Staggered Rollout for Innovation Adoption

Ricardo Fonseca*

July 2026

Abstract

Adoption generates public information, so forward-looking agents may wait for others to experiment. I study a principal who cannot force adoption or use transfers, but can commit at date zero to a path of anonymous supply releases and wants to reach a target adoption level as quickly as possible. Free availability preserves the option to wait and therefore delays completion. With a homogeneous good-state payoff, releasing exactly the target quantity at date zero instead reaches the target immediately. With heterogeneous payoffs, the principal must decide when adoption should proceed smoothly and when scarce supply should induce a wave in which many agents compete for units and adopt together. I characterize payoff-ordered rollouts implemented by anonymous staircase supply. Every such rollout decomposes into smooth adoption along an indifference path pinned down by its starting point and pooled waves that leave and later rejoin it. The design problem within this class therefore becomes choosing a compatible collection of adoption waves, a weighted interval-selection problem. The shape of the payoff distribution determines which pattern is optimal. A positive, nonincreasing density on the active payoff interval uniquely selects capped availability, whereas sufficiently rapid density growth makes staggered expansion profitable. Faster completion also redistributes welfare: agents moved forward lose the option value of waiting, while sufficiently late adopters benefit from earlier information. Transcatheter aortic valve replacement (TAVR) illustrates the model's combination of staged eligibility, continuous payoff heterogeneity, evidence collection, and credible supply commitments.

Keywords: innovation adoption, social learning, scarcity, dynamic mechanism design, rationing, diffusion.
JEL codes: C73, D82, D83, O33.

1 Introduction

Firms and public agencies routinely stage access to innovations. Microsoft introduced the new Bing through a limited preview and waitlist, explicitly seeking feedback before scaling access; Google opened Bard first to a waitlist in the United States and the United Kingdom and expanded access over time; and OpenAI initially released ChatGPT plugins to a small group drawn from

*Pontificia Universidad Javeriana, Carrera 7 No. 40-62, Bogotá, Colombia. E-mail: baricardo@javeriana.edu.co.
I am especially grateful to Bobak Pakzad-Hurson for his guidance and continued support. I also thank Jack Fanning, Teddy Mekonnen, Kareen Rozen, Roberto Serrano, and audiences at Brown University, LACEA–LAMES, EWMES, SBE, USP, Universidad de Chile, Universidad de los Andes, and the Econometric Society Microeconomic Theory Workshop for helpful comments. All remaining errors are mine.

a waitlist before a broader rollout (Microsoft, 2023; Google, 2023; OpenAI, 2023). Invitations, regional openings, account caps, eligibility rules, and usage tiers create scarcity by policy even when the underlying good is reproducible.

These practices matter because early use generates information. Initial users reveal whether an innovation is safe, reliable, useful, or vulnerable to misuse, so an agent who expects access later can wait for others to experiment first. Guaranteed future availability thereby creates an informational free-rider problem: everyone values the evidence, but someone must adopt to produce it. Limiting the option to wait can accelerate learning. At the same time, earlier adoption raises exposure to bad news, so the principal must balance calendar delay against the risk created by experimentation.

To study that tradeoff, I consider a principal who commits to a right-continuous cumulative-supply path before agents apply. An agent's private type is her payoff if the innovation is good, which I call her *good-state adoption payoff.* Agents choose when to apply while learning from a public bad-news process whose intensity rises with cumulative adoption. Because the principal cannot force adoption or use transfers, her only instrument is the timing and amount of anonymous supply. Rejected applicants remain eligible, so rationing changes continuation incentives as well as current access.

Because the policy problem concerns the timing of access, it is useful to state the principal's objective before solving for behavior. Her primitive objective is the discounted payoff from reaching the adoption target, not speed for its own sake. Reaching an adoption target may unlock operating experience, distribution infrastructure, demonstrated demand, authorization for a wider program, or entry into a second market. If completion at date T unlocks $V > 0$ and the principal discounts at rate $\rho > 0$, the calendar component of her payoff is $e^{-\rho T}V$. The lower-case type v introduced below denotes an adopter's good-state payoff.

The principal's objective depends on calendar time, whereas learning depends on cumulative exposure. To separate these two clocks, let M_t denote cumulative adoption, let B_t denote posterior bad-state odds conditional on no bad signal, and let $\beta > 0$ measure bad-news intensity per unit of cumulative adoption in the bad state. Define the exposure clock

$$X_t := \int_0^t M_s \, ds. \tag{1}$$

Along the no-bad-news path,

$$B_t = B_0 \exp\{-\beta X_t\}.$$

Fix a target $\bar{M} \in (0, 1)$. Single crossing makes payoff order the natural representation: whenever a type is willing to take a given anonymous opportunity, every higher-payoff type is more willing to do so. The continuum analysis therefore indexes payoff-ordered service by upper-tail rank. Among such rollouts that serve the upper $\bar{M}$ mass, participation of the marginal served agent and exclusion of lower-payoff agents pin down the same terminal posterior. They consequently have the same probability of surviving to completion, and the principal's primitive payoff is ranked exactly by target-completion time. Completion time is a derived exposure-clock representation of the principal's objective, not an independent welfare criterion.

With this reduction in place, the economic role of dated supply becomes transparent: it changes the option value of waiting. Fine staggering separates nearby payoff types: agents with larger gains accept early access, while agents with smaller gains wait for more information. Pooling sacrifices some sorting but produces a discrete increase in the adoption stock that generates public evidence. Unused units have no intrinsic value in the model. Scarcity matters because promised future access determines how much agents can free-ride on earlier experimentation.

I first isolate this force without heterogeneity. With a homogeneous good-state payoff, free availability sustains gradual take-up and the adoption stock crosses the principal's target only after positive delay. Under capped availability, the principal releases at date zero exactly the total supply needed to reach the target and commits not to release additional units. Waiting then risks exclusion. In the unique equilibrium, all agents apply immediately, a random mass $\bar{M}$ is served at date zero, and the target is reached at once. Because agents are identical in this benchmark, immediate application leads to random rationing at date zero. The heterogeneous analysis below instead studies payoff-ordered service. This benchmark shows that scarcity can accelerate adoption even when it performs no screening role.

Heterogeneity adds a sorting problem, and the main result solves it for continuously distributed adoption payoffs. I index agents from the highest payoff downward by *upper-tail rank*. The analysis focuses on payoff-ordered rollouts, in which higher-payoff agents are served no later than lower-payoff agents, although a positive interval of ranks may adopt together. This representation does not require the principal to observe types. Anonymous staircase capacity implements every schedule with finitely many pooling episodes. Anonymous rationing can also generate nondegenerate service-time distributions when rejected applicants reapply; those outcomes lie outside the rank-to-date representation and are not characterized here. Payoff-ordered rollouts are nevertheless a natural benchmark for announced access schedules because they are exactly what anonymous staircase supply implements when releases clear without rationing, as in launches that commit publicly to dated capacity expansions rather than deliberately using repeated rationing to vary individual service dates. The finite-type results make the boundary explicit. With two target-relevant types, an optimal anonymous finite plan needs at most two exhausted releases; with three types, rationing can make an additional release valuable. The continuum analysis therefore solves the payoff-ordered implementation problem globally while treating rationing-generated service-time distributions as a distinct extension. Between pools, service follows an indifference path pinned down by its starting point, the *separating ridge*. A wave departs from that path, raises the adoption stock while access is paused, and later returns to the same ridge. The decomposition theorem shows that every rollout in this representation consists of the required boundary service plus ridge passage, with a compatible collection of waves inserted along it.

Because each wave begins and ends on the same ridge states, it has an exact excess duration relative to the ridge segment it replaces. Within this ordered representation, the design problem is therefore a weighted interval-selection problem: disjoint waves can be combined, while overlapping waves compete for the same ranks. A countable optimum exists, finite staircase policies approximate its completion time arbitrarily closely, and a finite staircase attains the optimum whenever an optimal wave family is finite. Capped availability is optimal exactly when every feasible wave has nonnegative excess duration. A positive density that is nonincreasing on the active payoff interval makes capped availability uniquely optimal, while sufficiently rapid density growth makes small waves profitable. The same representation also identifies welfare incidence: agents moved into a departure pool lose the option value of waiting, whereas sufficiently late adopters gain from earlier information. For finitely many dated releases, carrying unused on-path capacity forward is without loss for the selected equilibrium. The Appendix records the finite-type boundary: two target-relevant payoff types require at most two releases, whereas three types can make a fourth release strictly valuable.

The mechanism is distinct from a terminal rush caused by anticipated expiration. An opening-day queue may reflect demand accumulated while access was unavailable. By contrast, a late rush before the remaining units or opportunities disappear reflects competition for the last opportunities, one force that can induce earlier adoption in the model. During the 2004 U.S. flu-vaccine shortage, de Janvry, Sadoulet, and Villas-Boas (2010) randomized university departments into receiving information about the sharply reduced number of remaining clinics and their schedule. The treatment raised clinic demand by 0.85 percentage points from a control mean of 0.77 percent, a 110 percent increase. The treatment bundled scarcity and approaching deadlines, so it does not separately identify the two channels, but it directly shows that information about a dwindling set of opportunities can induce earlier demand. A preregistered experiment with unvaccinated German respondents similarly found greater willingness to seek vaccination when local vaccine supply was described as scarce rather than abundant (Schnepf, 2022).

Field episodes display the same terminal timing, although they do not isolate the mechanism experimentally. In the Cayman Islands, public health authorities announced that the current vaccine batch would expire at the end of June 2021 and asked first-dose recipients to come before June 9 so that their second dose could be administered from the same batch (Radio Cayman, 2021). The stock was subsequently expected to be fully used before the next delivery (Connolly, 2021). Other incentives were present, including prize draws, so the episode is illustrative rather than causal. Its sequence nevertheless matches the mechanism: a known terminal constraint, accelerated acquisition, and exhaustion of the current stock.

This strategic-waiting mechanism also produces the familiar S-shaped diffusion pattern: early take-up is slow, adoption then accelerates as the learning stock grows, and diffusion eventually decelerates as the payoff of the marginal adopter falls (Ryan and Gross, 1943; Rogers, 2003; Young, 2009). Section 3 derives the pattern and illustrates it directly.

1.1 Contribution relative to the literature

The paper builds directly on the recent literature on forward-looking adoption with endogenous social learning, beginning with Frick and Ishii (2024). They characterize decentralized adoption when agents learn from the experience generated by earlier adopters. Their informational environment and strategic-waiting logic are the starting point here. The present paper takes that environment in a different direction: a committed principal controls anonymous access, and the object is the optimal rollout rather than the decentralized diffusion path. This line builds on earlier work on investment delay and strategic experimentation with informational spillovers, including Chamley and Gale (1994), Keller, Rady, and Cripps (2005), and Keller and Rady (2015).

Starting from that common learning environment, recent papers develop complementary margins of the adoption problem. Laiho, Murto, and Salmi (2025) compare decentralized incremental adoption with the path chosen by a planner who internalizes information production. Chen and Zhang (2025) and Laiho, Murto, and Salmi (2026) study adoption when social learning interacts with payoff or network externalities. In the latter, positive payoff externalities amplify informational free-riding: greater learning potential can reduce early adoption, slow information accumulation, and lower welfare. Their object is decentralized equilibrium under an exogenous learning technology, whereas the present paper holds the information process fixed and asks how a committed principal should design anonymous access over time. Most closely on the seller side, Laiho and Salmi (2026) study a durable-goods monopolist and forward-looking buyers who learn from past purchases. They show that informational free-riding produces too little and too slow selling in Markov equilibrium and that seller commitment can alleviate the distortion. Their instrument is dynamic pricing and their objective is profit; here the principal commits to aggregate access, learning depends on the stock of adopters, and the objective is to reach a target. The contribution lies in solving the anonymous ordered access-design problem generated by finite-batch incentives and showing how capacity trades off sorting against information production.

A closely related design problem appears in Lyu (2026), who studies a recommendation platform that corrects consumers' insufficient incentives to experiment through a dynamic recommendation policy. The optimal policy can temporarily suspend recommendations after unfavorable feedback. That pause is related in spirit to the waiting intervals studied here, but the instruments are different: Lyu designs information while holding access fixed, whereas I hold the information technology fixed and design anonymous access and rationing over time. Heterogeneous adoption payoffs then make the allocation of early service, as well as the amount of experimentation, part of the design problem.

The paper also relates to a separate literature in which scarcity accelerates purchase through other channels. Seller-induced excess demand can induce consumers to buy before learning their own valuations, while advance-selling and clearance-sale models use capacity and rationing risk to shift demand across dates (DeGraba, 1995; Möller and Watanabe, 2010; Nocke and Peitz, 2007). Parakhonyak and Vikander (2023) study a different form of social learning: observed sales aggregate buyers' private signals and thereby reveal information about the underlying state. Here the signal is public and is generated by aggregate adoption itself. Future access lets agents free-ride on the evidence produced by others, so scarcity accelerates adoption by limiting that option rather than by making sales informative or by relying on stockout risk alone.

The model also has a dynamic-mechanism interpretation. General dynamic mechanism design studies how intertemporal allocations and transfers screen evolving private information (Pavan, Segal, and Toikka, 2014). Here transfers are unavailable and types are fixed, so the principal instead uses the timing of anonymous capacity to screen willingness to adopt early. This places the paper closer to work on allocation without money, where future allocation opportunities provide incentives (Guo and Hörner, 2020), and to nonmarket allocation through costly waiting or other ordeals (Condorelli, 2012). The distinction is that agents also create public information for those who follow them. Single crossing then orders higher-payoff agents earlier, while the ridge-belief schedule

$$b(m) := \frac{r\nu(m)}{r + \beta m}$$

records the posterior odds that make rank m indifferent at the margin. The staircase theorem connects this direct rank-to-date representation back to the primitive instrument by implementing every finite-wave schedule with open access and cumulative inventory caps.

These ingredients lead to the paper's main methodological contribution. Within the payoff-ordered representation implemented by anonymous staircase supply, the continuum design problem can be written as weighted interval selection. The decomposition supplies the candidate intervals and their completion-time weights, while anonymous staircase capacity implements every finite selected family. For a general continuous distribution of adoption payoffs, the solution nests capped availability and arbitrary finite or countable wave patterns, gives an exact necessary-and-sufficient criterion for capped availability, and specifies when and how batching improves on it.

Roadmap. Section 2 presents the payoff and learning environment and reduces the principal's objective to minimizing the time required to reach the adoption target. Sections 3 and 4 study free availability and the pure effect of limiting total supply. Section 5 introduces payoff-ordered continuum rollouts and their implementation through anonymous staircase supply. Section 6 characterizes the optimal rollout within this class, derives the density conditions governing smooth adoption and waves, and studies welfare incidence. Section 7 presents TAVR as an institutional illustration of the paper's mechanism, and Section 8 concludes. The Appendix contains the finite-type boundary, the decomposition and interval-selection proofs, and self-contained proofs of the main density and welfare results. The Mathematical Supplement provides the operational foundations, the full finite-type arguments, including the numerical three-type result, the finite-approximation analysis, and the local calculations underlying the density conditions.

2 Environment

To formalize the mechanism, I now bring strategic waiting, information production, and the principal's control of access into a single dynamic environment. The finite analysis works directly with applications, anonymous rationing, and reapplication; the continuum analysis later uses an ordered rank-to-date representation built on the same primitives.

2.1 Agents, types, and adoption payoffs

A unit mass of agents is indexed by $i \in I = [0, 1]$. Throughout the paper, an agent's type v is her normalized payoff from adoption in the good state. The principal's payoff from reaching the adoption target is denoted by the separate upper-case variable V. In the finite model, the adoption-payoff types satisfy

$$v_1 > v_2 > \cdots > v_N > 0,$$

with type masses $q_n > 0$ summing to one. Later, adoption payoffs follow an atomless, full-support, absolutely continuous cumulative distribution function (cdf) F on the compact interval $[\underline{v}, \bar{v}]$, with density f.

There is a persistent state $\omega \in \{g, b\}$. Agents and the principal share the prior $\mu_0 = \mathbb{P}(\omega = g) \in (0, 1)$. Adoption is irreversible and yields v in the good state and -1 in the bad state. An agent who has not adopted receives zero flow payoff. All agents discount at rate $r > 0$. This rate is distinct from the principal's discount rate $\rho > 0$: r governs agents' adoption incentives, whereas ρ is used only to rank completion dates in the principal's objective.

Let μ_t be the posterior probability of the good state conditional on no bad signal. It is convenient to work with bad-state odds

$$B_t := \frac{1 - \mu_t}{\mu_t}.$$

Multiplying expected utility by the common positive factor $1/\mu_t$ leaves all comparisons unchanged. The normalized time-zero payoff from certain service at date t is therefore[1]

$$g(v, t) := e^{-rt}(v - B_t). \tag{2}$$

2.2 Public learning

Adoption generates publicly observed evidence through a Poisson bad-news process. Conditional on the good state, no bad signal arrives. Conditional on the bad state, bad signals arrive at an intensity proportional to the current adoption stock. This structure is natural for innovations evaluated through failures, adverse events, or unsuccessful trials. Medical and agricultural innovations are leading examples: use produces publicly relevant evidence, and a serious failure can sharply revise beliefs about the innovation's quality.

Given these payoffs, adoption matters dynamically because it produces public information. Throughout, M_t denotes cumulative successful adoption along the no-bad-news path. Bad news arrives only in the bad state. Conditional on $\omega = b$, its instantaneous intensity is βM_t, where $\beta > 0$ is the learning-rate parameter. A bad signal reveals the state, and no further adoption occurs. Along the no-signal path,

$$B_t = B_a \exp\left\{-\beta \int_a^t M_s \, ds\right\}, \qquad t \geq a. \tag{3}$$

Adoption is therefore productive experimentation: a larger stock generates faster information. An atom at date t does not change B_t instantaneously, but it increases the rate at which beliefs

[1] Heterogeneity in v can equivalently be interpreted as heterogeneity in prior beliefs. Suppose agents share a common good-state payoff $\bar{v}$ but an agent has initial bad-state odds B_{0i}, with all agents facing the same likelihood process. Multiplying that agent's payoff by the positive constant B_0/B_{0i} yields the common-prior formulation with $v = (B_0/B_{0i})\bar{v}$. Thus a higher v corresponds to greater initial optimism. The transformation preserves individual choices, equilibrium paths, and the principal's target-completion ranking.

move immediately afterward. Let σ denote the bad-signal time and define exposure over $[a,t]$ by $X_{a,t} := \int_a^t M_s\,ds$. Conditional on no earlier signal, Bayes' rule gives

$$\mathbb{P}(\sigma > t \mid \sigma > a) = \mu_a + (1-\mu_a)e^{-\beta X_{a,t}} = \frac{1+B_t}{1+B_a}. \tag{4}$$

The endpoint posterior and the probability of reaching it are therefore two representations of the same exposure statistic.

2.3 Supply, applications, and rationing

Since the principal cannot choose information directly, she affects this learning process only through access. At date zero, she commits to a right-continuous, nondecreasing cumulative supply path S and does not randomize over supply paths. Under a finite-release policy, S has finitely many jumps. At date t, let A_t denote the mass of applicants and M_t^- the stock of successful adoption immediately beforehand. The anonymous acceptance probability is

$$Q_t = \begin{cases} 0, & M_t^- = S_t, \\ 1, & M_t^- < S_t \text{ and } A_t \le S_t - M_t^-, \\ \dfrac{S_t - M_t^-}{A_t}, & A_t > S_t - M_t^- > 0. \end{cases} \tag{5}$$

Unsuccessful applicants remain eligible. Rationing outcomes are private, while the aggregate adoption stock and public signal are observed.

Agents use pure application strategies after every private history, including histories in which an earlier application was unsuccessful. Strategies need not be type-symmetric: agents with the same adoption payoff may apply at different dates. The equilibrium objects used below are the resulting aggregate applicant, service, adoption, and belief paths. Atoms in aggregate service correspond to batches; continuous increases correspond to gradual take-up while available capacity remains open.

2.4 Continuation payoffs and equilibrium

Definition 1 (Equilibrium)**.** Given a supply path S, an equilibrium consists of pure application strategies and aggregate paths (M, A) such that: (i) after every private history, each agent maximizes her continuation payoff; (ii) every application date prescribed by a strategy is optimal for that agent at that history; (iii) A and M are generated by the strategies and (5); and (iv) $M_t \le S_t$ for all t. When several equilibria exist, the principal selects her preferred equilibrium.

The equilibrium-selection convention matters only when an incentive inequality binds exactly. Principal-preferred selection chooses among equilibria supported by such weak inequalities; it does not impose type symmetry or require all indifferent agents to apply at once.

2.5 The principal's primitive objective and the exposure clock

Let $\bar{M} \in (0,1)$ denote the exogenous adoption target. Reaching it produces the payoff $V > 0$ provided decisive bad news has not arrived; the principal discounts at rate $\rho > 0$. For a policy S, let

$$T(S) := \inf\{t : M_t \geq \bar{M}\}$$

be the target-completion date on its selected no-bad-news equilibrium path; policies that do not attain the target have $T(S) = \infty$. The primitive problem is

$$\max_S \mathbb{E}\Big[e^{-\rho T(S)} V\, \mathbf{1}\{\sigma > T(S)\}\Big], \tag{6}$$

where the maximization is over supply paths and selected equilibrium outcomes satisfying: agents' sequential optimality after every private history; aggregate consistency with the application and rationing rule in (5); supply feasibility, $M_t \leq S_t$ for every t; Bayesian belief consistency as in (3); and target attainment, $T(S) < \infty$. This is a commercial, operational, or institutional milestone objective. It neither adds adopters' private utilities nor treats speed as a complete welfare criterion.

Calendar time and exposure time play different roles. Calendar time determines discounting, whereas

$$X_t := \int_0^t M_s\, ds \tag{7}$$

determines both the no-bad-news posterior and the probability of reaching it. The following reduction states exactly when the primitive objective becomes a completion-time problem.

Lemma 1 (Exposure-clock reduction)**.** *Fix a class $\mathcal{P}$ of policies that reach the target. Suppose their selected no-bad-news aggregate paths are fixed by the selected equilibrium and their posterior bad-state odds at completion equal the same number B^* for every $S \in \mathcal{P}$. Then*

$$\mathbb{E}\Big[e^{-\rho T(S)} V\, \mathbf{1}\{\sigma > T(S)\}\Big] = V e^{-\rho T(S)} \frac{1 + B^*}{1 + B_0}. \tag{8}$$

Maximizing the primitive payoff over $\mathcal{P}$ is therefore equivalent to minimizing no-bad-news target-completion time. Section 5 defines the continuum rollout model formally. Let $\nu(m) := F^{-1}(1-m)$ denote the adoption payoff at upper-tail rank m. For these rollouts, the upper $\bar{M}$ mass of payoff types is served, almost every type receives one date, and higher-payoff types are served weakly earlier. Participation of the marginal served type and exclusion of the lower tail imply $B^* = \nu(\bar{M})$.

Proof. On the selected no-bad-news path, $T(S)$ is fixed. Equation (4) with $a = 0$ gives

$$\mathbb{P}\{\sigma > T(S)\} = \frac{1 + B_{T(S)}}{1 + B_0} = \frac{1 + B^*}{1 + B_0}.$$

Multiplying this common survival probability by the principal's discounted payoff proves (8). The terminal-odds statement for the continuum rollout model follows from the boundary conditions in Proposition 4 and its appendix proof. □

Lemma 2 (A fixed experimentation budget). *Every ordered rollout that serves the upper $\bar{M}$ mass and satisfies participation and exclusion reaches the same terminal posterior and therefore uses the same total cumulative exposure,*

$$X^* = \frac{1}{\beta} \log\left(\frac{B_0}{\nu(\bar{M})}\right). \tag{9}$$

Hence minimizing completion time is equivalent to spending a fixed experimentation budget as quickly as incentives permit.

Proof. Terminal participation and exclusion imply $B_T = \nu(\bar{M})$. Substituting this equality into $B_T = B_0 e^{-\beta X_T}$ gives (9). □

At any given calendar date, a more front-loaded path is more likely to have produced bad news. Policies that reach the same terminal posterior nevertheless have the same survival probability at completion. Free availability generally ends at a different posterior and is not part of this comparison.

The Mathematical Supplement develops the finite-batch environment with positive-mass payoff types, individual application and rationing histories, and reapplication. The Appendix states the resulting two-type boundary and three-type counterexample. These finite-type results concern completion time under anonymous rationing and are not used to derive the payoff-ordered continuum characterization.

In the continuum analysis, the target $\bar{M} < 1$ selects an upper tail of the adoption-payoff distribution. Having reduced the principal's objective to minimizing completion time, the next step is to identify the source of delay that supply policy can correct. Free availability isolates that friction.

3 Free availability

The principal's problem matters because unrestricted access leaves agents free to wait for information produced by others. I therefore remove scarcity first and characterize that decentralized benchmark. Free availability means $Q_t = 1$ whenever an agent applies. Because every application is accepted, an individual applies at most once.

3.1 One payoff type and the free-availability ridge

Let all agents share the same good-state adoption payoff v. Their expected payoff from certain service at date t is $g_v(t) = e^{-rt}(v - B_t)$. On an interior interval of gradual adoption, nearby dates must be equally attractive, so $\dot{g}_v(t) = 0$. Using (3),

$$0 = e^{-rt}\left[\beta M_t B_t - r(v - B_t)\right].$$

The indifference ridge is

$$B_t = \frac{rv}{r + \beta M_t}. \tag{10}$$

Differentiating and substituting $\dot{B}_t = -\beta M_t B_t$ gives

$$\dot{M}_t = M_t(r + \beta M_t). \tag{11}$$

Cumulative adoption consequently rises strictly and convexly while one class is active. The duration required to move from stock m_a to m_b on the ridge is

$$\Delta(m_a, m_b) = \int_{m_a}^{m_b} \frac{dm}{m(r + \beta m)} = \frac{1}{r} \log \left[\frac{m_b(r + \beta m_a)}{m_a(r + \beta m_b)} \right]. \tag{12}$$

Proposition 1 (One-type equilibrium under free availability)**.** *If $v > B_0$, the equilibrium under free availability has the unique aggregate path obtained by jumping, when necessary, to the smallest feasible stock on* (10) *and then following* (11) *until the class is exhausted. If $v < B_0$, adoption never begins. At equality, weak tie selection permits both immediate-adoption and no-adoption equilibria.*

The convex path reflects strategic rather than technological acceleration. Early in the rollout, the stock generating information is small, so a substantial change in belief requires time. As adoption grows, learning speeds up, making earlier adoption increasingly attractive to the remaining agents.

3.2 Several discrete payoff types

The same ridge logic extends recursively across finitely many adoption-payoff types. With $v_1 > \cdots > v_N$, the free-availability path alternates between type-specific ridge segments and waiting intervals. By the single-crossing property, a higher-payoff type is exhausted before a lower-payoff type enters. When type n finishes, the stock is fixed until belief reaches the entry ridge of type $n + 1$. The resulting path is piecewise convex with flat segments.

Agents with the same payoff may adopt at different dates under asymmetric pure strategies, so the relevant object is the aggregate path rather than a common application date for each payoff type.

3.3 Continuous payoff heterogeneity

Passing from positive-mass types to the maintained atomless, full-support, absolutely continuous distribution removes the waiting gaps between adjacent payoff types and yields a scalar cutoff representation. If v_t is the marginal active payoff and $M_t = 1 - F(v_t)$ is the mass already served from the upper tail, local indifference gives

$$v_t = B_t \left(1 + \frac{\beta M_t}{r}\right), \qquad M_t = 1 - F(v_t). \tag{13}$$

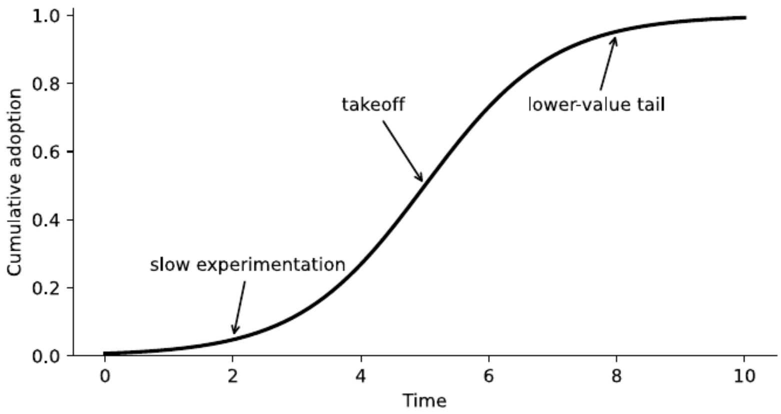

Figure 1: A strategic S-curve under free availability. Adoption is initially slow because little information is being generated. As adoption accumulates, learning accelerates and take-up speeds up. Later, the payoff of the marginal remaining adopter falls enough that diffusion slows again.

Equation (13) implicitly determines the aggregate free-availability path. It reduces the decentralized equilibrium to a scalar cutoff. Index the upper tail by cumulative mass m, measured downward from the highest payoff, and define the payoff at rank m by

$$\nu(m) := F^{-1}(1-m), \qquad b(m) := \frac{r\nu(m)}{r+\beta m}.$$

Because ν is positive and strictly decreasing, b is strictly decreasing.

Proposition 2 (Unique aggregate path under free availability)**.** *Free availability induces a unique aggregate no-signal adoption path in the continuum economy. In the interior-entry case $b(1) < B_0 < b(0)$, there is a unique initial mass $m_0^F \in (0,1)$ satisfying $b(m_0^F) = B_0$. Thereafter the path is continuous and uniquely solves*

$$B_t = b(M_t), \qquad \dot{M}_t = -\frac{\beta M_t b(M_t)}{b'(M_t)}, \qquad M_0 = m_0^F, \tag{14}$$

until the support is exhausted.

Section **??** of the Mathematical Supplement, "Free supply and aggregate-path uniqueness," proves Proposition 2 in full; see Proposition **??** there. Its concluding calculation derives the uniform-distribution curvature condition used below.

The unique free-availability path also yields a familiar aggregate pattern. For a uniform payoff distribution with enough dispersion, adoption is initially convex and eventually concave, with a unique inflection point under the explicit prior condition derived at the end of Section **??**, "Free supply and aggregate-path uniqueness," of the Mathematical Supplement. Early take-up is slow because little information is being generated; as the adoption stock grows, learning accelerates and pulls additional agents forward. Eventually the declining payoff of the marginal adopter dominates that learning effect, so diffusion slows. The model therefore produces an S-shaped adoption path without exogenous adoption opportunities or adjustment frictions.

This shape is useful for the design problem as well as for interpretation. It shows that the same option to wait that delays target completion can generate smooth diffusion in the decentralized benchmark. The target cap studied next changes that option directly, before heterogeneous payoffs create a separate sorting problem.

4 Homogeneous adoption payoffs: the pure scarcity force

Having characterized decentralized adoption, I next isolate the policy force behind scarcity. The free-availability benchmark shows that the source of delay is simple: each agent retains an unconditional option to wait for evidence produced by others. The homogeneous-payoff case shows how a target cap removes that option before screening across payoffs becomes relevant. To focus on the nontrivial case, suppose all agents who may be needed to reach the target share the same good-state adoption payoff $v > B_0$ and strictly prefer to wait rather than adopt immediately under free availability.

Under free availability, agents with the same payoff may adopt at different dates. The resulting stock rises gradually from its date-zero equilibrium mass, crosses the principal's target, and continues rising because access remains available to the rest of the population. Under capped availability, release exactly $\bar{M}$ units at date zero. In the unique equilibrium, all agents apply immediately: each applicant obtains the good with probability $\bar{M}$ and a strictly positive expected payoff, while an agent who waits obtains zero once the target inventory is exhausted. A random mass $\bar{M}$ is therefore served at date zero.

The same logic rules out delayed completion under the cap. Suppose a positive mass applies at a final rationing date $T > 0$. An agent scheduled to apply at T can instead apply at $T - \varepsilon$. If capacity is available immediately before T, then for sufficiently small $\varepsilon > 0$ this deviation guarantees service while sacrificing only an arbitrarily small amount of learning, and is therefore strictly profitable. If no capacity is available on any left-neighborhood of T, the target was already reached before T, contradicting that T is the completion date.

Theorem 1 (Homogeneous-payoff optimum)**.** *When all agents who may be needed to reach the target share one good-state adoption payoff $v > B_0$, releasing exactly $\bar{M}$ units at date zero reaches the target immediately. No policy can do better.*

Figure 2 isolates the pure scarcity force before heterogeneity introduces sorting. The Appendix gives the finite-type boundary. With two target-relevant payoff types, an unrestricted optimum can be implemented with at most two releases. With three types, however, rationing can keep the intermediate type indifferent across dates, making an additional release valuable; the Appendix provides a three-type example in which four releases strictly outperform every three-release plan. I now turn to continuous heterogeneity.

5 Continuously distributed adoption payoffs

The homogeneous-payoff result shows how a cap disciplines waiting when sorting plays no role. With continuously distributed adoption payoffs, access must also determine which agents move

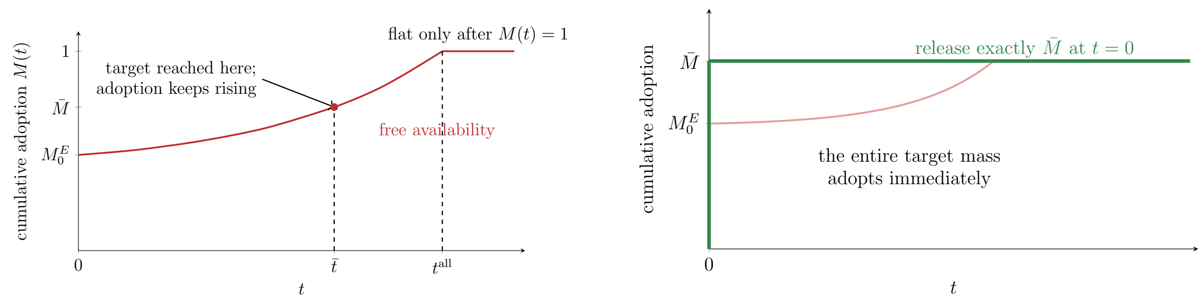

(a) Free availability: gradual adoption crosses the target and continues.

(b) Target cap: the target inventory is claimed at date zero.

Figure 2: One adoption payoff under free and capped availability. In panel (a), the unconditional option to adopt later sustains gradual take-up, so the target is reached only after positive delay and adoption continues afterward. In panel (b), making exactly the target mass available at date zero turns waiting into a risk of exclusion and reaches the target immediately.

first. Single crossing motivates a payoff-ordered class. The continuum representation records that class as a rank-to-date rollout: higher-payoff agents are served weakly earlier, almost every served rank receives one service date, and positive rank intervals may be pooled. Anonymous staircase supply implements every finite-wave rollout directly. Smooth portions arise while capacity remains open; pooled portions arise when capacity is exhausted or expanded.

5.1 Capped availability and its unique aggregate path

The simplest operational benchmark is to release the entire target inventory at date zero and let agents decide when to claim it.

Definition 2 (Capped availability)**.** Capped availability releases exactly the target inventory at date zero and leaves every unclaimed unit accessible:

$$S_t = \bar{M} \qquad \text{for every } t \geq 0.$$

While inventory remains, every applicant is served with certainty, so adoption follows the same local incentives as under free availability. The cap matters at the end because delay eventually risks exclusion.

Let

$$v_t^M := B_t, \qquad v_t^E := B_t \left(1 + \frac{\beta M_t}{r}\right)$$

be the myopic and strategic adoption-payoff thresholds while capacity remains. The latter is the good-state adoption payoff of the type whose certain-service payoff is locally flat at date t.

Proposition 3 (Unique aggregate path under capped availability)**.** *Suppose*

$$\nu(\bar{M}) < B_0 < \nu(0).$$

Define

$$b(m) = \frac{r\nu(m)}{r + \beta m},$$

and let m_0, m_1 be the unique pair satisfying $0 < m_0 < m_1 < \bar{M}$ and

$$b(m_0) = B_0, \qquad b(m_1) = \nu(\bar{M}). \tag{15}$$

Capped availability induces a unique aggregate no-signal equilibrium path. It jumps to m_0 at date zero, then follows

$$B_t = b(M_t), \qquad \dot{M}_t = -\frac{\beta M_t b(M_t)}{b'(M_t)} \tag{16}$$

from m_0 to m_1, and completes with an unrationed terminal stockout of mass $\bar{M} - m_1$. Equivalently,

$$m_0 = 1 - F(v_0^E), \qquad \bar{M} - m_1 = F(v_T^E) - F(v_T^M). \tag{17}$$

Uniqueness concerns the aggregate adoption and belief path, not the assignment of measure-zero cutoff types.

Section **??** of the Mathematical Supplement, "Capped availability," gives the full operational proof of Proposition 3; see Theorem **??** there. The proof establishes terminal zero rent, the unrationed stockout, and uniqueness of the aggregate path.

Even before solving the full design problem, a single up-front release can generate three phases: an initial wave of applications for the available units, smooth diffusion while supply remains available, and eventual exhaustion of the initial stock. No mid-course supply intervention is needed to implement this path. The rank interval $[m_0, m_1]$ between the initial and terminal pools is the *active rank interval*; its image $[\nu(m_1), \nu(m_0)]$ is the *active payoff interval.* The density conditions below apply to the range of payoffs served before the target is reached. Capped availability provides the benchmark with no intermediate adoption waves. The next subsection characterizes the broader payoff-ordered service patterns that can be implemented through anonymous capacity releases.

5.2 Rollout schedules, staircase supply, and the separating ridge

Capped availability shows that anonymous supply can decentralize a finely ordered service pattern without eliciting or observing payoff types. To characterize all such patterns, represent the principal's choice as a rank-to-date schedule. A single service date is assigned to almost every served rank, higher-payoff ranks are served weakly earlier, and positive-measure rank intervals may be pooled. This gives the continuum representation of the payoff-ordered rollout class. It contains the smooth-service and pooling patterns implemented by anonymous staircase supply without requiring type-specific access rules.

Index the served upper tail by cumulative mass rank $m \in [0, \bar{M}]$: $m = 0$ is the top of the payoff distribution, and $\nu(m) = F^{-1}(1 - m)$ is the good-state adoption payoff at rank m. A rollout assigns a date $\tau(m)$ to that rank, with τ weakly increasing; a positive interval may share one date. Let λ denote Lebesgue measure on $[0, \bar{M}]$. Because rank is measured in population mass, the adoption and belief paths are

$$M_t = \lambda\{m \in [0, \bar{M}] : \tau(m) \le t\}, \qquad B_t = B_0 \exp\left\{-\beta \int_0^t M_s \, ds\right\}.$$

Global direct incentive compatibility means that every served rank weakly prefers its assigned date to the date assigned to any other served rank. Formally,

$$e^{-r\tau(m)}\{\nu(m) - B_{\tau(m)}\} \geq e^{-r\tau(n)}\{\nu(m) - B_{\tau(n)}\} \qquad \text{for all served } m, n. \tag{18}$$

Although the schedule is written as a rank-to-date map, it need not be administered by observing ranks. Proposition 5 shows that every globally incentive-compatible schedule with finitely many waves is exactly decentralized by anonymous applications to a finite staircase cumulative-supply path.

Given this representation, local incentive compatibility on a smooth separating component makes the marginal type indifferent across neighboring service dates. The same calculation as under free availability gives

$$b(m) = \frac{r\nu(m)}{r + \beta m}. \tag{19}$$

If $x(m) = -\log b(m)$, the smooth time density is

$$\tau'(m) = \frac{x'(m)}{\beta m}. \tag{20}$$

Smooth passage from m_a to m_b therefore takes

$$\mathcal{S}(m_a, m_b) = \int_{m_a}^{m_b} \frac{x'(m)}{\beta m}\, dm. \tag{21}$$

Local indifference identifies the ridge, but a global characterization must also account for pools, pauses, and boundary behavior. The intervals served at the first and last service dates are the *boundary pools.* For $d < p < h$, a pooled excursion serves ranks $[d, p]$ at a departure date, holds cumulative adoption fixed at p while beliefs improve, then serves ranks $[p, h]$ at reentry and resumes ridge passage. An excursion is *exact* when it returns to the same stock–belief point as smooth passage and satisfies global incentive compatibility; Proposition 6 gives the contact equation. Boundary pools, ridge passage, and exact excursions exhaust the possible components.

Proposition 4 (Boundary pools and exact decomposition)**.** *Suppose ν is positive and strictly decreasing and*

$$\nu(\bar{M}) < B_0 < \nu(0).$$

Let τ be any finite-completion rollout satisfying global incentive compatibility, participation of the upper $\bar{M}$ mass, and exclusion below $\nu(\bar{M})$. Define $b(m) = r\nu(m)/(r + \beta m)$, and let m_0, m_1 be the unique pair satisfying $0 < m_0 < m_1 < \bar{M}$ and

$$b(m_0) = B_0, \qquad b(m_1) = \nu(\bar{M}).$$

If $t_0 = \tau(0)$ and $T = \tau(\bar{M})$, then $\tau = t_0$ on $[0, m_0]$ and $\tau = T$ on $[m_1, \bar{M}]$. Between these compulsory boundary pools, every continuously increasing segment is an absolutely continuous *passage along the ridge, and every jump is an exact pooled excursion* (d_j, p_j, h_j)*. The excursion supports have disjoint interiors; there are at most countably many of them; and*

$$T = t_0 + \mathcal{S}(m_0, m_1) + \sum_j \Gamma_F(d_j, p_j, h_j), \tag{22}$$

where

$$\Gamma_F(d, p, h) = \frac{1}{\beta} \int_d^h x'(m) \left(\frac{1}{p} - \frac{1}{m} \right) dm.$$

The initial pool may exceed m_0 *only by incorporating the departure block of the first excursion, and the terminal pool may begin below* m_1 *only by incorporating the reentry block of the last excursion.*

m_0 and m_1 are unique for every admissible adoption-payoff distribution; the optimal excursion family remains distribution-dependent. The decomposition also rules out a singular-continuous component, meaning continuous growth concentrated on a zero-measure set of ranks, and permits a finite or countable number of smooth and pooled components. In any optimum, the removable idle time satisfies $t_0 = 0$.

The direct schedule can also be implemented with the original supply instrument.

Proposition 5 (Anonymous staircase implementation)**.** *Let* τ *be a finite-completion, globally incentive-compatible rollout whose decomposition contains* $K < \infty$ *exact excursions* $e_k = (d_k, p_k, h_k)$ *in increasing mass order. Then* τ *is induced by a selected equilibrium of an anonymous nondecreasing cumulative-supply path with at most* $K + 1$ *positive jumps and no on-path rationing. If* $K = 0$*, the target inventory* $\bar{M}$ *is released at the first service date. If* $K \geq 1$*, cumulative supply is first raised to* p_1*, then to* p_{k+1} *at the reentry date of excursion* k *for* $k < K$*, and finally to* $\bar{M}$ *at the reentry date of excursion* K*. At each departure, the remaining inventory* $p_k - d_k$ *is exhausted; at reentry, mass* $h_k - p_k$ *is served immediately and the unclaimed balance supports the next smooth segment. Conversely, any no-rationing equilibrium of such a staircase path in which almost every payoff type receives one ordered certain service date induces a globally incentive-compatible ordered rollout.*

To see the implementation logic, consider one excursion (d, p, h). The principal makes p units available initially. The initial rush and smooth ridge passage use part of that inventory; the departure pool $[d, p]$ exhausts the remainder. Adoption then stops because the cap binds. At the prescribed reentry date, the principal raises cumulative supply to $\bar{M}$. The reentry pool $[p, h]$ adopts immediately, the remaining units stay available through the resumed ridge passage, and the terminal pool exhausts the target inventory. Accordingly, $p - d$ and $h - p$ are adoption jumps, not the two announced supply quantities: the supply increments are p and $\bar{M} - p$.

The displayed inequalities impose the generic interior case in which the initial boundary pool and the active ridge interval are both nondegenerate. Outside this case, the initial pool may collapse, so rollout begins directly on the ridge, or the active rollout may collapse into immediate completion. If $B_0 \geq \nu(0)$, even the highest-payoff type cannot initiate positive adoption, and no positive finite-completion rollout exists. The initial-batch, smooth-diffusion, terminal-batch pattern is only one possible payoff-ordered rollout. More generally, a rollout may contain several waves separated by smooth adoption.

6 Global characterization within the ordered class

Section 5 characterized the building blocks of a payoff-ordered rollout: smooth adoption along the ridge and discrete adoption waves. This section uses that decomposition to solve the principal's design problem. The key question is which waves should be inserted into the smooth benchmark and how their gains interact when they involve overlapping ranks.

Because each wave returns to the same ridge state that smooth passage would have reached, its effect on completion time can be evaluated independently. Gains add across disjoint rank intervals, while overlapping waves compete for the same agents. The design problem therefore becomes selecting a compatible finite or countable collection of waves. Capped availability and the optimal one-wave policy are special cases.

Recall

$$b(m) = \frac{r\nu(m)}{r + \beta m}, \qquad x(m) = -\log b(m),$$

and, in the generic interior case,

$$T^{\text{cap}} = \mathcal{S}(m_0, m_1) = \int_{m_0}^{m_1} \frac{x'(m)}{\beta m}\, dm.$$

This is the duration of the no-wave ridge passage between the two compulsory boundary pools.

6.1 Exact excursions and the global selection problem

The decomposition from Section 5 reduces each candidate wave to a comparison with the smooth ridge segment it replaces. Figure 3 illustrates this comparison. Starting from a point on the ridge, the smooth benchmark serves ranks continuously, whereas a wave serves an interval of ranks together, pauses while beliefs improve, and then returns to the same ridge state. Because the two paths share the same endpoints, their difference can be summarized by the change in completion time over that rank interval.

Consider the completion-time contribution of one departure–pause–reentry cycle relative to the smooth segment it replaces. Fix $m_0 \le d < p < h \le m_1$, put

$$a = \nu(p), \qquad k = \frac{r}{\beta p},$$

and define the three-argument pause duration

$$\Delta(d, p, h) = \frac{1}{\beta p} \log \frac{b(d)}{b(h)}. \tag{23}$$

The departure batch serves $[d, p]$, adoption remains at p for $\Delta(d, p, h)$, and the reentry batch serves $[p, h]$. The stock–belief pair after reentry is exactly the pair at which ridge passage reaches h. This endpoint matching is what makes the global problem tractable: a wave changes service timing only over the ranks it covers. As Figure 3 makes clear, disjoint waves affect separate rank intervals, whereas overlapping waves prescribe incompatible service patterns for some of the same agents.

Proposition 6 (Exact excursion criterion)**.** *Suppose F has a positive, continuously differentiable density on the active interval. The excursion is globally incentive compatible and returns to the ridge if and only if*

$$b(d)^k\{a - b(d)\} = b(h)^k\{a - b(h)\}. \tag{24}$$

Whenever (24) *holds, replacing smooth passage over $[d, h]$ changes completion time by*

$$\Gamma_F(d, p, h) = \frac{1}{\beta} \int_d^h x'(m) \left(\frac{1}{p} - \frac{1}{m} \right) dm. \tag{25}$$

Negative Γ_F means that the pooled excursion is faster than the ridge segment connecting the same endpoint stock–belief pairs.

This same-state comparison is what makes Γ_F the correct principal criterion. Ridge passage and the exact excursion begin at the same stock–belief pair $(d, b(d))$ and return at the same pair $(h, b(h))$, so Equation (4) gives the same conditional survival probability $(1 + b(h))/(1 + b(d))$ under both paths. A negative Γ_F therefore reaches the same stock-belief state, with the same probability of surviving to it, at an earlier calendar date. It does not purchase speed by lowering the chance of reaching reentry.

The contact equation determines feasibility and the gain formula prices the intervention. It makes the marginal payoff type $a = \nu(p)$ indifferent between the two batch dates, while single crossing assigns higher-payoff types to departure and lower-payoff types to reentry. A useful intervention must raise experimentation before the pause and then use a clearing release to return the schedule to the ridge. A lone delayed batch cannot help, because pausing at d and later jumping to $h > d$ has excess duration

$$\frac{1}{\beta} \int_d^h x'(m) \left(\frac{1}{d} - \frac{1}{m} \right) dm > 0.$$

To pass from one intervention to a complete rollout, collect all exact excursions that can be inserted into the ridge. Write $k = r/(\beta p)$ for the contact exponent associated with pivot p. Let $\overline{\mathcal{E}}_F$ denote the compact exact-excursion set characterized by Proposition 6:

$$\begin{aligned} \overline{\mathcal{E}}_F = \{&(m, m, m) : m \in [m_0, m_1]\} \\ &\cup \Big\{ (d, p, h) : m_0 \le d < p < h \le m_1, \\ &\qquad b(d)^k\{\nu(p) - b(d)\} = b(h)^k\{\nu(p) - b(h)\} \Big\}. \end{aligned} \tag{26}$$

Diagonal triples are null excursions. Two excursions are *compatible* when the interiors of their support intervals $[d, h]$ are disjoint; a finite or countable family is compatible when its members are pairwise compatible. Write $\mathfrak{A}_c$ for the set of compatible finite or countable families.

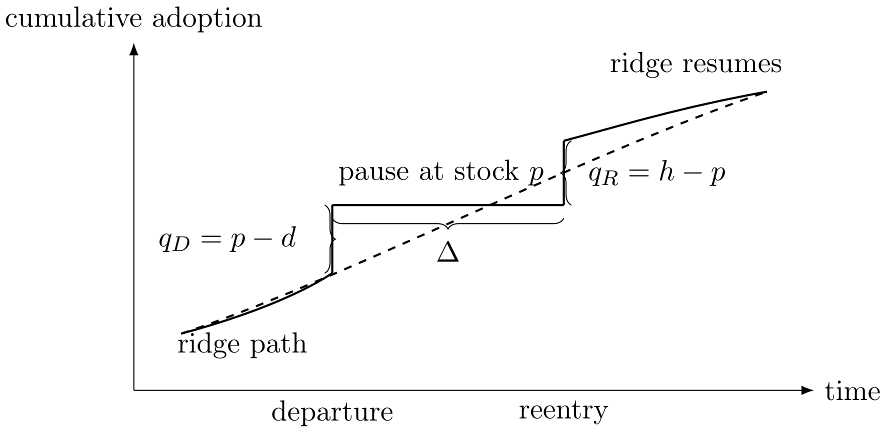

Figure 3: A pooled excursion. The departure batch advances mass $p - d$, the principal waits for $\Delta(d, p, h)$ while adoption remains at p, and the reentry batch serves mass $h - p$. The contact equation makes the two batch payoff lines meet at the marginal payoff $\nu(p)$.

Theorem 2 (Exact interval-selection reduction)**.** *Suppose the hypotheses of Proposition 4 hold. Among ordered rollouts with finitely many exact excursions, the infimum completion time is*

$$T^{\text{fin}} = T^{\text{cap}} + \inf_{\mathcal{A} \in \mathfrak{A}_f} \sum_{e \in \mathcal{A}} \Gamma_F(e). \tag{27}$$

Every finite-excursion rollout generates a compatible family with exactly this duration, and every finite compatible family reconstructs a globally incentive-compatible ordered schedule. Proposition 5 implements each such family by anonymous staircase supply.

Corollary 1 (Countable attainment and finite approximation)**.** *If, in addition, x is four times continuously differentiable on the active rank interval, then*

$$T^* = T^{\text{cap}} + \min_{\mathcal{A} \in \mathfrak{A}_c} \sum_{e \in \mathcal{A}} \Gamma_F(e).$$

The minimum is attained by a countable compatible family, its gain series is absolutely convergent, and for every $\varepsilon > 0$ a finite staircase policy reaches completion within ε of T^*. A finite staircase attains the optimum whenever an optimal family is finite.

The theorem is the exact reduction; the additional C^4 condition is used only to close the finite family class through the uniform small-excursion bound in the Appendix. Proposition 5 therefore identifies the optimum within the ordered class with the infimum over finite anonymous staircase implementations under the corollary's regularity.

The solution within the ordered class is an exact weighted interval-selection problem. Disjoint excursions can be combined because each returns to the ridge state inherited by the continuation, while overlapping excursions compete for the same mass interval. Their weights determine both the number and the endpoints of the selected waves. The Appendix gives the Bellman equation, the countable-attainment argument, and a finite-compression bound.

Optimal paths may alternate repeatedly between smooth separation and pooled waves. A boundary wave may merge with the initial or terminal pool, and capped availability is the excursion-free member of this larger class.

Regularity map. The decomposition, exact excursion criterion, and finite interval-selection reduction use the monotonicity and differentiability conditions stated with those results. Fourth-order smoothness is imposed only for countable attainment and finite-staircase approximation. The later local concentration results add their own reflected-contact, slack, and Jacobian conditions because they perturb excursion endpoints; those conditions are not part of the global reduction.

6.2 Who bears the cost of faster rollout

Before translating the global solution into primitive density conditions, the same representation answers a separate economic question: who bears the cost of acceleration? Fix a nondegenerate excursion $e = (d, p, h)$, normalize its departure date to zero and compare it with free availability beginning from the same inherited public history, summarized at departure by $(M, B) = (d, b(d))$. Under free take-up, rank $m \geq d$ would be served on the ridge after

$$\mathcal{S}(d, m) = \int_d^m \frac{x'(z)}{\beta z}\, dz$$

and would obtain

$$U_d^F(m) = e^{-r\mathcal{S}(d,m)}\{\nu(m) - b(m)\}. \tag{28}$$

Under the excursion, the departure pool $[d, p]$ is served immediately at belief $b(d)$, the reentry pool $[p, h]$ is served after $\Delta(d, p, h)$ at belief $b(h)$, and all later ranks resume the ridge from h. Let $U_d^e(m)$ denote the resulting normalized payoff.

Proposition 7 (Welfare incidence of a pooled wave)**.** *For every nondegenerate exact excursion $e = (d, p, h)$:*

(i) The departure rank is indifferent, $U_d^e(d) = U_d^F(d)$, while every agent advanced in the departure pool is strictly worse off:

$$U_d^e(m) < U_d^F(m) \qquad \text{for all } m \in (d, p].$$

(ii) Every rank at or after reentry receives the same belief-contingent payoff as under free take-up, shifted by the excursion's excess duration:

$$U_d^e(m) = e^{-r\Gamma_F(e)} U_d^F(m) \qquad \text{for all } m \geq h. \tag{29}$$

Every such follower is consequently weakly better off when $\Gamma_F(e) < 0$, and every follower with $U_d^F(m) > 0$ is strictly better off. A zero-rent terminal margin remains indifferent.

(iii) If $\Gamma_F(e) < 0$, there is a unique cutoff $\widehat{m} \in (p, h)$ within the reentry pool such that ranks in $[p, \widehat{m})$ are worse off than under free take-up; rank $\widehat{m}$ is indifferent; and ranks in $(\widehat{m}, h]$ are strictly better off.

Corollary 2 (Incidence along an optimal rollout)**.** *An optimal compatible family can be selected so that every retained nondegenerate wave has strictly negative excess duration. Relative to replacing any such wave by free ridge passage from the same inherited public history, its departure pool is worse off, every rank at or after reentry is weakly better off (strictly so when its benchmark payoff is positive), and a unique cutoff separates losers from gainers inside its reentry pool.*

A profitable wave is not a Pareto improvement. It accelerates learning by moving experimentation toward the higher-payoff agents in its departure pool, who accept earlier service because the staircase policy removes their option to wait under free availability. The gain passes down the payoff ranking: sufficiently late members of the reentry pool and every later adopter with a positive benchmark payoff receive the same informational terms earlier. A wave that lowers completion time thus imposes losses at its front and delivers gains at its back.

6.3 Capped availability as the excursion-free optimum

The welfare comparison uses the same excursion weights as the design problem. Returning to optimal policy, capped availability is optimal exactly when no admissible departure from the ridge has negative weight.

Theorem 3 (Exact capped-optimality criterion)**.** *Suppose the hypotheses of Proposition 4 hold. Capped availability minimizes completion time if and only if*

$$\Gamma_F(d,p,h) \geq 0 \qquad \textit{for every } (d,p,h) \in \overline{\mathcal{E}}_F. \tag{30}$$

It is the unique optimum if every nondegenerate contact triple satisfies the strict inequality.

Capped optimality is equivalent to a sign test over excursions. The criterion is exact but nonlocal: the contact equation links two ridge states, and the gain integrates over the interval between them. A pointwise density restriction provides a readable sufficient condition for the universal inequality in (30).

Write f for the density of good-state adoption payoffs. Since $\nu'(m) = -1/f(\nu(m))$,

$$x'(m) = \frac{\beta}{r+\beta m} + \frac{1}{\nu(m) f(\nu(m))}. \tag{31}$$

Corollary 3 (Capped availability under a nonincreasing active density)**.** *Suppose the interior boundary case holds, F has a positive, continuously differentiable density, and*

$$f'(v) \leq 0 \qquad \textit{for every } v \in [\nu(m_1), \nu(m_0)].$$

Then every nondegenerate feasible excursion has strictly positive excess duration. The unique optimum is therefore the excursion-free rollout: a date-zero block of mass m_0, smooth ridge rollout from m_0 to m_1, and a terminal block of mass $\bar{M} - m_1$.

Only the density on the active payoff interval matters for this conclusion. No monotonicity restriction is imposed above $\nu(m_0)$, among agents absorbed by the initial block; below $\nu(m_1)$, among agents absorbed by the terminal block; or below $\nu(\bar{M})$, among agents excluded from the target. More generally, a rising density need not make a wave profitable unless its concentration is sufficiently strong and compatible with the exact contact condition.

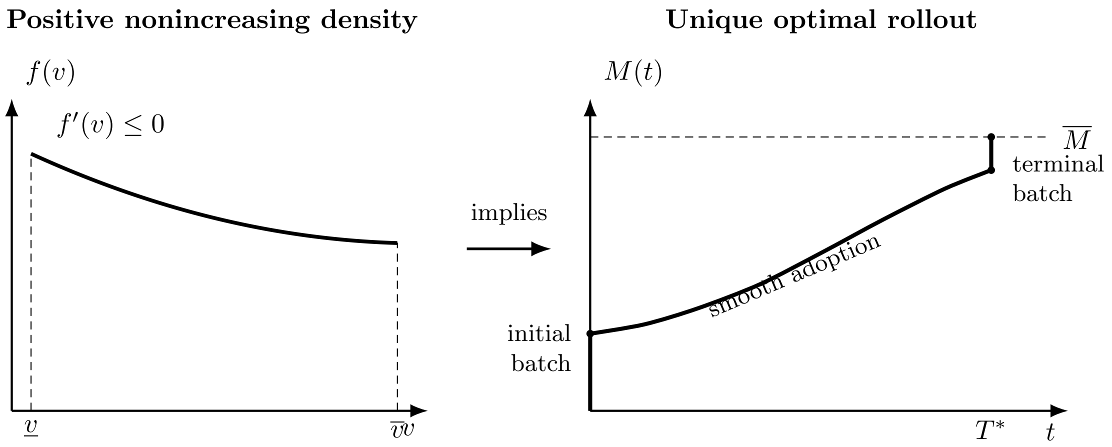

Figure 4: Nonincreasing active density and the unique optimal ordered rollout. When the payoff density is positive and nonincreasing on the active interval, every feasible pooled excursion increases completion time. The optimum therefore consists of the compulsory initial batch, smooth adoption along the separating ridge, and the terminal batch.

6.4 When waves help and how to design one

The exact cap criterion is global because a wave replaces an entire ridge segment, and it is nonlocal because its endpoints are linked by the contact equation. When it fails, the remaining task is to locate profitable departures in primitives and translate them into a usable intervention. Fix a plateau rank m. Move the departure endpoint a small distance $\varepsilon > 0$ to the left of m and use the contact equation to determine the corresponding reentry endpoint to the right. This locally paired excursion is the *reflected wave*. Let f denote the density of good-state adoption payoffs, let $v = \nu(m)$ be the adoption payoff at rank m, and write $f = f(v)$ and $f' = f'(v)$. Its gain has the expansion

$$\Gamma_F = K_F(m)\varepsilon^3 + O(\varepsilon^4), \tag{32}$$

where

$$K_F(m) = \frac{2rf - \beta m v f' + 2\beta r v f^2/(\beta m + r)}{3\beta^2 m^3 v^2 f^3}. \tag{33}$$

The denominator is positive. A small wave is profitable when the density rises rapidly enough:

$$\frac{f'(v)}{f(v)} > \frac{2r}{\beta m v} + \frac{2rf(v)}{m(\beta m + r)}. \tag{34}$$

Theorem 4 (Local primitive test for profitable batching)**.** *Let C be a compact subinterval of the interior of the active rank interval on which the contact equation locally determines the regular reflected reentry endpoint used in the expansion, and suppose the remainder in* (32) *is uniform on C.*

(i) If $K_F \leq -\eta < 0$ on C, every point of C supports sufficiently small exact waves with $\Gamma_F < 0$. Capped availability is therefore not optimal.

(ii) If $K_F \geq \eta > 0$ on C, every sufficiently small exact wave centered in C has $\Gamma_F > 0$.

The full proof is in Section **??** of the Mathematical Supplement, subsection "Local cubic coefficient and a uniform local pooling condition." Lemma **??** constructs the paired endpoint, Theorem **??** derives K_F, and Lemma **??** together with Proposition **??** proves the uniform sign conclusion.

Figure 5 illustrates the economic force behind the local test. When the density rises sharply near the current marginal payoff, a large mass of agents lies just below the margin. A departure batch advances part of this dense group and raises the adoption stock. The principal then pauses new service while the larger stock generates information, before a reentry batch returns the rollout to the separating path. The target can consequently be reached before it would be under capped availability. The figure is schematic: rising density creates the force favoring a wave, but profitability still requires the exact contact condition and the inequality in (34).

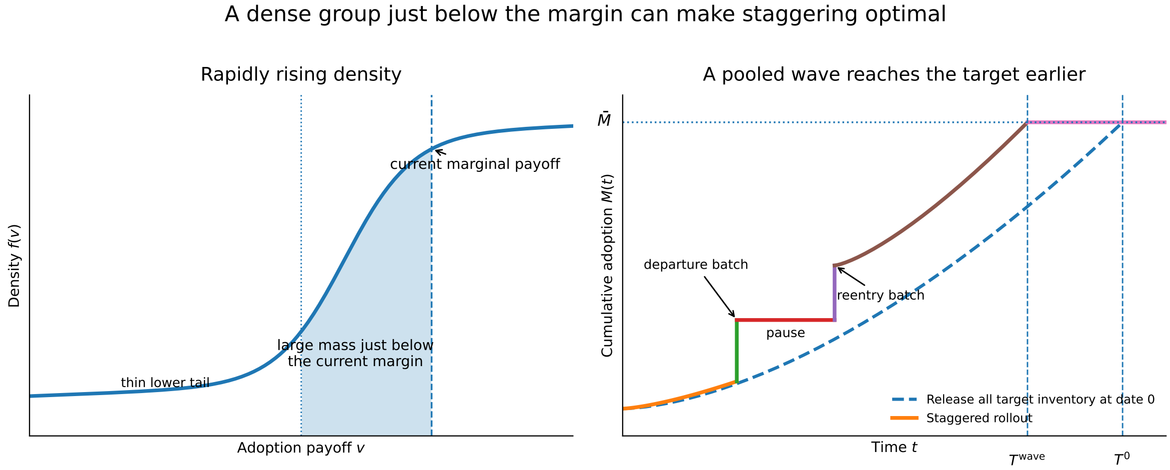

Figure 5: Rapidly rising density and a profitable pooled wave. The left panel shows a dense group of payoffs just below the current margin. In the right panel, a departure batch raises the experimentation stock, a pause allows the additional information to accumulate, and a reentry batch returns the rollout to the separating path. Under the local condition in (34) and exact contact, the resulting wave reaches the target earlier than capped availability.

The local test identifies where gains can arise; the global theorem determines which compatible gains can be selected together. For a single intervention, however, the choice is especially transparent.

Proposition 8 (Optimal one-wave design)**.** *Suppose F has a positive, continuously differentiable density on the active interval. Let $T^{1\mathrm{W}}$ denote the minimum completion time among schedules containing at most one pooled excursion. Then*

$$T^{1\mathrm{W}} = T^{\mathrm{cap}} + \min_{(d,p,h)\in\overline{\mathcal{E}}_F} \Gamma_F(d,p,h), \tag{35}$$

and a minimizing triple exists. If its gain is negative, the departure mass q_D^*, waiting time Δ^*, and reentry mass q_R^* are

$$q_D^* = p^* - d^*, \qquad \Delta^* = \frac{1}{\beta p^*} \log \frac{b(d^*)}{b(h^*)}, \qquad q_R^* = h^* - p^*. \tag{36}$$

If $d^* = m_0$, the departure block merges with the initial service pool; if $h^* = m_1$, the reentry block merges with the terminal service pool.

Once a profitable wave has been identified, the compact program converts the qualitative statement that "batching helps" into an operational prescription. The principal chooses where to leave the ridge, how much mass to advance, and where to reenter; matching the stock–belief endpoint then pins down the pause. For a fixed adoption plateau p, the contact equation has a reflected right endpoint, so the numerical problem is two-dimensional. Section **??** of the Mathematical Supplement derives the reflected-contact geometry and uniform local expansion; the one-wave first-order conditions follow from the compact program above. To first order, the departure and reentry blocks are equal, and the pause is proportional to their common scale and inversely proportional to the plateau stock. These policy quantities give an operational language for the staged-access practices used in actual launches.

7 Applications and institutional interpretation

TAVR provides an institutional illustration of the mechanism because its rollout combined staged eligibility, heterogeneous treatment benefits, capacity constraints, and public evidence generated through patient outcomes.

7.1 TAVR: staged clinical eligibility under registry learning

Transcatheter aortic valve replacement (TAVR) treats severe aortic stenosis without open surgical valve replacement. Its rollout combined staged access with organized evidence generation. The salient uncertainty during diffusion concerned safety, durability, and comparative performance. In TAVR, the main cross-patient linkage operates through information rather than through a direct effect of one patient's treatment on another patient's clinical payoff. Access was also administrative. Medicare coverage was tied to an FDA-approved indication, qualified hospitals and heart teams, and participation in a prospective national registry under Coverage with Evidence Development (Centers for Medicare & Medicaid Services, 2012, 2019). Early use thus generated the evidence on which broader use depended, and the coverage architecture required those outcomes to be recorded.

Consistent with that evidence-generating architecture, the rollout proceeded through explicit clinical-risk expansions. U.S. access began with patients who were inoperable or at extreme surgical risk, expanded to high-risk patients, then to intermediate-risk patients in 2016, and finally to low-risk patients in 2019 (Carroll et al., 2020; U.S. Food and Drug Administration, 2016, 2019). The central staging margin was eligibility under the device label and coverage rules rather than announced risk-specific patient prices. The expansion from higher-risk to lower-risk

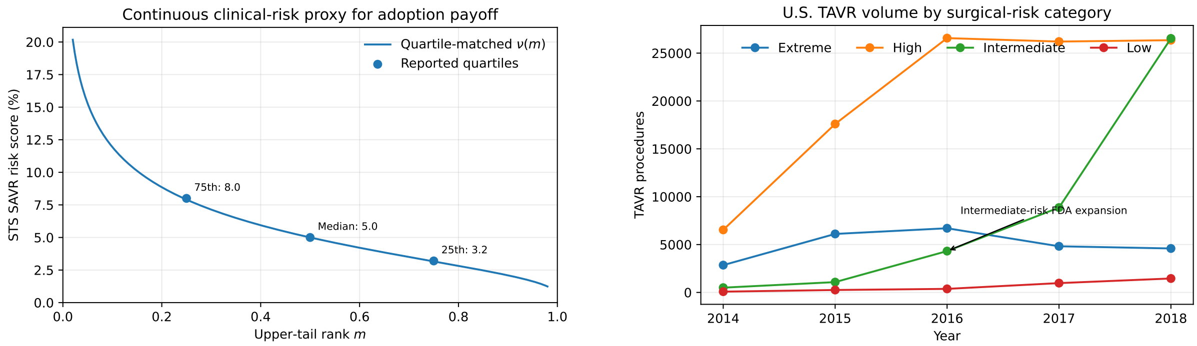

(a) Quartile-matched continuous risk proxy.

(b) TAVR procedures by surgical-risk category.

Figure 6: TAVR risk heterogeneity and indication expansion. Panel (a) fits a continuous distribution to the 2018 STS SAVR-risk-score quartiles reported in Bavaria (2019). Panel (b) replots the risk-category volumes reported in the same registry slide set. Source: author's calculations from Bavaria (2019). The vertical ordering in panel (a) is a proxy for incremental clinical value, not a claim that the risk score is itself utility or willingness to pay.

patients followed accumulating evidence on safety and efficacy, including the evidence supporting the 2019 low-risk indication. This sequence is consistent with staged eligibility making lower-risk patients more willing to accept TAVR after earlier use generated reassuring public evidence.

Because the eligibility sequence followed clinical risk, a patient's surgical risk provides a continuous proxy for the model's adoption payoff v. In this application, patient type corresponds to the incremental clinical benefit of TAVR relative to the available alternative. Because TAVR avoids open surgery, that benefit should generally rise with the expected risk of surgical aortic valve replacement (SAVR), although anatomy, life expectancy, and treatment preferences also matter. Patients with larger expected clinical gains are represented as higher types.

The mapping is not limited to the coarse labels used by the approval process, because the registry reports a continuous STS SAVR risk score. Among TAVR patients in 2018, the reported 25th percentile, median, and 75th percentile were 3.2, 5.0, and 8.0 percent, respectively (Bavaria, 2019). Panel a fits a lognormal distribution to those three quartiles and displays its upper-tail quantile $\nu(m) = F^{-1}(1 - m)$, the primitive object used in the model. This is a descriptive fit to treated patients, not an estimate of the payoff distribution among all eligible patients. Panel b shows the corresponding expansion in realized use. Intermediate-risk procedures increased from 498 in 2014 to 26,537 in 2018, while low-risk procedures remained comparatively rare before the 2019 label expansion (Bavaria, 2019).

Taken together, the registry and eligibility history establish continuous heterogeneity, staged access, and systematic evidence collection. They do not, however, identify the eligible-population distribution needed to sign Γ_F. TAVR therefore provides a close institutional application rather than a calibration of the optimal wave. FDA labels, Medicare coverage, registry requirements, and, in other settings, publicly announced waitlists can make the required access commitments credible.

8 Discussion and conclusion

The paper shows how committed scarcity can accelerate learning precisely because guaranteed future access creates strategic delay. A target-sized date-zero cap reaches the target immediately with one adoption payoff. With continuous heterogeneity, exact wave weights convert design within the ordered representation into interval selection. The underlying decomposition separates ridge service from pooled interventions and the boundary service they inherit, while anonymous staircase capacity implements every finite selected family. Density shape then determines whether capped availability survives the comparison with waves. Faster completion shifts the burden of experimentation toward agents advanced into departure pools and benefits later followers; carrying unused capacity forward is without loss when there are finitely many dated releases.

TAVR illustrates the institutional combination of staged eligibility, continuous clinical heterogeneity, systematic evidence collection, and credible commitment. Although registry summaries do not sign the optimal-wave statistic, they show why access timing is a meaningful design instrument in this setting. The baseline holds fixed commitment, beliefs about future access, and the allocation technology to isolate how supply timing corrects strategic waiting and organizes learning across heterogeneous adopters. Relaxing those assumptions would identify the value of commitment, uncertain future access, and richer rationing rules. Data on applications and remaining inventory would permit direct tests of the waiting and scarcity mechanisms.

References

Bavaria, J. E. (2019). TAVR update: New insights and perspectives from the U.S. National STS/ACC TVT Registry. STS/ACC TVT Registry presentation, October 24, 2019.

Carroll, J. D., M. J. Mack, S. Vemulapalli, H. C. Herrmann, T. G. Gleason, G. Hanzel, G. M. Deeb, V. H. Thourani, D. J. Cohen, N. Desai, A. J. Kirtane, S. Fitzgerald, J. Michaels, C. Krohn, F. A. Masoudi, R. G. Brindis, and J. E. Bavaria (2020). STS-ACC TVT Registry of transcatheter aortic valve replacement. *Journal of the American College of Cardiology* 76, 2492–2516.

Centers for Medicare & Medicaid Services (2012). National coverage determination for transcatheter aortic valve replacement under Coverage with Evidence Development. National Coverage Determination 20.32, May 1, 2012.

Centers for Medicare & Medicaid Services (2019). Decision memo for transcatheter aortic valve replacement. National Coverage Analysis CAG-00430R, June 21, 2019.

Chamley, C. and D. Gale (1994). Information revelation and strategic delay in a model of investment. *Econometrica* 62, 1065–1085.

Chen, W. and Q. Zhang (2025). Jump-start or Gradualism? Dynamic Incentives for Innovation Adoption. Working paper, September 2025.

Condorelli, D. (2012). What money can't buy: Efficient mechanism design with costly signals. *Games and Economic Behavior* 75, 613–624.

Connolly, N. (2021). Government: Vaccine supplies will be used up before expiration. *Cayman Compass*, June 1, 2021.

de Janvry, A., E. Sadoulet, and S. Villas-Boas (2010). Short on shots: Are calls for cooperative restraint effective in managing a flu vaccines shortage? *Journal of Economic Behavior & Organization* 76, 209–224.

DeGraba, P. (1995). Buying frenzies and seller-induced excess demand. *RAND Journal of Economics* 26, 331–342.

Frick, M. and Y. Ishii (2024). Innovation adoption by forward-looking social learners. *Theoretical Economics* 19, 1505–1541.

Google (2023). Try Bard and share your feedback. Company announcement, March 21, 2023.

Guo, Y. and J. Hörner (2020). Dynamic allocation without money. TSE Working Paper 20-1133, August 2020.

Keller, G. and S. Rady (2015). Breakdowns. *Theoretical Economics* 10, 175–202.

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

Laiho, T., P. Murto, and J. Salmi (2024). Slow social learning: Innovation adoption under network externalities. Working paper, July 2024.

Laiho, T., Murto, P., and Salmi, J. (2026). Slow social learning: Innovation adoption under network externalities. Working paper, May 6.

Laiho, T., P. Murto, and J. Salmi (2025). Gradual learning from incremental actions. *Theoretical Economics* 20, 93–130.

Laiho, T. and J. Salmi (2026). Coasian dynamics and endogenous learning. *Journal of Economic Theory* 232, 106136.

Lyu, C. (2026). Information design for social learning on a recommendation platform. *Journal of Economic Theory* 232, 106150.

Microsoft (2023). Reinventing search with a new AI-powered Microsoft Bing and Edge, your copilot for the web. Company announcement, February 7, 2023.

Möller, M. and M. Watanabe (2010). Advance purchase discounts versus clearance sales. *Economic Journal* 120, 1125–1148.

Nocke, V. and M. Peitz (2007). A theory of clearance sales. *Economic Journal* 117, 964–990.

OpenAI (2023). ChatGPT plugins. Company announcement, March 23, 2023.

Parakhonyak, A. and N. Vikander (2023). Information design through scarcity and social learning. *Journal of Economic Theory* 207, 105586.

Pavan, A., I. Segal, and J. Toikka (2014). Dynamic mechanism design: A Myersonian approach. *Econometrica* 82, 601–653.

Radio Cayman (2021). COVID-19 vaccination schedule, June 2021. Public Health Department announcement, May 31, 2021.

Rogers, E. M. (2003). *Diffusion of Innovations*, 5th ed. Free Press.

Ryan, B. and N. Gross (1943). The diffusion of hybrid seed corn in two Iowa communities. *Rural Sociology* 8, 15–24.

Schnepf, J. (2022). Does perceived scarcity of COVID-19 vaccines increase vaccination willingness? Results of an experimental study with German respondents in times of a national vaccine shortage. *PLOS ONE* 17, e0273441.

U.S. Food and Drug Administration (2016). Summary of safety and effectiveness data: Edwards SAPIEN 3 transcatheter heart valve, intermediate-risk indication. PMA Supplement P140031/S010, August 18, 2016.

U.S. Food and Drug Administration (2019). Summary of safety and effectiveness data: Edwards SAPIEN 3 and SAPIEN 3 Ultra, low-risk indication. PMA Supplement P140031/S085, August 16, 2019.

Young, H. P. (2009). Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. *American Economic Review* 99, 1899–1924.

Appendix

A Essential proofs and the finite-type boundary

This appendix records the finite-type boundary and proves the ordered-rollout decomposition, exact interval-selection reduction, countable-closure result, nonincreasing-density characterization, and welfare-incidence result. The Mathematical Supplement supplies the operational foundations, compactness and recovery results, reflected-contact expansion and uniform remainder, and longer algebraic certificates on which some premises below rest.

A.1 Operational preliminaries and finite release bounds

With cumulative supply S_t, prior adoption M_t^-, and applicant mass A_t, anonymous rationing admits each applicant with probability

$$Q_t = \begin{cases} 0, & M_t^- = S_t, \\ 1, & M_t^- < S_t \text{ and } A_t \le S_t - M_t^-, \\ (S_t - M_t^-)/A_t, & A_t > S_t - M_t^- > 0. \end{cases}$$

Rejected agents remain eligible. Fix any rule governing an applicant's future access after an unsuccessful application. Expected payoff under that rule is affine in the adoption payoff v. Willingness to apply is therefore an upper set: whenever one type is willing to apply with positive acceptance probability, every higher type strictly prefers the same application opportunity and continuation. This single-crossing property explains the payoff order represented by upper-tail rank in the continuum analysis.

For finitely many dated releases, destroying or withdrawing unused on-path capacity has no selected-equilibrium value. Given an outcome generated by expiring supply, retain at each date the cumulative amount actually served by then and carry every unused unit forward. Acceptance probabilities and continuation payoffs on the selected path are unchanged. This proves the persistence statement in the text. It does not imply immediate exhaustion or a bound on the number of distinct exhausted releases. Section **??** of the Mathematical Supplement, "Persistence of finite-batch supply," gives the full argument, including deviations after unsuccessful applications, in Proposition **??** there.

A.2 Two types and the three-type boundary

Let $v_H > v_L$, let the high-type mass be q_H, and suppose $v_H > B_0 > v_L$ and $q_H < \bar{M} < 1$. Terminal participation binds at $B_T = v_L$. Before T, only high types can be served. Fix any selected optimum with finitely many releases. Let a be its first high-type service date, put $m = M_{a^-}$ and $x = M_{T^-} - m$, and let $E = \int_a^T M_s \, ds$ be the exposure accumulated before terminal service. Define

$$z_0 = \frac{E}{T-a} - m, \qquad \widehat{T}(z) = a + \frac{E}{m+z}, \qquad z \in [z_0, x].$$

Replacing all intermediate high-type service by mass z at a and moving the residual high types to the terminal batch preserves the terminal posterior. At z_0, the replacement terminal date is T, and the original high type's option to wait until the certain terminal batch implies that its replacement continuation is no better than service at a. Continuity therefore yields either $z = x$ or a smallest z at which the high type is indifferent. The replacement is incentive compatible and completes weakly earlier. Compactness of the two-release parameter set then yields the following boundary result.

Theorem 5 (Two-type unrestricted optimum)**.** *With two target-relevant payoff types, atomless within-type populations, repeated eligibility after rejection, anonymous rationing, asymmetric strategies, principal-preferred equilibrium selection, and target $\bar{M} < 1$, the principal has an optimal finite supply plan with at most two positive releases. Both releases can be chosen to be immediately exhausted.*

Section **??** of the Mathematical Supplement, "Two-type unrestricted optimum," proves the theorem in full; see Theorem **??** there. It establishes terminal myopia, the no-shift inequality, the exposure-preserving consolidation, and attainment of the two-release optimum. The preceding paragraph is the proof roadmap used here.

The conclusion fails with three target-relevant types. For

$$(v_1, v_2, v_3) = (0.86, 0.8, 1/3), \quad (q_1, q_2, q_3) = (0.2, 0.4, 0.4),$$

$$B_0 = 0.75, \qquad \beta = 0.5, \qquad r = 0.4, \qquad \bar{M} = 0.98,$$

a feasible four-release plan completes at 4.1402437122. It serves type 1 at zero, offers type 2 a rationed opportunity at $s = 1.7489035062$ with acceptance probability 0.1104189201, clears rejected type-2 agents at $t = 2.2062097487$, and completes at $T = 4.1402437122$. Direct substitution verifies the relevant type-1, type-2, and type-3 incentive constraints.

Every three-release immediately exhausted plan reduces, by upper-block willingness and terminal service of type 3, to a program in the initial type-1 service x, the middle date s, and its acceptance probability p. Two elementary bounds eliminate plans that serve type 2 at zero or never serve type 2 before the terminal date. On the remaining domain, full type-2 application is without loss. Writing

$$m = x + p(0.6 - x), \qquad B_s = 0.75e^{-0.5xs},$$

completion and terminal acceptance are

$$T(x, s, p) = s + \frac{\log(B_s/v_3)}{0.5m}, \qquad Q_T(x, s, p) = \frac{0.98 - m}{1 - m}.$$

Let $g_n = e^{-0.4s}(v_n - B_s)$, $h_n = Q_T e^{-0.4T}(v_n - v_3)$, and $C_n = pg_n + (1-p)h_n$. The reduced program minimizes $T(x, s, p)$ over $0 \le x \le 0.2$, $s \ge 0$, and $0 \le p \le 1$, subject to

$$g_2 \ge h_2, \qquad C_1 \le 0.11, \qquad C_2 \ge 0.05.$$

For fixed (x, s), completion is decreasing in p. On the relevant sublevel set, the largest feasible p is therefore determined by type-2 participation, type-1 incentive compatibility, or $p = 1$. The certified branch minima are

Binding branch	Lower bound on completion time
Type-2 participation and type-1 IC	4.1700003382
Full middle acceptance, $p = 1$	4.2412342351
Remaining boundaries	> 4.1700003382

Section **??** of the Mathematical Supplement, "Three-type construction," derives the scalar equations on each branch. Subsection **??**, "Certified three-release benchmark," reports the root-isolation intervals and derivative certificates. The table is therefore an analytic lower-bound certificate, with computation used only to certify the displayed roots and signs.

Theorem 6 (Three-type release-count boundary)**.** *For the displayed primitives, every plan with at most three immediately exhausted releases has completion time at least* 4.1700003382*, while the feasible four-release plan completes at* 4.1402437122*. The two-type release-count theorem therefore does not extend to three target-relevant payoff types.*

Proof roadmap. The remaining Appendix proves the paper's central mathematical chain directly. Boundary normalization first pins down the compulsory initial and terminal pools. Continuous service then lies on the unique ridge, while every waiting interval is represented by an exact pooled excursion with virtual ridge endpoints. This yields the global decomposition and constructive staircase implementation. Exact excursion durations telescope over any finite compatible family, giving the interval-selection reduction. The final steps establish countable attainment, finite approximation, the nonincreasing-density characterization, and welfare incidence. Section **??** of the Mathematical Supplement supplies the completed-graph compactness and finite-recovery foundations. Section **??**, under "Local cubic coefficient and a uniform local pooling condition," supplies the reflected-contact construction and uniform fourth-order remainder used in finite compression and in the local profitability theorem.

B Ordered service schedules and exact excursion decomposition

Let $m \in [0, \bar{M}]$ index served agents from highest to lowest payoff, so the marginal payoff is $\nu(m) = F^{-1}(1 - m)$. Single crossing in the anonymous finite-batch model motivates a weakly increasing, right-continuous rank-to-date map $\tau : [0, \bar{M}] \to [0, T]$, which assigns the certain service date $\tau(m)$ to rank m. Flat intervals of τ pool different payoff types at one date. The representation does not cover rules that alter an agent's future access according to her own application and rationing history, because those rules can generate nondegenerate service-time distributions even under pure application strategies. Let λ denote Lebesgue measure on $[0, \bar{M}]$. The schedule induces

$$M_t = \lambda\{m \in [0, \bar{M}] : \tau(m) \le t\}, \qquad B_t = B_0 \exp\left\{-\beta \int_0^t M_s \, ds\right\}.$$

For each service date t, write

$$L_t(v) = e^{-rt}(v - B_t).$$

Global direct IC requires $L_{\tau(m)}(\nu(m)) \geq L_{\tau(n)}(\nu(m))$ for every served m, n, together with participation for served types and exclusion for types below $\nu(\bar{M})$.

Define

$$b(m) = \frac{r\nu(m)}{r + \beta m}, \qquad x(m) = -\log b(m),$$

and

$$\mathcal{S}(u, v) = \frac{1}{\beta} \int_u^v \frac{x'(m)}{m} \, dm.$$

Because ν is positive and strictly decreasing, b is strictly decreasing.

B.1 Boundary states and continuous service

Lemma 3 (Boundary normalization and terminal zero rent)**.** *Suppose $\bar{M} \in (0, 1)$ and*

$$\nu(\bar{M}) < B_0 < \nu(0).$$

Let τ be a finite-completion, globally incentive-compatible, ordered schedule serving exactly the upper $\bar{M}$ mass, with participation for served types and exclusion below $\nu(\bar{M})$. Put $t_0 = \tau(0)$ and let $T = \tau(\bar{M})$.

Subtracting t_0 from every service date preserves feasibility and reduces completion time by t_0. Also,

$$B_T = \nu(\bar{M}). \tag{37}$$

Consequently, there are unique ranks $0 < m_0 < m_1 < \bar{M}$ satisfying

$$b(m_0) = B_0, \qquad b(m_1) = \nu(\bar{M}). \tag{38}$$

Proof. No adoption occurs before t_0, so beliefs remain equal to B_0 on $[0, t_0)$. Shifting every service date backward by t_0 therefore leaves the belief attached to every corresponding service line unchanged and multiplies all dated payoffs by the common positive factor e^{rt_0}. Global IC, participation, and exclusion are preserved.

Let $v_* = \nu(\bar{M})$ be the marginal served type. Its assigned service line is the terminal line, because τ is ordered. Participation gives $L_T(v_*) \geq 0$. Every $v < v_*$ is excluded, so $L_T(v) \leq 0$. Letting $v \uparrow v_*$ yields $L_T(v_*) \leq 0$. Hence $L_T(v_*) = 0$, which is (37).

Finally, $b(0) = \nu(0)$ and

$$b(\bar{M}) = \frac{r\nu(\bar{M})}{r + \beta \bar{M}} < \nu(\bar{M}).$$

Strict monotonicity of b gives unique solutions to (38). Since $B_0 > \nu(\bar{M})$, the first solution lies strictly below the second. □

We call $[m_0, m_1]$ the *active rank interval* and $[\nu(m_1), \nu(m_0)]$ the corresponding *active payoff interval.*

Lemma 4 (Continuous service lies on the ridge)**.** *Under the hypotheses of Lemma 3, let μ_c denote the continuous part of the Stieltjes measure $d\tau$. At every continuity point m of τ in* $\operatorname{supp}\mu_c$*,*

$$B_{\tau(m)} = b(m). \tag{39}$$

The measure μ_c is absolutely continuous and is supported on $[m_0, m_1]$. Almost everywhere on its increasing part,

$$\tau'(m) = \frac{x'(m)}{\beta m}. \tag{40}$$

In particular, τ has no singular-continuous component.

Proof. Take a continuity point $m \in \operatorname{supp}\mu_c$. For each n, choose $a_n < m < b_n$ with $a_n, b_n \to m$ and $\tau(b_n) > \tau(a_n)$. Let c_n be the intersection of the service lines at these two dates. Their slopes are ordered, and global IC gives

$$\nu(b_n) \le c_n \le \nu(a_n),$$

so $c_n \to \nu(m)$. Put $\Delta_n = \tau(b_n) - \tau(a_n)$. Continuity gives $\Delta_n \to 0$. During the corresponding time interval cumulative adoption lies between a_n and b_n, hence its time average is $m + o(1)$ and

$$B_{\tau(b_n)} = B_{\tau(a_n)} \exp\{-\beta(m + o(1))\Delta_n\}.$$

The intersection formula is

$$c_n = \frac{B_{\tau(a_n)} - e^{-r\Delta_n} B_{\tau(b_n)}}{1 - e^{-r\Delta_n}}.$$

Expanding at $\Delta_n = 0$ and passing to the limit gives

$$\nu(m) = \frac{r + \beta m}{r} B_{\tau(m)},$$

which is (39).

Beliefs never exceed $B_0 = b(m_0)$ and never fall below the terminal belief $b(m_1)$. Hence (39) and strict monotonicity of b imply

$$\operatorname{supp}\mu_c \subseteq [m_0, m_1].$$

Let E be the set of continuity points in $\operatorname{supp}\mu_c$; the complement of E is countable and therefore μ_c-null. For $a < c$ in E, the belief law and (39) give

$$x(c) - x(a) = \beta \int_{\tau(a)}^{\tau(c)} M_s \, ds \ge \beta a\{\tau(c) - \tau(a)\}.$$

Since $a \ge m_0 > 0$ and x' is bounded on $[m_0, m_1]$, there is a finite constant

$$C = \frac{\sup_{[m_0,m_1]} x'}{\beta m_0}$$

such that $\tau(c) - \tau(a) \le C(c - a)$ for all $a < c$ in E.

Fix an interval $I = (u, v) \subseteq [m_0, m_1]$. If $\mu_c(I) = 0$, there is nothing to prove. Otherwise, because E has full μ_c-measure and μ_c is atomless, one can choose $a_n, c_n \in E \cap I$ with $a_n < c_n$ such that

$$\mu_c((u, a_n]) + \mu_c([c_n, v)) \longrightarrow 0.$$

Then

$$\mu_c(I) = \lim_{n\to\infty} \mu_c((a_n, c_n)) \le \liminf_{n\to\infty}\{\tau(c_n) - \tau(a_n)\} \le C(v - u).$$

Thus $\mu_c \ll \lambda$, ruling out a singular-continuous component.

At almost every point with $\tau'(m) > 0$, differentiating (39) and using $d[-\log B_t]/dt = \beta M_t$ gives

$$x'(m) = \beta m \tau'(m),$$

which is (40). □

Flat intervals of τ are batches. Jumps of τ are intervals of calendar time during which no new rank is served and the existing adoption stock continues to generate information. The next results identify every such jump with an actual pooled excursion.

B.2 Exact crossing, contact, and pooled blocks

Fix

$$e = (m_d, p, m_r), \qquad m_d < p < m_r,$$

and put $a = \nu(p)$ and $k = r/(\beta p)$. The excursion leaves the ridge at m_d, immediately releases mass $p - m_d$, holds cumulative adoption fixed at p, and rejoins the ridge at m_r. Its waiting time is

$$\Delta(e) = \frac{1}{\beta p} \log \frac{b(m_d)}{b(m_r)}. \tag{41}$$

The pivot type a is indifferent between the departure and reentry slots when

$$a - b(m_d) = e^{-r\Delta(e)}\{a - b(m_r)\}. \tag{42}$$

Equivalently, with

$$q_p(m) = \log\left[b(m)^k\{a - b(m)\}\right],$$

the contact condition is

$$q_p(m_d) = q_p(m_r). \tag{43}$$

Lemma 5 (Exact excursion feasibility)**.** *Suppose ν is positive, continuously differentiable, and strictly decreasing on a neighborhood of $[m_d, m_r]$. If (43) holds and the contact state lies in the interior ridge range, then the excursion clock is uniquely given by (41). The departure and reentry payoff lines cross exactly at $a = \nu(p)$, and the complete dated menu is globally incentive compatible for all participating types.*

Proof. Strict monotonicity of b makes (41) the unique positive clock connecting the two ridge beliefs while cumulative adoption is fixed at p. Substituting this clock into the equality of the two dated payoffs gives (42), so their unique affine intersection is a.

Every dated slot generates an affine line $L_j(v) = e^{-rt_j}(v - B_j)$ whose slope decreases strictly with its date. If c_j is the intersection of L_j and L_{j+1} and $c_1 > \cdots > c_{J-1}$, then

$$L_j(v) - L_{j+1}(v) = (e^{-rt_j} - e^{-rt_{j+1}})(v - c_j)$$

shows that the lines form the upper envelope in the required order. The smooth prefix meets the departure line at $\nu(m_d)$, the two excursion lines meet at $\nu(p)$, and the reentry line meets the shifted smooth continuation at $\nu(m_r)$. Since $m_d < p < m_r$ and ν is strictly decreasing,

$$\nu(m_d) > \nu(p) > \nu(m_r).$$

All other adjacent intersections retain their ridge order because the prefix is unchanged and every line in the continuation is shifted by the same calendar-time increment. The complete list of intersections is therefore ordered, so the affine lines form the global upper envelope. □

Lemma 6 (Every waiting interval has virtual ridge endpoints)**.** *Under the hypotheses of Lemma 3, let $p \in (0, \bar{M})$ be a jump point of τ, and write*

$$t^- = \tau(p^-), \qquad t^+ = \tau(p), \qquad \Delta = t^+ - t^- > 0.$$

Let $B^- = B_{t^-}$ and $B^+ = B_{t^+}$. There are unique $d < p < h$ such that

$$B^- = b(d), \qquad B^+ = b(h),$$

and the clock and crossing equations of the exact excursion (d, p, h) hold.

Proof. During the wait, cumulative adoption is p, so

$$B^+ = B^- e^{-\beta p \Delta}. \tag{44}$$

Taking one-sided limits of the IC inequalities at rank p makes the adjacent service lines intersect at $a = \nu(p)$:

$$a - B^- = e^{-r\Delta}\{a - B^+\}. \tag{45}$$

The right-continuous convention assigns rank p to the later line. The participation constraint for rank p, together with (45), implies $a - B^+ > 0$ and $a - B^- > 0$.

By Lemma 3, every service belief lies in $[b(m_1), b(m_0)]$. Strict monotonicity of b therefore gives unique $d, h \in [m_0, m_1]$ with the stated beliefs. Put $k = r/(\beta p)$. Equations (44)–(45) imply

$$(B^-)^k(a - B^-) = (B^+)^k(a - B^+).$$

On $(0, a)$, the function $z \mapsto z^k(a - z)$ has a unique maximum at

$$\frac{k}{k+1}a = b(p).$$

Since $B^- > B^+$ and the two function values agree, the beliefs lie on opposite sides of the maximum:

$$B^- > b(p) > B^+.$$

As b is strictly decreasing, $d < p < h$. Equations (44)–(45) are then exactly (41)–(42). □

Lemma 7 (A jump is implemented by its pooled blocks)**.** *For each jump at p with virtual endpoints $d < p < h$,*

$$\tau(m) = t^- \quad \text{for } d < m < p, \qquad \tau(m) = t^+ \quad \text{for } p < m < h. \tag{46}$$

If jump i precedes jump j, their virtual supports are ordered:

$$h_i \le d_j. \tag{47}$$

Thus every jump is an actual departure batch, a wait, and a reentry batch, and distinct excursion supports have disjoint interiors.

Proof. If jump i precedes jump j, beliefs are weakly lower at the later jump, so

$$B_i^+ \ge B_j^-.$$

Since $B_i^+ = b(h_i)$, $B_j^- = b(d_j)$, and b is strictly decreasing, this is (47).

Consider one jump (d, p, h). Another jump with pivot $q \in (d, p)$ would occur earlier and hence satisfy $h_q \le d$. But every jump has $q < h_q$, contradicting $q > d$. Likewise, a jump with pivot $q \in (p, h)$ would occur later and satisfy $h \le d_q < q$, contradicting $q < h$. Thus there are no other jumps inside either interval in (46).

There is also no continuous increase in (d, p). If a continuity point m in that interval belonged to $\operatorname{supp} \mu_c$, then $\tau(m) \le t^-$ and therefore $B_{\tau(m)} \ge B^- = b(d)$. Lemma 4 would instead require $B_{\tau(m)} = b(m) < b(d)$, a contradiction. The argument on (p, h) is symmetric: there $\tau(m) \ge t^+$ and hence $B_{\tau(m)} \le B^+ = b(h)$, whereas ridge service would require $b(m) > b(h)$.

The Stieltjes measure $d\tau$ has neither atoms nor continuous mass on (d, p) and (p, h). Hence τ is constant on each interval, and the one-sided endpoint values give (46). □

B.3 Global decomposition of arbitrary monotone schedules

Theorem 7 (Ordered-rollout decomposition)**.** *Assume the boundary and global-IC hypotheses above.*

Define the time-normalized schedule by

$$\bar{\tau} = \tau - t_0.$$

Then:

(i) $\bar{\tau}(m) = 0$ on $[0, m_0]$ and $\bar{\tau}(m) = T - t_0$ on $[m_1, \bar{M}]$;

(ii) there is a finite or countable chronological family of exact excursions

$$\mathcal{A} = \{(d_j, p_j, h_j)\}_{j\in J}$$

with pairwise disjoint support interiors in (m_0, m_1)*;*

(iii) outside those supports, every increase of $\bar{\tau}$ *is absolutely continuous ridge passage;*

(iv) completion time satisfies

$$T = t_0 + \mathcal{S}(m_0, m_1) + \sum_{j\in J} \Gamma_F(d_j, p_j, h_j), \tag{48}$$

where

$$\Gamma_F(d, p, h) = \frac{1}{\beta}\int_d^h x'(m)\left(\frac{1}{p} - \frac{1}{m}\right) dm. \tag{49}$$

The series in (48) *is absolutely convergent.*

Proof. The normalization is feasible by Lemma 3. We first establish the boundary batches. Continuous increase below m_0 is impossible: Lemma 4 would require a service belief $b(m) > b(m_0) = B_0$. A jump at $p < m_0$ is also impossible. Its departure belief satisfies $B^- \le B_0 = b(m_0)$, so its virtual departure rank obeys $d \ge m_0$, contradicting $d < p$. Thus $d\bar{\tau}$ vanishes on $(0, m_0)$ and $\bar{\tau} = 0$ there.

Similarly, continuous increase above m_1 would require $B = b(m) < b(m_1) = B_T$, although beliefs cannot fall below their terminal value before completion. A jump at $p > m_1$ would have $B^+ \ge B_T = b(m_1)$ and therefore $h \le m_1$, contradicting $p < h$. Hence $d\bar{\tau}$ vanishes on $(m_1, \bar{M})$ and the schedule is constant at its completion date there. This proves (i).

Lemma 4 makes the continuous part of $d\bar{\tau}$ absolutely continuous and pins every positive-density increase to the ridge. The atomic part contains at most countably many jumps. Lemmas 6 and 7 associate each jump with an actual exact excursion, order the supports chronologically, and make their interiors disjoint. Any same-date mass at the initial line beyond m_0 is the departure block of the first excursion; any same-date mass at the terminal line below m_1 is the reentry block of the last excursion. These are merely splits of one service line and do not change dates, beliefs, payoffs, or allocations.

There is no uncovered interior batch. Indeed, let $(a, c) \subset (m_0, m_1)$ be a maximal interval on which $\bar{\tau}$ is constant and whose interior is disjoint from all excursion supports. Each endpoint is approached either by continuous ridge service or by an excursion endpoint. In both cases the belief at the endpoint is the corresponding ridge belief. Since no time passes across (a, c), these beliefs are equal, so $b(a) = b(c)$. Strict monotonicity of b forces $a = c$. Thus every positive-mass interior flat block belongs to an excursion, and every remaining part of the interior is ridge passage. This proves (ii) and (iii).

On the complement of the excursion supports, ridge passage takes

$$\frac{1}{\beta}\int \frac{x'(m)}{m}\, dm.$$

Excursion (d_j, p_j, h_j) takes

$$\Delta_j = \frac{1}{\beta p_j}\{x(h_j) - x(d_j)\} = \frac{1}{\beta}\int_{d_j}^{h_j} \frac{x'(m)}{p_j}\, dm,$$

whereas smooth passage through the same ridge states takes $\mathcal{S}(d_j, h_j)$. Countable additivity over the disjoint supports gives (48). Finally,

$$\sum_j \Delta_j \le T - t_0, \qquad \sum_j \mathcal{S}(d_j, h_j) \le \mathcal{S}(m_0, m_1),$$

so

$$\sum_j |\Gamma_F(d_j, p_j, h_j)| \le T - t_0 + \mathcal{S}(m_0, m_1) < \infty.$$

□

Remark 1 (Compulsory boundary pools versus observed boundary batches). The ranks m_0 and m_1 are unique without any restriction on the sign of f'. Every admissible ordered rollout assigns the upper block $[0, m_0]$ to its first service line and the lower target block $[m_1, \bar{M}]$ to its terminal line. These are compulsory boundary pools. They need not equal the full observed boundary jumps: an excursion with departure rank $d = m_0$ adds its departure block to the initial batch, while an excursion with reentry rank $h = m_1$ adds its reentry block to the terminal batch. When all nondegenerate excursions are excluded, as under a nonincreasing density, the compulsory pools are exactly the full initial and terminal batches.

C Compatible-family selection, countable attainment, and finite compression

The compact exact-excursion set $\overline{\mathcal{E}}_F$ obtained from the contact equation, including diagonal null controls (m, m, m), closes the variational step used by the main global theorem. For $e = (d, p, h)$, write

$$\ell(e) = h - d.$$

A family is compatible when the interiors of its support intervals are pairwise disjoint. Let $\mathfrak{A}_f$ and $\mathfrak{A}_c$ denote the finite and countable compatible families, respectively.

Proposition 9 (Finite telescoping and reconstruction). *For every finite compatible family $\mathcal{A}$ of exact excursions,*

$$T(\mathcal{A}) = \mathcal{S}(m_0, m_1) + \sum_{e \in \mathcal{A}} \Gamma_F(e). \tag{50}$$

Conversely, chronological concatenation of ridge passage on uncovered intervals and the exact departure–pause–reentry bridge on each $I(e)$ produces a globally incentive-compatible, ordered schedule with duration (50).

Proof. Smooth time is additive over disjoint mass intervals. Replacing ridge passage on $I(e)$ by its exact excursion changes duration by $\Gamma_F(e)$, so summing over the pairwise disjoint supports gives (50). Conversely, each excursion leaves and returns at the ridge states $b(d)$ and $b(h)$. Exact feasibility orders the adjacent payoff-line intersections, while disjoint support interiors preserve that order across different excursions. Concatenation therefore preserves the global affine upper envelope and returns the continuation to the same ridge state inherited under smooth passage. □

For completeness, the corresponding Bellman equality can be written without assuming that an optimal countable family has a first excursion after every state. Let $J(x)$ be the infimum of total excursion gain from ridge mass x to m_1, and let $\mathcal{E}_d$ be the exact excursions departing at d. Then

$$J(x) = \min\left\{0, \inf_{x \le d < m_1} \inf_{e \in \mathcal{E}_d} [\Gamma_F(e) + J(h(e))]\right\}. \tag{51}$$

Equation (51) is an equality of infima. Immediate accumulation to the right of x may prevent the inner infimum from having a minimizing first action.

Proposition 10 (Uniform finite compression)**.** *Under the C^4 regularity of Corollary 1, there are $C < \infty$ and $\delta_0 > 0$ such that every exact excursion with $\ell(e) \le \delta_0$ satisfies*

$$|\Gamma_F(e)| \le C\ell(e)^3. \tag{52}$$

Let $M = m_1 - m_0$. For any compatible finite or countable family and any $0 < \delta \le \delta_0$, deleting excursions with span below δ changes total gain by at most $CM\delta^2$ and leaves at most M/δ excursions. Hence, for every $\varepsilon > 0$, every compatible family has a finite subfamily within ε of its gain.

Proof. The reflected-contact expansion and uniform fourth-order remainder in Section **??** of the Mathematical Supplement give (52) uniformly on the compact active rank domain. Compatibility implies $\sum_e \ell(e) \le M$. Therefore

$$\sum_{\ell(e)<\delta} |\Gamma_F(e)| \le C \sum_{\ell(e)<\delta} \ell(e)^3 \le CM\delta^2.$$

Disjoint support interiors also imply that at most M/δ excursions have span at least δ. □

Theorem 8 (Countable attainment and finite-staircase approximation)**.** *Suppose x is four times continuously differentiable on the active rank interval. The finite-family infimum is attained in the countable compatible closure:*

$$\inf_{\mathcal{A} \in \mathfrak{A}_f} \sum_{e \in \mathcal{A}} \Gamma_F(e) = \min_{\mathcal{A} \in \mathfrak{A}_c} \sum_{e \in \mathcal{A}} \Gamma_F(e). \tag{53}$$

The minimizing gain series is absolutely convergent. Nested finite subfamilies reconstruct globally incentive-compatible schedules whose cumulative service dates and beliefs converge uniformly to the countable schedule. Consequently, finite anonymous staircase policies attain completion times arbitrarily close to the ordered optimum; a finite staircase attains it whenever an optimal family is finite.

Proof. For a finite family, define the length-weighted marked measure

$$\mu_{\mathcal{A}} = \sum_{e \in \mathcal{A}} \ell(e)\delta_e, \qquad g(e) = \begin{cases} \Gamma_F(e)/\ell(e), & \ell(e) > 0, \\ 0, & \ell(e) = 0. \end{cases}$$

The measures have total mass at most M. Proposition 10 makes g continuous at diagonal null controls, and it is continuous on the remaining compact exact-excursion set. A minimizing sequence therefore has a weakly convergent subsequence.

On every set $\{\ell \geq 1/k\}$, compatibility bounds the number of atoms by kM. A limiting nonnull atom has weight exactly $\ell(e)$: a sufficiently small neighborhood of a nonnull excursion is a conflict neighborhood, so each approximating compatible family contributes at most one atom there. Taking the union over k yields a countable compatible nondegenerate family. Any diffuse residual is supported on $\{\ell = 0\}$, where $g = 0$. This proves attainment and (53). The cubic bound controls the small-span tail, while only finitely many supports exceed any fixed span, so the gain series is absolutely convergent.

Choose nested finite subfamilies whose omitted total support length tends to zero and delete any nonnegative-gain excursion. Replacing omitted excursions by ridge passage changes cumulative dates uniformly by a quantity tending to zero. Endpoint masses and ridge beliefs converge, and dated payoff lines converge uniformly on the compact adoption-payoff range. Passing to the limit in the finite upper-envelope inequalities preserves global incentive compatibility. The anonymous staircase implementation result in the main paper then implements every finite truncation. □

C.1 Nonincreasing density and uniqueness

This subsection supplies the short convex-quantile argument behind Corollary 3. Fix a nondegenerate exact excursion $e = (d, p, h)$, put $a = \nu(p)$ and $k = r/(\beta p)$, and define

$$q(m) = \log\{b(m)^k[a - b(m)]\}, \qquad L(m) = -\frac{b'(m)}{b(m)}.$$

The contact equation gives $q(d) = q(h)$. Direct differentiation yields

$$q'(m) = L(m)R(m), \qquad R(m) = \frac{(k+1)b(m) - ka}{a - b(m)}.$$

Thus $q' > 0$ on (d, p) and $q' < 0$ on (p, h). Write

$$\eta_p(m) = \frac{1}{p} - \frac{1}{m}, \qquad \psi_p(m) = \frac{\eta_p(m)}{R(m)},$$

where $\psi_p(p)$ is defined by continuity.

Lemma 8 (Convex quantiles make ψ_p decreasing)**.** *If ν is convex on $[d, h]$, then ψ_p is strictly decreasing.*

Proof. For $m \neq p$, set $s(m) = -(\nu(m) - a)/(m - p) > 0$. Convexity makes s nonincreasing. Algebra gives $\psi_p = -\phi$, where

$$\phi(m, s) = \frac{\beta}{r} \frac{a\beta + r(1 - p/m)s}{a\beta + (r + \beta p)s}.$$

Writing $D = a\beta + (r + \beta p)s$,

$$\partial_m \phi = \frac{\beta p s}{m^2 D} > 0, \qquad \partial_s \phi = \frac{\beta}{r} \frac{a\beta\{-rp/m - \beta p\}}{D^2} < 0.$$

Since s is nonincreasing, ϕ is strictly increasing and ψ_p is strictly decreasing. The value at p follows by continuous extension. □

Proof of Corollary 3. Because

$$\nu''(m) = -\frac{f'(\nu(m))}{f(\nu(m))^3},$$

the condition $f' \leq 0$ makes ν convex on the active interval. For the excursion above, the gain formula can be written as

$$\beta\Gamma_F(e) = \int_d^h L(m)\eta_p(m)\, dm = \int_d^h q'(m)\psi_p(m)\, dm.$$

Since $q(d) = q(h)$, $\int_d^h q'(m)\, dm = 0$. Hence

$$\beta\Gamma_F(e) = \int_d^h q'(m)\{\psi_p(m) - \psi_p(p)\}\, dm.$$

On (d, p) both factors are positive; on (p, h) both are negative. The product is positive on sets of positive measure, so every nondegenerate excursion has $\Gamma_F(e) > 0$.

The decomposition in Theorem 7 then gives

$$T = t_0 + \mathcal{S}(m_0, m_1) + \sum_j \Gamma_F(e_j) \geq \mathcal{S}(m_0, m_1),$$

with equality only when $t_0 = 0$ and the excursion family is empty. Proposition 3 supplies a feasible excursion-free schedule attaining this bound. With no excursion, the decomposition leaves only the compulsory boundary pools and ridge passage, whose dates are pinned down by $b(m)$; the aggregate path is therefore unique up to zero-mass cutoff conventions. □

C.2 Welfare incidence and scope

Proof of Proposition 7. Normalize the excursion's departure date to zero. For $m \in [d, p]$, excursion service gives

$$U_d^e(m) = \nu(m) - b(d).$$

Under ridge service, the line indexed by m is the unique maximizer for type $\nu(m)$, whereas the departure line is indexed by d. The two coincide at $m = d$ and the ridge line strictly dominates for $m > d$. This proves part (i).

For $m \geq h$, the excursion returns to the same ridge state $(h, b(h))$ as smooth passage and shifts the entire continuation by the excess duration $\Gamma_F(e)$. Beliefs and adoption payoffs are otherwise unchanged, so

$$U_d^e(m) = e^{-r\Gamma_F(e)} U_d^F(m),$$

which proves part (ii).

For $m \in [p, h]$, define the log payoff ratio

$$D(m) = \log \frac{U_d^e(m)}{U_d^F(m)} = -r\Delta + r\mathcal{S}(d, m) + \log \frac{\nu(m) - b(h)}{\nu(m) - b(m)}.$$

Using $\nu(m) - b(m) = \beta m b(m)/r$ and $\mathcal{S}'(d, m) = -b'(m)/(\beta m b(m))$ gives

$$D'(m) = \nu'(m) \left\{ \frac{1}{\nu(m) - b(h)} - \frac{1}{\nu(m) - b(m)} \right\} > 0,$$

because $\nu' < 0$ and $b(m) > b(h)$ for $m < h$. At p, the reentry line equals the departure line by exact contact, while the free ridge line strictly dominates the departure line, so $D(p) < 0$. At h, the two policies deliver the same belief and differ only by the time shift, so $D(h) = -r\Gamma_F(e) > 0$ when $\Gamma_F(e) < 0$. Continuity and strict monotonicity yield the unique cutoff in part (iii). □

Proof of Corollary 2. Starting from an optimal compatible family, delete every nondegenerate excursion with nonnegative gain. Compatibility is preserved and total completion time weakly falls, so an optimal family exists in which every retained nondegenerate excursion has strictly negative gain. Applying Proposition 7 excursion by excursion gives the stated incidence comparison. □

The finite operational results justify persistent anonymous supply and locate the two-type/three-type boundary. Single crossing motivates the payoff-ordered class, and the continuum theorems characterize its rank-to-date schedules and staircase implementations. They do not cover rules that change an agent's future access after her own unsuccessful applications. Such individual-history-dependent access is an additional dynamic-rationing instrument, and determining when it improves on anonymous rank-ordered rollout is a separate question.

Mathematical Supplement to “Staggered Rollout for Innovation Adoption”

Ricardo Fonseca

This Mathematical Supplement gives the full proofs, finite operational foundations, counterexample certificates, compactness arguments, reflected-contact calculations, and finite-staircase approximation omitted from the article Appendix. Terminology follows the article throughout: a *type* is indexed by its good-state adoption payoff.

1 Operational equilibrium and service-time distributions

With S_t denoting cumulative supply, M_t^- prior adoption, and A_t applicant mass, the anonymous-rationing acceptance probability is

$$Q_t = \begin{cases} 0, & M_t^- = S_t, \\ 1, & M_t^- < S_t \text{ and } A_t \le S_t - M_t^-, \\ (S_t - M_t^-)/A_t, & A_t > S_t - M_t^- > 0. \end{cases}$$

Rejected agents remain active. A realized application plan α specifies a sequence of dates contingent on the agent's own past failures; let T^α denote the resulting set of attempted dates. Conditional on the no-signal path, the probability of first success at date $s \in T^\alpha$ is

$$Q_s \prod_{k \in T^\alpha : k < s} (1 - Q_k).$$

The normalized time-zero payoff is therefore

$$U(v; \alpha) = \sum_{s \in T^\alpha} Q_s \prod_{k<s} (1 - Q_k) e^{-rs} (v - B_s).$$

This formula is affine in the adoption-payoff type v. For any two types $v > v'$, the difference between their payoffs from one fixed service-time distribution is

$$(v - v') \sum_s Q_s \prod_{k<s} (1 - Q_k) e^{-rs}.$$

This is the discounted-service moment that underlies upper-block willingness and the averaged-envelope interpretation. In particular, willingness to apply is an upper set in the payoff ranking: if one type is willing, every strictly higher-payoff type is willing as well.[1]

Lemma 1 (Upper-block willingness)**.** *Fix any continuation access conditioned on unsuccessful applications. If type v' weakly prefers applying now to that continuation, every type $v > v'$ strictly prefers applying now, unless the immediate acceptance probability is zero.*

[1]The affine-in-type comparison is the familiar single-crossing logic of screening models; see, for example, Myerson (1981).

Proof. Let A^+ be the continuation lottery after rejection and write its payoff as $av - c$. Applying now with probability $Q > 0$ gives

$$Qe^{-rt}(v - B_t) + (1 - Q)(av - c).$$

The difference from rejecting is $Q[(e^{-rt} - a)v - (e^{-rt}B_t - c)]$. If the difference is nonnegative at v', it is strictly larger at $v > v'$ because every possible service date in A^+ is later than t, so $a < e^{-rt}$. □

2 Persistence of finite-batch supply

The baseline model makes cumulative supply persistent. Allowing a finite collection of released tranches to expire, be withdrawn, or be destroyed does not improve the selected-equilibrium value. This is an on-path outcome-implementation result; the same continuation strategy need not remain optimal after every off-path private history.

Proposition 1 (No gain from on-path destruction)**.** *Consider an enlarged finite-batch policy class with finitely many labeled tranches. Each tranche has a release date and may later expire, be withdrawn, or be destroyed. Suppose a policy in this enlarged class admits a selected equilibrium with a well-defined accounting of on-path service to the labeled tranches. Then there exists a persistent finite-batch policy, using no more release dates, that admits an equilibrium with exactly the same on-path aggregate application, service, adoption, belief, and completion paths. Consequently, allowing expiration, withdrawal, or destruction cannot lower the principal's optimal completion time within the finite-batch class.*

Proof. Fix the selected equilibrium and a feasible accounting of each on-path unit of service to a labeled tranche. For tranche k, let a_k be the total mass allocated from that tranche before its withdrawal. Replace the original tranche by a persistent tranche of size a_k released at the same date; delete it if $a_k = 0$. Keep the same accounting of on-path service.

At every date before the original withdrawal of tranche k, cumulative on-path allocation from that tranche is no greater than a_k, its eventual allocated mass. Hence the trimmed tranche can finance every on-path allocation previously charged to it. By the original withdrawal date, exactly a_k has been allocated, so the trimmed tranche is exhausted. This argument is unaffected by overlapping tranches because the accounting fixes which tranche finances each on-path unit, and the same accounting remains feasible after trimming. Applying the construction to every tranche preserves the complete on-path service path and therefore the adoption, belief, and completion paths.

We next compare acceptance opportunities along that public path. At a date with positive on-path applicant mass, the original equilibrium is either underdemanded, in which case every applicant is served and the trimmed inventory remains sufficient for the same applicant mass, or rationed, in which case all capacity then available is allocated and none of it is later destroyed. The acceptance probability is therefore unchanged at every date with positive on-path applicant mass. At a date with zero on-path applicant mass, one atomless deviator is served with probability one whenever the trimmed policy has positive remaining capacity and with probability zero otherwise. The original policy offers probability one in the former case and weakly more capacity in the latter.

Take any deviation that conditions later applications on earlier unsuccessful attempts under the trimmed persistent policy. Delete every attempted date at which its acceptance probability is zero. At all remaining attempted dates, the original policy gives the same acceptance probability and the same public belief as the trimmed policy. The resulting deviation under the original policy therefore induces exactly the same distribution of service dates and the same payoff. Atomlessness is essential here: one agent's applications do not alter aggregate demand, the public path, or the dates at which the trimmed tranches are exhausted. Thus every payoff achievable by deviating under the persistent policy was already achievable under the original policy.

The prescribed on-path service-time distributions are unchanged, so the original equilibrium behavior remains optimal at every private history reached with positive probability. Retain that behavior on path and choose an optimal continuation after every other private history under the trimmed policy. The maintained best-response existence condition supplies these continuations. This gives an equilibrium implementing the same on-path outcome under persistent supply. □

Remark 1 (Persistence versus immediate exhaustion)*.* Removing destruction does not require every release to be absorbed at its release date. In a finite-type environment, a separate postponement argument can select an optimal representative with immediately exhausted positive releases. With a continuum of payoffs, one persistent tranche may instead be drawn down smoothly over an interval. Persistence, rather than instantaneous depletion, is the relevant without-loss conclusion here.

3 Free supply and aggregate-path uniqueness

Under free supply, successful service occurs at the first application. Here and below, v always denotes the adopter's good-state payoff, not the principal's launch payoff V. Let

$$g_v(t) = e^{-rt}(v - B_t).$$

Differentiation gives

$$\dot{g}_v(t) = e^{-rt}\{\beta M_t B_t - r(v - B_t)\}.$$

On an interior mixed-stopping interval, $\dot{g}_v = 0$, hence

$$B_t = \frac{rv}{r + \beta M_t}. \tag{1}$$

Differentiating the right-hand side and using $\dot{B}_t = -\beta M_t B_t$ yields

$$\dot{M}_t = M_t(r + \beta M_t). \tag{2}$$

The unique solution beginning at $M_a > 0$ is

$$M_t = \frac{rM_a e^{r(t-a)}}{r + \beta M_a(1 - e^{r(t-a)})},$$

up to the time at which the active payoff type is exhausted. Its second derivative is positive.

Proposition 2 (Aggregate uniqueness under free availability)**.** *Apart from knife-edge entry indifferences, free availability induces a unique aggregate no-signal adoption path.*

(a) For a finite distribution of adoption payoffs, payoff types begin adopting in descending order of payoffs. At date zero, all atoms strictly above a uniquely determined marginal atom adopt in full, and a uniquely determined fraction of that marginal atom may also adopt. Any residual mass of the marginal atom then follows (2)*. Each lower atom enters after a uniquely determined waiting interval and follows its own unique ridge segment until exhausted.*

(b) Suppose F has a positive, continuous density. Define $\nu(m) = F^{-1}(1-m)$ and

$$b(m) = \frac{r\nu(m)}{r + \beta m},$$

and assume $b(1) < B_0 < b(0)$. Then the unique aggregate path jumps at date zero to the unique $m_0^F \in (0,1)$ satisfying $b(m_0^F) = B_0$ and thereafter is continuous and solves

$$B_t = b(M_t), \qquad \dot{M}_t = -\frac{\beta M_t b(M_t)}{b'(M_t)}, \qquad M_0 = m_0^F, \tag{3}$$

until the support is exhausted.

Uniqueness is aggregate: asymmetric pure stopping assignments within a positive-mass payoff type may implement the same public path.

Proof. For part (a), the single-crossing property implies that no lower-payoff type can enter before a higher residual type. Fix a stage at which all payoff types above v_n have been exhausted and current adoption is q. While no one enters, q is constant and

$$\dot{g}_{v_n}(t) = e^{-rt}\{(r + \beta q)B_t - rv_n\}.$$

The term in braces falls strictly as B_t falls, so the payoff from service has at most one interior maximum on that waiting interval. The entry belief, if waiting is needed, is uniquely

$$B = \frac{rv_n}{r + \beta q}.$$

At date zero, proceed down the adoption-payoff ranking. Every atom whose service payoff is strictly decreasing even after both that atom and all higher atoms have entered adopts in full. There is then at most one marginal atom for which the date-zero derivative changes sign as its served fraction increases; strict monotonicity of $m \mapsto rv_n/(r + \beta m)$ pins down that fraction uniquely. At every positive interior date, however, a positive service jump is impossible: the jump in adoption raises the right derivative of g_{v_n} relative to the left derivative, whereas an interior maximizing date would require a downward kink. Thus each later atom begins continuously when its unique entry belief is reached and then mixes over the unique solution of (2). Exhaustion determines the next public state, and the argument iterates. Multiplicity survives only when a type is exactly indifferent at an entry boundary and weak tie selection permits both entry and delay.

For part (b), b is strictly decreasing because $\nu' < 0$ and $\nu > 0$. There can be no initial waiting: before adoption begins the belief is fixed at B_0, and every type with $v > B_0$ strictly prefers earlier service. There can be no positive interior adoption jump. If adoption jumped from a to $c > a$ at an interior date t, the one-sided derivatives of a type's certain-service payoff would satisfy

$$\dot{g}_v(t^+) - \dot{g}_v(t^-) = e^{-rt}\beta(c-a)B_t > 0.$$

A common service date for a positive interval of payoff types would require a downward kink, with the left derivative nonnegative and the right derivative nonpositive, which is impossible. There can also be no waiting interval after continuous service has begun: at its starting boundary the highest residual type lies on the ridge, and once the belief falls at a fixed stock that type strictly prefers service at the beginning of the pause. Equivalently, two continuous-service endpoints with the same stock would both require $B = b(m)$ even though the belief falls strictly during the pause.

Hence the path consists of one possible date-zero jump followed by continuous service. If the initial mass were below the solution of $b(m) = B_0$, the highest residual type would have a strictly declining service payoff at zero and would enter immediately; if it were above that solution, the lowest type in the jump would prefer a short delay. Thus the initial mass is the unique m_0^F. Every post-jump calendar interval contains service, so neighboring-date indifference applies on a dense set and continuity extends it to $B_t = b(M_t)$ throughout the active range. Since B is absolutely continuous and b^{-1} is continuously differentiable, $M = b^{-1}(B)$ is absolutely continuous as well; no singular-continuous adoption component remains. Combining the ridge identity with $\dot{B}_t = -\beta M_t B_t$ gives (3). Since b is continuously differentiable with $b' < 0$, the autonomous equation has a unique solution on every compact interior range. This proves aggregate uniqueness. The boundary cases are immediate: if $B_0 \geq b(0) = \nu(0)$, no adoption starts, while if $B_0 \leq b(1)$, the entire population adopts at date zero. □

For a continuous distribution of adoption payoffs, the marginal condition can also be written

$$v_t = B_t\left(1 + \frac{\beta M_t}{r}\right), \qquad M_t = 1 - F(v_t).$$

For a uniform distribution, substitution of $v_t = \bar{v} - (\bar{v} - \underline{v})M_t$ gives an autonomous scalar equation. The curvature of its solution changes sign at most once. The explicit threshold conditions in the main text follow by evaluating the derivative at the two support boundaries.

4 Capped availability

Capped availability sets $S_t = \bar{M}$ for every $t \geq 0$. Until inventory is exhausted, every applicant is accepted with probability one. The following theorem supplies the operational foundation for the capped path used in the main text.

Theorem 1 (Unique aggregate path under capped availability)**.** *Suppose F has a positive, continuous density, $\bar{M} \in (0,1)$, and*

$$\nu(\bar{M}) < B_0 < \nu(0).$$

Let

$$b(m) = \frac{r\nu(m)}{r + \beta m},$$

and let the unique pair $0 < m_0 < m_1 < \bar{M}$ solve

$$b(m_0) = B_0, \qquad b(m_1) = \nu(\bar{M}).$$

Capped availability has a unique aggregate no-signal equilibrium path. It jumps to m_0 at date zero, follows

$$B_t = b(M_t), \qquad \dot{M}_t = -\frac{\beta M_t b(M_t)}{b'(M_t)}$$

from m_0 to m_1, and completes with an unrationed terminal stockout of mass $\bar{M} - m_1$.

Proof. Every equilibrium reaches stockout in finite time. Some positive mass must adopt at date zero, as shown below. Thereafter, even if adoption were held fixed at that positive stock, the no-news belief would fall exponentially and reach the positive target payoff cutoff $\nu(\bar{M})$ in finite time, at which point enough residual types strictly prefer service. Let T denote the resulting stockout date. For every $t < T$, cumulative service is below $\bar{M}$, so an applicant is served with certainty. Upper-block willingness therefore orders successful service by adoption payoff before stockout.

There is no initial waiting. If the first service date were positive, beliefs would remain equal to B_0 beforehand, and every type receiving positive surplus at that date would strictly prefer applying at date zero. There is also no interior positive-mass service jump before stockout. If service jumped from a to $c > a$ at date $t \in (0, T)$, then for every payoff type v the derivative of certain-service payoff would jump upward by

$$e^{-rt}\beta(c - a)B_t > 0.$$

A positive interval of payoff types sharing an interior maximizing date would require the opposite sign pattern. Finally, there is no waiting interval between two continuous service segments. At a continuous-service endpoint with stock m, local indifference requires $B = b(m)$; while service is constant at m, the belief falls strictly, so the same ridge state cannot occur at both ends of a positive waiting interval.

Let the date-zero jump be m. If $m < m_0$, then $b(m) > B_0$, so the highest residual type has a strictly declining service payoff at zero and would enter immediately; the jump is too small. If $m > m_0$, then $b(m) < B_0$, so the lowest type included in the jump has a strictly increasing payoff immediately after zero and prefers a short delay; the jump is too large. Hence $m = m_0$. At that stock, $B_0 = b(m_0)$, and continuous service begins immediately. It cannot later stop while inventory remains: once the belief fell below the ridge at a fixed stock, the highest residual type would strictly prefer service at the start of the pause. Thus every preterminal interval after date zero contains service. Neighboring-date indifference and continuity give $B_t = b(M_t)$ throughout this range. Since B is absolutely continuous and b^{-1} is continuously differentiable, the adoption path is absolutely continuous and satisfies the displayed autonomous equation, which has a unique solution because $b' < 0$.

It remains to identify stockout. A positive rationed rush at an interior date cannot be an equilibrium. Any applicant with $v > B_T$ could apply an instant earlier, when inventory remains, obtain certain service, and approach the same belief with strictly higher acceptance probability. The zero-surplus cutoff has zero mass. Thus terminal demand equals remaining inventory and every terminal applicant is served.

The marginal target type $v_* = \nu(\bar{M})$ must weakly prefer terminal service, so $B_T \leq v_*$. Conversely, if $B_T < v_*$, a type just below v_* but above B_T is excluded in the proposed path yet can apply immediately before stockout and receive certain service with positive payoff. Hence $B_T \geq v_*$. Therefore

$$B_T = \nu(\bar{M}).$$

The unique preterminal stock is consequently m_1, defined by $b(m_1) = \nu(\bar{M})$, and the terminal stockout mass is $\bar{M} - m_1$.

Existence follows by prescribing the resulting path. Before T, every assigned applicant is served for sure. The smooth portion is the same certain-service ridge as under free availability, and it spans every calendar date in $[0, T]$. Hence every possible preterminal deviation date corresponds to a certain-service payoff line on the same affine upper envelope. At T, exactly the residual target mass applies and exhausts the residual inventory; after T no unit remains. Comparing each assigned line with that envelope verifies directly that every served payoff type prefers its prescribed date and every lower-payoff type prefers exclusion. Lemma 5 records the general batch–smooth–batch version of this argument. Principal-preferred tie selection assigns the zero-mass marginal cutoff to service. Thus the path is an operational equilibrium, and the preceding argument makes its aggregate path unique. □

Let

$$v_t^M = B_t, \qquad v_t^E = B_t\left(1 + \frac{\beta M_t}{r}\right).$$

At date zero, $\nu(m_0) = v_0^E$, so

$$m_0 = 1 - F(v_0^E).$$

At stockout, $v_T^E = \nu(m_1)$ and $v_T^M = B_T = \nu(\bar{M})$, hence

$$\bar{M} - m_1 = F(v_T^E) - F(v_T^M).$$

Thus the two boundary jumps have the threshold interpretation stated in the main text.

5 One-type unraveling

Release exactly $\bar{M} < 1$ at zero. Suppose a selected equilibrium reaches the target at $T > 0$. There must be a first successful adoption date $t_0 \leq T$. If $t_0 > 0$, beliefs are unchanged on $[0, t_0)$ and an agent scheduled at t_0 strictly prefers applying at zero. Thus $t_0 = 0$.

If capacity remains before T, any agent scheduled at a later date can apply slightly earlier and obtain certain service with less discounting and no worse belief. If capacity is exhausted before T, the target has already been reached because total supply equals $\bar{M}$. If instead the target is completed by a positive rationed mass at T, an applicant can move slightly earlier and face weakly more available stock before the terminal rush. Every delayed outcome is therefore contradicted. Immediate completion is optimal.

6 Two-type unrestricted optimum

Let the high and low adoption payoffs be $v_H > v_L$, let the high-type mass be q_H, and suppose

$$v_H > B_0 > v_L, \qquad q_H < \bar{M} < 1.$$

The low type is the marginal type needed to reach the target. Because $\bar{M} < 1$, only a submass $\bar{M} - q_H$ of low agents is needed. Call the high-type population, together with that selected low submass, the *target-relevant population.* When terminal participation binds at $B_T = v_L$, every low agent is indifferent between applying and remaining inactive. Atomlessness and principal-preferred equilibrium selection can therefore assign exactly the required low submass to the equilibrium that reaches the target and leave the remaining agents with the same low payoff inactive. Equivalently, the proof may normalize the target-relevant population to total mass $\bar{M}$. This normalization does not delete agents or change their incentives.

For a type $n \in \{H, L\}$, write its normalized payoff from certain service at date t as

$$u_n(t) := e^{-rt}(v_n - B_t).$$

Lemma 2 (Terminal myopia and target-relevant selection)**.** *Fix $K < \infty$ and a selected optimum among plans with at most K releases that has positive completion time T. Then*

$$B_T = v_L.$$

There is an optimal selected equilibrium in which the target-relevant low submass first applies at T and every residual target-relevant agent is served with certainty.

Proof. Terminal participation gives $B_T \leq v_L$. If the inequality were strict, the terminal release could be advanced while preserving positive terminal surplus, contradicting optimality; hence $B_T = v_L$. At every earlier history followed by positive exposure, $B_t > v_L$, so successful service gives the low type negative payoff. On a zero-exposure plateau, principal-preferred tie selection postpones low-type entry to the terminal date. At $B_T = v_L$, the low type is indifferent. Select exactly the residual target submass of low agents to apply at T, together with every residual high agent; the remaining agents with the same low payoff stay inactive. Terminal demand then equals remaining target capacity, so admission is certain for every target-relevant applicant. □

Fix K and take the selected K-release optimum. Let a be the first high-type service date and let T be the terminal low-type entry date from Lemma 2. Before T, only high types are served. Put

$$m = M_{a^-}, \qquad x = M_{T-} - m, \qquad E = \int_a^T M_s \, ds,$$

and define

$$z_0 = \frac{E}{T-a} - m, \qquad \widehat{T}(z) = a + \frac{E}{m+z}, \qquad z \in [z_0, x]. \tag{4}$$

Since $M_s \in [m, m+x]$ on $[a, T)$, $z_0 \in [0, x]$. The replacement serves mass z of the high type at a, has no further adoption until $\widehat{T}(z)$, and moves the remaining preterminal high-type service into the terminal batch. By construction,

$$(m+z)\{\widehat{T}(z) - a\} = E,$$

so the replacement reaches the same terminal belief as the original path.

Lemma 3 (Two-type no-shift inequality)**.** *Let $U_H = u_H(a)$ and let $V_H(z)$ be the high type's optimal continuation after skipping the initial replacement batch. Then*

$$V_H(z_0) \le U_H.$$

Proof. At z_0, equation (4) gives $\widehat{T}(z_0) = T$, and exposure matching leaves the terminal public state unchanged. Every residual target-relevant agent is admitted with certainty at T, so $V_H(z_0) = u_H(T)$. In the original plan, an unserved high agent who skips the opportunity at a can decline every intermediate batch and apply at the certain terminal batch. That agent's original optimal continuation is therefore at least $u_H(T)$. Since the high type is served at a, its current successful-service payoff is at least that optimal continuation. Combining the inequalities gives the result. □

Theorem 2 (Two-type unrestricted optimum)**.** *With two target-relevant payoff types, atomless within-type populations, repeated eligibility after rejection, anonymous rationing, asymmetric strategies, principal-preferred equilibrium selection, and target $\bar{M} < 1$, the principal has an optimal finite supply plan with at most two positive releases. Both releases can be chosen to be immediately exhausted.*

Proof. The map $z \mapsto V_H(z)$ is continuous. Lemma 3 gives $V_H(z_0) \le U_H$. If $V_H(x) \le U_H$, move all preterminal high-type service to a. Otherwise, the intermediate-value theorem gives a smallest $z^* \in [z_0, x]$ with $V_H(z^*) = U_H$. The high type then uses the initial batch; if equality binds, atomlessness and principal-preferred selection implement exactly the required high-type submass. There is no service until $\widehat{T}(z^*)$. Exposure matching gives the original terminal belief, the low type enters, and every residual target-relevant agent is admitted with certainty. Completion is weakly earlier because $\widehat{T}(z^*) \le T$.

This transformation applies to the selected optimum for every fixed release bound K. Hence no finite K-release plan can beat the best two-release plan. The two-release parameter set is compact after adjoining zero-capacity and coincident dummy releases, and its weak incentive constraints are closed, so its optimum is attained. Coincident or zero releases are merged or deleted. Finally, the first positive release occurs at date zero: before adoption starts, translating the entire continuation earlier changes no belief and multiplies all service payoffs by the same positive factor, preserving incentives while weakly improving completion. □

For the two-release representation, define

$$L = \log \frac{B_0}{v_L}, \qquad \Theta_H = \frac{1}{r} \log \frac{v_H - v_L}{v_H - B_0}. \tag{5}$$

If capacity c is released at date zero, stock remains c until the terminal date. Terminal myopia gives

$$T(c) = \frac{L}{\beta c}. \tag{6}$$

A high type who is not served at zero obtains certain terminal service at $T(c)$, so initial application requires

$$v_H - B_0 \ge e^{-rT(c)}(v_H - v_L), \tag{7}$$

which is equivalent to $T(c) \ge \Theta_H$.

7 Three-type construction

The three-type boundary is established by the following economy:

$$(v_1, v_2, v_3) = (0.86, 0.8, 1/3), \qquad (q_1, q_2, q_3) = (0.2, 0.4, 0.4),$$

$$B_0 = 0.75, \qquad \beta = 0.5, \qquad r = 0.4, \qquad \bar{M} = 0.98.$$

A feasible four-release plan serves type 1 at date zero, gives type 2 an early rationed opportunity at

$$s = 1.7489035062$$

with acceptance probability $Q = 0.1104189201$, clears rejected type-2 agents at

$$t = 2.2062097487,$$

and completes at

$$T = 4.1402437122.$$

The induced beliefs are

$$B_s = 0.75e^{-0.1s} = 0.6296618038,$$
$$B_t = B_s e^{-0.5(0.2+0.4Q)(t-s)} = 0.5954712920,$$
$$B_T = B_t e^{-0.3(T-t)} = 1/3.$$

Write $g_n(z) = e^{-0.4z}(v_n - B_z)$ for the payoff from certain service at date z. The relevant values are

$$g_1(0) = 0.11, \qquad g_1(s) = 0.1144327420,$$
$$g_1(t) = 0.1094497875, \qquad 0.95g_1(T) = 0.0955048525,$$
$$g_2(s) = g_2(t) = 0.95g_2(T) = 0.0846245529.$$

Moreover,

$$Qg_1(s) + (1-Q)g_1(t) = 0.11.$$

Thus type 1 accepts date-zero service and does not profit from the type-2 access rule after an unsuccessful application; type 2 is indifferent among the early application opportunity, clearing at t, and the terminal continuation; and type 3 has negative service payoff at s and t and zero payoff at T. Favorable tie selection therefore supports the displayed path.

7.1 Certified three-release benchmark

A plan with at most three immediately exhausted releases can be written with dates $0 \le s < T$, allowing zero-capacity or coincident dates. Type 3 is served only at T: if it applied earlier, upper-block willingness would bring every residual higher type, and expanding that release to the target would complete weakly sooner. At either preterminal date, applicants are therefore an upper block of types 1 and 2, with at most one partial cutoff. The sixteen ordered pairs of upper-block modes reduce analytically as follows.

If type 2 applies at date zero, its terminal option and $Q_T \ge 0.95$ imply

$$T \ge \frac{1}{0.4} \log \frac{0.95(v_2 - v_3)}{v_2 - B_0} = 5.4557473178. \tag{8}$$

If no type-2 mass is served before T, exposure is bounded by $0.2T$, so

$$T \ge \frac{\log(0.75/(1/3))}{0.5(0.2)} = 8.1093021622. \tag{9}$$

Hence every competitive three-release plan serves only type 1 at zero and type 2 at the middle date. If only a fraction of type 2 applies at s, increase the applicant share to one and lower the common acceptance probability to keep middle service fixed. Beliefs, terminal service, and type-2 utility are unchanged, while type 1 receives a lower probability of the attractive middle service. Full type-2 application is therefore without loss.

Let $x \in [0, 0.2]$ be type-1 service at zero and $p \in [0, 1]$ the middle acceptance probability. Then

$$m = x + p(0.6 - x), \qquad B_s = 0.75e^{-0.5xs}, \tag{10}$$

$$T(x, s, p) = s + \frac{\log(B_s/v_3)}{0.5m}, \qquad Q_T(x, s, p) = \frac{0.98 - m}{1 - m}. \tag{11}$$

Define

$$g_n = e^{-0.4s}(v_n - B_s), \qquad h_n = Q_T e^{-0.4T}(v_n - v_3), \qquad C_n = pg_n + (1 - p)h_n.$$

The remaining optimization problem is exactly

$$\min_{0 \le x \le 0.2,\ s \ge 0,\ 0 \le p \le 1} T(x, s, p) \tag{12}$$

subject to

$$g_2 \ge h_2, \qquad C_1 \le 0.11, \qquad C_2 \ge 0.05. \tag{13}$$

These inequalities state, respectively, that type 2 applies at s, type 1 accepts service at zero, and type 2 does not imitate date zero. Type 3's earlier constraints are automatic.

For fixed (x, s), completion is strictly decreasing in p. On the sublevel set $T \le 4.1701$, terminal continuation rises with p, C_1 rises with p, and the type-2 participation slack falls; the last constraint in (13) is slack. Thus the largest feasible p is determined by $g_2 = h_2$, $C_1 = 0.11$, or $p = 1$. Clearing positive denominators in the implicit derivatives gives the exhaustive branch comparison:

Binding case	Branch behavior	Lower bound
$g_2 = h_2$, $C_1 < 0.11$	Improves as x rises until type-1 IC binds	> 4.1700003382
$g_2 = h_2$, $C_1 = 0.11$	Lower branch improves to $x = 0.2$; upper branch is slower	4.1700003382
$p = 1$	Type-2 participation binds at the optimum	4.2412342351
Other boundaries	Covered by (8), (9), or a row above	> 4.1700003382

The derivative sign certificates and isolating intervals are reported in the replication supplement. They imply $x = 0.2$ on the minimizing branch. Put

$$A = \frac{v_1 - v_3}{v_2 - v_3} = \frac{79}{70}, \qquad B_s = 0.75e^{-0.1s}.$$

When type-2 participation and type-1 IC bind,

$$p(s) = \frac{(v_1 - B_0)e^{0.4s} - A(v_2 - B_s)}{(v_1 - B_s) - A(v_2 - B_s)}, \tag{14}$$

and exposure matching with type-2 indifference gives

$$\frac{\log(B_s/v_3)}{0.5[0.2 + 0.4p(s)]} = \frac{1}{0.4} \log \frac{Q_T(p(s))(v_2 - v_3)}{v_2 - B_s}. \tag{15}$$

The two isolated roots are

$$(s, p, T) = (1.0624057254, 0.6044630754, 4.2525967757)$$

and

$$(s, p, T) = (1.8812680292, 0.8605859690, 4.1700003382).$$

The latter is the minimizing root; the $p = 1$ boundary gives $T = 4.2412342351$.

Theorem 3 (Three-type release-count boundary)**.** *For the displayed primitives, every plan with at most three immediately exhausted releases has completion time at least* 4.1700003382*. The feasible four-release plan completes at* 4.1402437122*. Hence the two-type release-count theorem does not extend to three target-relevant payoff types.*

Proof. Upper-block willingness and terminal service of type 3 reduce the three-release class to the finite mode list. Bounds (8) and (9) eliminate modes with type 2 at zero or no preterminal type-2 service. The full-application transformation gives program (12)–(13). The branch comparison places the minimum at $x = 0.2$ and reduces it to the roots of (15) and the $p = 1$ boundary. The displayed certified values prove the lower bound, and the strict comparison with the verified four-release path proves the claim. □

Fixed release bounds. For any finite K, allowing weakly ordered dates and zero-capacity dummies makes the finite-dimensional equilibrium constraints closed on a bounded horizon, so a K-release optimum exists. Trimming each retained release to realized on-path service yields immediate exhaustion without changing the path. This observation does not consolidate distinct exhausted releases and does not create a general release-count bound.

8 Finite foundations and finite approximation

The continuum foundation has two parts: compactness and path convergence for selected ordered finite sequences, and conditional recovery of strongly regular ordered direct schedules from fine finite partitions.

8.1 Convergence and local-expansion conventions

The arguments below use three conventions that preserve economic objects that ordinary pointwise reasoning can lose.

Ordinary pointwise convergence can lose a pooled jump. We therefore represent a nondecreasing, right-continuous schedule h by its graph together with the vertical segment joining $h(x^-)$ and $h(x)$ at each jump. We call this the *completed graph* and use Hausdorff distance between these compact sets. For bounded monotone schedules, Helly selection gives a subsequence converging at every continuity point; when the vertical jump segments are retained, this is equivalent to convergence of completed graphs. The convention is used only to keep pooled service and terminal batches visible in the limit.[2]

The finite-approximation argument solves finitely many smooth contact equations. The proof displays the relevant Jacobian. When that Jacobian is bounded away from singularity on the compact set under consideration, the ordinary implicit-function theorem gives locally unique nearby contacts; a finite cover makes the bound uniform.[3]

The reflected-contact expansion also needs an exact way to pair the two sides of a strict local maximum. We construct that change of variable directly in Lemma 7: the loss from the maximum is written as $z^2 A(p,z)$ with $A>0$, and $z\sqrt{A(p,z)}$ turns equal-height reflection into ordinary sign reversal.[4]

8.2 Completed-graph compactness and path convergence

Let finite distributions F_N approximate F through ordered quantile cells $I_{j,N}$ of masses $q_{j,N}$, with mesh $\Delta_N = \max_j q_{j,N} \to 0$. Outside an exceptional set of mass $o(1)$, represent the selected ordered service assignment by a nondecreasing, right-continuous function τ_N on the served rank interval $[0,\bar{M}]$ and its completed graph. Suppose completion times are bounded by $\bar{T}$.

Lemma 4 (Completed-graph compactness)**.** *A subsequence of completed graphs converges in the Hausdorff metric to the completed graph of a nondecreasing correspondence τ. Right-continuous representatives converge at every continuity point and in L^1. Their generalized inverses M_N converge almost everywhere to M, exposure processes converge uniformly, and therefore $B_N \to B$ uniformly on $[0,\bar{T}]$.*

Proof. Helly selection gives pointwise convergence of monotone representatives at continuity points; under the completed-graph convention just stated, monotonicity makes this equivalent to Hausdorff convergence of the completed graphs. Boundedness gives L^1 convergence. Generalized inverses converge at every continuity point of M and hence almost everywhere. Dominated convergence yields

$$\int_0^{\bar{T}} |M_N(s) - M(s)|\, ds \to 0.$$

Consequently

$$\sup_{t\le \bar{T}} \left| \int_0^t (M_N - M)\, ds \right| \le \int_0^{\bar{T}} |M_N - M|\, ds \to 0.$$

The exponential belief law then gives uniform belief convergence. □

Completion-time convergence is not implied by graph convergence alone. We use the following endpoint condition.

[2] Related Helly-selection arguments for monotone strategies and allocations appear in Lauermann and Speit (2023) and Dilmé and Garrett (2026). For the completed-graph convention for paths with jumps, see Whitt (2002).

[3] For a standard parameterized statement of the implicit-function theorem, see Krantz and Parks (2002). No infinite-dimensional result is used here.

[4] This explicit construction is the one-dimensional quadratic normal form often associated with the Morse lemma; the proof below does not invoke the general theorem. For background, see Matsumoto (2002).

Definition 1 (Regular target crossing)**.** A limiting adoption path with completion date T crosses the target regularly if $M(t) < \bar{M}$ for every $t < T$ and either (i) M is continuous at T and crosses the target from below at a positive one-sided rate: there exist $c, \delta > 0$ such that $M(T) - M(t) \geq c(T - t)$ for every $t \in [T - \delta, T]$; or (ii) M has a jump at T whose size is bounded away from zero along the approximating sequence. Target-aligned terminal cells are allowed when their service dates converge to T.

Under this condition, completed-graph convergence implies convergence of hitting times. Without it, a vanishing terminal cell can determine every finite completion date while disappearing from the limiting graph.

8.3 Finite approximation of direct schedules

For the recovery direction, assume finitely many smooth components and pools, transverse endpoint contacts, adjacent-cutoff Jacobians uniformly bounded away from zero, and uniform slack in every nonbinding global direct-IC and lower-tail inequality.

Proposition 3 (Finite approximation of strongly regular schedules)**.** *Under these assumptions, target-aligned adjacent quantile partitions admit whole-cell finite direct menus with exact adjacent cutoff equalities, exhausted batches, global direct IC, and completion times converging to the continuum schedule.*

Proof. Assign each separating cell the date of its marginal continuum rank and assign all cells in a pooled block to the common pool date. The residual adjacent indifference equations and pool-contact equations form a finite system whose Jacobian converges to the block-diagonal continuum Jacobian. Uniform nonsingularity and the implicit-function theorem give date corrections tending to zero. Compactness turns the local bounds into a uniform bound over the finitely many component types. Affine single-crossing turns exact adjacent equalities into global ordering; uniform slack preserves every nonbinding comparison and lower-tail inequality. Aggregate graphs, exposure, beliefs, and completion times converge. □

The approximation is conditional and concerns ordered direct menus. It does not prove that every such menu is sequentially implementable in the original repeated-application game.

9 Ordered-schedule notation used in the local analysis

The working-paper Appendix proves the ordered-rollout decomposition, staircase reconstruction, exact interval-selection reduction, countable attainment, and finite compression. This section records only the notation and identities needed for the distinct density and reflected-contact calculations below.

Let $m \in [0, \bar{M}]$ index served agents from highest to lowest payoff, with marginal payoff $\nu(m) = F^{-1}(1 - m)$. Define

$$b(m) = \frac{r\nu(m)}{r + \beta m}, \qquad x(m) = -\log b(m), \qquad \mathcal{S}(u, v) = \frac{1}{\beta} \int_u^v \frac{x'(m)}{m}\, dm.$$

The active boundary ranks $0 < m_0 < m_1 < \bar{M}$ are the unique solutions of

$$b(m_0) = B_0, \qquad b(m_1) = \nu(\bar{M}). \tag{16}$$

The working-paper Appendix shows that continuous ordered service follows the ridge belief $b(m)$, while every waiting interval is represented by an exact excursion $e = (d, p, h)$ with $d < p < h$. Its pause at adoption stock p lasts

$$\Delta(e) = \frac{1}{\beta p} \log \frac{b(d)}{b(h)},$$

and its departure and reentry lines meet at the pivot payoff $\nu(p)$ exactly when

$$b(d)^{r/(\beta p)}\{\nu(p) - b(d)\} = b(h)^{r/(\beta p)}\{\nu(p) - b(h)\}.$$

Relative to smooth ridge passage over $[d, h]$, the excursion changes completion time by

$$\Gamma_F(d, p, h) = \frac{1}{\beta} \int_d^h x'(m) \left(\frac{1}{p} - \frac{1}{m} \right) dm. \tag{17}$$

All feasibility, decomposition, and reconstruction claims associated with these objects are proved in the working-paper Appendix; no part of those proofs is repeated here.

10 Diffuse and concentrated density regions

10.1 Canonical feasibility and exact wave gains

The working-paper Appendix contains the self-contained convex-quantile proof that a positive, nonincreasing active density makes every nondegenerate excursion costly and uniquely selects capped availability. To avoid duplicating that proof, this supplement retains only the distinct ingredients used elsewhere: feasibility of the canonical batch–smooth–batch schedule, the exact primitive gain criterion, and the payoff-ordered-versus-operational scope comparison.

Lemma 5 (Feasibility of the batch–smooth–batch schedule)**.** *Let m_0, m_1 solve* (16) *and define*

$$\tau^*(m) = \begin{cases} 0, & 0 \le m \le m_0, \\ \displaystyle\int_{m_0}^{m} \frac{x'(z)}{\beta z}\, dz, & m_0 < m < m_1, \\ \mathcal{S}(m_0, m_1), & m_1 \le m \le \bar{M}. \end{cases} \tag{18}$$

Then τ^* induces the ridge belief $B_{\tau^*(m)} = b(m)$ on $[m_0, m_1]$ and is globally incentive compatible. Served types participate; types below $\nu(\bar{M})$ prefer exclusion; and the completion time is $\mathcal{S}(m_0, m_1)$.

Proof. The initial batch leaves belief unchanged at $B_0 = b(m_0)$. Along the increasing part,

$$\beta \int_{m_0}^{m} z\, d\tau^*(z) = \int_{m_0}^{m} x'(z)\, dz = x(m) - x(m_0),$$

so the induced belief is $b(m)$. At completion it is $b(m_1) = \nu(\bar{M})$; the terminal batch does not change belief instantaneously.

For $m \in [m_0, m_1]$, let

$$L_m(v) = e^{-r\tau^*(m)}\{v - b(m)\},$$

where $m = m_0$ denotes the date-zero line and $m = m_1$ the terminal line. Since

$$\tau^{*\prime}(m) = -\frac{b'(m)}{\beta m b(m)},$$

direct differentiation gives

$$\frac{\partial L_m(v)}{\partial m} = e^{-r\tau^*(m)} b'(m)\left[\frac{r\{v - b(m)\}}{\beta m b(m)} - 1\right] = e^{-r\tau^*(m)} b'(m) \frac{r\{v - \nu(m)\}}{\beta m b(m)}. \tag{19}$$

Because $b' < 0$ and ν is strictly decreasing, a type $v = \nu(q)$ with $q \in [m_0, m_1]$ sees $L_m(v)$ increase up to $m = q$ and decrease afterward. If $q < m_0$, it decreases throughout $[m_0, m_1]$, so the date-zero line is optimal. If $q > m_1$, it increases throughout, so the terminal line is optimal. This proves global IC.

Participation holds because $\nu(m) > b(m)$ for $m > 0$, while every rank in the terminal batch has adoption payoff at least $\nu(\bar{M}) = b(m_1)$. An excluded type has $v < \nu(\bar{M})$; the terminal line is its best offered line by (19); and that line gives strictly negative payoff. Thus exclusion is optimal. The completion time follows from (18). □

Proposition 4 (Exact criterion for a profitable mid-rollout wave)**.** *Suppose F has a positive, continuously differentiable density on the relevant interval. Let τ^{cap} be the canonical schedule in* (18)*, with completion time $T^{\text{cap}} = \mathcal{S}(m_0, m_1)$. Fix*

$$m_0 < d < p < h < m_1, \qquad a = \nu(p), \qquad k = \frac{r}{\beta p},$$

and set

$$\Delta = \frac{1}{\beta p} \log \frac{b(d)}{b(h)}.$$

There is a globally incentive-compatible, one-wave schedule that follows τ^{cap} to d, serves $[d, p]$ at departure, waits Δ at stock p, serves $[p, h]$ at reentry, and then resumes the capped continuation if and only if

$$b(d)^k\{a - b(d)\} = b(h)^k\{a - b(h)\}. \tag{20}$$

When this condition holds, the completion time is

$$T^{\text{batch}} = T^{\text{cap}} + \Gamma_F(d, p, h),$$

where

$$\Gamma_F(d, p, h) = \frac{1}{\beta}\int_d^h \left[\frac{\beta}{r + \beta m} + \frac{1}{\nu(m) f(\nu(m))}\right]\left(\frac{1}{p} - \frac{1}{m}\right) dm. \tag{21}$$

Hence the episode beats capped availability if and only if (20) *holds and* (21) *is negative.*

Proof. The belief law while adoption is fixed at p forces the waiting time Δ. Adjacent IC between the departure and reentry pools requires their two affine payoff lines to cross at the cutoff payoff $a = \nu(p)$. Substitution of the forced clock gives exactly (20); hence the condition is necessary.

Assume it holds. Put

$$\Gamma = \Delta - \mathcal{S}(d, h).$$

Define the schedule by keeping the capped dates through d, assigning the departure line to ranks $[d, p]$, the reentry line to ranks $[p, h]$, and shifting every capped date after h by the common amount Γ. The departure belief is $b(d)$. The forced wait at stock p brings the belief to $b(h)$, so the shifted continuation begins at exactly the same public state as the capped continuation from h. It therefore reproduces all subsequent ridge beliefs and reaches the same terminal belief. The completion-time identity follows immediately, and (21) is the definition of the excursion gain after substituting

$$x'(m) = \frac{\beta}{r + \beta m} + \frac{1}{\nu(m)f(\nu(m))}.$$

It remains to verify global IC. Along the capped menu, the payoff to type $v = \nu(q)$ from the line indexed by m rises until $m = q$ and falls afterward, by (19). Thus, for every $q \le d$, the assigned capped line dominates the departure line, and the departure line dominates the reentry line because $\nu(q) > \nu(p)$ and the two lines cross at $\nu(p)$. Among all shifted continuation lines, the reentry line is best for such a high type. Hence no prefix type deviates.

For $q \in [d, p]$, the departure line is best among all prefix lines and dominates the reentry line by the same crossing argument. For $q \in [p, h]$, the reentry line dominates the departure line and is best among all continuation lines. For $q \ge h$, the common time shift preserves all IC comparisons inside the capped continuation; the reentry line dominates every earlier line for these lower-payoff types. Participation is preserved, and lower-tail exclusion is unchanged because the terminal belief remains $\nu(\bar{M})$, so shifting the terminal line changes only its positive discount factor, not the sign of any payoff. The schedule is therefore globally incentive compatible. □

Corollary 1 (A lone delayed clearing batch cannot beat capped availability)**.** *Fix $m_0 < d < h < m_1$. Consider a ridge-returning plan that follows capped availability to d, withholds service while adoption remains fixed at d, serves the entire interval $[d, h]$ in one later batch, and then resumes the capped continuation. Whenever this plan is incentive compatible, its excess duration is*

$$\frac{1}{\beta} \int_d^h x'(m) \left(\frac{1}{d} - \frac{1}{m} \right) dm > 0.$$

Hence a profitable mid-rollout departure from capped availability must first raise the learning stock before withholding and later clear a second block.

Proof. The waiting stock is d, so the belief law forces duration

$$\frac{1}{\beta d} \int_d^h x'(m) dm.$$

Subtracting smooth capped passage, $\int_d^h x'(m)/(\beta m)dm$, gives the displayed integral. Since $x'(m) > 0$ and $m > d$ throughout the interior, the difference is strictly positive. □

Remark 2 (Scope of payoff-ordered uniqueness). Let $V^D(F)$ be the infimum of completion times over the payoff-ordered direct schedules studied in Sections 9–10, and let $V^R(F)$ be the relaxed operational optimum allowing a disintegration K_v over service dates for almost every payoff type. The nonincreasing-density corollary proved in the working-paper Appendix identifies $V^D(F) = \mathcal{S}(m_0, m_1)$ and its unique aggregate optimizer when $f' \leq 0$. Theorem 1 shows that capped availability is an operational implementation of that schedule, so

$$V^R(F) \leq V^D(F).$$

Payoff-ordered schedules keep attention on the timing of anonymous supply, which is the mechanism studied in the paper. Rules that condition future access on an agent's own application and rationing history add a separate dynamic-rationing instrument and may permit strictly earlier completion. Determining when that instrument has value is a natural direction for future work. Compactness of finite operational outcomes, vanishing atomic masses, and terminal padding do not settle the comparison because they do not imply $K_v = \delta_{\tau(v)}$ almost everywhere.

10.2 Local cubic coefficient and a uniform local pooling condition

Fix an interior pivot p and put $a = \nu(p)$ and $k = r/(\beta p)$. For a reflected local excursion, write

$$m_d = p - h, \qquad m_r = p + \ell(h).$$

The exact contact equation is

$$q_p(p-h) = q_p(p + \ell(h)), \qquad q_p(m) = \log\left[b(m)^k\{a - b(m)\}\right].$$

Lemma 6 (Local reflected contact). *If ν is C^4, positive, and strictly decreasing near p, then q_p has a strict local maximum at p. For all sufficiently small $h > 0$, there is a unique $\ell(h) > 0$ satisfying the contact equation, and*

$$\ell(h) = h - \frac{q_p'''(p)}{3q_p''(p)}h^2 + O(h^3). \tag{22}$$

Proof. The function $z \mapsto z^k(a - z)$ has a strict maximum at

$$z = \frac{k}{k+1}a = b(p).$$

Since $b'(p) \neq 0$, composition with $b(m)$ gives $q_p'(p) = 0$ and $q_p''(p) < 0$. The two monotone branches around this strict maximum give a unique reflected point. Taylor expansion gives

$$q_p(p-h) = q_p(p) + \frac{1}{2}q_p''(p)h^2 - \frac{1}{6}q_p'''(p)h^3 + O(h^4),$$

$$q_p(p+\ell) = q_p(p) + \frac{1}{2}q_p''(p)\ell^2 + \frac{1}{6}q_p'''(p)\ell^3 + O(\ell^4).$$

Substitute $\ell = h + Ah^2 + O(h^3)$ and match cubic terms to obtain $A = -q_p'''(p)/(3q_p''(p))$. □

Theorem 4 (Cubic excursion expansion)**.** *Let*

$$y(m) = -\frac{\nu'(m)}{\nu(m)}.$$

For the reflected excursion of Lemma 6,

$$\Gamma_F(p-h, p, p+\ell(h)) = K_F(p; r, \beta)h^3 + O(h^4), \tag{23}$$

where

$$K_F(p; r, \beta) = \frac{(\beta p + 2r)y(p)^2 - \beta p y'(p) + \dfrac{2\beta r}{\beta p + r}y(p)}{3\beta^2 p^3}. \tag{24}$$

Writing $v = \nu(p)$, $f = f(v)$, and $f' = f'(v)$,

$$K_F(p; r, \beta) = \frac{2rf - \beta p v f' + \dfrac{2\beta r}{\beta p + r}vf^2}{3\beta^2 p^3 v^2 f^3}. \tag{25}$$

In particular,

$$K_F(p) < 0 \iff \frac{f'(v)}{f(v)} > \frac{2r}{\beta p v} + \frac{2rf(v)}{p(\beta p + r)}.$$

Proof. Write $b_j = b^{(j)}(p)$ and $b_0 = b(p) = ra/(r + \beta p)$. Expand the exact integral (17) over $[p-h, p+\ell(h)]$, substitute (22), and collect cubic terms. The constant, linear, and quadratic differences cancel because of same-state contact and reflection. Direct simplification gives

$$K_F = \frac{(\beta p + r)\{2a\beta b_1 + a\beta p b_2 + 2(\beta p + r)b_1^2\}}{3a^2\beta^2 p^3 r}. \tag{26}$$

Now substitute

$$b(m) = \frac{r\nu(m)}{r + \beta m}, \qquad \frac{\nu'}{\nu} = -y, \qquad \frac{\nu''}{\nu} = y^2 - y'.$$

Equation (26) reduces to (24). Since

$$\nu'(m) = -\frac{1}{f(\nu(m))}, \qquad y(m) = \frac{1}{\nu(m)f(\nu(m))},$$

we have

$$y'(m) = \frac{f(v) + vf'(v)}{v^2 f(v)^3}.$$

Substitution into (24) gives (25). □

Lemma 7 (Uniform fourth-order remainder)**.** *Let C be a compact interior rank interval. Suppose ν is C^4 on a neighborhood of C, the reflected contact remains on the simple local branch, and all* *rescaled controls remain in a common compact interior feasible set. Then there are $M < \infty$ and $\bar{h} > 0$ such that*

$$\left|\Gamma_F(p-h, p, p+\ell_p(h)) - K_F(p)h^3\right| \leq Mh^4$$

for every $p \in C$ and $0 < h \leq \bar{h}$.

Proof. The ordinary inverse branches of q_p are singular at their common maximum. To remove that singularity, write the loss from the maximum in a coordinate in which equal heights are exact mirror images. For z near zero, write

$$q_p(p) - q_p(p+z) = z^2 A(p, z),$$

where A is positive and C^3 on a neighborhood of the compact core. Uniform strictness of the maximum implies that A is bounded away from zero there. Define

$$\chi_p(z) = z\sqrt{A(p, z)}.$$

Then $q_p(p) - q_p(p+z) = \chi_p(z)^2$. Because A is positive and smooth, $(p, z) \mapsto \chi_p(z)$ is locally one-to-one, with its derivative bounded away from zero uniformly on the compact core. The reflected endpoint is therefore

$$\ell_p(h) = \chi_p^{-1}\big(-\chi_p(-h)\big).$$

The map $(p, h) \mapsto \ell_p(h)$ is C^3 with derivatives bounded uniformly on the compact core. Substitution into the exact gain integral makes the normalized remainder a continuous function of (p, h) after division by h^4. Compactness then yields a common constant C such that

$$\left|\Gamma_F(p, h) - K_F(p)h^3\right| \leq Ch^4$$

for all sufficiently small h, uniformly in p on the core. □

Proposition 5 (Uniform local sign test)**.** *Let $h_0 := \min\{\bar{h}, \eta/(2M)\}$.*

(i) If $K_F(p) \leq -\eta < 0$ throughout C, then every admissible reflected excursion based in C with $0 < h \leq h_0$ satisfies

$$\Gamma_F(p-h, p, p+\ell_p(h)) \leq -\frac{\eta}{2}h^3 < 0$$

and therefore beats capped availability.

(ii) If $K_F(p) \geq \eta > 0$ throughout C, then every such excursion satisfies

$$\Gamma_F(p-h, p, p+\ell_p(h)) \geq \frac{\eta}{2}h^3 > 0$$

and is slower than capped availability.

Proof. Combine Theorem 4 and Lemma 7. □

Selection and closure results. The working-paper Appendix proves finite telescoping and reconstruction, the Bellman representation, countable attainment, and finite-staircase approximation. Those arguments use the uniform local remainder established above but are not repeated in this supplement.

Scope of the finite foundation. The finite foundation supports the payoff-ordered analysis. It does not cover arbitrary individual-history-dependent access rules; Remark 2 treats that separate dynamic-rationing margin as a direction for future work.

References

Dilmé, F. and D. F. Garrett (2026). A dynamic theory of random price discounts. *Review of Economic Studies* 93, 1671–1709.

Krantz, S. G. and H. R. Parks (2002). *The Implicit Function Theorem: History, Theory, and Applications*. Birkhäuser.

Lauermann, S. and A. Speit (2023). Bidding in common-value auctions with an unknown number of competitors. *Econometrica* 91, 493–527.

Matsumoto, Y. (2002). *An Introduction to Morse Theory*. American Mathematical Society.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research* 6, 58–73.

Whitt, W. (2002). *Stochastic-Process Limits: An Introduction to Stochastic-Process Limits and Their Application to Queues*. Springer.