Staggered Rollout for Innovation Adoption[*]

Ricardo Fonseca[†]

June 2026

Abstract

When the quality of a good is uncertain, adoption generates public information, so forward-looking agents may wait for others to experiment. A principal who seeks to reach a fixed adoption target as quickly as possible can counter this delay by committing to a supply schedule. I first show that free availability induces a unique aggregate adoption path and can generate S-shaped diffusion when agents differ in their private adoption payoffs. Scarcity then changes the structure sharply: with one payoff type, it eliminates gradual adoption; with two, one initial and one terminal exhausted batch attain the unrestricted optimum; with three, rejection-contingent rationing can make four releases outperform every immediately exhausted plan with at most three releases. Atomlessness restores a tractable operational structure: with continuously distributed adoption payoffs, ordered rollouts with finitely many waves are exactly implemented by anonymous staircase supply. I characterize the globally optimal rollout in this continuum model by reducing the problem to weighted interval selection over pooled excursions around a unique smooth adoption path. Consequently, releasing the entire target inventory at date zero is uniquely optimal under a positive nonincreasing density, whereas a rapidly rising density creates profitable staggered releases. Such releases make agents accelerated into the initial pooled group strictly worse off but benefit every adopter served after access resumes. Thus, the distribution of adoption payoffs determines both the pace of rollout and who bears the welfare cost of accelerating it.

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

1 Introduction

In December 2020, the first COVID-19 vaccines were authorized. Access then expanded in waves: health workers and the elderly came first, followed by broader groups and eventually the general population. Medical risk and equity largely determined the order, but the rollout also changed incentives. Each vaccinated cohort generated public evidence about safety and effectiveness, so someone expecting access later could wait for that evidence rather than experiment early.

---

[*] 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, and Universidad de los Andes for helpful comments. All remaining errors are mine.

[†] Pontificia Universidad Javeriana, Carrera 7 No. 40-62, Bogotá, Colombia. E-mail: `baricardo@javeriana.edu.co`.

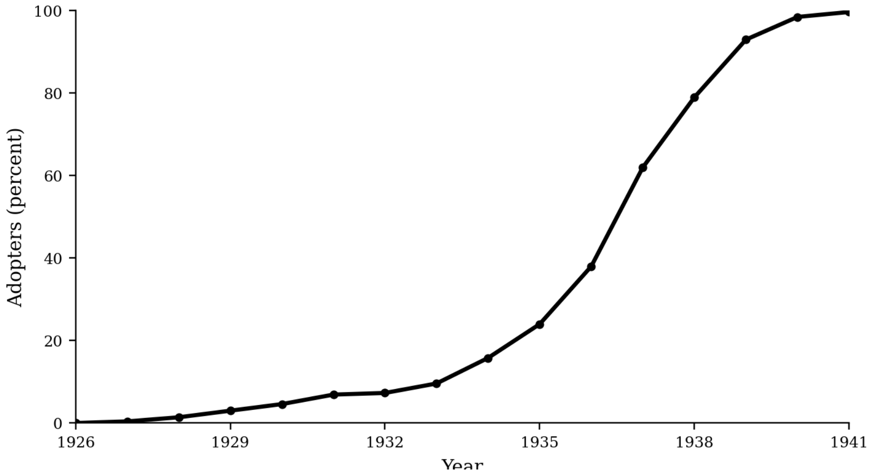

Figure 1: Hybrid-corn diffusion in two Iowa communities. Author's reconstruction from the annual adoption shares reported in Ryan and Gross (1943, Figure 4).

The same tension arises whenever adoption reveals quality and can be postponed. Farmers may wait for neighbors to try a new seed, physicians for broader experience with an unfamiliar treatment, and firms for early users to reveal the profitability of a new technology. Platforms and artificial-intelligence providers likewise use invitations, waitlists, and capacity tiers when releasing new products. In each case, availability determines not only who can adopt, but also how valuable it is to wait.

The empirical diffusion literature has long emphasized social learning in agriculture, health, and technology adoption (Ryan and Gross, 1943; Besley and Case, 1993; Foster and Rosenzweig, 1995; Munshi, 2004; Bandiera and Rasul, 2006; Conley and Udry, 2010; Dupas, 2014). Its best-known regularity is the S-shaped diffusion curve: take-up initially accelerates and later decelerates (Rogers, 2003; Young, 2009). Figure 1 reproduces the classic hybrid-corn pattern in Ryan and Gross (1943). Here such a path arises with fully forward-looking agents who strategically wait for public information rather than from myopia or an imposed contagion process.

This paper studies how a principal should release supply when agents differ in the payoff from successful adoption and learn from a public bad-news process whose intensity rises with cumulative adoption. The principal cannot force adoption or use transfers. She commits instead to a dated cumulative-supply path and minimizes the time required to reach a target adoption mass. The completion-time objective is deliberately narrow, but the model also identifies exactly who gains and loses when a rollout is accelerated.

The design problem combines learning, screening, and rationing. Fine staggering separates nearby values because high-value agents accept early access while lower-value agents wait. Pooling sacrifices some of that sorting but creates a discrete increase in the stock generating information. With positive-mass types, rationing adds a third instrument: an early chance of service and the continuation after rejection can jointly screen agents even when all rejected applicants remain eligible later. Scarcity is valuable not because unused units have intrinsic value, but because future access determines the option to free-ride on others' experimentation.

I first isolate the decentralized benchmark. Under free availability, apart from knife-edge entry indifferences, the aggregate no-signal equilibrium path is unique. Adoption is convex while one payoff class mixes along its indifference ridge; with several values, waiting intervals connect these class-specific segments. With a continuous distribution, a scalar cutoff equation governs the path, and the uniform case has a unique inflection point. Strategic waiting and heterogeneous payoffs therefore generate an endogenous S-shape.

Scarcity changes this benchmark and reveals where simple rollout ends and richer dynamic screening begins. With one payoff type, releasing the target mass at date zero eliminates gradual delay and completes adoption immediately. With two target-relevant types, one initial and one terminal exhausted batch attain the unrestricted optimum, even though the principal may use arbitrary finite calendars and agents may reapply after rejection. This simplicity fails with three types. A small early rationed opportunity for the intermediate type can preserve the top type's incentive to adopt immediately, while a later clearing opportunity for rejected agents raises the learning stock before the terminal date. A four-release plan then strictly outperforms every immediately exhausted plan with at most three releases. Rejection-contingent service is a genuine screening instrument, not a technical artifact.

The three-type result might suggest that rich rejection-contingent lotteries are unavoidable once the value distribution becomes more complex. Atomlessness changes the natural policy representation instead. The continuum analysis uses the same primitive instrument but no type-by-type lottery technology. With finitely many values, each type has positive mass and can be rationed by a positive release. Under a continuous distribution, an exact value has zero mass: a release directed only to that value is neither a supply increment nor an anonymous rationing opportunity. The planner therefore chooses an ordered schedule in which almost every value receives one service date, higher values are served weakly earlier, and positive-measure intervals may pool. This is not an abstract assignment technology introduced merely for tractability. Every feasible finite-wave schedule is exactly decentralized by anonymous staircase supply. An inventory cap remains open through smooth adoption, is exhausted by a departure pool, and is raised when the rollout reenters. The direct schedule is simply a representation of self-selection under aggregate capacity.

Once this operational bridge is established, the continuum problem becomes globally tractable. Every finite-completion rollout consists of optional boundary pools, smooth passage along a unique separating ridge, and a compatible finite or countable family of pooled excursions that leave the ridge and later return to it. Conversely, every compatible family reconstructs a globally incentive-compatible rollout. Completion time equals the excursion-free benchmark plus the sum of exact excursion gains, reducing the principal's problem to weighted interval selection. A countable optimum exists, and finite staircase policies approximate it arbitrarily closely. The normal form includes the boundary cases: adoption may begin smoothly with no initial pool, collapse into immediate completion, or fail to start when even the highest value is too low.

That variational representation, in turn, yields sharp primitive conclusions. Capped availability is optimal exactly when every feasible excursion has nonnegative excess duration, and it is uniquely optimal under a positive nonincreasing density. Rapidly rising density creates profitable waves. Conditional on one wave, a compact program determines the departure jump, the pause, and the reentry jump. The same geometry gives a precise welfare incidence. Relative to free take-up from the same public state, the departure pool is strictly worse off. If the wave speeds completion, every adopter after reentry is strictly better off, with a unique cutoff inside the reentry pool. Optimal staggering therefore transfers the burden of experimentation toward earlier, higher-value adopters and passes the resulting learning gains down the value ranking.

1.1 Contribution relative to the literature

Viewed from the benchmark side, the paper connects strategic experimentation and innovation diffusion to dynamic allocation. Chamley and Gale (1994), Keller, Rady, and Cripps (2005), and Keller and Rady (2015) study delay and experimentation when others' actions generate information. Frick and Ishii (2024) characterize waiting and adoption in a decentralized forward-looking social-learning game; their equilibrium is the natural free-supply benchmark here. Empirical and theoretical diffusion work documents social learning and explains S-shaped take-up through networks, contagion, or heterogeneous thresholds (Rogers, 2003; Young, 2009; Mossel, Sly, and Tamuz, 2015; Akbarpour and Jackson, 2018). I instead ask how a committed principal should control access when waiting itself is strategic.

Viewed from the design side, a related literature changes incentives through prices, recommendations, or externalities. Bonatti (2011) studies pricing and learning for a durable good, Che and Hörner (2018) recommendation systems, and Laiho, Murto, and Salmi (2025, 2026); Chen and Zhang (2025) endogenous learning from irreversible adoption. Artificial scarcity appears in DeGraba (1995); Nocke and Peitz (2007); Möller and Watanabe (2010) and, as a way to reshape observed actions, in Parakhonyak and Vikander (2023). Here the instrument is dated capacity, information arrives through public experimentation, and the principal trades screening against the speed of learning.

Finally, the finite model has a dynamic-mechanism interpretation that connects the two sides. Anonymous capacity and reapplication generate distributions over service dates without transfers, and the three-type example shows that rejection-contingent continuation can screen values. The continuum model exploits single crossing through an ordered direct representation, but the staircase theorem decentralizes that representation through open access and cumulative inventory caps. Thus the main contribution is not a convenient restriction to direct assignments: it is a global solution to an operational rollout problem, together with the distributional incidence of the resulting acceleration.

Roadmap. Sections 2–6 develop the environment, free-supply benchmark, and finite design results. Sections 7–8 establish staircase implementation, characterize globally optimal continuum rollout, and derive the welfare and density results. Section 9 discusses applications and interpretation. Core proofs are in Appendices A and B; numerical certification and auxiliary approximation results are collected in the Online Supplement.

2 Environment

The same primitives support both the finite operational analysis and the continuum characterization. I therefore describe the common payoff and learning environment first, and then distinguish the two policy representations when supply is introduced.

2.1 Agents, values, and payoffs

A unit mass of agents is indexed by $i \in I = [0,1]$. In the finite model, values belong to

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

with type masses $q_n > 0$ summing to one. Later, values are distributed according to a cdf F with compact support $[\underline{v}, \bar{v}]$.

There is a persistent state $\omega \in \{g, b\}$. Agents and the principal share prior $\mu_0 = \mathbb{P}(\omega = g) \in (0,1)$. Successful adoption is irreversible and yields v_i 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 $r > 0$.

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

$$g(v,t) := e^{-rt}(v - B_t). \tag{1}$$

[1]

2.2 Public learning

Given these payoffs, public information evolves through experimentation. Bad news arrives only in the bad state. Conditional on $\omega = b$, its instantaneous intensity is βM_t, where M_t is cumulative successful adoption and $\beta > 0$. A bad signal reveals the state and ends adoption. Along the no-signal path,

$$B_t = B_a \exp\left\{-\beta \int_a^t M_s \, ds\right\}, \qquad t \geq a. \tag{2}$$

---

[1] Heterogeneity in v can equivalently be interpreted as heterogeneity in prior beliefs. Suppose agents share a common good-state payoff $\bar{v}$ but type i has initial bad-state odds B_{0i}, with all types facing the same likelihood process. Multiplying type i's payoff by the positive constant B_0/B_{0i} yields the common-prior formulation with $v_i = (B_0/B_{0i})\bar{v}$. Thus higher v_i corresponds to greater initial optimism. The transformation preserves individual choices, equilibrium paths, and the principal's completion-time objective, though not an unweighted utilitarian welfare sum without a corresponding change in welfare weights.

Thus adoption is 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 move immediately afterward.

2.3 Supply, applications, and rationing

The principal affects this learning process only through access. She commits at date zero to a right-continuous nondecreasing cumulative supply path S. In the finite operational model, 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. Anonymous acceptance is

$$Q_t = \begin{cases} 0, & M_t^- = S_t, \\ 1, & M_t^- < S_t \text{ and } A_t \le S_t - M_t^-, \\ \dfrac{S_t - M_t^-}{A_t}, & A_t > S_t - M_t^- > 0. \end{cases} \tag{3}$$

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

Agents of one value need not choose the same realized application date. Under free supply, a value class can adopt over an interval. Under scarcity, a split type can be represented either by asymmetric one-shot dates or by common application followed by random rejection and reapplication. I allow behavioral mixed strategies after every private history. Atomlessness makes the induced aggregate paths deterministic conditional on the public history.[2]

2.4 Continuation values and equilibrium

Because rejection preserves eligibility, optimal behavior must be defined after every private rejection history. Fix conjectured aggregate paths (M, A) and the induced acceptance path Q. A pure continuation plan α is a sequence of future application dates, contingent on the agent's private rejection history. If T^α is the realized set of attempts after date t, the normalized continuation value of type v is

$$W_t^\alpha(v) = \sum_{s \in T^\alpha,\, s > t} Q_s \prod_{k \in T^\alpha :\, t < k < s} (1 - Q_k) e^{-r(s-t)} \left[v - B_t e^{-\beta \int_t^s M_u \, du} \right]. \tag{4}$$

The value of a behavioral strategy is obtained by integrating (4) over private randomization. Let $W_t(v)$ be the supremum over continuation plans.

Definition 1 (Equilibrium)**.** Given a supply path S, an equilibrium consists of behavioral application strategies and aggregate paths (M, A) such that: (i) after every private history each agent maximizes continuation value; (ii) every date used with positive probability is optimal for that type at that history; (iii) A and M are generated by the strategies and (3); and (iv) $M_t \le S_t$ for all t. When several equilibria exist, the principal selects her preferred equilibrium.

---

[2] Individual rationing outcomes remain private and random. Determinism refers only to aggregate applicant and service masses, which obey the corresponding acceptance probabilities almost surely in the atomless population.

The equilibrium-selection convention matters at weak inequalities. It is the same convention used in the one-value unraveling result and in the three-value construction, where type 2 and the terminal type are indifferent at selected dates.[3]

2.5 The principal

Against this equilibrium response, the principal fixes a target $\bar{M} \in (0,1)$ and minimizes

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

along the selected no-signal equilibrium path. In the finite-type results, the target is high enough to require positive service from every listed type; in particular, the two-type analysis assumes $q_H < \bar{M} < 1$. In the continuum sections, $\bar{M} < 1$ instead selects an upper portion of the value distribution. This distinction keeps the finite operational results separate from the continuous-value type-pure scheduling problem.

A release is *immediately exhausted* if all new capacity is allocated at its release date. A plan is K*-release* if it contains at most K positive jumps. In an N-type economy, call a plan *simple* if it has at most N positive releases and every release is immediately exhausted. The conjecture that some optimal finite plan is always simple is rejected below.

Lemma 1 (Upper-block willingness)**.** *Fix a public and private history and any continuation strategy after rejection. The gain from applying at a positive-capacity date rather than following that continuation is strictly increasing in v. Hence willing residual agents form an upper block, with at most one indifferent cutoff type.*

The lemma orders residual types' willingness to apply at a given history; it does not order instantaneous actions after some types have already left the market. It also does not imply one service date per type: a cutoff type can be rationed, rejected, and served later.

3 Free supply and endogenous S-shapes

Before asking how access should be restricted, I characterize the benchmark in which it is not restricted at all. Free supply means $Q_t = 1$ whenever an agent applies. An individual therefore applies at most once, although a value class can be spread across dates through within-type mixing.

3.1 One value and the free-supply ridge

Let all residual agents have value v. Their payoff from certain service 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 (2),

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

---

[3] Principal-preferred selection is used to choose among equilibria generated by weak inequalities; it is not a refinement requiring all indifferent agents to apply at once. Such a rule would eliminate the mixed stopping that supports gradual free-supply adoption.

The indifference ridge is

$$B_t = \frac{rv}{r + \beta M_t}. \tag{5}$$

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

$$\dot{M}_t = M_t(r + \beta M_t). \tag{6}$$

Thus cumulative adoption 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

$$\Delta(m_a, m_b) = \int_{m_a}^{m_b} \frac{dm}{m(r + \beta m)} = \frac{1}{r} \log \left[\frac{m_b(r + \beta m_a)}{m_a(r + \beta m_b)} \right]. \tag{7}$$

Proposition 1 (One-class free-supply equilibrium)**.** *If $v > B_0$, the free-supply equilibrium has the unique aggregate path obtained by jumping, when necessary, to the smallest feasible stock on* (5) *and then following* (6) *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 finite values

The one-class ridge extends recursively across finitely many values. With $v_1 > \cdots > v_N$, the free-supply path alternates between class-specific ridge segments and waiting intervals. By single crossing, a higher class is exhausted before a lower class 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.

This structure clarifies why a finite-type path need not be smooth and why the model must permit within-type mixing. A claim that all agents of one type adopt at one date would contradict the free-supply benchmark. The relevant object is the aggregate path, which may be generated by asymmetric pure dates or by a common mixed stopping rule.

3.3 An atomless benchmark and the S-shape

Passing from positive-mass classes to an atomless distribution removes the waiting gaps between adjacent values and yields a scalar cutoff representation. Let values be continuously distributed with cdf F. If v_t is the marginal active value and $M_t = 1 - F(v_t)$ is the mass already served from the upper tail, local indifference gives

$$v_t = B_t \left(1 + \frac{\beta M_t}{r} \right), \qquad M_t = 1 - F(v_t). \tag{8}$$

Equation (8) implicitly determines the aggregate free-supply path. It reduces the decentralized equilibrium to a scalar cutoff. Write

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

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

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

(a) With finitely many values, types enter in descending value order. At date zero, all classes strictly above a uniquely determined marginal class adopt in full, and a uniquely determined fraction of that marginal class may also adopt. Any residual mass of the marginal class then follows (6)*. Each lower class enters after a uniquely determined waiting interval and follows its own unique ridge segment until exhausted.*

(b) If F has a positive continuous density and $b(1) < B_0 < b(0)$, 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

$$B_t = b(M_t), \qquad \dot{M}_t = -\frac{\beta M_t b(M_t)}{b'(M_t)}, \qquad M_0 = m_0^F, \tag{9}$$

until the support is exhausted. Uniqueness concerns the aggregate path; individual stopping assignments within an atomless class may differ on null sets.

For a uniform distribution, differentiation gives a transparent curvature test. Early adoption accelerates because the learning stock is growing; late adoption decelerates because the remaining types become increasingly reluctant as the marginal value approaches the lower boundary of the support.

Proposition 3 (Uniform diffusion curves)**.** *Let $v \sim U[\underline{v}, \bar{v}]$. Suppose the initial marginal type is interior and value dispersion is large enough that adoption traverses a nondegenerate interior range. Then the free-supply equilibrium is strictly increasing and, under an explicit prior threshold, has a unique inflection point: it is initially convex and eventually concave.*

Figure 2 illustrates the resulting strategic S-shape. The uniform case is useful because the inflection condition is closed form. The economic force is more general: over an active interval with positive density and sufficient dispersion, a growing learning stock can initially dominate the decline in marginal value before that decline eventually slows adoption. Thus, free availability provides a unique diffusion benchmark and shows how strategic waiting can generate an aggregate S-shape. The same waiting incentive creates a role for supply policy, to which I now turn.

4 Finite-plan normalization

Having established the free-supply benchmark, I now turn to scarce design. Before solving that problem, I remove two inessential complications. First, for any fixed upper bound on the number

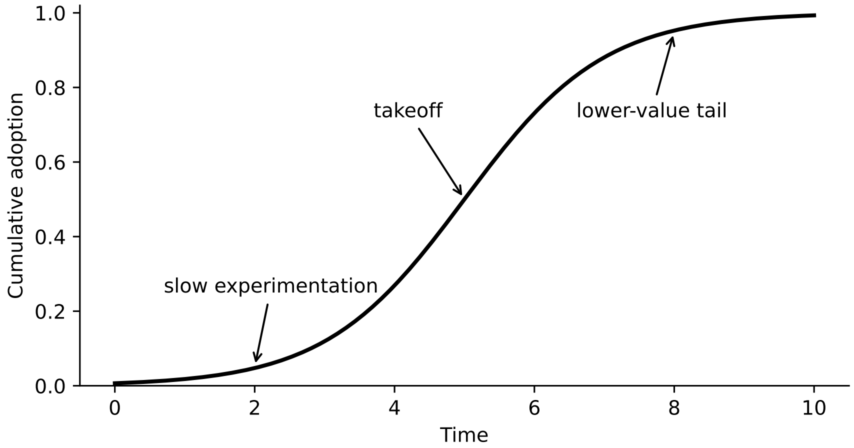

Figure 2: A strategic free-supply S-curve. Adoption is initially slow because little information is being generated. As adoption accumulates, learning accelerates and adoption speeds up. Later, the agents who have not yet adopted stand to gain less from successful adoption, so diffusion slows again.

of release dates, the selected-equilibrium problem is well posed. Second, unused inventory is never needed: any outcome can be represented by postponing unused capacity until the date at which it is actually allocated. The substantive complexity studied below therefore comes from the timing and rationing of exhausted releases, not from idle stock.

4.1 Fixed-release existence

Although the later counterexample rejects an N-release bound, finite optimization remains well posed once the calendar size is fixed. Fix K. Release dates, realized service masses, type-specific application shares, and rejection-contingent continuation values give a finite-dimensional formulation. Weakly ordered dates and zero-capacity dummy releases make the parameter set compact after imposing any feasible completion-time upper bound, while sequential optimality is represented by closed complementarity inequalities.[4]

Proposition 4 (Fixed-K existence)**.** *For every finite economy and fixed upper bound K on releases, the principal's selected-equilibrium problem has an optimal K-release plan.*

Thus finite-calendar optimization is well posed for every fixed K. The number of release dates remains an independent design margin.

4.2 Immediate-exhaustion representation

Existence alone does not yet identify a useful normal form. Suppose a release leaves inventory unused until a later date. Postpone the unused increment to the first date at which it is actually allocated. This leaves earlier acceptance opportunities and all realized service unchanged. Repeating this operation removes slack inventory.

---

[4] Writing realized service directly, rather than treating the acceptance ratio in (3) as a primitive at zero demand, removes the apparent discontinuity. The Appendix gives the compact parameterization.

Proposition 5 (Immediate-exhaustion representation). *For every optimal outcome among plans with at most K releases, there is an optimal representative with the same adoption path in which every positive retained release is immediately exhausted.*

The proposition preserves a useful normalization even though the batch-count bound fails. The principal may need more releases than values, but she never needs to leave released units idle in order to implement a given optimal path.[5]

5 Scarcity with one and two values

The preceding normalization lets us focus on the economics of scarce timing. In what follows, every positive release may be treated as immediately exhausted, and fixed-calendar optima exist. The remaining question is how many exhausted opportunities are useful and which values should face them.

5.1 One value: scarcity eliminates within-class delay

Start with the case in which scarcity has no screening role because all target-relevant agents have the same value. Suppose the unique target-relevant value satisfies $v > B_0$ and $\bar{M} < 1$. Release exactly $\bar{M}$ units at date zero. If all agents wait until a positive date, any one agent prefers applying slightly earlier because beliefs are unchanged before adoption begins and discounting is smaller. If adoption is gradual while inventory remains, an agent assigned to a later date prefers to apply just before that date. If the target is reached at a positive date through a final rationed rush, an applicant prefers an arbitrarily earlier chance at the same scarce stock. These deviations unravel every delayed equilibrium selected by the principal.

Theorem 1 (One-value optimum). *With one target-relevant value $v > B_0$ and target $\bar{M} < 1$, releasing exactly $\bar{M}$ units at date zero reaches the target immediately. No policy can do better.*

The result illustrates the core force of scarcity in its purest form. Free supply creates an unconditional option to wait. A capacity cap turns delay into a risk of exclusion and eliminates same-type free-riding.

5.2 Two values: exposure and incentive lower bounds

Adding a second value introduces a genuine tradeoff between early experimentation and the high type's incentive to wait. Let $v_H > B_0 > v_L$, let the high-type mass be q_H, and suppose $q_H < \bar{M} < 1$, so completion requires some low-type service. An optimal policy must satisfy two independent lower bounds.

First, even if all high agents adopt immediately, the low type cannot become willing before exposure drives belief from B_0 to v_L:

$$T \geq T_E := \frac{1}{\beta q_H} \log \frac{B_0}{v_L}. \tag{10}$$

---

[5] The operation postpones slack inventory to the date at which it is first allocated. It does not merge distinct exhausted releases and therefore does not imply any bound on their number.

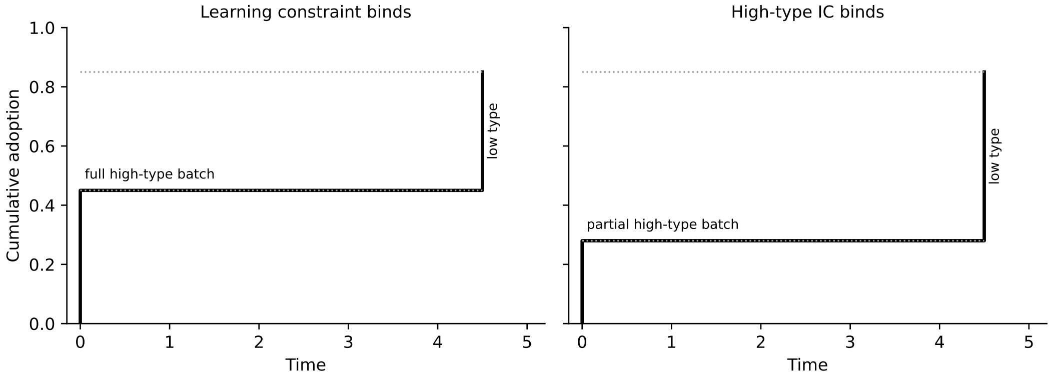

Figure 3: Optimal adoption paths with two values. Left: when learning is the binding constraint, all high-value agents adopt at date zero and the low-value agents enter at the earliest feasible date; the principal attains the first-best learning bound. Right: when the high type's incentive to wait is binding, only part of the high type is served at date zero, and the residual high-type mass joins the low type at the terminal batch.

Second, a high type served at zero must not prefer waiting for certain terminal service with the low type. Since terminal participation is fastest at $B_T = v_L$,

$$v_H - B_0 \geq e^{-rT}(v_H - v_L), \tag{11}$$

which is equivalent to

$$T \geq T_H := \frac{1}{r} \log \frac{v_H - v_L}{v_H - B_0}. \tag{12}$$

The principal can attain the larger of the two bounds. Let

$$T^* = \max\{T_E, T_H\}, \qquad m_0^* = \frac{1}{\beta T^*} \log \frac{B_0}{v_L}.$$

Release m_0^* at zero and the residual target mass at T^*. If $m_0^* < q_H$, all high agents may apply at zero and be rationed, or a specified mass m_0^* may be assigned to zero. Rejected or postponed high agents join the terminal batch. Indifference makes the two representations equivalent.

Theorem 2 (Unrestricted two-value optimum)**.** *Suppose $v_H > B_0 > v_L$ and $q_H < \bar{M} < 1$. The minimum completion time over all finite supply plans is*

$$T^* = \max\left\{ \frac{1}{r} \log \frac{v_H - v_L}{v_H - B_0}, \frac{1}{\beta q_H} \log \frac{B_0}{v_L} \right\}.$$

An optimal plan uses one initial and one terminal immediately exhausted release, with initial service mass m_0^* as above and certain terminal service for every residual target agent.

Figure 3 displays the two regimes. The theorem marks the positive boundary of the finite model. It does not rely on a universal compression theorem. With only one higher type to discipline and one lower entry condition, a single initial lottery and a terminal clearing opportunity span the relevant margins.

Comparative statics. Completion is weakly decreasing in β and q_H. A higher discount rate relaxes the high-type waiting constraint but also changes the free-supply ridge; in the two-type formula its direct effect is to reduce T_H. A more optimistic prior lowers both the exposure bound and the high-type incentive bound, and therefore weakly speeds completion.

5.3 Scarcity relative to free supply

Taken together, the one- and two-value results are not isolated curiosities. They reveal a general reason why dated capacity can strictly improve on unrestricted availability. Under free supply, the last target-relevant type enters only when its continuation value from waiting has fallen to its current service payoff. A small amount of scarcity at that entry margin changes the comparison discontinuously: an agent who postpones application may lose the unit rather than merely receive it later. The principal can therefore advance at least part of the marginal type without reducing the amount of prior experimentation.

The argument is easiest to see at a candidate free-supply completion date T^F. Let v_L be the marginal target-relevant value. Immediately before T^F, an active v_L agent is indifferent between service and continued waiting. Replace free availability near T^F by a terminal capacity equal to the residual target mass. If the applicant mass strictly exceeds that capacity, the batch rations and exhausts. Applying slightly earlier becomes strictly attractive because it increases the probability of receiving the scarce unit. If applicant mass equals capacity, an arbitrarily small reduction in terminal capacity creates the same force. The resulting extra adoption raises the learning stock before the original completion date.

Proposition 6 (Strict value of scarcity)**.** *Suppose the free-supply equilibrium reaches a target* $\bar{M} < 1$ *at a positive finite date and the last target-relevant type has positive residual mass at entry. Then there exists a finite dated-capacity policy whose selected equilibrium reaches the target strictly earlier.*

The proposition isolates the economic force: eliminating an unconditional future option can improve completion even before the global form of the optimal policy is known.

Private and social welfare. The principal's objective is completion time, not utilitarian welfare. Scarcity can nevertheless improve adopter welfare in some regions because it brings information production forward and can reduce the waiting costs borne by later agents. It can also hurt agents who are rationed or induced to experiment earlier. At this stage, however, the welfare comparison is only qualitative. The continuum geometry below sharpens it: a profitable wave has an ordered incidence, with losses at its front and gains at its back.

6 Three values: repeated rationing defeats the batch-count bound

The two-value theorem suggests that one exhausted release per value class might suffice. That inference is exactly what fails with three values. The reason is not idle inventory: by Proposition 5, every positive release can be taken to exhaust immediately. The new instrument is instead a *sequence of rationed opportunities* for the same intermediate type.

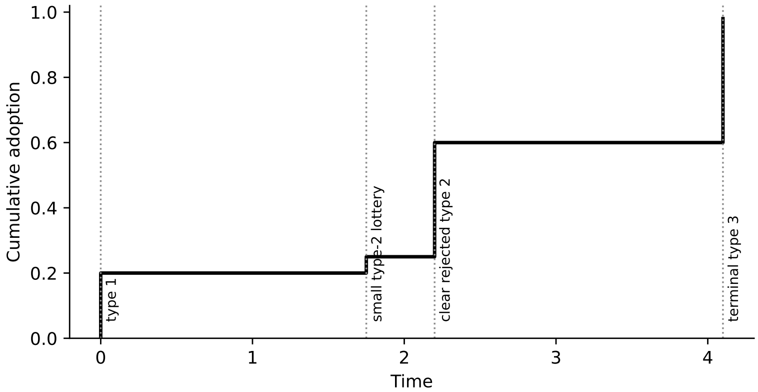

Figure 4: The four-release plan. Type 1 is served immediately. Type 2 receives a small early rationed opportunity and, after rejection, a later clearing opportunity. The final release serves the marginal lower type. Appendix A verifies the construction; the exhaustive binding-case enumeration and numerical certificate are in the Online Supplement.

In the counterexample, the top type is served immediately. The intermediate type then applies at two interior dates: a small fraction is accepted at the first date, while rejected agents remain eligible and are cleared at the second. The lowest target-relevant type enters only at the terminal date. The first intermediate opportunity is deliberately small because a larger early probability of service would attract the top type away from date zero. The second opportunity is deliberately large because postponing most of the intermediate type until the terminal date would sacrifice learning. Thus the two interior releases perform different incentive tasks and cannot generally be collapsed into one.

Theorem 3 (Simple-plan optimality fails with three values)**.** *There exists a three-value economy in which a feasible plan with four immediately exhausted releases completes strictly earlier than every immediately exhausted plan with at most three positive release dates. Consequently, the claim that every finite N-type economy has an optimal plan with at most N immediately exhausted releases is false already for $N = 3$.*

Figure 4 summarizes the mechanism. The comparison in Theorem 3 is global within the simple class, not merely a comparison with a single candidate three-date policy. Appendix A verifies the candidate construction. The Online Supplement gives the structural reduction, enumerates the upper-block applicant modes, reduces the best-simple-plan problem to a finite comparison of binding cases, and records the numerical certificate. The strict gap persists under sufficiently small perturbations of the primitives.

The lesson is narrow but important. Simple-plan optimality is not universal, and repeated eligibility can be economically substantive. The counterexample's operational equilibrium uses rejection-contingent rationing in an essential way: all intermediate agents apply at the first interior date, only a fraction are served, and rejected agents remain eligible for the later clearing release. Replacing that lottery by an ex ante partition into certain dates would change the anonymous-rationing deviation faced by the top type and is not part of the proof. This is also why the continuous-value analysis below takes ordered type-pure schedules as its primitive policy space rather than deriving them as unrestricted limits of the operational problem.

When additional finite releases are redundant. The counterexample does not imply that more release dates are always useful. The Online Supplement gives an exact continuation-value dominance condition: if rejecting an early opportunity preserves certain service at a later clearing date and every residual higher type weakly prefers that continuation to early service, the earlier rationed offer cannot separate the block and can be deleted. A primitive mass–gap inequality makes this condition directly checkable. These auxiliary results delimit the mechanism behind the counterexample without interrupting the transition to the continuum problem.

7 Continuous values: capped availability and type-pure scheduling

The finite results leave two lessons. Positive-mass rationing can be a genuine screening instrument, but it also makes a general finite-type solution combinatorially rich. Atomlessness changes the policy object and restores a global characterization. The finite operational model allows applications, anonymous rationing, rejection, and reapplication because each value class has positive mass. With a continuous value distribution, I first study capped availability, an operational aggregate-supply policy, and then solve the ordered type-pure scheduling problem. This continuum policy space retains positive-mass pools and smooth aggregate service without adding a separate type-by-type lottery technology.

7.1 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 thresholds while capacity remains. The latter is the value whose certain-service payoff is locally flat at date t.

Proposition 7 (Unique aggregate path under capped availability)**.** *Suppose F has a positive continuous density on the relevant support, $\bar{M} \in (0,1)$, and*

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

Define

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

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

$$b(m_0) = B_0, \qquad b(m_1) = \nu(\bar{M}). \tag{13}$$

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

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

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

$$m_0 = 1 - F(v_0^E), \qquad \bar{M} - m_1 = F(v_T^E) - F(v_T^M). \tag{15}$$

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

Thus, even before solving the full design problem, one up-front release already generates three endogenous adoption phases: an initial rush, smooth diffusion, and a final stockout. No mid-course supply intervention is needed to implement this path.

7.2 Ordered type-pure schedules, staircase supply, and the separating ridge

Capped availability suggests that a supply path can decentralize a finely ordered service pattern without observing values. To characterize all such patterns, the planner chooses an ordered type-pure schedule. A single service date is assigned to almost every served value, higher values are served weakly earlier, and positive-measure intervals of values may be pooled. This is the policy space generated by aggregate supply. A pool has positive mass and can be served by a positive release, while a smooth segment represents continuous aggregate service. A single value in an atomless distribution, however, has zero mass, so a release directed only to it is neither a supply jump nor an anonymous rationing opportunity under (3). Type-by-type service lotteries would require an additional allocation kernel. Relaxed joint value–service measures can represent limits of finite optimal plans, but those limits need not be type-pure and are not the object studied here.

To make this representation precise, an ordered type-pure schedule assigns a single date $\tau(m)$ to every served upper-tail rank $m \in [0, \bar{M}]$, with τ weakly increasing. Type purity permits a positive interval of values to share a date but rules out assigning a nondegenerate distribution of service dates to a positive set of values. Let λ denote Lebesgue measure on the rank interval $[0, \bar{M}]$. Since rank is measured in population mass, the adoption and belief paths are

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

Global direct incentive compatibility requires

$$e^{-r\tau(m)}\{\nu(m) - B_{\tau(m)}\} \geq e^{-r\tau(n)}\{\nu(m) - B_{\tau(n)}\} \qquad \text{for all served } m, n. \tag{16}$$

The ridge and density analysis below is global within this type-pure policy space. Although the schedule is written as a rank-to-date map, it need not be administered by observing ranks. Proposition 9 shows that every finite-wave feasible schedule is exactly decentralized by anonymous applications to a finite staircase cumulative-supply path.

Given this representation, local IC on a smooth separating component makes the marginal type indifferent across neighboring service dates. Writing $\nu(m) = F^{-1}(1 - m)$, the same calculation as under free supply gives

$$B^*(m) = \frac{r\nu(m)}{r + \beta m}. \tag{17}$$

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

$$\tau'(m) = \frac{x^{*\prime}(m)}{\beta m}. \tag{18}$$

Thus smooth passage from m_a to m_b takes

$$\mathcal{S}(m_a, m_b) = \int_{m_a}^{m_b} \frac{x^{*\prime}(m)}{\beta m} \, dm. \tag{19}$$

Local indifference identifies the ridge, but a global characterization must also account for pools, pauses, and boundary behavior. The next proposition extends the ridge identity to the entire type-pure schedule.

Proposition 8 (Boundary cores and exact type-pure decomposition)**.** *Suppose ν is positive and strictly decreasing, $\bar{M} \in (0, 1)$, and*

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

Let τ be any finite-completion ordered type-pure schedule 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 the unique $0 < m_0 < m_1 < \bar{M}$ solve

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

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 cores, 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*

$$T = t_0 + \mathcal{S}(m_0, m_1) + \sum_j \Gamma_F(d_j, p_j, h_j), \tag{20}$$

where

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

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

The proposition has two immediate implications. First, m_0 and m_1 are unique for every admissible density, without imposing $f' \leq 0$. What remains distribution-dependent is the optimal excursion family. The proposition also rules out a singular-continuous service-time assignment and permits a finite or countable number of smooth and pooled components. In any optimum the removable idle time satisfies $t_0 = 0$.

The second implication closes the apparent gap between the direct schedule and the original supply instrument.

Proposition 9 (Anonymous staircase implementation)**.** *Let* τ *be a finite-completion globally incentive-compatible ordered type-pure schedule 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 value receives one ordered certain service date induces a globally incentive-compatible ordered type-pure schedule.*

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. Thus $p - d$ and $h - p$ are adoption jumps, not the two announced supply quantities: the supply increments are p and $\bar{M} - p$.

Finally, the displayed inequalities impose the generic interior case in which the initial boundary core and the active ridge interval are both nondegenerate. Outside this case, the initial core 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 value cannot initiate positive adoption, and no positive finite-completion rollout exists. Accordingly, "initial batch–smooth diffusion–terminal batch" is one interior realization of the general normal form, not the characterization itself.

7.3 Smooth rollout and pooled waves

Having established both the schedule representation and its supply implementation, it remains to describe the only way an ordered rollout can depart from smooth separation. A pooled wave replaces a segment of smooth separation by a discrete increase in adoption, a period of waiting while beliefs improve, and a second increase that returns the schedule to the ridge. Pooling need not be absorbing: smooth rollout can resume after the wave.

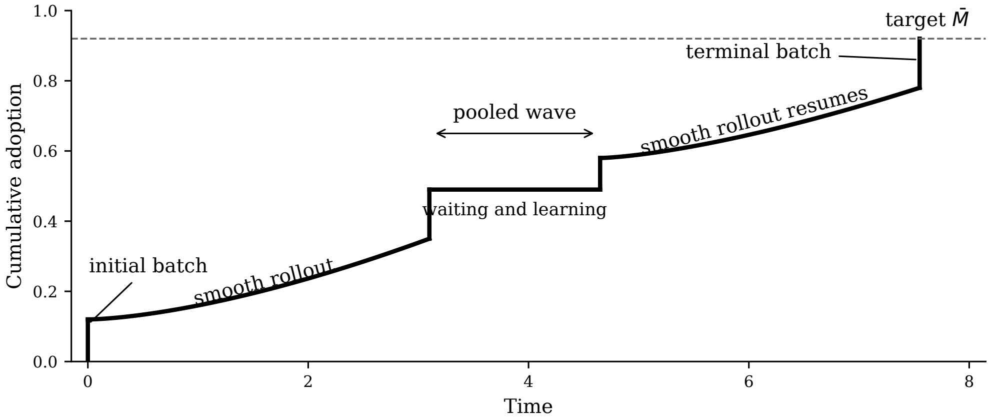

Figure 5: Smooth rollout and pooled waves in the type-pure continuum model. An interior wave produces a discrete increase in adoption, followed by a waiting interval while public beliefs adjust; smooth rollout may then resume. The final jump reaches the adoption target $\bar{M}$.

Figure 5 illustrates the possible components of an ordered service schedule. With the architecture now fixed, the next section compares the duration of a pooled excursion with smooth passage through the same public states.

This exact implementation result is distinct from finite-economy approximation. The continuum problem is defined directly in the atomless economy. The Online Supplement records a finite-menu recovery result for strongly regular type-pure schedules. Appendix B proves the result needed for optimal design here: every countable optimal excursion family is the limit of finite staircase implementations. Thus finite staircase policies attain the global continuum value arbitrarily closely without claiming that unrestricted finite operational optima purify into type-pure schedules.

8 Global characterization of optimal rollout

The decomposition in Proposition 8 identifies all possible components of a continuum rollout. Because every excursion returns to the same ridge state that smooth passage would have reached, those components can now be priced and selected. This section turns the architecture into a global solution. Each pooled excursion receives an exact completion-time weight, compatible excursions compete only through overlap of their mass supports, and the planner selects the best finite or countable family. Capped availability and the optimal one-wave policy then emerge as transparent special cases of the same problem.

Recall

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

and, in the generic interior case,

$$T^{\text{cap}} = \mathcal{S}(m_0, m_1) = \int_{m_0}^{m_1} \frac{x^{*\prime}(m)}{\beta m}\, dm.$$

This is the duration of the excursion-free ridge passage between the compulsory boundary contacts.

8.1 Exact excursions and the global selection problem

The first step is to price 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

$$\Delta(d, p, h) = \frac{1}{\beta p} \log \frac{b(d)}{b(h)}. \tag{21}$$

The departure batch serves $[d, p]$, adoption remains at p for $\Delta(d, p, h)$, and the reentry batch serves $[p, h]$. The public state after reentry is exactly the state at which ridge passage reaches h.

Proposition 10 (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*

$$b(d)^k\{a - b(d)\} = b(h)^k\{a - b(h)\}. \tag{22}$$

Whenever (22) *holds, replacing smooth passage over $[d, h]$ changes completion time by*

$$\Gamma_F(d, p, h) = \frac{1}{\beta} \int_d^h x^{*\prime}(m) \left(\frac{1}{p} - \frac{1}{m} \right) dm. \tag{23}$$

Negative Γ_F means that the pooled excursion is faster than the ridge segment connecting the same public states.

The proposition supplies both feasibility and price. The contact equation makes the marginal value $a = \nu(p)$ indifferent between the two batch lines. Single crossing assigns higher values to departure and lower values to reentry. A useful intervention therefore has three inseparable pieces: an early increase in experimentation, a pause while that larger stock generates information, and a

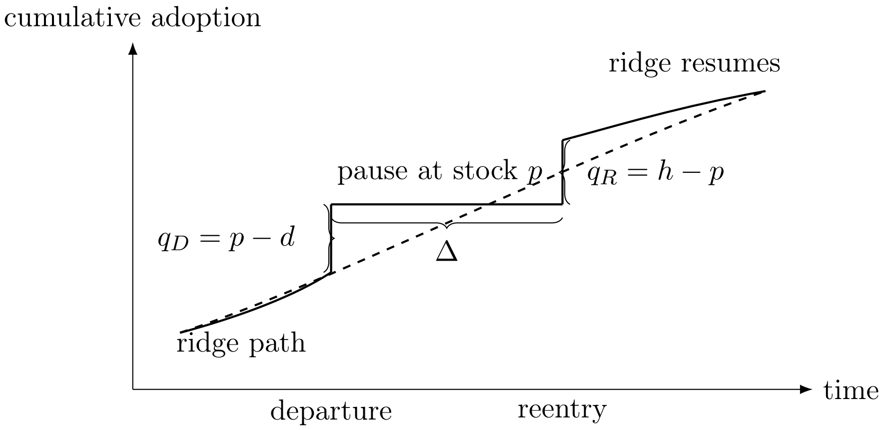

Figure 6: A pooled excursion. The departure batch advances mass $p - d$, the principal waits for $\Delta(d, p, h)$ while adoption remains at p, and the reentry batch serves mass $h - p$. The contact equation makes the two batch lines meet at the marginal value $\nu(p)$.

clearing release that returns the schedule to the ridge. A lone delayed batch cannot help, because pausing at d and later jumping to $h > d$ has excess duration

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

To pass from one intervention to a complete rollout, collect every exact excursion that can be inserted into the ridge.

Let

$$\overline{\mathcal{E}}_F = \Big\{(d, p, h) : m_0 \le d \le p \le h \le m_1,\ b(d)^{r/(\beta p)}\{\nu(p) - b(d)\} = b(h)^{r/(\beta p)}\{\nu(p) - b(h)\}\Big\}. \tag{24}$$

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

Theorem 4 (Global characterization of continuum rollout)**.** *Suppose the hypotheses of Proposition 8 hold and x^* is four times continuously differentiable on the active interval. Then the minimum completion time over all ordered type-pure schedules is*

$$T^* = T^{\text{cap}} + \min_{\mathcal{A} \in \mathfrak{A}_c} \sum_{e \in \mathcal{A}} \Gamma_F(e). \tag{25}$$

Every feasible schedule generates a compatible family with exactly this duration, and every compatible family reconstructs a globally incentive-compatible schedule. Every finite family is exactly implemented by an anonymous staircase supply path as in Proposition 9. The minimum is attained by a countable family $\mathcal{A}^*$. Its gain series is absolutely convergent, and for every $\varepsilon > 0$ a finite staircase policy implements a rollout whose completion time is within ε of T^*.

Together with Proposition 9, the theorem yields operational value equivalence: the global type-pure optimum is the infimum over finite anonymous staircase implementations, and it is attained by such a policy whenever the optimal excursion family is finite. Countable attainment adds closure, not a lower value unavailable to finite supply plans.

This equivalence is a complete variational characterization, not a claim that the optimizer always has a closed-form number of waves. Candidate excursions are weighted intervals. Disjoint candidates can be combined because each returns to the ridge state inherited by the continuation; overlapping candidates compete for the same mass interval. The number and endpoints of the selected waves are therefore determined by an exact weighted interval-selection problem. Appendix B gives the Bellman equation, countable-attainment argument, and an explicit finite-compression bound.

Accordingly, the continuum problem is globally solvable even though no universal "smooth then staircase" architecture exists. An optimal path may move smoothly, take a wave, return to smooth separation, and later take another wave. Boundary excursions may merge with the initial or terminal service line. The generic capped path is only the excursion-free member of this larger class.

8.2 Who bears the cost of faster rollout

The characterization does more than rank completion times: it also identifies who finances the acceleration. The same representation gives a precise distributional comparison. Fix a nondegenerate excursion $e = (d, p, h)$, normalize its departure date to zero, and compare it with free availability beginning from the same public state $(M, B) = (d, b(d))$. Under free take-up, rank $m \geq d$ would be served on the ridge after

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

and would obtain

$$U_d^F(m) = e^{-r\mathcal{S}(d,m)}\{\nu(m) - b(m)\}. \tag{26}$$

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 11 (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 weakly after reentry receives the same belief-contingent payoff as under free take-up, shifted by the excursion's excess duration:

$$U_d^e(m) = e^{-r\Gamma_F(e)} U_d^F(m) \qquad \text{for all } m \geq h. \tag{27}$$

Hence all such followers are strictly better off if and only if the excursion speeds completion, $\Gamma_F(e) < 0$.

(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 1 (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 state, its departure pool is worse off, every rank weakly after reentry is better off, and a unique cutoff separates losers from gainers inside its reentry pool.*

Hence a profitable wave is not a Pareto improvement. It accelerates learning by moving experimentation toward the higher-value agents in its departure pool. Those agents accept earlier service only because the staircase policy removes their free-supply option to wait. The gain is passed down the value ranking: sufficiently late members of the clearing pool, and every adopter who follows the wave, receive the same informational terms earlier. The theorem makes the incidence unusually sharp. A wave that lowers completion time imposes losses at its front and delivers gains at its back.

8.3 Capped availability as the excursion-free optimum

Having priced every feasible excursion, the principal's benchmark question has a simple answer: capped availability is optimal exactly when no admissible departure from the ridge has negative weight.

Theorem 5 (Exact capped-optimality criterion)**.** *Suppose the hypotheses of Proposition 8 hold. Capped availability minimizes completion time over the entire ordered type-pure policy space if and only if*

$$\Gamma_F(d, p, h) \geq 0 \qquad \text{for every } (d, p, h) \in \overline{\mathcal{E}}_F. \tag{28}$$

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

The theorem turns capped optimality into a sign test over excursions. That criterion is exact but nonlocal: the contact equation links two ridge states, and the gain integrates over the interval between them. A pointwise restriction on the density is therefore most useful as a readable sufficient condition for the universal inequality in (28).

Write f for the density. Since $\nu'(m) = -1/f(\nu(m))$,

$$x^{*\prime}(m) = \frac{\beta}{r + \beta m} + \frac{1}{\nu(m) f(\nu(m))}. \tag{29}$$

Corollary 2 (Excursion-free optimality under a nonincreasing density)**.** *Suppose* F *has a positive continuously differentiable density on the relevant support and* f *is nonincreasing. Then every nondegenerate feasible excursion has strictly positive excess duration. Hence the unique aggregate* *optimum is the excursion-free rollout determined by the boundary conditions. In the generic interior case, it is an initial batch of mass m_0, smooth ridge rollout from m_0 to m_1, and a terminal batch of mass $\bar{M} - m_1$.*

The density condition is sufficient, not necessary. A distribution may rise over some values and still satisfy (28): the increase may lie outside the active interval, may fail to generate an exact contact, or may be too weak to make its gain negative.

Remark 1 (Relation to the finite operational model). The finite and continuum analyses differ in whether a positive-mass type can be rationed, not in whether policy is operational. With finitely many values, anonymous rejection and reapplication can give one type a nondegenerate service-date lottery. In an atomless value distribution, a batch aimed only at one exact value has zero mass and is not an action of the aggregate supply process. The continuum model therefore rules out type-by-type lotteries, but Proposition 9 shows that its finite-wave schedules are implemented by the same anonymous cumulative-supply technology. Adding type-specific service lotteries would require an allocation kernel and would define a separate extension.

8.4 When waves help and how to design one

When the exact criterion fails, the next question is where profitable departures arise and how a planner should size them. Near a ridge contact point, reflected waves admit an expansion in primitives. Let $v = \nu(m)$ and write $f = f(v)$ and $f' = f'(v)$. For a wave of scale ε,

$$\Gamma_F = K_F(m)\varepsilon^3 + O(\varepsilon^4), \tag{30}$$

where

$$K_F(m) = \frac{2rf - \beta m v f' + 2\beta r v f^2/(\beta m + r)}{3\beta^2 m^3 v^2 f^3}. \tag{31}$$

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

$$\frac{f'(v)}{f(v)} > \frac{2r}{\beta m v} + \frac{2rf(v)}{m(\beta m + r)}. \tag{32}$$

Theorem 6 (Local primitive test for profitable batching)**.** *Let C be a compact interior ridge core on which the reflected-contact map is regular and the remainder in* (30) *is uniform.*

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

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

The local test identifies candidate gains, while the global theorem determines which compatible collection should be selected. For a single intervention, however, the choice is especially transparent.

Proposition 12 (Optimal one-wave design)**.** *Suppose F has a positive continuously differentiable density on the active interval. Among schedules containing at most one pooled excursion,*

$$T^{1\mathrm{W}} = T^{\mathrm{cap}} + \min_{(d,p,h)\in\overline{\mathcal{E}}_F} \Gamma_F(d,p,h), \tag{33}$$

and a minimizing triple exists. If its gain is negative, it implements

$$q_D^* = p^* - d^*, \qquad \Delta^* = \frac{1}{\beta p^*} \log \frac{b(d^*)}{b(h^*)}, \qquad q_R^* = h^* - p^*. \tag{34}$$

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

The compact program converts "batching helps" into an operational prescription. The principal chooses where to leave the ridge, how much mass to advance, and where to reenter; public-state matching then pins down the pause. For fixed plateau p, the contact equation has a reflected right endpoint, so the numerical problem is two-dimensional. Appendix B derives the interior first-order conditions and the local geometry of small waves. 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.

9 Applications and interpretation

The theory's empirical content is clearest in settings where access can be credibly staged and early use generates public evidence. Vaccination is the leading example. When early cohorts generate public evidence about safety or effectiveness, guaranteed future eligibility creates an option to wait. A credible inventory cap or dated expansion can bring experimentation forward. With a diffuse or nonincreasing density of adoption values, the excursion-free rollout is uniquely optimal; in the generic interior case, one up-front cap produces an initial rush, smooth sorting, and a terminal stockout. With a sharply rising density, many nearby agents have similar values, so advancing them together creates learning at relatively little screening cost. The interval-selection problem determines how many waves to use, and the one-wave program determines the two adoption jumps and the pause. Faster completion is not a Pareto improvement: the advanced experimenters lose relative to free take-up, while later followers gain from receiving the same information sooner.

Although vaccination provides the cleanest illustration, the same logic applies to agricultural extension, new treatments, digital platforms, cloud capacity, and artificial-intelligence services. Invitations, waitlists, eligibility rules, and usage tiers are staircase supply policies even when scarcity is administrative rather than physical. The implementation theorem makes the prescription literal: announce an access cap, leave unclaimed capacity available while adoption grows smoothly, hold the cap fixed after stockout, and raise it at reentry. A one-wave rollout requires only two supply announcements even though the induced adoption path contains two jumps, a pause, and smooth phases on either side.

These prescriptions rely on commitment. Public eligibility rules, procurement, visible waitlists, and reputational concerns can make an announced access path credible. Without commitment, a principal who can reopen access after slow take-up may be unable to sustain the scarcity that induces experimentation. The model should therefore be read as characterizing rollout when dated capacity can be committed, not as a prediction for every institution that informally announces a queue.

Subject to that institutional requirement, the comparative statics are most transparent through the policy formulas. Stronger learning shortens the pause for fixed contact ranks, and a larger plateau stock generates information faster. Discounting changes both the value of delay and the location of the separating ridge. Distributional shape is more fundamental: a nonincreasing density favors the excursion-free path, whereas a sufficiently steep local increase makes small waves profitable. Useful staggering is consequently not arbitrary stop-and-go rationing. It consists of a departure pool, a belief-matching pause, and a reentry pool whose sizes and location balance screening against learning.

10 Discussion and conclusion

The analysis begins from a simple tension: when adoption generates public information, abundant supply need not produce rapid diffusion. Forward-looking agents value the option to let others experiment first, and committed scarcity changes that option by making future access contingent.

Once access is scarce, the finite model isolates the dynamic-screening role of rationing. One type adopts immediately under a target-sized date-zero release, and two target-relevant types are optimally served by one initial and one terminal exhausted batch. With three types, however, an early chance of service and a later clearing opportunity can perform different incentive tasks. Rejection-contingent continuation is then a genuine design instrument, so the number of useful releases need not equal the number of values.

Atomlessness then changes the form of the problem without severing its operational link to supply. The continuous-value model yields a complete architecture. Almost every value receives one ordered service date, positive-measure intervals may pool, and every finite-wave schedule is exactly implemented by anonymous staircase supply. Every finite-completion rollout consists of optional boundary pools, smooth movement along a unique separating ridge, and a compatible finite or countable family of pooled excursions. Completion time is the excursion-free benchmark plus the sum of exact excursion gains, reducing global design to weighted interval selection. Capped availability is uniquely optimal under a positive nonincreasing density; rapidly rising density creates profitable waves; and the one-wave program determines the departure mass, pause, and reentry mass.

Finally, the same geometry that solves timing also delivers sharp welfare incidence. A completion-improving wave is not a Pareto improvement. The departure pool bears the cost of accelerated experimentation, while every adopter after reentry receives the same informational terms earlier; inside the reentry pool, one cutoff separates losers from gainers. Rollout design therefore determines not only how quickly information is produced, but also which adopters produce it for the benefit of whom.

Taken together, these results leave open questions about richer allocation technologies and simpler primitive descriptions, but not about the architecture solved here. A model with type-specific service lotteries would require an explicit random-allocation kernel, and a no-commitment analysis could overturn the scarcity that supports early experimentation. Within committed aggregate supply, however, the conclusion is complete: distributional shape determines whether rollout should be smooth, capped, or staggered, how multiple waves fit together, and who pays for the resulting speed.

References

Akbarpour, M. and M. O. Jackson (2018). Diffusion in networks and the virtue of burstiness. *American Economic Review* 108, 2458–2493.

Bandiera, O. and I. Rasul (2006). Social networks and technology adoption in northern Mozambique. *Economic Journal* 116, 869–902.

Banerjee, A. (1992). A simple model of herd behavior. *Quarterly Journal of Economics* 107, 797–817.

Besley, T. and A. Case (1993). Modeling technology adoption in developing countries. *American Economic Review Papers and Proceedings* 83, 396–402.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy* 100, 992–1026.

Bonatti, A. (2011). Menu pricing and learning. *American Economic Journal: Microeconomics* 3, 124–163.

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

Che, Y.-K. and J. Hörner (2018). Recommender systems as mechanisms for social learning. *Quarterly Journal of Economics* 133, 871–925.

Chen, W. and Q. Zhang (2025). Jump-start or gradualism? Dynamic incentives for innovation adoption. Working paper, April 2025.

Conley, T. and C. Udry (2010). Learning about a new technology: Pineapple in Ghana. *American Economic Review* 100, 35–69.

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

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

Dupas, P. (2014). Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment. *Econometrica* 82, 197–228.

Eyster, E. and M. Rabin (2014). Extensive imitation is irrational and harmful. *Quarterly Journal of Economics* 129, 1861–1898.

Foster, A. and M. Rosenzweig (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy* 103, 1176–1209.

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

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.

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

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

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

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.

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

Mossel, E., A. Sly, and O. Tamuz (2015). Strategic learning and the topology of social networks. *Econometrica* 83, 1755–1794.

Munshi, K. (2004). Social learning in a heterogeneous population: Technology diffusion in the Indian Green Revolution. *Journal of Development Economics* 73, 185–213.

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

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

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.

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.

Smith, L. and P. Sørensen (2000). Pathological outcomes of observational learning. *Econometrica* 68, 371–398.

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

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

Appendix

This appendix contains the core proofs cited in the main text. Detailed numerical certification of the three-value counterexample and auxiliary results on finite pooling, finite-menu recovery, benchmark competitors, additional curvature, and terminal padding are collected in the Online Supplement. Section numbers retain their original prefixes so that references remain stable. The results below establish aggregate-path uniqueness, capped-availability implementation, the finite operational benchmarks, the type-pure schedule decomposition, and the density comparisons. Throughout, 'service' means successful allocation after any rationing, and rejected applicants remain eligible unless a result states otherwise.

A.1 Operational equilibrium and lottery values

At date t, anonymous rationing gives acceptance 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 active. A realized application plan α specifies a sequence of dates contingent on the agent's own past failures. 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 value 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 v. For any two types $v > v'$, the difference between their values from one fixed lottery 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.[12]

Lemma A.1.1 (Upper-block willingness). *Fix any rejection-contingent continuation. If type v' weakly prefers applying now to that continuation, every type $v > v'$ strictly prefers applying now, unless immediate acceptance probability is zero.*

---

[12] 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 $a < e^{-rt}$. $\square$

A.2 Free supply and aggregate-path uniqueness

Under free supply, successful service occurs at the first application. 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

$$B_t = \frac{rv}{r + \beta M_t}. \tag{A.2.1}$$

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

$$\dot{M}_t = M_t(r + \beta M_t). \tag{A.2.2}$$

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 class is exhausted. Its second derivative is positive.

Proposition A.2.1 (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 value distribution, values enter in descending order. 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 (A.2.2)*. 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, put $\nu(m) = F^{-1}(1 - m)$ 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*

$$B_t = b(M_t), \qquad \dot{M}_t = -\frac{\beta M_t b(M_t)}{b'(M_t)}, \qquad M_0 = m_0^F, \tag{A.2.3}$$

until the support is exhausted.

Uniqueness is aggregate: asymmetric stopping assignments or mixing within a positive-mass atom may implement the same public path.

Proof. For part (a), single crossing implies that no lower value can enter before a higher residual value. Fix a stage at which all values 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 value ranking. Every atom whose service payoff is strictly decreasing even after all higher atoms and that atom 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 (A.2.2). 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 values 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 (A.2.3). 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, 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.

A.3 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 next result establishes the operational foundation of the capped path used in the main text.

Theorem A.3.1 (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 $0 < m_0 < m_1 < \bar{M}$ 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 value $\nu(\bar{M})$ 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 value 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 value v the derivative of certain-service payoff would jump upward by

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

A positive interval of values 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 $B_0 = b(m_0)$, 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 one of the certain-service payoff lines compared in Lemma A.11.2. At T, exactly the residual target mass applies and exactly exhausts the residual inventory; after T no unit remains. The same envelope calculation verifies that every served value prefers its prescribed date and every lower value prefers exclusion. 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.

A.4 One-value 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.

A.5 Two-value optimum

Let high mass be q_H. Any feasible plan completing at T must satisfy the exposure bound

$$\beta \int_0^T M_t dt \leq \beta q_H T$$

before low-type entry, because no low type has yet adopted and at most all high types can contribute. Since the earliest low-type participation has $B_T = v_L$,

$$\log \frac{B_0}{v_L} = \beta \int_0^T M_t dt \leq \beta q_H T.$$

This gives $T \geq T_E$.

A high type assigned positive probability of service at zero must obtain at least the value of waiting for the terminal certain opportunity. The fastest terminal belief is v_L, so

$$v_H - B_0 \geq e^{-rT}(v_H - v_L),$$

which gives $T \geq T_H$.

Set $T^* = \max\{T_E, T_H\}$ and

$$m_0^* = \frac{1}{\beta T^*} \log \frac{B_0}{v_L}.$$

Release m_0^* at zero. Keep the stock fixed until T^*, when belief equals v_L, and release the remaining target mass. The low type is willing. The high type's terminal deviation is weakly unprofitable by construction. If $m_0^* < q_H$, implement the initial mass by rationing all high agents or by deterministic splitting. This attains both lower bounds and proves optimality.

A.6 Three-value construction

The primitives are

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

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

The policy parameters are

$$s = 1.7489035062, \quad t = 2.2062097487, \quad Q = 0.1104189201, \quad T = 4.1402437122.$$

Beliefs are

$$B_s = 0.75e^{-0.1s} = 0.6296618038,$$

$$B_t = B_s e^{-0.5(0.2+0.4Q)(t-s)} = 0.5954712920,$$

$$B_T = B_t e^{-0.3(T-t)} = 1/3.$$

The service payoffs are

$$\begin{aligned} g_1(0) &= 0.11, \\ g_1(s) &= 0.1144327420, \\ g_1(t) &= 0.1094497875, \\ 0.95g_1(T) &= 0.0955048525, \end{aligned}$$

so

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

For type 2,

$$g_2(s) = g_2(t) = 0.95g_2(T) = 0.0846245529.$$

Type 3 obtains negative service payoff at s and t and zero at T. These calculations establish sequential feasibility under favorable tie selection.

The preceding calculations establish sequential feasibility of the four-release construction. The structural reduction, exhaustive binding-case comparison, and numerical certificate for all immediately exhausted plans using at most three release dates are reported in the Online Supplement (Sections A.6.1-A.6.3).

A.7 Fixed-K existence and exhaustion

Fix a horizon upper bound $\bar{T}$ obtained from any feasible policy. Permit K weakly ordered dates in $[0, \bar{T}]$ and zero-capacity dummy releases. Let y_{nk} be realized service mass of type n at date k, a_{nk} applicant mass, and z_{nk} residual mass. Beliefs are continuous functions of dates and cumulative realized service. Replace acceptance ratios by the identities

$$y_{nk} = Q_k a_{nk}, \qquad 0 \leq Q_k \leq 1,$$

with aggregate capacity constraints and complementarity. Sequential optimality is a finite family of weak inequalities comparing application and continuation. The feasible set is closed and bounded. Completion time is lower semicontinuous. Weierstrass gives a minimizer.

If release k contains unused inventory, reduce cumulative supply to realized service until the next release. No on-path acceptance probability or adoption changes. Repeating gives immediate exhaustion at every retained positive release. The operation postpones slack inventory; it does not consolidate distinct exhausted releases.

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 one 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. □

This proposition is conditional and concerns ordered type-pure menus. It does not prove that every such menu is sequentially implementable in the original repeated-application game.

A.10 Type-pure service schedules and exact excursion decomposition

Let ranks run from zero to the target, ordered from highest to lowest payoff. An ordered type-pure schedule is a weakly increasing right-continuous map assigning each rank one certain service date; flat intervals pool positive masses. This is the continuum policy space generated by aggregate supply. A release directed only to an exact atomless value has zero mass and is not represented by the cumulative supply path. Type-by-type lotteries require an additional allocation kernel. Relaxed finite-plan limits may be measure-valued and need not be type-pure; they are not used in the decomposition below.

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

For each service date t, write

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

Global direct IC means that $L_{\tau(m)}(\nu(m)) \ge L_{\tau(n)}(\nu(m))$ 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.

A.10.1 Boundary states and continuous service

Lemma A.10.1 (Boundary normalization and terminal zero rent). *Suppose* $\bar{M} \in (0, 1)$ *and*

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

Let the schedule be any finite-completion globally incentive-compatible ordered type-pure schedule serving exactly the upper target mass, with participation for served types and exclusion below the marginal target value. 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,*

$$B_T = \nu(\bar{M}). \tag{A.10.1}$$

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

$$b(m_0) = B_0, \qquad b(m_1) = \nu(\bar{M}). \tag{A.10.2}$$

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 (A.10.1).

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 (A.10.2). Since $B_0 > \nu(\bar{M})$, the first solution lies strictly below the second. $\square$

Lemma A.10.2 (Continuous service lies on the ridge)**.** *Under the hypotheses of Lemma A.10.1, let* μ_c *denote the continuous part of the Stieltjes measure* $d\tau$. *At every continuity point* m *of* τ *in* $\operatorname{supp} \mu_c$,

$$B_{\tau(m)} = b(m). \tag{A.10.3}$$

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

$$\tau'(m) = \frac{x'(m)}{\beta m}. \tag{A.10.4}$$

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

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

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 (A.10.3).

Beliefs never exceed $B_0 = b(m_0)$ and never fall below the terminal belief $b(m_1)$. Hence (A.10.3) 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 (A.10.3) give

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

Since $a \geq m_0 > 0$ 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 $\tau(c) - \tau(a) \leq C(c - a)$ 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

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

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

At almost every point with $\tau'(m) > 0$, differentiating (A.10.3) and using $d[-\log B_t]/dt = \beta M_t$ gives

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

which is (A.10.4). □

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.

A.10.2 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

$$\Delta(e) = \frac{1}{\beta p} \log \frac{b(m_d)}{b(m_r)}. \tag{A.10.5}$$

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

$$a - b(m_d) = e^{-r\Delta(e)}\{a - b(m_r)\}. \tag{A.10.6}$$

Equivalently, with

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

the contact condition is

$$q_p(m_d) = q_p(m_r). \tag{A.10.7}$$

Lemma A.10.3 (Exact excursion feasibility)**.** *Suppose ν is positive, continuously differentiable, and strictly decreasing on a neighborhood of $[m_d, m_r]$. If* (A.10.7) *holds and the contact state lies in the interior ridge range, then the excursion clock is uniquely given by* (A.10.5)*. The departure and reentry payoff lines cross exactly at $a = \nu(p)$. If the adjacent intersections before departure and after reentry are ordered, the complete dated menu is globally incentive compatible for all participating types.*

Proof. Strict monotonicity of b makes (A.10.5) 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 (A.10.6), 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. Apply this observation to the smooth prefix, departure line, reentry line, and smooth continuation. □

Lemma A.10.4 (Every waiting interval has virtual ridge endpoints)**.** *Under the hypotheses of Lemma A.10.1, 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

$$B^+ = B^- e^{-\beta p \Delta}. \tag{A.10.8}$$

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

$$a - B^- = e^{-r\Delta}\{a - B^+\}. \tag{A.10.9}$$

The right-continuous convention assigns rank p to the later line. Its participation constraint and (A.10.9) imply $a - B^+ > 0$ and $a - B^- > 0$.

By Lemma A.10.1, 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 (A.10.8)–(A.10.9) 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 (A.10.8)–(A.10.9) are then exactly (A.10.5)–(A.10.6). ☐

Lemma A.10.5 (A jump is implemented by its pooled blocks)**.** *For each jump at* p *with virtual endpoints* $d < p < h$*,*

$$\tau(m) = t^- \quad \text{for } d < m < p, \qquad \tau(m) = t^+ \quad \text{for } p < m < h. \tag{A.10.10}$$

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

$$h_i \le d_j. \tag{A.10.11}$$

Thus every jump is an actual departure batch, wait, and 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 (A.10.11).

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 (A.10.10).

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 A.10.2 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 (A.10.10). □

A.10.3 Global decomposition of arbitrary monotone schedules

Theorem A.10.1 (Type-pure schedule decomposition).
Assume the boundary and global-IC hypotheses above.

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

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

$$T = t_0 + \mathcal{S}(m_0, m_1) + \sum_{j \in J} \Gamma_F(d_j, p_j, h_j), \tag{A.10.12}$$

where

$$\Gamma_F(d,p,h) = \frac{1}{\beta} \int_d^h x'(m) \left(\frac{1}{p} - \frac{1}{m} \right) dm. \tag{A.10.13}$$

The series in (A.10.12) *is absolutely convergent.*

Proof. The normalization is feasible by Lemma A.10.1. We first establish the boundary batches. Continuous increase below m_0 is impossible: Lemma A.10.2 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

$B^+ \geq B_T = b(m_1)$ and therefore $h \leq m_1$, 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 A.10.2 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 A.10.4 and A.10.5 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 (A.10.12). Finally,

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

so

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

□

Remark 2 (Compulsory boundary cores versus observed boundary batches). The ranks m_0 and m_1 are unique without any restriction on the sign of f'. Every admissible type-pure schedule 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 cores. 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 cores are exactly the full initial and terminal batches.

A.11 Diffuse and concentrated density regions

A.11.1 Nonincreasing density and the *type-pure optimum*

For one feasible excursion abbreviate

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

The contact equation gives $q(m_d) = q(m_r)$. Direct differentiation yields

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

so q' is positive on (m_d, p) and negative on (p, m_r). Put

$$h(m) = \frac{1}{p} - \frac{1}{m}, \qquad \psi(m) = \frac{h(m)}{R(m)},$$

where $\psi(p)$ denotes the continuous limit.

Lemma A.11.1 (Convex quantiles make ψ decreasing)**.** *If ν is convex on the excursion interval, then ψ is strictly decreasing.*

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

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

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

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

Since $s'(m) \leq 0$ almost everywhere and s is nonincreasing, ϕ is strictly increasing and ψ strictly decreasing. The value at p follows by the continuous extension. $\square$

Lemma A.11.2 (Feasibility of the batch–smooth–batch schedule)**.** *Let m_0, m_1 solve* (A.10.2) *and define*

$$\tau^*(m) = \begin{cases} 0, & 0 \leq m \leq m_0, \\ \displaystyle\int_{m_0}^{m} \frac{x'(z)}{\beta z}\, dz, & m_0 < m < m_1, \\ \mathcal{S}(m_0, m_1), & m_1 \leq m \leq \bar{M}. \end{cases} \tag{A.11.1}$$

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 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

$$\frac{\partial L_m(v)}{\partial m} = e^{-r\tau^*(m)} b'(m) \left[\frac{r\{v - b(m)\}}{\beta m b(m)} - 1 \right] = e^{-r\tau^*(m)} b'(m) \frac{r\{v - \nu(m)\}}{\beta m b(m)}. \tag{A.11.2}$$

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 value 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 (A.11.2), and that line gives strictly negative payoff. Thus exclusion is optimal. The completion time follows from (A.11.1). □

Proposition A.11.1 (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* (A.11.1)*, 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

$$b(d)^k\{a - b(d)\} = b(h)^k\{a - b(h)\}. \tag{A.11.3}$$

When this condition holds, the completion time is

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

where

$$\Gamma_F(d,p,h) = \frac{1}{\beta}\int_d^h \left[\frac{\beta}{r+\beta m} + \frac{1}{\nu(m)f(\nu(m))}\right]\left(\frac{1}{p}-\frac{1}{m}\right)dm. \tag{A.11.4}$$

Hence the episode beats capped availability if and only if (A.11.3) *holds and* (A.11.4) *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 value $a = \nu(p)$. Substitution of the forced clock gives exactly (A.11.3); 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 (A.11.4) 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 (A.11.2). 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 values. 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 A.11.1 (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. $\square$

Theorem A.11.1 (Capped availability is uniquely optimal under a nonincreasing density). If the density is positive and nonincreasing on the relevant valuation range, every nondegenerate exact excursion has strictly positive excess duration. Under Theorem A.10.1, the unique optimal aggregate type-pure schedule is the path implemented by capped availability: the initial jump to the first boundary contact, smooth ridge passage to the second, and the terminal stockout.

Proof. Because

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

$f' \leq 0$ makes ν convex. In addition,

$$\beta\Gamma_F = \int L(m)h(m)dm = \int q'(m)\psi(m)dm.$$

Since $q(m_d) = q(m_r)$, $\int q' = 0$. Subtracting the continuous value $\psi(p)$ gives

$$\beta\Gamma_F = \int q'(m)\{\psi(m) - \psi(p)\}dm.$$

On the left of p, both factors are positive; on the right, both are negative. Their product is positive on sets of positive measure for every nondegenerate excursion. Hence $\Gamma_F > 0$.

Theorem A.10.1 gives

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

Therefore $T \geq \mathcal{S}(m_0, m_1)$, with equality only if $t_0 = 0$ and the excursion family is empty. Lemma A.11.2 shows that the schedule in (A.11.1) is feasible and attains this bound.

It remains to characterize equality. With no jumps and no singular-continuous component, the schedule is continuous and absolutely continuous. A nondegenerate interior flat interval is impossible: approaching its two endpoints through points carrying continuous increase, (A.10.3) would assign the same service belief to two distinct ranks, contradicting strict monotonicity of b. Hence the support of continuous increase is dense in $[m_0, m_1]$, and continuity extends (A.10.3) to the whole interval. Differentiation then gives (A.10.4) almost everywhere, which uniquely determines (A.11.1). The only remaining flat portions are the boundary batches established in Theorem A.10.1. Thus the aggregate completed graph is unique, apart from zero-mass cutoff conventions and relabeling of one same-date service line. $\square$

Remark 3 (Scope of type-pure uniqueness).

Theorem A.11.1 solves the continuous-value policy problem as specified: the planner chooses an ordered type-pure schedule. In the finite model, lotteries arise from positive applicant masses and positive releases. With atomless values, an exact type has zero mass, so a release directed only to it is not represented by cumulative supply and the anonymous-rationing rule has no applicant mass on which to operate. A type-by-type lottery requires an added allocation kernel. Relaxed joint value-service measures can arise as limits of finite optima, but need not be type-pure and define a separate problem.

A.11.2 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 A.11.3 (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*

$$\ell(h) = h - \frac{q_p'''(p)}{3q_p''(p)}h^2 + O(h^3). \tag{A.11.5}$$

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 A.11.2 (Cubic excursion expansion)**.** *Let*

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

For the reflected excursion of Lemma A.11.3,

$$\Gamma_F(p-h,p,p+\ell(h)) = K_F(p;r,\beta)h^3 + O(h^4), \tag{A.11.6}$$

where

$$K_F(p;r,\beta) = \frac{(\beta p+2r)y(p)^2 - \beta p y'(p) + \dfrac{2\beta r}{\beta p + r}y(p)}{3\beta^2 p^3}. \tag{A.11.7}$$

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

$$K_F(p;r,\beta) = \frac{2rf - \beta p v f' + \dfrac{2\beta r}{\beta p + r} v f^2}{3\beta^2 p^3 v^2 f^3}. \tag{A.11.8}$$

In particular,

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

Proof. Write $b_j = b^{(j)}(p)$ and $b_0 = b(p) = ra/(r+\beta p)$. Expand the exact integral (A.10.13) over $[p-h, p+\ell(h)]$, substitute (A.11.5), and collect cubic terms. The constant, linear, and quadratic differences cancel by same-state contact and reflection. Direct simplification gives

$$K_F = \frac{(\beta p + r)\{2a\beta b_1 + a\beta p b_2 + 2(\beta p + r)b_1^2\}}{3a^2\beta^2 p^3 r}. \tag{A.11.9}$$

Now substitute

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

Equation (A.11.9) reduces to (A.11.7). Since

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

we have

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

Substitution into (A.11.7) gives (A.11.8). □

Lemma A.11.4 (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*

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

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

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. Write, for z near zero,

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

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 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

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

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

Proposition A.11.2 (Uniform local sign test)**.** *Let* $h_0 := \min\{\bar{h}, \eta/(2M)\}$.

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

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

and therefore beats capped availability.

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

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

and is slower than capped availability.

Proof. Combine Theorem A.11.2 and Lemma A.11.4. □

Lemma A.11.5 (Slack-based insertion)**.** *Let a globally compatible ridge-contact family leave a smooth interval containing the support of a faster exact excursion* e*, and let* $\delta = -\Gamma_F(e) > 0$ *be the time saved. Suppose every type assigned before the departure state has slack at least* $\sigma > 0$ *against every post-reentry continuation alternative. If* $\bar{U}$ *bounds the absolute value of those later dated payoffs and*

$$(e^{r\delta} - 1)\bar{U} < \sigma,$$

then replacing the smooth interval by e *preserves global incentive compatibility.*

Proof. The excursion returns to the same public state and advances every post-reentry date by δ. Comparisons among post-reentry alternatives are unchanged because every corresponding payoff line is multiplied by the same positive factor $e^{r\delta}$. Comparisons of later assigned types against earlier alternatives are relaxed. The only potentially tighter constraints are those of types assigned before departure against post-reentry alternatives. Each such payoff rises by at most $(e^{r\delta} - 1)\bar{U}$, which is smaller than the maintained slack. □

Proposition A.11.1 shows that a negative-gain exact excursion is automatically a profitable deviation from capped availability. The lemma addresses a different question: whether the same excursion can replace components of an arbitrary globally compatible architecture. At an interior boundary of such an architecture, upstream waiting constraints may bind, so selecting a particular wave inside an optimum can require an explicit replacement condition.

Definition A.11.1 (Replacement condition)**.** Fix a globally compatible excursion e. For any globally compatible family $\mathcal{A}$ omitting e, let $\mathcal{C}_e(\mathcal{A}) \subseteq \mathcal{A}$ be a subfamily whose supports conflict with e. The excursion satisfies a replacement condition if, whenever $\mathcal{A}$ is optimal and omits e,

(i) $(\mathcal{A} \setminus \mathcal{C}_e(\mathcal{A})) \cup \{e\}$ is globally compatible; and

(ii)
$$\Gamma_F(e) \leq \sum_{a \in \mathcal{C}_e(\mathcal{A})} \Gamma_F(a).$$

Proposition A.11.3 (Conflict-dominance selection)**.** *If an excursion e satisfies a replacement condition, some optimal globally compatible ridge-contact family contains e.*

Proof. Take an optimal family $\mathcal{A}$. If it contains e, there is nothing to prove. Otherwise replace $\mathcal{C}_e(\mathcal{A})$ by e. Compatibility is preserved by the condition, and additivity of gains gives a weakly lower total duration. The replacement is therefore also optimal and contains e. □

A useful primitive sufficient case is supplied by Lemma A.11.5: a candidate wave whose support lies after a region of uniformly slack waiting constraints is globally insertable. Another is a wave at the beginning of the active region, where no upstream waiting constraints exist.

For a compact certified core C, define

$$W(C) = \inf\{\Gamma_F(e) : e \text{ is globally compatible and } C \subseteq I(e)\},$$

and let

$$H(C) = \inf\left\{\sum_{a \in \mathcal{A}} \Gamma_F(a) : \mathcal{A} \text{ is globally compatible, treats } C, \text{ and no } a \in \mathcal{A} \text{ covers } C\right\}.$$

The phrase "treats C" means that the family has the same prescribed outside components and replaces the smooth benchmark on the relevant core; this keeps the comparison within a common endpoint and compatibility problem.

Corollary A.11.2 (Whole-core selection)**.** *Suppose the outside components can be held fixed when replacing the treatment of* C*. If* $W(C)$ *is attained and* $W(C) \leq H(C)$*, some optimal globally compatible family contains a single excursion whose support covers* C*.*

Proof. Take an optimal treatment of the core. If it already contains a covering excursion, the claim holds. Otherwise its gain is at least $H(C)$. Replacing its treatment of C by an attaining covering excursion changes the gain to $W(C) \leq H(C)$ and preserves compatibility by hypothesis. The resulting family is optimal and covers C with one excursion. ☐

Theorem A.11.3 (Primitive concentration and conditional global batching)**.** *Let* C *be a compact interior core satisfying the regular reflected-contact hypotheses and* $K_F \leq -\eta < 0$*. Then every point of* C *supports sufficiently small negative-gain reflected excursions. If one such excursion has a replacement condition, some optimum contains a wave intersecting* C*. If, in addition, the whole-core hypotheses of Corollary A.11.2 hold, some optimum contains a wave whose support covers* C*.*

Proof. The local claim follows from Proposition A.11.2. The global claims follow from Proposition A.11.3 and Corollary A.11.2. ☐

The one-sandwich benchmark and the remaining auxiliary constructions are collected in the Online Supplement (Sections A.12-A.16).

B Global excursion selection, welfare incidence, and optimal one-wave design

This appendix completes the global continuum argument and then derives the quantitative one-wave conditions. Throughout, I maintain the notation of the continuum sections of the main text. In particular,

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

and $0 < m_0 < m_1 < \bar{M}$ denote the compulsory boundary contacts in the generic interior case. Boundary-degenerate cases are obtained by letting the initial contact or active interval collapse, with terminal service imposed by the target-screening condition.

For a prospective plateau rank $p \in [m_0, m_1]$, put

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

and define

$$q_p(m) = k \log b(m) + \log\{a - b(m)\}. \tag{1}$$

The exact crossing calculation in Appendix A shows that q_p has a strict maximum at p, is strictly increasing on its relevant left contact branch, and is strictly decreasing on its relevant right contact branch. An exact excursion therefore satisfies

$$q_p(d) = q_p(h), \tag{2}$$

which is equivalent to

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

Define the compact contact set

$$\overline{\mathcal{E}}_F = \{(d,p,h) : m_0 \le d \le p \le h \le m_1,\ q_p(d) = q_p(h)\}. \tag{3}$$

For $e = (d,p,h) \in \overline{\mathcal{E}}_F$, let

$$\Delta(e) = \frac{x(h) - x(d)}{\beta p} = \frac{1}{\beta p} \log \frac{b(d)}{b(h)}, \tag{4}$$

$$\Gamma_F(e) = \frac{1}{\beta} \int_d^h x'(m) \left(\frac{1}{p} - \frac{1}{m} \right) dm, \tag{5}$$

$$\ell(e) = h - d. \tag{6}$$

Diagonal triples are null excursions and have zero length and gain. Two excursions are compatible if the interiors of their support intervals $[d,h]$ are disjoint. Let $\mathfrak{A}_f$ and $\mathfrak{A}_c$ denote the finite and finite-or-countable compatible families, respectively.

B.1 Finite families and the interval-selection value

The exact decomposition in Appendix A gives the forward direction: every feasible ordered type-pure schedule produces a compatible family and the completion-time identity. The next lemma supplies the converse needed for optimization over families.

Lemma B.1 (Chronological concatenation)**.** *Every finite compatible family* $\mathcal{A} \in \mathfrak{A}_f$ *has a unique chronological concatenation with smooth ridge passage on the complement of its supports. The resulting schedule is globally incentive compatible and has completion time*

$$T(\mathcal{A}) = T^{\mathrm{cap}} + \sum_{e \in \mathcal{A}} \Gamma_F(e). \tag{7}$$

No additional inherited feasibility restriction is required beyond exact contact and compatibility.

Proof. Order the excursions by departure rank and fill every uncovered interval by ridge passage. State matching is exact: an excursion departing at d begins at belief $b(d)$ and, by (4), returns at belief $b(h)$. Hence the next smooth or pooled component inherits exactly the same public state it would inherit on the ridge.

For incentive compatibility, represent every dated opportunity by its affine payoff line

$$L(v) = e^{-rt}(v - B).$$

Along a smooth ridge segment, neighboring lines are tangent to the indirect-utility envelope in value order. An exact excursion replaces the ridge lines on $[d, h]$ by its departure and reentry lines. Equation (2) makes these lines cross at $\nu(p)$; their slope order assigns values above $\nu(p)$ to departure and values below $\nu(p)$ to reentry. At d and h, the retained lines meet the adjacent ridge envelopes. Thus the ordered list of envelope intersections is unchanged outside the replaced support. Chronological concatenation of disjoint supports preserves this ordered-crossing property, which implies every nonadjacent incentive inequality as well as the adjacent ones. The terminal line retains the lower-tail exclusion proved in Appendix A.

Smooth time is additive. Replacing ridge passage over $[d, h]$ by excursion e changes duration by exactly $\Gamma_F(e)$. Summing over disjoint supports gives (7). □

B.2 Anonymous staircase-supply implementation

The direct schedule generated by a finite family can be decentralized by the original anonymous application game. The construction uses cumulative inventory caps rather than type-specific invitations.

Proposition B.2 (Finite staircase implementation)**.** *Let* $\mathcal{A} = \{e_k = (d_k, p_k, h_k)\}_{k=1}^{K}$ *be a finite compatible family in increasing mass order, and let* $\tau_{R,k}$ *denote the reentry date of excursion* k *in its chronological concatenation. There is a right-continuous nondecreasing cumulative-supply path with at most* $K + 1$ *positive jumps whose selected equilibrium reproduces exactly the concatenated service schedule, adoption path, beliefs, and completion time. No intended applicant is rationed on the equilibrium path.*

If $K = 0$, cumulative supply is raised from zero to $\bar{M}$ at the first service date. If $K \geq 1$, after merging any coincident release dates, the path is

$$S_t = \begin{cases} 0, & t < t_0, \\ p_1, & t_0 \leq t < \tau_{R,1}, \\ p_{k+1}, & \tau_{R,k} \leq t < \tau_{R,k+1}, \quad k = 1, \ldots, K-1, \\ \bar{M}, & t \geq \tau_{R,K}, \end{cases} \tag{8}$$

where t_0 is the first service date. Conversely, any no-rationing equilibrium of a staircase path in which almost every value receives one ordered certain service date induces a globally incentive-compatible ordered type-pure schedule.

Proof. Consider first $K \geq 1$. Under the first cap p_1, the initial pool and the smooth ridge segment consume inventory until adoption reaches d_1. At the departure date, the intended applicant mass is exactly $p_1 - d_1$, equal to the unclaimed inventory. Every applicant is therefore accepted and the cap is exhausted: $M = p_1 = S$. During the following wait, the cap binds, so every application is rejected and adoption remains at p_1.

At $\tau_{R,1}$, cumulative supply rises to p_2 when $K > 1$ and to $\bar{M}$ when $K = 1$. The reentry applicant mass $h_1 - p_1$ is no larger than the new inventory. It is served with probability one. The unclaimed balance is, respectively, $p_2 - h_1$ or $\bar{M} - h_1$; it remains accessible while adoption follows the next smooth ridge segment. If another excursion follows, the balance is exhausted exactly by its departure pool $[d_2, p_2]$. Repeating the argument implements every excursion. After the last reentry, the final cap $\bar{M}$ supports the remaining smooth passage and is exhausted by terminal service. The case $K = 0$ is the same ledger with one cap: capped availability implements the excursion-free schedule. Excursions touching at one date merely combine consecutive cap increases and adoption pools.

It remains to verify incentives in the application game. At every date with unclaimed inventory, an application succeeds with probability one. Such dates are exactly the service dates represented by the direct schedule: pool dates and dates on smooth separating components. At dates at which the current cap is exhausted, an application fails with probability one and leaves eligibility unchanged. Hence any contingent application strategy is payoff-equivalent to choosing the date of its first successful application, possibly preceded by payoff-irrelevant rejected attempts. Global direct incentive compatibility makes the assigned successful date optimal among all service dates. Principal-preferred equilibrium selection resolves any weak equality in favor of the intended ordered application pattern. This proves exact implementation and no on-path rationing.

For the converse, take a no-rationing equilibrium with ordered certain service dates. Applying at another date with slack capacity would yield certain service at that date, while applying at a binding cap would have no effect. Sequential optimality therefore implies that every value weakly prefers its equilibrium service date to every other successful service date. These are precisely the global direct incentive inequalities, and the resulting rank-to-date map is ordered and type-pure. □

The construction clarifies the relation between adoption pools and supply releases. For one excursion (d, p, h), the principal initially releases p units and later raises cumulative supply to $\bar{M}$. The adoption jumps are $p - d$ at stockout and $h - p$ at reopening; they are not the two supply increments.

Consequently, the finite-family problem is already a weighted interval-selection problem:

$$T_f^* = T^{\text{cap}} + \inf_{\mathcal{A} \in \mathfrak{A}_f} \sum_{e \in \mathcal{A}} \Gamma_F(e). \tag{9}$$

A positive-gain excursion can always be deleted, so minimization may be restricted to nonpositive gains.

Let $J(x)$ denote the infimum of total excursion gain available from ridge rank x onward. Smoothly skipping to a later departure has zero incremental gain. Therefore

$$J(x) = \min \left\{ 0, \inf_{\substack{e=(d,p,h) \in \overline{\mathcal{E}}_F \\ d \geq x}} \left[\Gamma_F(e) + J(h) \right] \right\}. \tag{10}$$

This is an equality of infima. A countable optimizer may have excursions accumulating immediately to the right of x, so the equation need not possess a first minimizing action at every state.

B.3 Countable attainment and finite compression

The local contact expansion implies a cubic bound at null excursions. The C^4 condition imposed in the main theorem makes the estimate uniform over the compact active interval.

Lemma B.3 (Uniform cubic domination)**.** *There exist $C < \infty$ and $\delta_0 > 0$ such that every feasible excursion with $\ell(e) \leq \delta_0$ satisfies*

$$|\Gamma_F(e)| \leq C \ell(e)^3. \tag{11}$$

Moreover, if $q = p - d$ is its departure mass, then

$$h - d = 2q + O(q^2) \tag{12}$$

uniformly on the compact active interval.

Proof. Expand the contact equation around the diagonal $d = p = h$. The two contact branches meet at the strict maximum of q_p, and equality of their quadratic terms gives (12). Substituting the reflected endpoint into (5) cancels the first- and second-order terms; the first possible nonzero term is cubic. Uniform C^4 bounds and compactness give common constants C and δ_0. □

Put $M = m_1 - m_0$. Compatibility implies

$$\sum_{e \in \mathcal{A}} \ell(e) \leq M$$

for every finite or countable family. Hence deleting all excursions shorter than $\delta \le \delta_0$ changes total gain by at most

$$C \sum_{\ell(e)<\delta} \ell(e)^3 \le CM\delta^2, \tag{13}$$

and the retained family contains at most M/δ excursions.

Theorem B.4 (Countable attainment and finite compression)**.** *The finite-family infimum is attained in the countable compatible closure:*

$$\inf_{\mathcal{A}\in\mathfrak{A}_f} \sum_{e\in\mathcal{A}} \Gamma_F(e) = \min_{\mathcal{A}\in\mathfrak{A}_c} \sum_{e\in\mathcal{A}} \Gamma_F(e). \tag{14}$$

The gain series of every countable compatible family is absolutely convergent. If $\mathcal{A}^*$ is an attained optimum, then for every $\varepsilon > 0$ it has a finite subfamily whose induced rollout is within ε of the optimum. One may retain at most

$$N(\varepsilon) \le \left\lceil \frac{M}{\min\{\delta_0, \sqrt{\varepsilon/(CM)}\}} \right\rceil \tag{15}$$

economically relevant excursions.

Proof. Absolute convergence follows by separating the finitely many excursions with span above δ_0 from the remaining excursions and applying (11):

$$\sum_{\ell(e)\le\delta_0} |\Gamma_F(e)| \le C\delta_0^2 \sum_e \ell(e) < \infty.$$

For attainment, compactify $\overline{\mathcal{E}}_F$ by its null triples and associate with a finite family the length-weighted measure

$$\mu_{\mathcal{A}} = \sum_{e\in\mathcal{A}} \ell(e)\delta_e.$$

Its total mass is at most M. Take a minimizing sequence of finite families and a weakly convergent subsequence of these measures. Define

$$g(e) = \begin{cases} \Gamma_F(e)/\ell(e), & \ell(e) > 0, \\ 0, & \ell(e) = 0. \end{cases}$$

Lemma B.3 implies that g is continuous at null controls and hence on the compact contact set. Therefore

$$\int g\, d\mu_{\mathcal{A}} = \sum_{e\in\mathcal{A}} \Gamma_F(e)$$

passes to the weak limit.

On every set $\{\ell \ge \delta\}$, compatibility bounds the number of atoms by M/δ. The weak limit is therefore finitely atomic there. Taking the union over $\delta = 1/n$ yields a countable nondegenerate atomic part. Overlapping interiors are excluded because compatibility is closed under convergence.

Any residual diffuse mass is supported on null controls, where $g = 0$. The nondegenerate atoms thus form a countable compatible family attaining the limiting value.

Enumerate its excursions so that nested finite subfamilies exhaust their total support length. Fill omitted supports by smooth ridge passage. Equation (13) gives convergence of completion times and the cardinality bound (15). The finite schedules are uniformly Cauchy in service dates because the smooth time density is bounded on the compact active interval and omitted total support length tends to zero. Their limit is the countable concatenation. Continuity of dated payoff lines passes every finite global incentive inequality to the limit. By Proposition B.2, each finite schedule is exactly induced by a finite anonymous staircase supply policy. The countable optimizer is therefore sequentially implementable as the uniform limit of finite staircase equilibrium paths, including at accumulation points. ☐

Corollary B.5 (Operational value equivalence)**.** *The minimum completion time in the ordered type-pure continuum problem equals the infimum over its finite anonymous staircase implementations. For every $\varepsilon > 0$, a finite staircase policy attains completion time within ε of the global optimum. If an optimal excursion family is finite, the optimum is attained by a finite staircase policy.*

Proof. Lemma B.1 and Proposition B.2 map every finite compatible family to an exact finite staircase implementation. The converse part of Proposition B.2 maps every ordered no-rationing staircase equilibrium back to a feasible direct schedule. Theorem B.4 equates the finite-family infimum with the attained countable value and provides the ε-approximating finite families. ☐

Combining Lemma B.1, Proposition B.2, Theorem B.4, and the decomposition in Appendix A proves the global characterization and operational implementation statements in the main text.

B.4 Welfare incidence of a pooled excursion

Fix a nondegenerate exact excursion $e = (d, p, h)$ and normalize its departure date to zero. The free-availability continuation from the same public state $(d, b(d))$ assigns rank $m \geq d$ its ridge date

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

and payoff

$$U_F(m) = e^{-rt_F(m)}\{\nu(m) - b(m)\}. \tag{16}$$

The excursion assigns the payoff

$$U_e(m) = \begin{cases} \nu(m) - b(d), & d \leq m \leq p, \\ e^{-r\Delta}\{\nu(m) - b(h)\}, & p \leq m \leq h, \\ e^{-r[\Delta + \mathcal{S}(h,m)]}\{\nu(m) - b(m)\}, & m \geq h, \end{cases} \tag{17}$$

where $\Delta = \Delta(d, p, h)$. The first two expressions agree at p by the contact equation.

Proposition B.6 (Welfare incidence)**.** *For every nondegenerate exact excursion:*

(i) $U_e(d) = U_F(d)$ *and* $U_e(m) < U_F(m)$ *for every* $m \in (d, p]$*.*

(ii) For every $m \geq h$*,*

$$U_e(m) = e^{-r\Gamma_F(e)} U_F(m). \tag{18}$$

(iii) If $\Gamma_F(e) < 0$*, there is a unique* $\widehat{m} \in (p, h)$ *such that* $U_e < U_F$ *on* $[p, \widehat{m})$*,* $U_e(\widehat{m}) = U_F(\widehat{m})$*, and* $U_e > U_F$ *on* $(\widehat{m}, h]$*.*

Proof. For a fixed value v, let

$$R(z; v) = e^{-rt_F(z)}\{v - b(z)\}$$

be the payoff from selecting ridge rank z. Since $t'_F(z) = -b'(z)/(\beta z b(z))$ and the ridge identity gives

$$r\{\nu(z) - b(z)\} = \beta z b(z),$$

differentiation yields

$$\frac{\partial R(z; v)}{\partial z} = e^{-rt_F(z)} b'(z) \frac{r\{v - \nu(z)\}}{\beta z b(z)}. \tag{19}$$

Because $b' < 0$ and ν is strictly decreasing, a type $v = \nu(m)$ strictly prefers its own ridge rank m to every $z \neq m$. The departure payoff in (17) is the ridge line at $z = d$. This proves part (i), including equality at d. At p, the contact equation makes the departure and reentry lines equal, so the same strict inequality applies to either convention for the marginal rank.

For $m \geq h$, additivity of smooth time gives

$$t_F(m) = \mathcal{S}(d, h) + \mathcal{S}(h, m).$$

The excursion serves the same rank at the same ridge belief $b(m)$ after time $\Delta + \mathcal{S}(h, m)$. Since

$$\Gamma_F(e) = \Delta - \mathcal{S}(d, h),$$

equation (18) follows.

It remains to compare the reentry line with the free envelope on $[p, h]$. Put

$$H(m) = e^{-r\Delta}\{\nu(m) - b(h)\} - U_F(m).$$

Part (i) and the contact equation imply $H(p) < 0$. If $\Gamma_F(e) < 0$, part (ii) implies $H(h) > 0$. The envelope theorem applied to (16) gives

$$U'_F(m) = e^{-rt_F(m)} \nu'(m),$$

so

$$H'(m) = \nu'(m) \left[e^{-r\Delta} - e^{-rt_F(m)} \right]. \tag{20}$$

The function t_F is strictly increasing. Hence H is strictly increasing while $t_F(m) < \Delta$ and strictly decreasing while $t_F(m) > \Delta$. Because $H(p) < 0 < H(h)$, necessarily $t_F(p) < \Delta < t_F(h)$. Thus H crosses zero exactly once on its increasing branch and remains positive through h, proving part (iii). ☐

Corollary B.7 (Incidence along an optimal family)**.** *There is an optimal compatible family in which every retained nondegenerate excursion has strictly negative gain. For each such excursion, relative to free ridge passage from its inherited departure state, the departure pool loses, every rank weakly after reentry gains, and the reentry pool has the unique cutoff described in Proposition B.6.*

Proof. A positive-gain excursion can be deleted to lower completion time, and a zero-gain excursion can be deleted without changing it. Deletion preserves compatibility and the inherited ridge states of all remaining excursions. Apply Proposition B.6 to each retained excursion. ☐

The proposition separates the welfare incidence of speed from the principal's completion-time objective. The departure pool bears the cost of accelerated experimentation. If the wave is profitable in time, sufficiently late members of the reentry pool and every subsequent adopter obtain the same informational terms earlier.

B.5 Boundary contacts and existence of an optimal one-wave schedule

Proposition B.8 (Boundary extension and one-wave attainment)**.** *Suppose F has a positive continuously differentiable density on the active interval.*

(a) The set $\overline{\mathcal{E}}_F$ is nonempty and compact, and Γ_F is continuous on it.

(b) Every nondegenerate element of $\overline{\mathcal{E}}_F$ satisfies $d < p < h$. It implements a globally incentive-compatible one-excursion schedule with departure mass $p - d$, pause (4), and reentry mass $h - p$. The same statement holds when $d = m_0$ or $h = m_1$, with the corresponding block merged into the initial or terminal service line.

(c) The minimum one-wave completion time is

$$T^{\text{cap}} + \min_{\overline{\mathcal{E}}_F} \Gamma_F,$$

and the minimum is attained.

Proof. The map $(d, p, h) \mapsto q_p(d) - q_p(h)$ is continuous on the compact ordered cube. Its zero set is compact and is nonempty because every diagonal triple belongs to it. Continuity of (5) follows from continuity of x' and $p \geq m_0 > 0$.

Because q_p has a strict maximum at p, contact with $d \leq p \leq h$ implies either $d = p = h$ or $d < p < h$. The line-crossing argument in Lemma B.1 gives global incentive compatibility. When $d = m_0$ or $h = m_1$, the relevant excursion line coincides with the corresponding boundary line, so the block simply enlarges that boundary service mass. The duration identity then gives the displayed minimum, which is attained by compactness. ☐

B.6 First-order conditions

Let

$$C(d,p,h) = q_p(d) - q_p(h).$$

For a nondegenerate contact triple write $B_d = b(d)$ and $B_h = b(h)$ and define

$$R_d = \frac{B_d}{a - B_d} - k, \qquad R_h = k - \frac{B_h}{a - B_h}, \qquad D = \frac{1}{a - B_d} - \frac{1}{a - B_h}. \tag{21}$$

The strict increase and decrease of q_p imply $R_d, R_h > 0$. Since b is decreasing, $B_d > B_h$, so $D > 0$.

Proposition B.9 (Interior stationarity)**.** *Let (d,p,h) be an interior local minimizer of Γ_F subject to $C = 0$. Then*

$$\frac{1}{\beta}\frac{1/d - 1/p}{R_d} = \frac{1}{\beta}\frac{1/p - 1/h}{R_h} = \frac{\Delta(d,p,h)}{r\Delta(d,p,h) + pD/f(a)}. \tag{22}$$

At a boundary minimizer, the equality associated with the active boundary is replaced by the corresponding one-sided Kuhn–Tucker inequality.

Proof. Differentiate (5). The endpoint and plateau derivatives are

$$\Gamma_d = \frac{x'(d)}{\beta}\left(\frac{1}{d} - \frac{1}{p}\right), \tag{23}$$

$$\Gamma_h = \frac{x'(h)}{\beta}\left(\frac{1}{p} - \frac{1}{h}\right), \tag{24}$$

$$\Gamma_p = -\frac{1}{\beta p^2}\int_d^h x'(m)dm = -\frac{\Delta(d,p,h)}{p}. \tag{25}$$

For the contact constraint,

$$\frac{\partial q_p(m)}{\partial m} = b'(m)\left\{\frac{k}{b(m)} - \frac{1}{a - b(m)}\right\}.$$

Using $b' = -x'b$ and (21),

$$C_d = x'(d)R_d, \qquad C_h = x'(h)R_h. \tag{26}$$

Since $k_p = -r/(\beta p^2)$ and $a_p = -1/f(a)$,

$$\begin{aligned} C_p &= -\frac{r}{\beta p^2}\{\log b(d) - \log b(h)\} - \frac{1}{f(a)}\left\{\frac{1}{a - B_d} - \frac{1}{a - B_h}\right\} \\ &= -\frac{r\Delta(d,p,h)}{p} - \frac{D}{f(a)}. \end{aligned} \tag{27}$$

The gradient of C is nonzero, so the Lagrange multiplier theorem applies. Substitution of (23)–(27) into the stationarity equations yields (22). Standard one-sided differentiation gives the boundary inequalities. □

The first two ratios in (22) determine the relative expansion of the departure and reentry blocks. The final ratio places the wave in the value distribution, accounting for the effects of p on both the marginal value and the rate at which the plateau generates information.

B.7 Local geometry of small waves

The reflected-contact map is singular in the original level coordinate because its two branches meet at the maximum of q_p. A direct Taylor expansion nevertheless yields a useful policy approximation.

Lemma B.10 (Reflected endpoint and pause)**.** *Suppose* q_p *is* C^4 *near* p*,* $q_p'(p) = 0$*, and* $q_p''(p) < 0$*. Let* $d = p - s$ *and let* $h = p + \ell_p(s)$ *be the right-side reflected contact satisfying* $q_p(p - s) = q_p(p + \ell_p(s))$*. Then*

$$\ell_p(s) = s + A(p)s^2 + O(s^3), \qquad A(p) = -\frac{q_p'''(p)}{3q_p''(p)}. \tag{28}$$

Moreover,

$$\Delta(p - s, p, p + \ell_p(s)) = \frac{2x'(p)}{\beta p}s + \frac{x'(p)A(p)}{\beta p}s^2 + O(s^3). \tag{29}$$

Proof. Write $q_j = q_p^{(j)}(p)$. Taylor expansion on the left gives

$$q_p(p - s) = q_p(p) + \frac{q_2}{2}s^2 - \frac{q_3}{6}s^3 + O(s^4).$$

Seek $\ell_p(s) = s + As^2 + O(s^3)$. Expansion on the right gives

$$q_p(p + \ell_p(s)) = q_p(p) + \frac{q_2}{2}s^2 + \left(q_2 A + \frac{q_3}{6}\right) s^3 + O(s^4).$$

Equating cubic terms yields $A = -q_3/(3q_2)$. Finally,

$$\Delta = \frac{x(p + \ell_p(s)) - x(p - s)}{\beta p},$$

and Taylor expansion of x gives (29). □

Remark B.11 (Policy interpretation). To first order, a marginal wave uses two equal-sized batches. Its pause is linear in their common scale, is shorter when learning β is stronger, and is shorter when the plateau stock p is larger. The second-order term $A(p)$ determines whether the reentry block is slightly larger or smaller than the departure block.