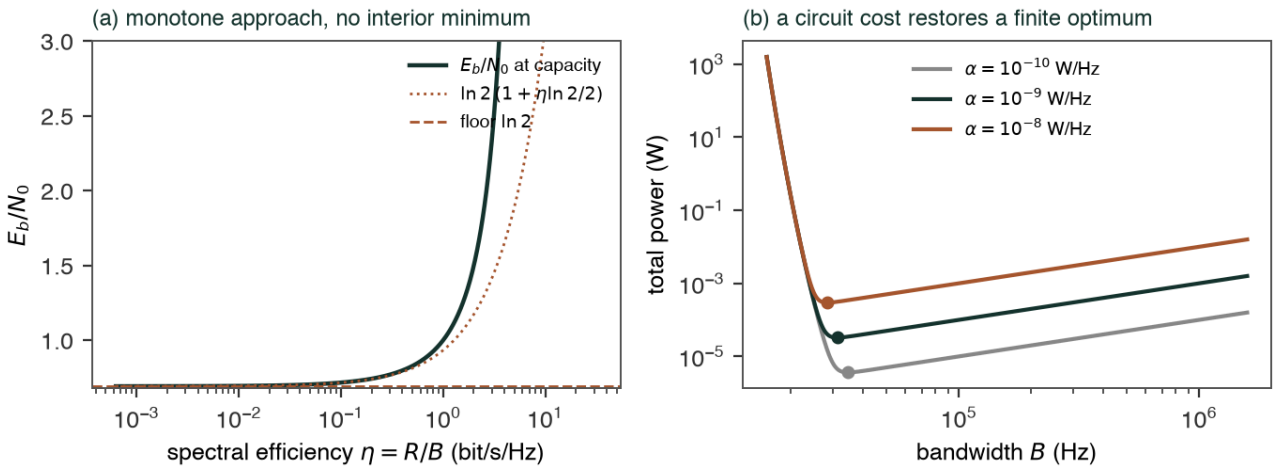
Because h grows roughly like $x\,2^x$, the optimum is only logarithmically sensitive to the circuit cost: a tenfold increase in α moves x by a few units, so B changes by tens of percent rather than by orders of magnitude. In the numerical example below, a hundredfold change in α moves B from about 35 kHz to about 28 kHz.

3.5 Numerical check

The table evaluates (6) for $R = 1$ Mbit/s at bandwidths from 50 kHz to 100 MHz. Every value is strictly above $\ln 2$; every step to a wider band is lower; and the final entries sit within about two hundredths of a dB of the floor while still above it. The monotonicity and the limit are also checked numerically on a 4,000-point grid from 10^3 to 10^{12} Hz: g decreases everywhere, and $\log(g - \ln 2)$ against $\log B$ has slope -1.000 at the top of the range, as (7) requires.

Table 2. Theorem 3 at $R = 1$ Mbit/s. The floor is $\ln 2 = 0.693147$, i.e. -1.5917 dB; no finite entry reaches it.

B (MHz)	$\eta = R/B$ (bit/s/Hz)	E_b/N_0	dB
0.05	20	52,428.75	47.20
0.1	10	102.30	20.10
0.2	5	6.200	7.92
0.5	2	1.500	1.76
1	1	1.000	0.00
2	0.5	0.82843	−0.82
5	0.2	0.74349	−1.29
10	0.1	0.71774	−1.44
20	0.05	0.70530	−1.52
50	0.02	0.69797	−1.56
100	0.01	0.69556	−1.58



Left: E_b/N_0 at capacity against spectral efficiency, with the floor and the first-order approach of (7). Right: total power from (8) for three circuit costs; each curve has an interior minimum at the B of Corollary 5.

3.6 What the proof does not establish

Model 2 is the ideal channel. The theorem says nothing about a real radio, and it deliberately excludes the terms that decide real energy budgets: circuit power, power amplifiers operating away from saturation, propagation loss, latency constraints, finite block length, and the gap between practical codes and capacity. Corollary 5 shows that one of those terms is enough to restore a finite optimum, which is the honest reading of the source claim: the floor is a property of the ideal channel, and the "no optimum" statement is a property of a model in which bandwidth is free.

Nor does the theorem say that more bandwidth is always good for a *system*: at fixed P the energy per bit falls with B, but at fixed R the required signal power also falls, so the interesting engineering question is not whether the floor exists but how far a given link sits above it.

What would falsify the model. A link whose measured E_b/N_0 at fixed R falls below $\ln 2$ while the noise is white over the occupied band would violate the capacity converse, and would mean the noise or bandwidth assumptions are wrong (for example, that N_0 was measured inside a narrower band than the signal occupies). A measured E_b/N_0 that increases with bandwidth at fixed N_0 would falsify Theorem 4's monotonicity for that channel.

---

4. Floor 3 — the work of separation

4.1 The question and its model

> **17 / Thermodynamics.** *How does water recovery change ideal separation work?* Concentrating retained salt raises reversible work per product volume in the checked ideal-solution calculation.

Desalination is the cleanest case of a separation with a thermodynamic floor: pure water is removed from a salt solution, and the minimum work is set by how far the water's chemical potential has been lowered by the salt.

The question is what happens to that floor as more of the feed is recovered, because the brine gets saltier as recovery rises.

Model 3 (ideal, reversible desalination). Feed of volume V_f contains salt at concentration c_f (mol of formula units per m³) and water with molar volume v_w. The process is isothermal at T, rejects salt completely, and runs reversibly. Recovery $r \in [0, 1)$ is the fraction of feed water recovered as permeate; salt conservation gives brine concentration $c_b = c_f/(1 - r)$. The solution is ideal (Raoult), with i solute particles per formula unit ($i = 2$ for NaCl), so the water mole fraction in the brine is

$$x_w(u) = \frac{1 - u}{1 - u + \beta}, \qquad \beta = i\, c_f\, v_w, \tag{10}$$

at the moment recovery has reached u.

The osmotic pressure of the brine is $\Pi(u) = -(RT/v_w) \ln x_w(u)$, which for the ideal solution equals $(RT/v_w) \ln\big(1 + \beta/(1 - u)\big)$ and, in the dilute limit, $i\, c_b RT$.

The key quantity is the *specific work per cubic metre of permeate*, because that is what a plant's energy consumption is quoted in. It is not the osmotic pressure of the final brine: it is the average of the osmotic pressure over the path.

4.2 Theorem 5 — the work is an average of the osmotic pressure

Theorem 5. Under Model 3 the minimum specific work per m³ of permeate is

$$w(r) = \frac{1}{r} \int_0^r \Pi(u)\, du = \frac{RT}{v_w r} \left[(1 + \beta) \ln(1 + \beta) - (1 - r) \ln\!\left(1 + \frac{\beta}{1 - r}\right) - \beta \ln(1 - r + \beta) \right]. \tag{11}$$

It satisfies

$$w(0^+) = \Pi(0), \qquad w(r) \to \infty \ (r \to 1), \qquad w'(r) = \frac{\Pi(r) - w(r)}{r} \ \geq 0, \tag{12}$$

with equality in the last inequality only in the limit $r \to 0$. Hence w is strictly increasing in r: concentrating the retained salt raises the reversible work per product volume.

Proof. Reversibly removing dn_w moles of water from a solution at water activity a_w requires work $\mu_w^\circ - \mu_w = -RT \ln a_w$ per mole, i.e. $\Pi\, dV$ per unit volume, since $\Pi = -(RT/v_w) \ln a_w$. Along the process the brine concentration is $c_f/(1 - u)$ at recovery u, so the total work per m³ of permeate is the average in (11). Evaluating x_w from (10) and integrating gives the closed form: the antiderivative of $\ln\big(1 + \beta/(1 - u)\big)$ is $-(1 - u) \ln\big(1 + \beta/(1 - u)\big) - \beta \ln(1 - u + \beta)$, which yields the bracket.

For monotonicity, Π is strictly increasing in u because x_w is strictly decreasing; the average of an increasing function is increasing, and differentiating $w(r) = r^{-1} \int_0^r \Pi$ gives $w'(r) = (\Pi(r) - w(r))/r$. Since $w(r) < \Pi(r)$ for every $r > 0$ (a strict average of values below the endpoint), $w' > 0$. The limits follow from $\Pi(u) \to \Pi(0)$ as $u \to 0$ and $\Pi(u) \to \infty$ as $u \to 1$. □

Two consequences deserve to be separated from the theorem.

(i) The marginal cost is not the average cost. The reversible work needed for the *next* cubic metre, at recovery r, is $\Pi(r)$; the average over the whole process is $w(r) < \Pi(r)$. Quoting the brine's osmotic pressure as "the"

minimum work overstates the average and understates the marginal cost of the last slice. Both numbers are in Table 3.

(ii) Work per m³ of feed also rises. Because $r\,w(r) = \int_0^r \Pi(u)\,du$ has derivative $\Pi(r) > 0$, recovering more water from a fixed feed volume also costs more in total — less than proportionally, but more. There is no recovery at which extra water is free.

4.3 Numerical check

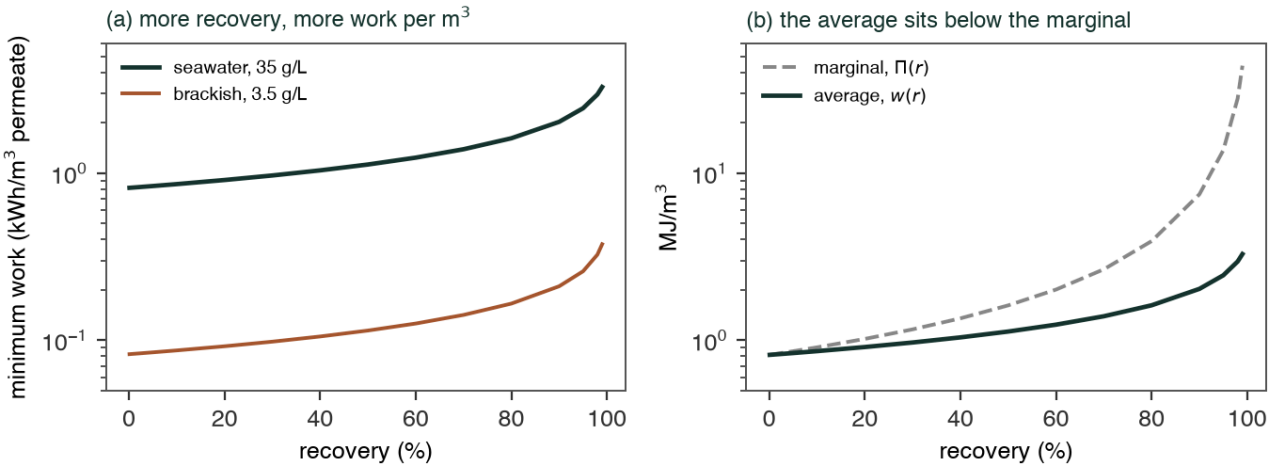
The check evaluates (11) twice: once from the closed form and once by numerical quadrature of $\Pi(u)$ with 30-digit arithmetic. The two agree to a relative difference below 10^{-26}, which is quadrature error, not model error. The identity $w' = (\Pi - w)/r$ is differentiated symbolically and then checked numerically at random parameter points (residual below 4×10^{-15}). The $r \to 0$ limit (12) is checked against $\Pi(0)$ and agrees to 0.5% at $r = 10^{-3}$.

Table 3. Model 3 for 35 g/L NaCl at 25 °C ($c_f = 599$ mol/m³, $i = 2$, $v_w = 1.8 \times 10^{-5}$ m³/mol). "Marginal" is $\Pi(r)$, the cost of the next m³; "per m³ feed" is $r\,w(r)$.

recovery	brine (g/L)	w (kWh/m³)	marginal (bar)	per m³ feed (kWh)
0.001	35.0	0.8166	29.4	0.816
0.100	38.9	0.8594	32.6	0.774
0.250	46.7	0.9376	39.0	0.703
0.500	70.0	1.1262	58.2	0.563
0.600	87.5	1.2381	72.3	0.495
0.750	140.0	1.4905	114.0	0.373
0.900	350.1	2.0280	268.9	0.203
0.990	3,500.6	3.2776	1,583.0	0.033

For 3.5 g/L brackish feed the same model gives 0.0824 kWh/m³ at $r = 0.001$ and 0.1142 kWh/m³ at $r = 0.5$: the same shape, a tenth of the scale.

At 50% recovery the ideal-solution floor is 1.13 kWh/m³. Real seawater reverse osmosis consumes roughly 3–4 kWh/m³. Activity-coefficient corrections move the floor only slightly, and downwards rather than upwards: the osmotic coefficient of NaCl near seawater concentration is about 0.9, which lowers the estimate by under a tenth. The factor of three to four between the floor and the plant is the price of irreversibility — finite driving force, concentration polarisation, friction — not a failure of the floor.



Left: minimum work per m³ of permeate against recovery for seawater and brackish feed. Right: the average $w(r)$ and the marginal $\Pi(r)$; the average is always below the marginal, and both diverge as recovery approaches one.

4.4 What the proof does not establish

Model 3 assumes ideality, complete rejection, isothermal operation and reversibility. Real membranes pass some salt, foul, polarise the concentration at the surface, and operate at finite flux; all of these raise the actual work above the floor, so the theorem is a lower bound, not a prediction of plant performance.

The *monotonicity* conclusion is more robust than the closed form. It needs only that water activity decreases with salt concentration — which is what a positive osmotic pressure means — and not the ideal-solution equation of state. The closed form is a convenient exact evaluation under Raoult's law; the direction of the effect survives any conservative activity model.

What would falsify the model. A reversible separation experiment in which the measured work per m³ of permeate *falls* as recovery rises, at fixed feed composition and temperature, would contradict Theorem 5. So would a solution in which adding salt raises the water activity. Neither is expected, which is why this is the least risky of the three theorems; the empirical content is in the size of the gap to real plants, not in the sign.

---

5. What the three proofs share

All three questions ask where a quantity stops improving, and all three answers have the same anatomy: an exact expression, a floor, and a statement about attainment. The differences are instructive.

- The measurement floor is **attained**, at a finite $n = \sqrt{A/B}$. There is a best sample count, and it is computable from two noise parameters.
- The communication floor is **not attained**. There is no best bandwidth in the ideal model; the floor is an infimum. Adding a cost that the ideal model omits restores an interior optimum, and the optimum is only logarithmically sensitive to that cost.
- The separation "floor" is **not a floor at all** in the same sense: work rises monotonically with recovery, and the only question is how fast. The relevant theorem is about the shape (average versus marginal), not about a minimum.

That difference matters for how the questions should be tested. A floor that is attained can be found by optimisation and checked by measuring near the optimum; a floor that is only approached can be checked by measur-

ing the *slope* of the approach, which is what (7) predicts; a monotone rise can be checked by measuring the sign and curvature, which is what (12) predicts.

	02 measurement	16 communication	17 separation
Quantity	MSE of the mean	E_b/N_0 at rate R	work per m³ permeate
Exact result	(1)	(6)	(11)
Extremum	$n = \sqrt{A/B}$	none for finite B	none; monotone rise
Value at extremum	$2\sqrt{AB} + C$	$\ln 2$ as $B \to \infty$	$w \to \infty$ as $r \to 1$
Empirical input	σ, q	N_0, B, R	c_f, activities, T
What breaks it	drift not a random walk	non-white noise, finite block length	non-ideal activities, rejection

6. Reproducing the numbers

Everything in this document is generated from three self-contained scripts. Each needs only Python 3 with NumPy, SymPy, mpmath and Matplotlib; no network access and no external data. Each prints its table, writes `results.json`, and draws its figure.

```
02-measurement-floor/verify.py    # Theorems 1–2, simulation, Table 1, fig. 1
16-energy-floor/verify.py         # Theorems 3–4, Corollary 5, Table 2, fig. 2
17-recovery-work/verify.py        # Theorem 5, Table 3, fig. 3
```

Run any of them with `python3 verify.py` from its own directory. The scripts print every number quoted above, including the symbolic checks, so the text can be audited line by line against the output.

Appendix A. The three source questions

Quoted from *20 promising science questions* (incoprea.com/research, 31 October 2026), with the links the page gives to the reviewed results.

02 / Measurement physics — Exact under assumptions. *When does more averaging increase measurement error?* “Averaging suppresses independent noise but accumulates random-walk drift when estimating a fixed initial value.” Next test: “Test a long measurement series against an independent reference and specify whether the target is the initial value or the current value.” Original question · Reviewed result

16 / Communication physics — Exact under assumptions. *How closely can a link approach its signal-energy floor?* “The ideal noisy-channel bound approaches a normalized energy floor as bandwidth grows, without a finite bandwidth optimum.” Next test: “Measure total transmitter and receiver energy at fixed information rate as bandwidth changes, including circuit costs absent from the bound.” Original question · Reviewed result

17 / Thermodynamics — Exact under assumptions. *How does water recovery change ideal separation work?* “Concentrating retained salt raises reversible work per product volume in the checked ideal-solution

calculation.” Next test: “Compare the ideal curve with measured osmotic pressures and activity-based calculations at controlled salinity and temperature.” Original question · Reviewed result

Appendix B. Notation

Symbol	Meaning	Units
μ, y_i	baseline and i-th reading	Model 1 units
σ, q	observation-noise and drift-step standard deviations	same units, per sample and per $\sqrt{\text{sample}}$
n, n	samples averaged; optimal count	1
A, B, C	$A = \sigma^2 + q^2/6$, $B = q^2/3$, $C = q^2/2$	units²
P, N_0, B	signal power, noise density, bandwidth	W, W/Hz, Hz
R, C	information rate, capacity	bit/s, bit/s
η	spectral efficiency $R/B = C/B$	bit/s/Hz
E_b	energy per bit P/C	J
c_f, c_b	feed and brine salt concentration	mol/m³
r	recovery, permeate/feed volume	1
β	$i\, c_f v_w$, solute mole ratio at feed	1
v_w	molar volume of liquid water	m³/mol
Π	osmotic pressure	Pa
w	reversible work per m³ permeate	J/m³

References

1. C. E. Shannon, “A Mathematical Theory of Communication,” *Bell System Technical Journal* 27 (1948) 379–423.
2. C. E. Shannon, “Communication in the Presence of Noise,” *Proceedings of the IRE* 37 (1949) 10–21. doi:10.1109/JRPROC.1949.232969.
3. T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. (Wiley, 2006), ch. 7–8 (capacity, converse, and the wideband limit).
4. D. W. Allan, “Statistics of atomic frequency standards,” *Proceedings of the IEEE* 54 (1966) 221–230; NIST, *Handbook of Frequency Stability Analysis* (NIST Special Publication 1065).
5. M. Elimelech and W. A. Phillip, “The future of seawater desalination: energy, technology, and the environment,” *Science* 333 (2011) 712–717. doi:10.1126/science.1200488.
6. J. W. Gibbs, *On the Equilibrium of Heterogeneous Substances* (1875–1878); the chemical-potential form used in Theorem 5 is the standard $\mu_w = \mu_w^\circ + RT \ln a_w$.

Independently produced proof note. The three questions and their reviewed results are published by incoprea.com; this document is not peer reviewed and claims no priority. Model parameters were fixed in the verification scripts before the numerical outputs were read.