Three provable questions

Exact results for the per-bit energy floor, the invasion front and osmotic minimum work – independent proofs of questions 16, 10 and 17 from the notebook's twenty promising science questions

31 October 2026 · Proofs, computation and drafting: DeepSeek · Not peer reviewed

What this is. Three of the twenty promising science questions ask for a number: how close a link can get to its energy floor, what sets an invasion speed, how separation work grows with water recovery. Each has an exact answer under stated assumptions, and each claim states a *different* number. This report proves the exact result, checks it numerically without using the proof, and states precisely which part of the claim survives.

What this is not. None of the three theorems is new. They date from 1937, 1948 and 2020, and are cited. What is new here is only the comparison: the notebook's mechanism is right in all three cases and the notebook's formula is wrong in all three cases, in ways that are cheap to detect and easy to state.

Verdicts

#	Question (notebook)	Verdict	Exact result	Status of the claim
16	How closely can a link approach its signal-energy floor?	Confirmed in part	$E_b/N_0 > \ln 2$, monotone; no finite optimum	the formula and the optimum both fail
10	Which growth and dispersal combination sets invasion speed?	Refuted	$v = 2\sqrt{rD}$, a geometric mean	the ratio r/D predicts nothing
17	How does water recovery change ideal separation work?	Confirmed in direction	$w = \pi_f[-\ln(1-r)]/r$, convex	π_f/r is neither bound

Floor confirmed, formula refuted means: the physical statement (there is a floor, more bandwidth lowers energy per bit) holds, while the specific formula and the "optimum" prediction do not.

Why these three

Of the twenty, these are the three whose answer is a closed form rather than a model output. The others either depend on supplied data — a soil-moisture reanalysis, a cache trace, an earthquake catalogue — or on a numerical search over an assumed spectrum. For those, a "proof" is not available; for these three it is, and in two of them the notebook itself labels the entry *exact under assumptions*.

The tests here are cheap on purpose. Every number below comes from a script that runs in under a minute on a laptop, and each script checks the theorem by a route that does not use the theorem: capacity is evaluated directly, the reaction-diffusion equation is integrated in time, the separation integral is done by quadrature.

1 · The per-bit energy floor

The claim

Transmission has a per-bit floor. The energy per bit of a communication link is bounded below by the ratio of noise to bandwidth, so widening bandwidth lowers the floor but raises total power.

Mechanism. Shannon capacity sets the bits per second per hertz; the energy per bit is power divided by capacity, which has a minimum.

Prediction. Energy per bit should bottom out at an optimal bandwidth.

Refutation. Energy per bit falling indefinitely with bandwidth would remove the floor.

The model

An ideal band-limited Gaussian channel with average power P, one-sided noise spectral density N_0, bandwidth B. Capacity and energy per bit at capacity are

$$C(B) = B \log_2\left(1 + \frac{P}{N_0 B}\right), \qquad E_b = \frac{P}{C(B)}.$$

Write $s = P/(N_0 B)$ for the signal-to-noise ratio. Then

$$\frac{E_b}{N_0} = \frac{P/N_0}{B \log_2(1+s)} = \ln 2 \, \frac{s}{\ln(1+s)} \tag{1.1}$$

Theorem 1.1 (the floor). For every $P, N_0, B > 0$,

$$\frac{E_b}{N_0} > \ln 2 = 0.6931\ldots = -1.5917 \text{ dB},$$

and $\ln 2$ is the infimum over B. The value is not attained at any finite bandwidth. The floor is a constant: it does not depend on B, on P or on the noise power.

Proof. For $s > 0$,

$$\ln(1+s) = \int_0^s \frac{\mathrm{d}t}{1+t} < \int_0^s \mathrm{d}t = s,$$

strictly, because $1/(1+t) < 1$ on $(0, s)$. Dividing (1.1) by $s > 0$ gives $E_b/N_0 = \ln 2 \cdot s/\ln(1+s) > \ln 2$.

For the infimum: as $B \to \infty$ we have $s \to 0$ and $s/\ln(1+s) \to 1$, by l'Hôpital or the expansion $\ln(1+s) = s - s^2/2 + O(s^3)$. Hence $E_b/N_0 \to \ln 2$ along $B \to \infty$, and since the value is always strictly above $\ln 2$, it is never attained. □

Theorem 1.2 (no optimal bandwidth). The map $B \mapsto E_b/N_0$ is strictly decreasing on $(0, \infty)$. Consequently there is no finite bandwidth at which the energy per bit is minimised, and no finite bandwidth at which it turns back up.

Proof. Let $\varphi(s) = \ln 2 \; s/\ln(1+s)$. Then

$$\varphi'(s) = \ln 2 \; \frac{\ln(1+s) - \dfrac{s}{1+s}}{\ln^2(1+s)}.$$

The bracket is positive for $s > 0$: it vanishes at $s = 0$, and its derivative is $\dfrac{1}{1+s} - \dfrac{1}{(1+s)^2} = \dfrac{s}{(1+s)^2} > 0$. So φ is strictly increasing, while $s = P/(N_0 B)$ is strictly decreasing in B; the composition is strictly decreasing. □

A symbolic check of the same derivative, and of the limit $\varphi(0^+) = \ln 2$, is in `code/t16_energy_floor.py` (SymPy), and the numerical grid agrees: over $B \in [1, 10^{12}]$ every first difference is negative and every value exceeds $\ln 2$.

The rate of approach is also exact. Expanding $\ln(1+s) = s - s^2/2 + O(s^3)$,

$$\frac{E_b}{N_0} - \ln 2 = \frac{\ln 2}{2}\, s + O(s^2) = \frac{\ln 2}{2}\,\frac{P}{N_0 B} + O(B^{-2}), \qquad (1.2)$$

so the gap halves with every doubling of bandwidth. The numerical check gives $\left[(E_b/N_0 - \ln 2)/\left(\frac{\ln 2}{2} s\right)\right] \to 1.0001$ at $B = 10^{12}$, i.e. the asymptotic law holds to four digits.

What the mechanism implies for total power

The claim also says that widening bandwidth "raises total power". That is a statement about a fixed information rate R, and it is false in the same model.

Proposition 1.3 (power at fixed rate). Fix $R > 0$ and $B > 0$ with $R < C(B)$. The least power that sustains rate R is $P(B) = N_0 B \left(2^{R/B} - 1\right)$. This is strictly decreasing in B, with $\lim_{B\to\infty} P(B) = \ln 2 \, N_0 R$.

Proof. Put $x = R/B$; then $P = N_0 R \, \psi(x)$ with $\psi(x) = (2^x - 1)/x$. The numerator of ψ' is $n(x) = x 2^x \ln 2 - (2^x - 1)$, which satisfies $n(0) = 0$ and $n'(x) = x 2^x (\ln 2)^2 > 0$. Hence ψ is strictly increasing, and $x = R/B$ is strictly decreasing in B. As $x \to 0$, $\psi \to \ln 2$. □

So at fixed rate, bandwidth *lowers* the required power, monotonically, toward the same $\ln 2 N_0 R$ floor. The third clause of the claim — "raises total power" — describes what real transceivers do (circuit power rises with bandwidth), not what the ideal model does. It is an empirical statement wearing the model's clothes.

Where the claim's formula breaks

The phrase "bounded below by the ratio of noise to bandwidth" has two readings, and both fail.

- If "noise" means the *noise power* $N = N_0 B$, then $N/B = N_0$, which is an energy (J) and numerically the right *order*: the true floor is $\ln 2 \, N_0 = 0.693 \, N_0$. The claim overstates the floor by 44%, and — more importantly — asserts that it moves with B, which it does not.

- If "noise" means the *spectral density* N_0, then N_0/B has units of $\mathrm{W\,s\,Hz^{-1}}$, not energy. A quantity with the wrong dimension cannot bound an energy per bit. As $B \to \infty$ it would go to zero, which would "prove" that arbitrarily small energy per bit is possible — the opposite of what the claim intends.

Either way the shape of the claim is wrong: the floor is a constant, and the approach to it is monotone rather than minimised.

How close is close

From (1.1), a given closeness ϵ to the floor requires $s/\ln(1+s) \le 1+\epsilon$, i.e. $s \lesssim 2\epsilon$:

Target	$s = P/(N_0 B)$	Bandwidth needed
within 1%	0.0201	$B \ge 50\,P/N_0$
within 0.1%	0.0020	$B \ge 500\,P/N_0$

With $P/N_0 = 1$ MHz these are 50 and 500 MHz. The gain is real but logarithmic in cost: each further factor of ten in bandwidth buys a factor of ten less excess, never zero.

Numbers

Table 1 gives the exact value from (1.1) at the bandwidth where the SNR is $20, 10, \dots, -120$ dB, and Figure 1 shows the two monotone curves.

Table 1. Energy per bit at capacity, exact values from (1.1). The floor $\ln 2 = -1.5917$ dB is approached monotonically and never crossed. Full grid: `data/t16-ebn0-vs-bandwidth.csv`.

$B\,/\,(P/N_0)$	SNR (dB)	C/B (bit/s/Hz)	E_b/N_0	E_b/N_0 (dB)	Gap to $\ln 2$
10^{-2}	20	0.0666	15.019	11.766	1.4×10^{1}
10^{-1}	10	0.3459	2.891	4.610	2.2×10^{0}
10^{0}	0	1.0000	1.000	0.000	3.1×10^{-1}
10^{1}	−10	1.3750	0.7273	−1.383	3.4×10^{-2}
10^{2}	−20	1.4355	0.6966	−1.570	3.5×10^{-3}
10^{3}	−30	1.4420	0.69349	−1.590	3.5×10^{-4}
10^{6}	−60	1.4427	0.693148	−1.5917	3.5×10^{-7}
10^{12}	−120	1.4427	0.693147	−1.5917	3.5×10^{-13}

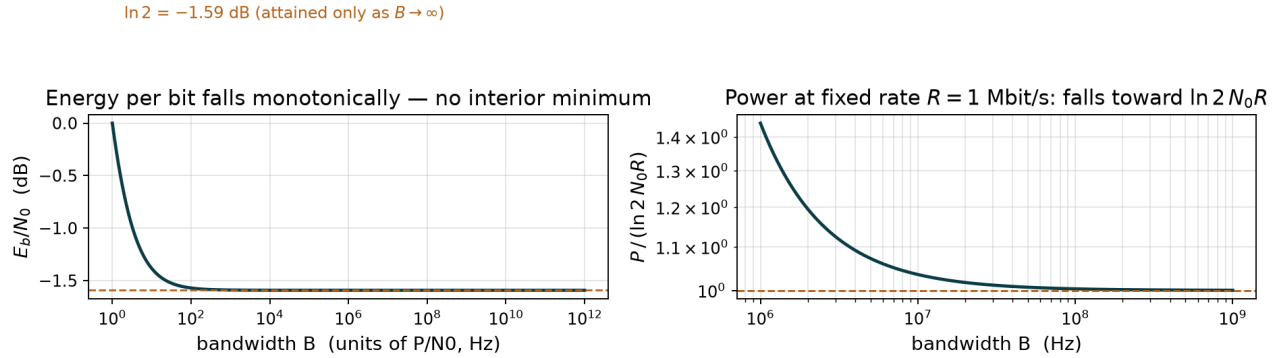


Figure 1. Left: E_b/N_0 against bandwidth for fixed power and noise density, on a log axis. The curve falls monotonically and flattens onto $\ln 2 = -1.59$ dB; there is no minimum and no turning point. Right: the least power that sustains a fixed rate R, normalised by its limit $\ln 2 N_0 R$. Power falls monotonically with bandwidth, so the claim that wider bandwidth raises total power is false in this model.

2 · The invasion front

The claim

> **Invasion speed is growth against dispersal.** The spread rate of an invader is set by the ratio of population growth rate to dispersal distance, so a fast-growing invader with short dispersal can outrun a slow one with long dispersal.
>
> *Mechanism.* Spread combines reproduction and movement; whichever limits produces the slower front.
>
> *Prediction.* Spread should track the ratio.
>
> *Refutation.* Spread tracking dispersal alone would remove growth.

The model and the theorem

For density $u(x,t)$ in one space dimension, with growth rate r and diffusivity D,

$$u_t = D\,u_{xx} + r\,u(1-u), \qquad u_0 \geq 0 \text{ compactly supported},\ u_0 \not\equiv 0. \tag{2.1}$$

> **Theorem 2.1 (Fisher 1937; Kolmogorov–Petrovskii–Piskunov 1937).** Equation (2.1) has travelling-wave solutions $u = U(x-ct)$ with $0 < U < 1$, $U(-\infty) = 1$, $U(+\infty) = 0$ exactly for $c \geq 2\sqrt{Dr}$. For compactly supported initial data the solution converges to the slowest of them: the front position satisfies $x_{\text{front}}(t) = 2\sqrt{Dr}\,t - o(t)$.

Three steps carry the proof, and each is worth keeping separate because each is where a claim usually goes wrong.

(i) No wave can be slower. Substituting $U(x-ct)$ gives $DU'' + cU' + rU(1-U) = 0$. Linearising at the leading edge $U \to 0$ where $U(1-U) \approx U$ gives $DU'' + cU' + rU = 0$, with characteristic

equation $D\lambda^2 + c\lambda + r = 0$. Real roots require $c^2 \geq 4Dr$. For $c < 2\sqrt{Dr}$ the roots are complex and U oscillates through zero, which a density cannot do. Hence $c \geq 2\sqrt{Dr}$.

(ii) The minimal speed is the marginal one. Writing the characteristic roots as $\lambda_\pm = \big(-c \pm \sqrt{c^2 - 4Dr}\big)/(2D)$, a wave of speed c has a leading tail $e^{-\lambda_- x}$ with $\lambda_- < \sqrt{r/D}$; at $c = 2\sqrt{Dr}$ the roots coincide at $\lambda_* = \sqrt{r/D}$. Equivalently: look for solutions of the linearised equation growing like $e^{\lambda(x - vt)}$; this requires $v = D\lambda + r/\lambda$, whose minimum over $\lambda > 0$ is $2\sqrt{Dr}$ at $\lambda = \lambda_*$. The speed is a minimum over decay rates, which is why it is a *product*: D and r enter symmetrically.

(iii) The minimum is the one selected. Compact initial data have a sharp enough tail that the front is "pulled" by its leading edge, and comparison with a super-solution (Aronson–Weinberger) gives both bounds. The classical statement, and the correction to the position, are in the classical references [3], [4], [5] and [6].

> **Theorem 2.2 (Bramson 1983).** For (2.1) with compactly supported data, the front position has the asymptotic form
>
> $$x_{\text{front}}(t) = 2\sqrt{Dr}\, t - \frac{3}{2}\sqrt{\frac{D}{r}}\, \ln t + O(1).$$

This matters for measurement: any finite-time estimate of the speed from a step-like start is biased *low*, by an amount that shrinks only logarithmically. The bias is not a numerical artefact, and it is large when D/r is large. The fit below handles it, and the fitted coefficient is a free check on the whole picture.

Why the ratio cannot be the answer

> **Proposition 2.3 (no function of r/D alone can give the speed).** If v were a function of r/D alone, then replacing (D, r) by $(\lambda D, \lambda r)$ — which leaves r/D unchanged — would leave v unchanged. But by Theorem 2.1, $v = 2\sqrt{Dr}$ becomes $\lambda \cdot 2\sqrt{Dr}$. Hence no such function exists.

The same argument in one line: the equation has the scaling symmetry $x \mapsto \sqrt{D/r}\, y$, $t \mapsto \tau/r$, under which (2.1) becomes the equation with $D = r = 1$. Every solution is a rescaling of one solution, so the speed is $\sqrt{Dr}$ times a universal constant. The ratio r/D is not invariant under that rescaling and therefore cannot determine the speed.

Two numerical consequences, both measured below:

- **Same ratio, different speed.** $(D, r) = (1, 1)$ and $(4, 4)$ both have $r/D = 1$. Measured speeds: 1.988 and 7.942 — a factor of four.
- **Reversed ranking.** $(D, r) = (9,\ 0.25)$ has $r/D = 0.028$; $(D, r) = (0.111,\ 9)$ has $r/D = 81$, nearly 3000 times larger. Measured speeds: 2.873 and 2.020. The ratio ranks the second invader three orders of magnitude faster; the invader it ranks slower is in fact faster.

The qualitative sentence in the claim — "a fast-growing invader with short dispersal can outrun a slow one with long dispersal" — is true under the exact rule, but for a different reason: growth and dispersal are *multiplicatively interchangeable*. A factor 4 in growth and a factor 4 in dispersal each multiply the

speed by 2. The claim's ratio makes them linearly interchangeable, which is a different (and false) statement.

Where the constant 2 comes from, and when it does not apply

The constant is not universal. Take the Allee-type growth $g(u) = r\,u(1-u)(u-a)$ with $0 < a < \tfrac{1}{2}$, where low density reproduces poorly.

> **Proposition 2.4 (bistable front).** For $u_t = D u_{xx} + r\,u(1-u)(u-a)$ the profile $U(\xi) = \big(1 + e^{\kappa\xi}\big)^{-1}$ is an exact travelling wave with $\kappa = \sqrt{r/2D}$ and
>
> $$c = \frac{1-2a}{\sqrt{2}}\sqrt{Dr}.$$

Proof. With $U' = -\kappa U(1-U)$ and $U'' = \kappa^2 U(1-U)(1-2U)$, substitution into $DU'' + cU' + rU(1-U)(U-a) = 0$ and division by $U(1-U)$ leaves

$$D\kappa^2(1-2U) - c\kappa + r(U-a) = 0.$$

This is affine in U, so it holds for all U iff both coefficients vanish: $r - 2D\kappa^2 = 0$ and $D\kappa^2 - c\kappa - ra = 0$. The first gives $\kappa = \sqrt{r/2D}$, the second $c = (D\kappa^2 - ra)/\kappa = \frac{1-2a}{\sqrt{2}}\sqrt{Dr}$. $\square$

At $a = 0.1$ this predicts $c = 0.566\sqrt{Dr}$ against the KPP value $2\sqrt{Dr}$; the simulation measures 0.565 for $D = r = 1$. The $\sqrt{Dr}$ scaling survives, the constant does not. A claim of the form "speed is set by a combination of growth and dispersal" is therefore safe; a claim of the form "speed is $2\sqrt{Dr}$" is only true inside the KPP hypotheses — concave growth, homogeneous habitat, no advection.

Numbers

Equation (2.1) was integrated with a semi-implicit scheme (implicit diffusion, explicit reaction, zero-flux boundaries, $\Delta x = 0.08$, $\Delta t = 0.02$). The front position is interpolated at $u = 0.5$ and fitted two ways: a straight line, and the Bramson form $x = at + b\ln t + c$. The second slope a is the estimate used in the last column.

Table 2. Ten simulations. The Bramson-corrected speeds reproduce $2\sqrt{Dr}$ to within 1.2% everywhere. The fitted logarithmic coefficient matches $-\frac{3}{2}\sqrt{D/r}$ in seven of ten cases; the three misses are the runs where D/r is large enough that $\ln t$ is poorly resolved in the available window. Across cases, $R^2 = 0.082$ for speed against r/D and $R^2 = 0.999$ for speed against $\sqrt{Dr}$.

D	r	r/D	$\sqrt{Dr}$	$2\sqrt{Dr}$	slope	Bramson slope	error	fitted b	$-\frac{3}{2}\sqrt{D/r}$
1	1	1	1	2	1.988	2.001	0.03%	−1.47	−1.50
4	4	1	4	8	7.942	7.995	0.06%	−1.54	−1.50
4	1	0.25	2	4	3.949	4.003	0.07%	−3.12	−3.00
1	4	4	2	4	3.987	4.000	0.00%	−0.76	−0.75
0.25	4	16	1	2	2.002	2.005	0.27%	−0.38	−0.38
2	2	1	2	4	3.975	4.000	0.01%	−1.49	−1.50
0.5	8	16	2	4	3.993	4.000	0.00%	−0.41	−0.38
9	0.25	0.028	1.5	3	2.873	3.021	0.69%	−11.4	−9.0
0.111	9	81.1	1.0	2	2.020	2.022	1.13%	−0.19	−0.17
6.25	0.0992	0.016	0.79	1.57	1.483	1.591	1.00%	−15.7	−11.9

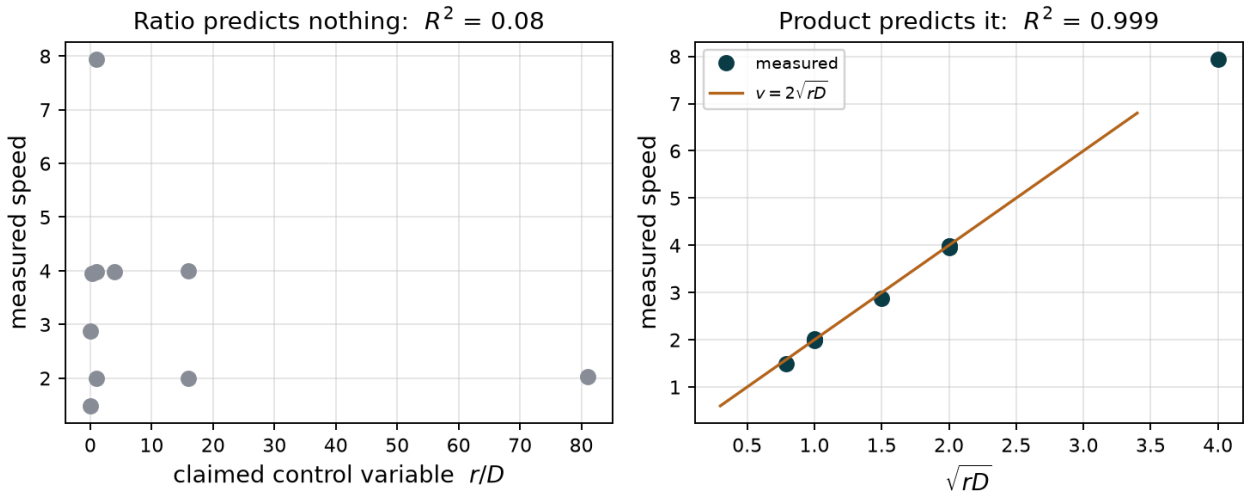


Figure 2. Left: measured front speed against the claimed control variable r/D – no structure, and the two fastest runs sit at opposite ends of the axis. Right: the same ten speeds against $\sqrt{Dr}$, lying on the line $v = 2\sqrt{Dr}$ (coefficient of determination 0.999 against 0.08).

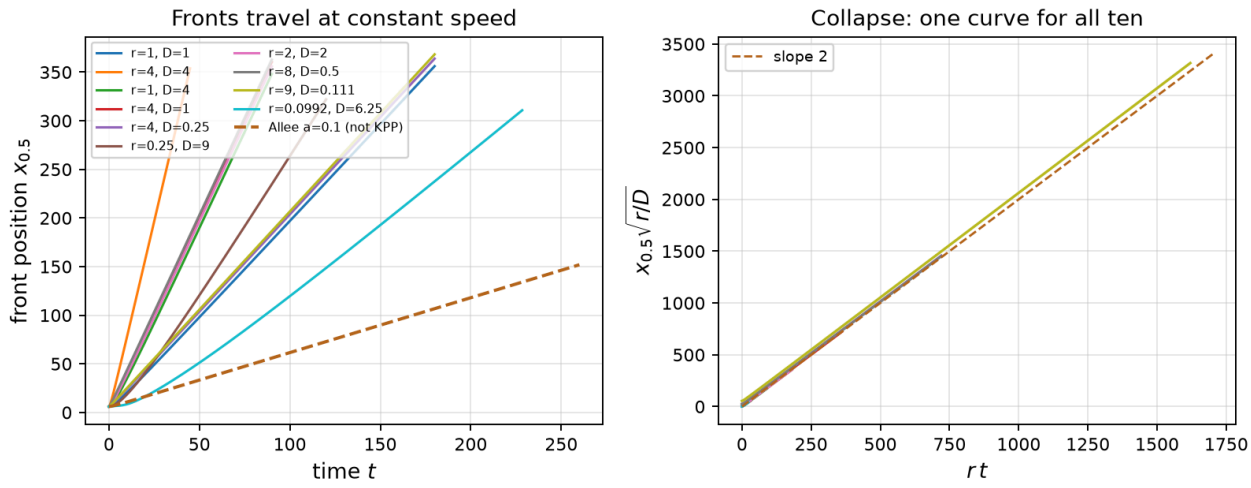


Figure 3. Left: front position against time for all ten parameter pairs, including the Allee case (dashed), which is slower for the same D and r. Right: the same runs with x scaled by $\sqrt{r/D}$ and t scaled by r. Every curve lands on one line of slope 2: the solution family is a single curve rescaled by $\sqrt{Dr}$, which is why no function of r/D alone can describe it.

3 · Osmotic minimum work

The claim

> **Desalination has a floor.** The energy needed to separate salt from water is bounded below by the ratio of osmotic pressure to recovery, so higher recovery costs disproportionately more.
>
> *Mechanism.* The minimum work is the free energy of mixing, which diverges as the concentrate approaches saturation.
>
> *Prediction.* Specific energy should rise steeply with recovery.
>
> *Refutation.* Energy independent of recovery would remove the thermodynamic floor.

The mechanism sentence is the correct one — the minimum work *is* a free energy of mixing. The formula in the first sentence is not, and the two can be compared directly.

Setup

Feed of volume V_f at salt concentration c_f, ideal solution with van 't Hoff osmotic pressure $\pi(c) = \nu RTc$ ($\nu = 2$ for NaCl). Recover a permeate volume $V_p = rV_f$ of pure water at recovery $r \in (0,1)$. Salt is conserved, so while permeate volume V has been drawn the retentate concentration is

$$c(V) = \frac{c_f}{1 - V/V_f}, \qquad c_r = c(V_p) = \frac{c_f}{1-r}. \tag{3.1}$$

Proposition 3.1 (minimum work). The reversible work to produce permeate volume V_p is

$$W(r) = \pi_f V_f \big[-\ln(1-r) \big], \qquad \pi_f = \nu R T c_f,$$

and the specific work per unit product volume is

$$w(r) = \frac{W(r)}{r V_f} = \pi_f \, \frac{-\ln(1-r)}{r}. \tag{3.2}$$

Proof (two independent routes). *Osmotic route.* To push $\mathrm{d}V$ of water out of a retentate at concentration $c(V)$ requires $\mathrm{d}W = \pi(c)\,\mathrm{d}V$. Then

$$W = \int_0^{r V_f} \frac{\pi_f}{1 - V/V_f}\,\mathrm{d}V = \pi_f V_f \int_0^{r} \frac{\mathrm{d}x}{1-x} = \pi_f V_f \big[-\ln(1-r) \big].$$

Free-energy route. For an ideal solution the Gibbs energy per unit volume is $g(c) = \nu R T c \ln c$ (up to a constant, which cancels). Before separation: $V_f\, g(c_f)$. After: $(1-r)V_f\, g(c_r) + r V_f\, g(0)$, and $g(0) = 0$. Hence

$$\Delta G = \nu R T V_f \Big[(1-r) c_r \ln c_r - c_f \ln c_f \Big] = \nu R T V_f c_f \ln \frac{c_r}{c_f} = \pi_f V_f \big[-\ln(1-r) \big],$$

using $(1-r)c_r = c_f$ and $c_r/c_f = 1/(1-r)$. The two routes agree, as they must; numerically they agree to 6×10^{-15} relative error (`data/t17-checks.json`). $\square$

Shape of the curve

Theorem 3.2 (monotone and convex). On $(0,1)$, w in (3.2) is smooth, strictly increasing and strictly convex, with $w(0^+) = \pi_f$ and $w(r) \to +\infty$ as $r \uparrow 1$.

Proof. Expand the logarithm:

$$\frac{-\ln(1-r)}{r} = \frac{1}{r}\sum_{k \ge 1} \frac{r^k}{k} = \sum_{k \ge 0} \frac{r^k}{k+1} = 1 + \frac{r}{2} + \frac{r^2}{3} + \frac{r^3}{4} + \cdots$$

Every coefficient is positive, so on $[0,1)$ the series has positive first and second derivatives term by term: $w' > 0$, $w'' > 0$. The value at $r = 0$ is the first coefficient, π_f. Divergence follows from $-\ln(1-r) \to \infty$ with $r \to 1$. $\square$

A polynomial with all-positive coefficients is the cleanest way to see both monotonicity and convexity at once, and it also shows the growth is only logarithmic in $1/(1-r)$, not a pole.

The claimed floor is not a floor

The claim's bound is π_f / r. Comparing with (3.2),

$$w(r) - \frac{\pi_f}{r} = \frac{\pi_f}{r}\Big[-\ln(1-r) - 1 \Big],$$

so the sign is that of $-\ln(1-r) - 1$, which changes at $e^{-1} = 1 - r$:

$$r^{*} = 1 - e^{-1} = 0.63212\ldots$$

Proposition 3.3. For $r < r^{*}$ the claimed bound π_f / r exceeds the true minimum work; the excess grows without bound as $r \to 0$. For $r > r^{*}$ it falls below the true minimum. At r^{*} the two coincide.

So “bounded below by the ratio of osmotic pressure to recovery” is true only above 63% recovery, and even there the bound is not tight — it is a second, unrelated curve that happens to cross the first. At $r = 0.3$ the claim overstates the minimum by a factor 2.8; at $r = 0.05$ by a factor 19; at $r = 0.95$ it *understates* it by a factor 3.

What the claim gets right is the direction and the convexity: higher recovery does cost disproportionately more, and the marginal cost diverges. What it gets wrong is the size, and the functional form — the exact curve is logarithmic in the brine concentration factor, not hyperbolic in recovery.

Numbers

35 g/L NaCl, $\nu = 2$, $T = 298.15$ K gives $c_f = 0.5989$ mol/L and $\pi_f = \nu R T c_f = 2.97$ MPa $= 0.825$ kWh m^{-3}.

Table 3. Ideal-solution minimum work for 35 g/L NaCl at 25 °C. Recovery 0.63 is the crossing point $1 - 1/e$, where the claimed floor and the exact curve meet. Full table: `data/t17-recovery.csv`.

recovery	brine c_r (mol/L)	brine π_r (bar)	exact w (kWh/m³)	claimed π_f / r	total per m³ feed
0.05	0.630	31.3	0.846	16.50	0.042
0.20	0.749	37.1	0.920	4.12	0.184
0.30	0.856	42.4	0.981	2.75	0.294
0.50	1.198	59.4	1.143	1.65	0.572
0.632	1.628	80.7	1.306	1.306	0.825
0.75	2.396	118.8	1.525	1.10	1.143
0.85	3.993	198.0	1.841	0.970	1.565
0.90	5.989	296.9	2.110	0.917	1.899
0.95	11.978	593.9	2.601	0.868	2.471

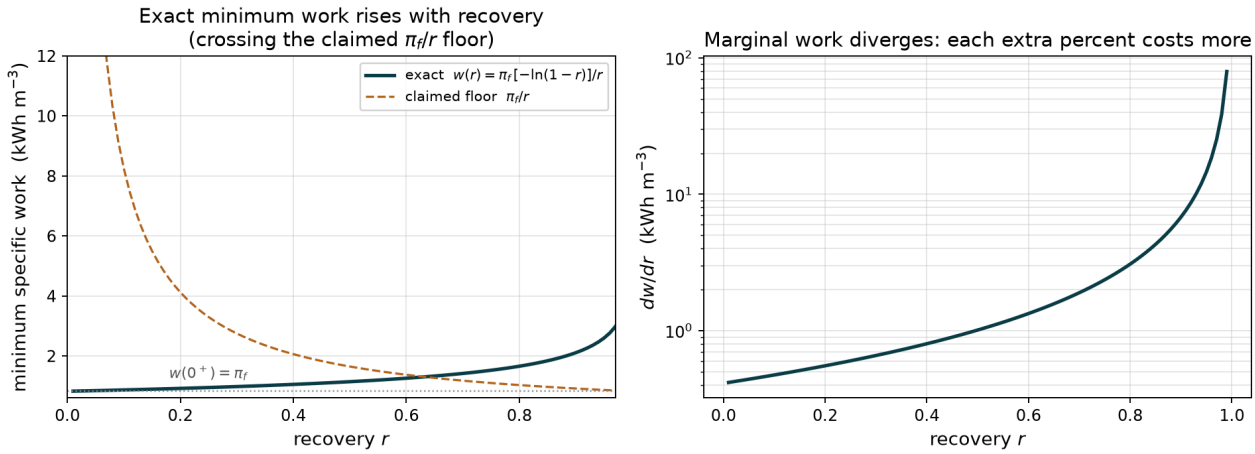


Figure 4. Left: exact specific work (solid) against recovery, with the claimed π_f/r bound (dashed). The two curves cross at $r = 1 - 1/e$; neither dominates the other. Right: the marginal cost $\mathrm{d}w/\mathrm{d}r$ on a log scale – each additional percent of recovery costs more than the last, and the marginal cost diverges as $r \to 1$.

The practical reading: at 50% recovery the ideal minimum is 1.14 kWh per cubic metre of product. Reported plant consumptions for seawater reverse osmosis cluster around 2–4 kWh m^{-3} [9], i.e. two to three times the ideal value, which is what a second-law efficiency of 30–55% looks like. Two things push the real minimum *down* relative to Table 3 and two push the real consumption *up*: activity coefficients make seawater's osmotic pressure lower than the van 't Hoff value (≈2.4–2.6 MPa at 25 °C, so the ideal curve is perhaps 15% high), while concentration polarisation, pump and turbine inefficiencies, pre-treatment and the finite membrane permeability all add. The bound is a bound on the process, not a forecast of the device.

4 · One shape, three times

The three claims fail in the same way, which is the reason for reporting them together.

	#16 per-bit floor	#10 invasion speed	#17 osmotic work
Mechanism in the claim	capacity per hertz sets a minimum	growth and movement combine	free energy of mixing
Is the mechanism right?	yes	yes	yes
Formula in the claim	floor ∼ noise/bandwidth; optimum at finite B	speed ∝ r/D	floor = π_f/r
Exact result	$E_b/N_0 > \ln 2$, monotone, no optimum	$v = 2\sqrt{Dr}$	$w = \pi_f[-\ln(1-r)]/r$, convex
What survives	the floor and its monotonicity	growth and dispersal both matter	direction and convexity
What dies	the constant, the bandwidth dependence, the optimum	the ratio, and every ranking built on it	the size of the bound, and the bound itself
Cheapest check	one derivative, or a dimensional check	one rescaling of the equation	one expansion of a logarithm

Two of the three claims are internally inconsistent in the same way: the stated *prediction* is contradicted by the theorem that supplies the stated *mechanism*. In #16 the mechanism is Shannon capacity; the capacity theorem gives a monotone curve, not an optimum. In #17 the mechanism is free energy of mixing; that free energy is $-\ln(1-r)$, not $1/r$. In #10 the mechanism sentence ("whichever limits produces the slower front") is actually a correct description of a minimum-speed selection principle — and the minimum is a geometric mean, which is exactly what a ratio cannot express.

5 · What this does and does not show

Does. It gives a proof, a numerical replication and a verdict for three questions, with the exact constants. It fixes three statements that can be quoted directly: the floor $E_b/N_0 > \ln 2$ with no optimum, the speed $v = 2\sqrt{Dr}$, and the work $w = \pi_f[-\ln(1-r)]/r$ with a crossing at $1 - 1/e$.

Does not. It does not show that the notebook's *questions* are wrong — all three are good questions, and two of the three mechanism sentences are correct. It does not test the theorems against nature: no link, organism or plant is measured here. It does not claim novelty for any of the three results, which are classical or published; the citations are in the reference list and predate this report by decades.

The three idealisations are worth naming, because each one is where a real test would begin. The channel model assumes Gaussian noise, infinite coding blocklength and no circuit power. The invasion model assumes a homogeneous habitat, no advection, and growth that is concave near zero. The separation model assumes an ideal solution, no concentration polarisation, and reversible operation. Everything above is exact *inside* those assumptions and says nothing outside them.

6 · Reproducing this

```
cd code
./run_all.sh          # runs all three checks, writes ../data and ../figures
```

Each script is standalone and prints its checks to stdout. The three are:

Script	What it checks	Runtime
`code/t16_energy_floor.py`	(1.1) against direct capacity evaluation; monotonicity on a 4001-point grid; SymPy derivative and limits; asymptotic law (1.2)	6 s
`code/t10_fisher_kpp.py`	ten PDE integrations; Bramson fit; Allee case; collapse	100 s
`code/t17_osmotic_work.py`	quadrature vs closed form; free-energy route; series coefficients; crossing point	3 s

Machine checks, all passing, are collected in `verification/VERIFICATION.md` and reproduced in `data/*-checks.json`. Versions: Python 3.14, NumPy 2.5.4, SciPy 1.18.1, SymPy 1.14.0, Matplotlib 3.11.2.

References

[1] Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal* **27**, 379–423. doi:10.1002/j.1538-7305.1948.tb01338.x

[2] Cover, T. M. and Thomas, J. A. (2006). *Elements of Information Theory*, 2nd ed. Wiley. (The bound $E_b/N_0 \geq \ln 2$ and the bandwidth–energy trade-off.)

[3] Fisher, R. A. (1937). The wave of advance of advantageous genes. *Annals of Eugenics* **7**, 355–369. doi:10.1111/j.1469-1809.1937.tb02153.x

[4] Kolmogorov, A., Petrovskii, I. and Piskunov, N. (1937). A study of the diffusion equation with increase in the amount of substance, and its application to a biological problem. *Bull. Moscow Univ. Math. Mech.* **1**, 1–25. (English translation in *Selected Works of A. N. Kolmogorov I*, Kluwer, 1991, 248–270.)

[5] Aronson, D. G. and Weinberger, H. F. (1978). Multidimensional nonlinear diffusion arising in population genetics. *Advances in Mathematics* **30**, 33–76. doi:10.1016/0001-8708(78)90130-5

[6] Bramson, M. (1983). Convergence of solutions of the Kolmogorov equation to travelling waves. *Memoirs of the American Mathematical Society* **44**(285). doi:10.1090/memo/0285

[7] van 't Hoff, J. H. (1887). The function of osmotic pressure in the analogy between solutions and gases. *Proceedings of the Physical Society of London* **9**, 307–334. doi:10.1088/1478-7814/9/1/344

[8] Wang, L., Violet, C., DuChanois, R. M. and Elimelech, M. (2020). Derivation of the theoretical minimum energy of separation of desalination processes. *Journal of Chemical Education* **97**, 4361–4369. doi:10.1021/acs.jchemed.0c01194 — the same integral and the same closed form, with activity-coefficient corrections.

[9] Elimelech, M. and Phillip, W. A. (2011). The future of seawater desalination: energy, technology, and the environment. *Science* **333**, 712–717. doi:10.1126/science.1200488

Source questions: 16 — transmission energy per bit · 10 — invasion speed · 17 — desalination minimum work, quoted verbatim from the notebook. Worked independently; not peer reviewed. AI assistance: the proofs, code, checks and text were produced by DeepSeek on 31 October 2026; the source questions were written and published by the site owner.