Online Appendix for Technological Revolutions and Stock Prices Revisited

Roberto C. Gutierrez Jr. Robert C. Ready

Section A gives additional results for the endogenous-adoption model. Sections B–E consider Poisson learning, finite review calendars, heterogeneous conversion costs, and the numerical solution, respectively. Sections F–H give the expected-adoption-time approximation, parameter robustness, and the selection of the lead boundary.

A Additional Perspective on the Model

This section supplements Figure G1 in Online Appendix G and shows why the exceptional calibration that produces bubble-like paths is empirically implausible.

A.1 Adoptions at various times across parameterizations

Figure A1 disaggregates the parameter-robustness results in Figure G1 by true technology quality, $\psi = k\sigma_J$ for $k = 1, \ldots, 6$. Higher-quality technologies are adopted sooner, but expected adoption dates still adjust smoothly and none of the paths exhibits a bubble.

A.2 Producing bubbles with extreme parameterizations of the endogenous model

Under the benchmark of Pástor and Veronesi (2009) and nearby parameterizations, the downward-sloping adoption boundary makes $\mathbb{E}_t[t^{**}]$ adjust smoothly. The associated increase in systematic risk is therefore spread across the revolution rather than concentrated near adoption. A broad parameter search produces bubble-like paths only when ϕ is extreme and learning is exceptionally rapid.

For each price benchmark $j \in \{N, NA, IA\}$, define the time-t constant yield r_t^j by $M_t^j = E_t[B_T^N] \exp[-r_t^j(T-t)]$. These yields summarize current valuations rather than realized returns. The approximation in the main text expresses the new-economy yield as a time-weighted average of the current never-adopt (NA) and immediate-adopt (IA) yields:

$$\int_t^T r_t^N \, ds \;\approx\; \frac{\mathbb{E}_t[t^{**}] - t}{T-t} \int_t^T r_t^{NA} \, ds \;+\; \frac{T - \mathbb{E}_t[t^{**}]}{T-t} \int_t^T r_t^{IA} \, ds,$$

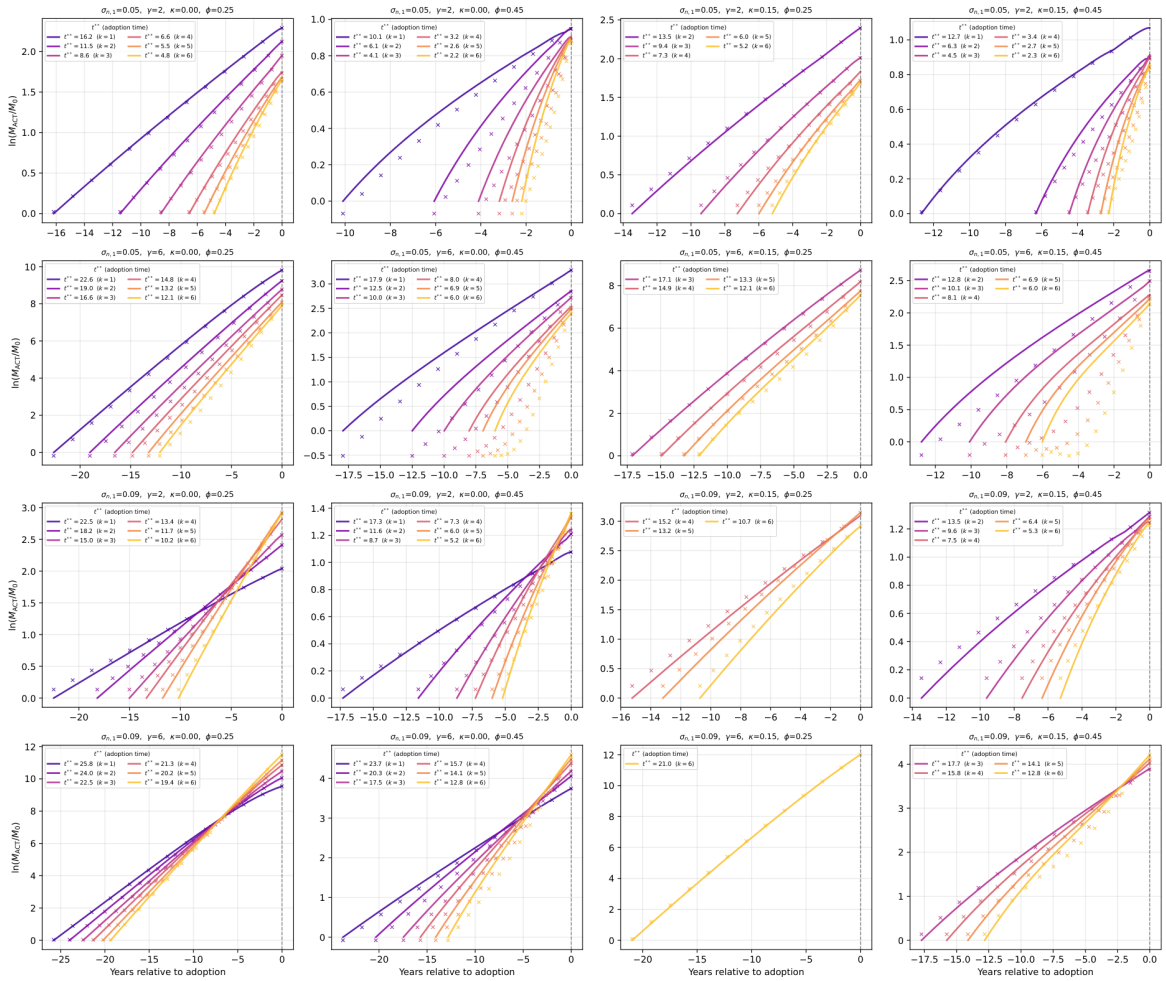


Figure A1: **Cumulative unexpected returns in revolutions across parameterizations: Various adoption times.** Each colored line shows the log price path $\log(M_t^{ACT}/M_0)$ in the endogenous adoption-time model for a different value of the true technology quality $\psi_{\text{true}} = k \cdot \sigma_J$, $k = 1, 2, \ldots, 6$. The 16 panels span the combinations $\sigma_{N,1} \in \{0.05, 0.09\}$, $\gamma \in \{2, 6\}$, $\kappa \in \{0, 0.15\}$, and $\phi \in \{0.25, 0.45\}$; all remaining parameters are set to their baseline values. Paths are deterministic (Kalman filter with no idiosyncratic noise). Vertical markers indicate the adoption time t^{**} for each draw.

which produces the weighted log-price formula

$$\log M_t^N \;\approx\; w_t \log M_t^{NA} \;+\; (1-w_t)\log M_t^{IA}, \qquad w_t = \frac{\mathbb{E}_t[t^{**}]-t}{T-t}.$$

The approximation treats the current IA expected return as a constant yield. In fact, at $s > t$ the NA–IA expected-return gap is $\gamma A_2 (T-s)^2 \hat{\sigma}_s^2/(T-s)$ and declines as learning resolves uncertainty. Later adoption therefore carries a lower IA expected return.

This distinction matters when learning is rapid. Since

$$\hat{\sigma}_t^2 = \left(\sigma_J^{-2} + \frac{\phi^2}{\sigma_{N,1}^2}(t-t^*)\right)^{-1},$$

large ϕ makes uncertainty and the NA–IA return gap collapse early. A fall in $\mathbb{E}_t[t^{**}]$ then shifts adoption toward dates at which $\hat{\sigma}_t$ and the relevant IA yield are unusually high. The constant-yield approximation understates this discount-rate effect, which can dominate the cash-flow news on sufficiently early adoption paths.

Figure A2 shows that high ϕ or high γ alone is insufficient; their combination generates bubbles only on very early adoption paths. Figure A3 shows the corresponding failure of the constant-yield approximation as the IA–NA price gap rapidly narrows.

However, we think that there are two features of this mechanism that limit its empirical relevance. First, the rapid collapse of $\hat{\sigma}_t^2$ produces a pronounced U-shape in return volatility, visible in panel (c) of Figure A3. Annualized return volatility begins near 35%, falls to roughly 5% as the posterior variance collapses, and then rises back to near 35% as the remaining gap Δ_t is traversed near the adoption event. This pattern arises because return volatility through the learning about cash-flows channel is proportional to $\hat{\sigma}_t$, which is quickly falling due to rapid learning. Additionally, positive productivity shocks $d\tilde{Z}_{1,t}$ that push $\hat{\psi}_t$ toward the adoption boundary can generate a large negative discount-rate response by bringing aggregate adoption forward. This response offsets the positive cash-flow effect. It does not require the immediate-versus-never adoption discount $\mathcal{D}(t) = \gamma A_2 (T-t)^2 \hat{\sigma}_t^2$ to rise: that discount declines with time, while news changes how much of it is incorporated into the current price. This channel is eliminated after adoption allowing volatility to rebound quickly (it then falls quickly again with learning). This extreme volatility path, and in particular the large drop in volatility prior to the price drop, is not characteristic of observed technology booms.

Second, the calibration $\phi = 1.50$ is not empirically plausible. The baseline value $\phi = 0.3551$ is estimated from profitability data in Pástor and Veronesi (2006). A value of $\phi = 1.50$ implies a mean-reversion half-life of $\log(2)/1.50 \approx 0.46$ years, meaning that profitability shocks are almost entirely transitory within a single fiscal year. That degree of mean-reversion is far outside the range supported by the data on earnings dynamics and implies an implausibly high level of short-run predictability in firm profits.

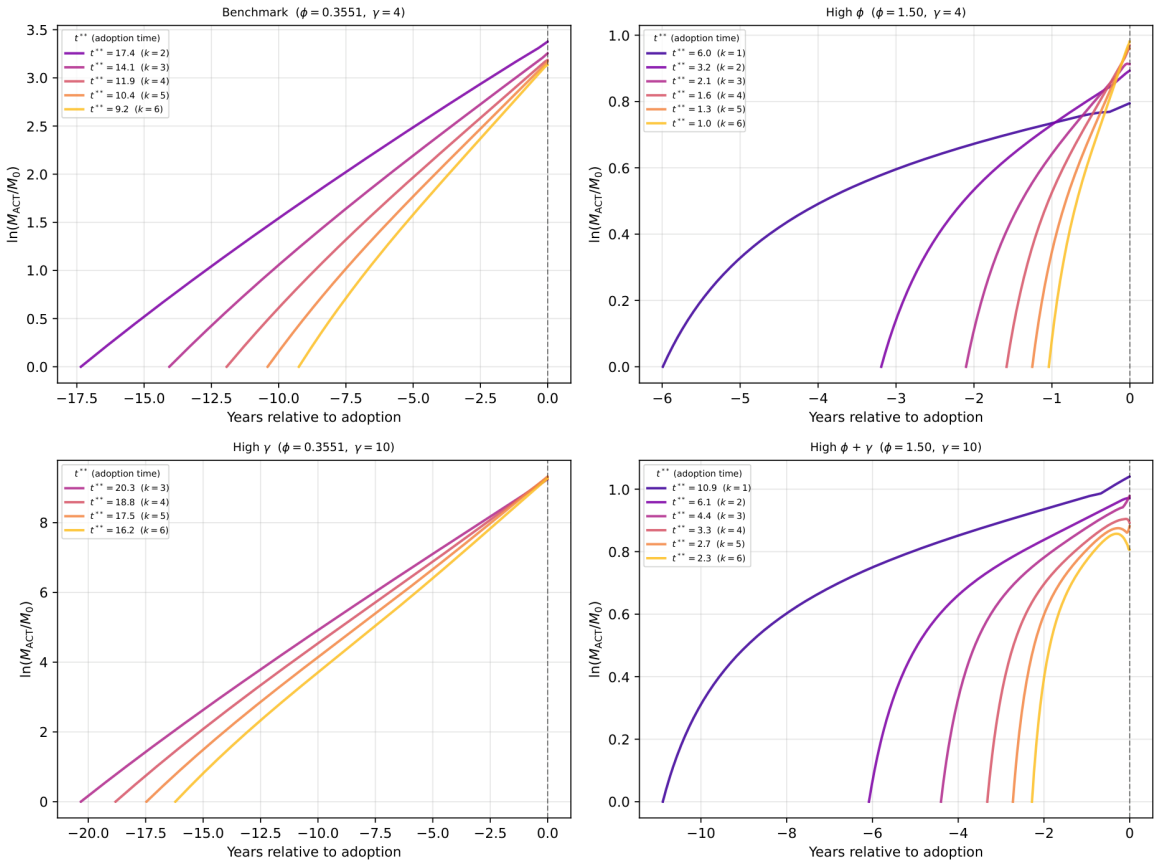


Figure A2: **Deterministic price paths: 2×2 parameterization grid (endogenous model).** Each panel plots $\log(M_t^{ACT}/M_0)$ paths for $\psi_{\text{true}} = k{\cdot}\sigma_J$, $k = 1, \ldots, 6$ (colored lines), as in Figure A1. The four panels vary the speed of mean reversion ϕ and the coefficient of relative risk aversion γ. Top left: baseline ($\phi = 0.3551$, $\gamma = 4$); top right: high ϕ ($\phi = 1.50$, $\gamma = 4$); bottom left: high γ ($\phi = 0.3551$, $\gamma = 10$); bottom right: high ϕ and γ ($\phi = 1.50$, $\gamma = 10$). All other parameters are set as in the benchmark of Figure A1.

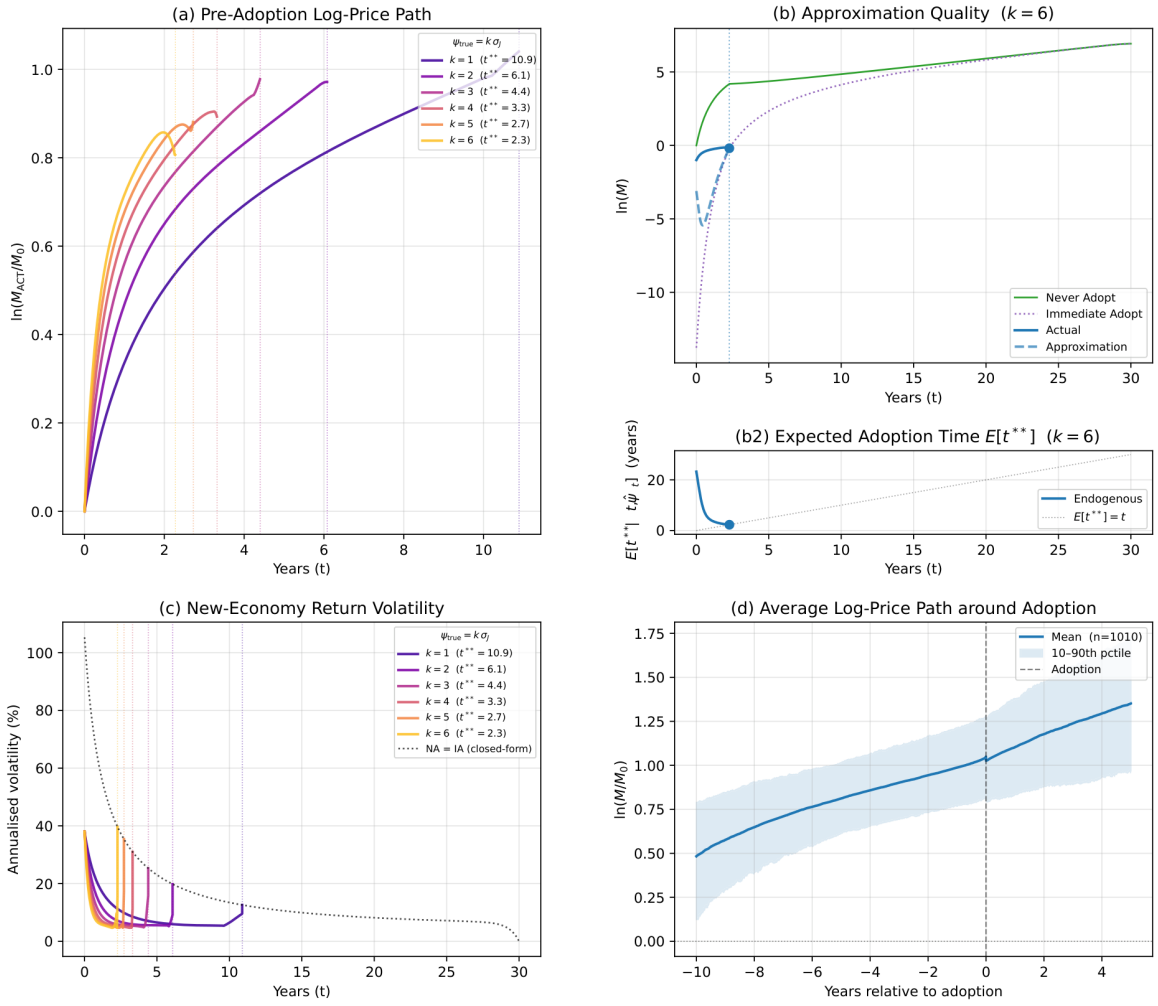


Figure A3: **High-ϕ, high-γ parameterization: price paths, approximation quality, volatility, and Monte Carlo average (endogenous model, $\phi = 1.50$, $\gamma = 10$).** Panel (a): $\log(M_t^{ACT}/M_0)$ for $\psi_{\text{true}} = k \cdot \sigma_J$, $k = 1, \ldots, 6$ (colored lines). Vertical dotted lines mark adoption times t^{**}. Panel (b): actual and approximate log prices for the $k = 6$ path, together with the never-adopt and immediate-adopt benchmarks; the lower subpanel shows the expected adoption time $E_t[t^{**}]$. Panel (c): annualized return volatility for each path. Panel (d): average $\log(M/M_0)$ path aligned to adoption from a Monte Carlo simulation with $N = 5{,}000$ total paths. The mean and 10th–90th percentile band use the $n = 1{,}010$ adopting paths, where n denotes the number of paths that adopt. Adoption cost $\kappa = 0.1$.

B Poisson Technology Shocks and Lumpy Learning

In this section we replace the Gaussian productivity shocks in the endogenous adoption time setting with random Poisson arrivals and find that lumpy information arrival does not restore the PV bubble. A bubble requires the relative strength of the discount-rate effect of favorable news to exceed the associated cash-flow effect. Replacing Gaussian learning with Poisson arrivals makes adoption news more discrete, so that a single signal can move beliefs from well below the adoption boundary to above it and therefore generate a large discount-rate response. But the same signal also contains large cash-flow news. It raises current new-economy productivity and updates the posterior mean of long-run technology quality. We find that, as in the Gaussian model, the cash-flow component dominates for every positive signal that induces adoption by year 8, so all such signals produce positive new-economy returns.

This result shows that the absence of a bubble does not depend simply on Gaussian paths being smooth. The downward-sloping endogenous boundary remains important because the threshold falls as learning resolves uncertainty and the value of waiting declines. Unlike the exogenous take-it-or-leave-it decision, the endogenous rule does not allow a small amount of favorable cash-flow news to generate a disproportionately large change in discount rates. A Poisson jump can make the change in adoption probability abrupt, but a larger positive jump simultaneously produces a larger cash-flow response. Small or zero-surprise arrivals can also induce adoption by lowering posterior variance, but we find that, aside from a terminal fixed-cost effect examined below, these variance-driven crossings likewise do not produce negative returns.

To keep the model as similar as possible to the benchmark in PV, the jump process is calibrated to have the same expected flow of information and the same local productivity variance as the Gaussian benchmark. As the arrival intensity increases, signals become more frequent and smaller, and the Poisson model converges to the Gaussian environment. At low arrival intensities, information arrives infrequently through larger signals. Each arrival changes the posterior mean, reduces posterior variance, and moves current new-economy productivity through the same underlying innovation.

B.1 Gaussian benchmark

Let ψ denote the unknown difference between the long-run productivity of the new and old technologies, and let

$$\widehat{\psi}_t = E_t[\psi], \qquad v_t = \mathrm{Var}_t(\psi).$$

Before aggregate adoption, learning about ψ is generated by the new-economy-specific productivity innovation. In the Gaussian model, posterior variance evolves according to

$$\frac{d}{dt} v_t^{-1} = \left(\frac{\phi}{\sigma}\right)^2, \tag{B1}$$

and the posterior mean has innovation volatility

$$d\widehat{\psi}_t = \frac{v_t\phi}{\sigma}\, d\widetilde{Z}_t^N. \tag{B2}$$

Here ϕ controls mean reversion and σ is the volatility of the technology-specific productivity shock, denoted $\sigma_{N,1}$ in the main text. Thus the local variance and covariance rates of the productivity and belief innovations are

$$\frac{\mathrm{Var}_t(d\rho_t^N)}{dt} = \sigma^2, \qquad \frac{\mathrm{Var}_t(d\widehat{\psi}_t)}{dt} = \left(\frac{v_t\phi}{\sigma}\right)^2, \qquad \frac{\mathrm{Cov}_t(d\rho_t^N, d\widehat{\psi}_t)}{dt} = v_t\phi. \tag{B3}$$

The common productivity shock is unchanged in the exercise below. We replace only the new-economy-specific Gaussian innovation that generates learning before adoption.

B.2 Poisson signals and posterior updating

Signals arrive according to a Poisson process N_t with intensity λ. Conditional on an arrival and the current posterior, the signal is

$$Y_{n+1} = \psi + \varepsilon_{n+1}, \qquad \varepsilon_{n+1} \sim \mathcal{N}(0, R_\lambda), \qquad R_\lambda = \lambda\left(\frac{\sigma}{\phi}\right)^2, \tag{B4}$$

where $n = N_t$ is the number of signals observed by time t. The signal variance rises with λ because a high-arrival-intensity calibration divides a fixed information flow across many relatively uninformative signals. The expected precision flow is exactly the same as in the Gaussian benchmark.

$$\lambda R_\lambda^{-1} = \left(\frac{\phi}{\sigma}\right)^2. \tag{B5}$$

Normal conjugacy gives

$$v_n = \left(v_0^{-1} + nR_\lambda^{-1}\right)^{-1}, \qquad v_{n+1} = \left(v_n^{-1} + R_\lambda^{-1}\right)^{-1}. \tag{B6}$$

Let

$$Z_{n+1} = \frac{Y_{n+1} - \widehat{\psi}_n}{\sqrt{v_n + R_\lambda}} \sim \mathcal{N}(0,1)$$

denote the standardized predictive innovation. The posterior mean update can then be written as

$$\widehat{\psi}_{n+1} = \widehat{\psi}_n + \sqrt{v_n - v_{n+1}}\, Z_{n+1}. \tag{B7}$$

Every arrival therefore reduces posterior variance, while the realization of Z_{n+1} determines whether the posterior mean rises or falls. A zero-surprise signal has $Z_{n+1} = 0$. Such a signal leaves the posterior mean unchanged but still reduces posterior variance from v_n to v_{n+1}.

To preserve the joint cash-flow and learning content of the Gaussian shock, the same

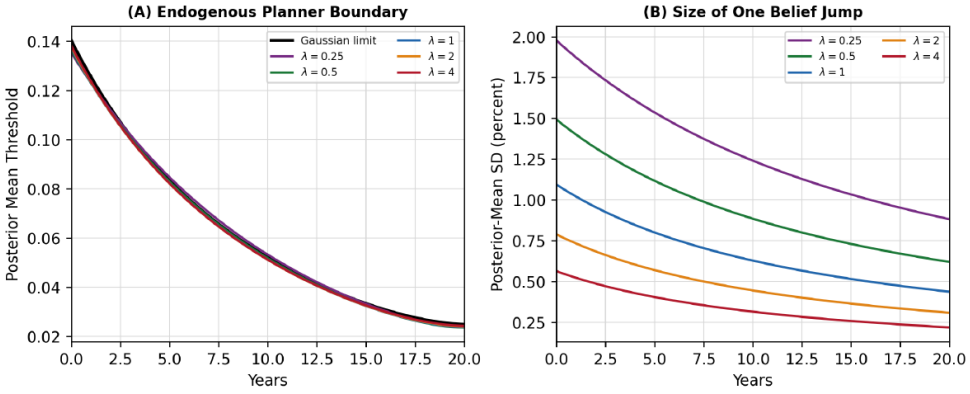


Figure B1: **Poisson learning and adoption boundaries.** Panel A plots the Gaussian planner boundary and Poisson-average planner boundaries for alternative arrival intensities λ. Panel B plots the posterior-mean standard deviation associated with one signal arrival. The Poisson-average boundary at each date averages count-specific boundaries using the time-t Poisson distribution of the number of signals.

standardized innovation changes current new-economy productivity by

$$\Delta\rho^N_{n+1} = \frac{\sigma}{\sqrt{\lambda}} Z_{n+1}. \tag{B8}$$

This specification matches the Gaussian productivity variance rate exactly,

$$\lambda \operatorname{Var}(\Delta\rho^N_{n+1}) = \sigma^2.$$

Moreover, as λ becomes large,

$$\lambda(v_n - v_{n+1}) \longrightarrow \left(\frac{v_n\phi}{\sigma}\right)^2, \qquad \lambda \operatorname{Cov}(\Delta\rho^N_{n+1}, \Delta\widehat{\psi}_{n+1}) \longrightarrow v_n\phi.$$

The compound-Poisson process therefore converges to the Gaussian filter in equations (B2)–(B3). Lower values of λ preserve the same expected information and variance flows but concentrate them in fewer, larger innovations.

Figure B1 shows the resulting planner boundaries. Because the expected information flow is held fixed, the boundaries averaged over the Poisson distribution of signal counts remain close to the Gaussian boundary. The main change is at the path level. Between arrivals the signal count and posterior variance are constant. When a signal arrives, the relevant boundary changes from $b_n(t)$ to $b_{n+1}(t)$ at the same time that beliefs and current productivity respond to Z_{n+1}.

B.3 Adoption and pricing

The planner's stopping problem is solved on the state $(t, n, \widehat{\psi})$. Let

$$\theta = 1 - \gamma, \qquad A_1(\tau) = \frac{1 - e^{-\phi\tau}}{\phi}, \qquad A_2(\tau) = \tau - A_1(\tau), \qquad \tau = T - t.$$

The log certainty-equivalent value of immediate adoption, after removing terms common to adoption and continuation, is

$$S_{t,n}(\widehat{\psi}) = \log(1 - \kappa) + A_2(T - t)\widehat{\psi} + \frac{\theta}{2}A_2(T - t)^2 v_n. \tag{B9}$$

Let $V_{t,n}(\widehat{\psi})$ be the planner's optimal log certainty-equivalent value after the same normalization, so $V_{t,n} = \max\{S_{t,n}, C_{t,n}\}$; never adopting has normalized value zero. Over an interval of length Δt, let K be the number of signal arrivals. The continuation value is

$$C_{t,n}(\widehat{\psi}) = \frac{1}{\theta}\log\left[\sum_{k=0}^{\infty}\Pr(K = k)E\left[\exp\left\{\theta V_{t+\Delta t,n+k}\left(\widehat{\psi} + \sqrt{v_n - v_{n+k}}Z\right)\right\}\right]\right], \tag{B10}$$

where Z is standard normal. The planner adopts when the stopping value is at least as large as the continuation value. The count-specific boundary $b_n(t)$ is therefore defined by

$$S_{t,n}(b_n(t)) = C_{t,n}(b_n(t)). \tag{B11}$$

Asset values are computed from the corresponding pricing moments. Let

$$D_t = E_t[W_T^{-\gamma}], \qquad N_t^N = E_t[W_T^{-\gamma}B_T^N]. \tag{B12}$$

The log value of the new market can be written in the PV log-linear form as

$$\log M_t^N = \log N_t^N - \log D_t + A_1(T - t)\rho_t^N + \log B_t^N. \tag{B13}$$

The pricing recursion integrates jointly over equations (B7) and (B8). In particular, the contemporaneous response to a signal contains the direct productivity term

$$A_1(T - t)\Delta\rho_t^N = A_1(T - t)\frac{\sigma}{\sqrt{\lambda}}Z, \tag{B14}$$

in addition to the effect of the posterior-mean update on expected future growth. Omitting equation (B14) incorrectly treats the arrival as pure information that changes the adoption decision without changing current technology cash flows.

Figure B2 reports a representative large positive signal that induces adoption at year 8 in the low-intensity calibration $\lambda = 0.25$. The signal raises both current productivity and the posterior mean. In this path, the direct productivity contribution is 0.291 log points, total cash-flow news is 0.459 log points, and the adoption and discount-rate component is -0.049 log points. The resulting new-market return is positive, at 0.409 log points. The simulation contains 162 signal-induced crossings by year 8. Of these, 155 are caused by

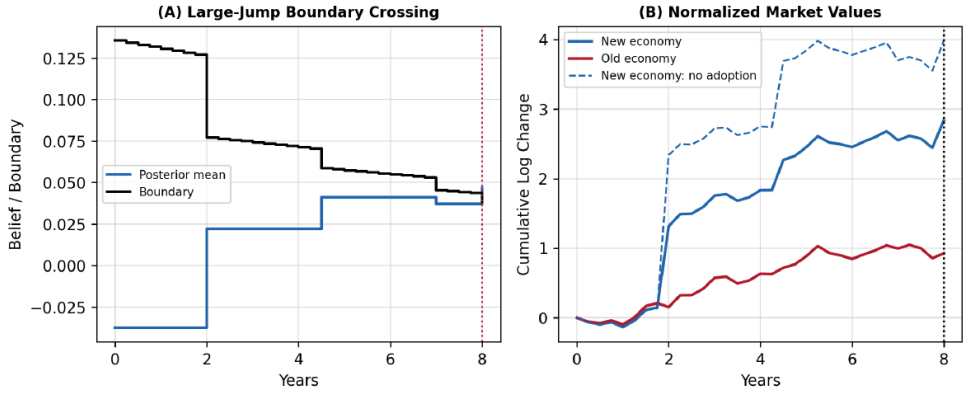


Figure B2: **A large Poisson signal and market values.** Panel A plots the posterior mean and the count-specific planner boundary along a representative path that adopts at year 8. Panel B plots cumulative log changes in new- and old-economy market values and the corresponding no-adoption new-economy value. The vertical dotted line marks the realized adoption time.

positive productivity shocks, and every one of those 155 crossings produces a positive new-market return. The remaining seven crossings follow negative productivity shocks: six have negative returns and one has a positive return. The minimum return among the positive-shock crossings is 0.042 log points and the median is 0.464 log points. Restricting the sample to crossings strictly before year 8 does not change this result.

The six negative returns in the unrestricted crossing sample all follow negative productivity shocks. In these observations, the signal lowers both current productivity and the posterior mean, but the accompanying reduction in posterior variance lowers the adoption boundary by enough to induce adoption nonetheless. These observations are therefore adverse-news returns rather than bubbles generated by positive technology news. They are also distinct from the late-horizon fixed-cost effect considered below. In the calibration, a positive signal large enough to move beliefs rapidly through the boundary always contains enough cash-flow news to dominate the associated discount-rate effect.

B.4 Zero-surprise learning events

A Poisson arrival can affect adoption even when it contains no directional news about technology quality. Consider a zero-surprise signal, $Z = 0$. From equations (B7) and (B8),

$$\Delta\widehat{\psi} = 0, \qquad \Delta\rho^N = 0, \qquad \Delta v = v_{n+1} - v_n < 0. \tag{B15}$$

If

$$b_{n+1}(t) < \widehat{\psi} < b_n(t), \tag{B16}$$

the signal induces immediate adoption entirely through learning. Beliefs do not change, but the reduction in uncertainty lowers the adoption boundary through the interval containing current beliefs.

This experiment distinguishes posterior uncertainty from stock-return risk. Posterior

variance necessarily falls at the signal arrival. The conditional variance of new-market returns, however, need not fall. Before adoption, signals move both expected cash flows and the expected timing of adoption. The discount-rate response associated with adoption timing partially offsets the cash-flow response. Once adoption occurs, the stopping-time channel disappears and the new-market price loads more directly on subsequent technology shocks. As a result, the posterior variance of ψ can fall while the conditional variance of new-market returns rises.

For example, in the $\kappa = 0.10$ calibration at year 8, a zero-surprise signal reduces posterior variance from 0.000972 to 0.000813 and moves the relevant boundary from 0.0596 to 0.0501. Setting beliefs between these two boundaries induces adoption without any change in the posterior mean or current productivity. The new-market return is positive, at approximately 7.3 percent. At the same time, the annualized quadratic variation attributable to Poisson signals rises from approximately 0.049 to 0.065. Thus the conditional variance of the latent technology parameter falls, but the conditional variance of the new market rises when the adoption-timing discount-rate channel is shut down.

Figure B3 compares these zero-surprise value and variance effects for $\kappa = 0.10$ and $\kappa = 0$. With a positive fixed adoption cost, some zero-surprise events generate negative new-market returns late in the horizon, near the inflection in the adoption boundary created by the terminal condition. These negative returns are not a robust consequence of lumpy learning. When $\kappa = 0$, zero-surprise returns are nonnegative up to numerical precision throughout the same interval. Removing the fixed cost eliminates the terminal pull on the boundary and the associated late-horizon negative returns.

The negative zero-surprise observations with $\kappa > 0$ are therefore best interpreted as an artifact of combining a finite terminal date with a fixed level cost of adoption. They do not provide a general jump-learning mechanism for restoring the bubble. Away from this terminal feature, positive Poisson signals produce positive returns, and pure variance-reducing arrivals do not produce the sustained negative price response required by the PV mechanism.

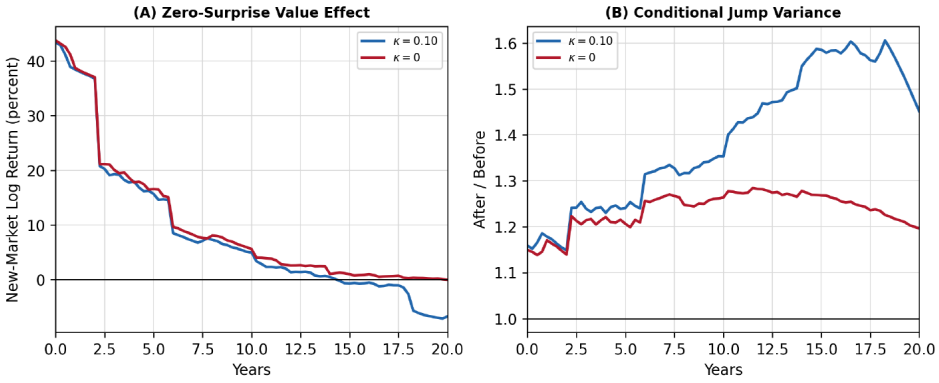


Figure B3: **Zero-surprise learning events and the fixed adoption cost.** Panel A plots the new-market log return produced by a signal that leaves the posterior mean and current productivity unchanged but reduces posterior variance enough to induce adoption. Panel B plots the ratio of conditional Poisson-jump quadratic variation after adoption to that before adoption. The two lines compare $\kappa = 0.10$ with $\kappa = 0$.

C Finite calendar of review dates

Section 2 finds a price bubble when adoption is decided at one exogenous date, and no bubble when adoption timing is chosen optimally. This appendix considers calendars of review dates that bridge these extremes and reports the supporting results. A later review removes the reversal at year eight, and the reversal reappears at whichever review is the last one.

C.1 Model

Everything is as in Section 2 except the timing of the adoption decision. Let $T_1 < T_2 < \cdots < T_K < T$ be the review dates. At each T_k the planner either adopts, converting all capital at cost κ as before, or waits until T_{k+1}. After T_K the technology is abandoned. With $K = 1$ and $T_1 = \bar{t}^{**}$, this is the exogenous-date model, and as the spacing between reviews shrinks to zero, this converges to the endogenous-timing model. Write $\mathcal{V}(t, \hat{\psi})$ for the planner's continuation value and $\mathcal{V}^A(t, \hat{\psi})$ for the value of adopting now. The threshold $\bar{\psi}_k$ at review k makes the planner indifferent between adopting and waiting,

$$\begin{aligned} \mathcal{V}^A(T_k, \bar{\psi}_k) &= E\big[\mathcal{V}(T_{k+1}, \hat{\psi}_{T_{k+1}}) \,\big|\, \hat{\psi}_{T_k} = \bar{\psi}_k\big], \\ \mathcal{V}(T_{k+1}, \hat{\psi}) &= \max\big\{\mathcal{V}^A(T_{k+1}, \hat{\psi}),\ E\big[\mathcal{V}(T_{k+2}, \hat{\psi}_{T_{k+2}}) \,\big|\, \hat{\psi}_{T_{k+1}} = \hat{\psi}\big]\big\}, \end{aligned} \tag{C1}$$

where the never-adopt value closes the recursion at T_K and the belief $\hat{\psi}$ evolves between reviews as in Section 2. This is the endogenous-timing problem restricted to the calendar, solved by the same backward induction with the adoption decision available only on review dates. Prices are computed exactly as in Section 2, with the switch to the adopted economy imposed on the review dates. At the last review, the threshold is the exogenous-date threshold with the remaining horizon. At every earlier review, the threshold is higher by the option value of the reviews still to come. At the baseline $\sigma = 0.07$, the year-eight threshold is 0.040 when that review is the only one, 0.063 to 0.068 when a second review follows at year nine to twelve, 0.064 to 0.073 under calendars that review every four years to every quarter, and 0.074 under endogenous timing. The prior mean is $\mu_J = 0$ throughout, and every calendar is simulated on one million paths.

C.2 The year-eight cohort

Table C1 and Figure C1 evaluate each calendar on the paths that adopt at the year-eight review, in event time. That is the timing of the exogenous-date benchmark, so the horizon and the residual uncertainty are held fixed and only the calendar varies. Under the single review, the new-economy price peaks a year before adoption, falls in every month after, and loses 6.0% in the final month alone, 18.7% from the peak. With a second review at year nine, ten or twelve, or with reviews every two or four years, the price rises in every month into the year-eight adoption, gaining 0.9% to 4.3% in the final month.

Two forces produce the rise. The first is the pricing of the second chance. Under a single review, the value of the claim on the review date steps at the threshold from M^{NA} just below to M^{IA} just above, with the full adoption discount $\mathcal{D}(8) \equiv \gamma A_2 (T-8)^2 \hat{\sigma}_8^2 = 1.77$. With another review at year twelve, adoption at that later date is highly likely just below

Table C1: **The year-eight adoption cohort under review calendars.**

Review dates	year-eight threshold	final month	peak to -1
$\{8\}$	0.040	-6.0%	-18.7%
$\{8, 9\}$	0.068	$+4.3\%$	0.0%
$\{8, 10\}$	0.066	$+3.2\%$	0.0%
$\{8, 12\}$	0.063	$+1.3\%$	0.0%
every 4 years	0.064	$+0.9\%$	0.0%
every 2 years	0.068	$+2.4\%$	0.0%

Notes. $\sigma = 0.07$, $\mu_J = 0$. Paths adopting at the year-eight review, aligned in event time on the last pre-adoption monthly observation ($h = -1$). "Final month" is the mean log change of the new-economy price over the final month before adoption, from $h = -2$ to $h = -1$, which excludes the switch to the adopted economy at $h = 0$. "Peak to -1" is the change from the mean path's pre-adoption peak to $h = -1$, and is 0.0% when the path peaks at adoption. All entries are significant at $|z| > 3$. Cohorts of 2,300 to 19,600 paths.

the year-eight threshold, so the continuation value already reflects much of the systematic risk associated with later adoption. A rough illustration compares the adoption discounts at the two dates: $\mathcal{D}(8) = 1.77$ and $\mathcal{D}(12) = 0.99$. Their difference, 0.78, illustrates the smaller remaining repricing, although it is not the exact year-eight price gap, which depends on the continuation value. The pricing effect can be isolated in the simulations. Holding the year-eight threshold at its single-review value of 0.040 and adding only the year-twelve review shrinks the final-month price fall from 6.0% to 0.5%.

The second force is selection. The option to wait until year twelve raises the optimal year-eight threshold to 0.063, changing the set of paths that adopt at year eight. The paths that clear this higher threshold have experienced stronger favorable news and exhibit a steeper price run-up. Together, the pricing effect and the change in the adopting cohort turn the price fall into the rise shown in Figure C1.

C.3 All adopting paths in calendar time

A path below the threshold at year eight may adopt at a later review, and Table C2 follows the price fall there, reporting the mean new-economy price change across all adopting paths in two windows around each calendar's last review. The final month before the review is the window of Table C1 and ends before the switch to the adopted economy, so under the single review it repeats that table's 6.0 percent fall. The review month contains the switch on the paths adopting there, and the same paths lose 13.8 percent. The two windows agree on where the fall lives and differ only in how much of the repricing they contain.

Three facts stand out. First, the price reversal is at the last review and at no earlier one. The earlier review months are not in Table C2 because there is no fall there to report. Under $\{8, 10\}$, the price gains in the year-eight review month, even on the paths that adopt there, because the paths that wait carry a priced chance at year ten. The rise of those year-eight adopters into their review is the $\{8, 10\}$ line of Figure C1. Second, the number of reviews does nothing. A single review at year ten, two ending there, three, or nine all produce the

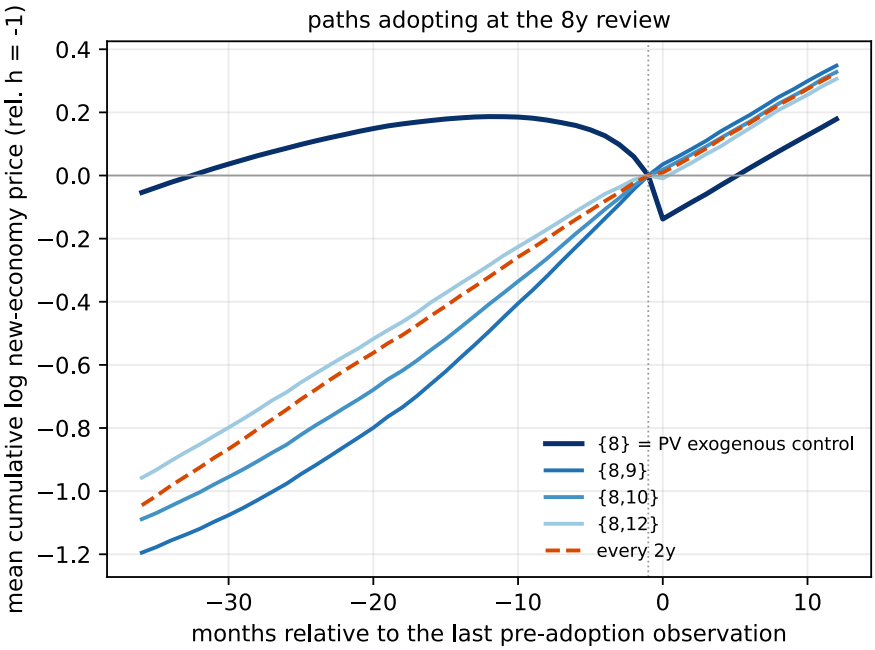


Figure C1: **New-economy prices into a year-eight adoption under review calendars.**

Mean cumulative log price of the new-economy claim for paths adopting at the year-eight review, relative to its value at the last pre-adoption observation ($h = -1$), in months around adoption. The dark line is the single year-eight review of Section 2. The lighter solid lines add one later review at year nine, ten or twelve, and the dashed line reviews every two years. $\sigma = 0.07$, $\mu_J = 0$, each cohort above 2,000 paths. Normalizing at $h = -1$ places the benchmark's trough and the calendars' peaks at the same point, so the figure understates the benchmark's earlier run-up relative to the calendars'.

Table C2: **The fall lives at the calendar's last review.**

Review calendar	last review s	final month before	review month	$\mathcal{D}(s)$
$\{8\}$, $\{4,8\}$, or $\{2,4,6,8\}$	8	-6.0%	-13.8%	1.77
$\{10\}$, $\{8,10\}$, or $\{8,8.25,\ldots,10\}$	10	-2.3%	-6.7%	1.34
$\{8,10,12\}$	12	-0.5%	-2.9%	0.99
$\{8,10,12,14\}$	14	$+0.5\%$	-0.8%	0.71
$\{8,10,12,14,16\}$	16	$+1.1\%$	$+0.3\%$	0.48
every two years to T	none	$+1.8\%$	$+1.8\%$	—
endogenous timing	none	$+1.8\%$	$+1.9\%$	—

Notes. Mean log change of the new-economy price, in percent, across every path that adopts at some review by year sixteen, in two windows around the calendar's last review. "Final month before" ends at the last pre-review observation and excludes the switch to the adopted economy; it is the window of Table C1, and under $\{8\}$ it is that table's -6.0 percent on the same paths. "Review month" contains the switch on the paths adopting there. The final two rows have no last review and report the year-eight review; every review month through year twelve is likewise positive. $\mathcal{D}(s)$ is the adoption discount, the log gap between M^{NA} and M^{IA} for an adoption at s. $\sigma = 0.07$, $\mu_J = 0$, one million simulated paths per calendar, standard errors below 0.3 percent. Adopting paths number 19,600 under $\{8\}$, 47,800 under $\{8,10\}$ and 151,000 under the five-date calendar.

same 6.7 percent fall at year ten, and $\{4,8\}$ and $\{2,4,6,8\}$ reproduce the single review's 13.8 percent fall at year eight, with almost no path adopting at the earlier dates since the option value of the year-eight review puts their hurdles at 0.078 to 0.109. Earlier reviews remove the paths that adopt before the last date and leave the step at that date intact. Third, the last date does everything. Across calendars that review every two years from year eight and end at years ten through sixteen, the review month changes by -6.7, -2.9, -0.8 and $+0.3$ percent, tracking the remaining discount $\mathcal{D}(s)$ of 1.34, 0.99, 0.71 and 0.48. The final month before the review fades a step sooner, positive by year fourteen, because with a later last review more of the smaller remaining discount is already in the price. Almost all of that decline is the shorter remaining horizon under $T = 30$ rather than resolved uncertainty, since $\hat{\sigma}_s$ falls only from 0.034 to 0.031 between years ten and sixteen. A longer-lived technology would carry the fall to later dates.

What removes the fall at an earlier review is that a near miss there is very likely to adopt at the next one, whose hurdle is lower as the horizon shortens and uncertainty resolves. The hurdle turns back up only in the final years, where too little horizon remains for the productivity advantage to repay the fixed conversion cost κ, and by then the remaining discount is too small to matter. Calendars that run to the end of the horizon and the endogenous-timing model show no reversal at any review because no review is the last one.

D Partial adoption with endogenous timing

This appendix allows old-economy firms with different conversion costs to choose when to adopt. Firms face no costs enforcing coordination, and waiting remains an available choice at every decision date. Adopted firms inherit the new economy's additional productivity shock, dZ_1, so that partial adoption does not reveal technology quality and leaves public learning unchanged. Market prices then align firms' adoption incentives with the planner's, as in the static partial-adoption case.

D.1 Why partial adoption requires an information specification

Let unadopted old-economy productivity satisfy

$$\mathrm{d}\rho_t^O = \phi(\bar{\rho} - \rho_t^O)\,\mathrm{d}t + \sigma_0\,\mathrm{d}Z_{0,t}. \tag{D1}$$

Under the large-scale adoption specification, a firm adopting at date s keeps the same common shock but changes its long-run productivity to $\bar{\rho} + \psi$. Productivity is continuous at adoption. Its subsequent gap relative to old technology, $g_{s,t} = \rho_{s,t}^A - \rho_t^O$, therefore satisfies

$$\mathrm{d}g_{s,t} = \phi(\psi - g_{s,t})\,\mathrm{d}t, \qquad g_{s,s} = 0, \qquad g_{s,t} = \psi\big(1 - e^{-\phi(t-s)}\big). \tag{D2}$$

If both productivities are observed separately, the common shock cancels and any observation at $t > s$ identifies ψ. Even one separately observable adopter suffices when its productivity has no additional noise. Thus information becomes perfect arbitrarily soon after the first adoption.

To avoid this immediate revelation, adopted old-economy firms inherit the additional shock carried by the new economy. Firms use public information and have no additional private signals about technology quality.

D.2 Costs, cohorts, and the adoption decision

A unit mass of old-economy firms is indexed by $x \in [0, 1]$, with conversion cost

$$\kappa(x) = \bar{\kappa} + \varepsilon(2x - 1), \qquad 0 < \bar{\kappa} - \varepsilon < \bar{\kappa} + \varepsilon < 1. \tag{D3}$$

Adoption destroys the fraction $\kappa(x)$ of the firm's capital, but does not change its productivity level on impact. Subsequent capital grows at the firm's productivity. Adoption is irreversible. Write τ_x for a firm's adoption time, with $\tau_x = \infty$ if it never adopts, and q_t for the adopted mass. Define retained capital across the first q types by

$$K(q) = \int_0^q [1 - \kappa(x)]\,\mathrm{d}x = (1 - \bar{\kappa} + \varepsilon)q - \varepsilon q^2. \tag{D4}$$

The representative investor has terminal utility $u(W_T) = W_T^{1-\gamma}/(1-\gamma)$. The terminal pricing kernel is proportional to $M_T = u'(W_T) = W_T^{-\gamma}$. A firm takes the aggregate adoption policy and this kernel as given, and chooses its stopping time to maximize the market value of its

terminal capital. The planner chooses the entire adoption allocation to maximize expected terminal utility. Prices of terminal-capital claims are computed as $P_t^j = \mathbb{E}_t[M_T B_T^j]/\mathbb{E}_t[M_T]$, using the same normalization as the main text. Both problems include the continuation value of waiting.

Cost ordering supports a cutoff in firm type. For a fixed aggregate policy, the value of a stopping strategy can be written as $aA+B$, where $a = 1-\kappa(x)$, $A \geq 0$ is the priced payoff on paths with adoption, and $B \geq 0$ is the priced payoff on paths without adoption. Immediate adoption has value aA_0. If immediate adoption dominates every waiting strategy for some a, then $a(A_0 - A) \geq B$ for each such strategy. The inequality also holds at any larger a. Consequently, whenever a type adopts, cheaper remaining types also prefer to adopt. The aggregate decision can therefore be represented by a new cutoff $r \in [q, 1]$.

D.3 Adoption-invariant learning

In this specification, the new economy and adopted old-economy firms share both shocks:

$$\mathrm{d}\rho_{s,t}^A = \phi(\bar{\rho} + \psi - \rho_{s,t}^A)\,\mathrm{d}t + \sigma_0\,\mathrm{d}Z_{0,t} + \sigma_1\,\mathrm{d}Z_{1,t}. \tag{D5}$$

The public observes new-economy productivity and the old-technology reference process ρ_t^O throughout the horizon, including after full adoption. The reference continues to follow the old-technology productivity equation regardless of the mass of unadopted firms. Subtracting the reference's innovation from the new economy's gives the signal

$$\mathrm{d}S_t = \phi\psi\,\mathrm{d}t + \sigma_1\,\mathrm{d}Z_{1,t}. \tag{D6}$$

Adopted old-economy cohorts provide the same signal, so their entry leaves learning unchanged. With prior $\psi \sim N(\psi_0, v_0)$, posterior mean m_t and variance v_t satisfy

$$\mathrm{d}m_t = v_t \frac{\phi}{\sigma_1}\,\mathrm{d}\widehat{Z}_t, \qquad v_t^{-1} = v_0^{-1} + \left(\frac{\phi}{\sigma_1}\right)^2 t. \tag{D7}$$

Cohort history has an exact recursive representation in this environment. Let $a_t = (1 - e^{-\phi(T-t)})/\phi$, let B_t^0 be capital under continued old production, and define productivity-adjusted capital by $H_t^0 = B_t^0 e^{a_t(\rho_t^O - \bar{\rho})}$. Define H_t^x similarly for each adopted firm, and set $y_t = \int_{x \leq q_t} H_t^x / H_t^0\,\mathrm{d}x$. Between adoption decisions, all adopted adjusted-capital ratios have the same proportional evolution,

$$\mathrm{d}\log y_t = a_t\phi\psi\,\mathrm{d}t + a_t\sigma_1\,\mathrm{d}Z_{1,t}. \tag{D8}$$

An increase from q to r adds $K(r) - K(q)$ to y. At the terminal date,

$$W_T = B_T^0(1 - q_T + y_T). \tag{D9}$$

At $y = 0$, adopted capital remains zero until the first adoption. The recursive state is therefore (t, m, q, y), with deterministic v_t. This state preserves the capital and productivity consequences of earlier adoption dates.

Proposition D1 (Common optimal stopping conditions)**.** *When public information is independent of adoption and firms have no additional private information, a planner-optimal allocation is supported by firms' optimal stopping decisions under the kernel $M_T = u'(W_T)$. Conversely, an allocation at which every firm optimizes under that kernel is planner optimal. Thus the planner and decentralized problems share optimal adoption boundaries, with a common treatment of indifference.*

Proof. Consider a change in the stopping strategy of an infinitesimal mass of firms. Because the public filtration and each firm's production opportunities are unchanged, its effect on welfare is the expected change in terminal capital weighted by $u'(W_T)$. This is precisely the firm's market-value objective. A profitable firm-level deviation would therefore contradict planner optimality. Conversely, let W_T^* be generated by an allocation of individually optimal strategies and let $\widetilde{W}_T$ be generated by any other feasible allocation. Concavity gives

$$\mathbb{E}[u(\widetilde{W}_T) - u(W_T^*)] \leq \mathbb{E}[u'(W_T^*)(\widetilde{W}_T - W_T^*)] \leq 0. \tag{D10}$$

The last inequality follows by adding the firms' optimality inequalities under the common kernel. Alternative stopping strategies are admissible under the same public filtration, which is the key requirement for this argument. □

D.4 Recursive solution and equilibrium conditions

We work backward from terminal wealth and terminal marginal utility on the state space (t, m, q, y). At each date and public state, a candidate aggregate cutoff r determines the post-decision adopted mass and adjusted cohort capital. Conditional continuation payoffs are integrated using the joint Gaussian transition for beliefs and adjusted capital, retaining the covariance between productivity news, belief updates, and the pricing kernel. The recursion includes the option to wait and adopt at a later date.

Write $s = (m, q, y)$ for the public state at date t. Let $A_t(s; r)$ be the value of immediate adoption per unit of retained capital and let $C_t(x, s; r)$ be type x's value of waiting, including its optimal future adoption option. The individual Bellman step is

$$V_t(x, s; r) = \max\{[1 - \kappa(x)]A_t(s; r),\ C_t(x, s; r)\}. \tag{D11}$$

The cutoff must agree with these choices. Defining $G_t(x, s; r) = [1-\kappa(x)]A_t(s; r) - C_t(x, s; r)$, consistency requires

$$G_t(x, s; r) \geq 0 \quad (q \leq x < r), \qquad G_t(x, s; r) \leq 0 \quad (r < x \leq 1), \tag{D12}$$

with marginal indifference for an interior cutoff. Corner choices use the corresponding one-sided inequalities. We evaluate candidate cutoffs against the continuation policies of all other firms; an individual deviation does not move the aggregate state. The planner instead chooses r to maximize expected continuation utility, updating the same adopted mass and adjusted cohort capital.

An equilibrium satisfies individual Bellman optimality, the cutoff conditions in (D12), and consistency between individual choices and aggregate adoption.

D.5 Adoption and stock-price dynamics

The calibration is $T = 30$, $\gamma = 4$, $\phi = 0.3551$, $\bar{\rho} = 0.1217$, $\sigma_0 = \sigma_1 = 0.05$, and prior standard deviation 0.04. Conversion costs are uniformly distributed between 5% and 15%, and decisions occur annually. We set the initial mean to $\psi_0 = 0.0522$ to target a probability of 0.1% of full adoption by year 10. The planner and decentralized firms use the same initial beliefs. Full adoption is measured as $q \geq 0.999$; initial adoption is endogenous and is included in the figure.

In related work, Gutierrez and Ready (2026) use a version of this framework in a quantitative model of technology adoption, valuations, and disaster risk.

Figure D1 shows mean adoption and stock-price paths conditional on full adoption by year 10. The proposition establishes that decentralized adoption is socially optimal. The figure presents the adoption and price paths for firms and the planner. Both new-economy and aggregate old-economy prices rise smoothly through the adoption wave, with no bubble pattern in the conditional mean paths. The old-economy claim includes adopted and unadopted firms. Gradual adoption spreads the increase in systematic exposure over time while favorable productivity news supports rising prices.

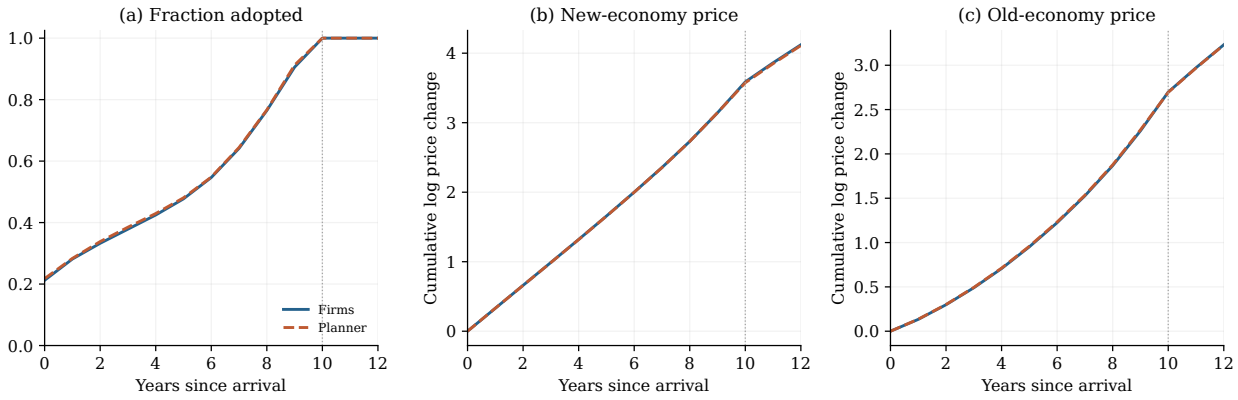


Figure D1: **Endogenous partial adoption and stock prices.** Mean paths conditional on at least 99.9% adoption by year 10 when adopted old-economy firms inherit the new-economy shock dZ_1. Solid and dashed lines show decentralized firms and the planner, respectively. Both use the same initial beliefs; each mean conditions on completion under its own policy. Stock-price changes are measured from date zero. Initial adoption is included in panel (a), and the dotted line marks year 10.

E Numerical Solution of the Decentralized Adoption Boundaries

For the join and planner boundaries, the value-matching condition is a single optimal-stopping problem in $(t, \hat{\psi}_t)$ with exercise coefficient $c = 2\gamma - 1$ (join) or $c = \gamma - 1$ (planner). This coefficient multiplies the uncertainty adjustment in the normalized log exercise payoff, $\log(1-\kappa) + A_2(T-t)\hat{\psi}_t - \frac{1}{2}cA_2(T-t)^2\hat{\sigma}_t^2$. We solve the corresponding free-boundary PDE on a grid that moves with the static threshold; no additional state variable is required.

For the lead equilibrium, an individual firm's incentive to adopt early depends on the entire future aggregate adoption boundary. We therefore evaluate candidate aggregate boundaries rather than obtain the boundary from a single optimal-stopping problem. We restrict candidates to the two-parameter family in equation (H1). For each candidate, we compute the value of adopting immediately and the value of waiting, taking the remaining firms' adoption rule and the resulting pricing kernel as given. The individual continuation problem retains the option to adopt at a later date.

A candidate is admissible if immediate adoption is no more valuable than waiting at every pre-adoption state evaluated on the numerical grid, as required by equation (4). Among admissible candidates, we select the boundary with the highest value at year 15. This criterion selects an equilibrium within the specified family; it does not establish a unique equilibrium or a pointwise highest feasible boundary.

For the bubble exercise, we first solve the individual optimal-stopping problem when technological risk remains idiosyncratic, yielding the $\gamma = 0$ boundary. We then hold that aggregate boundary fixed and evaluate firms' adoption and continuation values under $\gamma = 4$. These values determine the lag penalty and the preemption award described in Section 4.2.1.

F Expected adoption time approximation

This section motivates the expected-adoption-time approximation and explains the probability measures used to summarize fixed and endogenous adoption timing. Exact prices retain state-price weighting in both cases. First, suppose that, conditional on information at time t, adoption is known to occur at a future date $s \in [t, T]$, where $s = T$ denotes never adopting. Write

$$A_t \equiv A_2(T-t), \qquad A_s \equiv A_2(T-s).$$

The terms involving the unknown productivity gain ψ satisfy $B_T^N \propto \exp(A_t\psi)$, while adoption at s contributes $\exp(-\gamma A_s\psi)$ to marginal utility. Since $\psi|\mathcal{F}_t \sim N(\hat{\psi}_t, \hat{\sigma}_t^2)$, the log-affine solution therefore gives

$$\log M_t^N(s) = \log M_t^{NA} - \gamma A_t A_s \hat{\sigma}_t^2. \tag{F1}$$

Defining

$$\mathcal{D}_t \equiv \log M_t^{NA} - \log M_t^{IA} = \gamma A_t^2 \hat{\sigma}_t^2,$$

equation (F1) becomes

$$\log M_t^N(s) = \log M_t^{NA} - \mathcal{D}_t \frac{A_2(T-s)}{A_2(T-t)}.$$

Approximating the exposure ratio by calendar time,

$$\frac{A_2(T-s)}{A_2(T-t)} \approx \frac{T-s}{T-t},$$

and replacing the uncertain adoption date by its conditional mean gives

$$\log M_t^N \approx \log M_t^{IA} + \frac{E_t^{\mathbb{M}}[t^{**}] - t}{T-t}\left(\log M_t^{NA} - \log M_t^{IA}\right), \tag{F2}$$

where $\mathbb{M}$ denotes the measure used to summarize adoption timing. Replacing a random adoption date by its mean also abstracts from the covariation between adoption timing and state-price-weighted payoffs.

To see why, define

$$G_t \equiv E_t[W_T^{-\gamma}], \qquad H_t \equiv E_t[B_T^N W_T^{-\gamma}],$$

so that $M_t^N = H_t/G_t$. Since H_t is a martingale under the physical measure, for any bounded stopping time τ,

$$M_t^N = \frac{E_t^P[H_\tau]}{G_t}. \tag{F3}$$

This identity holds for both fixed and endogenous stopping dates; by itself it does not justify using an unweighted physical mean adoption time. Endogenous adoption occurs at first passage of the boundary $\bar{\psi}(s)$, so $\hat{\psi}_s = \bar{\psi}(s)$ conditional on adoption at s. The component of H_s generated by adoption is then pinned down by s. Let H_s^A and H_s^0 denote the pricing numerators $E_s[B_T^N W_T^{-\gamma}]$ under immediate aggregate adoption at s and a commitment never

to adopt, respectively, evaluated at the same current state. Gaussian integration gives

$$\log \frac{H_s^A}{H_s^0} = -\gamma \log(1-\kappa) - \gamma A_s \hat{\psi}_s + \frac{1}{2}\gamma(\gamma-2)A_s^2\hat{\sigma}_s^2.$$

Using

$$\bar{\psi}(s) = -\frac{\log(1-\kappa)}{A_s} + \frac{1}{2}(\gamma-1)A_s\hat{\sigma}_s^2 + \chi(s)$$

at the first-passage boundary yields

$$\log \frac{H_s^A}{H_s^0} = -\frac{\gamma}{2}A_s^2\hat{\sigma}_s^2 - \gamma A_s \chi(s),$$

which is a deterministic function of the hitting date. Write this ratio as $f(s)$. On paths that hit the adoption boundary, the payoff entering (F3) is $H_{t^{**}}^0 f(t^{**})$. Although the boundary pins down the ratio, $H_{t^{**}}^0$ remains random and can covary with the hitting date. Replacing this weighted expectation by a summary based on $E_t^P[t^{**}]$ is an approximation, not a consequence of the martingale identity. The boundary restriction helps explain why the physical first-passage mean can nevertheless provide a useful summary, as it does in the numerical comparisons.

The fixed-review-date model is different. At the prespecified date $\bar{t}^{**}$, adoption occurs on the event

$$\mathcal{A} = \{\hat{\psi}_{\bar{t}^{**}} \geq \bar{\psi}\}.$$

Conditional on $\mathcal{A}$, the state is not pinned down at a boundary: $\hat{\psi}_{\bar{t}^{**}}$ varies throughout the adoption region. State-price weighting therefore remains within the adoption event. Since $H_{\bar{t}^{**}} = G_{\bar{t}^{**}} M_{\bar{t}^{**}}$,

$$\frac{E_t[\mathbf{1}_{\mathcal{A}} G_{\bar{t}^{**}}]}{G_t} = Q_t(\mathcal{A}),$$

so the natural one-dimensional summary is the risk-neutral probability of adoption. Representing failure to adopt as adoption at T gives

$$E_t^Q[t^{**}] = Q_t(\mathcal{A})\bar{t}^{**} + [1 - Q_t(\mathcal{A})]\, T.$$

Because the fixed-date adoption time has only these two possible values, this mean identifies its entire distribution.

Thus risk-neutral adoption probabilities summarize the state-price weight of the fixed-date adoption event, while physical first-passage times provide a useful approximation for the endogenous model in our numerical results. Neither timing summary alone determines exact prices: the approximation in (F2) also abstracts from payoff variation and its covariation with adoption timing. Exact prices in both models are generated by the same state-price density.

F.1 Price responses to productivity news

Stock Price Response to Productivity Shocks. Now consider the effect of a positive shock $d\tilde{Z}_{1,t} > 0$, which represents good news about the new technology's productivity growth $\hat{\psi}_t$ and productivity ρ_t^N. This shock has two opposing effects, a cash-flow effect and a discount-rate effect.

Cash-Flow Effect (positive): The exposure of M_t^{IA} (the immediately-adopt value of the new-economy) to productivity shocks is given by

$$\frac{\partial \log(M_t^{IA})}{d\tilde{Z}_{1,t}} = A_2(\tau)\frac{\partial \hat{\psi}_t}{d\tilde{Z}_{1,t}} + A_1(\tau)\frac{\partial \rho_t^N}{d\tilde{Z}_{1,t}} \geq 0 \tag{F4}$$

A positive unexpected productivity shock increases both ρ_t^N and $\hat{\psi}_t$ and therefore increases expected terminal book value, raising the price level of M_t^{IA} (and M_t^{NA} equally as well).

Discount-Rate Effect (negative): From the approximation above:

$$\frac{\partial \left(\frac{\mathbb{E}_t[t^{**}]-t}{T-t}A_2(\tau)^2\gamma\hat{\sigma}_t^2\right)}{d\tilde{Z}_{1,t}} = \frac{\partial \mathbb{E}_t[t^{**}]}{d\hat{\psi}_t}\frac{\partial \hat{\psi}_t}{d\tilde{Z}_{1,t}}\frac{1}{T-t}\mathcal{D}(t) \leq 0 \tag{F5}$$

Here the inequality holds since $\mathbb{E}_t[t^{**}]$ decreases with an increase in the subjective belief about productivity. Therefore, through the discount-rate channel, a positive productivity shock decreases the stock price as an earlier adoption time raises the discount rate. The relative size and timing of this negative effect across the exogenous and endogenous models are entirely determined by the behavior of $\frac{\partial \mathbb{E}_t[t^{**}]}{d\hat{\psi}_t}$.

Boundary slope and expected adoption timing. The planner's incentive to wait and the gradual response of expected adoption time are connected through the adoption boundary. In Equation (2), the compensation required for bearing unresolved systematic risk falls as learning reduces posterior variance. Over the economically relevant region, the resulting optimal boundary slopes downward. A belief just below today's boundary therefore faces a lower hurdle later, so missing the boundary need not imply abandoning the technology. The two-date example in Section 2.2 illustrates the same point: just below the initial threshold, later adoption is highly likely, and the continuation value already reflects much of the systematic risk associated with that adoption.

A simple first-passage calculation illustrates the role of boundary slope. For a belief process $X_s = x + \sigma W_{s-t}$ and a linear boundary $b(s) = b(t) - c(s-t)$ with $c > 0$, the expected time to reach the boundary is $(b(t) - x)/c$. Its derivative with respect to the current belief is $-1/c$, independent of the distance to the boundary. Favorable news therefore advances expected adoption without becoming disproportionately influential just before the boundary is reached. This calculation is an illustration rather than an exact formula for our finite-horizon model, whose boundary is nonlinear and whose learning volatility changes over time. The numerical solution exhibits the same smoothing of expected adoption timing and the discount-rate response in Equation (F5), leaving the cash-flow effect dominant in Figure 2.

G Bubbles across parameterizations

To show that the failure to produce a bubble under endogenous timing in Section 2 of the main text generalizes across reasonable parameters, Figure G1 shows the cumulative unexpected new-economy returns under parameter permutations of the most important discount-rate parameters $(\sigma_{N,1}, \gamma, \kappa, \phi)$. This is analogous to panels E and F in Figure 1, with the exception that we focus on the endogenous revolutions that occur in the two-year band around the median endogenous adoption times to obtain a reasonable number of adoptions.[1] As the plot shows, for each specification, there is a clear bubble pattern in the exogenous case, and this bubble pattern is notably more pronounced when risk aversion is high. However, even in high risk-aversion calibrations, there is no bubble pattern in any of the endogenous adoption time cases.

Note that this plot only shows that there is no bubble pattern for adoptions occurring around the median adoption time. However, we also find that this lack of a bubble pattern holds across all adoption times, which, as argued above, is related to the fact that the approximation in equation (3) of the main text is performing well across all these parameterizations.[2]

These results show that the lack of a bubble in the endogenous case holds for a wide set of reasonable parameters. While it is possible to choose parameters to generate a bubble pattern in some subset of adoptions, we find that these parameterizations are likely to be empirically implausible.[3] The systematic-risk channel for price bubbles is not reasonably attainable within the planner's endogenous-adoption problem.

[1] For parameterizations with very few exogenous adoptions after eight years, we move the exogenous adoption time $\bar{t}^{**}$ to 12 years and then 16 years to get enough adoptions. This occurs in all panels in the bottom row.

[2] See Figure A1.

[3] For instance, in Section A.2 we show that with implausibly fast mean reversion of productivity (ϕ), bubble patterns can be generated, but these fast-learning cases exhibit return volatility which is counterfactually collapsing prior to adoption.

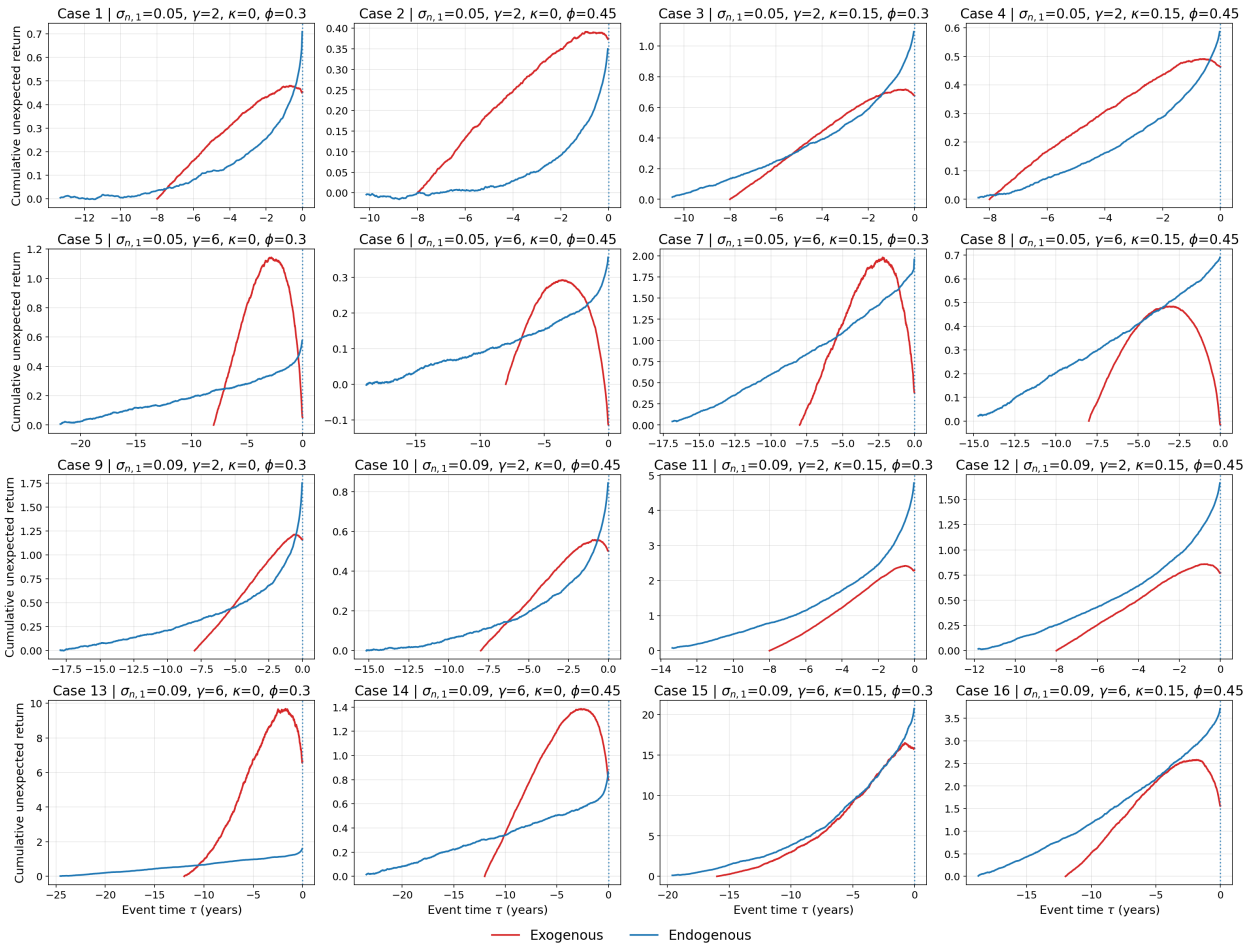


Figure G1: **Cumulative unexpected returns in revolutions across parameterizations**

Each panel shows new-economy cumulative unexpected returns conditional on a technological revolution in the endogenous and exogenous cases for a given set of parameters. Exogenous adoptions occur at time $t^{**} = 8$. For the endogenous adoption time case, we show adoptions in the two-year band around the median adoption time. The 16 panels represent the 16 possible combinations of parameters ($\kappa \in \{0, 0.15\}, \gamma \in \{2, 6\}, \sigma_{N,1} \in \{0.05, 0.09\}, \phi \in \{0.3, 0.45\}$), with remaining parameters set as in the baseline calibration of Pástor and Veronesi (2009). For parameter combinations where very few exogenous adoptions occur after eight years, we move the exogenous adoption decision to year 12 or 16 to generate enough exogenous adoptions.

H Selecting the lead boundary

Unlike the join problem, the feasible lead set does not have a unique pointwise maximum: admissible boundaries can cross because current deviation incentives depend on the future boundary. We therefore restrict the search to the two-parameter family[4]

$$b(t; x_1, x_2) = \frac{-\log(1-\kappa)}{A_2(T-t)} + x_1\hat{\sigma}_t + x_2 A_2(T-t)\hat{\sigma}_t^2 \tag{H1}$$

and select

$$(x_1^*, x_2^*) \in \arg\max_{x_1,x_2} b(15; x_1, x_2) \quad \text{subject to Equation (4).}$$

Maximizing the boundary at year 15 is an equilibrium-selection convention, not an additional equilibrium condition. Other scalar objectives—including the average boundary or its value at another date—select different members of the feasible set. We have verified that the qualitative results in the main text are robust to these alternatives.

[4]This family is flexible enough to fit the numerical planner and join boundaries very closely. The first term is the fixed-cost portion of the exogenous-time threshold when there is no continuation value. The term $x_1\hat{\sigma}_t$ allows an option-value correction: locally, near an exercise boundary with approximately Gaussian uncertainty, the time value of the option is first order in the relevant standard deviation. The $x_2 A_2(T-t)\hat{\sigma}_t^2$ term then allows a variance adjustment as is present in the exogenous threshold, and can provide additional curvature in the dynamic boundary.

References

Gutierrez, Jr., Roberto C. and Robert C. Ready, "Do Markets Discipline the Adoption of Dangerous Technology?," 2026. Working paper, University of Oregon.

Pástor, Ľuboš and Pietro Veronesi, "Was There a Nasdaq Bubble in the Late 1990s?," *Journal of Financial Economics*, 2006, *81* (1), 61–100.

Pástor, Ľuboš and Pietro Veronesi, "Technological Revolutions and Stock Prices," *American Economic Review*, 2009, *99* (4), 1451–1483.