Revenue Without Returns: Platform Income Under Negative Financial Returns

Flyxion
Independent Researcher

October 2026

Abstract

A platform that sells access to experimentation can be profitable while the average experiment loses money, because the platform is paid for attempts while participants receive financial returns according to their outcomes. This essay establishes that compatibility as an accounting result and distinguishes the mean participant outcome, the typical participant outcome, platform operating profit, and social surplus. Negative expected financial returns need not imply irrational participation, since experiments may also yield consumption value, learning, and information.

A participation model with heterogeneous nonfinancial value characterizes voluntary entry and relates its price elasticity to the hazard rate of the value distribution. Introducing heterogeneous success probabilities permits a distinction between average and marginal entrants, yielding conditions under which additional participation produces less than proportional growth in expected successes. A continuation model then distinguishes services in which success encourages further purchases from those in which success terminates demand. Combining entry, selection, and continuation produces a covariance decomposition of the price elasticity of lifetime transaction volume. Under increasing-hazard selection and independence, the covariance correction changes sign according to whether success or failure produces greater continuation. The composition of active customers can consequently change over successive attempts even when individual success probabilities remain fixed.

The analysis extends to computation sold by usage, including language-model inference. Near a user's optimal computation budget, additional computation reduces user net benefit at second order while increasing the provider's immediate contribution at first order under positive per-unit margins. A conditional quantity-control result identifies the largest allocation compatible with participation, while subscription pricing and retention effects demonstrate that the direction of the provider's incentive is not universal.

The propositions are accompanied by explicit assumptions, potentially refuting observations, and numerical checks that illustrate both the predicted regimes and their failure when the relevant conditions are relaxed. The resulting account does not presume deception, irrationality, or socially wasteful experimentation. It identifies circumstances in which the profitability of experimental infrastructure can diverge from the financial returns of the activities it enables.

1 The Claim and Its Scope

Consider a population of independent projects that each purchase an enabling service from a common provider. The service might be cloud computing, advertising, a software subscription, a course, a marketplace listing, or access to an AI development tool. The question is whether the provider's profitability implies anything about the profitability of the projects it serves. It implies nothing about the sign of the participants' average profit, and the reason is worth stating carefully because the opposite inference is common in discussion of technology markets. A provider is paid per unit of use, whereas a project is paid only if it succeeds. The provider's revenue is therefore a function of the number of attempts, while the project's revenue is a function of their outcomes. The literature on two-sided markets treats the pricing of the intermediary as a problem distinct from the payoffs of the parties it connects, including the choice between fixed and per-transaction charges [6, 1, 7], and the literature on entrepreneurship as experimentation treats attempts as purchases of information whose value need not be financial [4].

The statement developed here is deliberately modest. It asserts that a certain pair of inequalities is jointly satisfiable, that the pair is not exotic, and that moving from this observation to claims about the effect of cheaper experimentation requires additional assumptions that can be written down and tested. It does not assert that customers are irrational or that any particular industry has the structure described. Throughout, negative financial return is kept distinct from negative total utility, since a participant may lose money on an experiment that was nonetheless worth undertaking.

The formal content is organized as follows. Section 2 fixes notation and the maintained assumptions. Sections 3 and 4 establish the accounting claim and separate the three quantities, and Section 6 supplies the participation model that sustains it voluntarily. Sections 7 and 8 develop selection and continuation, and Section 9 combines entry, selection, and continuation in a single result on the price elasticity of lifetime revenue. Section 10 applies the same accounting logic to computation sold by usage, Section 11 addresses readings the results do not support, and Section 12 states the observations that would refute each prediction and summarizes the results and their conditions.

2 Notation and Maintained Assumptions

The participant pays a price c per attempt to the platform, whose marginal cost per attempt is k and whose fixed cost is F. An attempt yields a gross return X_i with mean μ and profit $\pi_i = X_i - c$. There is a continuum of potential participants of mass M, of whom $N(c)$ is the expected number who enter. A participant's nonfinancial value of attempting is v, with distribution function H, density h, survival function $\bar{H}$ and hazard rate $\lambda = h/\bar{H}$. When types differ, a type has attempt-level success probability p with density f and receives a payoff G on success, so that its expected return is pG. The expected number of successes is $\Sigma(c)$, and $\bar{p}$ and p_m are the average and marginal success probabilities among entrants, whose distribution at price c is ν_c. After an attempt a participant continues with probability q_s following a success and q_f following a failure, $q(p)$ is the resulting continuation probability of a type, $w(p) = 1/(1-q(p))$ is its expected number of attempts, $B(c)$ is the cohort's lifetime volume of attempts, and ε_N and ε_B are the price elasticities of entry and of lifetime volume. In the computation model T is units of computation, $Q(T)$ is expected answer quality, V is the value per unit of quality, W is

the user's net benefit, and $\bar{W}$ is the user's reservation level. The symbol Π denotes platform profit or contribution in each setting, with its argument naming the choice variable.

Several assumptions are maintained throughout unless a section states otherwise. The price is exogenous and uniform across participants, with no price discrimination or volume discounts, except in the passages that optimize it. Participants are risk neutral over financial returns, and their projects are independent except where correlated shocks are mentioned. Platform capacity does not bind, and there are no congestion or network effects. Participants' perceived returns equal true returns except where the discussion of signaling says otherwise. Continuation probabilities are reduced-form primitives. All expectations exist, the distributions have densities where used, and differentiation under the integral sign is permitted.

3 A Minimal Model

Let N participants each undertake one experiment. Each experiment requires a payment c to the platform and yields a random gross return X_i with finite mean $\mu = \mathbb{E}[X_i]$. The participant's profit is $\pi_i = X_i - c$, with expectation $\mu - c$. This is the return on the modeled transaction, net of the platform charge and of any other costs already counted in X_i. When X_i excludes other project expenditures such as wages, equipment, financing, and opportunity costs, π_i is not the complete economic profit of the underlying venture, and the accounting result is not a complete theory of entrepreneurial returns. The platform bears a marginal service cost k per experiment and a fixed cost F, so its profit is

$$\Pi_P = N(c - k) - F.$$

Proposition 1 (Joint satisfiability)**.** *If $\mu < c$, $c > k$, and $N > F/(c - k)$, then the expected profit of each participant is negative while the platform's profit is positive.*

Proof. The first condition gives $\mathbb{E}[\pi_i] = \mu - c < 0$ directly. The second and third give $N(c-k) > F$, hence $\Pi_P > 0$. The first condition concerns the participants' returns and the others concern the platform's cost structure and scale, so the three can be chosen independently. □

The proposition is an accounting statement. It shows that the two inequalities are compatible at some values of (N, c), and it makes no claim that such values are an equilibrium, since in the long run N depends on c and on the participants' beliefs. Section 6 supplies the conditions under which the arrangement can be sustained by voluntary entry. For a concrete instance of the accounting result, suppose $N = 10{,}000$, $c = 200$, and $k = 80$, with each project succeeding independently with probability 0.06 and returning 1,500 on success and nothing otherwise. These numbers are assumptions chosen for illustration and are not measurements of any industry. Then $\mu = 0.06 \times 1{,}500 = 90$, so the average participant loses 110, while the platform collects 2,000,000 in revenue and 1,200,000 in contribution before fixed costs. About 600 projects succeed, and each of them returns more than it cost, so the existence of visible winners is fully compatible with a negative average.

4 Three Quantities That Should Be Kept Apart

Public argument about platforms tends to slide between three distinct statements. The first concerns the mean outcome, $\mu - c$. The second concerns the typical outcome, better described

by the median of $X_i - c$. The third concerns the platform's operating profit. Proposition 1 relates the first and third. The second behaves differently, and the difference matters.

With the Bernoulli structure above, the median participant receives zero gross return whenever $p < 1/2$, so the median profit is $-c$ however large the success payoff is. If the success payoff is allowed to be heavy tailed, for example Pareto with a tail index above one so that the mean exists, then μ can exceed c while most participants still lose money. In that case the participants as a group are profitable on average, the platform is profitable, and the typical participant still loses. Whether experimentation is unprofitable on average is determined by the mean, whereas whether the typical participant loses money is determined by the median. These statements are logically distinct, since neither implies the other, although skewness and heavy tails can make their divergence especially pronounced. The empirical literature on entrepreneurial returns has been concerned with exactly this distinction. Moskowitz and Vissing-Jørgensen document that private business returns are highly skewed and that the mean return is low relative to the risk borne [5], and Hall and Woodward find that a substantial share of venture-backed founders earn less than their outside options once nondiversifiable risk is priced [3]. Excess entry relative to what expected profit would justify has likewise been documented in laboratory settings [2]. None of these findings requires the platform to do anything other than sell a service at a price.

Remark 1. A third distinction concerns social value. Suppose an experiment yields expected financial output μ, private nonfinancial value v, and an external informational benefit e to others, and that the real cost of servicing it is k. Ignoring fixed costs and other external costs, social surplus is $S = \mu + v + e - k$, since the payment c is a transfer between participant and platform, while private participation requires $\mu + v - c \geq 0$. The difference $S - (\mu + v - c) = e + c - k$ shows two wedges between private and social value, the informational externality and the markup of price over real marginal cost. A privately unprofitable experiment can therefore be socially valuable if e is large, and a privately profitable one can be socially wasteful if external costs are large. A welfare analysis would also need fixed costs, congestion, displacement of other projects, and a criterion for which outcomes create information for others, and it is outside the present model.

5 Risk Aggregation and What the Platform Bears

Let K denote the number of successful projects. If successes are independent with common probability p, then $K \sim \text{Binomial}(N, p)$, with $\mathbb{E}[K] = Np$ and $\text{Var}(K) = Np(1-p)$. The coefficient of variation is

$$\frac{\sqrt{\text{Var}(K)}}{\mathbb{E}[K]} = \sqrt{\frac{1-p}{Np}},$$

which falls as N grows. A single participant faces an outcome whose variability is of order one, whereas the total number of successes in a large population is nearly deterministic in relative terms.

This is the diversification argument, and it is often stated as if it described the platform's position. It does not quite do so. The platform's revenue in a single period is Nc, which does not depend on K at all. What the platform bears is risk in N and c over time, and these depend on K only through the decisions participants make after observing outcomes. Diversification across independent projects therefore benefits the platform mainly by making the aggregate

pattern of continuation predictable. It does not protect against correlated shocks, such as a change in a common downstream market that lowers p for every project at once, or against a competitor lowering the price of the same service.

6 Participation, Aspiration, and Learning

The question of why people participate when $\mu - c < 0$ is answered most simply by allowing utility to contain more than expected financial profit. Let participant i have

$$U_i = \mu - c + a_i + \ell_i,$$

where a_i is the immediate consumption value of attempting something, including the satisfaction of pursuing a goal, and ℓ_i is the value of what is learned. The two terms are different in kind. The first is a benefit received when the attempt is made. The second is instrumental, because learning changes the participant's later prospects, for instance through an evolving success probability $p_{i,t+1} = p_{i,t} + \Delta p_{i,t}$, and so alters the financial value of continuing and not merely the nonfinancial utility of the present attempt. Write $v_i = a_i + \ell_i$ for the reduced form. Participant i enters if and only if $v_i > c - \mu$. This requires no mistake of fact: a participant who knows that the expected financial return is negative may rationally pay for an experiment whose nonfinancial value exceeds the expected loss, in the way a person may rationally pay for a course or a trip. It also shows how the arrangement of Proposition 1 can be sustained voluntarily, which the accounting result alone does not.

This section treats the case in which all participants share the same expected return μ, and Section 7 lets returns differ across types. Let v_i be distributed over a continuum of potential participants of mass M, so that $N(c)$ is the expected number of entrants, with distribution function H, density h and survival function $\bar{H} = 1 - H$, and let the support of v be the whole real line, since nonfinancial value can be negative for people who dislike attempting. With μ fixed, the number of entrants at price c is

$$N(c) = M\,\bar{H}(c - \mu).$$

Proposition 2 (Elasticity of entry)**.** *With μ independent of c, platform revenue $R(c) = cN(c)$ satisfies*

$$\frac{dR}{dc} = N(c)\big[1 - \varepsilon(c)\big], \qquad \varepsilon(c) = -\frac{cN'(c)}{N(c)} = c\,\lambda(c - \mu),$$

where $\lambda = h/\bar{H}$ is the hazard rate of H. Revenue rises as the price falls if and only if $\varepsilon(c) > 1$.

Proof. Differentiating gives $N'(c) = -Mh(c - \mu)$, so $\varepsilon = cMh(c-\mu)/[M\bar{H}(c-\mu)] = c\lambda(c-\mu)$. Then $dR/dc = N + cN' = N(1 - \varepsilon)$. □

The elasticity of entry is thus the price times the hazard rate of the nonfinancial-value distribution at the participation threshold. A population whose value is concentrated near the threshold is elastic, while a population of committed enthusiasts with large v_i and few participants near the margin is inelastic.

Two qualifications limit how the result should be read. First, the proposition holds the expected return fixed. If $\mu = \mu(c)$, perhaps because a dearer service offers better capabilities, then $N'(c) = -Mh(c - \mu(c))[1 - \mu'(c)]$ and

$$\varepsilon(c) = c\,\lambda\big(c - \mu(c)\big)\big[1 - \mu'(c)\big],$$

which separates a change in price from a change in quality. The same form covers signaling: if participants read a higher price as evidence of capability or downstream validation, their perceived return $\tilde{\mu}(c)$ rises with c (entry depends on the perceived return, while realized profitability depends on the true return distribution, and the two need not agree), the factor $1 - \tilde{\mu}'(c)$ lowers the measured elasticity, and when $\tilde{\mu}' > 1$ entry rises with price locally. Second, H and the participants' beliefs about μ are held fixed. If participants revise their beliefs about success from the observed outcomes of their peers, then H shifts as c and the population change, and the hazard-rate formula describes a condition at a given belief state and not an equilibrium path. The formula is a comparative static and should be read as one.

Remark 2. The effect of a falling cost of attempting a project depends on who bears the reduction. If a new tool lowers a component of the participant's cost that is not paid to the platform, the threshold $c - \mu$ falls while the platform's price is unchanged, and the platform gains entrants at no cost to itself. If the tool lowers the platform's own price, Proposition 2 governs. The two cases should not be merged into a single claim that cheaper experimentation raises upstream revenue.

7 Attempts, Successes, and Selection

A fall in the cost of an attempt draws in participants who would not previously have entered. If the new entrants differ from earlier ones, the number of successes need not rise in proportion to the number of attempts. To state this precisely, let successes pay a common amount G, so that participant i has $\mu_i = p_i G$ and enters if and only if $v_i > c - p_i G$. Let p have density f on $[0, 1]$, independent of v. The numbers of entrants and of expected successes are

$$N(c) = M \int f(p)\, \bar{H}(c - pG)\, dp, \qquad \Sigma(c) = M \int p\, f(p)\, \bar{H}(c - pG)\, dp.$$

Differentiating under the integral, $-N'(c) = M \int f\, h(c - pG)\, dp$ is the density of marginal entrants and $-\Sigma'(c) = M \int p\, f\, h(c - pG)\, dp$ is the density of their expected successes, so the marginal entrant has success probability $p_m = \Sigma'(c)/N'(c)$, and $d\Sigma/dN = p_m$. With $\bar{p} = \Sigma/N$ the average success probability among entrants, the logarithmic elasticity is

$$\frac{d \log \Sigma}{d \log N} = \frac{p_m}{\bar{p}}.$$

Lemma 1 (Subproportional successes)**.** *Suppose the distribution of p among entrants is nondegenerate, the marginal-entry density has positive mass, and the hazard rate λ of H is defined and positive on the relevant range of $c - pG$. If $G > 0$ and $\lambda(c - pG)$ is strictly decreasing in p over the support of the entrants, which holds in particular when λ is strictly increasing there, then $p_m < \bar{p}$, and consequently $d \log \Sigma / d \log N < 1$: a given proportional increase in attempts produces a smaller proportional increase in expected successes. If the hazard rate of H is constant, $p_m = \bar{p}$ and the elasticity equals one.*

Proof. Among entrants the density of p is proportional to $g(p) = f(p)\bar{H}(c - pG)$. The density of p among marginal entrants is proportional to $f(p)h(c - pG) = g(p)\lambda(c - pG)$. Hence $p_m = \mathbb{E}_g[p\, \lambda(c - pG)]/\mathbb{E}_g[\lambda(c - pG)]$ and $p_m < \mathbb{E}_g[p] = \bar{p}$ if and only if $\mathrm{Cov}_g\big(p, \lambda(c - pG)\big) < 0$. The argument $c - pG$ is decreasing in p, so $\lambda(c - pG)$ is decreasing in p exactly when λ is

increasing, which gives a strictly negative covariance under a strictly increasing hazard. For a constant hazard the covariance is zero, provided the relevant thresholds $c - pG$ lie in the region where the hazard is constant; if some thresholds fall outside the support, the effective hazard is zero for those types and the covariance can change. The nondegeneracy condition is needed because the entrant distribution can be degenerate even when the population is not, for instance when only one success-probability type has positive probability of entry at the price considered, and the covariance then vanishes. □

The lemma replaces an ordering assumption with a condition on H. Entrants with high success probability need a smaller nonfinancial value to enter, so they are over-represented among those already inside, and a fall in price adds participants from the margin where the over-representation of high-p types is weaker. Common families with increasing hazard on their support include the normal, the uniform on a bounded interval, and gamma distributions with shape above one, and the exponential (gamma with shape one) is the borderline case. Gamma and exponential distributions have nonnegative support, so the survival function and density must be extended appropriately where the threshold $c - pG$ is negative, and the lemma applies only on ranges where the hazard is defined. The result fails under other structure. If more capable participants place lower nonfinancial value on attempting, so that v and p are negatively dependent, perhaps because they have better outside options, the independence-based proof fails and, depending on the joint distribution, the elasticity can reach or exceed one. Positive dependence can strengthen subproportional selection, as it does in the conditional-normal example of Appendix A, but the direction generally depends on the conditional hazard functions and not on the sign of dependence alone. A new tool that raises individual success probabilities is a separate mechanism, since it changes p itself and so combines a change in participation with a change in the success rate, and it is not part of the entry comparative static. The sign of the selection is therefore an empirical question, and the lemma identifies the assumptions under which it is negative.

Remark 3. Independence can be relaxed without a copula by working with the conditional hazard $\lambda(\cdot \mid p)$ of v given p. The entrant density of p is then proportional to $f(p)\bar{H}(c - pG \mid p)$, the marginal-entrant density to $f(p)h(c - pG \mid p)$, and the proof of Lemma 1 goes through with $\lambda(c - pG \mid p)$ in place of $\lambda(c - pG)$. Subproportional growth of successes holds if and only if $\mathrm{Cov}_g\big(p, \lambda(c - pG \mid p)\big) < 0$. Negative dependence between v and p can push this covariance toward zero or above it, which is the sense in which the conclusion can fail, and the condition shows how much dependence is needed; positive dependence makes the covariance more negative in the numerical specification examined there. Appendix A gives a numerical example in which the ratio $p_m/\bar{p}$ crosses one at a sufficiently strong negative dependence.

8 Continuation Matters at Least as Much as Success

Over many periods the platform's revenue depends on how long participants continue to attempt projects. In this section and the next, success is an attempt-level event: each purchase is a distinct attempt, and a success indicator is observed before the next purchase decision. It is not the final commercial success of an entire venture, since a participant who succeeds in one attempt, for instance by launching a product, may continue to buy, while one who fails may try again. Suppose that after each attempt a participant continues with probability q_s following a success and q_f following a failure, and that continuation depends on nothing else about the

participant's history. The probability of continuing after an attempt is $q = p\,q_s + (1-p)\,q_f$. The expected number of attempts per entrant is $1/(1-q)$, which is finite only when $q < 1$, and lifetime platform revenue per entrant is

$$L = \frac{c}{1-q}.$$

Proposition 3 (Sign of the success effect)**.** *Under one-step memory,*

$$\frac{\partial L}{\partial p} = \frac{c\,(q_s - q_f)}{(1-q)^2}.$$

Lifetime revenue is increasing in the success probability when $q_s > q_f$, independent of it when $q_s = q_f$, and decreasing in it when $q_s < q_f$.

Proof. Since $\partial q/\partial p = q_s - q_f$ and $\partial L/\partial q = c/(1-q)^2$, the chain rule gives the stated derivative. □

The third regime is not by itself evidence of perverse incentives. A participant who completes a project may stop purchasing the service, whereas an unsuccessful participant may keep experimenting. The second regime is consistent with a service whose value lies in the activity of experimenting itself, including aspiration, learning, or enjoyment, because continuation and revenue are independent of the immediately preceding outcome. The first and third regimes correspond to a distinction between services. Call a service expansion-oriented when it helps a customer grow something and continue purchasing, so that success raises continuation and $q_s > q_f$, and completion-oriented when it helps a customer finish something and leave, as when a course that teaches a skill removes the need for further instruction while an unsuccessful student buys more, so that $q_s < q_f$ and the provider's commercial interest diverges from the customer's progress even though both parties act rationally. Whether success is terminal or generative may therefore be more informative than whether a customer is successful.

The proposition rests on one-step memory, and the restriction is not innocuous. A participant who has failed five times and succeeded once may behave differently from one who has failed once, even if $q_s = q_f$ at each step. With history-dependent continuation the expected number of attempts is that of a longer chain, and the three regimes need no longer follow from the sign of $q_s - q_f$ alone. The one-step formulation is an approximation suited to a first treatment and not a description of behavior.

The continuation model deliberately holds attempt-level success probabilities fixed, so the learning term ℓ_i of Section 6 is read there as the anticipated value of information or transferable skill, and learning-induced changes in $p_{i,t}$ are excluded; if they were included, the stationary formula $w(p) = 1/(1-q(p))$ would no longer be exact. Two channels that the model treats together would then separate: learning can raise the probability that an attempt succeeds, and it can reduce the number of attempts needed to reach a goal, so that improved competence could raise successful outcomes while lowering platform revenue.

The probabilities q_s and q_f are reduced-form primitives and are not derived from the entry rule of Section 6. If the per-attempt payoff and nonfinancial value were stationary, attempts were indefinitely repeatable, and there were no budget, horizon, or terminal-goal constraints, a participant who found the first attempt worthwhile would find every repetition worthwhile, and continuation would be certain. Positive stopping therefore requires a departure from that stationary, unconstrained benchmark, so that the value of continuing varies across attempts,

through period-specific shocks to v, revision of beliefs about p after each outcome, budget constraints, or exhaustion of the goal, and the pair (q_s, q_f) summarizes these. Deriving it from an explicit repeated-choice problem is left open, and with a general state space the expected total attempts from each starting state are $\mathbf{t} = (I - \Gamma)^{-1}\mathbf{1}$ when the spectral radius of the substochastic continuation matrix Γ is below one; with history dependence the sign result above need not survive. One consequence is that the same price c governs entry and every later purchase, and that those who continue are in general a selected subset of those who entered.

The model also leaves an identification problem. Continuation after failure may reflect immediate enjoyment, pursuit of an aspiration, acquisition of information, or a rational investment in improved future performance, since the same observed behavior can arise from a_i or from ℓ_i in the earlier utility. Continuation after failure therefore does not, by itself, establish that aspiration is the main source of customer value, and the aspiration reading is one mechanism among several that the mathematics accommodates.

Proposition 4 (Composition of active customers)**.** *Let types have success probability p with entrant density proportional to $\varphi(p) = f(p)\bar{H}(c - pG)$, and let each type continue with probability $q(p) = q_f + (q_s - q_f)p \in (0, 1)$ after every attempt. The type-p mass making its t-th attempt is proportional to $\varphi(p)q(p)^{t-1}$, and the average success probability among customers making attempt t is*

$$\bar{p}_t = \frac{\int p\,\varphi(p)\,q(p)^{t-1}\,dp}{\int \varphi(p)\,q(p)^{t-1}\,dp}, \qquad \bar{p}_{t+1} - \bar{p}_t = \frac{(q_s - q_f)\,\mathrm{Var}_t(p)}{\mathbb{E}_t[q(p)]},$$

where $\mathbb{E}_t$ and Var_t are taken under the weights $\varphi(p)q(p)^{t-1}$. Hence if the distribution of p among those active at attempt t is nondegenerate, $\bar{p}_t$ is strictly decreasing in t when $q_s < q_f$, constant when $q_s = q_f$, and strictly increasing when $q_s > q_f$.

Proof. A type reaches attempt t after continuing $t - 1$ times, each with probability $q(p)$. Then $\bar{p}_{t+1} = \mathbb{E}_t[pq]/\mathbb{E}_t[q]$, so $\bar{p}_{t+1} - \bar{p}_t = \big(\mathbb{E}_t[pq] - \mathbb{E}_t[p]\mathbb{E}_t[q]\big)/\mathbb{E}_t[q] = \mathrm{Cov}_t(p, q(p))/\mathbb{E}_t[q]$. Since $q(p)$ is affine in p with slope $q_s - q_f$, $\mathrm{Cov}_t(p, q(p)) = (q_s - q_f)\mathrm{Var}_t(p)$. □

Remark 4. The persistence result is stronger than an ordering of means. The density of types making attempt $t + 1$, relative to the density at attempt t, is proportional to $q(p)$. When $q_s > q_f$ this likelihood ratio is increasing in p, so the distribution of active types shifts upward in the monotone likelihood-ratio order, and when $q_s < q_f$ it shifts downward. Consequently the expectation of every increasing function of p, whenever it exists, moves in the corresponding direction. The size of the change in the mean is governed by $\mathrm{Var}_t(p)$, so the change in the mean success probability diminishes as the active population becomes homogeneous, provided $\mathbb{E}_t[q]$ stays bounded away from zero, although the likelihood-ratio ordering itself is unchanged, since successive cohorts remain weighted by the same function $q(p)$.

The proposition separates selection at entry from selection through persistence. Entry selects participants whose expected value exceeds the price, and continuation then selects among them according to their success histories, so the active population after several attempts need not resemble the entrant population even when prices and individual success probabilities are unchanged. When unsuccessful participants continue more often, the active population concentrates among low-success types, and the success rate observed among active customers falls over time with no deterioration in anyone's competence or in the technology.

9 A Decomposition of Lifetime Revenue

The selection and continuation results combine once types are allowed to differ. Work in the heterogeneous setting of Section 7, where a participant with success probability p enters if $v > c - pG$, and let each type continue as in Section 8, so that its continuation probability is $q(p) = q_f + (q_s - q_f)p$ and its expected number of attempts is $w(p) = 1/(1 - q(p))$, finite when $q_s, q_f < 1$. Let ν_c denote the distribution of p among entrants at price c, with density proportional to $f(p)\bar{H}(c - pG)$. Expected lifetime transaction volume of the cohort is

$$B(c) = M \int f(p)\, \bar{H}(c - pG)\, w(p)\, dp = N(c)\, \mathbb{E}_{\nu_c}[w(p)],$$

and lifetime revenue is $\mathcal{R}(c) = cB(c)$. A cohort-average shortcut, $cN(c)/(1 - \bar{q}(c))$ with $\bar{q}$ the average continuation probability, is not exact, because w is convex in q and $\mathbb{E}[1/(1 - q)] \geq 1/(1 - \mathbb{E}[q])$ with equality if and only if the continuation probability is constant almost surely among entrants, which holds when entrants do not differ in p and also when $q_s = q_f$.

Proposition 5 (Lifetime revenue elasticity)**.** *Let $\varepsilon_N = c\, \mathbb{E}_{\nu_c}[\lambda(c - pG)]$ be the elasticity of entry and $\varepsilon_B = -cB'(c)/B(c)$ the elasticity of lifetime volume. Then*

$$\frac{d \log \mathcal{R}}{d \log c} = 1 - \varepsilon_B(c), \qquad \varepsilon_B(c) = \varepsilon_N(c) + c\, \frac{\mathrm{Cov}_{\nu_c}\big(w(p),\, \lambda(c - pG)\big)}{\mathbb{E}_{\nu_c}[w(p)]}.$$

Proof. Taking logarithms, $\log \mathcal{R} = \log c + \log B(c)$, which gives the first expression. Differentiating under the integral, $-B'(c) = M \int f\, h(c - pG)\, w\, dp = M \int f\, \bar{H}(c - pG)\, \lambda(c - pG)\, w\, dp$, so $\varepsilon_B = c\, \mathbb{E}_{\nu_c}[\lambda w]/\mathbb{E}_{\nu_c}[w]$. Likewise $\varepsilon_N = c\, \mathbb{E}_{\nu_c}[\lambda]$. The difference is $c\big(\mathbb{E}[\lambda w]/\mathbb{E}[w] - \mathbb{E}[\lambda]\big) = c\, \mathrm{Cov}(w, \lambda)/\mathbb{E}[w]$. □

Lemma 2 (Sign of the covariance term)**.** *Suppose $G > 0$, ν_c is nondegenerate, and $\lambda(c - pG)$ is strictly decreasing in p over the support of ν_c, which holds in particular when λ is strictly increasing there. Then $\varepsilon_B < \varepsilon_N$ when $q_s > q_f$, $\varepsilon_B = \varepsilon_N$ when $q_s = q_f$, and $\varepsilon_B > \varepsilon_N$ when $q_s < q_f$.*

Proof. The weight $w(p)$ is increasing in p when $q_s > q_f$ and decreasing when $q_s < q_f$, and constant when they are equal. By assumption $\lambda(c - pG)$ is strictly decreasing in p over the support. The covariance of two functions of the same variable that are monotone in opposite directions is negative, and the covariance of two functions that are monotone in the same direction is positive, by Chebyshev's covariance inequality, with strict inequality for a nondegenerate distribution and strictly monotone factors. □

The economic content is that the participants easiest to lose when the price rises are not the ones who would have generated the most lifetime revenue. When successful participants continue more often, the high-success types are both more likely to enter and more valuable over their lifetimes, and a price increase disproportionately removes entrants who would have made fewer later purchases, so lifetime revenue is less elastic in price than initial participation. When unsuccessful participants continue more often, the relationship reverses. The result therefore does not support a simple claim that platforms profit from unsuccessful participants. It shows how the relationship between participant success and platform revenue depends on the interaction of selection, continuation, and pricing. If the price is also allowed to affect

continuation through channels other than composition, such as price-dependent q_s and q_f, those effects enter additively and are not modeled here.

The revenue decomposition does not carry over unchanged to contribution. With a constant marginal cost k per attempt, contribution is $(c-k)B(c)$, so that $d\log[(c-k)B]/d\log c = c/(c-k) - \varepsilon_B$, and a rise in lifetime volume raises contribution only when $c > k$. If the marginal cost depends on type, contribution is $\mathcal{C}(c) = N(c)\,\mathbb{E}_{\nu_c}[u_c]$ with $u_c(p) = (c - k(p))\,w(p)$, which is positive for participating types when $c > k(p)$, and

$$\frac{d\log \mathcal{C}}{d\log c} = c\,\frac{\mathbb{E}_{\nu_c}[w] - \mathbb{E}_{\nu_c}[\lambda u_c]}{\mathbb{E}_{\nu_c}[u_c]},$$

so the covariance structure of the revenue result no longer describes profit directly. Heterogeneous costs introduce a second selection mechanism, since a price increase changes the distribution of servicing costs as well as of persistence, and the most persistent customers need not carry the highest margins.

Profit differs from revenue in the obvious way,

$$\Pi(c) = (c-k)\,B(c) - F,$$

and the fixed cost does not affect the optimal price. The first-order condition for revenue at an interior optimum c_R is $\varepsilon_B(c_R) = 1$, equivalently $B(c_R) + c_R B'(c_R) = 0$. The first-order condition for profit at an interior optimum is $B + (c-k)B' = 0$, equivalently $\varepsilon_B = c/(c-k)$. When $k = 0$ the two conditions coincide. When $k > 0$ and lifetime volume is decreasing in price, evaluating the profit derivative at the revenue optimum gives

$$\Pi'(c_R) = B(c_R) + (c_R - k)B'(c_R) = -k\,B'(c_R) > 0,$$

so profit is still rising at the revenue-maximizing price. If profit is unimodal, its maximizer lies above c_R. The argument needs interior optima and unimodality and not a monotone elasticity, and it does not rule out the two objectives sharing a boundary optimum in degenerate cases.

Remark 5. Because $B(c) = N(c)\,\mathbb{E}_{\nu_c}[w]$, the elasticities satisfy

$$\frac{d\log \mathbb{E}_{\nu_c}[w]}{d\log c} = \varepsilon_N - \varepsilon_B = -c\,\frac{\mathrm{Cov}_{\nu_c}\big(w(p), \lambda(c - pG)\big)}{\mathbb{E}_{\nu_c}[w]},$$

which is positive under the conditions of Lemma 2 with $q_s > q_f$. A higher price can therefore raise expected lifetime volume per entrant while lowering the number of entrants, producing smaller but more persistent cohorts. This is a statement about the composition of the remaining customers and does not by itself show that total revenue or profit rises. It suggests an observable quantity, namely whether customers entering at higher prices show greater subsequent transaction volume, meaning the number of purchases and not monetary expenditure, since a higher-priced cohort could make fewer purchases and still generate more revenue per entrant, after controlling for service quality and other determinants of continuation.

10 Usage Pricing and the Allocation of Computation

The structure of the argument applies with little change when the service sold is computation itself, as in language-model inference priced per token. Let a task be answered with T units of

reasoning computation, producing an answer of quality $Q(T)$ that is, in expectation and over the range of interest, increasing and concave, so that additional computation helps but with diminishing effect. Outside that range quality can be nonmonotone, since further reasoning can reinforce an erroneous premise or repeat what was already established, and the local results below require only that Q be twice differentiable with negative curvature near the interior optimum. The user values quality at V per unit and pays c per unit of computation, while the provider's marginal cost is $k < c$. The user's net benefit is $W(T) = VQ(T) - cT$, maximized where

$$V\,Q'(T^*) = c,$$

which is the stopping rule that further reasoning should continue only while the expected improvement in quality from another increment exceeds its cost, written $\mathbb{E}[\Delta Q \mid \mathcal{H}_t] > (c/V)\,\mathbb{E}[\Delta T \mid \mathcal{H}_t]$, with $\mathcal{H}_t$ the reasoning history so far and c/V the quality that one unit of computation must buy to be worth its price. The rule is conceptual, since the expected improvement is not generally computable exactly, but it fixes the quantity at issue. The optimum is problem dependent, as a straightforward factual question may have Q' near zero after little computation, whereas a proof with several interacting constraints may sustain a substantial Q' over a longer interval. It is also distinct from maximizing the ratio $\mathbb{E}[Q]/\mathbb{E}[C]$, which favors cheap and poor answers, and from the constrained form that minimizes expected cost subject to $\mathbb{E}[Q] \geq q_{\min}$.

Proposition 6 (Local divergence)**.** *Let Q be twice differentiable with $Q'' < 0$ near an interior optimum T^*. For a small increase $\delta > 0$ in computation, the user's net benefit changes by $\Delta W = \frac{1}{2}VQ''(T^*)\delta^2 + o(\delta^2)$, which is second order and negative, while the provider's per-task contribution $\Pi(T) = (c-k)T$ changes by $\Delta\Pi = (c-k)\delta$, which is first order and positive.*

Proof. Since $W'(T^*) = 0$, the second-order Taylor expansion of W at T^* gives the first expression. The second follows from linearity of Π in T and $c > k$. □

The proposition compares the orders of two effects and says nothing about what any provider does. A statement about behavior needs an explicit objective and constraint.

Proposition 7 (Quantity-control boundary)**.** *Suppose the provider chooses $T \in [0, \infty)$ at fixed $c > k$ to maximize the single-task contribution $(c-k)T$ subject to the participation constraint $W(T) \geq \bar{W}$, where $\bar{W}$ is the user's reservation utility. Let W be strictly concave with interior maximum at T^*, let $W(T^*) > \bar{W}$, and let $T_+ > T^*$ satisfy $W(T_+) = \bar{W}$. Then the optimal allocation is $T_+ = \max\{T \geq 0 : W(T) \geq \bar{W}\}$, the largest feasible allocation, which exceeds the user's optimum T^*, and the user is left indifferent between participating and declining.*

Proof. Since W is strictly concave with an interior maximum at T^*, it is strictly decreasing to the right of T^*. The assumed $T_+ > T^*$ with $W(T_+) = \bar{W}$ is therefore the largest feasible allocation, and since $(c-k)T$ is strictly increasing in T, the provider chooses T_+. □

This separates technological inefficiency from excess computation that is sustained only by the participation constraint. The computation between T^* and T_+ is not valuable to the user at the margin, yet it generates revenue for the supplier while leaving the user enough benefit to remain. The result is conditional on provider control of T, on the stated objective, and on a binding participation constraint, and it does not establish that any actual platform behaves in this way. It is a statement about quantity control at a fixed unit price and not a general result

about monopoly pricing, since a provider choosing both c and T faces a different problem. The point T_+ exists, for instance, whenever Q is bounded and the domain is unrestricted, because cT grows without limit. Several features of real services work against it. The objective above is the contribution from a single transaction and not long-run profit, and retention, reputation, competition, latency, and capacity constraints can make a smaller T preferable. The provider's default may already lie below T^*, and the provider may not control the allocation at all.

Retention can be built into the objective by letting computation affect continuation, so that lifetime contribution (not to be confused with the cohort contribution $\mathcal{C}(c)$ of Section 9) is

$$\mathcal{P}(T) = \frac{(c-k)\,T}{1-q(T)} - F, \qquad \frac{d\mathcal{P}}{dT} = (c-k)\,\frac{1-q(T)+T\,q'(T)}{\big(1-q(T)\big)^2}.$$

Additional computation raises lifetime contribution if and only if $1-q(T)+Tq'(T) > 0$. If more computation lowers continuation enough, for instance through latency or cost, the inequality fails and the provider's interest points to a smaller T. If more computation raises quality and retention, the retention effect reinforces the immediate revenue effect. The combined model links the continuation result of Section 8 to the allocation of computation within an attempt, though the two mechanisms are different: continuation earns revenue through repeated attempts, and computation earns it through consumption within one.

The direction of the incentive also depends on what is priced. Under usage pricing at a positive price per unit, $\partial\Pi/\partial T = c - k > 0$. Under a fixed subscription fee A, the provider's contribution is $\Pi(T) = A - kT$ with $\partial\Pi/\partial T = -k < 0$, so the immediate incentive points toward less computation. The sign flip is an instance of the choice between fixed and per-transaction charges that the two-sided-markets literature treats as a design variable [7, 1], and both statements follow from the same accounting structure without any assumption about intentions. The incentives governing experimentation therefore depend not only on whether an experiment succeeds but on which observable activity is priced, who controls its quantity, and how that activity affects future participation.

Remark 6. The same structure explains why expenditure is easily mistaken for value. Computation spent, hours worked, and pages written are observable, whereas the counterfactual quality obtainable with less is not. If an evaluator rewards visible reasoning, a system can learn to produce traces that are not causally necessary for its conclusions, so that perceived quality $\tilde{Q}(T)$ keeps rising after true quality $Q(T)$ has flattened, and the gap $\tilde{Q} - Q$ is the room in which unnecessary computation goes unnoticed. The pattern of treating effort as evidence of desert has been discussed in the literature on meritocracy [8], but the mechanism here is the $\tilde{Q}$ versus Q distinction and does not depend on that literature. It also separates productivity from efficiency, since a model using ten times the computation may produce more while producing less per unit of computation. Under usage pricing at a fixed positive unit price, revenue is increasing in T, so the decomposition of the preceding section has no term that rewards the elimination of unnecessary computation, whereas under subscription pricing the immediate incentive reverses.

11 Readings the Results Do Not Support

Several readings of the results are tempting and are not supported by them, and stating them separately may prevent misunderstanding.

The platform profits from failure. The accounting result requires only that revenue depend on attempts. How revenue depends on success is determined by the sign of $q_s - q_f$ in Proposition 3, which can be either, and the sign of the covariance term in Lemma 2 reverses with it. A platform whose customers continue after success gains from their success.

Participants are deceived or irrational. The participation model sustains the arrangement with fully informed participants whose nonfinancial value exceeds their expected financial loss. Departures between perceived and true returns are modeled separately and are not needed for any proposition.

Proposition 1 describes an equilibrium or a real industry. It states that two inequalities can hold together and nothing more, and the numerical example is illustrative.

Mean and median losses are the same claim. Whether the average participant loses is a statement about the mean, and whether the typical participant loses is a statement about the median. Section 4 treats them as distinct, and the refuting observations in Section 12 are stated for each separately.

Cheaper experimentation raises successes in proportion to attempts. Lemma 1 gives the conditions under which successes grow more slowly than attempts, and a tool that raises success probabilities is a separate mechanism that changes the success rate itself.

The quantity-control boundary predicts provider behavior. Proposition 7 is a conditional optimization result under fixed price, provider control of the quantity, a single-transaction objective, and a binding participation constraint. It says what such a provider would choose, not what any provider does, and the surrounding discussion lists the features of real services that work against it.

The revenue results apply unchanged to profit. Contribution carries the margin $c - k$, and with type-dependent costs the covariance structure changes, as shown in Section 9.

Negative financial returns mean negative total utility or social loss. Neither follows. The participation model allows an experiment to lose money and remain worthwhile to the participant, and Remark 1 shows that private and social value are separated by an informational externality and by the markup of price over real cost.

12 Conditional Predictions and What Would Refute Them

Each prediction below follows only from the assumptions named with it, and each comes with an observation that would contradict it. The first concerns the three-way distinction. If the median gross return in a market is found to exceed the price paid, the claim that typical participants lose money in that market is refuted, and if the mean gross return exceeds the price paid, the claim that participants lose money on average is refuted, in each case whatever the platform's profitability. The second concerns selection. Under the independence and increasing-hazard conditions of Lemma 1, the success probability of price-induced marginal entrants should

be below the average among all entrants, $p_m < \bar{p}$, and the elasticity of expected successes with respect to attempts should be below one. The cleanest test randomizes prices, compares success probabilities among all entrants assigned to each price, and separately estimates the success probability of the entrants induced by a lower price. A chronological comparison of cohorts is weaker, since later cohorts can include returning participants and a changed pool of potential entrants. An estimated elasticity at or above one, with the endogenous composition of entrants left free to respond to price and exogenous changes in technology and market conditions controlled, contradicts the lemma's conclusion for that market and indicates that one of its assumptions fails there. The third concerns continuation. Repeat spending after a failed attempt is a potentially important determinant of platform revenue that cannot be inferred from the success rate alone, and Proposition 3 says which sign of $q_s - q_f$ makes the success rate matter. Proposition 4 adds that the success rate among customers making their t-th attempt should move with t in the direction of $q_s - q_f$ when prices and individual competence are unchanged, and a movement in the opposite direction, measured on a fixed cohort, counts against the one-step continuation model for that market. The fourth concerns elasticity. Under the maintained assumption that perceived and true returns coincide and that service quality does not vary with price, the elasticity of entry with respect to price, estimated from exogenous price variation, should equal $c\lambda(c - \mu)$. A systematic departure indicates that at least one maintained assumption or measurement condition fails. Possible explanations include price-dependent perceived returns, as in signaling, correlated changes in service quality, heterogeneous financial returns, changes in the distribution of nonfinancial value, and measurement error, and some of these can be separated by other means, such as price variation that is announced as arbitrary or that leaves quality unchanged.

A fifth item is a measurement program and not a test of the same kind as the first four, since it cannot be run until the user's net benefit function has been estimated. A difference between user-selected and provider-selected budgets is not by itself evidence of inefficient allocation, since users may misjudge the quality function and providers may be constrained by latency or capacity. The relevant comparison estimates the user's net benefit function, or a defensible proxy for it, and compares the observed allocation with the estimated optimum T^*. Under the conditions of Proposition 7, observed allocations should lie above T^* where the provider controls the budget and the participation constraint binds, and an observed allocation at or below T^* in such a market contradicts the quantity-control reading for that market.

The tests also require that the quantities entering each restriction be identified independently of the outcome being tested. If the distribution of nonfinancial value is recovered from the same entry decisions used to estimate the elasticity, the equality $\varepsilon = c\lambda(c - \mu)$ is nearly an identity and its agreement with the data validates nothing, so the distribution should be measured separately, for example through willingness-to-pay elicitation or randomized incentives, and its implied hazard compared with the elasticity from experimentally varied prices. Likewise, lower success rates among later cohorts support the selection lemma only if the distribution of success probabilities can be measured before entry, since changing market conditions, project types, or technology could otherwise produce the same pattern. Without this, some of the falsifications above would test definitions and not economic behavior.

Table 1 collects the results, the conditions on which each depends, and the observation that would count against each.

Table 1: Results, conditions, and refuting observations.

Result	Conditions	Counted against by
Proposition 1, joint satisfiability	$\mu < c$, $c > k$, $N > F/(c-k)$; no equilibrium claim	Not an empirical claim; an accounting statement
Proposition 2, elasticity of entry	Common μ and fixed H; perceived equals true return; quality fixed in price	Elasticity from exogenous price variation differing from $c\lambda(c-\mu)$, with λ measured independently
Lemma 1, subproportional successes	Independence of v and p; $G > 0$; $\lambda(c - pG)$ strictly decreasing in p; nondegenerate entrants	Success elasticity at or above one under randomized prices, with $p_m \geq \bar{p}$ measured among price-induced marginal entrants specifically
Proposition 3, sign of success effect	One-step memory; fixed p; attempt-level success; $q < 1$	Lifetime attempts moving with p against the sign of $q_s - q_f$, with p observed and q_s, q_f estimated separately
Proposition 4, composition of active customers	Heterogeneous p; one-step continuation; $q(p) \in (0,1)$; fixed prices and competence	Success rate among a fixed cohort's active customers moving against the sign of $q_s - q_f$
Proposition 5 and Lemma 2, lifetime revenue elasticity	Heterogeneous p; $G > 0$; $\lambda(c - pG)$ strictly decreasing in p; $q_s, q_f < 1$	$\varepsilon_B \geq \varepsilon_N$ when $q_s > q_f$, estimated with independent measures of λ and w
Proposition 6, local divergence	Q twice differentiable with $Q'' < 0$ near T^*; $c > k$	Evidence that quality lacks the assumed local curvature, or that provider contribution is not increasing in computation; a mathematical consequence of the assumed functions otherwise
Proposition 7, quantity-control boundary	Fixed unit price; provider controls T; single-transaction objective; binding participation constraint	Observed allocations at or below the estimated user optimum where these conditions hold

The generalization suggested by the model is that technological progress can increase the profitability of experimental infrastructure without producing a corresponding increase in the profitability of the experiments it supports. It is not a universal law. Competition can drive c toward k, correlated shocks can reduce platform revenue, and some tools raise p enough to change the sign of $\mu - c$. Negative average private returns, positive platform profit, and positive social surplus are jointly possible, as Remark 1 shows, and a profitable platform with voluntary participation does not by itself establish positive social surplus. What the model establishes is that a platform's revenue depends on the volume and persistence of paid experimentation, whereas the financial success of participants depends on the distribution of experimental outcomes, and that the relationship between the two is mediated by entry,

selection, learning, and continuation and is not guaranteed by the platform's profitability. As Section 11 notes, the thesis does not imply that the underlying transactions are wasteful.

A Numerical Verification

The formulas were checked numerically on a grid. The baseline specification takes nonfinancial value $v \sim N(0,1)$, which has an increasing hazard rate, success probability $p \sim \text{Beta}(2,3)$ independent of v, price $c = 1$, and success payoff $G = 3$. Table 2 compares each closed form with a central finite difference of the corresponding aggregate for three continuation regimes, and the persistence column compares the step from attempt 3 to attempt 4 with the formula of Proposition 4. In every row the closed form and the numerical derivative agree to the printed precision. The marginal and average success probabilities are $p_m = 0.3742$ and $\bar{p} = 0.4726$, so $p_m/\bar{p} = 0.792 < 1$, as Lemma 1 requires.

(q_s, q_f)	ε_N	ε_B (formula)	ε_B (numerical)	$\bar{p}_4 - \bar{p}_3$	Formula
(0.8, 0.4)	0.5955	0.5374	0.5374	+0.02362	+0.02362
(0.6, 0.6)	0.5955	0.5955	0.5955	0.00000	0.00000
(0.4, 0.8)	0.5955	0.6616	0.6616	−0.02262	−0.02262

Table 2: Baseline check of Proposition 5, Lemma 2, and Proposition 4.

Three further checks bear on the assumptions. First, with a constant hazard (an exponential shifted to cover the thresholds, $q_s = 0.8$, $q_f = 0.4$) the model gives $\varepsilon_N = \varepsilon_B = 0.5000$ and $p_m = \bar{p} = 0.4619$, the borderline case of the lemmas. Second, with decreasing hazards (a shifted Lomax and a shifted lognormal, same q_s, q_f) the conclusions fail as the lemmas warn: $p_m > \bar{p}$ in both cases, with $(p_m, \bar{p}) = (0.4465, 0.4346)$ and $(0.4301, 0.4234)$ respectively, and $\varepsilon_B > \varepsilon_N$ even though $q_s > q_f$. Third, with dependence introduced through $v \mid p \sim N(\rho(p - 0.4), 1)$, the ratio $p_m/\bar{p}$ is 0.691 at $\rho = 6$, 0.709 at $\rho = 3$, 0.792 at $\rho = 0$, 1.000 at $\rho = -3$, and 1.263 at $\rho = -6$, and in every case $p_m < \bar{p}$ holds exactly when $\text{Cov}_g\big(p, \lambda(c - pG \mid p)\big) < 0$. In this specification positive dependence strengthens subproportional selection and sufficiently negative dependence reverses it. In a sweep of forty sampled parameter configurations (random seed 5, trapezoidal integration on a grid of 40,001 points, central differences with step 10^{-4}, and tolerance 10^{-5} for the elasticities and 10^{-9} for the persistence step), the numerical calculations agreed with the analytical identities. The configurations varied the value distribution (normal, logistic, and Gumbel), the success-probability distribution, c, G, and (q_s, q_f), and the agreement covered the entry elasticity, the covariance form of ε_B, the log-elasticity identity $d \log \Sigma / d \log N = p_m/\bar{p}$, the elasticity of expected attempts per entrant, and the persistence step in all draws, together with the sign of $\varepsilon_B - \varepsilon_N$ and $p_m < \bar{p}$ in every increasing-hazard draw. Closed forms for the contribution elasticities, the profit-versus-revenue price ordering, the quantity-control optimum, and the retention derivative were checked in the same way, the contribution elasticities for algebraic accuracy only, since no sign result is claimed for them. These checks verify the algebra and the stated sign conditions on the chosen distributions and do not prove the results beyond them. The sign results are proved in the body, from Chebyshev's covariance inequality under the stated monotonicity and nondegeneracy conditions and from the identity $p_m/\bar{p} = \mathbb{E}_g[p\,\lambda]/(\mathbb{E}_g[p]\,\mathbb{E}_g[\lambda])$ for selection, and the numerical work confirms agreement and shows how the conclusions fail when the conditions do.

References

[1] Armstrong, M. (2006). Competition in Two-Sided Markets. *RAND Journal of Economics*, 37(3), 668–691.

[2] Camerer, C., and Lovallo, D. (1999). Overconfidence and Excess Entry: An Experimental Approach. *American Economic Review*, 89(1), 306–318.

[3] Hall, R. E., and Woodward, S. E. (2010). The Burden of the Nondiversifiable Risk of Entrepreneurship. *American Economic Review*, 100(3), 1163–1194.

[4] Kerr, W. R., Nanda, R., and Rhodes-Kropf, M. (2014). Entrepreneurship as Experimentation. *Journal of Economic Perspectives*, 28(3), 25–48.

[5] Moskowitz, T. J., and Vissing-Jørgensen, A. (2002). The Returns to Entrepreneurial Investment: A Private Equity Premium Puzzle? *American Economic Review*, 92(4), 745–778.

[6] Rochet, J.-C., and Tirole, J. (2003). Platform Competition in Two-Sided Markets. *Journal of the European Economic Association*, 1(4), 990–1029.

[7] Rochet, J.-C., and Tirole, J. (2006). Two-Sided Markets: A Progress Report. *RAND Journal of Economics*, 37(3), 645–667.

[8] Sandel, M. J. (2020). *The Tyranny of Merit: What's Become of the Common Good?* New York: Farrar, Straus and Giroux.