Clay Shooting:
Equilibrium Effort Solving Outstanding Problems

Peter Cotton

October 31, 2026

Abstract

The Clay Mathematics Institute posts one million dollars on each Millennium problem. We ask what solving one is worth to the person attacking it, and what the purse contributes to the effort. A Brownian attempt run between a target and an abandonment floor buys success probability p at expected cost $\zeta^2 p/(1-p)$, where ζ^2 is the cost of diffusing across the remaining distance, and this price list recovers the reward from the cost of an attempt, the number of attackers and the observed rate at which problems fall. The level of the reward is not identified by one solution in twenty-six years: the point estimate is \$96M, and the hazard interval carries it from \$18M to \$3.8B. The surplus over the threshold at which anyone would start is identified, at about twice the cost of one attempt, and the purse is 55% of it at every point of the interval.

1 What a Millennium problem is worth

The Clay Mathematics Institute posted one million dollars on each of seven problems in May 2000 (Clay Mathematics Institute, 2000). One has been solved, in 2002–2003, and the person entitled to the purse declined it. That refusal says the reward is not the money in escrow. It does not say what the reward is.

This paper computes it. Suppose two people are at work on a given problem at any time, an attempt takes about three years, a senior mathematician costs about three hundred thousand dollars a year all in, and problems fall at the observed rate of one per hundred and fifty-nine problem-years. Then solving a Millennium problem is worth about \$96M to the person attacking it, and the purse is 1.0% of that.

The inversion rests on a price list for persistence. An attempt run between a target at distance d and an abandonment floor buys success probability p at expected cost

$$k(p) = \zeta^2 \frac{p}{1-p}, \qquad \zeta^2 = \frac{c\,d^2}{\sigma^2}, \tag{1}$$

where ζ^2 is the cost rate times the diffusive time to cover the remaining distance. Optimizing $Ap - k(p)$ against a reward A gives $p = 1 - \zeta/\sqrt{A}$ and $k = p(1-p)A$, which inverts to $A = k/(p(1-p))$. The hazard supplies p through an accounting identity, and what comes out is

$$A = \frac{cN}{h(1-p)}, \qquad p = \frac{hT}{N}, \tag{2}$$

in the cost rate c, the number N of concurrent attempts, and the hazard h. This is Proposition 2. The duration T of an attempt cancels, so how long people persist does not enter the level, only how many of them there are and how fast problems fall.

One event does not identify that level. The exact Poisson interval for the hazard is a factor of two hundred wide, and it carries A from \$18M to \$3.8B. Anyone quoting \$96M is quoting the midpoint of an interval that contains almost anything.

The surplus is a different matter, and it is the paper's point. Optimal sizing forces

$$A - \zeta^2 = k\,\frac{2-p}{1-p} \;\geq\; 2k, \tag{3}$$

so the reward exceeds the threshold at which anyone would start by about twice the cost of a single attempt, however large the reward. Here T survives and N and h drop out. The surplus is \$1.8M, the purse is 55% of it, and that share moves from 55.5% to 54.1% across the entire hazard interval. Deleting the purse from the reward costs 55% of the equilibrium solvers.

So the purse is a rounding error against the value of a Clay problem and a majority stake in the surplus that makes attacking one worthwhile. The two ratios differ by a factor of fifty-three because the level and the surplus depend on disjoint sets of parameters. Dividing the purse by the value of a solution answers the wrong question.

The rest of the paper proves (1) to (3) and then puts them to work. Section 2 proves the price list, whose last clause is the useful one: it holds at every floor, optimal or not, at any discount rate, so it identifies ζ^2 from an observed attempt without assuming the attempt was well run. Section 3 does the Millennium inversion. Section 4 closes the competitive loop, where the reward is shared rather than won outright and the equilibrium is a scalar root; in a large pool it collapses to $G(y^*) = V_e/V$ with $G(y) = (1 - e^{-y})/y$, and competition bids the effective reward down to break-even exactly. Section 6 designs the purse: a funder splitting escrowed capital posts 1.7564 times break-even on each problem whatever the capital, and solves 71.5% of them.

Scope. Sections 2–6 treat the lottery allocation: every entrant pays its attempt to completion and one winner is drawn among the successes. The first-solution race, in which costs cease when a rival announces and the rival hazard varies with calendar time, is the companion note's coupled system (Cotton, 2026b), and it is not solved here. The constants below are lottery constants.

2 The price list

Fix one researcher and one approach. While working, the researcher observes $dX_t = \sigma\, dW_t$ with $\sigma > 0$, starting at x_0. First passage to a target $b > x_0$ is a solution, working costs $c > 0$ per unit time, and abandonment at a floor $a < x_0$ is irreversible. Write $d = b - x_0$ for the distance still to travel, $L = b - a$ for the continuation width, and

$$\zeta^2 = \frac{c\,d^2}{\sigma^2} \tag{4}$$

for the *hardness*, the cost rate times the diffusive time d^2/σ^2 to cover the remaining distance.

Proposition 1 (Persistence is bought at a convex price). For any width $L > d$,

$$p = 1 - \frac{d}{L}, \qquad k = \frac{c\,d\,(L-d)}{\sigma^2} = \zeta^2\,\frac{p}{1-p}. \tag{5}$$

The map $k(\cdot)$ is strictly increasing and strictly convex on $[0,1)$, with $k(0)=0$ and $k(p)\to\infty$ as $p\to 1$, and an observed pair (k,p) determines

$$\zeta^2 = k\,\frac{1-p}{p} \tag{6}$$

with no reference to the prize, the discount rate, the rival hazard, or whether L was chosen well.

Proof. Optional stopping applied to X and to $X_t^2-\sigma^2 t$ at $\tau=\tau_a\wedge\tau_b$, legitimate since X is bounded on (a,b) and $\mathbb{E}[\tau]<\infty$, gives the gambler's-ruin identities $p=(x_0-a)/(b-a)=1-d/L$ and $\mathbb{E}[\tau]=(x_0-a)(b-x_0)/\sigma^2=d(L-d)/\sigma^2$, so $k=c\mathbb{E}[\tau]$. Substituting $L=d/(1-p)$ gives $k=c\,d^2p/\big(\sigma^2(1-p)\big)$. Convexity and the limits follow. □

The last clause of Proposition 1 is what makes the rest work. Standard race models specify a cost-to-hazard technology and then compute; here the technology is a theorem about Brownian motion, and it holds off the optimal path as well as on it.

Corollary 1 (Optimal attempt size). Let $A>0$ be the expected reward conditional on solving. The attempt maximizing $Ap-k(p)$ has

$$p=\Big(1-\frac{\zeta}{\sqrt{A}}\Big)_+,\qquad k=p(1-p)A,\qquad \textit{value}=Ap^2=\big(\sqrt{A}-\zeta\big)_+^2, \tag{7}$$

and is worth mounting exactly when $A>\zeta^2$. Expected spend never exceeds $A/4$, with equality only at $p=1/2$.

Proof. $Ap-\zeta^2p/(1-p)$ has derivative $A-\zeta^2/(1-p)^2$, decreasing in p and vanishing at $(1-p)^2=\zeta^2/A$. That stationary point lies in $(0,1)$ exactly when $A>\zeta^2$; otherwise the derivative is negative throughout and $p=0$. Substituting $1-p=\zeta/\sqrt{A}$ gives $k=\zeta p\sqrt{A}=p(1-p)A$ and $Ap-k=Ap^2=(\sqrt{A}-\zeta)^2$. Finally $k/A=p(1-p)\le 1/4$. □

The threshold ζ^2 has a plain reading: before anyone starts, the reward must cover the cost of an undirected diffusive crossing of the remaining gap. The bound $k\le A/4$ says that the expensive attempts are the genuinely uncertain ones; cheap long shots and near-certainties both spend little.

Corollary 2 (The surplus is a small multiple of one attempt). If the attempt is optimally sized,

$$A-\zeta^2 = k\,\frac{2-p}{1-p}\;\ge\;2k, \tag{8}$$

with the bound approached as $p\downarrow 0$ and equal to $3k$ at $p=1/2$.

Proof. $A=\zeta^2/(1-p)^2$, so $A-\zeta^2=\zeta^2p(2-p)/(1-p)^2$, and $k=\zeta^2p/(1-p)$. □

Corollary 2 is the paper's central observation. However large the reward, if attempts are long shots then the reward only just clears the threshold ζ^2, and the margin is roughly twice the cost of one attempt. The level of the reward and the surplus it generates are different quantities, and only the second governs behavior. Section 3 rests on that distinction.

A verification. A seeded Euler simulation, given only the boundary location and the increment law, reproduces (7). At $A = 100$, $c = \sigma = 1$ and $d = 4$, so $\zeta = 4$, $L = 10$ and $p = 0.6$, fifty thousand paths at $dt = 2 \times 10^{-3}$ give $\hat{p} = 0.6051$ (s.e. 0.0022) and $\hat{k} = 24.05$ (s.e. 0.092) against $p = 0.6$ and $k = 24$. The upward bias in $\hat{p}$ is Euler overshoot at the upper boundary and shrinks with dt. The price identity (5) is checked against k computed directly from the width at five widths, with residuals at 10^{-14}.

3 The implied reward

Proposition 1 identifies ζ^2 from an attempt's cost and success rate through (6). Corollary 1 inverts an optimally sized attempt to the reward it was chasing. What is missing is p, and the historical record supplies it.

Proposition 2 (The implied reward). Let N attempts run per problem at all times, each of mean duration T at cost rate c, each succeeding with probability p and sized optimally against a reward A. If problems are solved at hazard h per problem-year, then

$$p = \frac{hT}{N}, \qquad A = \frac{cN}{h(1-p)}, \qquad \zeta^2 = \frac{cN(1-p)}{h}, \qquad A - \zeta^2 = cT\,\frac{2-p}{1-p}. \tag{9}$$

The level A is free of T up to the factor $(1-p)$; the surplus $A - \zeta^2$ is free of N and h, and tends to $2cT$ as $p \downarrow 0$.

Proof. Completed attempts arrive at rate N/T per problem-year and each solves with probability p, so $h = pN/T$, which is the first identity. One attempt costs $k = cT$. Corollary 1 inverted gives $A = k/(p(1-p))$ and Proposition 1 gives $\zeta^2 = k(1-p)/p$. Substituting $k = cT$ and $p = hT/N$,

$$A = \frac{cT}{(hT/N)(1-p)} = \frac{cN}{h(1-p)}, \qquad \zeta^2 = \frac{cT(1-p)N}{hT} = \frac{cN(1-p)}{h},$$

and (8) gives $A - \zeta^2 = k(2-p)/(1-p)$ with $k = cT$. The limits are immediate. □

The two scalings are disjoint, and everything that follows uses that. The level is cN/h: what a solution is worth rises with the cost of the people chasing it and with how many are chasing, and falls with how often problems fall. The surplus is $2cT$: what makes chasing worthwhile rises with the cost of one attempt and nothing else. A funder who moves the purse moves the second and leaves the first alone.

The Millennium numbers. One solution since May 2000, with the solved problem exposed only until 2003, gives an exposure of $6 \times 26 + 3 = 159$ problem-years and $\hat{h} = 0.0063$ per problem-year. At $c = \$300\text{k}$ per year, $N = 2$ and $T = 3$ years, (9) returns $p = 0.0094$, an attempt cost of \$0.9M, $\zeta^2 = \$94.5\text{M}$ and

$$A = \$96.3\text{M}, \qquad A - \zeta^2 = \$1.81\text{M}, \tag{10}$$

against $2k = \$1.80\text{M}$. The purse is 1.04% of the level and 55.3% of the surplus, a ratio of fifty-three between the two shares. Dilution is negligible at these odds: the reward before sharing is \$96.8M, since $\mathbb{E}[(1+K)^{-1}] = 0.9953$ with two attackers at $p = 0.0094$.

	implied level A, \$M		
cost rate	$N = 1$	$N = 2$	$N = 5$
\$200k	32	64	160
\$300k	49	96	239
\$400k	65	128	319

The level is not identified and the surplus is. The exact Poisson interval for one event in 159 problem-years is $[1.6 \times 10^{-4}, 3.5 \times 10^{-2}]$, which is to say a solution every 29 to 6280 years. Since A is proportional to $1/h$, that interval is inherited whole.

h	A, \$M	purse/A	surplus, \$M	purse/surplus
1.6×10^{-4}	3769	0.027%	1.800	55.5%
6.3×10^{-3}	96.3	1.04%	1.809	55.3%
3.5×10^{-2}	18.1	5.53%	1.850	54.1%

The level moves by a factor of two hundred and the purse's share of the surplus moves by one and a half percentage points. This is not a coincidence: h enters A and ζ^2 the same way and cancels in the difference, exactly as (9) says. The quantity everyone would quote is the one the record cannot pin down, and the quantity that governs behavior is the one it can.

The purse's share of the effort. Deleting the purse from the reward and re-solving the large-pool equilibrium of Proposition 4 takes expected solvers from 0.0380 to 0.0171, a fall of 55.1%, tracking the purse's 55.3% share of the surplus. Near threshold $y^* \approx 2(V/V_e - 1)$, so solvers are proportional to the excess of the reward over break-even and the purse's share of that excess is its share of the effort. Across the hazard interval the fall runs 53% to 56%.

Two of the four inputs come from outside the record. The cost rate c and the number N of concurrent attempts are stipulated, and A is proportional to both, so the level inherits their error too. The duration T is stipulated and does not matter for the level. The surplus needs only c and T, and the cost of a serious attempt is the quantity to go and measure.

Setting $V_e = \zeta^2$ assumes no cost of entering beyond the flow cost of working. A setup cost raises V_e to $(\zeta + \sqrt{K})^2$ and shrinks the excess the purse is measured against, so every share above is a lower bound on the purse's.

Abandonment and rivalry. The remaining Millennium question is not what the reward is but why anyone abandons it. That question belongs to the first-solution race of the companion note rather than to the lottery, because abandonment is driven by a rival's solution destroying the option, not merely diluting it. In the race an attempt is discounted at $\rho = r + h$, the interest rate plus the rival hazard, so rivalry accounts for a share $h/(r + h)$ of the impatience to finish. Two adjustments both shrink that share. The observed hazard is for the problem as a whole and includes the focal researcher's own attempts, so with N symmetric attackers the rival hazard is only $(N - 1)/N$ of it. And any positive discount rate enlarges the denominator.

r	$h/(r+h)$ at $\hat{h}$	rival share, $N=3$	$h/(r+h)$ at upper h
0.03	0.173	0.123	0.539
0.05	0.112	0.077	0.412
0.10	0.059	0.040	0.260

A solution hazard of order 10^{-2} per year is small beside any plausible discount rate, so rivalry is unlikely to be the main reason anyone abandons a Millennium problem. They stop because the shells cost money and the years pass.

The fast-solution regime. All the interesting variation sits in $\zeta = d\sqrt{c}/\sigma$. An attempt run at machine cost lowers c, and a genuinely broader search raises σ; both lower ζ, which lowers the break-even purse quadratically and raises expected solvers. If that compresses expected time-to-solution from decades to a few years, h rises by two orders of magnitude, and by (9) the implied level falls in proportion while the rivalry share moves from under a tenth to above four fifths. Only then is the purse a clay pigeon, thrown up for several guns with one hit counting. This is a graded, falsifiable prediction.

4 Equilibrium attempt size

We turn to the contest. There are n researchers, a prize V, and a common hardness ζ^2. Each researcher chooses a width, equivalently a success probability $p_i \in [0,1)$, and pays $k(p_i)$; choosing $p_i = 0$ is staying out. Successes are independent, and if several researchers succeed one is drawn uniformly for the prize.

Conditional on solving, researcher i's expected reward is

$$A_i = V\,\mathbb{E}\Big[\frac{1}{1+K_{-i}}\Big] = V\int_0^1 \prod_{j\neq i}\big(1-p_j+p_j z\big)\,dz, \tag{11}$$

by $1/(m+1) = \int_0^1 z^m dz$, where K_{-i} counts successful rivals. Payoffs are $A_i p_i - k(p_i)$, so Corollary 1 makes each best response $p_i = (1-\zeta/\sqrt{A_i})_+$. Rivals enter through one number. In the symmetric case (11) collapses to

$$A(p) = V\,\frac{1-(1-p)^n}{np}. \tag{12}$$

Proposition 3 (The symmetric equilibrium is a scalar root). A in (12) is strictly decreasing in p on $(0,1]$, with $A(0^+) = V$ and $A(1) = V/n$, and strictly decreasing in n. If $V \le \zeta^2$ the only symmetric equilibrium is $p = 0$. If $V > \zeta^2$ there is a unique symmetric equilibrium $p_n \in (0,1)$, the root of

$$p = 1 - \frac{\zeta}{\sqrt{A(p)}}, \tag{13}$$

it decreases in n, and each researcher earns $A(p_n)p_n^2 > 0$.

Proof. Write $w = 1-p$. Then $\partial_p\big[(1-w^n)/p\big]$ has the sign of $g(p) = npw^{n-1} - 1 + w^n$, and $g(0) = 0$ with $g'(p) = -n(n-1)pw^{n-2} < 0$, so $g < 0$ and A decreases in p. For monotonicity in

n, the difference $(n+1)(1-w^n) - n(1-w^{n+1}) = 1-(n+1)w^n + nw^{n+1}$ vanishes at $w=1$ with derivative $n(n+1)w^{n-1}(w-1) < 0$, hence is positive for $w<1$, giving $A_n > A_{n+1}$.

Set $\Phi(p) = (1-\zeta/\sqrt{A(p)})_+ - p$, continuous and strictly decreasing on $[0,1]$: where the positive part is active it is $1-\zeta/\sqrt{A(p)} - p$ with A strictly decreasing, and where it is clamped it is $-p$. Now $\Phi(0) = 1-\zeta/\sqrt{V}$ and $\Phi(1) = -\zeta\sqrt{n}/\sqrt{V} < 0$. If $V \le \zeta^2$ then $\Phi(0) \le 0$ and $p=0$ is the unique root, and it is an equilibrium, since a deviator faces $A = V \le \zeta^2$ and has best response 0. If $V > \zeta^2$ then $\Phi(0) > 0 > \Phi(1)$ and there is exactly one root, decreasing in n because A_n is. □

Proposition 4 (Large pool). Let $n \to \infty$ with $np \to y$. Then $A \to V\,G(y)$ with $G(y) = (1-e^{-y})/y$, strictly decreasing from 1 to 0, and the equilibrium condition becomes

$$\boxed{G(y^*) = \frac{\zeta^2}{V}} \tag{14}$$

with a unique positive root whenever $V > \zeta^2$. At that root the expected number of solvers is y^*, the solve probability is $1-e^{-y^*}$, the effective reward collapses to $A^* = \zeta^2$, individual effort and payoff vanish, and aggregate effort equals the expected payout,

$$\lim n\,k(p_n) = y^*\zeta^2 = V\big(1-e^{-y^*}\big). \tag{15}$$

With a setup cost $K > 0$ payable on entry,

$$G(y^*) = \frac{V_e}{V}, \quad V_e = \big(\zeta+\sqrt{K}\big)^2, \quad p^* = \frac{\sqrt{K}}{\zeta+\sqrt{K}}, \quad N^* = y^*\,\frac{\zeta+\sqrt{K}}{\sqrt{K}}, \tag{16}$$

where V_e is the purse at which a lone researcher exactly breaks even, $A^* = V_e$, each attempt costs $\zeta\sqrt{K}$, and $N^* \to \infty$ as $K \downarrow 0$.

Proof. $(1-y/n)^n \to e^{-y}$ gives $A \to VG(y)$. Since $p = 1-\zeta/\sqrt{A}$ and $p \to 0$, we need $\zeta/\sqrt{A^*} \to 1$, so $A^* = \zeta^2$ and (14) follows; uniqueness is strict monotonicity of G. Then $k = p(1-p)A \to 0$ and the payoff $Ap^2 \to 0$, while $nk \to y\,A^* = y^*\zeta^2 = V\,y^*G(y^*) = V(1-e^{-y^*})$. With a setup cost, free entry requires $Ap^2 = K$, and substituting $p = 1-\zeta/\sqrt{A}$ gives $(\sqrt{A}-\zeta)^2 = K$, so $\sqrt{A^*} = \zeta+\sqrt{K}$. Then $p^* = 1-\zeta/(\zeta+\sqrt{K})$, $k = p^*(1-p^*)A^* = \zeta\sqrt{K}$ and $N^* = y^*/p^*$. Finally V_e solves $(\sqrt{V}-\zeta)^2 = K$, the lone researcher's break-even condition in Corollary 1. □

V/V_e	y^*	$\mathbb{P}$(solved)	np_n at $n = 10^4$
1.02	0.03974	0.0390	0.03972
1.5	0.87422	0.5828	0.87387
2	1.59362	0.7968	1.59286
4	3.92069	0.9802	3.91742
10	9.99955	0.99995	9.97959

The last column solves Proposition 3 exactly at $n = 10^4$. A deviation certificate backs the fixed points: at three computed equilibria, no point of a two-million-point grid over own success probability beats the equilibrium payoff against equilibrium rivals, the best falling short by 4×10^{-12} or less. The grid does not know the closed-form best response, so the check can fail for the right reason.

Solutions are pinned, attempts are not. Both y^* and the solve probability depend only on V/V_e. The entrant count $N^* = y^*(\zeta + \sqrt{K})/\sqrt{K}$ does not, and it diverges as the setup cost vanishes, since a great many researchers each buy a vanishing probability. A claim about how many people attempt a famous problem is a claim about setup costs.

Rents are competed away exactly. $A^* = V_e$. In a large pool the effective reward equals the purse at which a lone researcher would exactly break even, so every researcher behaves as though facing no competition and a prize of V_e, and aggregate effort equals the expected payout. This is the classical free-entry dissipation result with the cost curve derived rather than posited.

5 The solo width is not a best response

A tempting shortcut fixes attempt lengths at the solo-optimal $L = \sigma\sqrt{V/c}$ and lets researchers choose only whether to enter. The shortcut is wrong, and the size of the error can be computed.

Take $V = 100$, $c = \sigma = 1$, $d = 5$, and two entrants at the solo width $L = 10$, so $p = 1/2$ each. Then $A = V(1 - p/2) = 75$, and the payoff at the solo width is $75 \cdot \frac{1}{2} - 25 = 12.5$. The best response is $L = \sigma\sqrt{A/c} = \sqrt{75} \approx 8.6603$, worth 13.3975, a gain of 7.2%. Competition shortens attempts, because the prize is diluted and a diluted prize does not justify the same persistence.

The shortcut also misreports how many people attempt. Holding $(p, k = p(1-p)V)$ exogenous, any three designated entrants form a pure Nash equilibrium whenever $p < 2 - \sqrt{2} \approx 0.5858$: each incumbent earns $V\big[1 - (1-p)^3\big]/3 - p(1-p)V = Vp^3/3 > 0$, while a fourth would earn $Vp^2\big(-\frac{1}{2} + p - \frac{p^2}{4}\big) < 0$. Two entrants is never a pure equilibrium, since a third always gains $Vp^3/3$ by joining.

The companion note's exact potential selects three entrants on exactly the same range, and not by coincidence. An exact potential moves by the deviator's own payoff, so its maximiser leaves three at the p where the fourth entrant breaks even, which is again $2 - \sqrt{2}$. The committed increments agree with the entrant payoffs to 2×10^{-16}.

The symmetric mixed equilibrium of the same game gives $nq \to 2n/(n-1)$ as $p \downarrow 0$, which is 2.667 at $n = 4$ and 2.041 at $n = 50$, reaching two only in the pool limit. Any headline count from the shortcut is an equilibrium selection in disguise.

6 Purse design

A funder commits escrowed capital B to identical problems with break-even purse V_e, choosing the number M of problems and posting $V = B/M$ on each. Free entry then sets each problem's solve probability by Proposition 4. Write $H(y) = 1 - e^{-y}$, so $G = H/y$.

Proposition 5 (The optimal purse is a fixed multiple of break-even). Under Proposition 4, $V = V_e\, y/H(y)$, so $M = B\, H(y)/(V_e y)$ and the expected number of problems solved is

$$\Psi = M\, H(y) = \frac{B}{V_e} \cdot \frac{H(y)^2}{y}. \tag{17}$$

Over the continuous relaxation of M, Ψ has a unique positive maximizer, characterized by

$$\boxed{e^y = 1 + 2y} \tag{18}$$

at which, exactly,

$$\frac{V^*}{V_e} = y^* + \frac{1}{2}, \qquad \mathbb{P}(\textit{solved}) = \frac{2y^*}{1+2y^*}, \qquad \Psi^* = \frac{B}{V_e} \cdot \frac{H(y^*)^2}{y^*}. \tag{19}$$

Numerically $y^* = 1.25643121$, $V^* = 1.75643121\,V_e$, $\mathbb{P}(\textit{solved}) = 0.71533186$, *and* $\Psi^* = 0.40726438\,B/V_e$ *on* $M^* = 0.56933627\,B/V_e$ *problems.*

Proof. $G(y) = V_e/V$ gives $V = V_e y/H(y)$, hence M and Ψ. Maximizing $\log \Psi = 2\log H - \log y$ up to a constant gives $2H'/H = 1/y$, and with $H' = e^{-y}$ this is $2ye^{-y} = 1 - e^{-y}$, that is $e^{-y}(2y+1) = 1$, that is $e^y = 1 + 2y$. The function $e^y - 1 - 2y$ is convex with value 0 and derivative -1 at the origin and tends to infinity, so it has exactly one positive zero. At that zero $H = 1 - 1/(1+2y) = 2y/(1+2y)$, so $y/H = (1+2y)/2 = y + \frac{1}{2}$. □

The purse multiple is independent of B: capital chooses how many problems to post, not how much to post on each. Running the same program through the committed-width shortcut of Section 5, where $V = V_{\min}(y/H)^2$, gives $e^y = 1 + \frac{3}{2}y$, $y^* = 0.76268856$, $V^* = \left(y^* + \frac{2}{3}\right)^2 V_{\min} = 2.04305637\,V_{\min}$ and $\mathbb{P}(\text{solved}) = 0.53358923$. Both constants agree with a two-million-point brute-force maximization of the exact objective. For a purse exponent q, meaning $V = V_e(y/H)^q$, the condition is $e^y = 1 + (1+1/q)y$ and the optimum is $V^*/V_e = \left(y^* + q/(q+1)\right)^q$. Widths that adjust give $q = 1$ and widths that are committed give $q = 2$.

Integer M reintroduces budget-dependent rounding, so budget independence belongs to the relaxation. And expected count solved is one objective among several: a funder who cares most about the single hardest problem should not spread at all.

Sensitivity near threshold. For an attempt sized against V, $p = 1 - \zeta/\sqrt{V}$, so $dp/d\log V = (1-p)/2$ and doubling V sends $p \mapsto 1 - (1-p)/\sqrt{2}$. From $p = 0.01$ that is $p = 0.29996$, and a tenfold purse gives 0.68694. Sensitivity is largest where p is smallest, because a small p means V only just exceeds ζ^2. Equation (14) inherits the steepness, which is why the purse carried 55% of the Millennium effort in Section 3 on a 1% share of the level, and why Proposition 5 places the optimum at 1.76 times break-even, inside the steep region by construction.

Far above threshold the picture reverses. The companion note's continuation width at discount-plus-hazard $\rho > 0$ is $L = \sigma(2\rho)^{-1/2}\,\mathrm{arcosh}(1+\rho V/c)$, which grows like $\log V$ once $\rho V \gg c$, against $\sigma\sqrt{V/c}$ at $\rho = 0$. Far from threshold a purse buys persistence only logarithmically, and spreading capital across more problems beats concentrating it on one. Neither regime applies universally; the ratio V/V_e says which one a contest is in.

7 Assumptions

The lottery and the race. Sections 2–6 solve the lottery. In the first-solution race costs cease when a rival announces, the rival hazard is a function of calendar time, and the equilibrium is the companion note's coupled stopping and forward-distribution system. The race destroys the option where the lottery only dilutes it, so the constants here do not transfer.

Asymmetric equilibria. Proposition 3 gives uniqueness inside the symmetric class, and it needs only that A decreases. Outside that class the question is open. The tempting route is Banach: the symmetric best-response map has modulus $|T'| = \zeta|A'(p)|/(2A^{3/2})$, and a modulus below one

would settle uniqueness outright. It is 0.18 at $n = 2$ and 0.77 at $n = 10$, then 12.0 at $n = 100$ and 500.3 at $n = 10^4$, so the argument fails exactly in the large pool the paper spends most of its time in. Best responses do fall in rivals' success probabilities, which makes this a game of strategic substitutes, and that is all we use.

Pausing. The companion note's intensity extension permits $u = 0$, and one might expect a pause to soften the walk-away threshold. It does not. With progress frozen while idle, constant discounting and rival hazard, and no information arriving, the value function decays at rate ρ while nothing is gained, so waiting strictly destroys value and the optimal control is bang-bang. A mathematician who returns to a problem after twenty years needs something to have changed while it sat: new tools lowering c, or a broader method raising σ. That is a model in which ζ drifts down, which is the mechanism of Section 3.

Independence. The product survival formula behind (11) and behind the rivalry share needs independence across researchers. In mathematics the dominant uncertainty is shared: whether the statement is true, whether the method can work at all. A rival's visible failure is evidence about one's own prospects, and pricing that requires filtering and a correlated state.

Public progress and partial credit. Preprints and talks reveal rival states, which makes a joint-state game. Partial results are published and rewarded, so an all-or-nothing first-passage payoff overstates the waste. Verification takes years, the Poincaré purse being paid seven after the preprints, so a claimed proof that later fails restarts the race and the hazard conflates a claim with a solution.

Measurement. Of the four inputs to Proposition 2 only the hazard comes from data, and it comes from one event. The cost rate and the number of concurrent attackers are stipulated, and the implied level is proportional to both. Measuring what a serious attempt costs is harder than any mathematics here.

8 Related work

Racing for a prize. Loury (1979) and Lee and Wilde (1980) set the pattern: a flow of R&D spending buys a Poisson hazard of success, and free entry dissipates the rents. Dasgupta and Stiglitz (1980) added industrial structure, Reinganum (1981) the dynamic game, Fudenberg et al. (1983) preemption and leapfrogging. Harris and Vickers (1985, 1987) put the race on a state space with observable positions and asked who gives up. In all of them the cost-to-hazard technology is specified; Proposition 1 derives it, which is what turns Section 4 into a scalar fixed point.

Private progress and preemption. Hopenhayn and Squintani (2011) is the closest neighbor, with privately observed progress, an option to stop, and preemption. Bobtcheff et al. (2017) model publishing early against polishing under priority risk. Moscarini and Squintani (2010) and Keller et al. (2005) study experimentation when a rival's information is itself a motive to continue. Those papers close the loop in the harder allocation and leave the cost curve in reduced form; this one closes the curve in the easier allocation.

Contest design. Tullock (1980) and Baye et al. (1996) price rent dissipation. Taylor (1995) and Che and Gale (2003) design research tournaments, Fullerton and McAfee (1999) auctions entry to select contestants, and Halac et al. (2017) design contests when the technology must be discovered. Proposition 5 uses a cruder instrument than any of these, namely how many purses and how large, and returns a fixed multiple of break-even.

Entry games and potentials. Morgan et al. (2012) and Fu et al. (2015) study endogenous entry with subsequent effort. The companion note's fixed-attempt lottery and its exact potential sit in the tradition of Rosenthal (1973) and Monderer and Shapley (1996), and Section 5 uses that potential to locate the three-entrant equilibria.

Budgeted races. The forward-backward architecture of the companion note comes from the budgeted Brownian race (Cotton, 2026a), where one controller prunes many paths under a path-time budget and is paid the terminal maximum. There a shadow price closes the loop. Here it is a consistency condition between stopping rules and the reward they share, so the concavity certificates of that problem say nothing about equilibrium in this one.

9 Closing

An attempt on a hard problem has a price list: success probability p costs $k(p) = \zeta^2 p/(1-p)$. Most of the paper follows from that one formula. It makes the cost curve a theorem about Brownian motion rather than an assumption. It reduces the competitive equilibrium to a single scalar equation, in which competition drives the effective reward down to exactly the break-even level. And it fixes the optimal purse at 1.7564 times break-even, whatever capital the funder has.

Whether a funder should post fewer, larger purses or more, smaller ones depends on where a contest sits relative to break-even. Below about twice break-even, each doubling of the purse raises success probability by roughly $(1-p)/2$, and the purse is the lever that matters. Far above break-even, persistence grows only like $\log V$, so the same dollar does more good posted on another problem. The ratio V/V_e decides which regime applies, and the optimal design, at 1.76, sits just inside the steep one.

The same quantity governs a different problem. A controller allocating test-time compute across parallel rollouts of a language model also pays ζ to traverse the distance to a target, and faces the same convex price for reliability. That is why the budgeted Brownian race (Cotton, 2026a) and this contest share an architecture: a forward equation for how attempts evolve and a backward equation for when to stop. They differ in what ties the two together: a shadow price on compute there, and here a consistency condition between stopping rules and the reward they share.

The first-solution race, in which a rival's announcement ends everyone else's attempt, is the problem the title literally describes. Its existence and selection theory remain open, and we leave them to future work.

References

Baye, M. R., Kovenock, D., and de Vries, C. G. (1996). The all-pay auction with complete information. *Economic Theory*, 8(2):291–305.

Bobtcheff, C., Bolte, J., and Mariotti, T. (2017). Researcher's dilemma. *Review of Economic Studies*, 84(3):969–1014.

Che, Y.-K., and Gale, I. (2003). Optimal design of research contests. *American Economic Review*, 93(3):646–671.

Clay Mathematics Institute (2000). The Millennium Prize Problems. Rules and problem descriptions, `claymath.org`. Accessed September 2026.

Cotton, P. (2026a). Wavefront pruning in budgeted Brownian races. Working draft, `brownianbandit.microprediction.org`.

Cotton, P. (2026b). Research prizes with uncertain compute costs: from entry games to Brownian stopping races. Working note, `brownianbandit.microprediction.org`.

Dasgupta, P., and Stiglitz, J. (1980). Uncertainty, industrial structure, and the speed of R&D. *Bell Journal of Economics*, 11(1):1–28.

Fu, Q., Jiao, Q., and Lu, J. (2015). Contests with endogenous entry. *International Journal of Game Theory*, 44(2):387–424.

Fudenberg, D., Gilbert, R., Stiglitz, J., and Tirole, J. (1983). Preemption, leapfrogging and competition in patent races. *European Economic Review*, 22(1):3–31.

Fullerton, R. L., and McAfee, R. P. (1999). Auctioning entry into tournaments. *Journal of Political Economy*, 107(3):573–605.

Halac, M., Kartik, N., and Liu, Q. (2017). Contests for experimentation. *Journal of Political Economy*, 125(5):1523–1569.

Harris, C., and Vickers, J. (1985). Perfect equilibrium in a model of a race. *Review of Economic Studies*, 52(2):193–209.

Harris, C., and Vickers, J. (1987). Racing with uncertainty. *Review of Economic Studies*, 54(1):1–21.

Hopenhayn, H., and Squintani, F. (2011). Preemption games with private information. *Review of Economic Studies*, 78(2):667–692.

Keller, G., Rady, S., and Cripps, M. (2005). Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68.

Lee, T., and Wilde, L. L. (1980). Market structure and innovation: a reformulation. *Quarterly Journal of Economics*, 94(2):429–436.

Loury, G. C. (1979). Market structure and innovation. *Quarterly Journal of Economics*, 93(3):395–410.

Monderer, D., and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior*, 14(1):124–143.

Morgan, J., Orzen, H., and Sefton, M. (2012). Endogenous entry in contests. *Economic Theory*, 51:435–463.

Moscarini, G., and Squintani, F. (2010). Competition and experimentation with private information. *Journal of Economic Theory*, 145(2):639–660.

Reinganum, J. F. (1981). Dynamic games of innovation. *Journal of Economic Theory*, 25(1):21–41.

Rosenthal, R. W. (1973). A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67.

Taylor, C. R. (1995). Digging for golden carrots: an analysis of research tournaments. *American Economic Review*, 85(4):872–890.

Tullock, G. (1980). Efficient rent seeking. In J. M. Buchanan, R. D. Tollison, and G. Tullock, eds., *Toward a Theory of the Rent-Seeking Society*. Texas A&M University Press.