Staggered Rollout for Innovation Adoption

Ricardo Fonseca*

October 10, 2026

Abstract

A farmer considering a new perennial crop, or someone considering a new vaccine, may prefer to wait and learn from others' experience. I show that publicly committing to a temporary pause in access can accelerate adoption: some agents adopt earlier, and their experience informs those who wait. I characterize the access calendar that reaches an adoption target as soon as possible, for general smooth distributions of good-state payoffs. A pause must last long enough to keep early adopters from waiting until reopening. Each pause is compared with the gradual adoption it replaces. These time effects add across nonoverlapping pauses, so the optimal calendar selects the pauses that together save the most time. Continuous access until a final round is optimal when the payoff density is nonincreasing, whereas sufficiently strong local concentration can make pauses beneficial. For a fixed adoption population, faster completion need not raise welfare: a pause induces some agents to bear the risks of earlier experimentation and makes others wait longer. Heterogeneity in prior beliefs is equivalent to payoff heterogeneity, allowing the principal to use data on beliefs or gains to choose access dates.

Keywords: innovation adoption, social learning, dynamic games, strategic delay.
JEL codes: C73, D82, D83, O33.

*Pontificia Universidad Javeriana, Bogotá, Colombia. Email: baricardo@javeriana.edu.co. I am especially grateful to Bobak Pakzad-Hurson for his guidance and continued support. I also thank Jack Fanning, Teddy Mekonnen, Kareen Rozen, Roberto Serrano, and audiences at Brown University, LACEA–LAMES, EWMES, SBE, USP, Universidad de Chile, Universidad de los Andes, and the Econometric Society Microeconomic Theory Workshop for helpful comments. All remaining errors are mine.

1 Introduction

Consider a program that helps farmers establish a new perennial crop. A farmer may find adoption worthwhile today and still prefer to wait: fields planted by others can reveal whether the crop is ill suited to local conditions. The program wants the crop to be widely grown locally and sets an uptake target. It cannot pay farmers to experiment early or prevent them from observing others' experience, but it can decide when support is available. Vaccination programs face a similar issue: people can benefit from a new vaccine yet prefer to wait for evidence from earlier recipients. Could announcing a temporary pause in access induce earlier adoption and improve learning enough to reach the target sooner?

A pause changes the alternatives available to potential adopters. Someone who would have adopted during the pause must instead adopt before closure or wait until access resumes. Those with more to gain are more willing to act early, and their experience keeps generating evidence during the pause. Continuous access initially has fewer adopters producing information, but keeps adding new adopters during the same period. A pause helps when the information gained from bringing adoption forward outweighs the adoption and learning forgone during closure. The distribution of gains determines both sides of this comparison.

I study a principal who commits to an access calendar to reach an adoption target as soon as possible. Agents differ in their private payoff from a good innovation; a bad innovation yields a common loss. Experience after adoption can reveal bad quality. Until that happens, the absence of bad news makes agents more optimistic and encourages further adoption. Each agent decides whether and when to adopt, retaining access to all later opportunities if she declines an earlier one. The principal therefore has to induce early adoption through the opportunities it promises later.

I characterize the optimal calendar for general smooth distributions of good-state payoffs: whether to pause access, when, and for how long. Earlier reopening reduces the delay for those who wait, but weakens the incentive to adopt before closure. To induce a given group to adopt early, the principal must keep access closed long enough that the marginal adopter is indifferent between adopting before closure and waiting until reopening.

The characterization evaluates a pause by comparing it with gradual adoption between the same initial and final states. Both paths reach the same stock of adopters and the same public belief, leaving the remaining agents with the same adoption problem. Reaching this endpoint sooner advances completion by the same amount; reaching it later delays completion. The effects of successive pauses add. Choosing the optimal calendar amounts to choosing pauses that do not overlap and save the most time in total, a choice that can be computed by a dynamic program. Calendars with finitely many pauses and unique

equilibria come arbitrarily close to the optimal completion time, so the principal need not rely on favorable equilibrium selection.

Figure 1 contrasts uniform payoffs with concentration. In the top row, good-state payoffs are uniform (panel A). A pause can draw some agents into early adoption, but reopening must be delayed enough to keep them from waiting. During this delay, continuous access keeps adding adopters who produce information. With equal numbers of agents in equal payoff intervals, this later expansion outweighs the early group's advantage. Every pause delays completion, so the optimal calendar offers continuous access until a final round (panel B). The same result holds with a nonincreasing density: adoption passes into payoff intervals containing at least as many agents, making the expansion prevented by closure still harder to compensate for.

The bottom row instead has a concentration of payoffs in the middle of the distribution (panel C). As adoption passes through the concentration, a pause can bring a large group forward while delaying relatively few agents with lower payoffs. Panel D shows a pause whose early informational gain outweighs the learning forgone during closure. Reopening is also the final round, and the target is reached sooner than under continuous access.

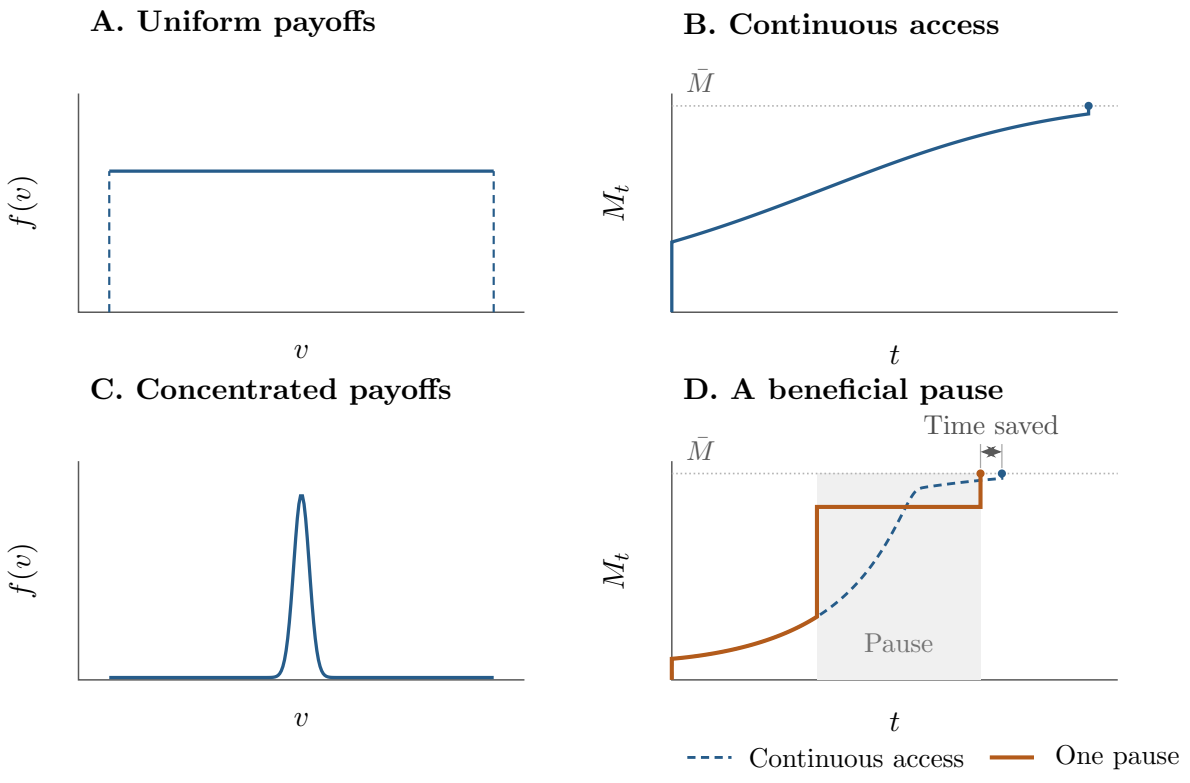


Figure 1: Payoff distributions and adoption calendars. Here $f(v)$ is the good-state payoff density, M_t is the stock of adopters, and $\bar{M}$ is the target. Panel D compares a feasible pause with continuous access until a final round. Appendix G.1 specifies the examples.

For a small pause, the comparison has a simple expression. Along the continuous-access path, let v_t be the good-state payoff of the marginal adopter before the final round, f the payoff density, and r the discount rate. Define

$$\Lambda_t = -\frac{d}{dt} \log f(v_t).$$

As learning brings lower-payoff agents into adoption, Λ_t measures how quickly the density they encounter falls. A small pause brings completion forward where $\Lambda_t > 2r$. For uniform payoffs, $\Lambda_t = 0$. On the lower-payoff side of a sufficiently strong concentration, the density falls fast enough for the early informational gain to dominate. Learning enters through the speed at which v_t falls; discounting enters through the incentive to adopt at closure rather than wait until reopening. The condition identifies where small pauses help; the characterization determines their length and how to combine them.

Holding the final population of adopters fixed, faster completion need not make them better off. Agents drawn forward to closure lose intermediate dates they would have preferred, while earlier information can benefit agents who adopt later. Either force can dominate: a pause that accelerates adoption can increase welfare, and a pause that saves still more time can reduce it.

The same calendar design applies when agents differ in their prior beliefs about quality rather than in their gains from success. People may value protection similarly yet differ in their confidence that a new vaccine is safe and effective. A change of variables maps more optimistic agents into higher-payoff types and preserves adoption choices under every calendar. A principal can therefore use an estimated distribution of beliefs or gains, together with the learning and discount rates and the target, to choose access dates. A numerical illustration follows this calculation from initial beliefs to closing and reopening dates.

Farmers learn from others' experience (Foster and Rosenzweig, 1995; Conley and Udry, 2010), and adoption patterns can reflect strategic waiting (Bandiera and Rasul, 2006). An email reporting scarcity and a reduced clinic schedule increased attendance at the next clinic during the 2004 influenza-vaccine shortage (de Janvry et al., 2010). In the experiment of Cason et al. (2025), reducing the probability of a second opportunity raises early adoption among intermediate-cost participants but lowers total adoption. An announced reopening preserves a later opportunity to adopt while inducing some agents to experiment before closure.

Related literature. The incentive to wait for others' experience connects this paper to work on irreversible investment and social learning. In Décamps and Mariotti (2004), private investment costs determine which of two firms moves first, and that investment improves the public information available to the follower. Here, too, private types concern

adoption returns rather than information about quality. With a continuum of agents, the principal must induce an entire early cohort to change the public information path. Frick and Ishii (2024) study informational free-riding among forward-looking consumers whose adoption opportunities arrive stochastically. Their analysis explains adoption patterns and the limits of improving the learning technology. I introduce a principal who designs these opportunities to reach an adoption target sooner.

Past adopters continue to generate information in Laiho, Murto, and Salmi (2025), as they do here. Their Brownian-learning model compares decentralized adoption with a planner who directly chooses adoption dates to maximize welfare. Here the principal chooses an access calendar to reach an adoption target sooner, while agents decide whether and when to adopt. Laiho, Murto, and Salmi (2026) show how positive payoff externalities can reinforce informational free-riding, so that greater learning potential slows learning in their extremal equilibria. Here the incentive to wait arises from information alone.

The closest policy comparison is with Knoepfle and Salmi (2025), who choose when to disclose hard evidence generated by investment. A promise of useful information later can undermine the early investment needed to produce it in both papers. Under pure bad news, transparency maximizes the cumulative probability of disclosed evidence at every date. Censoring news can induce faster investment, but disclosing the additional evidence earlier would make some early investors prefer to wait. With both good and bad news, an optimal policy discloses bad evidence immediately and delays favorable evidence, accelerating experimentation and the generation of bad evidence. The principal here keeps all bad news public and instead restricts adoption dates. Existing adopters keep producing evidence during a closure; in their model, further evidence requires additional investment. Their optimal disclosure policy is a Pareto improvement over transparency, whereas reaching an adoption target sooner can lower adopter welfare by removing valuable intermediate dates.

Recommendation policies also influence who experiments. Che and Hörner (2018) choose what information to send to consumers who decide upon arrival. Lyu (2026) shows that recommendations may be temporarily suspended following negative feedback. With long-lived, heterogeneous adopters, Chen et al. (2025) address the incentive to wait directly: their optimal feedback policy can permanently recommend against adoption even when the principal remains relatively optimistic. In these papers, adoption generates information that the principal decides whether and how to disclose. Here evidence remains public, and the principal commits to closing and reopening access.

Commitment to future terms has a similar role in Laiho and Salmi (2026). In their durable-good monopoly, buyers can wait for both information and lower prices. Commitment to higher future prices limits informational free-riding and eliminates the sales delays that can arise without commitment. Here the principal commits to access dates, with fixed adoption terms and guaranteed service. Suspending sales stops information

generation in their model, whereas earlier adopters continue producing information during a closure here.

Section 2 introduces the model. Section 3 characterizes the optimal access calendar and explains its computation. Section 4 derives the theoretical implications for payoff distributions and welfare. Section 5 develops the practical implications, relating calendar design to survey inputs and existing evidence. Section 6 discusses extensions with alternative objectives, transfers, and observable types. Section 7 concludes.

2 Model

2.1 The innovation and access

A principal offers an innovation to a nonatomic population of mass one. It can choose the dates on which the innovation is available, but each agent decides whether and when to adopt. The principal wants to reach an uptake target $\bar{M}$ as soon as possible. For example, a vaccination program may have committed to reaching a given coverage, while a supplier may receive a reward when sales reach a specified level. The objective concerns the time needed to reach that target; Section 4.3 examines the consequences for adopter welfare.

The innovation has unknown quality $\omega \in \{g, b\}$. The principal and agents share the prior $\mu_0 = \mathbb{P}(\omega = g) \in (0, 1)$, and agents discount at rate $r > 0$. Adoption is irreversible and yields a payoff v in the good state and -1 in the bad state. An agent who never adopts receives zero.[1] The good-state payoff v is privately known and has distribution F, with full support on $V = [\underline{v}, \bar{v}] \subset (0, \infty)$ and a strictly positive, continuously differentiable density. The payoff type v is independent of ω, and agents receive no private signals about quality. Heterogeneity in values has an equivalent interpretation as heterogeneity in initial subjective beliefs, developed in Section 2.3. Each agent knows her payoff before the calendar is announced; the principal knows its population distribution but observes neither individual payoffs nor quality. The distribution and all other primitives are common knowledge.

At date zero, the principal commits to a closed set of available dates $D \subseteq \mathbb{R}_+$. Every agent who has not yet adopted can do so at any date $t \in D$, on the same fixed terms. An agent who declines one opportunity remains eligible at every subsequent date. A calendar can provide continuous access, a sequence of isolated rounds, or continuous access interrupted by pauses. Access is guaranteed at an available date: the principal does not choose how many of those who wish to adopt will be served. There are no transfers or capacity constraints.

[1] The zero payoff normalizes the value of an existing alternative, such as another crop or remaining unvaccinated. The prior incorporates information available before the program begins.

These assumptions distinguish a temporary closure from a threat that an agent will lose her place in a queue. Under the policies considered here, waiting until reopening is a choice the agent knows she can make. Any adoption induced before closure must therefore be preferable to that later opportunity. A vaccination program would need to guarantee service on the announced dates while withholding it during the pause. In the crop example, the corresponding commitment concerns comparable technical support. The role of this commitment is discussed in Section 5.

After adoption, experience with the innovation produces public evidence about quality. Conditional on bad quality, a conclusive signal arrives at rate β times cumulative adoption, where $\beta > 0$; it never arrives under good quality.[2] All agents observe the calendar, aggregate adoption, and the arrival or absence of this public evidence. There are no other aggregate shocks or private signals. Write σ for the first signal date. Once bad quality is revealed, every agent who has not adopted strictly prefers never to do so. News takes precedence over adoption at the same date, although a tie has probability zero.

The information available to a waiting agent depends on the experience of all previous adopters. In particular, closing access does not stop the arrival of evidence. Earlier vaccine recipients can develop adverse reactions while new vaccinations are suspended; fields already planted can reveal poor crop performance while support is unavailable. Conversely, the continued absence of such evidence makes potential adopters more optimistic. A larger stock of adopters makes silence more informative because a bad innovation would be more likely to have produced a signal. This is why the adoption induced at closure can be useful during the pause. Each agent takes this information process as given, since her individual action has no effect on aggregate adoption.

2.2 Adoption choices and the completion target

A pure strategy is a jointly Borel-measurable stopping rule $\theta(v, \sigma) \in D \cup \{\infty\}$, adapted to the public signal history. An agent can decline any number of opportunities and adopt later. Since adopting after bad news is strictly dominated, every undominated pure rule can be represented by a measurable intended date $\tau(v) \in D \cup \{\infty\}$: adopt on that date if no news has arrived, and never adopt otherwise. The resulting adoption path before news is deterministic in equilibrium. Observing who adopts gives no further information about quality, because the private payoffs contain no information about the state. Appendix A.1 proves this stopping-rule representation.

[2]The conclusive signal idealizes evidence sufficient to end further use. In 1999, evidence linking RotaShield to an increased risk of bowel obstruction led to the withdrawal of its U.S. vaccination recommendation (Centers for Disease Control and Prevention, 1999).

Write cumulative adoption on the no-news path as

$$M_t = \int_V \mathbf{1}\{\tau(v) \le t\}\, dF(v) \tag{1}$$

Conditional on bad quality, the probability that no signal has arrived through t is $\exp\{-\beta \int_0^t M_s\, ds\}$.

Bayes' rule therefore gives the bad-state odds

$$B_t := \frac{1-\mu_t}{\mu_t} = B_0 \exp\left\{-\beta \int_0^t M_s\, ds\right\}, \qquad B_0 = \frac{1-\mu_0}{\mu_0}. \tag{2}$$

Beliefs depend on the amount of experience accumulated, measured by $\int_0^t M_s\, ds$. An adoption pool changes the rate of learning but does not cause a jump in beliefs. For example, if cumulative adoption remains fixed at p for a pause of length Δ, bad-state odds are multiplied by $e^{-\beta p \Delta}$. A larger early pool allows the same change in beliefs to occur over a shorter pause.

Consider an agent who plans to adopt at t unless news arrives first. In the good state, she adopts for certain and receives v. In the bad state, she incurs the loss only if the public signal has not arrived by t. Her time-zero expected payoff is consequently

$$e^{-rt}\left\{\mu_0 v - (1-\mu_0)\exp\left(-\beta \int_0^t M_s\, ds\right)\right\} = \mu_0 e^{-rt}(v - B_t).$$

Dropping the common positive factor μ_0, define

$$L_t(v) = e^{-rt}(v - B_t), \qquad L_\infty(v) = 0. \tag{3}$$

An available date offers an affine payoff in v. Earlier adoption discounts the good-state payoff less, whereas later adoption reduces the probability of investing in a bad innovation. The choice depends on the entire future calendar, since an agent may prefer to skip several opportunities.

A *pure no-news equilibrium* for D is a measurable τ with (M, B) given by (1)–(2), such that, for F-almost every v,

$$\tau(v) \in \arg\max_{t \in D \cup \{\infty\}} L_t(v). \tag{4}$$

The maximum exists because D is closed, dated payoffs are continuous, and $L_t(v) \to 0$ as $t \to \infty$. This is the pure Nash outcome concept for the adoption game. Aggregate observations are deterministic on the no-news path and unaffected by a unilateral deviation, so the stopping-rule representation retains every individual adoption choice. Throughout,

equilibrium refers to this concept, and uniqueness concerns the aggregate outcome; indifferent cutoff types and unused calendar dates do not affect it.

At any no-news date u, an agent who has not adopted chooses among the remaining dates $D \cap [u, \infty)$ and nonadoption. Her conditional payoff from $s \geq u$ is $\mu_u e^{ru} L_s(v)$. The positive factor does not depend on s, so she maximizes the same payoff over the remaining opportunities, including after declining her intended date. An individual's deviation leaves aggregate adoption and beliefs unchanged.

Higher-payoff agents adopt weakly earlier. To see why, take two available dates $s < t$. The payoff difference $L_s(v) - L_t(v)$ is strictly increasing in v. If a lower type prefers s, a higher type cannot prefer t. Moreover, a higher type earns strictly more from any date at which a lower type is willing to adopt, so adopters form an upper interval of the payoff distribution. Write $\nu(m) = F^{-1}(1 - m)$, so that a mass m of agents has payoffs at least $\nu(m)$, and $\tau(m) = \tau(\nu(m))$ for the adoption date of an agent with payoff $\nu(m)$. I choose a right-continuous version at null cutoffs. A flat part of τ is a pool of agents adopting together.

For a calendar and an equilibrium, the completion time is

$$T_{\bar{M}}(D, \tau) = \inf\{t : M_t \geq \bar{M}\}, \qquad \inf \varnothing = \infty.$$

I maintain the interior-prior assumption

$$0 < 1 - F(B_0) < \bar{M} < 1. \tag{5}$$

Equivalently, $\nu(\bar{M}) < B_0 < \nu(0)$. Some agents find immediate adoption worthwhile, so experimentation can start. But even if access ended immediately, the agents willing to adopt would not be enough to reach the target. Further evidence is needed to bring in the remaining adopters.

The principal chooses a calendar to minimize completion time. If it can select a favorable equilibrium, its value is

$$V^{\mathrm{C}}(\bar{M}) = \inf\{T_{\bar{M}}(D, \tau) : (D, \tau, M, B) \text{ is a pure no-news equilibrium}\}. \tag{6}$$

If it instead evaluates each calendar by its worst equilibrium, the value is

$$V^{\mathrm{wc}}(\bar{M}) = \inf_{D} \sup\{T_{\bar{M}}(D, \tau) : \tau \text{ is a pure no-news equilibrium of } D\}, \tag{7}$$

where the supremum is infinite if there is no pure equilibrium or an equilibrium does not complete. The second criterion requires the calendar to perform well even when agents coordinate on its least favorable equilibrium. Theorem 1 will show that the two values

coincide: the principal can approach its best completion time with calendars that induce a unique aggregate adoption path.

It suffices to consider outcomes in which access ends with adoption exactly at the target. Proposition A.1 in Appendix A establishes that allowing adoption to exceed the target cannot improve the minimum completion time. Simply ending access when adoption first reaches the target can enlarge the final pool. The proof constructs a calendar that stops exactly at the target no later than the original outcome, while preserving incentives to adopt at earlier dates. Write $V(\bar{M}) = V^{\mathrm{C}}(\bar{M})$. At the last opportunity of an outcome ending exactly at the target, the marginal target type must earn zero surplus. Otherwise, slightly lower types would also want to adopt. Terminal odds thus satisfy $B_T = \nu(\bar{M})$, regardless of the calendar used to get there.

This observation also gives the objective an ex ante interpretation. Since initial and terminal odds are fixed, every such outcome accumulates the same amount of evidence before completion:

$$e^{-\beta \int_0^T M_s \, ds} = \frac{\nu(\bar{M})}{B_0}.$$

The probability of completing before bad news is consequently

$$\mathbb{P}(\sigma > T) = \mu_0 + (1 - \mu_0)\frac{\nu(\bar{M})}{B_0} = \mu_0\{1 + \nu(\bar{M})\}. \tag{8}$$

Suppose the principal receives a nonnegative, strictly decreasing reward $R(T)$ upon completion before news. Its expected reward is the fixed probability in (8) times $R(T)$, so it chooses the calendar with the earliest completion. A discounted prize, $R(T) = e^{-\rho T}$, is one example. The target reduction both advances completion and weakly raises this probability when the original outcome exceeds the target, so the interpretation also applies to the objective of reaching at least $\bar{M}$. Completion can occur before bad quality is revealed; the principal's reward need not coincide with adopters' payoffs.

2.3 Heterogeneous beliefs

The model also has an equivalent interpretation in which agents value a successful innovation equally but differ in their initial optimism about its quality.[3] Suppose the good-state payoff is a common $g > 0$, and an agent initially assigns a privately known subjective probability $q \in (0, 1)$ to good quality. Agents share the conditional law of public evidence. Their different priors express different initial assessments, rather than private signals whose realization could reveal quality to others.

[3]Vaccine surveys document substantial heterogeneity in confidence about safety and effectiveness; see Larson et al. (2016).

The corresponding payoff type is

$$\frac{v}{B_0} = \frac{gq}{1-q}, \qquad q = \frac{v}{v + B_0 g}. \tag{9}$$

In this interpretation, B_0 is a reference scale: one may set $B_0 = 1$ and express all equivalent payoffs in those units. Changing this scale changes neither adoption choices nor the optimal dates. More optimistic agents correspond to higher-payoff types.

To see why their adoption choices coincide, let $S_t = \exp\{-\beta \int_0^t M_s \, ds\}$ be the probability of no news conditional on bad quality. Her time-zero expected payoff from adopting at t unless bad news arrives first is

$$e^{-rt}\{qg - (1-q)S_t\} = \frac{1-q}{B_0} e^{-rt}(v - B_t) = \frac{1-q}{B_0} L_t(v). \tag{10}$$

The positive factor is the same for every date. The agent makes the same choice between early adoption, waiting, and never adopting; the same holds after any history without news. Every payoff distribution F can thus be read as a distribution of initial assessments of quality. If G is the distribution of initial optimism, the two population distributions satisfy $F(v) = G(v/(v + B_0 g))$. Transforming the population by (9) preserves the aggregate adoption path under each calendar, its completion time, and the optimal calendar. Agents observe the same public evidence throughout, with initially more optimistic agents remaining more optimistic after silence.

The same argument allows gains and beliefs to vary together. For an agent with good-state gain $g_i > 0$ and prior q_i, define $v_i = B_0 g_i q_i/(1-q_i)$. Equation (10) then holds agent by agent. The distribution of these equivalent payoffs summarizes the joint heterogeneity relevant for adoption choices. The transformation rescales agents' utilities by different positive factors; welfare is evaluated in their original utility units in Section 4.3.

3 Optimal access calendars

To accelerate completion, a pause must induce enough early experimentation to compensate for the adoption postponed during closure. I characterize the optimal calendar by comparing pauses with continuous adoption between the same initial and final stocks of adopters and public beliefs. Matching these endpoints makes each pause's time gain or loss carry through to completion.

For the calendar comparisons below, access ends at a finite completion date with adoption exactly at the target, and the first used date is normalized to zero. Proposition A.1 justifies this restriction. I first derive adoption under continuous availability, then consider calendars that interrupt access before the final round.

3.1 Continuous availability and the final round

An agent may be willing to adopt at the current belief and still prefer to wait. A short delay costs discounted adoption surplus, but it also gives the existing adopters time to reveal bad quality. Under continuous access, the marginal agent must be indifferent between these choices; otherwise she would move to a neighboring date.

Differentiating $L_t(\nu(M_t))$ while holding her type fixed gives

$$-r\{\nu(M_t) - B_t\} + \beta M_t B_t = 0.$$

The first term is the cost of delay. The second is its informational benefit: bad-state odds fall as earlier adopters continue to produce evidence. A larger stock makes waiting more valuable, so it must be accompanied by better beliefs before the next agent is willing to adopt. Solving this indifference condition for the odds at stock m gives

$$b(m) = \frac{r\nu(m)}{r + \beta m}, \qquad x(m) = -\log b(m). \tag{11}$$

Since ν is positive and decreasing, b is strictly decreasing. Lower-payoff agents require better beliefs before adopting. Learning determines how quickly the economy moves along this path:

$$\frac{d[-\log B_t]}{dt} = \beta M_t.$$

Increasing adoption by dm requires $-\log B_t$ to rise by $x'(m)\,dm$. With m existing adopters, this takes $x'(m)\,dm/(\beta m)$ units of time. At a flatter part of b, a small improvement in beliefs induces a larger group to adopt. The same stock produces evidence at the same rate across distributions, but the evidence needed to bring in the next group differs.

The highest-payoff agents adopt immediately, forming an initial pool through m_0. A final deadline creates another pool for a different reason: once all later opportunities have been removed, agents need only prefer adoption to never adopting. As established in Section 2, ending exactly at the target requires the marginal target type to earn zero surplus. This fixes the terminal belief at

$$B_T = \nu(\bar{M}). \tag{12}$$

Define the initial and final points of gradual adoption by

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

Condition (5) implies $0 < m_0 < m_1 < \bar{M}$. If access ends when the continuous path reaches m_1, the last opportunity is at

$$T^{\mathrm{R}} = \frac{1}{\beta}\int_{m_0}^{m_1} \frac{x'(m)}{m}\,dm.$$

At this belief, agents with payoffs between $\nu(\bar{M})$ and $\nu(m_1)$ have nonnegative expected adoption payoffs but would prefer to wait under perpetual access. They join the terminal pool when access ends. Comparing temporary closures with perpetual access would attribute this deadline effect to the pauses. I therefore use continuous access until a final round as the benchmark.

Proposition 1 (Continuous availability)**.** *Perpetual availability $D = \mathbb{R}_+$ has a unique equilibrium. It begins with the pool m_0 and, until the population is exhausted, follows*

$$B_t = b(M_t), \qquad \dot{M}_t = -\frac{\beta M_t b(M_t)}{b'(M_t)} > 0. \tag{14}$$

The calendar $D^{\mathrm{R}} = [0, T^{\mathrm{R}}]$ has a unique equilibrium that follows the same path through m_1, then ends with the pool $[m_1, \bar{M}]$.

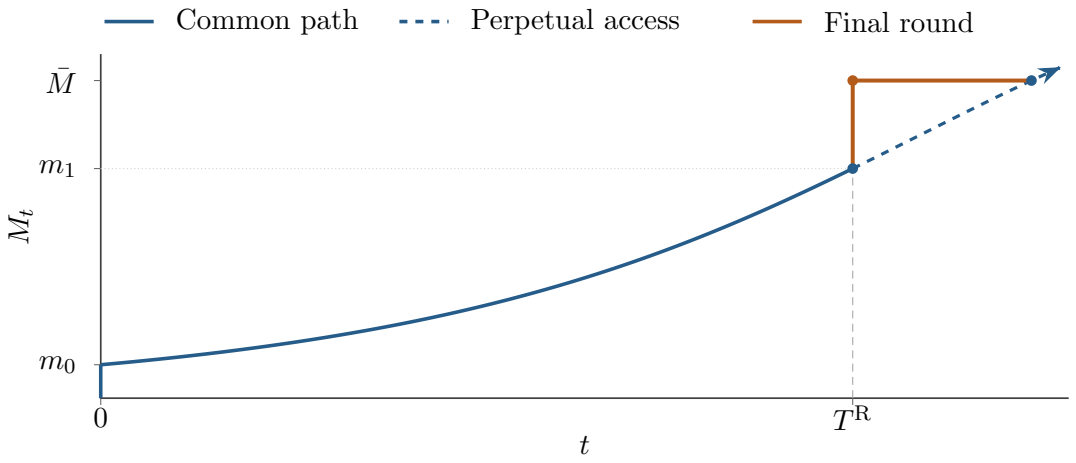


Figure 2: Continuous access and the final round. Both calendars follow the same path through m_1. Ending access at T^{R} brings adoption to $\bar{M}$. The example uses $\nu(m) = 4 - 3m$, $r = \beta = 1$, $m_0 = 0.1$, and $\bar{M} = 0.95$.

At positive dates before the final round in Figure 2, continuous access cannot support a pool: the additional adopters would raise learning immediately afterward, making a short delay more attractive just when the pooled agents are supposed to act. Nor can adoption stop for an interval while access remains open. The type separating the last and next adopters would prefer an intermediate date. These arguments, made precise in Appendix B, force gradual adoption onto $B = b(M)$.

The gradual path under perpetual access can be S-shaped even though agents are forward-looking and can adopt at any date. In the uniform example of Figure 2, adoption accelerates along the common path as a growing stock makes silence more informative. Along the perpetual-access continuation beyond m_1, the lower valuations of the remaining

agents eventually outweigh this faster learning, and adoption slows. The slowdown is modest over the continuation shown.[4]

3.2 One pause: incentives and completion time

A pause makes some agents adopt before closure rather than wait until reopening. The principal must keep access closed long enough to sustain that choice, even though a longer pause postpones adoption by those who wait.

For $m_0 \leq d < p < h \leq m_1$, agents with payoffs between $\nu(p)$ and $\nu(d)$ adopt at closure, and those with payoffs between $\nu(h)$ and $\nu(p)$ at reopening. Adoption jumps from d to p, remains there during closure, then jumps to h. The endpoint odds are $b(d)$ and $b(h)$. With stock p, learning occurs at rate βp, so the pause lasts

$$\Delta(d, p, h) = \frac{x(h) - x(d)}{\beta p}. \tag{15}$$

The agent with payoff $\nu(p)$ must be indifferent between the two dates:

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

Waiting allows her to avoid adoption if bad news arrives, but delays her good-state payoff. Combining indifference with $b(h) = b(d)e^{-\beta p\Delta}$ gives

$$b(d)^{r/(\beta p)}\{\nu(p) - b(d)\} = b(h)^{r/(\beta p)}\{\nu(p) - b(h)\}. \tag{16}$$

This condition determines which early and late groups can be induced by a pause. Holding the early pool fixed, slightly earlier reopening makes waiting preferable for its marginal adopter, so a shorter pause requires a different pool.

Compare the pause with continuous adoption between the same states, $(d, b(d))$ and $(h, b(h))$. The benchmark takes

$$\mathcal{S}(d, h) = \frac{1}{\beta}\int_d^h \frac{x'(m)}{m}\, dm. \tag{17}$$

[4]S-shaped adoption is widely documented in diffusion research (Rogers, 2003), and the shape of adoption can help distinguish different diffusion mechanisms (Young, 2009). Frick and Ishii (2024) also obtain S-shaped adoption under bad news with forward-looking agents and homogeneous valuations; in their endogenous bad-news setting, the concave phase disappears as the rate of Poisson adoption opportunities tends to infinity.

The difference in duration is

$$\begin{aligned}\Gamma_F(d,p,h) &= \Delta(d,p,h) - \mathcal{S}(d,h) \\ &= \frac{1}{\beta}\int_d^h x'(m)\left(\frac{1}{p}-\frac{1}{m}\right)dm. \end{aligned} \tag{18}$$

A negative Γ_F means that the pause saves time; a positive Γ_F means that it takes longer than continuous access. For $m < p$, the pause has more adopters producing evidence; for $m > p$, continuous access has more. Let M_t^R denote benchmark adoption measured from $(d, b(d))$. Reaching the same belief requires $p\Delta = \int_0^{\mathcal{S}(d,h)} M_t^R\, dt$, hence

$$p\Gamma_F(d,p,h) = \int_0^{\mathcal{S}(d,h)} [M_t^R - p]\, dt.$$

The pause is faster when p exceeds average benchmark adoption.

At reopening, adoption and beliefs agree under both paths. Using the same continuation shifts every later adoption date, including the final round, by Γ_F.

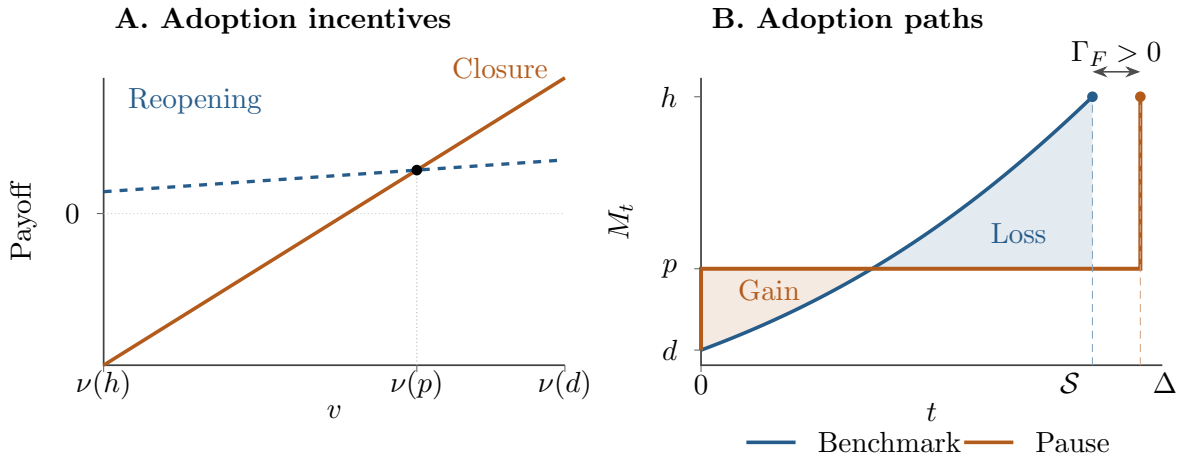


Figure 3: Incentives and learning during a pause. The payoff lines in A intersect at $\nu(p)$. In B, shading measures stock differences over the benchmark duration $\mathcal{S}(d,h)$; later loss minus early gain equals $p\Gamma_F$. The economy is as in Figure 2, with $d = 0.2$, $h = 0.6$, and p determined by indifference.

In panel B of Figure 3, the later loss is larger than the early gain, so the pause delays completion.

3.3 The general calendar problem

The principal must now choose which pauses to combine to reach the target fastest. Theorem 1 shows that their effects on completion time add across feasible, nonoverlapping pauses. The optimal calendar minimizes this sum, leaving access open between the chosen pauses.

Pauses are nonoverlapping when the interiors of their intervals $[d, h]$ are disjoint. Each then replaces a different passage under continuous access. Shared endpoints are allowed: one opportunity can reopen access after a pause and close it before another.

Let $\overline{\mathcal{E}}_F$ contain the feasible pauses of positive duration in $[m_0, m_1]$ and their zero-length limits $(d, p, h) = (m, m, m)$, for which $\Gamma_F = 0$. Write $\mathfrak{A}_f$ for finite families of nonoverlapping pauses.

Where adoption is continuous, the option to move to a neighboring date forces the marginal type onto $B = b(M)$. Where adoption dates are separated by a pause, the types on its two sides impose the indifference condition (16). Beliefs improve as one passage follows another, ordering the corresponding intervals $[d, h]$. Appendix D shows that these restrictions account for the whole schedule, including calendars without a finite list of pauses.

With several pauses, an agent can skip the next reopening and wait for a later round. The indifference condition for each pause must therefore support adoption choices across the whole calendar.

Theorem 1 (Calendar characterization)**.** *Every equilibrium outcome consists of initial and terminal pools, continuous adoption, and a finite or countable family $\mathcal{A}$ of nonoverlapping pauses. Its completion time is*

$$T = \mathcal{S}(m_0, m_1) + \sum_{e \in \mathcal{A}} \Gamma_F(e), \tag{19}$$

with an absolutely convergent sum. The optimal completion time is

$$V(\bar{M}) = \mathcal{S}(m_0, m_1) + \inf_{\mathcal{A} \in \mathfrak{A}_f} \sum_{e \in \mathcal{A}} \Gamma_F(e). \tag{20}$$

Every finite such family can be implemented with a unique equilibrium. An optimal equilibrium exists, and calendars with finitely many pauses and unique equilibria approach its completion time arbitrarily closely. Hence

$$V^{\mathrm{wc}}(\bar{M}) = V^{\mathrm{C}}(\bar{M}). \tag{21}$$

An agent who prefers closure to the next reopening also prefers it to later reopenings. Earlier payoff lines have larger slopes, and successive lines intersect at successively lower types. An agent on the early side of one intersection therefore prefers the earlier date at each subsequent comparison. Together with the payoff derivative during continuous access, this ordering verifies her choice against all announced dates.

Denote the bad-state odds at closure by D_0 and the pause duration by Δ. The stock P at closure in any equilibrium must satisfy

$$\nu(P) = \frac{D_0 - e^{-r\Delta} D_0 e^{-\beta P \Delta}}{1 - e^{-r\Delta}}.$$

As the closing stock increases, the marginal payoff falls, while more evidence during closure makes waiting more attractive. The left side therefore decreases in P and the right side increases, fixing the only closing stock consistent with the marginal agent's choice. Continuous access determines adoption between pauses. Appendix C applies this argument successively, including opportunities shared by two pauses.

3.4 Computing the calendar

I use Theorem 1 and Γ_F to compute the optimal calendar numerically. The pause that saves the most time on its own may displace several pauses that together save more time. The principal must therefore compare collections of nonoverlapping intervals $[d, h]$.

Given F, r, β, B_0, and $\bar{M}$, the benchmark determines m_0, m_1, and b. For each pair of endpoints $d < h$, indifference determines a unique closing stock $p(d, h) \in (d, h)$. Write $g(d, h)$ for $\Gamma_F(d, p(d, h), h)$, the same time cost expressed as a function of the pause's endpoints:

$$g(d, h) = \frac{x(h) - x(d)}{\beta p(d, h)} - \mathcal{S}(d, h), \qquad g(d, d) = 0. \tag{22}$$

Choosing (d, h) fixes both the cohort brought forward and the length of the pause. The cost is continuous in the endpoints, including when they coincide.

Restrict pause endpoints to a grid $\Pi = \{a_0, \ldots, a_n\}$ with $m_0 = a_0 < \cdots < a_n = m_1$. Let J_j be the lowest total cost of nonoverlapping pauses with endpoints among $a_0, \ldots, a_j$. Then

$$J_0 = 0, \qquad J_j = \min \left\{ J_{j-1},\ \min_{0 \le i < j} \left[J_i + g(a_i, a_j) \right] \right\} \quad (j = 1, \ldots, n) \tag{23}$$

On a grid with three points, the principal first compares a pause over $[a_0, a_1]$ with continuous access. At a_2, it compares keeping that choice, adding a pause over $[a_1, a_2]$, and replacing the earlier choice with a single pause over $[a_0, a_2]$. The longer pause therefore competes with two adjacent pauses, either pause alone, and no pause. An earlier choice affects when adoption reaches the last pause's initial state, but neither that state nor the pause's cost. After evaluating candidate costs, the recursion uses a quadratic number of comparisons.

Recording the minimizing choices recovers the intervals. Their indifference conditions and learning rates determine closing and reopening dates, with continuous access between

pauses and a final round at the target belief. Completion occurs at $T^{\mathrm{R}} + J_n$. The recursion chooses the number of pauses as well as their endpoints. As the grid becomes finer, its best completion time converges to the unrestricted optimum. Continuity of the pause cost and the finite-calendar approximation in Theorem 1 give this convergence; Appendix D supplies the argument.

4 Distributional implications

I use Theorem 1 to study how the payoff distribution F shapes the optimal calendar and how pauses affect adopter welfare. When $f'(v) > 0$, agents are more densely represented just above an adoption margin than just below it. A sufficiently steep increase in density can make the informational gain from earlier experimentation outweigh the adoption postponed during closure. Holding final adoption fixed, faster completion can benefit some adopters and hurt others; aggregate welfare can rise or fall.

Theorem 2 (Payoff distributions). *1. Continuous access until the final round is optimal if and only if*

$$\Gamma_F(e) \geq 0 \qquad \text{for every } e \in \overline{\mathcal{E}}_F. \tag{24}$$

Its aggregate outcome is uniquely optimal if the inequality is strict for every pause.

2. If $f'(v) \leq 0$ on $[\nu(m_1), \nu(m_0)]$, every pause has $\Gamma_F > 0$, and continuous access until the final round gives the unique optimal aggregate outcome.

3. Let $m \in (m_0, m_1)$ and $v = \nu(m)$, with ν of class C^4 near m. The condition

$$\frac{f'(v)}{f(v)} > \frac{2r}{\beta m v} + \frac{2r f(v)}{m(\beta m + r)} \tag{25}$$

produces arbitrarily small pauses around m with $\Gamma_F(e) < 0$.

Part 3 is local: a beneficial pause can exist even if the density decreases elsewhere. Combining it with continuous access before and after closure gives a calendar with a unique equilibrium.

4.1 Nonincreasing density

Suppose the density is nonincreasing. Moving from high to low payoffs, equal payoff intervals contain weakly more agents. Continuous access keeps bringing these agents into the stock as beliefs improve. A closure can advance some of them, but the principal must leave enough time before reopening to keep that early group from waiting. Over this interval, continuous adoption draws on the increasing mass of lower-payoff agents. The

early gain is always too small to compensate for the later loss, however the endpoints of the pause are chosen.

A nonincreasing density makes the upper quantile ν convex: the payoff reduction associated with an additional mass of adopters becomes weakly smaller at lower types. The proof in Appendix E uses this convexity and the separating type's indifference to pair the early gain with a strictly larger later loss. The uniform case shows that the conclusion does not rely on more agents at lower payoffs: the duration needed to sustain the early pool already makes the later loss exceed the early gain.

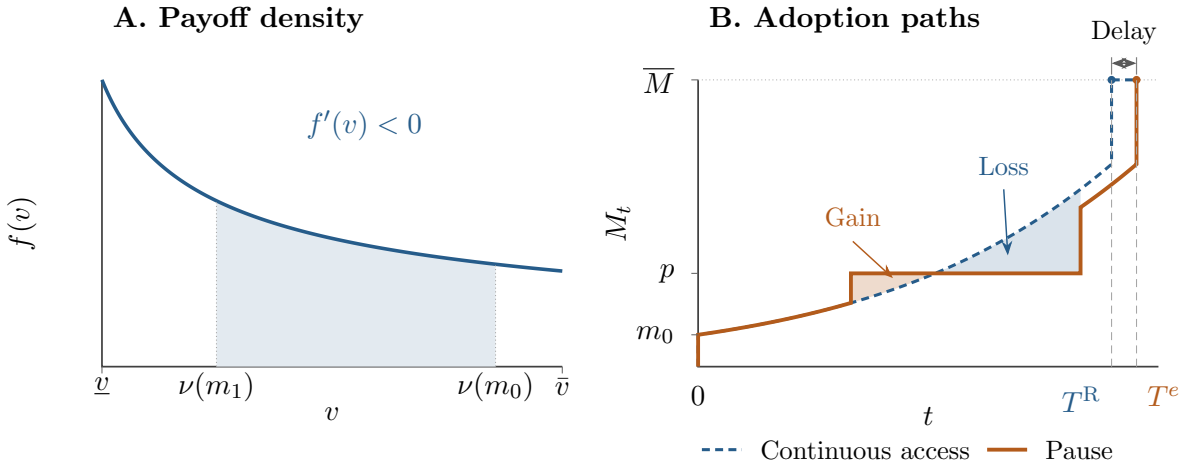


Figure 4: Decreasing density and a slower pause. Panel A shades the payoff interval $[\nu(m_1), \nu(m_0)]$ of agents who adopt gradually under continuous access. Panel B shades the early stock gain and later loss from a feasible pause in the same economy. Here $\nu(m) = 5 - 6m + 2m^2$, $r = \beta = 1$, $m_0 = 0.1$, $\bar{M} = 0.9$, $d = 0.2$, and $h = 0.5$.

For a common good-state gain g and initial probabilities with density φ, the transformation in Section 2.3 gives payoff density

$$f(v) = \frac{(1-q)^2}{B_0 g}\varphi(q), \qquad q = \frac{v}{v + B_0 g}.$$

When initial optimism is uniform on an interval inside $(0, 1)$, this density is strictly decreasing. A given change in probability corresponds to a larger change in the equivalent payoff at higher optimism, so equal probability intervals spread the most optimistic agents over wider payoff intervals. Continuous access until a final round is therefore optimal under both uniform payoffs and uniform optimism. More generally, a nonincreasing optimism density also induces a decreasing payoff density.

4.2 Concentration

The comparison can reverse when many agents lie on the higher-payoff side of an adoption margin. Drawing a group from this concentration into early adoption can produce a substantial stock before continuous access would have brought it into use. Yet those agents

have nearby payoffs and similar preferences over dates. A short interruption may not induce enough of them to act early, while a longer interruption sacrifices more expansion. The rise in density must be strong enough to overcome this constraint on the two pools.

For a small pause, closure and reopening must adjust together to keep the separating agent indifferent. This adjustment gives the threshold in (25). The condition identifies beneficial small pauses; finite pauses are evaluated directly by Γ_F.

Along the continuous-access path, with adoption stock $m = M_t^{\mathrm{R}}$, put $v = \nu(m)$. Since $\dot{m} = \beta m / x'(m)$ and $\nu'(m) = -1/f(v)$,

$$-\frac{d}{dt} \log f(\nu(M_t^{\mathrm{R}})) = \frac{\beta m f'(v)}{f(v)^2 x'(m)}. \tag{26}$$

This is Λ_t, the rate introduced in the illustration. Condition (25) is therefore $\Lambda_t > 2r$. As adoption proceeds, the marginal payoff falls along the fixed distribution. Where the condition holds, continuous adoption moves out of a densely populated payoff range toward a less densely populated one. The principal can benefit by bringing part of that nearby concentration forward, allowing its experience to accumulate before the benchmark would have made full use of it.

Holding the payoff margin fixed, an increase in β/r makes the information from early adoption arrive faster relative to the cost of waiting. A smaller rise in density then suffices for a profitable local pause.

Figure 5 illustrates the computation in Section 3.4 for a distribution concentrated at high payoffs. Comparing collections of pauses selects a calendar with one pause, which reopens at the final round.[5] At reopening, the final-round incentive draws in the terminal pool as well as the agents who waited through the pause.

4.3 Welfare

Faster completion can raise or lower adopter welfare. To compare calendars serving the same population, hold final adoption fixed at $\bar{M}$. Aggregate expected payoff under the common prior, omitting the common positive factor μ_0, is

$$\mathcal{W}(\tau) = \int_0^{\bar{M}} e^{-r\tau(m)} \{\nu(m) - B_{\tau(m)}\}\, dm. \tag{27}$$

Under heterogeneous priors, the corresponding welfare criterion weights each agent by $(1-q)/B_0$, as in (10). When different (g, q) pairs give the same equivalent payoff, their mean weight is used for that payoff type.

[5]The calendar is selected by grid searches followed by endpoint refinement. These searches provide no numerical bound on the gap to the unrestricted optimum.

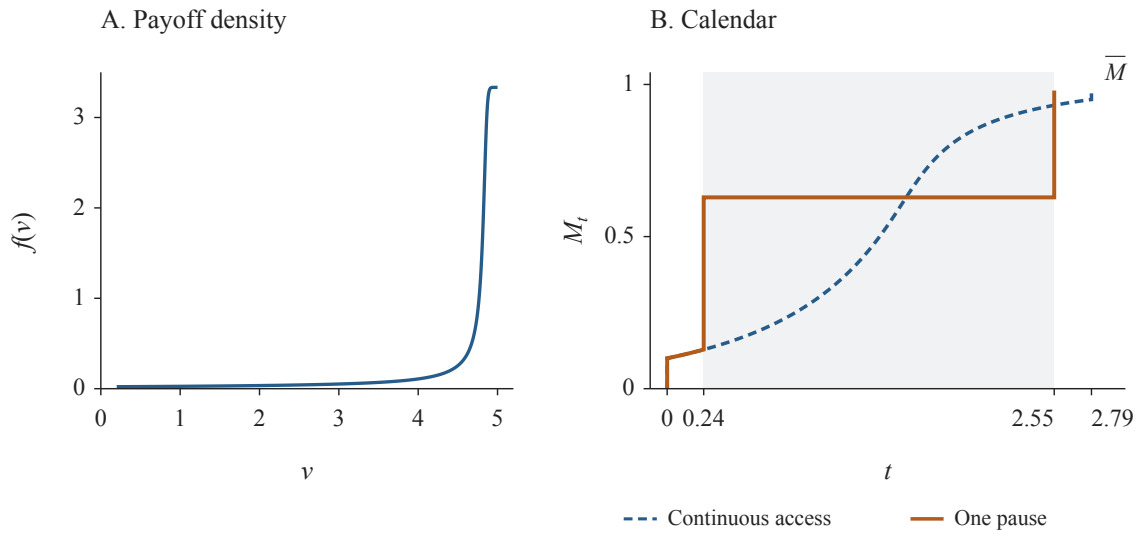


Figure 5: Payoff density and the computed calendar. Panel A shows the density corresponding to $\nu(m) = 5 - 0.3m - 4.5m^{10}$. Panel B compares the selected one-pause calendar with continuous access until a final round; shading marks the pause. Both calendars use $r = \beta = 1$, $m_0 = 0.1$, and $\bar{M} = 0.98$.

Figure 6 compares two feasible pauses for the payoff distribution underlying Figure 5, evaluated under the common prior. These curves concern the payoff interpretation of the distribution; welfare for the heterogeneous-belief interpretation developed in Section 5.1 requires the individual weights above. Both begin at the same closing date and speed completion. The closing cohort loses in both cases: these agents prefer closure to reopening, but would choose an intermediate date under continuous access. Some agents at reopening also lose; others gain from faster learning. Later adopters receive the same information sooner. In panel A, the gains outweigh the losses. The longer pause in panel B brings completion forward by more but lowers welfare.

If a pause takes at least as long as the continuous adoption it replaces, the removed options are not offset by an earlier continuation, and every such pause strictly lowers welfare. With a nonincreasing density, every pause is of this kind. Among outcomes ending with adoption exactly at the target, continuous access until the final round then uniquely maximizes both speed and adopter welfare.

The gains to later adopters depend on their surplus, including that of the final pool: the same advance in adoption dates is more valuable when their surplus is larger. Appendix E derives this comparison and gives smooth examples with each welfare sign.

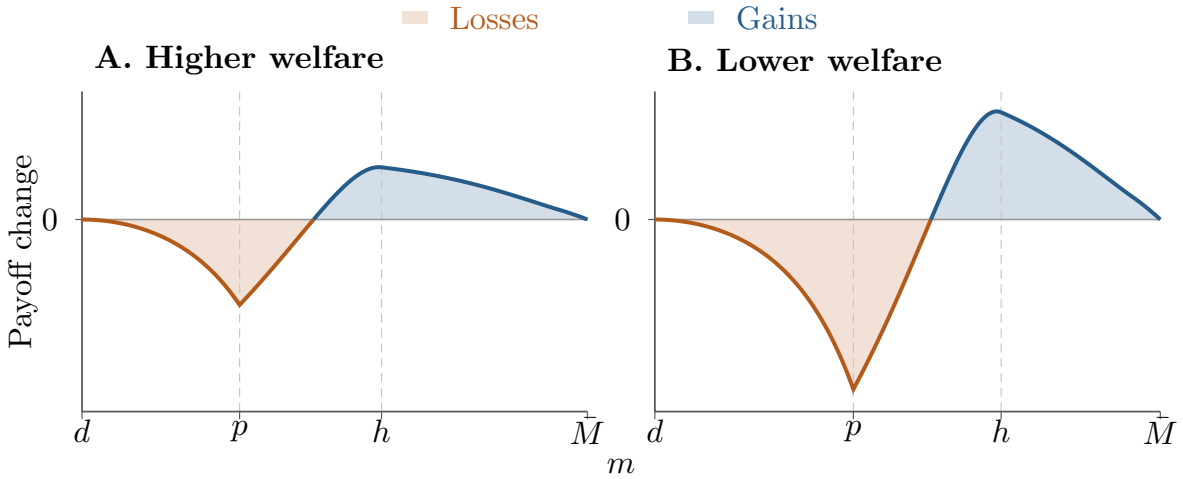


Figure 6: Welfare effects of two faster pauses for the payoff distribution of Figure 5. Both pauses begin from adoption $d = 0.44$; reopening brings adoption to $h = 0.76$ (A) and $h = 0.81$ (B). Curves show date-zero expected payoff changes relative to continuous access under the common prior, on identical axes. Since m measures population mass, signed area gives aggregate welfare change.

5 Practical implications

5.1 From beliefs and gains to a calendar

Before introducing a new vaccine, a program can ask prospective recipients what probability they assign to its meeting a specified standard of safety and effectiveness. Stated choices can provide estimates of their gains if the vaccine meets that standard. Together with the learning and discount rates and an adoption target, the resulting population distribution supplies the inputs for choosing closing and reopening dates. Here the learning rate measures how the experience of existing adopters produces public evidence of poor quality, as in Section 2.[6]

For a common good-state gain, apply (9) to each respondent's initial probability and use the resulting equivalent payoff distribution in the computation of Section 3.4. Numerical probability questions elicit the assessments represented by q; qualitative confidence categories, illustrated in Appendix G.2, require a cardinal mapping. When both beliefs and gains vary, the calculation uses their joint distribution across respondents, as in Section 2.3.

Figure 5 works through a hypothetical vaccination program. Suppose all agents gain one if the vaccine works and lose one if it fails. About three quarters assign success a probability between 0.50 and 0.53, while the remaining agents are more skeptical. Applying (9) gives the equivalent payoff density in panel A. Those beliefs correspond to payoffs

[6]There is an established literature on eliciting subjective probabilities from surveys; see Manski (2004). For identifying distributions of returns and willingness to pay from probabilistic stated choices, see Méango and Girsberger (2026). To use a valuation estimate as a good-state payoff, the survey must condition on the innovation working or specify the mapping from the measured valuation to that payoff.

from about 4.5 to the top of the support at 5, accounting for the concentration on the right; the left tail contains the more skeptical agents. This density is used to compute the calendar in panel B.[7] For a target of 98 percent, the numerical exercise selects access until 0.24 and a final reopening at 2.55. About 63 percent have adopted when access closes, and their experience supplies information during the pause. Continuous access until a final round would instead reach the target at 2.79.

5.2 Access restrictions and early experimentation

The relevant principal controls opportunities to adopt, while adopters continue to generate information after access closes. A public program might control grants or technical support for a new technology; a vaccination program might control clinic dates. The instrument is particularly useful when the principal can commit to dates but cannot readily pay agents to experiment early.

During the 2004 influenza-vaccine shortage, de Janvry et al. (2010) randomly sent emails announcing the shortage and the dates of the two remaining vaccination clinics. Receiving this email roughly doubled attendance at the next clinic.[8]

In the perennial-crop example, suppose support is available this planting season, suspended next season, and offered again in the following one. High-return farmers may plant now rather than give up a planting opportunity. Their fields continue to provide experience during the unsupported season, helping other farmers decide whether to plant when support returns. Evidence that farmers learn from neighbors' experience supports this channel (Conley and Udry, 2010; Bandiera and Rasul, 2006).

The experiment of Cason et al. (2025) provides evidence for the incentive effect of restricting later access. Three participants face different adoption costs, and early adopters produce public signals. In Figure 7, a lower probability of a second opportunity increases early adoption by middle-cost participants, while total adoption falls. In the experiment, access may not return. Here reopening is guaranteed, so early adoption must be induced by the delay until that later opportunity.

Once the early cohort has adopted, the principal may wish to reopen before the announced date. Anticipating this change would lead some agents to wait, weakening the initial adoption response. Support budgets and operating dates fixed in advance can help make closure and reopening credible. In agriculture, planting seasons also constrain the available dates. The calendar computation here allows arbitrary dates; an application to seasonal planting would require restricting the principal's choice to feasible seasons.

[7] The example uses $g = 1$, $r = \beta = 1$, and the reference scale $B_0 = b(0.1)$ for the payoff quantile in Figure 5. Its descending prior quantile is therefore $q(m) = \nu(m)/[\nu(m) + B_0]$. Dates are in the resulting normalized time units.

[8] The experiment combined scarcity and clinic-date information.

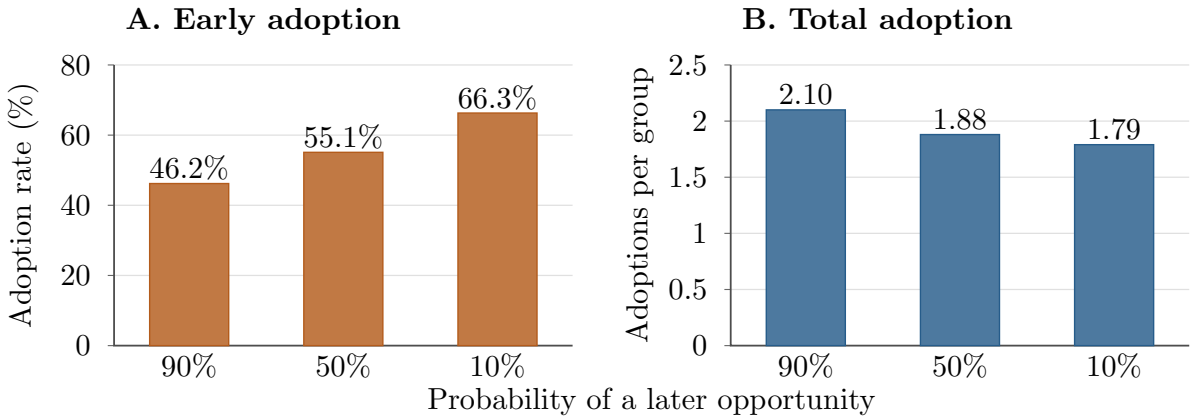


Figure 7: Early and total adoption in Cason et al. (2025). Panel A shows first-stage adoption by the middle-cost participant; panel B shows total adoption by all three group members across both stages, including games that end after stage one. Bars are reported means from Table 1 and Figure 4. The experiment has 144 participants in 12 sessions and varies the probability of a second opportunity.

6 Discussion

Different objectives. An uptake target can reflect a social benefit that individuals do not take into account. A vaccination program, for example, may seek coverage sufficient for herd immunity. Individuals can wait for evidence about vaccine safety while considering only their own protection. Reaching the target sooner can then be socially valuable even if some adopters are worse off. An access pause could improve expected social welfare if the expected gain from earlier collective protection exceeds the costs of experimentation and restricted access.[9]

A funding agreement or a reward tied to an uptake target can also make completion speed valuable to the principal even when adopters would prefer another calendar. A principal concerned with adopters' payoffs may also want to choose the adoption target. The welfare comparison in Section 4.3 holds the final adoption population fixed; with an endogenous target, additional adoption must be worth the risk it entails, given the evidence available when it occurs. Even with the target held fixed, choosing a welfare-maximizing calendar requires more than comparing completion times. Matching adoption and beliefs at reopening leaves a time difference for all later adopters, but it does not erase the gains and losses of those who adopted by reopening.

A supplier may instead value the timing of every sale. Two calendars with the same completion date can then differ because they distribute adoption differently before that date. A public program may incur a fixed cost whenever support resumes, giving it a reason to prefer fewer rounds even when additional pauses would accelerate completion.

[9]This interpretation keeps perceived private adoption gains fixed, as in the model. Allowing agents to adjust those gains as vaccination lowers their infection risk would introduce an additional payoff externality.

Transfers. If payments are unrestricted and the sole objective is completion speed, a sufficiently large initial subsidy induces at least the target population to adopt at date zero. By compensating agents for taking the risk, the principal no longer needs to wait for learning to make adoption attractive.

Immediate adoption need not maximize welfare. A subsidy can compensate an early vaccine recipient for bearing safety risk; it does not make the vaccine safer. If the entire target population adopts at date zero, none of its members can use others' experience to avoid a bad innovation. A principal concerned with adopters' payoffs may instead subsidize an early cohort and let the remaining agents wait for its experience. Choosing whom to subsidize and when would require comparing the informational benefit with the cost of experimentation and the social cost of funding the payments.

A payment raises the return to early adoption; a pause changes what an agent can do after declining it. A program with limited or costly financial support could therefore choose both its terms and its dates.

Observed payoffs. Would knowing which patients gain most from a medical innovation allow a principal to accelerate its adoption? Transcatheter aortic valve replacement (TAVR) initially became available to inoperable patients and later to patients at lower surgical risk (Centers for Medicare & Medicaid Services, 2019; U.S. Food and Drug Administration, 2019). Medicare coverage required registry follow-up, including mortality, stroke and quality of life (Centers for Medicare & Medicaid Services, 2019). The relevant gain is relative to the patient's treatment alternative, which surgical risk alone does not rank. Historical TAVR approvals responded to clinical evidence; the comparison below concerns eligibility dates committed in advance.

Knowing adoption gains allows the principal to invite high-return agents first. But an invitation that remains valid also allows its recipient to wait for others' experience. Identifying agents who are most willing to adopt does not by itself make them willing to experiment early.

For eligibility that expands from high-return to lower-return agents, with access retained after admission, observing payoffs does not improve the optimal completion time. An access calendar offered to everyone can match or improve the completion time of every equilibrium outcome under such a priority policy.[10] Adoption choices already put higher-return agents first. What the principal needs to change is their opportunity to wait, which a calendar can do without identifying each agent's payoff.

[10]Appendix F gives the argument for expanding eligibility with persistent invitations.

Population data can consequently inform the calendar even when invitations cannot depend on individual survey answers. The distribution determines the dates, and agents select themselves into adoption. Expiring personalized invitations would remove future opportunities for particular agents and give the principal an additional instrument. Whether that instrument allows faster completion remains open beyond the persistent priority policies considered here.

7 Conclusion

An announced closure changes adoption before it takes effect. Some agents move early because declining the current opportunity means waiting until reopening, and their experience continues to inform others during the pause. Evaluating a closure only by the adoption it prevents would miss this response. The principal must also honor its dates after the early cohort has adopted. Reopening sooner may then be attractive, but the expectation of doing so would weaken the incentive to experiment before closure.

An average gain or confidence level cannot reveal how large a cohort a closure can bring forward or how much adoption it would postpone. Data on the distribution of gains or beliefs can inform both sides of this comparison. The general characterization uses that distribution to choose closing and reopening dates, including the duration needed to sustain each pause. The dates depend on how quickly existing adopters produce evidence and how costly agents find waiting.

The principal's interest in speed need not coincide with adopters' interests, even when the final adoption population is unchanged. A closure shifts experimentation risk toward those who adopt early and removes dates they would otherwise prefer. Later adopters can benefit from earlier information, but a larger time saving need not deliver a larger welfare gain. An evaluation based only on uptake would miss both the cost of generating the evidence and who bears it.

A Strategies and adoption targets

A.1 Strategies and continuation choices

For a fixed D, a pure strategy is a jointly Borel-measurable stopping rule on $V \times [0, \infty]$, $\theta(v, \sigma) \in D \cup \{\infty\}$, adapted to the public signal history: for each v, the event $\{\theta(v, \sigma) \leq t\}$ depends only on news observed by t. The rule specifies adoption on every signal history, with ∞ meaning nonadoption. An individual's deviation can be any such rule, including one that declines any number of opportunities and adopts at a later one, and measurability in v makes the population aggregates well defined.

Stopping-date representation. Since adopting after news is strictly dominated, consider rules that stop all adoption after news. On the no-news history such a rule specifies a measurable intended date $\tau(v) \in D \cup \{\infty\}$ and takes the form

$$\theta_\tau(v, \sigma) = \begin{cases} \tau(v), & \tau(v) < \sigma, \\ \infty, & \sigma \leq \tau(v). \end{cases}$$

Every measurable τ induces an admissible stopping rule, and every undominated pure rule has this representation. To see the latter, evaluate the rule on a history where news never arrives. Adaptation requires the same adoption decision on any history where news arrives after the intended date, and if news arrives first, nonadoption is optimal. The representation therefore captures every undominated stopping rule. For completeness, an equilibrium outcome can be supplied with optimal continuations at every remaining-unadopted history. At any no-news date u, before that date's adoption decision, let the continuation select a date in

$$\arg \max_{s \in (D \cap [u, \infty)) \cup \{\infty\}} L_s(v). \tag{28}$$

The conditional expected payoff from the intended date $s \geq u$ is $\mu_u e^{ru} L_s(v)$, so this is the correct continuation problem. If $u \leq \tau(v)$, the assigned date remains a maximizer and is retained. If the agent is still unadopted at a subsequent date $u > \tau(v)$ after having skipped her assigned opportunity, use an optimal remaining date, with the earliest maximizer as a fixed tie-breaking rule. Repeating this rule as long as she remains unadopted is consistent with the chosen continuation; after bad news prescribe never to adopt, and after adoption no decision remains. These prescriptions cover counterfactual histories arising from an individual's own missed opportunities without altering the equilibrium aggregate path. Silence has positive probability at every finite date and news reveals bad quality, so beliefs at all these histories follow the stated updating rule. This supplies the sequentially rational continuation interpretation of the outcome definition; no additional

equilibrium-selection restriction is imposed. Throughout, uniqueness refers to aggregate adoption, disregarding F-null cutoff choices.

Adoption order. Suppose $v_H > v_L$ but $t_H := \tau(v_H) > t_L := \tau(v_L)$, and define

$$\Pi(v) = L_{t_L}(v) - L_{t_H}(v).$$

Because $t_L < t_H$,

$$\Pi'(v) = e^{-rt_L} - e^{-rt_H} > 0.$$

Optimality for v_L gives $\Pi(v_L) \geq 0$, and strict single crossing gives $\Pi(v_H) > 0$, contradicting optimality of t_H for v_H. Moreover, if v_L adopts with a nonnegative payoff, v_H obtains a strictly positive payoff from the same line and cannot choose nonadoption, so adopters form an upper interval of types.

A.2 Adoption targets and direct schedules

We first truncate the calendar at the first date the target is reached. For a terminal mass $q < 1$, let $V(q)$ be the infimum completion time over weakly increasing, right-continuous adoption schedules $\tau : [0, q] \to \mathbb{R}_+$ with beliefs given by (2) and satisfying, for all $m, n \in [0, q]$,

$$L_{\tau(m)}(\nu(m)) \geq L_{\tau(n)}(\nu(m)), \qquad L_{\tau(m)}(\nu(m)) \geq 0,$$

and $L_{\tau(m)}(v) \leq 0$ for every $v < \nu(q)$ in V and every assigned date $\tau(m)$. These are incentive compatibility, participation, and exclusion. This is an analytical representation rather than the primitive policy. An initial unused time interval may be removed by translating all used dates to the left, which changes neither the belief attached to any date nor any IC comparison.

The following proposition establishes the calendar representation used in Section 2.

Proposition A.1 (Adoption targets and access calendars)**.** *The value of the calendar problem is*

$$V^{\mathrm{C}}(\bar{M}) = \inf_{q \in [\bar{M}, 1)} V(q) = V(\bar{M}). \tag{29}$$

Access-calendar outcomes ending at the target and feasible direct schedules for $V(\bar{M})$ implement one another.

The proof has three steps. First, access can end when the target is first reached without changing earlier adoption or the belief at completion. The resulting terminal mass is $q \geq \bar{M}$; if $q < 1$, terminal odds equal $\nu(q)$ and the outcome gives a feasible direct schedule for $V(q)$. Second, $V(q_1) < V(q_2)$ whenever $\bar{M} \leq q_1 < q_2 < 1$. Third, an outcome

ending with full adoption can be cut at the target state $(m_1, b(m_1))$, where $b(m_1) = \nu(\bar{M})$, and completed strictly earlier.

The key restriction comes from participation at the last opportunity. After dates beyond completion are removed, the lowest type in an interior terminal pool must receive zero surplus, because a strictly positive surplus would also attract slightly lower types. Terminal odds therefore equal $\nu(q)$, and reaching a larger terminal mass requires more favorable beliefs and hence more evidence before completion.

Obtaining an earlier completion from this observation requires care about the agents who produce that evidence. If the target belief is reached during a pause, merely moving reopening to that instant can make waiting more attractive to agents who were supposed to adopt before closure. The proof instead adjusts the pool adopting at closure together with the final date, restoring the boundary agent's indifference while preserving her comparisons with other dates, and this adjustment underlies the strict comparison of terminal masses. This appendix first establishes truncation and the calendar representation, then gives the pause-cutting argument and treats full adoption, where no excluded type pins down terminal beliefs.

A.3 Ending access at completion

Let an equilibrium first reach $\bar{M}$ at the finite date T, and write $D_T = D \cap [0, T]$ for the calendar ending then.

Lemma A.1 (Truncation)**.** *Ending access at T preserves earlier adoption and the belief B_T. Its terminal adoption mass is*

$$q = 1 - F(B_T) \geq \bar{M}$$

when B_T is in the support, with $q = 1$ below the support and $q = 0$ above it.

Proof. Ordered self-selection implies that the types adopting weakly before T form an upper interval. Let v_T be its lower boundary. If an atom of adoption occurs at T, its lower boundary chooses the line at T and weakly prefers it to every line $s < T$. If adoption reaches the boundary continuously, take a sequence of assigned types and dates converging to (v_T, T); continuity of dated payoffs gives the same inequalities. Hence, for every $s < T$,

$$L_T(v_T) \geq L_s(v_T).$$

For $v < v_T$,

$$\frac{d}{dv}\{L_s(v) - L_T(v)\} = e^{-rs} - e^{-rT} > 0,$$

so $L_T(v) \geq L_s(v)$. Thus a lower type whose original plan used a date after T can, after deletion, choose only T or nonadoption; it cannot newly choose a strict pre-T date.

Every type originally assigned weakly before T still maximizes after dates are removed. Keep those choices. Assign each remaining type to T if the line at T is nonnegative and to nonadoption otherwise. New adoption occurs only at T, so (M_t, B_t) is unchanged for $t < T$, and simultaneous terminal adoption does not change B_T. Earlier adopters also have $v \geq B_T$, since participation required $v \geq B_{\tau(v)} \geq B_T$. At the final opportunity all remaining types with $v \geq B_T$ adopt. Together with the previously adopted upper interval, this gives terminal mass $1 - F(B_T)$ and preserves at least the original target mass. The stated boundary cases are immediate. □

If $q < 1$ and F is continuous and strictly increasing, the terminal identity in Lemma A.1 is $B_T = \nu(q)$. Ordered choice then supplies global IC, participation, and exclusion for a fixed-q direct schedule.

A.4 Calendars and direct schedules

Suppose first that a selected equilibrium terminates exactly at mass $q < 1$. Define its generalized inverse

$$\tau(m) = \inf\{t : M_t > m\} \quad (0 \leq m < q), \qquad \tau(q) = T.$$

After normalizing the first used date to zero, this is a weakly increasing, right-continuous version of the inverse; alternative choices at cutoff types have zero mass. Optimality initially holds almost everywhere. Monotonicity and continuity of dated payoffs extend its inequalities to every type by limits. Every assigned date belongs to the common menu, so individual optimality gives, for all m, n,

$$L_{\tau(m)}(\nu(m)) \geq L_{\tau(n)}(\nu(m)).$$

Comparison with nonadoption gives participation, and the choices of types below $\nu(q)$ give exclusion. The inverse schedule is therefore feasible for $V(q)$.

Conversely, take any finite-completion direct schedule satisfying those conditions and put

$$D = \overline{\tau([0, q])}. \tag{30}$$

Global IC makes each assigned date optimal among the image. If a limit date is added by closure, choose image dates converging to it. The induced belief path and (3) are continuous, so the same inequality holds at the limit. Exclusion makes nonadoption optimal for every lower type. The schedule and its belief path therefore form a selected equilibrium of D. Unused dates in any implementing calendar may be deleted; deletion removes deviations and closure adds only payoff limits. This proves outcome equivalence for schedules ending with adoption exactly at the specified terminal mass.

A.5 Truncating a pause

For $q < 1$, define the active endpoint $z(q)$ by

$$b(z(q)) = \nu(q). \tag{31}$$

Under (5), $m_0 < z(q) < q$, and $z(q)$ is strictly increasing in q. Appendix D proves decomposition and reconstruction for any fixed $q < 1$ from global IC, participation, exclusion, and the belief law alone; it does not assume that the principal prefers to end adoption exactly at the target. Write $W(z)$ for the infimum completion time over families of nonoverlapping pauses with endpoint z. Thus $V(q) = W(z(q))$. We use this fixed-endpoint result to compare different terminal masses.

Lemma A.2 (Pause truncation)**.** *Fix $d < z$ in the active range of the continuous-access path. There is a unique $p_z \in (d, z)$ such that (d, p_z, z) is a pause. Its duration*

$$\Delta(d, z) = \frac{x(z) - x(d)}{\beta p_z} \tag{32}$$

is continuously differentiable and strictly increasing in z. Consequently, if (d, p, h) is a pause and $d < z < h$, replacing it by (d, p_z, z) strictly shortens the pause.

Proof. Put $D = b(d)$, $Z = b(z)$, and, for a candidate p, $a = \nu(p)$. On the domain $d < p < z$ and $a > D$, define

$$K(p; d, z) = \frac{r}{\beta p} \log \frac{D}{Z} - \log \frac{a - Z}{a - D}. \tag{33}$$

The equation $K = 0$ is exactly the indifference equation.

For fixed a and p, the function

$$\phi_p(B) = B^{r/(\beta p)}(a - B)$$

has its unique maximum at $b(p) = r\nu(p)/(r + \beta p)$. At $p = d$, $D = b(d)$ is the maximizer, while $Z < D$, so $K(d; d, z) > 0$ in the right limit. If the domain ends before z because $\nu(p) \downarrow D$, the second logarithm diverges and $K \to -\infty$. Otherwise p can approach z; then $Z = b(z)$ is the unique maximizer and $K(z; d, z) < 0$ in the left limit. A root exists.

Differentiate while holding (d, z) fixed. With $a' = \nu'(p) < 0$,

$$K_p = -\frac{r}{\beta p^2} \log \frac{D}{Z} + a' \left\{ \frac{1}{a - D} - \frac{1}{a - Z} \right\} < 0. \tag{34}$$

Indeed, $D > Z$ makes the expression in braces positive. The root is therefore unique.

At a root, equality $\phi_p(D) = \phi_p(Z)$ with $D > Z$ implies

$$D > b(p) > Z. \tag{35}$$

Differentiating with respect to Z gives

$$K_Z = -\frac{r}{\beta p Z} + \frac{1}{a - Z} < 0, \tag{36}$$

where the inequality is equivalent to $Z < b(p)$. Thus $dp/dZ = -K_Z/K_p < 0$. Since $Z'(z) = b'(z) < 0$, the implicit pivot satisfies $p'_z > 0$.

Under indifference, the crossing equality can be written

$$\Delta(d, z) = \frac{1}{r} \log \frac{\nu(p_z) - Z}{\nu(p_z) - D}. \tag{37}$$

Differentiation yields

$$r\Delta_z = \nu'(p_z)p'_z \left\{ \frac{1}{\nu(p_z) - Z} - \frac{1}{\nu(p_z) - D} \right\} - \frac{Z'}{\nu(p_z) - Z} > 0. \tag{38}$$

The bracket is negative, $\nu' < 0$, $p'_z > 0$, and $Z' < 0$. This proves strict monotonicity and all asserted signs. □

A.6 Larger adoption targets take longer to reach

Let $q < 1$, let $b(z) = \nu(q)$, and consider a reconstructed prefix ending with a pause (d, p, z).

Lemma A.3 (Terminal pools)**.** *The reopening pool $[p, z]$ and terminal pool $[z, q]$ can use one final line while preserving incentive compatibility, participation, and exclusion.*

Proof. The payoff line at reopening in the pause has belief $b(z)$ and date equal to the completion date of the reconstructed prefix. The terminal line required by target zero rent has the same date and belief because $b(z) = \nu(q)$. They are therefore one affine line, not two distinct alternatives, and the two flat blocks merge into $[p, q]$.

Indifference makes the preceding payoff line at closure and this merged line cross at $\nu(p)$. All earlier adjacent intersections remain strictly above $\nu(p)$, so the descending-intersection argument continues to generate the global affine upper envelope. No new dated line has been added relative to the reopening in the pause, and hence no prefix type acquires a remote deviation. Agents with payoffs $\nu(m)$ for $m \in [p, q]$ have $\nu(m) \geq \nu(q) = b(z)$ and obtain a nonnegative payoff from the merged line; the agent with payoff $\nu(q)$ obtains zero. For $v < \nu(q)$, the merged final line is negative. Since it is the upper-envelope line for every value below the last intersection, every earlier line is weakly worse, so such a type

chooses nonadoption. The same argument applies when continuous adoption reaches z: the endpoint and terminal pool use one final payoff line. □

Lemma A.4 (Monotonicity)**.** *For $m_0 < z_1 < z_2$,*

$$0 \leq W(z_2) - W(z_1) \leq \mathcal{S}(z_1, z_2). \tag{39}$$

W is continuous and strictly increasing, and a minimizing family exists for every endpoint.

Proof. Take any rollout formed from nonoverlapping pauses ending at z_2. If z_1 lies outside all open pause supports, truncate at its state on the continuous-access path $(z_1, b(z_1))$ and delete all later components. If it lies inside a pause (d, p, h), replace the pause by the shortened pause (d, p_{z_1}, z_1) from Lemma A.2, then delete all later components. The retained pause supports have disjoint interiors. Chronological reconstruction places adjacent affine-line intersections at descending types and therefore preserves global IC through the payoff line at the revised reopening date. Add the target terminal pool whose lower type has zero rent at $b(z_1)$. Lemma A.3 shows explicitly that it merges with the pool adopting at the revised reopening date, preserves participation and exclusion, and creates no remote deviation. The new schedule is feasible for endpoint z_1 and is no longer than the original. Taking infima gives $W(z_1) \leq W(z_2)$.

Conversely, take an ε-optimal family for z_1, retain it, and add continuous adoption from z_1 to z_2. Reconstruction is again globally IC and adds exactly $\mathcal{S}(z_1, z_2)$. Letting $\varepsilon \downarrow 0$ gives the upper bound. Because $\mathcal{S}(z_1, z_2) \to 0$ as $z_2 \downarrow z_1$, these two bounds give continuity.

Appendix D supplies a minimizing countable family for z_2 under the maintained regularity. Its truncation is strictly shorter. Inside a pause, Lemma A.2 strictly lowers the duration. At a point outside all open pause supports, the deleted belief change is $x(z_2) - x(z_1) > 0$. Since $dX/dt = \beta M_t \leq \beta$, it requires at least $[x(z_2) - x(z_1)]/\beta > 0$ units of time. This also covers residual accumulation points with no adjacent interval of continuous adoption. Hence $W(z_1) < W(z_2)$. □

Since $z(q)$ is increasing, Lemma A.4 proves the strict terminal-mass comparison used in Proposition A.1 and

$$\inf_{q \in [\bar{M}, 1)} V(q) = V(\bar{M}).$$

A.7 The full-rollout boundary

For $q = 1$, exclusion no longer pins terminal zero rent; the lowest type may obtain positive terminal rent. We therefore do not define $V(1)$ by imposing $B_T = \nu(1)$. We first establish a local decomposition using an interior belief boundary.

Consider an ordered, incentive-compatible schedule with finite completion and first adoption at date zero. Fix $B^\star \in (0, B_0)$ and $z > m_0$ with $B^\star = b(z)$, and suppose the terminal belief is below $B^\star$.

Lemma A.5 (Local decomposition)**.** *Before beliefs reach $B^\star$, adoption consists of the initial pool through m_0, continuous-access adoption below z, and a finite or countable family of pauses reopening above $B^\star$. Continuous adoption is absolutely continuous; pause supports have disjoint interiors and accumulate only at or below z. At the boundary,*

1. *continuous adoption or pauses approach $(z, B^\star)$, possibly through countable accumulation;*

2. *a pause ends or begins at belief $B^\star$; or*

3. *one pause straddles $B^\star$.*

In the third case, the pause begins at a belief $D > B^\star$ with a unique adoption stock $d < z$ on the continuous-access path. The replacement pause $(d, p^\star, z)$ is unique and strictly shorter; the original post-pause belief need not be inverted. Any adoption atom at $B^\star$ uses one payoff line and merges with the final pool when the schedule is truncated there.

Proof. The neighboring-line argument in Appendix D proves $B_{\tau(m)} = b(m)$ at continuity points of τ in the support of its continuous Stieltjes part, using only global IC and the belief law. Before the present boundary, service beliefs satisfy

$$B_0 \geq B_{\tau(m)} > B^\star = b(z).$$

Strict monotonicity of b therefore places every such point of continuous support in $[m_0, z)$. The Lipschitz argument excluding a singular-continuous part also uses only the positive lower bound m_0 and boundedness of x' on the compact interval $[m_0, z]$. Thus the lower belief bound supplied by the terminal identity $B_T = b(m_1)$ is replaced, locally, by $B^\star$.

Consider next a jump of τ with adjacent beliefs $B^- > B^+$. Its clock and adjacent IC conditions are

$$B^+ = B^- e^{-\beta p \Delta}, \qquad \nu(p) - B^- = e^{-r\Delta}\{\nu(p) - B^+\}. \tag{40}$$

If $B^+ \geq B^\star$, both beliefs belong to $[b(z), b(m_0)]$. They therefore have unique inverses $B^- = b(d)$ and $B^+ = b(h)$ with $m_0 \leq d < p < h \leq z$. The proof in Appendix D then applies verbatim: (40) is indifference, the schedule is flat on the blocks assigned to closure and reopening, and chronologically distinct pauses have disjoint virtual support interiors. No terminal rent or exclusion condition is used. The same local argument rules out an

uncovered flat interval strictly below z: equal endpoint beliefs would give $b(a) = b(c)$ for $a < c$.

There can be at most one jump with $B^- > B^\star > B^+$ because beliefs are continuous in calendar time and strictly decrease while the positive adoption stock is held fixed. For that crossing pause only the belief at its beginning is inverted:

$$D := B^- = b(d), \qquad d < z.$$

The original post-pause belief may lie below the range of b used in the fixed- target decomposition, so we do not define an endpoint on the continuous-access path for the end of the original pause. Set $Z = B^\star = b(z)$. Lemma A.2, whose existence proof uses only $d < z$, supplies a unique pivot $p^\star \in (d, z)$ satisfying indifference between D and Z.

To compare clocks directly in beliefs, define $p(Z')$ by the indifference equation (33) for fixed D and a variable lower endpoint belief Z'. At every root, (35)–(36) give $D > b(p) > Z'$, $K_p < 0$, $K_{Z'} < 0$, and hence $p_{Z'} < 0$. The pause duration is

$$\Delta(D, Z') = \frac{1}{r} \log \frac{\nu(p(Z')) - Z'}{\nu(p(Z')) - D},$$

so

$$r\Delta_{Z'} = \nu'(p) p_{Z'} \left\{ \frac{1}{\nu(p) - Z'} - \frac{1}{\nu(p) - D} \right\} - \frac{1}{\nu(p) - Z'} < 0. \tag{41}$$

The first term is negative because $\nu' < 0$, $p_{Z'} < 0$, and the bracket is negative; the second is negative. The original pivot and beliefs satisfy the same indifference equation by (40). Let p_0 be the original pivot. Its positive payoff at either endpoint gives $\nu(p_0) > D$. For $B^+ < Z' < D$, strict unimodality implies $\phi_{p_0}(Z') > \phi_{p_0}(D) = \phi_{p_0}(B^+)$, and hence $K(p_0; D, Z') < 0$. At $p = d$ the opposite strict sign holds. Thus every $Z' \in (B^+, B^\star]$ has a unique root in (d, p_0), entirely inside the participation domain. At $Z' = B^+$ the root is p_0. The implicit-function theorem therefore connects the original and cut pivots. This bracketing does not require the equation $\nu(p) = D$ to have a solution inside the population support, nor a corresponding point on the continuous-access path for B^+. Since $B^\star > B^+$, equation (41) makes the replacement pause to $B^\star$ strictly shorter than the original straddling pause. This proves the third case without inverting B^+.

It remains to close the prefix. The components wholly before the boundary are chronological and have disjoint supports in $[m_0, z]$. A finite truncation, replacing omitted components by continuous adoption, reconstructs a globally IC upper envelope. If countably many components accumulate at the boundary, their total pause duration is bounded by the original finite completion time and their supports have total length at most $z - m_0$. Taking nested finite subfamilies and filling omitted supports with continuous adoption, service dates and beliefs converge to the line with belief $b(z) = B^\star$;

the global upper-envelope inequalities are closed under this convergence. Thus countable accumulation creates no missing last component and no remote deviation.

Finally suppose a positive adoption atom occurs exactly when the belief is $B^\star$. Beliefs do not jump at adoption, so every agent in that atom faces the same affine line. The local no-uncovered-flat argument implies that the atom comprises the agents adopting when a pause ends at the boundary, when a pause begins there, at the shared date of two such pauses, or in a boundary pool on the limiting payoff line under continuous access. Truncation discards every component after the boundary and keeps only this one line. Assigning further terminal adopters to the same date and belief merely enlarges the same flat interval; it does not add a new deviation. This treats an atom at the crossing and completes the local decomposition. □

Suppose an equilibrium first reaches $\bar{M}$ at T, access ends then, and the terminal mass is one.

Lemma A.6 (Full adoption)**.** *There is an equilibrium ending with adoption at $\bar{M}$ strictly before T.*

Proof. Full terminal participation implies $B_T \le \underline{v} < \nu(\bar{M})$. Put $z = z(\bar{M})$, so $b(z) = \nu(\bar{M})$. Beliefs are continuous and fall from $B_0 > b(z)$ to $B_T < b(z)$; hence they cross the target belief on the continuous-access path before T. Ordered self-selection and individual optimality over the common menu give the outcome with access ending at completion a globally IC inverse schedule on $[0, 1]$; only lower-tail exclusion, which is not used in Lemma A.5, is absent.

Apply Lemma A.5 with $B^\star = b(z)$. Keep all components before the boundary. If a pause straddles the boundary, use its strictly shorter replacement pause; otherwise close the continuous-access path or the incident pause at z. Countable left accumulation and any adoption atom at the boundary are covered by the lemma. Assign $[z, \bar{M}]$ to the final line. Lemma A.3 gives participation, target-tail exclusion, the global affine upper envelope, and absence of remote deviations. The terminal line has belief $b(z) = \nu(\bar{M})$ and therefore zero rent for the marginal target type.

If there is no straddling pause, the constructed completion date is the original boundary date, which is strictly before T. If a pause straddles the boundary, the replacement pause ends strictly before the original pause and hence before T. The calendar implementation follows from (30). □

Combining Lemmas A.1, A.4, and A.6 gives the lower bound $V^{\mathrm{C}}(\bar{M}) \ge V(\bar{M})$. The calendar implementation of any ε-optimal direct schedule gives the reverse inequality. This completes the proof of Proposition A.1.

B Continuous availability

We use two local observations that also apply to available intervals followed by a future pause. First, for a fixed type,

$$L_t'(v) = e^{-rt}\{(r + \beta M_t)B_t - rv\}. \tag{42}$$

A positive adoption atom at an interior available date raises the right derivative relative to the left. It cannot give the downward kink required at a maximum. Hence no positive atom occurs at such a date.

Second, consider an open pause between used dates $a < c$, with stock $m > 0$ and limiting boundary type $\nu(m)$. Limits of the two adjacent groups' IC inequalities give

$$\nu(m) - B_a = e^{-r(c-a)}\{\nu(m) - B_a e^{-\beta m(c-a)}\}.$$

The fixed-stock payoff $u(s) = e^{-rs}\{\nu(m) - B_a e^{-\beta m s}\}$ is strictly unimodal. Equal values at $s = 0$ and $s = c - a$ imply $B_a > b(m)$ and $u'(0) > 0$. If a right neighborhood of a is available, an immediately later date therefore strictly dominates a for the boundary type and for nearby types assigned on the left. Such a pause is impossible. This argument compares equilibrium endpoint payoffs before subsequent adoption; it does not assume that the stock stays fixed after the pause.

At the start of a positive-length available interval beginning at date zero, there cannot be initial waiting: prior to the first adoption beliefs remain B_0, and types above B_0 prefer the available date zero. Let m be the date-zero mass. If $m > m_0$, the bottom of that pool has a strictly positive right derivative in (42) and delays. If $m < m_0$, the derivative for the highest residual types is strictly negative at zero. Right continuity then rules out their adopting at interior dates accumulating at zero. If adoption resumes later, it would create the impossible pause just described; if it never resumes, residual types sufficiently close to $\nu(m) > B_0$ have positive date-zero payoffs and cannot remain out. Thus $m = m_0$.

At every subsequent point in the support of continuous adoption, limits of neighboring types' first-order conditions give $B_t = b(M_t)$. A last pause with no subsequent adoption is also impossible: at a state on the continuous-access path $(m, b(m))$, nearby residual types have positive payoffs and fixed-stock maxima arbitrarily close to that date. Thus, while the solution remains in the interior range of adoption masses, every available right neighborhood contains adoption. The preceding no-atom and no-pause arguments make its support the full available interval. Since $b' < 0$, $M = b^{-1}(B)$ is absolutely continuous on each compact interior range, and

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

This also excludes singular-continuous adoption. The scalar equation, or equivalently the clock $\int_{m_0}^{M_t} x'(m)/(\beta m)\,dm = t$, has a unique solution. This proves Proposition 1.

Proof of the finite-horizon assertion in Proposition 1. The same arguments apply strictly before T^{R} even though access ends there. They fix the initial pool and entire initial continuous-access path. Its left limit at the horizon is m_1, and its belief is $b(m_1) = \nu(\bar{M})$. For any residual type $v \le \nu(m_1)$, along the continuous-access path

$$L_t'(v) = re^{-rt}\{\nu(M_t) - v\} \ge 0.$$

The horizon line is therefore its best dated line. It is positive exactly for $v > \nu(\bar{M})$, forcing the terminal pool through $\bar{M}$ and excluding lower types. The proposed schedule is globally IC, so this establishes existence as well as uniqueness, up to null cutoff ties. □

C Unique implementation of finite pause calendars

Let $e = (d, p, h)$ be a pause with $m_0 \le d < p < h \le m_1$.

Lemma C.1 (Uniqueness)**.** *The calendar*

$$D_e = [0, t_d] \cup [t_h, T^e], \qquad t_d = \mathcal{S}(m_0, d), \quad t_h = t_d + \Delta(d, p, h), \quad T^e = t_h + \mathcal{S}(h, m_1)$$

uniquely implements the pause path ending with adoption at the target.

Either available interval may collapse to an endpoint, and the pause may increase or decrease completion time.

Proof. Existence follows from the global-IC reconstruction in Appendix D. We show that every equilibrium has the same path. Write $D_0 = b(d)$ and $\Delta = t_h - t_d > 0$; D_0 is a belief, not a calendar. If $t_d > 0$, Appendix B's local arguments fix the initial pool and continuous-access path strictly before t_d, with left state $(d, b(d))$. They remain valid in the presence of the later pause. If $t_d = 0$, there is no prefix and $D_0 = B_0$.

Let $P = M_{t_d}$ include any atom at closure. Ordered choice makes $\nu(P)$ the cutoff payoff. We have $P > 0$: if $t_d = 0$ and nobody adopts then, beliefs remain B_0 until the first later adoption, contradicting the date-zero preference of types above B_0. For $t_d > 0$, already $P \ge d > 0$. Participation and the limiting earlier adopters' IC give $\nu(P) > D_0$. Indeed, equality would make a deviation to reopening strictly profitable, since the positive stock reduces beliefs during the pause. Consequently $P < \bar{M}$, because $D_0 > b(m_1) = \nu(\bar{M})$. There must be positive post-pause adoption: residual types close to $\nu(P)$ have positive payoffs at closure, so they cannot all choose nonadoption.

Measure elapsed time from t_d and let $s_* \ge \Delta$ be the infimum of post-pause adoption dates. For agents with payoffs $\nu(q_n)$ where $q_n \downarrow P$, ordering implies $\tau(q_n) - t_d \to s_*$. Until

that date the stock is P; continuity of beliefs, including at atoms, makes the limiting boundary payoff

$$u_P(s) = e^{-rs}\{\nu(P) - D_0 e^{-\beta P s}\}, \qquad 0 \le s \le s_*. \tag{43}$$

Its derivative is

$$u_P'(s) = e^{-rs}\{(r + \beta P) D_0 e^{-\beta P s} - r\nu(P)\}.$$

If $P \le d$, then $D_0 \le b(P)$ and $u_P(s) < u_P(0)$ for every $s > 0$. Taking limits in the first post-pause adopters' comparison with adoption at closure gives the opposite inequality at s_*. Thus $P > d$. This also means a positive mass adopts at closure, so types assigned exactly at t_d approach $\nu(P)$ from above.

For $P > d$, the fixed-stock payoff has a unique strict maximum at

$$s_0(P) = \frac{1}{\beta P} \log \frac{D_0}{b(P)} > 0.$$

If $s_0(P) \ge \Delta$, adopting at reopening strictly beats adopting at closure for the boundary type, contradicting limits of the early pool's IC. Therefore $s_0(P) < \Delta$. If $s_* > \Delta$, the stock remains P up to s_* and u_P strictly decreases on $[\Delta, s_*]$. Limits of the first post-pause adopters' best responses would require $u_P(s_*) \ge u_P(\Delta)$, a contradiction. Hence

$$s_* = \Delta. \tag{44}$$

We now account explicitly for learning *after* reopening. Let t be any later available date, evaluated under the actual equilibrium belief path. For the same agents with payoffs $\nu(q_n)$ where $q_n \downarrow P$, global optimality gives

$$L_{\tau(q_n)}(\nu(q_n)) \ge L_t(\nu(q_n)).$$

Their chosen dates converge to t_h by (44). Continuity in date and type therefore implies

$$L_{t_h}(\nu(P)) \ge L_t(\nu(P)). \tag{45}$$

Thus the payoff line at reopening is the boundary type's best later line without extending the fixed-stock calculation beyond its valid interval. Taking limits from the pool at closure and the types adopting at or after reopening yields equality of those two lines:

$$\nu(P) = C(P) := \frac{D_0 - e^{-r\Delta} D_0 e^{-\beta P \Delta}}{1 - e^{-r\Delta}}. \tag{46}$$

The left side is strictly decreasing, whereas

$$C'(P) = \frac{e^{-r\Delta}\beta\Delta D_0 e^{-\beta P\Delta}}{1 - e^{-r\Delta}} > 0.$$

There is at most one root. Indifference supplies the intended root p, so $P = p$ and $B_{t_h} = b(h)$.

If $h = m_1$, reopening is the horizon. Single crossing and terminal participation then force the remaining target types to adopt and all lower types to remain out. Suppose instead $h < m_1$, and put $Q = M_{t_h}$. Participation at reopening bounds $Q < \bar{M}$. If $Q > h$, the bottom of its pool has a strictly positive right payoff derivative and gains from a short available delay. If $Q < h$, the limiting residual type has a strictly negative right derivative. Right continuity first rules out adoption dates accumulating immediately after t_h: types approaching $\nu(Q)$ at such interior dates could not satisfy their local first-order conditions. If adoption resumes at a later first support date, the stock is Q until then, and the fixed-stock payoff strictly decreases after reopening, contradicting the limiting first adopters' IC. If it never resumes, nearby residual types have positive payoffs at reopening and prefer that line to all earlier lines by single crossing at $\nu(p)$. They cannot remain out. Thus $Q = h$.

The local continuous-availability argument now fixes the continuous-access path from h to m_1. At T^e the belief is $\nu(\bar{M})$ and the same horizon-line argument as in Proposition 1 forces exactly the terminal pool through $\bar{M}$. This exhausts the calendar and proves uniqueness. □

Finite-family induction. Now let $\mathcal{A} = \{(d_j, p_j, h_j)\}_{j=1}^{J}$ be any finite chronological family of nonoverlapping pauses, and use its associated calendar: continuous availability on the uncovered intervals and the intended open calendar pause for each pause. We show successively that every equilibrium has the reconstructed aggregate path.

Continuous availability fixes the initial continuous-access path up to the first closing date, exactly as in the preceding proof. Suppose inductively that the intended path and hence the belief at closure $D_j = b(d_j)$ have been fixed through the beginning of pause j. Let P_j be adoption immediately after its closing date. The argument leading to (44) applies without reference to the form of later calendar pauses: there is residual participation, the first post-pause adoption must occur at the announced reopening date, and the boundary types on its two sides make the payoff lines at closure and reopening intersect at $\nu(P_j)$. Thus

$$\nu(P_j) = \frac{D_j - e^{-r\Delta_j} D_j e^{-\beta P_j \Delta_j}}{1 - e^{-r\Delta_j}}. \tag{47}$$

The left side is strictly decreasing in P_j and the right side is strictly increasing. Indifference supplies the unique root $P_j = p_j$, so the belief at reopening is $b(h_j)$.

If $h_j < d_{j+1}$, there is an interval of continuous availability after reopening. The derivative and first-residual-type argument used above forces the adoption stock at reopening to be h_j, after which the local continuous-access indifference condition uniquely fixes the path through d_{j+1}. If $h_j = d_{j+1}$, the shared reopening and closing date may have a larger atom. Apply (47) directly to the next pause, whose belief at closure is $b(h_j) = b(d_{j+1})$; it uniquely fixes that atom's terminal stock at p_{j+1}. This covers adjacent pauses without artificially inserting an interval of continuous adoption. After the last reopening, continuous availability and the terminal-line argument fix the final segment of continuous adoption and terminal pool. Induction proves aggregate uniqueness for the associated calendar of every finite family of nonoverlapping pauses.

For each such family, best- and worst-equilibrium completion times therefore coincide with (19). The finite-family infimum in (20) and the elementary inequality $V^{\mathrm{wc}} \geq V^{\mathrm{C}}$ now give (21).

For one pause, exact accounting gives $T^e = T^{\mathrm{R}} + \Gamma_F(e)$. Hence a pause that saves time is a uniquely implemented improvement, proving the implementation claim in Theorem 2(3). The finite-family argument does not claim that a calendar generated by a countably infinite family has a unique equilibrium; its finite subfamilies are enough for equality of the two value criteria.

D Proof of the global rollout characterization

The characterization includes the following boundary and regularity properties. The initial pool contains $[0, m_0]$ and the terminal pool contains $[m_1, \bar{M}]$. Adoption between pause supports follows absolutely continuous adoption. The infimum is attained by a finite or countable family of nonoverlapping pauses whose reconstruction is a calendar equilibrium. If an optimal family is finite, its associated calendar uniquely implements an optimum. No aggregate-uniqueness assertion is made for a countably infinite minimizing calendar.

We work directly with a finite-completion ordered schedule satisfying global dated-line IC, participation, exclusion, and the belief law. The terminal mass is fixed at $\bar{M}$. Nothing in this proof assumes that exact attainment is optimal in the primitive at-least-target problem. The same argument applies to any fixed terminal mass $q \in [\bar{M}, 1)$ after replacing the active endpoint by the solution of $b(z(q)) = \nu(q)$; Appendix A uses that fixed-mass geometry to establish the exact-target reduction.

D.1 Boundary normalization and the continuous part

If the first used opportunity is $t_0 > 0$, no adoption occurs before it and beliefs remain B_0. Subtracting t_0 from every used date leaves the belief attached to each line unchanged and multiplies every dated payoff by the same positive factor. We therefore normalize the first

used date to zero. An unnormalized outcome has an additional t_0 in (19). Participation of the terminal marginal type and exclusion immediately below it give $B_T = \nu(\bar{M})$. Write

$$X_t = -\log B_t, \qquad B_0 = b(m_0), \qquad B_T = b(m_1).$$

Beliefs attached to used dates consequently lie in $[b(m_1), b(m_0)]$.

Cutoff comparisons are understood by one-sided limits from the full-measure set on which equilibrium IC holds; the bracketing points below are chosen from that set. Let $d\tau$ be the Stieltjes measure of the monotone right-continuous schedule of adoption dates, let J be its countable jump set, and let μ_c be the continuous part of $d\tau$. At an interior continuity point $m \in \operatorname{supp} \mu_c$, choose $a_n < m < b_n$ converging to m with $\tau(b_n) > \tau(a_n)$. Global IC puts the intersection c_n of the two dated lines between $\nu(b_n)$ and $\nu(a_n)$, so $c_n \to \nu(m)$. Put $\Delta_n = \tau(b_n) - \tau(a_n)$. Continuity gives $\Delta_n \to 0$, and ordering bounds the adoption stock between a_n and b_n at almost every intervening date. Hence

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

The affine-line intersection formula and expansion at $\Delta_n = 0$ give $c_n \to B_{\tau(m)}(r + \beta m)/r$. Therefore

$$B_{\tau(m)} = b(m). \tag{48}$$

The boundary belief bounds and strict monotonicity of b place all such points in $[m_0, m_1]$. Because μ_c is nonatomic and J is countable, μ_c is supported on this active range.

For two continuity points $a < c$ of τ in $\operatorname{supp} \mu_c$, equation (48) implies

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

Thus

$$\tau(c) - \tau(a) \leq C_0(c - a), \qquad C_0 = \frac{\|x'\|_\infty}{\beta m_0}.$$

This also bounds μ_c on an arbitrary interval. Indeed, continuity points in its support carry all of its mass: choose such points approaching the first and last portions of the interval that carry continuous mass. The mass omitted at the two ends tends to zero, while the mass between the chosen points is at most their difference in τ. Consequently $\mu_c(I) \leq C_0|I|$ for every interval I. It follows that μ_c is absolutely continuous. There is no singular-continuous component of the schedule.

D.2 Adoption at closure and reopening

Let $p \in J$ have adjacent used dates $t^- < t^+$, beliefs $B^- > B^+$, and duration $\Delta = t^+ - t^-$. The stock is p during the pause, so

$$B^+ = B^- e^{-\beta p \Delta}. \tag{49}$$

One-sided limits of global IC make the adjacent lines cross at $a = \nu(p)$:

$$a - B^- = e^{-r\Delta}(a - B^+). \tag{50}$$

The common net payoff is strictly positive. For example, solving this equality for a gives $a = B^- + e^{-r\Delta}(B^- - B^+)/(1 - e^{-r\Delta}) > B^-$. The function

$$B \longmapsto B^{r/(\beta p)}\{a - B\}$$

is strictly increasing up to its unique maximum at $b(p)$ and strictly decreasing thereafter. Equations (49)–(50) therefore imply $B^- > b(p) > B^+$. The boundary belief range gives unique virtual endpoints d, h such that

$$B^- = b(d), \qquad B^+ = b(h), \qquad m_0 \le d < p < h \le m_1.$$

The pause duration and indifference condition are (15)–(16).

The virtual supports are already ordered before we identify their pools. If pause i precedes pause j, chronological belief monotonicity gives $B_i^+ \ge B_j^-$, and hence

$$h_i \le d_j.$$

Their support interiors are thus disjoint. In particular, no other jump pivot can lie in (d, p) or (p, h).

There is also no continuous Stieltjes mass in either interval. A continuous support point $m \in (d, p)$ would, by (48), require belief $b(m) < b(d)$ at a date no later than t^-, which is impossible. A continuous support point $m \in (p, h)$ would require belief $b(m) > b(h)$ at a date no earlier than t^+, also impossible. Since the jump part and continuous part both vanish on these intervals, τ is flat on (d, p) at t^- and on (p, h) at t^+. Each pause therefore has the two adoption pools used in its characterization. If $h_i = d_j$, the pool at the end of one pause and the pool at the start of the next are assigned to one shared opportunity.

D.3 Boundary pools and continuous adoption between pauses

There is no continuous Stieltjes mass below m_0 or above m_1. A jump pivot below m_0 would require a pause beginning at a virtual endpoint $d < p < m_0$ despite $d \ge m_0$; a pivot

above m_1 gives the symmetric contradiction. The initial pool therefore contains $[0, m_0]$, and the terminal pool contains $[m_1, \bar{M}]$. Either pool can be enlarged by an incident boundary pause. There is no jump at m_0 or m_1, so the service beliefs at these points are $b(m_0)$ and $b(m_1)$.

Let

$$R = [m_0, m_1] \setminus \bigcup_j (d_j, h_j)$$

be the closed residual set. We now show that $B_{\tau(m)} = b(m)$ on all of R, not merely at points supporting a positive continuous derivative. Let K consist of the continuity points of $\operatorname{supp} \mu_c$, all virtual pause endpoints, and m_0, m_1. We have proved $K \subseteq R$ and the continuous-access indifference condition on K.

To see that K is dense in R, suppose a component (a, c) of $[m_0, m_1] \setminus \overline{K}$ meets R. It contains no continuous Stieltjes mass and no virtual endpoint. It cannot contain a jump pivot: the corresponding support would have to contain all of (a, c), since both of its endpoints lie outside that component, contrary to its meeting R. Thus τ is constant on (a, c). The endpoints belong to $\overline{K}$ and are not jump pivots, because each pivot has a neighborhood inside its own open pause support. Continuity of the service belief at the endpoints gives

$$b(a) = B_{\tau(a)} = B_{\tau(c)} = b(c),$$

contradicting $a < c$ and strict monotonicity of b. This proves density and the continuous-access indifference condition on R. In particular, no positive-length interior pool is left uncovered by the pause supports and the boundary blocks.

At almost every density point m of R, the monotone schedule is differentiable. Each positive-mass pool contains at most one point in R, since the service belief is constant on a pool whereas b is strictly decreasing. There are at most countably many such pools. Outside this exceptional set, a density point of R is not pooled, and therefore $M_{\tau(m)} = m$ and X is differentiable at $\tau(m)$ with derivative βm. At such a density point, the derivatives of $X_{\tau(m)}$ and $x(m)$ agree because the two functions agree on R. The chain rule therefore gives

$$\tau'(m) = \frac{x'(m)}{\beta m} \qquad \text{for Lebesgue-almost every } m \in R. \tag{51}$$

The continuous derivative vanishes in the pause pools and boundary pools. This proves absolutely continuous adoption along the continuous-access path on the residual set, even when that set is not a union of intervals.

The only remaining Stieltjes mass is the pause duration at each pivot. Integrating (51) and summing these atoms gives

$$T = \int_R \frac{x'(m)}{\beta m}\, dm + \sum_j \Delta_j = \mathcal{S}(m_0, m_1) + \sum_j \Gamma_F(d_j, p_j, h_j).$$

This is (19). Moreover,

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

so the gain series is absolutely convergent.

D.4 Finite reconstruction and global IC

Take a finite family of nonoverlapping pauses in chronological order. Follow the continuous-access path on uncovered intervals and insert the pause on each support. Each pause begins and ends at the same belief states on the continuous-access path as the segment it replaces, and its duration satisfies the belief law. Boundary and incident pools use one line whenever their dates coincide.

Here is a direct verification of all dated-line comparisons. On an interval of continuous access, a fixed type has derivative

$$L_t'(v) = re^{-rt}\{\nu(M_t) - v\}.$$

Across a pause with pivot p, indifference gives

$$L_{t^-}(v) - L_{t^+}(v) = \{e^{-rt^-} - e^{-rt^+}\}\{v - \nu(p)\}.$$

Cumulative adoption on the continuous-access path and at pause pivots is weakly increasing over time. Thus, for any fixed type, movement through the available dates improves its payoff until its assigned line and weakly lowers it thereafter. This proves its comparison with every earlier and later available date, including remote dates across several pauses. Equivalently, the chronological dated lines form their global affine upper envelope.

For all agents in the adopting population, comparison with the terminal line gives participation because $B_T = \nu(\bar{M})$. For $v < \nu(\bar{M})$, every move along the continuous-access path and pause crossing weakly favors the later line; the terminal line is therefore the best dated line and is negative. These types choose nonadoption. Taking the closure of the used image adds only payoff limits, as in (30). The reconstruction is a calendar equilibrium. Appendix C proves uniqueness of its aggregate equilibrium for the calendar associated with the finite family.

For any original countable decomposition, finite subfamilies approximate its total gain by absolute convergence. Each such subfamily is feasible by the preceding construction. Conversely, every finite family gives a feasible outcome ending with adoption exactly at the target. These two observations prove (20) under the maintained regularity. Finite-family uniqueness and $V^{\mathrm{wc}} \geq V^{\mathrm{C}}$ then give (21).

D.5 Uniform bounds and countable reconstruction

Write $A = \|x'\|_\infty$ and $L = m_1 - m_0$. For a pause $e = (d, p, h)$ with span $\ell(e) = h - d$, equation (18) directly gives

$$|\Gamma_F(e)| \leq \frac{A}{\beta m_0^2}\ell(e)^2, \qquad \left|\frac{1}{p} - \frac{1}{m}\right| \leq \frac{\ell(e)}{m_0^2} \quad (d \leq m \leq h). \tag{52}$$

The pause duration also gives

$$0 < \Delta(e) \leq \frac{A}{\beta m_0}\ell(e), \qquad \mathcal{S}(d, h) \leq \frac{A}{\beta m_0}\ell(e). \tag{53}$$

Disjoint interiors imply $\sum_e \ell(e) \leq L$. Deleting every pause of span below δ changes total gain by at most $AL\delta/(\beta m_0^2)$ and leaves at most L/δ pauses. Finite approximation therefore requires no extra smoothness.

We give an explicit countable reconstruction because pauses can accumulate without a first chronological member. Let τ^{R} be the continuous-access schedule ending with adoption exactly at the target, including its boundary pools. For one pause define the right-continuous correction to adoption dates

$$c_e(m) = \begin{cases} 0, & m < d, \\ -\mathcal{S}(d, m), & d \leq m < p, \\ \Delta(e) - \mathcal{S}(d, m), & p \leq m \leq h, \\ \Gamma_F(e), & m > h. \end{cases}$$

The reopening date is chosen at the null pivot. For a finite family of nonoverlapping pauses, the reconstructed schedule is $\tau^{\mathrm{R}} + \sum_e c_e$. The uniform estimate

$$\|c_e\|_\infty \leq \Delta(e) + \mathcal{S}(d, h) \leq \frac{2A}{\beta m_0}\ell(e)$$

makes these corrections summable for every countable family of nonoverlapping pauses. Choose nested finite subfamilies exhausting that family and fill each omitted support with continuous adoption. Their schedules converge uniformly to a finite monotone right-continuous schedule τ; the limit is independent of the enumeration.

Belief consistency passes to the same limit. Fubini's theorem gives, for any schedule for the fixed adopting population,

$$X_\tau(t) = X_0 + \beta \int_0^{\bar{M}} (t - \tau(m))_+ \, dm.$$

Consequently

$$\sup_{t \geq 0} |X_{\tau_n}(t) - X_\tau(t)| \leq \beta \bar{M} \|\tau_n - \tau\|_\infty.$$

Each X_τ is $\beta\bar{M}$-Lipschitz, so beliefs at assigned dates converge uniformly as well. Dated payoffs then converge uniformly on the compact type range. Global IC, participation, and exclusion are closed inequalities and pass from the finite reconstructions to τ. The closure of its used image implements the countable schedule as an access calendar. This argument needs neither uniform convergence of cumulative adoption at its atoms nor a first or last pause at an accumulation point.

D.6 Existence of an optimal family

The closed pause set, augmented by total diagonals $d = p = h$, is compact. Indeed, the indifference condition is continuous on $m_0 \leq d \leq p \leq h \leq m_1$, where p is bounded away from zero. A partial diagonal $p = d < h$ or $d < p = h$ cannot satisfy indifference: one endpoint would be the strict maximum of the function in the indifference condition and the other would have a strictly smaller value. Thus all nonnull controls in the closure remain pauses.

Attach to each finite family of nonoverlapping pauses the marked measure and gain density

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

The function g is continuous away from the diagonal by its defining formula, and (52) makes it continuous at the diagonal. The measures have mass at most L, so a minimizing sequence has a weakly convergent subsequence on the compact control set, with limit μ.

We verify that its nonnull part describes an actual family of nonoverlapping pauses. For every $\varepsilon > 0$, the region $\ell \geq \varepsilon$ contains only finitely many points of $\operatorname{supp} \mu$. Otherwise, choose more than $2L/\varepsilon$ distinct such points and disjoint neighborhoods on which $\ell > \varepsilon/2$. Each neighborhood has positive limiting measure and therefore eventually contains a control from the approximating family, contradicting its length bound. Taking the union over $\varepsilon = 1/k$ shows that the nonnull part of μ is countably atomic and locally finite away from the diagonal.

Fix a nonnull limiting atom e. Choose a sufficiently small neighborhood U containing no other limiting atom and no diagonal, with $\mu(\partial U) = 0$, such that all supports of controls

in U contain one common interior point. An approximating family of nonoverlapping pauses places at most one control in U. Since $\mu_{\mathcal{A}_n}(U) \to \mu(\{e\}) > 0$, it eventually places exactly one there, and that control converges to e. Its marked weight converges to $\ell(e)$, proving

$$\mu(\{e\}) = \ell(e).$$

Moreover, two nonnull limiting atoms cannot have overlapping support interiors. Small neighborhoods of such a pair would require two conflicting controls in every sufficiently late approximating family. The nonnull atoms therefore form a countable family of nonoverlapping pauses. Any remaining measure lies on the total diagonal and has zero gain.

Finally, weak convergence and continuity of g give

$$\lim_n \sum_{e \in \mathcal{A}_n} \Gamma_F(e) = \int g \, d\mu = \sum_{e \in \mathcal{A}_*} \Gamma_F(e),$$

where $\mathcal{A}_*$ is the countable family of nonoverlapping pauses represented by the nonnull limiting atoms. Its reconstruction was proved above, so it attains the finite-family infimum. Its finite subfamilies approximate its value. The entire attainment argument uses only the maintained regularity and the compact active range; additional smoothness is reserved for the local expansion in Appendix E. This completes Theorem 1.

Adoption at closure for fixed endpoints. For the computation in Section 3.4, we verify that each pair of endpoints determines a unique stock of adopters at closure. Fix $d < h$, put $D = b(d) > Z = b(h)$, and define

$$R(d, h, p) = \nu(p) - D - \frac{D - Z}{\exp\{r \log(D/Z)/(\beta p)\} - 1}.$$

The last fraction is strictly increasing in p, whereas ν is strictly decreasing, so $R_p < 0$. At $p = d$, write $L = \log(D/Z) > 0$ and $k = r/(\beta d)$. Since $\nu(d) = D(1 + 1/k)$,

$$R(d, h, d) = D \left\{ \frac{1}{k} - \frac{1 - e^{-L}}{e^{kL} - 1} \right\} > 0.$$

At $p = h$, now with $k = r/(\beta h)$,

$$R(d, h, h) = Z \left\{ \frac{1}{k} - \frac{e^{L} - 1}{1 - e^{-kL}} \right\} < 0.$$

Both signs follow from $1 - e^{-u} < u < e^{u} - 1$ for $u > 0$. There is therefore a unique root $p(d, h) \in (d, h)$. It is the stock of adopters at closure, and is continuously differentiable off the diagonal by the implicit function theorem. Set $p(d, d) = d$. The inequalities

$d \le p(d,h) \le h$ give continuity at the diagonal. Since $m_0 > 0$ and x' is continuous, (22) is continuous on the compact endpoint triangle, including $g(d,d) = 0$. To obtain convergence of the grid calculation, take a finite family whose cost is within ε of the unrestricted infimum. On a sufficiently fine grid, move its distinct endpoints to nearby grid points while preserving their order and shared endpoints. The intervals remain nonoverlapping, and continuity of g makes the change in their total cost tend to zero. The grid optimum is no worse than this family and cannot beat the unrestricted value. Letting the mesh and then ε tend to zero gives the claim.

E Distributional shape and adopter welfare

E.1 Nonincreasing density

For one pause abbreviate $a = \nu(p)$, $k = r/(\beta p)$, and

$$\mathcal{Q}(m) = \log\{b(m)^k(a - b(m))\}, \quad L(m) = -\frac{b'(m)}{b(m)}, \quad H(m) = \frac{1}{p} - \frac{1}{m}.$$

Indifference gives $\mathcal{Q}(d) = \mathcal{Q}(h)$. Differentiation yields

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

Thus $\mathcal{Q}' > 0$ on (d,p) and $\mathcal{Q}' < 0$ on (p,h). Put $\psi(m) = H(m)/R(m)$, with its continuous value at p. Then

$$\beta\Gamma_F = \int_d^h \mathcal{Q}'(m)\psi(m)\,dm. \tag{54}$$

If $f' \le 0$, then $\nu''(m) = -f'(\nu(m))/f(\nu(m))^3 \ge 0$. To see what this convexity does, write for $m \ne p$

$$s(m) = -\frac{\nu(m) - a}{m - p} > 0.$$

Convexity makes s nonincreasing. Algebra gives $\psi = -\phi$, where

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

If $D_s = a\beta + (r + \beta p)s$, then

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

Since s is nonincreasing, ϕ is strictly increasing and ψ strictly decreasing. Also $\int_d^h \mathcal{Q}' = 0$. Subtracting $\psi(p)$ in (54) gives

$$\beta\Gamma_F = \int_d^h \mathcal{Q}'(m)\{\psi(m) - \psi(p)\}\, dm > 0 :$$

both factors are positive left of p and both are negative right of p. The decomposition theorem proves parts (1)–(2) of Theorem 2.

E.2 Small pauses

Suppose ν is C^4 near $m \in (m_0, m_1)$, and write $v = \nu(m)$, $f = f(v)$, and $f' = f'(v)$. Fix the pivot at m, set $d = m - \varepsilon$, and choose $h > m$ by indifference. This reflected pause has

$$\Gamma_F = K_F(m)\varepsilon^3 + O(\varepsilon^4), \tag{55}$$

where

$$K_F(m) = \frac{2rf - \beta m v f' + 2\beta r v f^2/(\beta m + r)}{3\beta^2 m^3 v^2 f^3}. \tag{56}$$

Fix an interior pivot p and define

$$\mathcal{Q}_p(m) = \log\left[b(m)^{r/(\beta p)}\{\nu(p) - b(m)\}\right].$$

The inner function $B^{r/(\beta p)}\{\nu(p) - B\}$ has a strict maximum at $B = b(p)$; because $b'(p) \neq 0$, $\mathcal{Q}_p'(p) = 0$ and $\mathcal{Q}_p''(p) < 0$. Hence, for each small left span $\varepsilon > 0$, there is a unique right span $\ell(\varepsilon) > 0$ satisfying $\mathcal{Q}_p(p - \varepsilon) = \mathcal{Q}_p(p + \ell)$. Taylor expansion gives

$$\ell(\varepsilon) = \varepsilon - \frac{\mathcal{Q}_p'''(p)}{3\mathcal{Q}_p''(p)}\varepsilon^2 + O(\varepsilon^3). \tag{57}$$

Substitute (57) into the exact integral (18) and expand. The constant, linear, and quadratic terms cancel by indifference and local reflection. For a directly reproducible intermediate identity, put $a = \nu(p)$ and $b_j = b^{(j)}(p)$. The coefficient is

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

Writing $y = -\nu'/\nu$ gives

$$K_F(p) = \frac{(\beta p + 2r)y(p)^2 - \beta p y'(p) + 2\beta r y(p)/(\beta p + r)}{3\beta^2 p^3}. \tag{58}$$

Since

$$y(p) = \frac{1}{\nu(p)f(\nu(p))}, \qquad y'(p) = \frac{f(v) + vf'(v)}{v^2 f(v)^3},$$

(58) reduces to (56).

For a uniform remainder on a compact interior core contained in the neighborhood where ν is C^4, write

$$\mathcal{Q}_p(p) - \mathcal{Q}_p(p+z) = z^2 A(p,z),$$

where A is positive and C^2 under the maintained C^4 hypothesis. The coordinate $\chi_p(z) = z\sqrt{A(p,z)}$ maps points with the same function value to opposite coordinates: $\ell_p(\varepsilon) = \chi_p^{-1}(-\chi_p(-\varepsilon))$. The inverse function theorem and the uniform strict quadratic maximum give a unique reflected branch. A uniform fourth-order Taylor expansion of the indifference equation then gives

$$\ell_p(\varepsilon) = \varepsilon + c(p)\varepsilon^2 + O(\varepsilon^3)$$

uniformly in p. A second-order Taylor expansion of the exact gain integrand, whose third derivative is uniformly bounded on the compact core, then makes the remainder in (55) uniformly $O(\varepsilon^4)$. A strictly negative K_F therefore produces pauses that save time for all sufficiently small spans, completing part (3).

Learning and impatience. The amount of concentration needed for this effect depends on how much information the agents adopting at closure can produce relative to agents' cost of waiting. Write $\lambda = \beta/r$ for learning relative to impatience, $\eta_f(v) = vf'(v)/f(v)$ for the local elasticity of density, and $H_+(v) = vf(v)/m$ for the payoff-scaled upper-tail hazard at $m = 1 - F(v)$. Condition (25) becomes

$$\eta_f(v) > \frac{2}{\lambda m} + \frac{2H_+(v)}{1+\lambda m}. \tag{59}$$

For a fixed payoff margin in the active range, an increase in λ lowers both terms on the right, so a less pronounced rise in density can support a faster pause: the additional early adopters generate information more quickly relative to the discounting cost of the pause. This interpretation concerns the local test at a given margin. A change in β or r also changes which points lie in the active range and which pauses satisfy indifference, so the inequality does not by itself determine how the optimal calendar changes.

E.3 Welfare comparison

Consider a pause $e = (d, p, h)$ with the same preceding path and the same continuation as the segment of continuous adoption it replaces. Write $D = b(d)$, $Z = b(h)$, $S = \mathcal{S}(d,h)$, and let $C_h \geq 0$ be continuation welfare measured from the common endpoint $(h, b(h))$.

Measured at closure, the welfare change is

$$\Omega_F(e;C_h) = \int_d^p \{\nu(m) - D\}\,dm + e^{-r\Delta}\int_p^h \{\nu(m) - Z\}\,dm$$
$$- \int_d^h e^{-r\mathcal{S}(d,m)}\{\nu(m) - b(m)\}\,dm + \{e^{-r\Delta} - e^{-rS}\}C_h. \tag{60}$$

The first three integrals compare the payoffs of the agents whose adoption dates under continuous access are removed, each of whom joins one of the two pools under the pause. The last term accounts for everyone who adopts after the common endpoint, whose payoffs change only through the date at which the continuation begins. Multiplying by $\mu_0 e^{-rt_d}$ converts the expression to a date-zero welfare change.

Suppose $\Delta \geq S$. A type adopting at closure receives the payoff line at the start of the segment of continuous adoption rather than its own weakly preferred payoff line under continuous access. A type adopting at reopening receives belief Z at date Δ. The endpoint of the continuous-access path offers the same belief at the weakly earlier date S, and this type's own payoff line under continuous access is weakly preferred to that endpoint line. The inequalities are strict on a positive-measure set for a pause. Finally, every continuation adopter has nonnegative net payoff and receives the same belief line weakly later under the pause. Hence $\Omega_F(e;C_h) < 0$ whenever $\Gamma_F(e) \geq 0$: a welfare-improving pause must bring completion forward. Deleting such pauses one at a time preserves disjoint interiors and reconstructs a globally IC calendar. Theorem 2(2) therefore makes the continuous-access path the unique aggregate welfare optimum among outcomes ending with adoption exactly at the target when $f' \leq 0$. For a countable family, remove any one pause first, yielding a strict welfare gain, and then remove the others along nested finite subfamilies. The uniform schedule convergence in Appendix D and boundedness of discounted payoffs let these welfare comparisons pass to the continuous-access limit.

The agreement can fail when a pause is faster. Advancing later adopters' payoffs produces a gain, but this gain need not compensate for the timing options lost within the pools, and the local expansion makes the comparison explicit. Suppose ν is C^4 near $p \in (m_0, m_1)$ and consider the reflected pause with pivot p and left span ε. Let C_p be welfare along the continuation under continuous access, including its terminal pool, discounted from the date of state $(p, b(p))$:

$$C_p = \int_p^{m_1} e^{-r\mathcal{S}(p,m)}\{\nu(m) - b(m)\}\,dm$$
$$+ e^{-r\mathcal{S}(p,m_1)}\int_{m_1}^{\bar{M}} \{\nu(m) - \nu(\bar{M})\}\,dm. \tag{61}$$

Then

$$\Omega_F = r\left\{-K_F(p)C_p - \frac{x'(p)[-\nu'(p)]}{3\beta p}\right\}\varepsilon^3 + O(\varepsilon^4). \tag{62}$$

The term $-K_F(p)C_p$ values the change in the continuation date, and the second term is the option value lost by forcing the nearby agents into the two pools. For a locally faster pause, $K_F(p) < 0$, the cubic continuation gain outweighs the cubic loss exactly when

$$[-K_F(p)]C_p > \frac{x'(p)[-\nu'(p)]}{3\beta p}. \tag{63}$$

Under this strict inequality, sufficiently small reflected pauses improve welfare as well as completion time. Reversing the inequality makes them welfare-reducing, and at equality higher-order terms determine the sign. Section E.4 provides smooth polynomial examples with each strict welfare sign. The dependence on C_p also explains why the principal's additive problem does not extend to welfare maximization: the same change in completion time is more valuable to adopters when the continuation contains more surplus, and that continuation itself depends on the later calendar.

For the local formula, let

$$t'_R(z) = \frac{x'(z)}{\beta z}$$

be time per unit of additional adoption under continuous access, and normalize $t_R(p) = 0$. The payoff to type $\nu(m)$ from the payoff line under continuous access indexed by z is

$$U(z; m) = e^{-rt_R(z)}\{\nu(m) - b(z)\}.$$

Indifference under continuous access gives $U_z(p; p) = 0$. Differentiating that identity along the diagonal, or differentiating U directly and using $r\{\nu(p) - b(p)\} = \beta p b(p)$, gives

$$U_{zz}(p; p) = rt'_R(p)\nu'(p) < 0. \tag{64}$$

The reflected branch satisfies $d = p - \varepsilon$ and $h = p + \varepsilon + O(\varepsilon^2)$. Assigning the left half to the start of the segment of continuous adoption and the right half to its endpoint therefore has aggregate option-value effect

$$\frac{1}{2}U_{zz}(p; p)\left\{\int_0^\varepsilon z^2\,dz + \int_0^\varepsilon z^2\,dz\right\} = -\frac{rx'(p)[-\nu'(p)]}{3\beta p}\varepsilon^3 + O(\varepsilon^4).$$

Replacing the time at the endpoint of the continuous-access path by the time at reopening changes the payoffs of the $O(\varepsilon)$ mass adopting at reopening only at order $O(\varepsilon^4)$, because $\Delta - S = O(\varepsilon^3)$. For the positive-mass continuation,

$$\{e^{-r\Delta} - e^{-rS}\}C_h = -rK_F(p)C_p\varepsilon^3 + O(\varepsilon^4),$$

using $C_h = C_p + O(\varepsilon)$ and (55). Adding the two terms proves (62) and (63).

E.4 Two polynomial-quantile examples

The two signs in (63) occur in smooth, strictly positive-density environments. Set $r = 1$, $\beta = 5$, $\bar{M} = 0.95$, choose $B_0 = b(0.1)$ so that $m_0 = 0.1$, and consider local reflected pauses at $p = 0.3$. A strictly decreasing C^∞ upper-tail quantile on $m \in [0, 1]$ induces the density $f(\nu(m)) = -1/\nu'(m) > 0$.

Write $H_p = x'(p)[-\nu'(p)]/(3\beta p)$. Rounded evaluations of (56) and (61) give the two welfare signs:

Quantile $\nu(m)$	m_1	$K_F(p)$	C_p	H_p	$[-K_F(p)]C_p - H_p$
$5 - 0.1m - 4m^3$	0.43	-0.56	1.24	0.59	0.10
$5 - 0.5m - 4m^5$	0.46	-0.04	1.38	0.31	-0.26

Both quantiles are strictly decreasing and positive. Sufficiently small pauses are faster in both cases, but raise welfare only in the first.

F Observed types and persistent priority

Keep the model's primitives and let the principal choose a closed calendar D and a nondecreasing eligibility cutoff $e(t)$. At a date $t \in D$, an agent with payoff $\nu(m)$ can adopt if $m \le e(t)$, with guaranteed access. Thus eligibility expands from higher to lower payoffs, and admission preserves every later opportunity. Every finite-completion equilibrium under such a policy can be matched or improved by an access calendar ending exactly at the target.

Proof. If a lower type adopts before a higher type, both can choose both dates: upper-tail eligibility gives the higher type access to the earlier date, and persistence gives the lower type access to the later one. Strict single crossing contradicts their choices. Participation and the same eligibility comparison ensure that adopters form an upper interval. Cap eligibility at $\bar{M}$ and delete dates after first attainment. This leaves earlier choices and beliefs unchanged, trims the final mass to $\bar{M}$, and preserves every forward comparison $L_{\tau(m)}(\nu(m)) \ge L_{\tau(n)}(\nu(m))$ for $n > m$.

Write $H = \nu(\bar{M}) = b(m_1)$ and define $\mathfrak{m}(B) = b^{-1}(B)$ for $B \in [H, B_0]$. For a no-adoption interval with beliefs $D > Z \ge H$ and stock p, limiting forward incentives at its endpoints give

$$\nu(p) - D \ge e^{-r\Delta}\{\nu(p) - Z\} > 0, \qquad \Delta = \frac{\log(D/Z)}{\beta p}.$$

Equivalently, $K(p; D, Z) \ge 0$, where

$$K(p; D, Z) = \frac{r}{\beta p} \log \frac{D}{Z} - \log \frac{\nu(p) - Z}{\nu(p) - D}.$$

This function is strictly decreasing in p on its participation domain. Lemma A.2 supplies its unique zero $P(D,Z)$, with $\mathfrak{m}(D) < P(D,Z) < \mathfrak{m}(Z)$; hence $p \leq P(D,Z)$. Replacing the interval by the corresponding pause therefore accelerates its belief change. Endpoint constraints remain valid when endpoints carry no atoms: take adopters with payoffs approaching $\nu(p)$ from above and used dates approaching the other endpoint, and pass to the limit.

Elsewhere in this belief range, $M_t \leq \mathfrak{m}(B_t)$ at every right accumulation point of adoption support. Otherwise $(r+\beta M_t)B_t - r\nu(M_t) > 0$. Right continuity would make $L_s'(v) > 0$ in a small right neighborhood for all agents with payoffs in a small interval below $\nu(M_t)$. Two adoption dates in that neighborhood, and an adopter from that band at the earlier date, would contradict persistence and forward optimality. Support points without right accumulation are countable; thus the bound holds at almost every remaining belief, including under singular continuous adoption.

Discard initial delay. The belief clock then satisfies $dt = d(-\log B)/(\beta M(B))$. Compare paths up to H, or up to the start of an interval straddling H. On each complete no-adoption interval use stock $P(D,Z)$; on their complement use $\mathfrak{m}(B)$. Both substitutions are weakly faster. Chronological endpoint beliefs order their supports, so Appendix D reconstructs them as an access calendar.

At most one interval straddles H. Let its initial belief, stock, and full duration be $D > H$, p, and Δ_o, and let $u(s) = e^{-rs}\{\nu(p) - De^{-\beta ps}\}$ and $s_H = \log(D/H)/(\beta p)$. The endpoint incentive limit gives $\nu(p) > D$ and $u(\Delta_o) \leq u(0)$. If $u(s_H) \leq u(0)$, the preceding stock bound applies at H, and the replacement ends within $s_H < \Delta_o$. Otherwise unimodality gives a positive return time Δ_p with $u(\Delta_p) = u(0)$ and $s_H < \Delta_p \leq \Delta_o$. Its terminal belief Z_p is below H. The root-bracketing argument and equation (41) show that raising this endpoint from Z_p to H strictly shortens the indifference pause. Thus its replacement again ends before the original interval ends.

Terminal participation ensures the original path reaches H by completion. At H, append the final pool through $\bar{M}$ using Lemma A.3. For countably many intervals, the disjoint supports, summable duration bounds, and uniform reconstruction in Appendix D preserve all incentive inequalities in the limit. The belief-clock comparison remains valid as an improper integral at the initial state. The resulting calendar completes weakly earlier. □

G Numerical example and survey illustration

G.1 The introductory example

Figure 1 in the main paper compares distributions on the same payoff support $[0.1, 5]$, with $r = \beta = 1$, target $\bar{M} = 0.98$, and the same prior. The upper row uses a uniform distribution. The lower row uses a mixture with weight 0.2 on that uniform distribution and weight 0.8 on a normal distribution with mean 2.55 and standard deviation 0.1, truncated to the same support. The prior odds are $B_0 = b(0.1)$ in the concentrated economy, giving an initial adoption mass of 0.1 there. Under this prior, the uniform economy has its own initial adoption mass. The illustrated pause begins from adoption $d = 0.3$ and reopens at the final round, $h = m_1$. The stock p during closure is determined by (16); equations (15) and (18) give its duration and its effect on completion time.

G.2 Survey willingness and the measurement of incentives

The distributional question is visible in survey evidence on willingness to adopt. Figure G.1 compares responses in Japan and the United Kingdom to the same vaccine-willingness statement in the February 2021 Imperial College London–YouGov Behaviour Tracker (Jones et al., 2020). Responses in Japan are relatively diffuse across the five categories. In the United Kingdom, roughly seven in ten respondents strongly agree. The comparison illustrates a difference in how willingness is distributed across a population, beyond the fraction classified as willing.

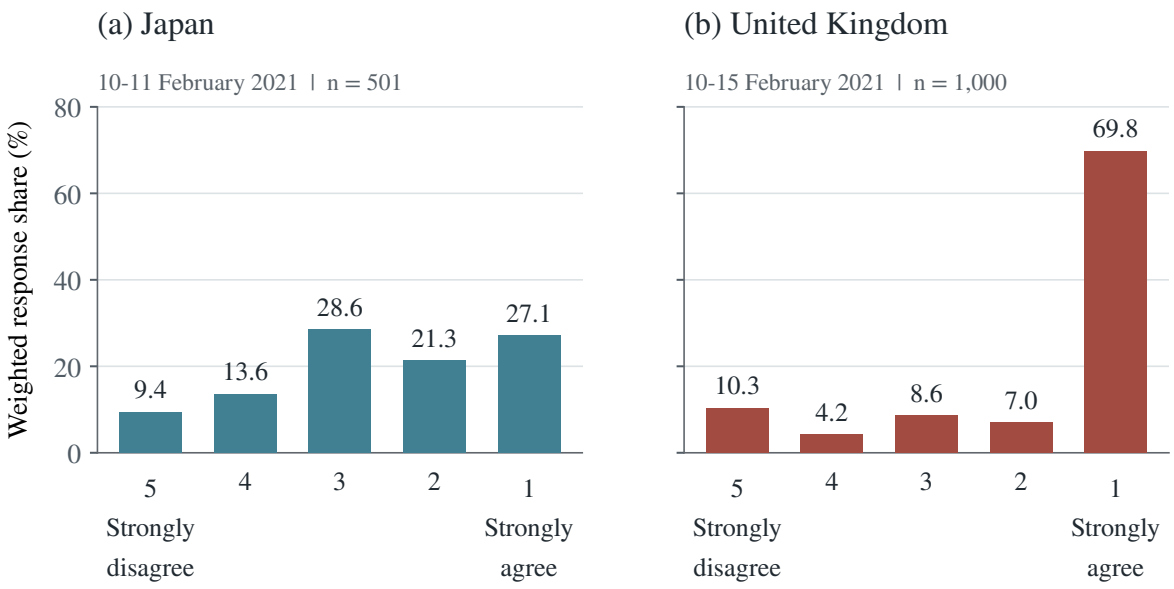


Figure G.1: Weighted responses to "If a Covid-19 vaccine were made available to me this week, I would definitely get it." Country survey rounds fall within the February 2021 report's fieldwork window; all responses are valid. Categories 2–4 retain their original numeric codes because only the endpoints have verbal labels in the source. Source: Jones et al. (2020), item `vac_1`.

These responses illustrate why the distribution of assessments can matter for design. Willingness reflects both beliefs about quality and the value of adoption. Numerical probability questions or stated choices of the kind discussed in Section 5 can supply the corresponding cardinal inputs. Given an estimated distribution, the target, and learning and discount rates, Section 3.4 chooses the calendar. The ordinal categories in this figure motivate that measurement exercise; they are not treated as numerical priors or used to calculate either country's calendar.

References

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

Cason, T. N., S. Hudja, and Y. Rosokha (2025). Early innovation adoption and information free-riding: An experimental investigation. Working paper, November 6.

Centers for Disease Control and Prevention (1999). Withdrawal of rotavirus vaccine recommendation. *Morbidity and Mortality Weekly Report* 48(43), 1007.

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

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

Chen, C.-H., H. Eraslan, J. Ishida, and T. Yamashita (2025). Optimal feedback dynamics against free-riding in collective experimentation. ISER Discussion Paper 1247, revised September 4, 2025.

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

Décamps, J.-P. and T. Mariotti (2004). Investment timing and learning externalities. *Journal of Economic Theory* 118, 80–102.

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

Foster, A. D. and M. R. 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.

Jones, S. P., Imperial College London Big Data Analytical Unit, and YouGov Plc (2020). Imperial College London YouGov Covid Data Hub, version 1.0. YouGov Plc, April 2020. February 2021 survey rounds.

Knoepfle, J. and J. Salmi (2025). Dynamic evidence disclosure: Delay the good to accelerate the bad. Working paper, arXiv:2406.11728, revised November 18, 2025.

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

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 6, 2026.

Larson, H. J., A. de Figueiredo, Z. Xiahong, W. S. Schulz, P. Verger, I. G. Johnston, A. R. Cook, and N. S. Jones (2016). The state of vaccine confidence 2016: Global insights through a 67-country survey. *EBioMedicine* 12, 295–301.

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

Manski, C. F. (2004). Measuring expectations. *Econometrica*, 72(5), 1329–1376.

Méango, R., and E. M. Girsberger (2026). Just ask them twice: Choice probabilities and identification of ex ante returns and willingness-to-pay. *The Economic Journal*, advance publication, ueag018.

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

U.S. Food and Drug Administration (2019). Edwards SAPIEN 3 and SAPIEN 3 Ultra: Expansion to low-risk patients. PMA Supplement P140031/S085, August 16.

Young, H. P. (2009). Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. *American Economic Review* 99(5), 1899–1924.