Technological Revolutions and Stock Prices Revisited*

Roberto C. Gutierrez Jr.† Robert C. Ready†

September 2, 2026

Abstract

Pástor and Veronesi (2009) (PV) develop a rational explanation for stock-price bubbles during technological revolutions. We show the bubble to depend on two assumptions: the planner faces a single adoption decision at a prespecified date, and all firms adopt simultaneously. The bubble does not survive optimal adoption timing because the planner prefers to wait until learning reduces technological uncertainty. Under decentralized adoption, firms anticipate that adoption by others makes their own exposure systematic, encouraging them to wait, which again avoids the bubble. Strong coordination and preemption incentives can restore the bubble through a socially inefficient, value-destroying rush to adopt.

*We thank Ľuboš Pástor and Pietro Veronesi for making their replication code available. We also thank Ľuboš Pástor for his constructive comments, in particular the suggestion to analyze decentralized adoption. Generative AI tools were used in the writing of this paper. All AI output was verified by the authors. Any mistakes are our own. Correspondence can be directed to rready@uoregon.edu.

†Lundquist College of Business, University of Oregon.

1 Introduction

Pástor and Veronesi (2009) (hereafter PV) develop a rational explanation for stock price bubbles during technological revolutions. In their framework, a social planner adopts a new technology on an economy-wide scale when its social benefits exceed those of the old technology. Positive shocks to the productivity of the new technology therefore have two opposing effects on the prices of new-technology firms. They raise expected cash flows, which pushes prices up. They also make economy-wide adoption more likely, which raises systematic risk and discount rates. Early in a revolution, the cash-flow effect dominates. Later, as adoption becomes likely, the discount-rate effect can take over. Conditional on a successful revolution, the stock prices of new-technology firms first rise and then fall, creating a "bubble" pattern *ex post*.

The intuitive appeal of this mechanism has made it a natural reference point in recent discussions of generative-AI stocks.[1] Its key implication is that positive news about a new technology can become bad news for stock prices, even though the approaching adoption is socially optimal.

This paper revisits and extends PV's model of technological adoption. We focus on the mechanics of the adoption decision, which converts new-technology risk from idiosyncratic to systematic. Because systematic risk is priced, the new-economy claim is worth less if adoption occurs than if it never does. The PV bubble arises when this decline in value is incorporated into prices abruptly, creating a price crash upon imminent adoption. Two assumptions make it abrupt: the planner faces a single adoption decision at a prespecified date, and all firms adopt at once. We show that relaxing either one is enough to remove the bubble.

Our first contribution concerns the timing of adoption. PV's planner faces a one-time, all-or-nothing decision at an exogenously fixed date, but PV also consider a planner who

[1] The citation for the 2025 Stephen A. Ross Prize in Financial Economics discusses the influence and implications of PV (FARFE, 2025). The Bloomberg article "This is how the AI stock boom plays out" (Wallace, 2025) also notes implications for AI.

chooses the adoption time optimally. While PV conclude a bubble arises in both settings, we show it does so only in the first one. The confusion is due to PV reporting a hump-shaped pattern in the market-to-book ratio, which indeed survives endogenous timing. We show, however, that this pattern reflects expected mean reversion rather than negative stock returns, so that stock prices in the endogenous case rise steadily throughout the revolution without ever falling.

The key feature of the one-time adoption decision is that the prespecified date is the planner's final opportunity to adopt. Falling just short of the adoption threshold therefore means never adopting, so a small amount of positive cash-flow news near the decision date can increase discount rates abruptly and generate a crash. With endogenous timing, failing to adopt today means a delay rather than abandonment. Moreover, the planner has an incentive to delay, because learning reduces the uncertainty that adoption would make systematic. The expected adoption time therefore gradually draws nearer as beliefs improve, with prices adjusting steadily along the way rather than crashing late in the revolution.

Our second contribution is to decentralize the adoption decision. Technological adoption in practice is chosen by firms, not by a planner, and decentralization reveals a feedback between adoption and the nature of technological risk. A firm considering adoption in isolation will value unresolved uncertainty about the new technology because its payoff is convex, as emphasized in Pástor and Veronesi (2006). But this value depends on the risk remaining idiosyncratic. Adoption by other firms, whether current or anticipated, makes the technological shock affect aggregate wealth, so the firm's exposure becomes systematic, making adoption less attractive while uncertainty is high.

To formalize, we first retain PV's fixed decision date but allow firms to adopt separately. An arbitrarily small amount of heterogeneity in conversion costs generates partial adoption. The lowest-cost firms adopt at lower productivity thresholds, at which their exposure remains largely idiosyncratic. As more firms adopt, new-technology risk becomes increasingly systematic, raising the threshold required for each additional firm. Heterogeneity determines the particular firms that adopt, but the endogenous change in the pricing kernel generates a nondegenerate spread in adoption thresholds even as heterogeneity vanishes. The resulting allocation is socially optimal, because the marginal firm's private adoption condition coincides with the planner's condition for the aggregate adoption fraction. Because the fraction of the economy expected to adopt rises smoothly with beliefs rather than jumping between none and all, no single moment carries a large repricing, and the bubble disappears even though the adoption date remains fixed exactly as in PV.

We next ask whether the bubble survives when firms are instead forced to adopt simultaneously. In PV's exogenous-timing model, PV show that the planner's simultaneous-adoption threshold is one of many decentralized equilibria that can be supported by off-equilibrium costs, such as network effects, that penalize firms for moving before or after the rest of the economy. We extend this enforcement approach to endogenous timing. Even with simultaneous adoption enforced, the same feedback operates through expectations. A firm that could adopt in isolation may prefer to move early and capture the convexity benefit of high uncertainty. But an actual first mover anticipates that the rest of the economy may subsequently adopt, making its own exposure systematic and reducing its continuation value. Across the equilibria we study, in which adoption is not deliberately hastened, firms wait for uncertainty to resolve, much as the planner does; expected adoption time responds smoothly to news; and no bubble arises.

The bubble can be restored if firms are forced to coordinate and incentivized not to wait for uncertainty to resolve. Equilibrium multiplicity leaves room to construct such an outcome. The boundary an individual firm would choose in the absence of systematic risk can be supported as a coordinated equilibrium using a penalty for lagging behind the aggregate wave in combination with an award for moving first. Firms then race to adopt as though future aggregate adoption would not make their exposure systematic, even though that systematic risk materializes once the boundary is reached. The sudden transition revives the bubble, but it destroys value. Adoption occurs while substantial uncertainty remains, welfare falls well below the planner benchmark, and old-economy firms are worse off than if they had never adopted. This highlights a tension in the PV bubble: the same systematic risk that causes the crash gives firms an incentive to avoid the adoption that produces it.

Our results connect to three literatures. The first is work on stock-price bubbles, much of which builds on disagreement, speculation, short-sale constraints, or sentiment, as in Harrison and Kreps (1978), Scheinkman and Xiong (2003), and Barberis et al. (1998). Giglio and Severo (2012) link technological change to asset-price bubbles through its effects on intangible capital and financial frictions. PV offer a distinct rational mechanism in which positive news about a new technology raises systematic risk. We show that this mechanism relies on two features, a last-chance opportunity to adopt and a simultaneous economy-wide transition, and that relaxing either removes the price decline necessary to complete the bubble. The paper thereby also connects the bubble literature to work on investment, technological innovation, and asset prices, including Carlson et al. (2004), Garleanu et al. (2012), Kung and Schmid (2015), and Kogan et al. (2020).

The second is the literature on technology adoption under uncertainty and strategic real options, including Fudenberg and Tirole (1985), Reinganum (1981), Grenadier (1996), Grenadier and Weiss (1997), Grenadier (2002), Weeds (2002), and Lambrecht and Perraudin (2003). This literature centers on how preemption, first-mover advantages, coordination, and strategic interaction alter the value of waiting. Closely related is Andrei and Carlin (2023), in which Schumpeterian competition induces excessive experimentation with a new technology, affecting technological uncertainty, systematic risk, and asset prices. Like them, we find that competition can lead firms to bear technological risk in socially inefficient ways. Our mechanism differs in that learning remains exogenous and strategic interaction operates through the price of risk rather than product-market payoffs. What competitors may do in the future determines whether the uncertainty underlying a firm's own adoption decision is idiosyncratic or systematic.

Finally, the analysis speaks to current discussions of AI adoption and valuations, as in

Bonaparte (2024) and Eisfeldt et al. (2023). The PV framework highlights an important consequence of widespread adoption: risk that begins as idiosyncratic becomes systematic as use of the technology broadens. Our results show that this effect also changes firms' incentives to adopt. The prospect of bearing systematic risk encourages firms to delay adoption, smoothing the transition that could otherwise lead to a crash. A PV-style stock price collapse therefore depends on competitive or institutional forces strong enough to overcome these incentives and thereby produce a sudden and unexpected economy-wide adoption wave.

2 Model with socially optimal technological adoption

We first briefly review the model of Pástor and Veronesi (2009) and then focus on the underlying mechanism that generates bubble dynamics in stock prices.[2] The economy has a finite horizon $[0, T]$ with a representative agent (social planner) who has power utility over terminal wealth W_T with risk aversion $\gamma > 1$.

$$u(W_T) = \frac{W_T^{1-\gamma}}{1-\gamma}$$

The agent is endowed with initial capital B_0, which produces output $Y_t = \rho_t B_t$. This output is used to grow the capital stock so that $dB_t = Y_t dt = \rho_t B_t dt$. Shocks to productivity therefore do not impact the current level of capital, but they do impact its future growth. The capital stock is fully consumed at time T so that $B_T = W_T$.

PV assume that the value of new-economy firms is infinitesimal in size relative to the old economy. Hence, the new economy only affects terminal wealth through its impact on the productivity process of the old economy ρ_t, and old economy capital B_T entirely determines terminal wealth.

Technology and Productivity. Initially, only the old technology is available. At time t^*, a new technology becomes available. Old economy productivity ρ_t is a mean-reverting

[2]For details regarding the model, we refer the reader to Pástor and Veronesi (2009) and its technical appendix.

process whose mean may be changed by "adoption" of a new technology at a time $t^{**} \geq t^*$. This adoption increases the long-run mean of the productivity process by an amount ψ, so that

$$\begin{aligned} d\rho_t &= \phi(\bar{\rho} - \rho_t)dt + \sigma dZ_{0,t}, \;\; 0 < t < t^{**}, \\ d\rho_t &= \phi(\bar{\rho} + \psi - \rho_t)dt + \sigma dZ_{0,t}, \;\; t^{**} \leq t < T \end{aligned}$$

Here ϕ is the speed of mean reversion, σ is the exposure to an old-economy productivity shock generated by Brownian increments $dZ_{0,t}$. The new technology's productivity gain ψ is unobservable. When the new technology appears at t^*, ψ is drawn from $N(\mu_J, \sigma_J^2)$ with known variance. After t^*, the new-economy capital stock (B_t^N) and productivity (ρ_t^N) are observable and evolve according to

$$\begin{aligned} dB_t^N &= \rho_t^N B_t^N dt \\ d\rho_t^N &= \phi(\bar{\rho} + \psi - \rho_t^N)dt + \sigma_{N,0} dZ_{0,t} + \sigma_{N,1} dZ_{1,t} \end{aligned}$$

Here $Z_{1,t}$ is a Brownian motion uncorrelated with $Z_{0,t}$, and $\sigma_{N,0}$ and $\sigma_{N,1}$ are the new economy's productivity loadings on the two shocks. By observing ρ_t^N and ρ_t, the agent learns about ψ. The posterior distribution is $\psi | \mathcal{F}_t \sim N(\hat{\psi}_t, \hat{\sigma}_t^2)$, where the posterior mean $\hat{\psi}_t$, conditional on the filtration $\mathcal{F}_t$ generated by the observable productivity levels, is a martingale and the posterior variance $\hat{\sigma}_t^2$ declines deterministically over time with learning. Practically, unexpected shocks to new-economy productivity lead to upward revisions in $\hat{\psi}_t$. The orthogonalized unanticipated shocks to new-economy productivity (controlling for shocks to old-economy productivity) are $d\tilde{Z}_{1,t}$.

Adoption Decision. The planner chooses to adopt the new technology if doing so increases expected utility $\mathbb{E}_t\left[\frac{W_T^{1-\gamma}}{1-\gamma}\right]$. In the *exogenous* adoption time scenario, the planner decides at a prespecified time $\bar{t}^{**}$ whether to adopt the new technology on a large scale.

Adoption occurs if and only if the posterior mean exceeds a threshold:

$$\hat{\psi}_{\bar{t}^{**}} \geq \bar{\psi} = -\frac{\log(1-\kappa)}{A_2(\bar{\tau}^{**})} + \frac{1}{2}(\gamma - 1)A_2(\bar{\tau}^{**})\hat{\sigma}^2_{\bar{t}^{**}} \tag{1}$$

where $\bar{\tau}^{**} = T - \bar{t}^{**}$, κ is a proportional conversion cost which decreases current capital B_t, and $A_2(\tau) = \tau - (1 - e^{-\phi\tau})/\phi$. This $\bar{\psi}$ is the level of subjective belief about the productivity at $\bar{t}^{**}$ for which the agent is indifferent, in terms of expected terminal utility, between the adoption and no-adoption productivity process for the old economy.

In the *endogenous* adoption time scenario, adoption occurs optimally at the time when adoption maximizes the agent's expected utility. As PV shows, this problem is akin to the optimal exercise of an American option, whereby the social planner considers the benefit from adoption relative to the continuation benefit from waiting to adopt. Here there is a threshold $\bar{\psi}(t)$ which now depends on time t (only t since $\hat{\sigma}_t$ is deterministic in time) where the planner adopts if $\hat{\psi}_t \geq \bar{\psi}(t)$. This threshold can be written as the sum of the static adoption threshold for a given time t and a continuation value, so that

$$\bar{\psi}(t) = -\frac{\log(1-\kappa)}{A_2(\tau)} + \frac{1}{2}(\gamma - 1)A_2(\tau)\hat{\sigma}^2_t + \chi(t). \tag{2}$$

Here $\chi(t) \geq 0$ is the continuation option value of waiting to adopt. A closed form for this term is not available and the endogenous adoption boundary is therefore obtained through the numerical solution of the PDEs laid out in PV. We note that $\chi(t) = 0$ when $t = T$ or $\hat{\sigma}^2_t = 0$, so that the value of $\chi(t)$ is generally falling through time as the terminal time approaches, and as more is learned about productivity. The last two terms therefore lead to an adoption boundary which is downward sloping in time, as the uncertainty about the new economy resolves. The first term in contrast creates a positive slope, as the fixed cost of adoption in a finite-horizon model creates a "use it or lose it" effect; there must be enough time remaining to benefit from adoption and justify the fixed cost. Generally, the last term dominates and the adoption boundary is downward sloping in the relevant regions, as we discuss after presenting the baseline specifications of the PV model.

This downward slope will be important for the asset pricing implications of adoption, as it implies that expected adoption time, and therefore systematic risk, responds smoothly to news throughout the revolution.

Asset Pricing. As PV show, the state price density is uniquely given by

$$\pi_t = \frac{1}{\lambda}\mathbb{E}_t[W_T^{-\gamma}],$$

where λ is the Lagrange multiplier from the planner's utility maximization problem. PV assumes that there is a money-market account earning a risk-free rate. Since the model has no intermediate consumption, this is a free parameter and may be normalized to zero. The market values of the old- and new-economy stocks, denoted by M_t and M_t^N respectively, are given by the standard pricing formulas:

$$M_t = \mathbb{E}_t\left[\frac{\pi_T}{\pi_t}B_T\right] \quad \text{and} \quad M_t^N = \mathbb{E}_t\left[\frac{\pi_T}{\pi_t}B_T^N\right],$$

where B_T and B_T^N are the terminal book values (the only cash flows in the model). PV consider market-to-book (M/B) ratios to normalize the market values, where M_t/B_t and M_t^N/B_t^N are old-economy and new-economy ratios respectively. Despite the simplicity of the setup, solving the model is quite involved. Details for the solution can be found in PV and its technical appendix.

Stock price bubbles in the model. Throughout the paper, we use "bubble" to describe an ex post pattern in new-economy prices in which, conditional on adoption by a specified horizon, average prices rise for a period and then fall as adoption approaches. We solve the model using the replication code provided on the *American Economic Review* website under the baseline calibration. We are most interested in how the characteristic bubble pattern in new-economy stock prices, conditional on technological adoption, differs across the exogenous and endogenous adoption time models. Figure 1 plots the relevant results. Following PV, we consider adoptions in the endogenous case that occur within one year of the exogenous adoption time of $\bar{t}^{**}$.

The left-hand panels represent the exogenous adoption time case, and the right-hand panels represent the endogenous case. As Panels A and B show, a technology adoption under either model version is characterized, *ex post*, by a long string of unexpectedly positive productivity shocks. These raise the subjective belief $\hat{\psi}_t$ enough to warrant adoption. This in turn has an impact on stock prices conditional on observing a revolution.

As Panels C and D show, in both the exogenous and endogenous cases, the M/B ratio of the new economy initially rises and falls, generating a hump-shaped pattern. PV reports the plot for the endogenous case (its Figure 6), and after observing the shape in $\frac{M_t^N}{B_t^N}$, states that the "conclusions (...) are unaffected by endogenizing t^{**}". However, the economically relevant object is a price bubble.

To examine whether the stock price bubble follows from the pattern in M/B, Panel E plots cumulative returns to the new economy under an exogenous adoption time. Stock prices in this case do display a bubble pattern, with a period of abrupt negative returns just prior to adoption. This is consistent with the realized, though non-cumulative, returns reported in Panel C of Figure 4 in PV. In contrast, when adoption is endogenous, returns behave very differently. As Panel F shows, cumulative returns rise steadily throughout the revolution in the baseline calibration, without ever falling. Importantly, PV does not report returns or stock prices in the endogenous case, so this lack of a bubble is not immediately apparent.

The fact that market prices do not exhibit a bubble pattern in the endogenous case may seem surprising given that productivity shocks affect book values only through their drift, so that any unexpected movement in M/B is immediately reflected in returns. However, both $\frac{M_t^N}{B_t^N}$ and B_t^N have their own drifts, and over this period, B_t^N rises from high productivity faster than $\frac{M_t^N}{B_t^N}$ falls. Moreover, the decline in $\frac{M_t^N}{B_t^N}$ reflects expected mean reversion rather than newly arriving shocks. Conditional on a sequence of very positive productivity shocks, productivity is well above $\bar{\rho} + \hat{\psi}_t$, so expected productivity growth and M/B are anticipated to decline. Panels G and H make this distinction explicit by decomposing changes in M/B into drift and shock components. In the exogenous case, contemporaneous shocks drive the decline in M/B and generate a bubble pattern in prices. In the endogenous case, the downward movement in M/B solely reflects expected mean reversion.

The takeaway of Figure 1 is that, for the baseline calibration, there is no stock-price bubble in the endogenous adoption time case. We now turn to a discussion of why no bubble arises. We show that the difference between the endogenous and exogenous models is due to a fundamental difference in discount-rate dynamics across the two models.

2.1 Discount-rate effects in technological revolutions

To understand the discount rate effects of technological adoption, we first consider the new economy at any time t, and define $\frac{M_t^{NA}}{B_t^N}$ ("NA" for "Never Adopt") as the M/B ratio of the new economy at time t, given a level of $\hat{\psi}_t$ and ρ_t^N and under the assumption that adoption can never occur, regardless of the subjective belief.[3] Likewise define $\frac{M_t^{IA}}{B_t^N}$ ("IA" for "Immediately Adopt") as the value of the stock if adoption were to occur immediately, regardless of its optimality, given a level of $\hat{\psi}_t$ and ρ_t^N. These values are given by

$$\frac{M_t^{NA}}{B_t^N} = e^{C_0(\tau)+A_1(\tau)\rho_t^N+A_2(\tau)\hat{\psi}_t+\frac{1}{2}A_2(\tau)^2\hat{\sigma}_t^2}$$

$$\frac{M_t^{IA}}{B_t^N} = e^{C_0(\tau)+A_1(\tau)\rho_t^N+A_2(\tau)\hat{\psi}_t+\frac{1}{2}A_2(\tau)^2(1-2\gamma)\hat{\sigma}_t^2}$$

Here $\tau = T - t$, $A_1(\tau) = \tau - A_2(\tau)$, and C_0 is a constant that is defined in PV.[4] Note that the log-difference between these extremes, which we refer to as the "adoption discount,"

[3] This ratio is an endogenous function of the time-t state variables: (ρ_t^N, ρ_t, $\hat{\psi}_t$, and t). We suppress this in the notation for simplicity of exposition throughout this section and simply use the subscript t to denote a value conditional on the time-t state of the model.

[4] These equations are of the same form given in Corollary 2 in PV, which considers valuations just above and below the optimal threshold at time t^{**}, but are generalized to any time $t \in [t^*, t^{**}]$.

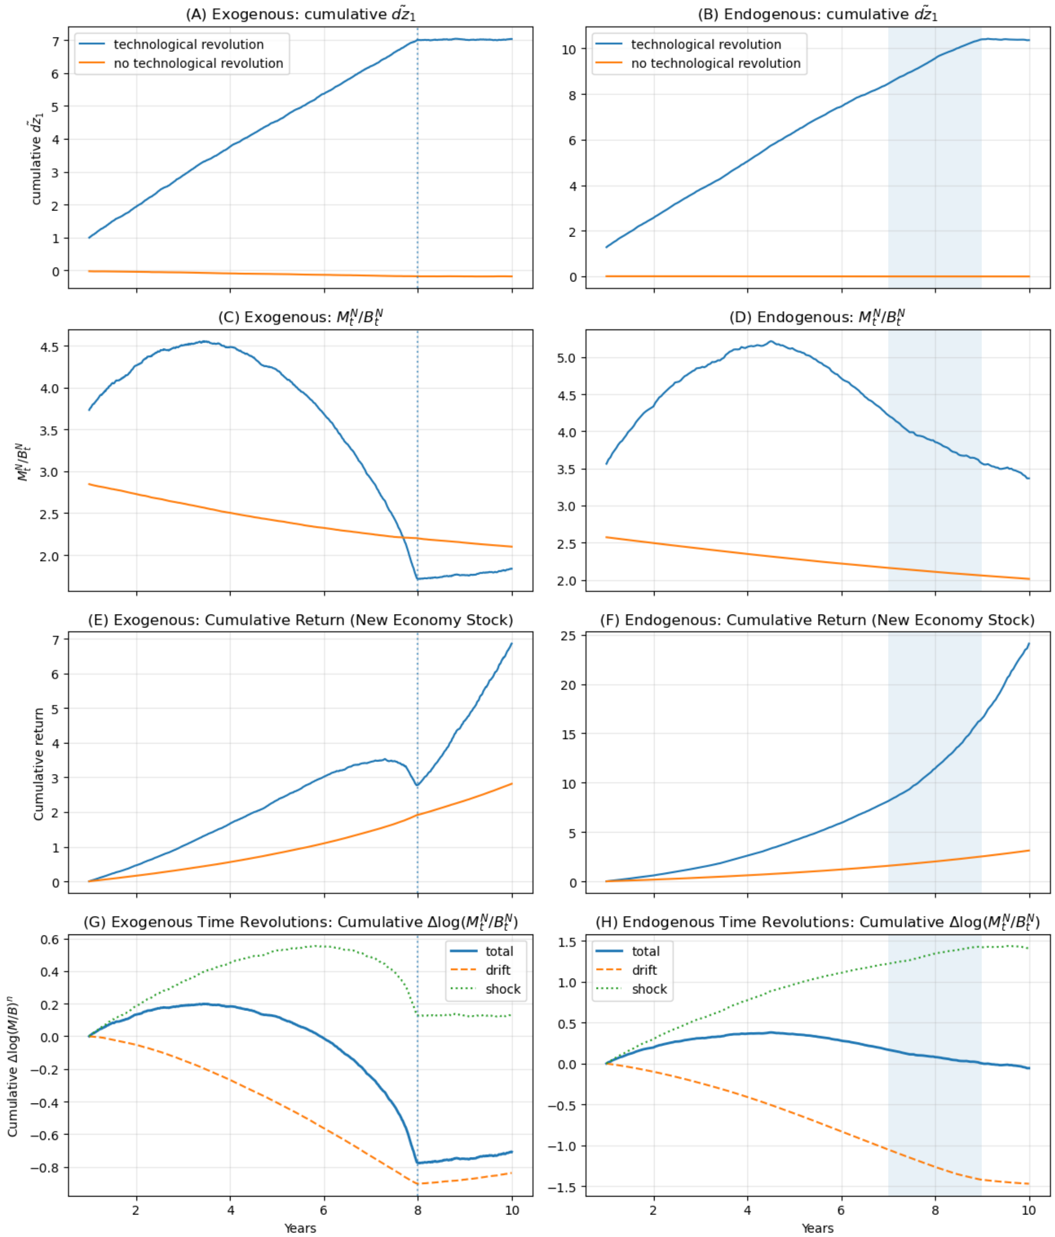

Figure 1: **New-economy bubbles in baseline calibration**

The figure shows means of simulated time-series, conditional on adoption (blue line) and on no adoption (orange line), for the new economy under baseline calibration of the model in Pástor and Veronesi (2009): $\gamma = 4$, $\phi = 0.3551$, $\bar{\rho} = 0.1217$, $\mu_J = 0$, $\sigma_J = 0.04$, $\sigma = \sigma_{N,0} = \sigma_{N,1} = 0.07$, and $\kappa = 0.1$. The left-hand panels show the model where adoption decisions occur at an exogenous time $\bar{t}^{**} = 8$. The right-hand panels show the model where endogenous adoption time is chosen optimally. Following PV, in the endogenous time case, we report adoptions that occur at $t^{**} \in [7, 9]$. Panels A and B show the series of cumulative unexpected productivity shocks to the new economy. Panels C and D show the M/B ratio for the new economy. Panels E and F show the cumulative returns to the new-economy stock. Panels G and H decompose the log of M/B ratios in Panels C and D into unexpected innovations and conditionally expected drift.

is $\mathcal{D}(t) \equiv \gamma A_2(T-t)^2\hat{\sigma}_t^2$, representing the lower price due to the systematic risk incurred if the new technology were adopted today versus never being adopted. Since $\gamma > 1$, $A_2(\tau) > 0$ and B_t^N is the same in both scenarios, we have that $M_t^{IA} \leq M_t^{NA}$.

The adoption discount is decreasing in t. As time passes, subjective risk falls, and time until the terminal cash flow falls as well. Equality occurs at time T ($\mathcal{D}(T) = 0$) when adoption no longer has any effect on the economy. Therefore, the "never adopt" case is equivalent to adopting at time T.

Also note that immediate adoption, from a discount-rate perspective, is the worst possible case for systematic risk. The longer you wait, the less time remains for shocks to the new economy to impact the old economy. Likewise, committing to never adopt is the best possible case for systematic risk. Lastly, since the adoption decision does not impact the cash flow of the new economy, the true market value M_t^N under uncertain adoption satisfies $M_t^{IA} \leq M_t^N \leq M_t^{NA}$.

2.1.1 Asset pricing implications of the option to wait

We first provide a simple example to illustrate how the move from an exogenous one-time adoption to an endogenous optimally timed adoption removes the new-economy price bubble. The key ingredients are the option to wait along with the planner's desire to wait for new-technology risk to resolve through learning, and a single later date is sufficient to see the effect.

In PV's benchmark exogenous model, the new-economy price falls throughout the year before adoption and, conditional on adoption, on average loses 6.0% in the final month alone. However, if a single second review date is added, so that the planner can adopt either at $T_1 = 8$ or, failing that, at $T_2 = 12$, we find that the average price conditional on a year-eight adoption *rises* into the adoption, and by 1.3% in the final month.[5]

The reason is that a later adoption date alters the meaning of a miss at year eight.

[5] A preset-calendar model of adoption is in the Online Appendix, showing the above results and others. All numbers use the baseline calibration with $\mu_J = 0$.

When year eight is the last chance, a non-adoption is permanent. Just below the threshold, the claim keeps its never-adopt value, and just above it, takes its adopted value. The step between the two is the adoption discount $\mathcal{D}(t) \equiv \gamma A_2 (T-t)^2 \hat{\sigma}_t^2$, equal to 1.77 at year eight, so a claim adopted then is worth 83% less than one that will never adopt. Price lies between these two values, reflecting how likely adoption has become. When beliefs are near the year-eight threshold, arbitrarily small news can move beliefs above the threshold, pulling price toward the full discount and creating the crash that completes the bubble.

When year twelve is an additionally available adoption time, a near miss is a delay. The planner's desire to wait for uncertainty to resolve means that the adoption hurdle is lower at year twelve, when learning has further reduced new-technology risk. Therefore, when beliefs are just below the year-eight adoption threshold, adoption at year twelve is now very likely. The discount $\mathcal{D}(12)$ has therefore already been impounded into the price approaching year eight. Adopting at year eight then lowers log prices by only the increment $\mathcal{D}(8) - \mathcal{D}(12) = 0.78$. The increment is still large but considerably smaller, and it shrinks as the time between T_1 and T_2 shrinks. For example, it falls to 0.23 when $T_2 = 9$.

Hence, PV's exogenous model assumes not only a fixed adoption date but a final one. Under continuous review, no date is last, and the thresholds tend to fall as uncertainty resolves, so the discount enters the price incrementally as positive cash-flow news arrives throughout the revolution. These incremental increases in discount rates therefore never overwhelm the cash-flow effect to generate negative returns.

2.1.2 Discount-rate effects in the exogenous and endogenous models

In order to understand how this logic applies to the endogenous model presented in PV, we use a simple approximation that links the value of the new economy to the expected timing of adoption:

$$\log(M_t^N) \approx \log(M_t^{IA}) + \left(\frac{\mathbb{E}_t[t^{**}] - t}{T - t} \right) \mathcal{D}(t) \tag{3}$$

Again, M_t^{IA} is the value under immediate adoption and $\mathcal{D}(t) = \gamma A_2(\tau)^2 \hat{\sigma}_t^2 = \log(M_t^{NA}) - \log(M_t^{IA})$ is the discount associated with making new-technology risk systematic immediately rather than never. Expected adoption time therefore summarizes how much of this discount is incorporated into the current price. Appendix A derives the approximation.

The appropriate probability measure for expected adoption time differs across the two versions of the model. In the exogenous case, the natural summary is the risk-neutral expected adoption time,

$$\mathbb{E}_t^Q[t^{**}] = p_t^Q \bar{t}^{**} + (1 - p_t^Q)T,$$

where p_t^Q is the risk-neutral probability of adoption. With endogenous timing, the physical expected first-passage time $\mathbb{E}_t^P[t^{**}]$ provides the more natural approximation. Intuitively, a fixed review date selects among a range of future states, whereas endogenous first passage largely pins down the belief at adoption and instead randomizes the date at which the boundary is reached. We show this distinction formally in Appendix A and confirm numerically that risk-neutral expectations perform better in the exogenous model while physical expectations perform better with endogenous timing.

For the planner problem studied in this section, however, physical and risk-neutral adoption probabilities are quantitatively close. We therefore use physical probabilities for both specifications below, allowing us to summarize both cases using the same object, $\mathbb{E}_t[t^{**}]$, with little loss of accuracy. This shortcut will not generally be innocuous once adoption is decentralized. In Section 3, physical and risk-neutral adoption probabilities can diverge substantially, and we distinguish between the two measures when needed.

This approximation of expected timing in equation 3 closely matches the numerically computed market value for the baseline calibration and most parameterizations we consider. It also provides the main intuition for the price dynamics. The discount-rate component of the new-economy price is approximately linear in expected adoption time. A positive productivity shock can therefore generate a large negative discount-rate effect only if it causes a large decline in expected time until adoption. The distinction between the exogenous and endogenous cases then depends primarily on how sharply $\mathbb{E}_t[t^{**}]$ responds to news. With endogenous timing, the stopping rule allows expected adoption time to adjust continuously as beliefs evolve. With a fixed review date, by contrast, a relatively small change in beliefs can sharply change the probability of crossing a single future hurdle, concentrating the associated increase in systematic risk over a short interval.

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

$$\frac{\partial \log(M_t^{IA})}{d\tilde{Z}_{1,t}} = A_2(\tau)\frac{\partial \hat{\psi}_t}{d\tilde{Z}_{1,t}} + A_1(\tau)\frac{\partial \rho_t^N}{d\tilde{Z}_{1,t}} \geq 0 \tag{4}$$

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:

$$\frac{\partial \left(\frac{\mathbb{E}_t[t^{**}]-t}{T-t} A_2(\tau)^2 \gamma \hat{\sigma}_t^2\right)}{d\tilde{Z}_{1,t}} = \frac{\partial \mathbb{E}_t[t^{**}]}{d\hat{\psi}_t}\frac{\partial \hat{\psi}_t}{d\tilde{Z}_{1,t}}\frac{1}{T-t}\mathcal{D}(t) \leq 0 \tag{5}$$

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 $\frac{\partial \mathbb{E}_t[t^{**}]}{d\hat{\psi}_t}$.

Figure 2 visualizes the dynamics of expected adoption times and prices for the two model cases using the baseline calibration of PV. To produce this figure, we construct a single "representative" simulation of productivity beliefs in the model. We draw a positive true value for ψ that is sufficient to generate a typical adoption in both model cases. We then construct a simulation path where the true values of $dZ_{1,t}$ are equal to zero. This leads to positive values of $d\tilde{Z}_{1,t}$ as productivity consistently outperforms expectations, and as a consequence, the subjective belief $\hat{\psi}_t$ rises through learning.

In Panel A, we plot this subjective belief along with the adoption threshold at t^{**} for the exogenous adoption time case, as well as the downward-sloping continuous threshold for the endogenous case. As the figure shows, in this simulation the subjective belief rises quickly enough to trigger an adoption in both models. If adoption time is endogenous, adoption occurs slightly later due to the value of waiting to learn.[6]

Panel B shows the evolution of $\mathbb{E}_t[t^{**}]$ in both models along this simulation path. In the endogenous case, productivity shocks have an immediate effect on $\mathbb{E}_t[t^{**}]$, which begins to decline steadily early in the revolution. Thus, even when expected adoption remains far in the future, shocks early in the revolution still have noticeable effects on the expected adoption time in the endogenous case, and hence affect prices. Along this path, the sensitivity $\partial\mathbb{E}_t[t^{**}]/\partial\hat{\psi}_t$ remains relatively stable until adoption occurs.

This stability is related to the classical first-passage problem in which a martingale crosses a downward-sloping linear boundary. In that case, the sensitivity of the expected stopping time to shocks is determined by the slope of the boundary and is independent of the current distance from the threshold (see Karatzas and Shreve (1998)). Here, the same learning about $\hat{\psi}_t$ that determines adoption also generates a downward-sloping adoption threshold (Equation 2). Therefore, when the approximation in Equation 3 is accurate, the exposure of discount rates to productivity shocks tends to remain relatively stable over the course of the revolution. Since the approximation is generally accurate for parameterizations near the baseline specification of PV, discount-rate shocks are incremental and insufficient to overwhelm the positive cash-flow shocks.

[6] We set the $d\tilde{Z}_{1,t} = 0$ after the endogenous adoption to be consistent with conditioning on adoption, but this does not impact pre-adoption dynamics for either model.

The exogenous specification behaves differently because the adoption threshold is not downward sloping, but instead takes the form of a discrete hurdle at $\bar{t}^{**}$. As Panel B shows, early in the revolution, increases in productivity have only a small effect on the probability of adoption because adoption at $\bar{t}^{**}$ remains highly unlikely. As a result, the initial decline in expected adoption time for the exogenous model is modest. As the preset adoption decision nears, however, even small increases in beliefs about productivity have a large effect on the probability of adoption when adoption is near the threshold. Expected adoption time in the exogenous case therefore falls rapidly late in the revolution.

This difference in the behavior of expected adoption times across the two models leads to very different stock price paths, as Panel C shows. The panel plots the two extremes, M_t^{NA} and M_t^{IA}, which are the same in the endogenous and exogenous models, and the true market values for the respective model cases. The approximate values from Equation 3 are also shown and perform extremely well. In the endogenous model, the steady drop in expected adoption time leads to a smooth transition from M_t^{NA} to M_t^{IA}, effectively spreading out the discount-rate effect. The cash-flow effect of increased productivity dominates throughout the entire revolution to produce a steadily rising price in the endogenous case.[7] In contrast, the concentrated decrease in expected adoption time near the revolution in the exogenous case leads to a large discount-rate effect late in the revolution, which overwhelms the cash-flow effect and creates a bubble pattern.

Panel D provides yet another perspective on the two models. The panel graphs the exposure of new-economy stock prices to new-economy productivity shocks ($\frac{\partial \log(M_t^N)}{d\tilde{Z}_{1,t}}$). In the exogenous case, this value is strongly negative just prior to adoption, meaning the *positive* productivity shocks that are necessary to create an adoption are the actual triggers of *negative* stock returns. This is the core driver of the systematic-risk mechanism. The existence of this region in the exogenous case, which is proved in Corollary 2 of PV, is necessary to produce a bubble pattern when conditioning on an adoption. Since this region does not exist in this

[7] The endogenous model has a lower price initially as its expected adoption time is initially lower than the exogenous model, and it therefore has more systematic risk.

simulation for the endogenous model, no bubble occurs.

Additionally, one may wonder whether relaxing the assumption of the smooth arrival of news can produce an abrupt change in the expected timing of adoption, and hence trigger a bubble. The non-existence of the region above also means that lumpy news arrival does not generally imply a bubble. We solve the endogenous model but replace Gaussian productivity shocks with discrete Poisson arrivals, holding the expected flow of information fixed.[8] While beliefs can then jump through the adoption boundary and generate sudden increases in discount rates, these jumps also generate large increases in expected cash flows. We find that the relative cash-flow effect still dominates, so that every positive signal that triggers adoption carries cash-flow news that outweighs the discount-rate effect.

To show that this failure to produce a bubble generalizes across reasonable parameters, Figure 3 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.[9] 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 is performing well across all these parameterizations.[10]

These results show that the lack of a bubble in the endogenous case holds for a wide set

[8] See Online Appendix.

[9] 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.

[10] See Figure A1 in the Online Appendix.

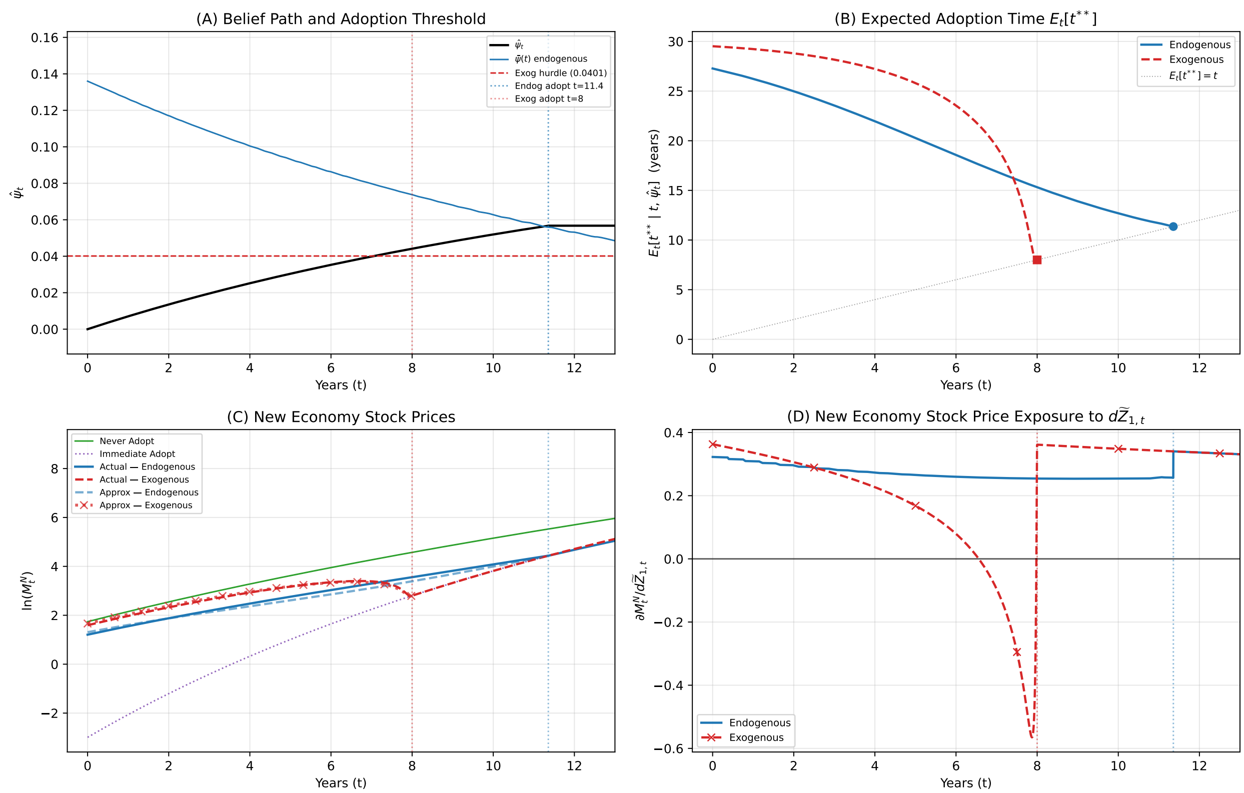

Figure 2: **Expected adoption times and new-economy stock market prices.**

Each panel shows model dynamics in the exogenous and endogenous model specifications for a single representative simulation of new-economy productivity that generates an adoption in both cases. Panel A shows the simulated new-technology productivity belief ($\hat{\psi}_t$) as well as the adoption thresholds. Panel B shows the conditional expected adoption times $\mathbb{E}_t[t^{**}]$. Panel C shows the actual prices $\log(M_t^N)$ of the new economy in the two model specifications, which fall between the two extremes of $\log(M_t^{NA})$ and $\log(M_t^{IA})$, and also shows the approximate price using $\log\left(M_t^N\right) \approx \log\left(M_t^{IA}\right) + \left(\frac{\mathbb{E}_t[t^{**}]-t}{T-t}\right) A_2(\tau)^2 \gamma \hat{\sigma}_t^2$. Panel D shows the conditional exposure of the new-economy stock price to new-economy productivity shocks $\left(\frac{\partial \log(M_t^N)}{d\widetilde{Z}_{1,t}}\right)$.

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.[11] In short, the systematic-risk channel for price bubbles is not reasonably attainable within the planner's endogenous-adoption problem.

We next explore whether decentralized adoption decisions can restore a price bubble by delivering a boundary that suppresses the option to delay, so that the price at adoption bears the full discount.

3 Decentralized adoption and asset prices

The prior section focused on the temporal dynamics of the PV model, establishing that a last-chance adoption date is required to produce a price bubble when the planner chooses adoption. We turn now from the timing of adoptions to how adoption is decided across firms. PV show that, when the adoption date is fixed and all firms adopt at once, the planner's decision is one of a range of decentralized equilibria, each supported by costs on firms that deviate from the aggregate adoption. These costs both enforce simultaneous adoption and select the belief at which it occurs. We extend the PV analysis in three steps. First, Section 3.1 characterizes the range of simultaneous-adoption equilibria at PV's fixed date, including its two extremes, and then removes the deviation costs so that a subset of firms may adopt. Second, Section 3.2 carries the simultaneous-adoption equilibria to endogenous timing. Third, Section 3.3 constructs an equilibrium that restores the bubble.

One theme runs through all three cases. An individual firm values unresolved technological risk while that risk is idiosyncratic, but adoption by others, current or anticipated, makes the same risk systematic and gives the firm a reason to wait. This incentive to delay keeps decentralized adoption close to the planner's and spreads the repricing of new-technology risk, so the bubble survives only when coordination and preemption incentives override it,

[11] For instance, in the Online Appendix 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.

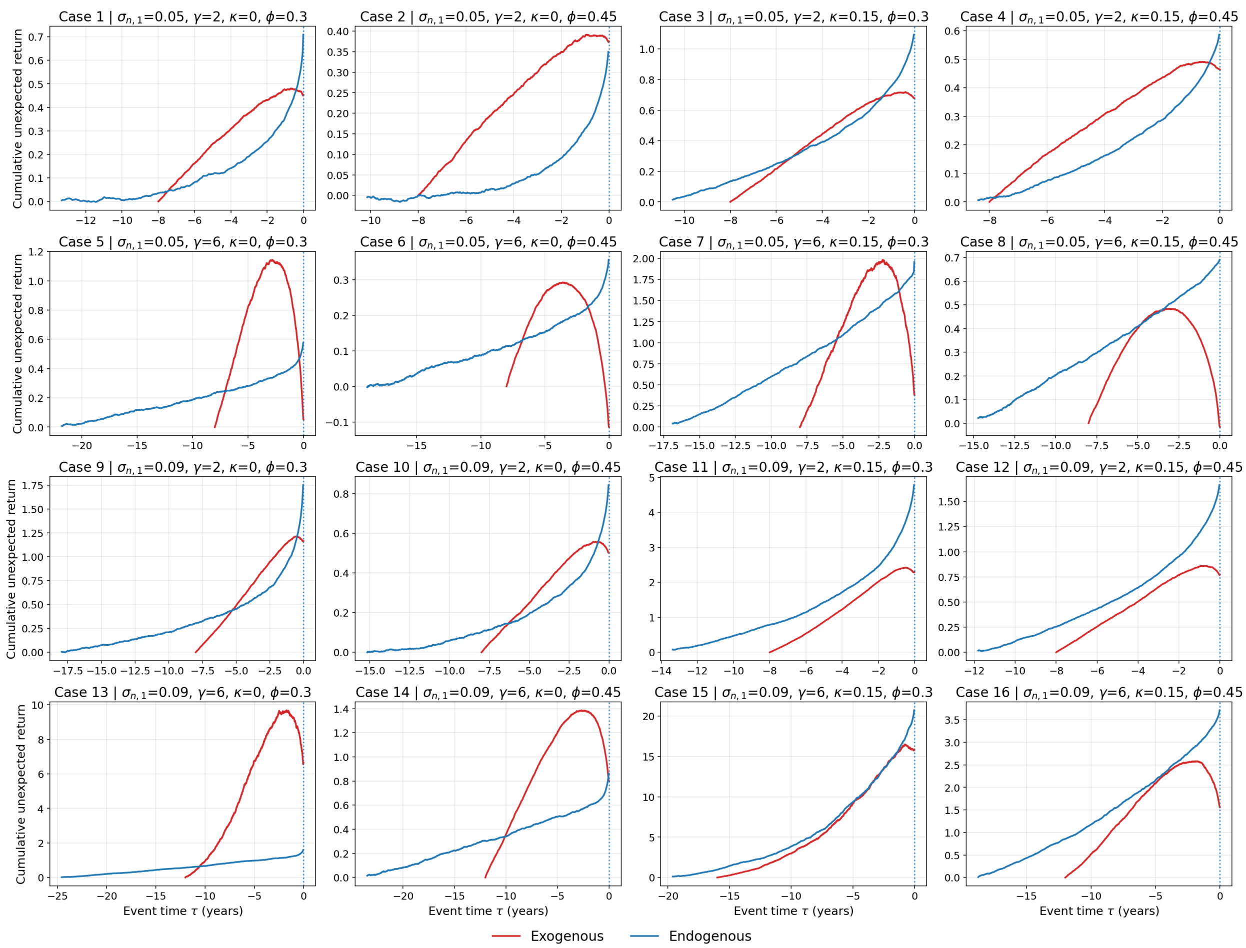

Figure 3: **Cumulative unexpected returns in revolutions across parameterizations**

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.

and then at a large cost both to social welfare and to the value of an adopting firm.

3.1 Decentralized adoption thresholds in the exogenous model

We first characterize decentralized adoption rules in PV's exogenous adoption-time setting, where the different coordination assumptions are especially transparent. There are three simultaneous adoption rules we study, which we term the planner, join, and unilateral rules. The planner rule is the utility-maximizing decision of Section 2. The unilateral rule adopts at the belief where a single firm is indifferent between adopting alone and remaining in the old technology while the rest of the economy stays old, so its new-technology exposure is priced as idiosyncratic. The join rule adopts at the belief where a firm is indifferent between joining an economy-wide adoption and lagging behind it, so the same exposure is priced as systematic. We then remove coordination costs and allow a small amount of heterogeneity in conversion costs to study a partial adoption equilibrium.

Proposition 1 states the resulting static thresholds and the lead and lag costs needed to support each simultaneous equilibrium as an aggregate-adoption rule. The proof, together with the underlying exponential-affine building blocks from PV's equations (14)–(16) and Technical Appendix, is provided in our Appendix.

For an aggregate regime $r \in \{0, 1\}$ (old or adopted economy) and a firm's own technology $a \in \{0, 1\}$ (old or new), let $V^{a|r}$ denote the value of a firm's claim under technology a, evaluated under the regime-r pricing kernel. The separation of the firm's own technology a from the aggregate regime r reflects the maintained assumption that an individual firm's adoption changes its own payoff but does not alter the aggregate regime or its pricing kernel. For any candidate adoption rule, define the lead cost

$$c^{lead} = \max\big\{0,\ \log V^{1|0} - \log V^{0|0}\big\}$$

as the penalty needed to deter a firm from adopting unilaterally while the aggregate economy is still old ($r = 0$), and the lag cost

$$c^{lag} = \max\bigl\{0,\ \log V^{0|1} - \log V^{1|1}\bigr\}$$

as the penalty needed to deter a firm from remaining in the old technology once the aggregate economy has adopted ($r = 1$). In the static problem,

$$V^{a|r} = E_{\bar{t}^{**}}\left[\pi^{r}_{\bar{t}^{**},T} B^{a}_{i,T}\right],$$

where π^r is the regime-r pricing kernel and $B^a_{i,T}$ is firm i's terminal book value under technology a. Because the costs are defined as log differences, they enter the underlying level values multiplicatively. In each case, the cost is exactly large enough to restore indifference with the non-deviating value at the boundary.

Proposition 1 (Static thresholds and support costs)**.** *Fix an adoption date $\bar{t}^{**}$; let $\tau = T - \bar{t}^{**}$; and write $\hat{\sigma}^2 \equiv \hat{\sigma}^2_{\bar{t}^{**}}$. The planner, join (J), and unilateral (U) thresholds for $\hat{\psi}_{\bar{t}^{**}}$ are*

$$\bar{\psi}_{planner,\bar{t}^{**}} = \frac{-\log(1-\kappa) + \frac{1}{2}(\gamma - 1)A_2(\tau)^2\hat{\sigma}^2}{A_2(\tau)}.$$

$$\bar{\psi}_{J,\bar{t}^{**}} = \frac{-\log(1-\kappa) + \frac{1}{2}(2\gamma - 1)A_2(\tau)^2\hat{\sigma}^2}{A_2(\tau)}.$$

$$\bar{\psi}_{U,\bar{t}^{**}} = \frac{-\log(1-\kappa) - \frac{1}{2}A_2(\tau)^2\hat{\sigma}^2}{A_2(\tau)}.$$

For $\gamma > 1$, these satisfy

$$\bar{\psi}_{U,\bar{t}^{**}} < \bar{\psi}_{planner,\bar{t}^{**}} < \bar{\psi}_{J,\bar{t}^{**}}.$$

The lead and lag costs required to support each threshold as the aggregate-adoption rule are

$$\bigl(c^{lead}_J, c^{lag}_J\bigr) = \bigl(\gamma A_2(\tau)^2\hat{\sigma}^2, 0\bigr),$$

$$\bigl(c^{lead}_U, c^{lag}_U\bigr) = \bigl(0, \gamma A_2(\tau)^2\hat{\sigma}^2\bigr),$$

and

$$\left(c^{lead}_{planner}, c^{lag}_{planner}\right) = \left(\tfrac{1}{2}\gamma A_2(\tau)^2\hat{\sigma}^2, \tfrac{1}{2}\gamma A_2(\tau)^2\hat{\sigma}^2\right).$$

Proof. See the Appendix. □

The join equilibrium therefore needs only a lead cost, to deter premature unilateral adoption; the unilateral equilibrium needs only a lag cost, to induce the rest of the economy to follow; and decentralizing the planner's threshold needs both, each equal to half the join or unilateral magnitude. The planner threshold lies exactly halfway between the unilateral and join thresholds. These three thresholds correspond to different ways of sustaining a binary aggregate transition rather than to different technological environments.

The variance terms also reveal the private force pushing toward early adoption. Holding the decision date fixed, the unilateral threshold is decreasing in posterior variance:

$$\frac{\partial \bar{\psi}_{U,\bar{t}^{**}}}{\partial \hat{\sigma}^2} = -\frac{1}{2}A_2(\tau) < 0.$$

A firm evaluating adoption against the old-economy pricing kernel benefits from unresolved technology risk because its adoption payoff is convex. The firm treats the additional risk created by its own infinitesimal adoption as idiosyncratic: it captures the convexity benefit but does not internalize the systematic-risk cost that would arise if the whole economy adopted. Greater uncertainty therefore raises the option value of adopting and lowers the belief required to do so. This force works in the opposite direction from the planner's desire to wait until learning has reduced the systematic risk created by aggregate adoption.

No coordination costs. We now allow firms to adopt separately. Firms differ slightly in conversion costs, which determine their order of adoption, and face no lead or lag costs. As adoption expands, new-technology risk becomes increasingly systematic and raises the threshold of the next firm.

Proposition 2 (Static adoption without coordination costs)**.** *Suppose firms face no lead or lag costs and have conversion costs uniformly distributed on $[\kappa - \epsilon, \kappa + \epsilon]$, where $\epsilon > 0$. The competitive adoption fraction equals the planner's optimal fraction under the conditions stated in the Appendix.*

Additionally, under the partial-adoption aggregation specified in the Appendix, as ϵ converges to zero, let $\bar{\psi}_0(q)$ denote the limiting belief at which the marginal firm q adopts. Then

$$\bar{\psi}_0(0) = \bar{\psi}_U, \qquad \bar{\psi}_0\left(\frac{1}{2}\right) = \bar{\psi}_{planner}, \qquad \bar{\psi}_0(1) = \bar{\psi}_J,$$

and

$$\bar{\psi}_0(1) - \bar{\psi}_0(0) = \gamma A_2(\tau)\hat{\sigma}^2.$$

Proof. See the Appendix. □

The partial adoption rule changes the interpretation of the three simultaneous thresholds. They are points on the adoption schedule that emerge when firms are allowed to adopt separately: the unilateral threshold governs the first adopter, the planner threshold corresponds to half of the economy, and the join threshold governs the last adopter. Importantly, the interval between the first and last adoption thresholds does not vanish with the dispersion in conversion costs. An arbitrarily small amount of heterogeneity orders firms, while the change in the stochastic discount factor as q rises spreads adoption across a nondegenerate range of beliefs.

Without lead or lag costs, decentralized firms choose to spread out as systematic risk makes adoption a strategic substitute across firms. Moreover, this decentralization yields the same adoption fraction as a planner who can choose q. The all-or-nothing economy-wide transition relies on the coordination forces that overcome this effect to keep firms moving together. Because a PV-style price decline requires a sudden economy-wide transition from idiosyncratic to systematic risk, this distinction will be central to the asset-price analysis below.

3.1.1 Asset-price implications

We next compare asset prices under synchronized and partial adoption.[12] For each synchronized rule — unilateral, planner, and join — we choose a separate prior mean μ_J so that the simulation generates approximately 500 economy-wide adoptions by $\bar{t}^{**} = 8$, and we average across those adoption paths.[13] We apply the same calibration to the partial-adoption economy, defining the event as complete adoption at the decision date. This normalization lets us compare rare technological revolutions across adoption rules without mechanically changing their ex ante frequency.

Panels A and B of Figure 4 report physical and risk-neutral adoption probabilities in the synchronized and partial-adoption economies, respectively.[14] In the partial-adoption economy, this is equivalently the probability that a randomly selected firm has adopted, or the expected aggregate adoption fraction. Because the prior mean for each rule places its threshold at the same percentile of the terminal-belief distribution, the physical expected adoption fraction is identical across the three rules and is shown once. The risk-neutral paths differ because the stochastic discount factor assigns different weights to the adoption state under the unilateral, planner, and join calibrations. Panels C and D report cumulative log returns of the new economy, while Panels E and F report cumulative log returns of the old economy.

The curvature of the risk-neutral probability as a function of the physical probability summarizes the welfare content of each rule. Under the unilateral rule, aggregate adoption occurs at such a low belief that it produces a large welfare loss. The stochastic discount

[12] We use the benchmark PV calibration with the one exception that the volatility of old and new productivity shocks is 0.05 instead of 0.07. This allows for clearer illustration of the mechanism in the figures, but does not affect qualitative implications.

[13] We use rule-specific values of μ_J to preserve the PV experiment in which adoption is initially unlikely but becomes likely as a technological revolution unfolds. A negative μ_J can be interpreted literally as implying that the typical new technology does not improve long-run growth. Alternatively, it can stand in for an additional installation cost measured as a reduction in subsequent growth, rather than as the one-time level cost κ.

[14] Recall from Section 2 that for the exogenous model, the risk-neutral probabilities better approximate the discount rate effect than the physical probabilities.

factor is therefore high in the adoption state, and the risk-neutral adoption probability is a concave function of the physical probability: it rises rapidly when the true probability is still low and then flattens. Under join, firms wait beyond the planner's threshold, so eventual adoption is a favorable aggregate shock. Its state price is low, and the risk-neutral adoption probability is instead convex in the physical probability: it initially responds little and accelerates only near the decision date. The planner lies between these extremes.

The synchronized economy retains the central asset-pricing feature of PV. The physical probability of adoption rises slowly early in the revolution and then accelerates as the fixed decision date approaches. Conditional on a revolution, the new-economy price rises with favorable productivity news but declines near the adoption date as the probability of an economy-wide shift to systematic risk increases sharply. The rule-specific calibration of μ_J equalizes the cumulative cash-flow news for the new economy, so its cumulative return reaches the same endpoint under all three rules. The old-economy paths instead depend on the productivity belief at which aggregate adoption occurs. That belief is lowest under unilateral adoption and highest under join, making the transition a large adverse aggregate shock in the former case and a favorable shock in the latter. These paths make the welfare ordering especially transparent: unilateral adoption produces a large loss for the old economy, the planner leaves its value approximately unchanged, and join adoption produces a gain because firms have delayed a valuable transition.

The same probability curvature determines the strength of the bubble. In the unilateral case, the adoption state is overweighted under the risk-neutral measure well before the decision date, spreading its price effect over time and producing the weakest reversal. In the join case, the favorable adoption state is initially underweighted and the risk-neutral probability catches up only near the decision date, concentrating the repricing and producing the strongest bubble. The planner case lies between the two.

The right column repeats the exercise when firms are heterogeneous and face no coordination costs. Conditional on eventual complete adoption, the expected fraction adopting begins to rise well before the decision date and increases steadily as information arrives. The risk-neutral expected fraction lies below the physical expectation through most of the revolution, reflecting the lower state price attached to outcomes in which a larger share of the economy adopts. Because the adoption fraction at the common decision date adjusts smoothly with beliefs, the associated change in expected systematic risk is spread across the pre-adoption period. The new- and old-economy cumulative returns therefore rise smoothly rather than displaying the concentrated late-revolution reversal generated by synchronized adoption.

Panels G and H compare welfare with the corresponding planner allocation. With CRRA preferences, conditional welfare is

$$U_t^r = \frac{1}{1-\gamma}\overline{E}_t[(W_T^r)^{1-\gamma}],$$

where $\overline{E}_t$ averages conditional expected utility across the selected complete-adoption paths. We express this welfare level as the certainty-equivalent ratio

$$\Omega_t^r = \left(\frac{\overline{E}_t[(W_T^r)^{1-\gamma}]}{\overline{E}_t[(W_T^{planner})^{1-\gamma}]} \right)^{1/(1-\gamma)},$$

where the common factor $1/(1-\gamma)$ cancels before applying the certainty-equivalent transformation. This normalization sets the planner benchmark equal to one, with values below one indicating welfare losses. The planner line in Panel G is therefore one by construction. The unilateral rule adopts in states in which the planner would wait, whereas the join rule delays adoption in states in which the planner would adopt. Both deviations lower welfare during the revolution. In the partial-adoption economy, decentralized firms choose the same adoption fraction as the planner at every belief, so the welfare ratio in Panel H is also one by construction.

The static comparison therefore isolates the role of how adoption is aggregated. When the economy still moves simultaneously at a fixed date, the PV bubble survives under the

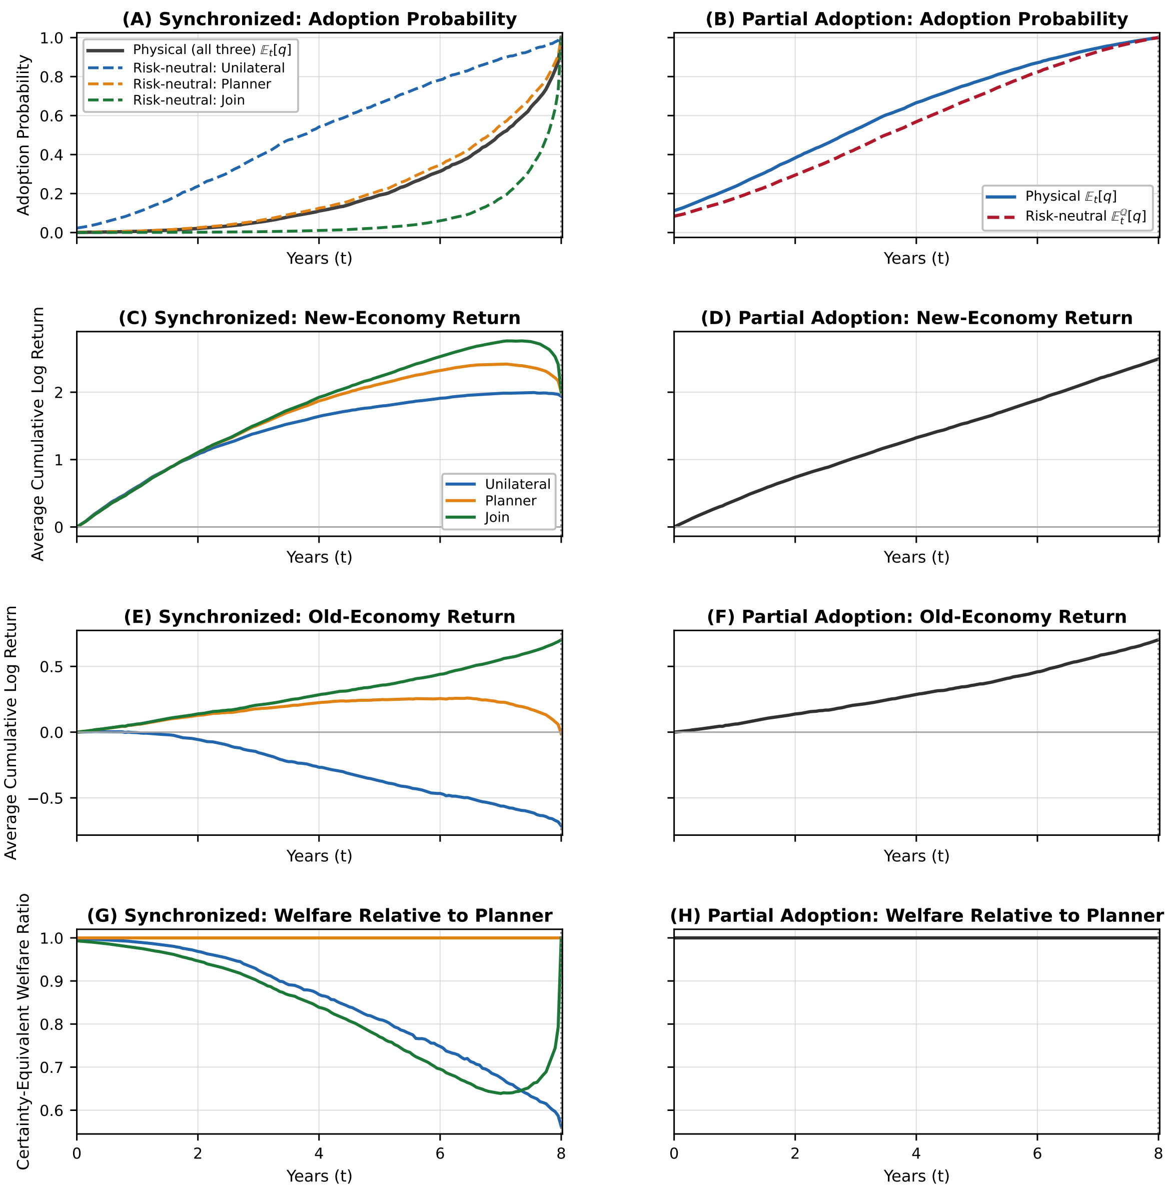

Figure 4: **Asset prices and welfare under synchronized and partial adoption.** The left column shows the synchronized unilateral, planner, and join rules; the right column shows partial adoption by heterogeneous firms with no coordination costs. Panels A and B show physical and risk-neutral adoption probabilities; in the partial-adoption economy, these are the probabilities that a randomly selected firm has adopted. Solid and dashed lines distinguish the two measures. Panels C and D show average cumulative new-economy log returns, Panels E and F show average cumulative old-economy log returns, and Panels G and H show the certainty-equivalent ratio associated with $U_t^r = (1-\gamma)^{-1}E_t[(W_T^r)^{1-\gamma}]$, normalized by the corresponding planner allocation. Colors identify the synchronized rules as indicated by the legends. For each rule, μ_J is calibrated as described in the text. Simulations use $dt = 0.05$. Parameters include $\sigma = \sigma_{N,0} = \sigma_{N,1} = 0.05$ and $\sigma_J = 0.04$.

unilateral, planner, and join decentralizations: changing the equilibrium threshold changes the strength and welfare content of the reversal, but not the underlying mechanism. Once heterogeneous firms adopt at the common date only when their own conversion costs justify doing so, adoption is partial, the introduction of systematic risk is spread across belief states, and the bubble disappears.

In the Online Appendix, we also solve the endogenous-timing problem with heterogeneous firms and partial adoption, where gradual adoption additionally makes aggregate output a second source of information about ψ. Its asset-pricing implications are similar to those of the static partial-adoption case: gradual adoption spreads the increase in systematic risk over time and does not produce a bubble. We therefore retain simultaneous aggregate adoption going forward and next ask whether the PV bubble survives when the date of that transition is endogenous.

3.2 Decentralized adoption with endogenous timing

We now allow the date of simultaneous adoption to be chosen endogenously by firms. Accordingly, we study symmetric threshold equilibria in which all firms adopt when a common boundary is reached. For an equilibrium rule $e \in \{J, U\}$, let $\bar{\psi}_e(t)$ denote its aggregate adoption boundary and define

$$t_e^{**} = \inf\{s \geq t : \hat{\psi}_s \geq \bar{\psi}_e(s)\}.$$

As in the exogenous model, an individual firm takes aggregate wealth, the stochastic discount factor, the public belief process, and the adoption rule of the remaining firms as given. We do not solve for the lead and lag costs themselves. Instead, we assume that whichever cost is used to support a selected equilibrium can be made sufficiently large.

This leaves an equilibrium-selection problem. In the join case, we impose no lag cost. A lead cost rules out adoption before the common boundary, and firms adopt when joining an immediate aggregate transition is at least as valuable as continuing to wait. The lowest boundary satisfying this condition is unique at every date. Higher boundaries can also be supported with larger lead costs, but we select the unique minimum no-lag boundary. Writing $V_J^{0|0}(t, \hat{\psi}; \bar{\psi}_J)$ for the value of waiting while the economy remains old and $V^{1|1}(t, \hat{\psi})$ for the value after immediate aggregate adoption, the selected boundary satisfies

$$V^{1|1}(t, \bar{\psi}_J(t)) = V_J^{0|0}(t, \bar{\psi}_J(t); \bar{\psi}_J).$$

In the unilateral case, we impose no lead cost. A firm is free to adopt before the proposed aggregate boundary, while a sufficiently large lag cost induces the remaining firms to adopt when that boundary is reached. A candidate boundary b is feasible only if no firm prefers early adoption at any belief strictly below it:

$$V_U^{1|0}(t, \hat{\psi}; b) \leq V_U^{0|0}(t, \hat{\psi}; b) \qquad \text{for every } t \text{ and } \hat{\psi} < b(t). \tag{6}$$

Here the early adopter anticipates that the rest of the economy may still adopt at the future boundary b. Thus both continuation values depend on the entire future path of the proposed aggregate rule.

Unlike the join problem, the feasible unilateral 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[15]

$$b(t; x_1, x_2) = \frac{-\log(1-\kappa)}{A_2(T-t)} + x_1\hat{\sigma}_t + x_2 A_2(T-t)\hat{\sigma}_t^2 \tag{7}$$

[15]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.

and select

$$(x_1^*, x_2^*) \in \arg\max_{x_1, x_2} b(15; x_1, x_2) \quad \text{subject to Equation (6).}$$

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 below are robust to these alternatives.

The planner's endogenous boundary provides the welfare benchmark. It solves the optimal stopping problem for aggregate utility and can be implemented as a decentralized equilibrium by assuming sufficiently large lead and lag costs. The three boundaries coincide when $\gamma = 0$. This common boundary is also the stopping rule an individual firm would choose in isolation, without pricing the effect of adoption on aggregate systematic risk. Unlike the static threshold, it embeds the option to wait and learn before exercising. Its upward slope over the economically relevant region shows that private convexity can, by itself, produce an upward-sloping endogenous adoption boundary: as posterior variance falls, the private benefit of unresolved idiosyncratic risk disappears, making the firm less eager to adopt. With a constant stochastic discount factor, the planner and individual firm nevertheless solve the same stopping problem. Risk aversion separates their boundaries because an aggregate transition makes previously idiosyncratic new-technology risk systematic.

Figure 5 illustrates this separation. The join boundary rises most with risk aversion because a firm must be willing to adopt even when all other firms adopt at the same time. The planner boundary lies below it. The unilateral boundary is lowest and is the flattest of the three. Although a potential first mover does not internalize the effect of its own infinitesimal action on aggregate risk, it does not behave as though future aggregate adoption were impossible. It anticipates that the remaining firms may subsequently adopt, so the future aggregate transition enters its continuation value and makes new-technology risk systematically priced. This force offsets the upward-sloping private boundary visible at $\gamma = 0$, restores a downward slope over the relevant pre-adoption region, and preserves much of the planner's incentive to wait for learning.

3.2.1 Asset-price and welfare implications

Figure 6 compares the three rules at $\gamma = 4$. For each rule, we calibrate μ_J separately using the same adoption-frequency target as above and average across the resulting adoption paths. This normalization compares rare technological revolutions across equilibrium rules while allowing the realized adoption date to be endogenous and to differ across selected paths. We align each selected path on its final pre-adoption observation and report dynamics in event time, measured as years remaining until adoption.

Panel A reports the expected time remaining until adoption under the physical measure, $E_t[t_e^{**} - t]$, with paths that do not adopt before T assigned an adoption date of T. This is the expected stopping time that enters the expected-adoption-time approximation in Section 2. Expected time remaining declines gradually under all three rules. It responds somewhat more strongly under the unilateral rule and more slowly under join, but none exhibits the abrupt collapse required for the PV bubble mechanism. The unilateral rule comes closest. Its relatively flat boundary allows favorable news to pull adoption forward more quickly than under the planner or join rules, but anticipated adoption by the rest of the economy still smooths the decline in expected time remaining.

Panels B and C report new- and old-economy cumulative log returns relative to each path's final pre-adoption observation. The paths are aligned before averaging, so zero on the horizontal axis is the last observation before adoption for every simulation. The new-economy returns rise toward adoption under all three rules, without the synchronized-adoption reversal found at a fixed decision date. The old-economy paths reflect the different timing of the transition across rules. Increasing risk aversion raises the boundaries in Figure 5 and reinforces the incentive to wait, rather than creating an abrupt adoption wave. In the calibrations we study, the new-economy bubble disappears under the planner, join, and unilateral rules.

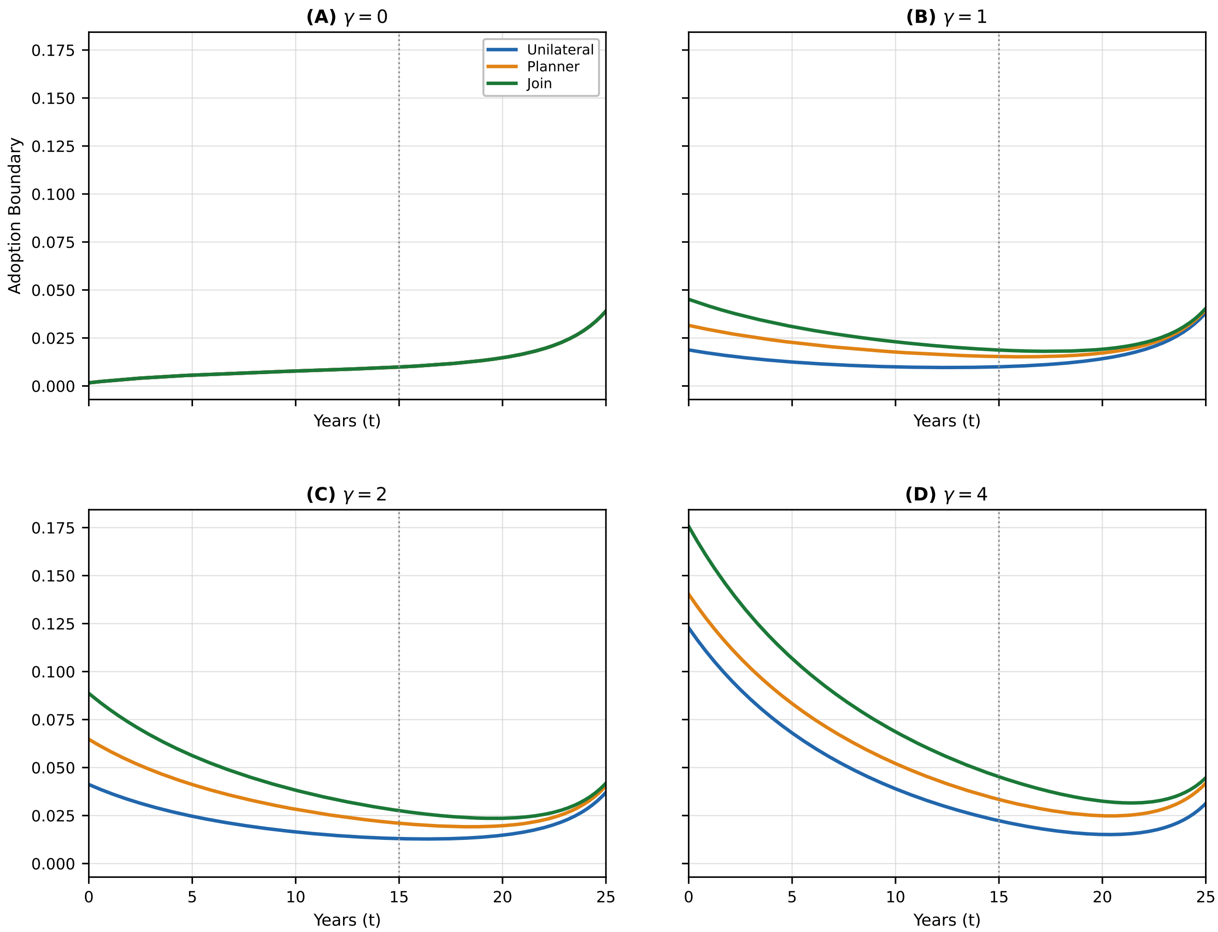

Figure 5: **Endogenous adoption boundaries and risk aversion.** Panels A–D show the unilateral, planner, and join adoption boundaries for $\gamma \in \{0, 1, 2, 4\}$. The unilateral boundary is selected from Equation (7) by maximizing its value at year 15 subject to the no-profitable-early-adoption condition. The dotted vertical line denotes year 15. Parameters other than γ are $\kappa = 0.10$, $\sigma = \sigma_{N,0} = \sigma_{N,1} = 0.05$, $\sigma_J = 0.04$, $\phi = 0.3551$, and $T = 30$.

Panel D reports certainty-equivalent welfare relative to the planner, averaging the relevant conditional terminal-wealth moments over the same selected paths. The welfare calculation retains the technology growth generated after either rule adopts; in particular, if the planner adopts before join, the planner continues to accrue new-technology growth while join remains in the old economy. The planner is one by construction. The unilateral rule adopts too early relative to the planner, while join adopts too late, so both rules generate welfare losses.

Endogenizing the adoption date therefore removes the new-economy bubble even when adoption remains simultaneous. Across all three decentralized rules, firms retain an option to wait for uncertainty to resolve. Under risk aversion, anticipated aggregate adoption makes unresolved new-technology risk systematically priced, generating downward-sloping boundaries over the relevant region and preventing small news from producing the abrupt acceleration in adoption required by PV. The unilateral boundary is the flattest and comes closest, but the same systematic risk that could generate the price reversal also discourages firms from adopting early. To see what is required to recover the bubble, we next exploit the multiplicity of aggregate-adoption equilibria and support the upward-sloping $\gamma = 0$ boundary in the risk-averse economy.

3.3 Selecting an equilibrium with a bubble

Producing a bubble requires a sudden increase in the risk-neutral probability of aggregate adoption while substantial uncertainty about the new technology remains. This is unlikely when firms wait for uncertainty to resolve and the adoption boundary slopes downward. The upward-sloping $\gamma = 0$ boundary is therefore a natural candidate for an equilibrium in which a small improvement in beliefs produces the sharp increase in adoption probability needed for a bubble.

Let $b_{\gamma=0}(t)$ denote the exact $\gamma = 0$ stopping boundary. We implement this boundary in the $\gamma = 4$ economy in three steps.

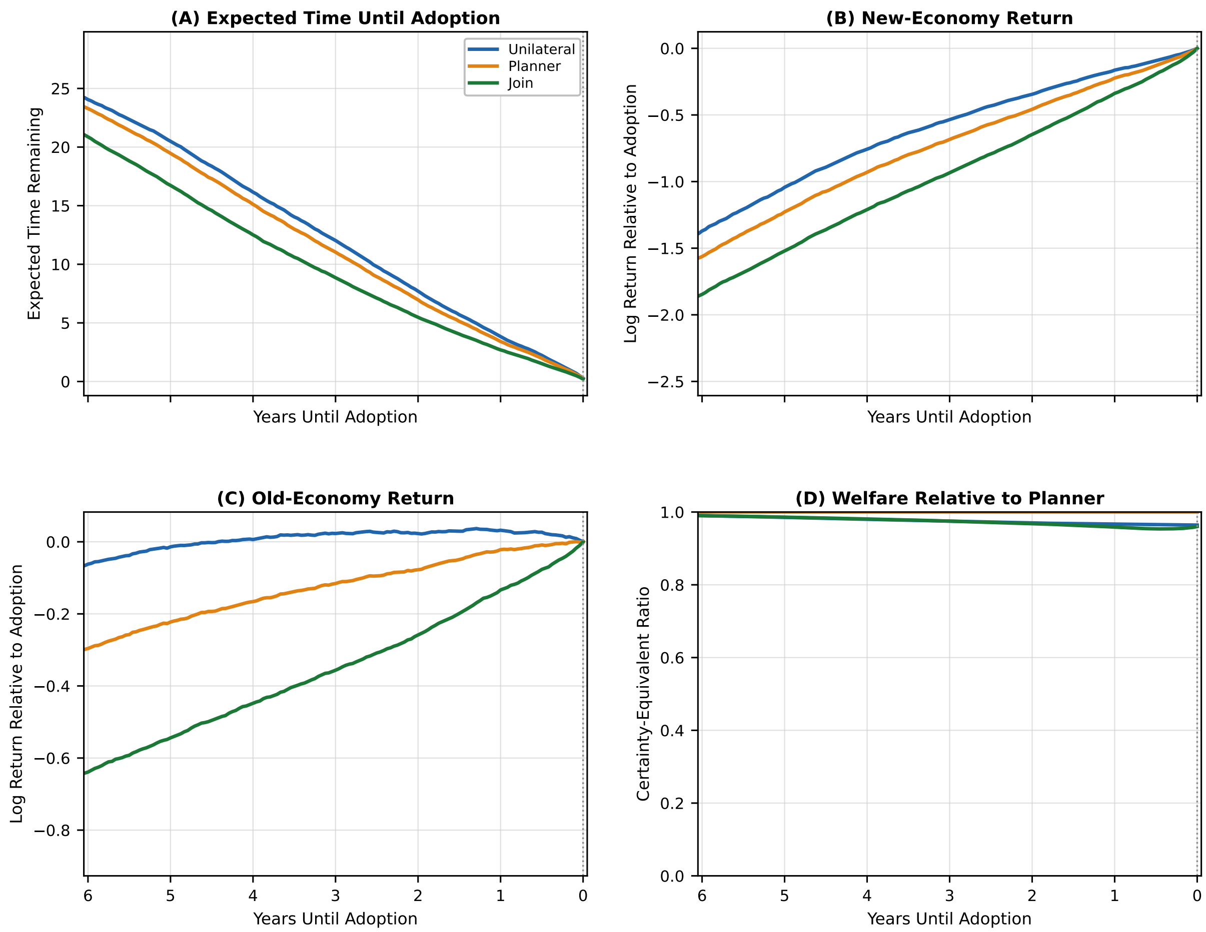

Figure 6: **Asset prices and welfare with endogenous adoption timing.** Panel A shows the physical-measure expected time remaining until adoption under the unilateral, planner, and join rules. Panels B and C show cumulative new- and old-economy log returns, normalized to zero at each path's final pre-adoption observation before averaging. All panels are aligned in event time, with zero denoting the final pre-adoption observation. Panel D shows the certainty-equivalent welfare ratio relative to the planner. For each rule, μ_J is calibrated using the adoption-frequency target described in the text. The vertical dotted line denotes the final pre-adoption observation. Simulations and asset-pricing objects use a semi-monthly time grid. Parameters are $\gamma = 4$, $\kappa = 0.10$, $\sigma = \sigma_{N,0} = \sigma_{N,1} = 0.05$, $\sigma_J = 0.04$, $\phi = 0.3551$, and $T = 30$.

Lag costs. We first solve for the penalty required to make a firm join the aggregate wave at $b_{\gamma=0}$. Using the value notation from Section 3.1, the minimum lag penalty is

$$c^{lag}_{\gamma=0}(t) = \max\left\{0, \log V^{0|1}(t, b_{\gamma=0}(t); b_{\gamma=0}) - \log V^{1|1}(t, b_{\gamma=0}(t); b_{\gamma=0})\right\}.$$

The penalty is large because a firm would otherwise remain in the old economy, continue learning, and preserve the option to join later. It is almost the firm's entire value near the beginning of the horizon and remains about 30 percent at year 15. In the computation, joining becomes more attractive relative to lagging as beliefs rise. Consequently, this same lag-cost schedule can also support aggregate boundaries above $b_{\gamma=0}$.

Admissibility with zero lead costs. Evaluated at $b_{\gamma=0}$ in the $\gamma = 4$ economy, the lead constraint satisfies

$$\log V^{1|0}(t, b_{\gamma=0}(t); b_{\gamma=0}) - \log V^{0|0}(t, b_{\gamma=0}(t); b_{\gamma=0}) < 0.$$

Thus a firm already prefers waiting to leading, and a positive lead cost is not necessary. With the lag costs above and zero lead costs, $b_{\gamma=0}$ is an equilibrium, but it is not uniquely selected. Higher boundaries can also be equilibria: firms do not want to lead early, and the lag penalty remains sufficient once a higher boundary is reached. Within the two-parameter boundary family considered in the previous subsection, the endogenous unilateral boundary is the highest member of this zero-lead-cost set when ranked by its year-15 level.

Maximal preemption award. The lag penalty leaves the level of adoption indeterminate: it supports $b_{\gamma=0}$, but it also permits firms to coordinate on later boundaries. The award addresses the remaining selection question—why adopt this early rather than coordinate on a later boundary? We set

$$a_{\gamma=0}(t) = \log V^{0|0}(t, b_{\gamma=0}(t); b_{\gamma=0}) - \log V^{1|0}(t, b_{\gamma=0}(t); b_{\gamma=0}) > 0. \tag{8}$$

This award exhausts the slack in the lead constraint at $b_{\gamma=0}$. It is the largest award consistent with preserving that boundary, conditional on continuation along $b_{\gamma=0}$: a larger award would make leading strictly preferable there and push adoption earlier. The calculation does not establish that every higher boundary is eliminated. It does, however, increase the payoff from moving first and therefore creates preemption pressure against coordination on later boundaries.

We calculate the award recursively, starting at T and moving backward while holding the future aggregate rule at $b_{\gamma=0}$. At each date, equation (8) converts the slack between waiting and leading into the award. The resulting schedule is declining: approximately 2.5 percent of firm value at the beginning of the horizon, 1.5 percent at year 10, and 0.8 percent at year 15. This is consistent with the first-mover award being more pronounced early in the technological revolution. A bubble relies on firms wanting to adopt while uncertainty is high rather than waiting for learning, and an award that declines with uncertainty accomplishes this.

Figure 7 compares the $\gamma = 0$ schedules with those for the planner boundary. Unlike the positive lead cost for the planner, the negative signed lead cost for $b_{\gamma=0}$ is the boundary-preserving preemption award described above. The lag cost is the minimum penalty required to induce firms to join the aggregate wave, and again is larger as firms adopt earlier when systematic risk is higher.

3.3.1 Asset-price and welfare implications

Figure 8 compares the supported $\gamma = 0$ rule with the $\gamma = 4$ planner. We use the same simulation, calibration, and event-time construction as in the previous subsection, aligning the resulting adoption paths on their final pre-adoption observation.

Panel A reports the expected time remaining until adoption under the physical measure. Under the $\gamma = 0$ boundary, adoption is generally far in the future or never occurs, but expected adoption time collapses along paths that ultimately cross the boundary. Panel

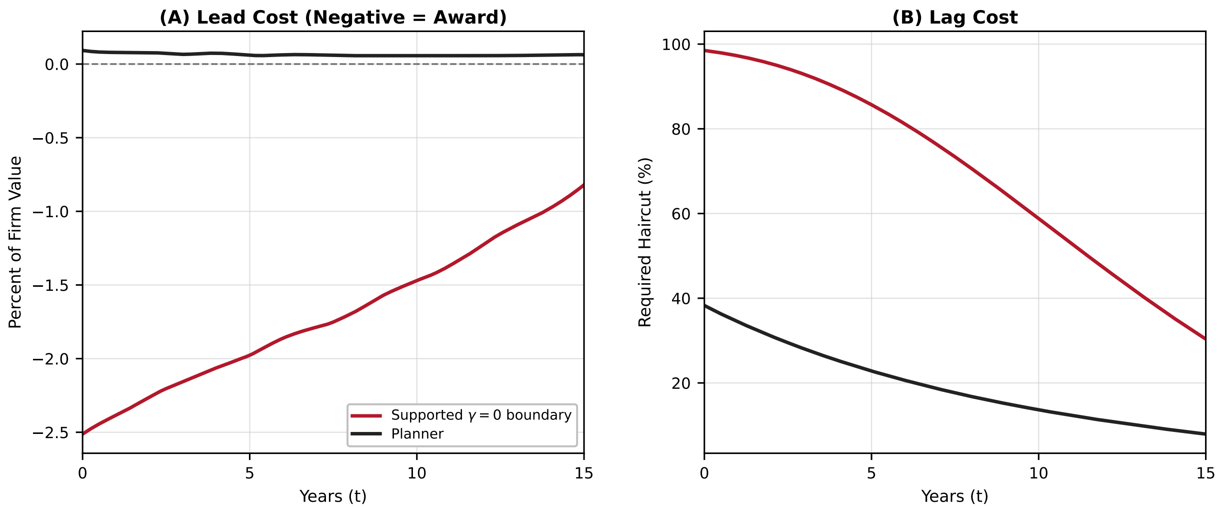

Figure 7: **Lead and lag costs for the $\gamma = 0$ and planner boundaries.** Panel A reports the signed lead cost; negative values are first-mover awards. For the $\gamma = 0$ boundary, the absolute value of the negative line is the largest award consistent with preserving the boundary. Panel B reports the minimum proportional haircut on the value of a firm that refuses to join the aggregate wave and subsequently follows its optimal private join rule. Both boundaries are valued using $\gamma = 4$. Parameters are $\kappa = 0.10$, $\sigma = \sigma_{N,0} = \sigma_{N,1} = 0.05$, $\sigma_J = 0.04$, $\phi = 0.3551$, and $T = 30$.

B shows the corresponding new-economy return. Its value rises as adoption becomes more likely, peaks shortly before the transition, and then falls as the final increase in adoption probability makes new-technology risk systematic. Under the definition in Section 2, the supported equilibrium therefore generates a new-economy bubble. The peak-to-adoption decline is modest in this calibration. We do not emphasize its particular magnitude, which can vary with the calibration. Our central result is that the supported equilibrium generates a downturn at all, a pattern that is otherwise difficult to obtain when adoption timing is endogenous.

Panel C shows the corresponding old-economy return. Old-economy values initially rise, peak well before adoption, and then fall steeply as the transition approaches. This decline is substantially larger than the new-economy reversal in Panel B. The high lag costs prevent firms from escaping the aggregate wave once the boundary is reached.

Panel D shows the resulting welfare loss. Welfare falls to roughly 53 percent of the planner benchmark at adoption. The first-mover award creates pressure to front-run later adoption boundaries, while the lag penalty locks firms into the aggregate transition. Firms collectively adopt at a state in which the planner places substantial value on continued waiting. The equilibrium force that produces the bubble therefore destroys both old-economy firm value and aggregate welfare.

In summary, the supported equilibrium generates a new-economy bubble with a modest award for moving first, under three percent of firm value, and a large penalty for being left behind that approaches the firm's entire value early in the horizon and remains near a third of it at year 15. Together they turn adoption into a race for the award. Every firm therefore moves at $b_{\gamma=0}$, the penalty keeps it there, and every firm is worse off for it. Adoption occurs while substantial uncertainty remains, the option to wait and learn is destroyed, old-economy values fall more sharply than new-economy values, and aggregate welfare falls to roughly half the planner benchmark. The same adoption race that produces the bubble makes it costly for both social welfare and the adopting firms.

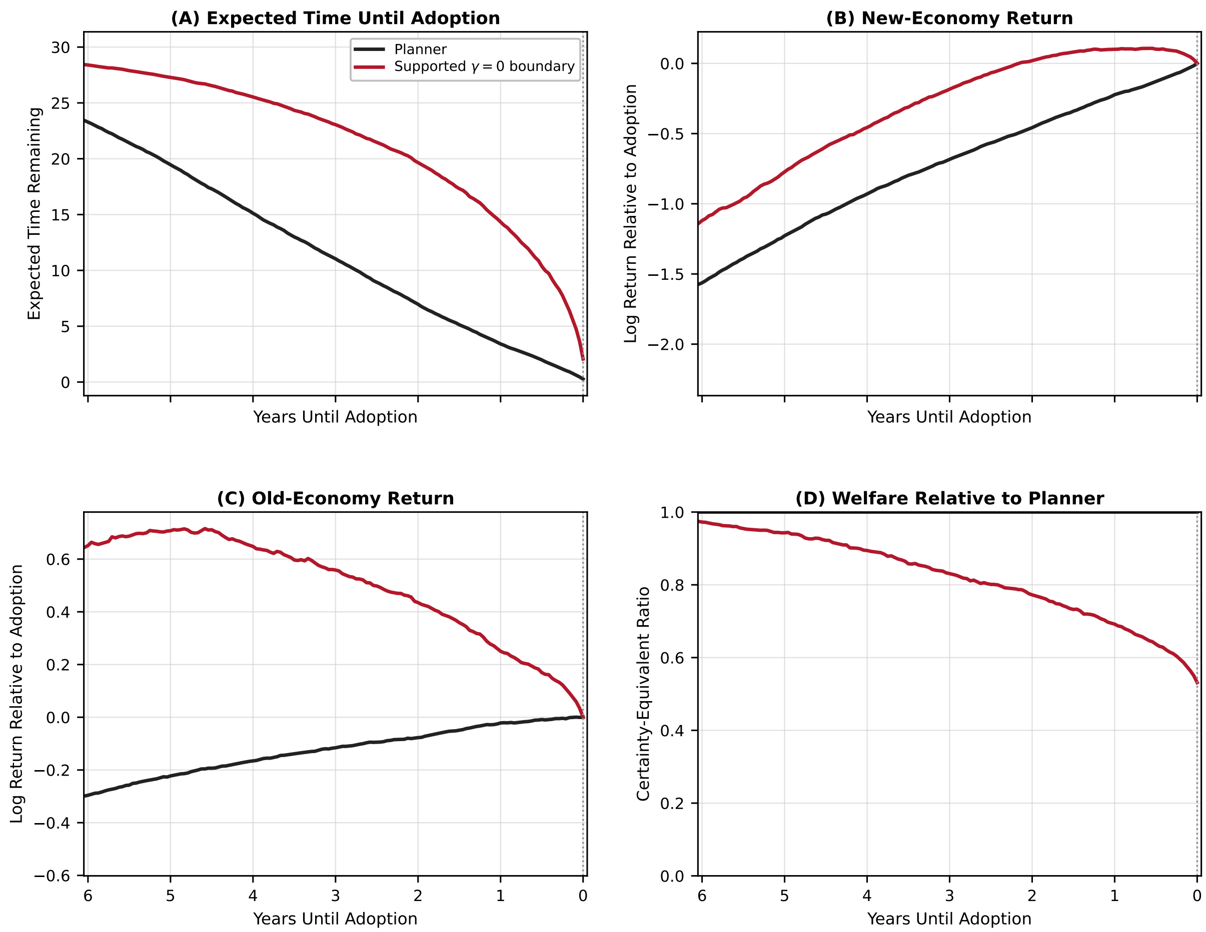

Figure 8: **Asset prices and welfare under the implemented $\gamma = 0$ boundary.** Panel A shows the physical-measure expected time remaining until adoption under the planner and $\gamma = 0$ rules. Panels B and C show cumulative new- and old-economy log returns, normalized to zero at each path's final pre-adoption observation before averaging. Panel D shows the certainty-equivalent welfare ratio relative to the planner. For each rule, μ_J is calibrated using the adoption-frequency target described in the text. Event time zero denotes the final pre-adoption observation. Simulations and asset-pricing objects use a semi-monthly grid. Parameters are $\gamma = 4$, $\kappa = 0.10$, $\sigma = \sigma_{N,0} = \sigma_{N,1} = 0.05$, $\sigma_J = 0.04$, $\phi = 0.3551$, and $T = 30$.

4 Conclusion

In this paper, we reexamine stock-price bubbles during technological revolutions in the framework of Pástor and Veronesi (2009). In that framework, positive shocks to new-technology productivity raise expected cash flows, but also increase the likelihood of adoption, making new-technology risk systematic and raising discount rates. We show that the bubble depends critically on how the adoption decision is modeled.

We first revisit the planner's endogenous adoption problem already considered by Pástor and Veronesi (2009). Although their discussion suggests that the bubble mechanism survives when adoption timing is chosen optimally, we find that the endogenous stopping rule substantially smooths expected adoption. The planner's option to wait and learn prevents the sharp increase in systematic risk needed to overturn the cash-flow effect, and the bubble disappears.

We then extend the model to decentralized adoption and allow firms to make separate adoption decisions. At PV's fixed adoption date, even arbitrarily small heterogeneity in conversion costs generates partial adoption. Low-cost firms have lower adoption thresholds, but the threshold rises with the adoption fraction as new-technology risk becomes increasingly systematic. The resulting spread in adoption thresholds persists even as heterogeneity vanishes, so the repricing of systematic risk is dispersed rather than concentrated in a single economy-wide threshold, and the bubble therefore disappears. Moreover, the decentralized adoption fraction coincides with the planner's optimal partial-adoption allocation.

When decentralized adoption timing is endogenous, the bubble remains absent even when deviation costs enforce simultaneous adoption. A potential first mover may value unresolved idiosyncratic risk, but it anticipates that subsequent adoption by others will make its exposure systematic. This reduces the incentive to move early and preserves the value of waiting for uncertainty to resolve. Expected adoption time therefore adjusts gradually with news, as it does under the planner. Strong preemption and coordination incentives can restore an abrupt transition and revive the bubble, but only through inefficiently early adoption that destroys both adopting firm value and welfare.

Our results therefore change how the bubble mechanism in Pástor and Veronesi (2009) should be interpreted. A stock-price bubble generated by technological adoption is not a benign implication of learning about a socially valuable technology. It requires a competitive environment that induces firms to adopt before uncertainty has resolved and causes relatively small positive productivity shocks to trigger a sudden coordinated wave of simultaneous adoption that reduces both social welfare and adopting-firm value. For technologies such as AI, the relevant question is therefore not simply whether adoption will become widespread, but whether the competitive environment can generate the necessary unanticipated and value-destroying simultaneous adoption wave.

Appendix

A Expected adoption time approximation

This section derives the expected-adoption-time approximation and explains why its natural probability measure differs between fixed and endogenous adoption timing. 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

$$\log M_t^N(s) = \log M_t^{NA} - \gamma A_t A_s \hat{\sigma}_t^2. \tag{9}$$

Defining

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

equation (9) 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

$$\log M_t^N \approx \log M_t^{IA} + \frac{E_t^{\mathbb{M}}[t^{**}] - t}{T-t} \left(\log M_t^{NA} - \log M_t^{IA}\right), \tag{10}$$

where the appropriate measure $\mathbb{M}$ depends on how adoption timing is determined.

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,

$$M_t^N = \frac{E_t^P[H_\tau]}{G_t}. \tag{11}$$

Thus an expectation over endogenous stopping dates is naturally a physical expectation. In particular, 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. To see this, 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. Hence the discount-rate component associated with the timing of adoption is averaged over the physical distribution of first-passage times, motivating $\mathbb{M} = P$ in (10).

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 fixed-date adoption selects among states at a common horizon, making risk-neutral probabilities natural, whereas endogenous first-passage adoption selects the date at which a boundary state is reached, making physical hitting-time probabilities natural. This distinction concerns the one-dimensional approximation in (10); exact prices in both models are generated by the same state-price density.

B Proofs

B.1 Exponential-affine building blocks

Let $B^0_{i,T}$ denote firm i's terminal book value under the old technology, and let ψ denote the realized productivity gain. If the firm adopts at $\bar{t}^{**}$, its terminal book value is instead $B^1_{i,T} = (1-\kappa)B^0_{i,T}\exp(A_2(\tau)\psi)$, where $\tau = T-\bar{t}^{**}$. This relation holds state by state, because adoption shifts only the drift anchor of the firm's productivity from $\bar{\rho}$ to $\bar{\rho}+\psi$ while the shock path is unchanged, so the productivity gap between the two technologies accumulates deterministically to $A_2(\tau)\psi$ by time T. The regime-r kernel factors in the same way. Aggregate terminal wealth is $W_T = (1-\kappa)^r B^0_T \exp(rA_2(\tau)\psi)$, so $\pi^r_{\bar{t}^{**},T} \propto (B^0_T)^{-\gamma}\exp(-\gamma r A_2(\tau)\psi)$, the product of a component driven by old-economy shocks and a ψ-component. Because ψ is independent of the old-economy shocks, the old-economy components factor out of numerator and denominator and cancel from the ratio $V^{a|r}/V^{0|r} = E_{\bar{t}^{**}}[\pi^r_{\bar{t}^{**},T}B^a_{i,T}]/E_{\bar{t}^{**}}[\pi^r_{\bar{t}^{**},T}B^0_{i,T}]$, leaving $(1-\kappa)^a\, E_{\bar{t}^{**}}[e^{(a-\gamma r)A_2(\tau)\psi}]/E_{\bar{t}^{**}}[e^{-\gamma r A_2(\tau)\psi}]$. Integrating over $\psi \mid \mathcal{F}_{\bar{t}^{**}} \sim N(\hat{\psi}_{\bar{t}^{**}}, \hat{\sigma}^2_{\bar{t}^{**}})$ then gives, for $a, r \in \{0,1\}$,

$$\frac{V^{a|r}}{V^{0|r}} = \frac{E_{\bar{t}^{**}}\left[\pi^r_{\bar{t}^{**},T}B^a_{i,T}\right]}{E_{\bar{t}^{**}}\left[\pi^r_{\bar{t}^{**},T}B^0_{i,T}\right]} = (1-\kappa)^a \exp\Big(aA_2(\tau)\hat{\psi}_{\bar{t}^{**}} + \tfrac{a}{2}\big(a-2\gamma r\big)A_2(\tau)^2\hat{\sigma}^2_{\bar{t}^{**}}\Big). \qquad (*)$$

Taking logs and setting $a=1$ gives the lead and lag cost building blocks

$$D^0(\bar{t}^{**},\hat{\psi}) \equiv \log V^{1|0} - \log V^{0|0} = \log(1-\kappa) + A_2(\tau)\hat{\psi} + \tfrac{1}{2}A_2(\tau)^2\hat{\sigma}^2_{\bar{t}^{**}},$$

$$D^1(\bar{t}^{**},\hat{\psi}) \equiv \log V^{1|1} - \log V^{0|1} = \log(1-\kappa) + A_2(\tau)\hat{\psi} - \tfrac{1}{2}(2\gamma-1)A_2(\tau)^2\hat{\sigma}^2_{\bar{t}^{**}}, \qquad (**)$$

so that $c^{lead} = \max\{0, D^0\}$ and $c^{lag} = \max\{0, -D^1\}$. The opposite sign on the variance term across the two cases is the source of the wedge between the join and unilateral thresholds.

A second building block is the planner's welfare and marginal-utility comparison, which uses the representative agent's terminal wealth $W_T = B_T$ directly rather than an individual firm's claim. If the aggregate economy adopts at $\bar{t}^{**}$, then relative to no adoption,

$$\frac{E_{\bar{t}^{**}}\left[W_T^{1-\gamma} \mid \text{adopt}\right]}{E_{\bar{t}^{**}}\left[W_T^{1-\gamma} \mid \text{no adopt}\right]} = (1-\kappa)^{1-\gamma} \exp\Big((1-\gamma)A_2(\tau)\hat{\psi}_{\bar{t}^{**}} + \tfrac{1}{2}(1-\gamma)^2 A_2(\tau)^2 \hat{\sigma}^2_{\bar{t}^{**}}\Big),$$

$$\frac{E_{\bar{t}^{**}}\left[W_T^{-\gamma} \mid \text{adopt}\right]}{E_{\bar{t}^{**}}\left[W_T^{-\gamma} \mid \text{no adopt}\right]} = (1-\kappa)^{-\gamma} \exp\Big(-\gamma A_2(\tau)\hat{\psi}_{\bar{t}^{**}} + \tfrac{1}{2}\gamma^2 A_2(\tau)^2 \hat{\sigma}^2_{\bar{t}^{**}}\Big).$$

The planner threshold sets the first ratio equal to one; the second ratio gives the corresponding change in marginal utility (and hence in the pricing kernel) that adoption induces.

B.2 Proof of Proposition 1

Thresholds. The *join* threshold is the value of $\hat{\psi}_{\bar{t}^{**}}$ at which an individual firm is indifferent between joining aggregate adoption and lagging behind once the rest of the economy has adopted, i.e. $E_{\bar{t}^{**}}[\pi^1_{\bar{t}^{**},T} B^1_{i,T}] = E_{\bar{t}^{**}}[\pi^1_{\bar{t}^{**},T} B^0_{i,T}]$, or $D^1 = 0$. By (**) this is

$$\log(1-\kappa) + A_2(\tau)\bar{\psi}_{J,\bar{t}^{**}} - \tfrac{1}{2}(2\gamma-1)A_2(\tau)^2\hat{\sigma}^2_{\bar{t}^{**}} = 0 \quad\Longleftrightarrow\quad \bar{\psi}_{J,\bar{t}^{**}} = \frac{-\log(1-\kappa) + \frac{1}{2}(2\gamma-1)A_2(\tau)^2\hat{\sigma}^2_{\bar{t}^{**}}}{A_2(\tau)}.$$

The *unilateral* threshold is the value of $\hat{\psi}_{\bar{t}^{**}}$ at which a single firm is indifferent between adopting alone and remaining in the old technology, taking the old-economy kernel as given, i.e. $E_{\bar{t}^{**}}[\pi^0_{\bar{t}^{**},T} B^1_{i,T}] = E_{\bar{t}^{**}}[\pi^0_{\bar{t}^{**},T} B^0_{i,T}]$, or $D^0 = 0$. By (**) this is

$$\log(1-\kappa) + A_2(\tau)\bar{\psi}_{U,\bar{t}^{**}} + \tfrac{1}{2}A_2(\tau)^2\hat{\sigma}^2_{\bar{t}^{**}} = 0 \quad\Longleftrightarrow\quad \bar{\psi}_{U,\bar{t}^{**}} = \frac{-\log(1-\kappa) - \frac{1}{2}A_2(\tau)^2\hat{\sigma}^2_{\bar{t}^{**}}}{A_2(\tau)}.$$

The *planner* threshold sets the welfare ratio of the previous subsection equal to one:

$$(1-\gamma)\log(1-\kappa) + (1-\gamma)A_2(\tau)\bar{\psi}_{planner,\bar{t}^{**}} + \tfrac{1}{2}(1-\gamma)^2 A_2(\tau)^2 \hat{\sigma}^2_{\bar{t}^{**}} = 0.$$

Dividing by $(1-\gamma)$ gives $\log(1-\kappa) + A_2(\tau)\bar{\psi}_{planner,\bar{t}^{**}} + \frac{1}{2}(1-\gamma)A_2(\tau)^2\hat{\sigma}^2_{\bar{t}^{**}} = 0$, i.e.

$$\bar{\psi}_{planner,\bar{t}^{**}} = \frac{-\log(1-\kappa) + \frac{1}{2}(\gamma-1)A_2(\tau)^2\hat{\sigma}^2_{\bar{t}^{**}}}{A_2(\tau)}.$$

The planner internalizes the increased aggregate variance generated by adoption, so with $\gamma > 1$ a higher $\hat{\sigma}^2_{\bar{t}^{**}}$ raises the threshold.

Support costs. By the definitions of Section 3.1, $c^{lead} = \max\{0, D^0\}$ and $c^{lag} = \max\{0, -D^1\}$, evaluated at the candidate threshold using $(**)$.

At the join threshold, $D^1(\bar{\psi}_{J,\bar{t}^{**}}) = 0$ by construction, so $c_J^{lag} = 0$. To evaluate D^0 at $\bar{\psi}_{J,\bar{t}^{**}}$, substitute $\log(1-\kappa) + A_2(\tau)\bar{\psi}_{J,\bar{t}^{**}} = \frac{1}{2}(2\gamma-1)A_2(\tau)^2\hat{\sigma}^2$ into the first expression in $(**)$:

$$D^0(\bar{\psi}_{J,\bar{t}^{**}}) = \tfrac{1}{2}(2\gamma-1)A_2(\tau)^2\hat{\sigma}^2 + \tfrac{1}{2}A_2(\tau)^2\hat{\sigma}^2 = \gamma A_2(\tau)^2\hat{\sigma}^2 > 0,$$

so $c_J^{lead} = \gamma A_2(\tau)^2\hat{\sigma}^2$.

Symmetrically, at the unilateral threshold $D^0(\bar{\psi}_{U,\bar{t}^{**}}) = 0$ by construction, so $c_U^{lead} = 0$. Substituting $\log(1-\kappa) + A_2(\tau)\bar{\psi}_{U,\bar{t}^{**}} = -\frac{1}{2}A_2(\tau)^2\hat{\sigma}^2$ into the second expression in $(**)$,

$$D^1(\bar{\psi}_{U,\bar{t}^{**}}) = -\tfrac{1}{2}A_2(\tau)^2\hat{\sigma}^2 - \tfrac{1}{2}(2\gamma-1)A_2(\tau)^2\hat{\sigma}^2 = -\gamma A_2(\tau)^2\hat{\sigma}^2 < 0,$$

so $c_U^{lag} = -D^1(\bar{\psi}_{U,\bar{t}^{**}}) = \gamma A_2(\tau)^2\hat{\sigma}^2$, exactly equal to c_J^{lead}. This is why the join and unilateral equilibria require support costs of opposite type but identical static magnitude.

Finally, at the planner's threshold, $\log(1-\kappa) + A_2(\tau)\bar{\psi}_{planner,\bar{t}^{**}} = \frac{1}{2}(\gamma-1)A_2(\tau)^2\hat{\sigma}^2$.

Substituting into $(**)$,

$$D^0(\bar{\psi}_{planner,\bar{t}^{**}}) = \tfrac{1}{2}(\gamma-1)A_2(\tau)^2\hat{\sigma}^2 + \tfrac{1}{2}A_2(\tau)^2\hat{\sigma}^2 = \tfrac{1}{2}\gamma A_2(\tau)^2\hat{\sigma}^2 > 0,$$

$$D^1(\bar{\psi}_{planner,\bar{t}^{**}}) = \tfrac{1}{2}(\gamma-1)A_2(\tau)^2\hat{\sigma}^2 - \tfrac{1}{2}(2\gamma-1)A_2(\tau)^2\hat{\sigma}^2 = -\tfrac{1}{2}\gamma A_2(\tau)^2\hat{\sigma}^2 < 0,$$

so decentralizing the planner's threshold requires both a lead and a lag cost,

$$c^{lead}_{planner} = c^{lag}_{planner} = \tfrac{1}{2}\gamma A_2(\tau)^2\hat{\sigma}^2,$$

each equal to half of $c^{lead}_J = c^{lag}_U$. This completes the proof.

B.3 Proof of Proposition 2

Partial-adoption aggregation and corner conditions. Index firms by $x \in [0,1]$, let

$$\kappa_\epsilon(x) = \kappa + \epsilon(2x-1), \qquad K_\epsilon(q) = \int_0^q [1-\kappa_\epsilon(x)]\,dx,$$

and suppose the lowest-cost fraction q adopts. The partial-adoption aggregation used in the text is

$$W_T(q) = B^0_T H_\epsilon(q,\psi), \qquad H_\epsilon(q,\psi) = 1-q+K_\epsilon(q)e^{A_2(\tau)\psi}.$$

Define the marginal gain

$$\mathcal{G}_\epsilon(q,\hat{\psi}) \equiv E_{\bar{t}^{**}}\left[H_\epsilon(q,\psi)^{-\gamma}\left\{[1-\kappa_\epsilon(q)]e^{A_2(\tau)\psi}-1\right\}\right]. \tag{12}$$

Assuming the marginal gain crosses zero once, the adoption fraction is characterized explicitly by

$$q^* = \begin{cases} 0, & \mathcal{G}_\epsilon(0,\hat{\psi}) \le 0, \\ q \in (0,1) \text{ satisfying } \mathcal{G}_\epsilon(q,\hat{\psi}) = 0, & \mathcal{G}_\epsilon(0,\hat{\psi}) > 0 > \mathcal{G}_\epsilon(1,\hat{\psi}), \\ 1, & \mathcal{G}_\epsilon(1,\hat{\psi}) \ge 0. \end{cases}$$

Without the single-crossing assumption, the same inequalities describe the corners, while an interior equilibrium is any stable zero of $\mathcal{G}_\epsilon$; the planner selects the global welfare maximum.

We first prove the efficiency statement without using the partial-adoption approximation. Let $W_T(q)$ denote exact aggregate terminal wealth when the lowest-cost fraction q adopts, and let

$$\Delta_T(q) \equiv \left.\frac{\partial W_T(q)}{\partial q}\right|_q$$

denote the exact increment to terminal aggregate wealth from converting the marginal firm. This notation permits all shocks and nonlinearities in the exact model; no exponential-affine representation is required. Conditional on the information available at $\bar{t}^{**}$, the planner chooses q to maximize

$$E_{\bar{t}^{**}}[u(W_T(q))].$$

At an interior optimum its marginal condition is

$$E_{\bar{t}^{**}}[u'(W_T(q))\Delta_T(q)] = 0. \tag{13}$$

For power utility, $u'(W_T) = W_T^{-\gamma}$.

Now consider the marginal, atomistic firm. Its adoption changes its own terminal payoff by $\Delta_T(q)$, but has a measure-zero effect on aggregate wealth and therefore takes the equilibrium stochastic discount factor as given. Up to a positive normalization known at $\bar{t}^{**}$, that pricing kernel is

$$\pi_{\bar{t}^{**},T}(q) = \frac{u'(W_T(q))}{E_{\bar{t}^{**}}[u'(W_T(q))]}.$$

Consequently the firm's value gain from adoption has the sign of

$$E_{\bar{t}^{**}}[\pi_{\bar{t}^{**},T}(q)\Delta_T(q)] \quad \text{and hence the sign of} \quad E_{\bar{t}^{**}}[u'(W_T(q))\Delta_T(q)]\,. \tag{14}$$

Equations (13) and (14) are the same marginal condition. The corresponding inequalities give the same corner choices at $q = 0$ and $q = 1$. Thus the competitive and planner adoption fractions coincide in the exact model. The argument requires that the marginal firm's private payoff increment is also its contribution to aggregate wealth and that adoption has no separate technological, network, learning, or informational externality. It does not require a closed-form solution for $W_T(q)$.

We next establish the closed-form schedule statements under the partial-adoption aggregation used in the text. There,

$$W_T(q) = B_T^0 H_\epsilon(q,\psi), \qquad \frac{\partial H_\epsilon(q,\psi)}{\partial q} = [1-\kappa_\epsilon(q)]e^{A_2(\tau)\psi} - 1,$$

where the second equality uses $K_\epsilon'(q) = 1 - \kappa_\epsilon(q)$. Differentiating expected utility therefore gives, up to the positive factor $(B_T^0)^{1-\gamma}$,

$$E_{\bar{t}^{**}}\left[H_\epsilon(q,\psi)^{-\gamma}\left\{[1-\kappa_\epsilon(q)]e^{A_2(\tau)\psi} - 1\right\}\right],$$

which is exactly $\mathcal{G}_\epsilon(q,\hat{\psi})$ in (12). Pricing the marginal firm's payoff gain with the same kernel yields the identical expression, proving the first statement of the proposition in this aggregation.

When $\epsilon = 0$, $H_0(0,\psi) = 1$, so the first adopter is priced by the all-old kernel and its zero-value-gain condition is the unilateral condition; hence $\bar{\psi}_0(0) = \bar{\psi}_U$. At the other endpoint, $H_0(1,\psi) = (1-\kappa)e^{A_2(\tau)\psi}$, so the final adopter is priced by the all-adopted kernel and $\bar{\psi}_0(1) = \bar{\psi}_J$. At $q = 1/2$, the lognormal payoff symmetry in this aggregation makes the marginal condition equal to the all-or-nothing planner indifference condition, yielding

$\bar{\psi}_0(1/2) = \bar{\psi}_{planner}$. Finally, subtracting the unilateral threshold from the join threshold in Proposition 1 gives

$$\bar{\psi}_0(1) - \bar{\psi}_0(0) = \bar{\psi}_J - \bar{\psi}_U = \gamma A_2(\tau)\hat{\sigma}^2.$$

The endpoint, midpoint, and width identities use this particular aggregation; unlike the efficiency result, they are not asserted as general closed-form properties of the exact partial-adoption model.

References

Andrei, Daniel and Bruce I Carlin, "Schumpeterian competition in a Lucas economy," *Journal of Economic Theory*, 2023, *208*, 105613.

Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, "A Model of Investor Sentiment," *Journal of Finance*, 1998, *53* (2), 307–343.

Bonaparte, Yosef, "Artificial intelligence in finance: Valuations and opportunities," *Finance Research Letters*, 2024, *60*, 104851.

Carlson, Murray, Adlai Fisher, and Ron Giammarino, "Corporate Investment and Asset Price Dynamics: Implications for the Cross-Section of Returns," *Journal of Finance*, 2004, *59* (6), 2577–2603.

Eisfeldt, Andrea L, Gregor Schubert, and Miao Ben Zhang, "Generative AI and firm values," Technical Report, National Bureau of Economic Research 2023.

FARFE, "Prize Citation for the 9th Stephen A. Ross Prize in Financial Economics," `https://farfe.org/FARFE_9th_Ross_Prize_Citation.pdf` 2025. Unpublished award citation document.

Fudenberg, Drew and Jean Tirole, "Preemption and Rent Equalization in the Adoption of New Technology," *Review of Economic Studies*, 1985, *52*, 383–401.

Garleanu, Nicolae, Stavros Panageas, and Jianfeng Yu, "Technological growth and asset pricing," *Journal of Finance*, 2012, *67* (4), 1265–1292.

Giglio, Stefano and Tiago Severo, "Intangible Capital, Relative Asset Shortages and Bubbles," *Journal of Monetary Economics*, 2012, *59* (3), 303–317.

Grenadier, Steven R., "The Strategic Exercise of Options: Development Cascades and Overbuilding in Real Estate Markets," *Journal of Finance*, 1996, *51*, 1653–1679.

_ , "Option Exercise Games: An Application to the Equilibrium Investment Strategies of Firms," *Review of Financial Studies*, 2002, *15*, 691–721.

_ **and Allen M. Weiss**, "Investment in Technological Innovations: An Option Pricing Approach," *Journal of Financial Economics*, 1997, *44*, 397–416.

Harrison, J Michael and David M Kreps, "Speculative investor behavior in a stock market with heterogeneous expectations," *The Quarterly Journal of Economics*, 1978, *92* (2), 323–336.

Karatzas, Ioannis and Steven E Shreve, *Methods of mathematical finance*, Vol. 39, Springer, 1998.

Kogan, Leonid, Dimitris Papanikolaou, and Noah Stoffman, "Left behind: Creative destruction, inequality, and the stock market," *Journal of Political Economy*, 2020, *128* (3), 855–906.

Kung, Howard and Lukas Schmid, "Innovation, growth, and asset prices," *Journal of Finance*, 2015, *70* (3), 1001–1037.

Lambrecht, Bart and William Perraudin, "Real Options and Preemption under Incomplete Information," *Journal of Economic Dynamics and Control*, 2003, *27*, 619–643.

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.

Reinganum, Jennifer F., "On the Diffusion of New Technology: A Game Theoretic Approach," *Review of Economic Studies*, 1981, *48*, 395–405.

Scheinkman, José A. and Wei Xiong, "Overconfidence and Speculative Bubbles," *Journal of Political Economy*, 2003, *111* (6), 1183–1219.

Wallace, Matthew C., "This Is How the AI Stock Boom Plays Out," *Bloomberg Opinion*, October 2025.

Weeds, Helen, "Strategic Delay in a Real Options Model of R&D Competition," *Review of Economic Studies*, 2002, *69*, 729–747.

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.

A Additional Perspective on the Model

This section supplements Figure 3 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 Figure 3 by true technology quality, $\psi = k\sigma_J$ for $k = 1, \dots, 6$. 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.

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:

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

which produces the weighted log-price formula

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

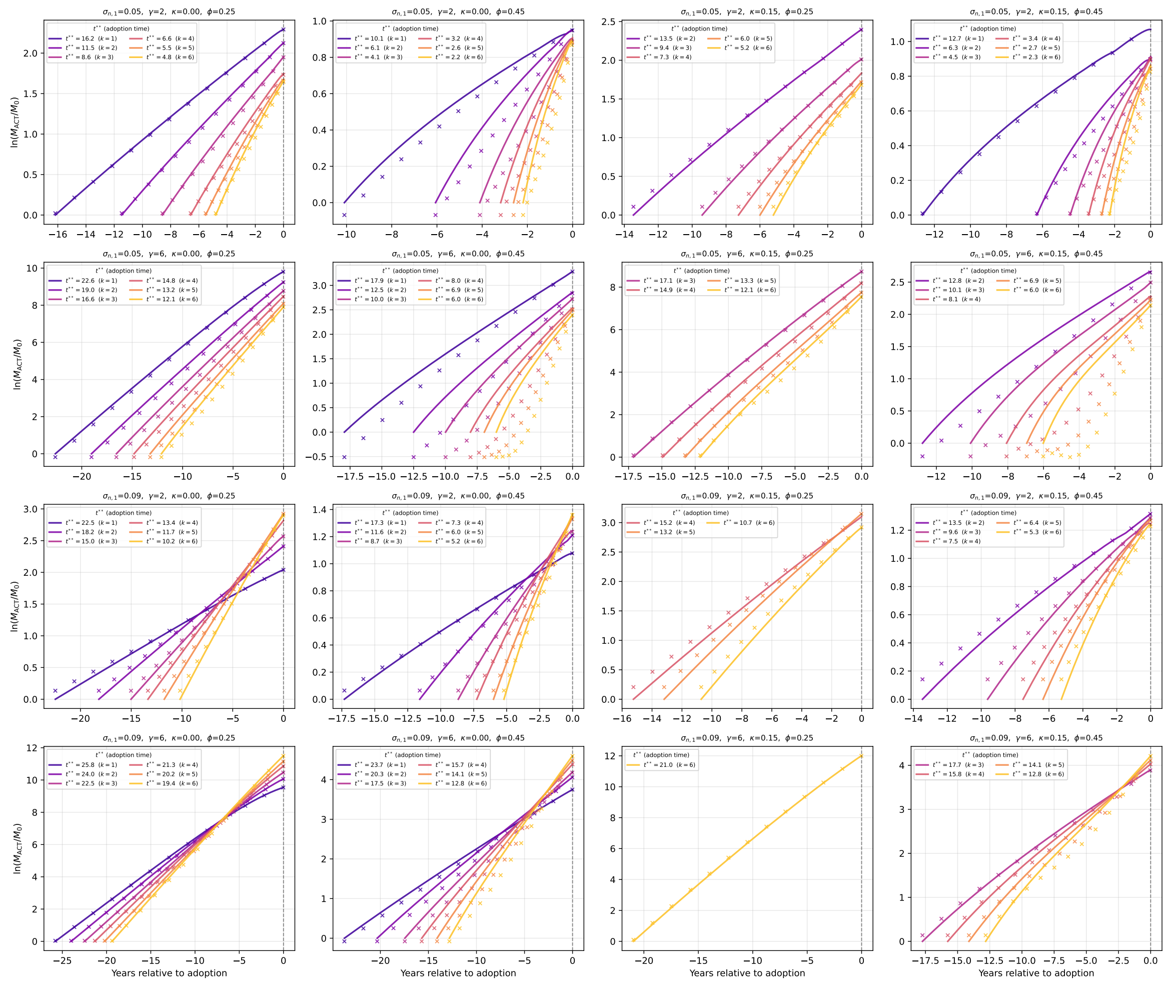

Figure A1: **Cumulative unexpected returns in revolutions across parameterizations: Various adoption times** Each colored line shows the log price path $\log(M_t^{ACT}/M_0)$ in the endogenous adoption-time model for a different value of the true technology quality $\psi_{\text{true}} = k \cdot \sigma_J$, $k = 1, 2, \ldots, 6$. 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.

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, the positive productivity shocks $d\tilde{Z}_{1,t}$ that push $\hat{\psi}_t$ toward the adoption generate large negative price effects near the boundary through the rising adoption discount when $\hat{\sigma}_t$ is still large further attenuating the cash-flow effect. 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.

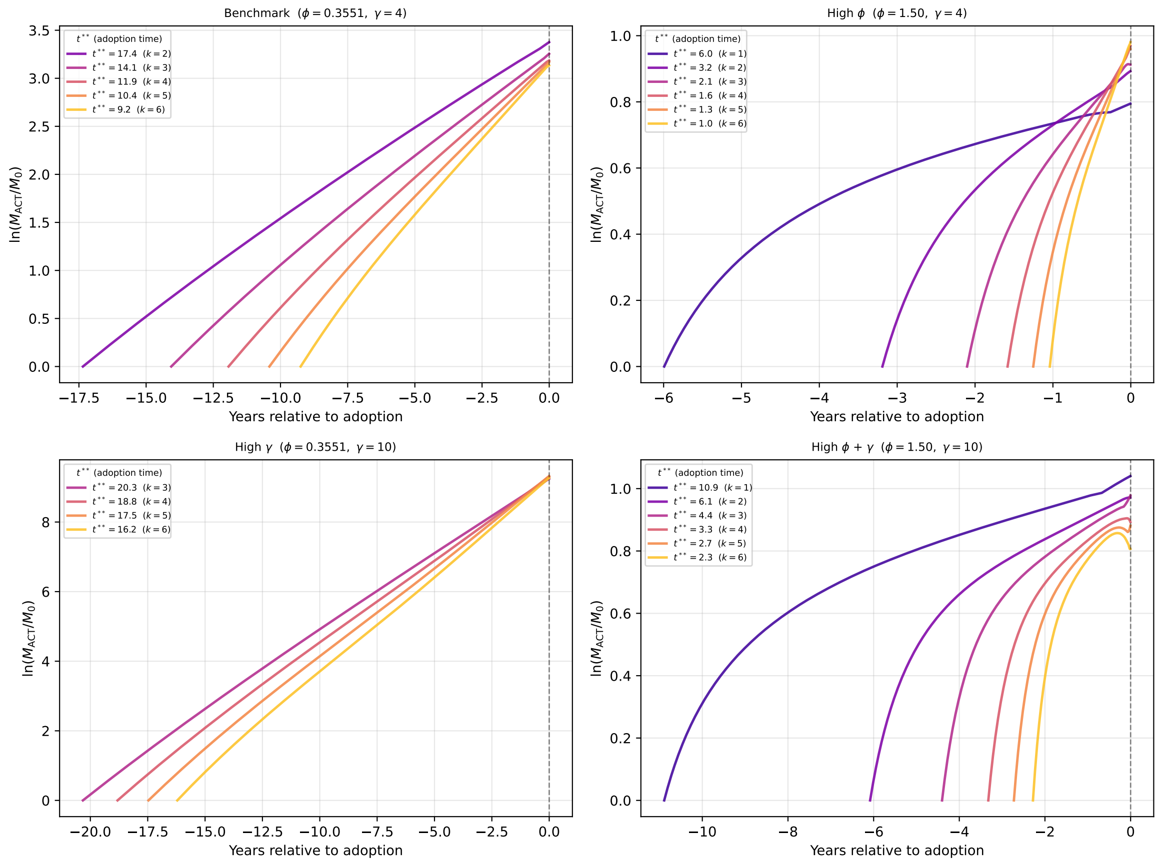

Figure A2: **Deterministic price paths: 2×2 parameterization grid (endogenous model).** Each panel plots $\log(M_t^{ACT}/M_0)$ paths for $\psi_{\text{true}} = k \cdot \sigma_J$, $k = 1, \ldots, 6$ (colored lines), as in Figure A1. The four panels vary the speed of mean reversion ϕ and the coefficient 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.

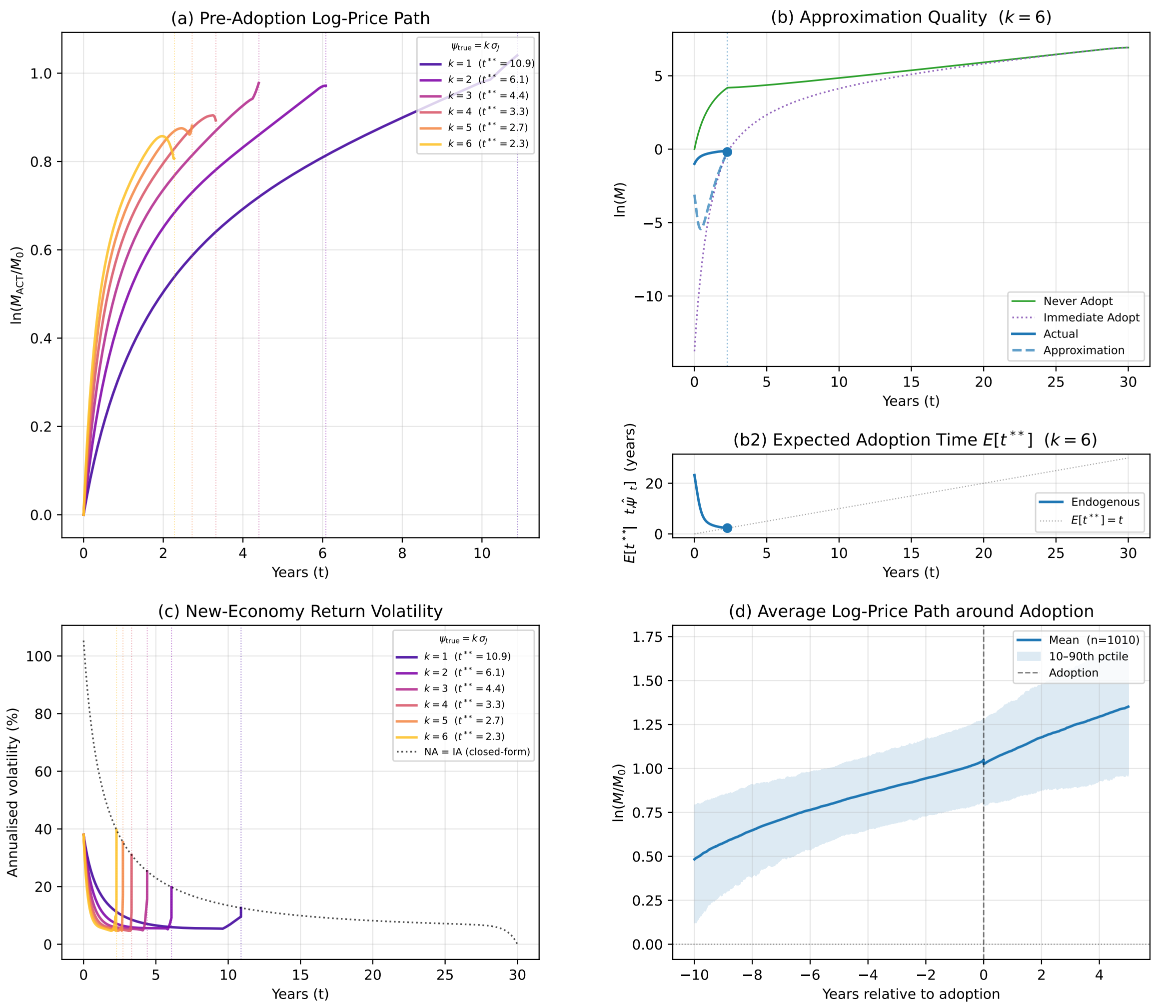

Figure A3: **High-ϕ, high-γ parameterization: price paths, approximation quality, volatility, and Monte Carlo average (endogenous model, $\phi = 1.50$, $\gamma = 10$).** Panel (a): $\log(M_t^{ACT}/M_0)$ for $\psi_{\text{true}} = k \cdot \sigma_J$, $k = 1, \ldots, 6$ (colored lines). Vertical dotted lines 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 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

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

and the posterior mean has innovation volatility

$$d\widehat{\psi}_t = \frac{v_t \phi}{\sigma}\, d\widetilde{Z}_t^N. \tag{2}$$

Here ϕ controls mean reversion and σ is the volatility of the technology-specific productivity shock. Thus the local variance and covariance rates of the productivity and belief innovations are

$$\frac{\mathrm{Var}_t(d\rho_t^N)}{dt} = \sigma^2, \qquad \frac{\mathrm{Var}_t(d\widehat{\psi}_t)}{dt} = \left(\frac{v_t\phi}{\sigma}\right)^2, \qquad \frac{\mathrm{Cov}_t(d\rho_t^N, d\widehat{\psi}_t)}{dt} = v_t\phi. \tag{3}$$

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

$$Y_{n+1} = \psi + \varepsilon_{n+1}, \qquad \varepsilon_{n+1} \sim \mathcal{N}(0, R_\lambda), \qquad R_\lambda = \lambda\left(\frac{\sigma}{\phi}\right)^2, \tag{4}$$

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.

$$\lambda R_\lambda^{-1} = \left(\frac{\phi}{\sigma}\right)^2. \tag{5}$$

Normal conjugacy gives

$$v_n = \left(v_0^{-1} + nR_\lambda^{-1}\right)^{-1}, \qquad v_{n+1} = \left(v_n^{-1} + R_\lambda^{-1}\right)^{-1}. \tag{6}$$

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

$$\widehat{\psi}_{n+1} = \widehat{\psi}_n + \sqrt{v_n - v_{n+1}}\, Z_{n+1}. \tag{7}$$

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 standardized innovation changes current new-economy productivity by

$$\Delta\rho_{n+1}^N = \frac{\sigma}{\sqrt{\lambda}} Z_{n+1}. \tag{8}$$

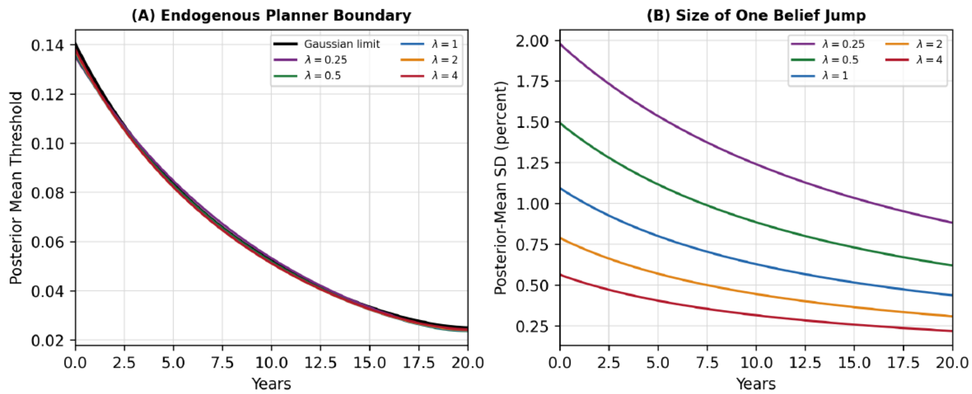

Figure B1: **Poisson learning and adoption boundaries.** Panel A plots the Gaussian planner boundary and Poisson-average planner boundaries for alternative arrival intensities λ. Panel B plots the posterior-mean standard deviation associated with one signal arrival. The Poisson-average boundary at each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

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 (2)–(3). 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 certainty-equivalent value of immediate adoption, after removing terms common to adoption and continuation, is

$$S_{t,n}(\widehat{\psi}) = \log(1-\kappa) + A_2(T-t)\widehat{\psi} + \frac{\theta}{2}A_2(T-t)^2 v_n. \tag{9}$$

Over an interval of length Δt, let K be the number of signal arrivals. The continuation value is

$$C_{t,n}(\widehat{\psi}) = \frac{1}{\theta}\log\left[\sum_{k=0}^{\infty}\Pr(K=k)E\left[\exp\Big\{\theta V_{t+\Delta t,n+k}\left(\widehat{\psi}+\sqrt{v_n - v_{n+k}}Z\right)\Big\}\right]\right], \tag{10}$$

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

$$S_{t,n}(b_n(t)) = C_{t,n}(b_n(t)). \tag{11}$$

Asset values are computed from the corresponding pricing moments. Let

$$D_t = E_t[W_T^{-\gamma}], \qquad N_t^N = E_t[W_T^{-\gamma}B_T^N]. \tag{12}$$

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

$$\log P_t^N = \log N_t^N - \log D_t + A_1(T-t)\rho_t^N + \log B_t^N. \tag{13}$$

The pricing recursion integrates jointly over equations (7) and (8). In particular, the contemporaneous response to a signal contains the direct productivity term

$$A_1(T-t)\Delta\rho_t^N = A_1(T-t)\frac{\sigma}{\sqrt{\lambda}}Z, \tag{14}$$

in addition to the effect of the posterior-mean update on expected future growth. Omitting equation (14) 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 positive productivity shocks, and every one of those 155 crossings produces a positive new-market 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 adop-

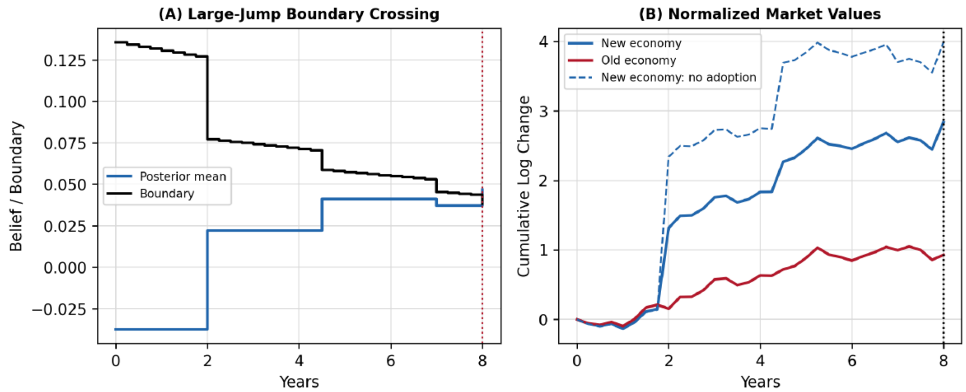

Figure B2: **A large Poisson signal and market values.** Panel A plots the posterior mean and the count-specific planner boundary along a representative path that adopts at year 8. Panel B plots cumulative log changes in new- and old-economy market values and the corresponding no-adoption new-economy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

tion 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 (7) and (8),

$$\Delta\widehat{\psi} = 0, \qquad \Delta\rho^N = 0, \qquad \Delta v = v_{n+1} - v_n < 0. \tag{15}$$

If

$$b_{n+1}(t) < \widehat{\psi} < b_n(t), \tag{16}$$

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.

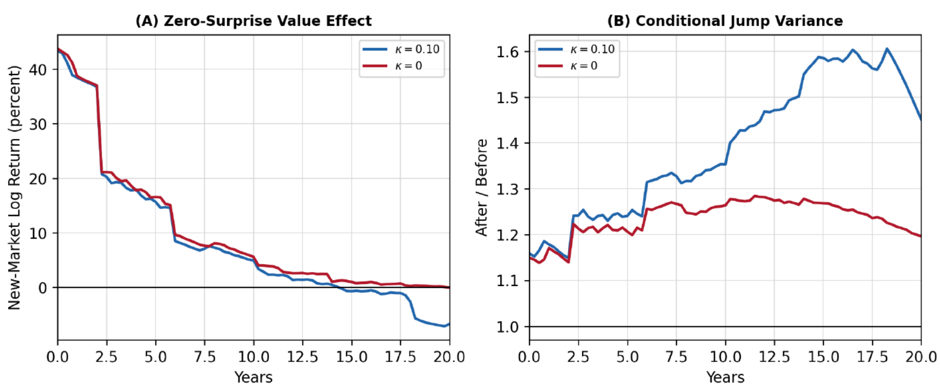

Figure B3: **Zero-surprise learning events and the fixed adoption cost.** Panel A plots the new-market log return produced by a signal that leaves the posterior mean and current productivity unchanged but reduces posterior variance enough to induce adoption. Panel B plots the ratio of conditional Poisson-jump 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 the threshold to M^{IA} just above, and the step is the full adoption discount $\mathcal{D}(8) \equiv \gamma A_2(T-8)^2\hat{\sigma}_8^2 = 1.77$. Both values are in closed form, and $\hat{\sigma}_s$ is deterministic, so

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.

$\mathcal{D}(T_2)$ is already known at year eight. A claim certain to adopt at T_2 is worth $M^{NA}e^{-\mathcal{D}(T_2)}$, and a path that narrowly misses at year eight adopts at T_2 with probability close to but below one. So a near miss is then a delay rather than a never, and the step shrinks from $\mathcal{D}(8)$ toward the increment $\mathcal{D}(8) - \mathcal{D}(T_2)$, which is 0.23, 0.78 and 1.29 for T_2 at nine, twelve and sixteen. This force can be isolated. 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 year-twelve review raises the year-eight threshold to 0.063, the paths that clear it are those whose beliefs rose most steeply, and their run-up is steeper. Together the two forces 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 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

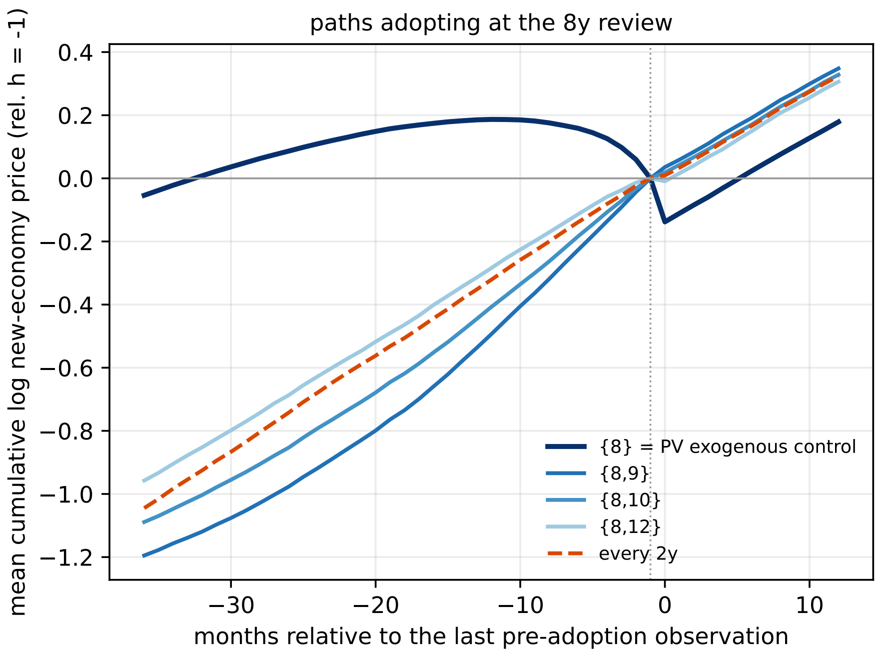

Figure C1: **New-economy prices into a year-eight adoption under review calendars.**

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.

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 Heterogeneous conversion costs and endogenous timing

Section 3 of the paper shows that when firms with different conversion costs adopt at a common date, their costs determine the order of their adoptions, their choices match the planner's, and the repricing of new-technology risk is spread across the queue, which diffuses a price bubble. This appendix lets these heterogeneous firms choose *when* they each will optimally adopt. Allowing endogenous adoption timing brings several issues. What does the market learn about ψ while adoption is under way? Is a sorted queue an equilibrium in stopping time? How does spreading an adoption wave over years affect the price pattern of new-technology stocks?

We keep every primitive of the model in Section 2 and add only the cost schedule already used in Section 3 of the paper. For clarity in this appendix, we label the new-economy firms as "venture" firms, and label the old-economy firms that adopt the new technology as "converted" firms.

D.1 Setup

Firms are indexed by $x \in [0,1]$ with conversion cost $\kappa_\varepsilon(x) = \kappa + \varepsilon(2x-1)$, $\varepsilon > 0$. Lower-x firms convert first, so the adopted fraction q_t identifies the marginal type $\kappa_\varepsilon(q_t)$. Conversion is irreversible. A firm converting at date s pays $\kappa_\varepsilon(x)$ out of capital, and its productivity thereafter mean-reverts to $\bar{\rho} + \psi$ while keeping the common exposure $\sigma\, dZ_{0,t}$, with the level continuous at conversion.[1] Because a converted cohort's productivity gap to the old vintage evolves deterministically given ψ, terminal wealth aggregates exactly across cohorts,

$$W_T \;=\; B_T^0\Big[(1-q_T) + \int_0^{q_T} \big(1-\kappa_\varepsilon(x)\big)\, e^{A_2(T-\tau(x))\,\psi}\, dx\Big], \tag{D1}$$

where B_T^0 is the all-old-technology capital and $\tau(x)$ is firm x's conversion date. If every conversion occurs at one date, (D1) collapses to the terminal-wealth expression used in Section 3, which is therefore exact in that case. When adoption is spread over time, early converters load more of ψ by T than late ones, and (D1) keeps that mixture. Prices are computed as in Section 2, with the pricing kernel $\pi_t \propto E_t[W_T^{-\gamma}]$ built from (D1). The new-economy claim is, as before, the infinitesimal claim on the venture's productivity ρ^N.

D.2 Learning as the adoption wave runs through the economy

In Section 2, the public learns about ψ from one source, the venture's productivity ρ_t^N, and the posterior variance falls along a deterministic path. Once firms convert, total output in

[1] Converted firms inherit the old economy's shock rather than the venture's loadings. This is the large-scale adoption structure of Pástor and Veronesi (2009). What adoption transfers is the drift anchor ψ. The venture's idiosyncratic shock dZ_1 is implementation noise at pilot scale, and it washes out when a continuum of incumbents adopts. Keeping the common shock isolates the mechanism of interest, since the entire repricing at conversion then runs through ψ-uncertainty becoming systematic. It also makes each cohort's productivity gap to the old vintage deterministic given ψ, which is what delivers the exact aggregation that follows.

the economy becomes a second source, because converted firms produce with mean $\bar{\rho} + \psi$. A converted incumbent runs new-technology production commingled with old-technology operations. Its conversion is a discrete, announced act (an investment, a license), but the return on it accrues inside aggregate output. Agents can observe aggregate output but cannot split output by vintage.

Suppose instead that the market could observe old-vintage and converted-cohort output *separately*. Differencing the two would cancel the common shock dZ_0 and reveal ψ exactly the moment the first firm converts, because neither sector carries idiosyncratic noise. That degeneracy is a property of the shock structure. We now derive the learning filter.

Assumption D1 (Public information during the wave)**.** *The public observes (i) the venture's productivity ρ_t^N; (ii) every adoption event, hence the adopted fraction q_t and the marginal cost type; and (iii) the* total *output of the converting sector, but not its decomposition into old-vintage and converted output.*

The assumption gives up nothing about the economy's composition. Agents know the cost schedule, the adopted fraction and every cohort's age, and these enter the filter below as known coefficients. What they cannot see are ψ and dZ_0. The alternative of letting agents learn nothing from converted output is not available. Aggregate consumption is public because it prices the kernel, and by the lemma below its drift carries ψ whenever $0 < q_t < 1$.

Lemma D1 (Total output is a one-dimensional signal)**.** *Let $\rho_{O,t}$ be old-vintage productivity and let g_v be a converted vintage's productivity gap to it. The common shock cancels inside each gap, so $dg_v = \phi(\psi - g_v)\,dt$ with $g_v = 0$ at conversion. Total output Y_t of the converting sector then obeys*

$$dY_t = \phi(\bar{\rho} - Y_t)\,dt + q_t\,\phi\,\psi\,dt + \sigma\,dZ_{0,t}.$$

Total output is a single observable whose drift loads on ψ with the known coefficient $q_t\phi$ and whose noise is the common shock. Filtering (Y_t, ρ_t^N) jointly, with noises correlated through dZ_0, leaves the posterior mean $\hat{\psi}_t$ a martingale as in Section 2 and turns the posterior variance into a state variable that obeys a scalar Riccati equation,

$$\frac{d\hat{\sigma}_t^2}{dt} = -\,\hat{\sigma}_t^4\,r(q_t), \qquad r(q) = \phi^2\left[\frac{q^2}{\sigma^2} + \frac{\left(1 - q\,\sigma_{N,0}/\sigma\right)^2}{\sigma_{N,1}^2}\right]. \tag{D2}$$

Proposition D1 (The learning rate along the wave)**.** *(i) $r(0) = (\phi/\sigma_{N,1})^2$ is the pre-adoption precision rate of Section 2. (ii) $r(1) = \phi^2[\sigma^{-2} + (1 - \sigma_{N,0}/\sigma)^2\sigma_{N,1}^{-2}]$ is the post-adoption rate of Section 2, cross-term included. (iii) r is U-shaped and is minimized at $q^* = \sigma\sigma_{N,0}/(\sigma_{N,0}^2 + \sigma_{N,1}^2)$. At the baseline $\sigma = \sigma_{N,0} = \sigma_{N,1}$, $q^* = \frac{1}{2}$ and $r(q) \propto q^2 + (1-q)^2$ runs at half the endpoint speed there.*

The two endpoints say that (D2) nests both filters of Section 2 exactly. Before any conversion the market learns from the venture alone. After complete conversion it learns from the adopted economy at precisely the rate assumed there. The shape in between is the economic content (Figure D1, panel (a)). With few adopters, conversion *contaminates* what used to be a clean read on the old economy. Once the aggregate contains enough adopters, total output becomes a good signal of ψ in its own right. Early adoption therefore slows public learning, and the information externality changes sign at q^*, negative on the first half of the wave and positive on the second. This follows from the shock structure rather than from a modeling choice. Because prices are functions of public beliefs and firms adopt on public information, neither prices nor adoption events carry information beyond Assumption D1. The state of the economy is $(t, \hat{\psi}_t, \hat{\sigma}_t^2, q_t)$.

D.3 The sorted schedule

Adoption policies are type-indexed threshold rules. Firm x converts when $\hat{\psi}_t$ first crosses a boundary $b(t; \kappa_\varepsilon(x))$, and the cost ordering makes the queue monotone. Fixing a conversion date, the marginal type is indifferent between converting and staying out when the market-priced gain from conversion is zero, which is the condition of Proposition 2. Solving it date by date gives the static schedule shown in Figure D1, panel (b). The front of the queue is the unilateral threshold of the lowest-cost type, the back is the join threshold of the highest-cost type, and the midpoint is the planner's threshold. The range of beliefs the queue sweeps, in log-value units $\gamma A_2(\tau)^2 \hat{\sigma}^2$, is exactly the gap between the two support costs of Proposition 1, and it does not depend on ε. Two consequences follow. First, the sorted schedule is the planner's allocation. Because $u'(W_T) = W_T^{-\gamma}$, a planner choosing the adopted fraction has the same first-order condition as the marginal firm, so heterogeneity in conversion costs resolves the coordination problem of Section 3 with no coercion at all in the static problem. Second, since the belief band is fixed, no choice of heterogeneity re-concentrates the repricing; it only spreads it out.

D.4 Coordination costs under endogenous timing

We do not solve the dynamic game in which every firm chooses its conversion date. Once the posterior variance depends on the path of adoption, the pricing state carries the history of the queue, and the frictionless mean-field equilibrium becomes challenging to compute. What we can do is ask whether a given schedule survives when firms choose their own timing, and if it does not, what it costs to hold it in place. We start from the static schedule of the previous subsection, which is the planner's allocation, and consider a two-parameter family of schedules around it in which each type's threshold carries its own loading on $A_2(\tau)\hat{\sigma}^2$. At every state visited by the wave, we price the value of a firm that follows the schedule and of one that shifts its own threshold, using the wave's own kernel and beliefs, and we look for a schedule in the family at which no shift pays. In a continuum, an individual deviator does not move the wave, so one nested simulation prices every shift on the same paths.[2]

The answer is negative at both calibrations and every ε. From any schedule in the family, a measure-zero firm gains by waiting, at every date and every point of the queue. The motive is the same one that makes lag costs necessary for the synchronized rules of Section 3. As firms ahead convert, the ψ-risk a converting firm takes on is being made systematic, and the kernel charges progressively more for it, so the marginal firm would rather stay with the old

[2] The deviation values are computed by quadrature over the posterior of ψ inside a Monte Carlo over the shocks. The pricing kernel is exponential in ψ with a log standard deviation of $\gamma A_2 \hat{\sigma} \approx 3.8$ in mid-wave, so a Monte Carlo average over draws of ψ is dominated by rare extreme draws and does not settle at any feasible sample size. Integrating over ψ by quadrature removes that source of variance.

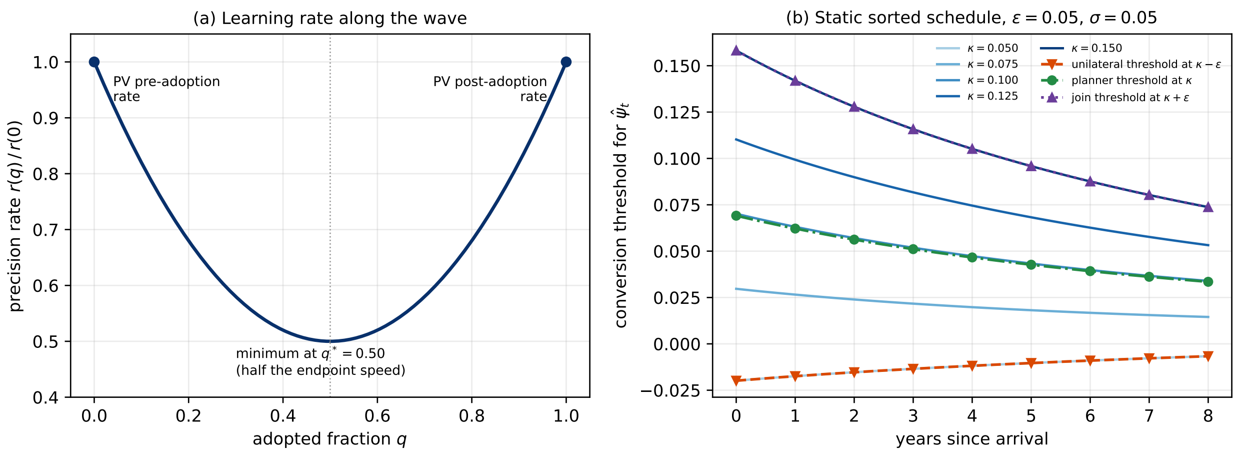

Figure D1: **Learning and the sorted schedule under heterogeneous conversion costs.**

Panel (a) plots the precision rate $r(q)$ of (D2) relative to its pre-adoption value at the baseline calibration ($\sigma = \sigma_{N,0} = \sigma_{N,1}$). The endpoints are the pre- and post-adoption learning rates of Section 2. The rate is minimized at $q^* = 1/2$, where the market learns at half the endpoint speed. Panel (b) plots the static schedule $b(t; \kappa)$ for $\varepsilon = 0.05$ and $\sigma = 0.05$, that is, the belief at which a firm of cost κ converts when the adopted fraction is the mass of cheaper firms, over the first eight years after arrival. The three marked lines are the unilateral threshold of the lowest-cost type, the planner's threshold at $\kappa = 0.10$, and the join threshold of the highest-cost type. The schedule's ends coincide with the first and last, and its midpoint sits at the planner's threshold up to a term of order ε.

technology a while longer. In a continuum this delay is free of strategic cost, since a single firm's delay changes neither the adopted mass nor the kernel, and no rival takes its slot. What differs from the synchronized case is the size of the temptation. There, sitting the wave out avoids the whole repricing at once. Here, delay avoids only the slice of repricing that the next stretch of the queue delivers, with the firm re-entering the queue later if desired.

Since no schedule in the family is self-enforcing, any queue we price has to be supported, and we support the static schedule. It is the natural choice because it is the planner's allocation, so whatever it costs to hold in place buys the efficient outcome. Other schedules could be supported by larger lag costs. The instrument is the one that supports the synchronized rules of Section 3, a lag cost imposed on a firm that misses its slot, and we compute the minimum lag cost that makes every visited state incentive compatible. It is one-sided, since no lead cost is required anywhere and no one wants to go early. It is substantial. Expressed as a share of firm value it is 28 to 29 percent at $\sigma = 0.05$ and 22 to 23 percent at $\sigma = 0.07$, varying by less than two percentage points across $\varepsilon \in \{0.01, 0.025, 0.05\}$. It is nonetheless roughly one third of what a synchronized rule needs, which is $\gamma A_2^2 \hat{\sigma}^2 \approx 4.7$ log units at arrival, essentially the whole firm, for the unilateral and join rules of Proposition 1. It binds at the front of the queue, declines by about a quarter along it, and is lower at the higher volatility. Because the supported allocation is the planner's, this is coercion without misallocation, and the enforcement is the only cost of decentralization in the static problem.

Learning feeds back on timing. Compared with a wave in which the market learned from the venture alone, the lowest-cost firms at the front of the queue convert when they otherwise would, but the higher a firm's cost, the later it converts, by 1.0 to 1.3 years at the highest-cost types the wave reaches, at both calibrations. Early conversion slows public learning, so the beliefs that would carry the higher-cost types over their thresholds arrive later. The static efficiency identity also breaks under learning. Firms at the front of the queue do not internalize that their conversion slows public learning, and firms at the back do not internalize that theirs speeds it up, so there is a wedge between the marginal firm's condition and the planner's, negative while $q < \frac{1}{2}$ and positive after, in line with Proposition D1. Its net effect over a wave is ambiguous in sign, and we report its structure rather than a welfare number, since the states we observe cover only the first half of the wave. Finally, the posterior variance, which falls monotonically in Section 2, falls most slowly around $q \approx \frac{1}{2}$, at half its endpoint speed.

D.5 Price dynamics

We now examine the price patterns of new-economy ("venture") stocks under this adoption schedule. With synchronized adoptions, the event is the single adoption date. With dispersed adoptions, we examine returns anchored to aggregate adoption milestones, the first date T_a at which the adopted fraction crosses $a^* = 0.25$ or 0.50. Mean price run-ups over months $[T_a - 24, T_a - 6]$ across successful adoption paths are reported. We examine returns over pre-adoption windows $[T_a - h, T_a - 1]$ with h from 2 to 12 months, and for each cohort we select the window with the lowest mean return. The window is chosen on the cohort mean, so all paths in a cohort contribute to the same window, and the choice picks up the deepest average pre-adoption fall whatever its length. The price series we use is the new-economy price relative to old-economy book. With old-economy book growing at about 12 percent per year, this normalization removes the common growth trend, so a price that carries no adoption news is roughly flat, which is what the placebo shows. As a power check, the same pipeline is run on the synchronized planner rule conditioned on adopting at the anchor date, and on a placebo that anchors on calendar dates unrelated to adoption.

Adoption begins at arrival on every path, because the cheapest types sit below the prior mean, but complete conversion within the eight-year window occurs on essentially no path. The quarter milestone is reached on 37 (23) percent of paths at $\sigma = 0.05$ (0.07), at a median of 4.4 (5.4) years. The half milestone is reached on 4.9 (0.6) percent, at 6.8 (7.2) years, and three quarters on well under one percent. The results below describe the first quarter and first half of the wave. That is where the repricing is largest. The belief band the queue must sweep, $\gamma A_2^2 \hat{\sigma}^2$, is the total repricing a complete wave can deliver, and since both A_2 and $\hat{\sigma}^2$ fall as time passes, the part still to come shrinks with every year of transit. The last adopters convert at beliefs so high that the transition is a favorable aggregate shock (the join end of the band), which is the opposite of the state that produces a crash.

Table D1: **Bubble criteria at aggregate adoption milestones.**

Cell	n	(i) run-up, $[T_a{-}24, T_a{-}6]$	(ii) fall, $\min_h [T_a{-}h, T_a{-}1]$	Verdict
$\sigma = 0.05$, T_{50}	2200	+0.147 (z=61)	none (+0.189 at h=12)	boom, no crash
$\sigma = 0.05$, T_{25}	1200	+0.136 (z=42)	none (+0.214 at h=12)	boom, no crash
$\sigma = 0.07$, T_{25}	2200	+0.168 (z=56)	none (+0.275 at h=12)	boom, no crash
Sync. planner, 0.05	259	−0.207	−1.945 at h=2 (z=−6.6)	crash
Sync. planner, 0.07	630	−0.086	−1.020 at h=2 (z=−8.5)	crash
Placebo (either σ)		−0.01 to −0.04	−0.045 at h=2, n.s.	nothing

Notes. New-economy price relative to old-economy book, $\varepsilon = 0.05$, sorted schedule supported by the lag cost of the previous subsection. T_a is the first date the adopted fraction reaches a. Cohorts are paths on which the milestone occurs with a full 24-month history. (i) is the mean annualized log change over $[T_a - 24, T_a - 6]$ months. (ii) is the minimum over $h \in \{2, \ldots, 12\}$ of the cohort-mean annualized log change over $[T_a - h, T_a - 1]$, with the selected h shown, reported as "none" when the minimum is positive. The synchronized control conditions the planner rule of Section 2 on adopting at the anchor date and runs the identical pipeline. The placebo anchors on calendar dates. Old-economy book grows at $\bar{\rho} \simeq 12$ percent per year, so absolute price changes are the reported rates plus about 0.12 per year.

Table D1 gives the result. At both calibrations and both milestones the new-economy price rises into the anchor at 0.14 to 0.17 per year over the preceding two years, cumulatively +0.35 to +0.50 log units, and keeps rising through and after the crossing. There is no window, at either anchor or either calibration, in which the price falls. The window that attains the minimum is the longest one, twelve months, in every dispersed cell and the shortest one, two months, in the control. This is the difference between a rise that accelerates into the anchor and a fall concentrated in the last two months before adoption. The synchronized control run through the same pipeline crashes, −1.0 to −1.9 per year into adoption, so the instrument has power, and the placebo shows nothing. Figure D2 shows the anatomy, a steady rise into the milestone with a narrow interquartile band, a flat old-economy market claim, and a flat placebo. The mechanism is the one Section 2 identifies for the crash, run in reverse. A crash needs a state in which adoption raises marginal utility, that is, adoption as a systematic-risk shock at a low belief, which is what the synchronized planner rule delivers. The queue never visits that state. By construction its first adopters are those for whom the risk is idiosyncratic, and its last adopt when the transition is good news for aggregate wealth. Along a successful wave the kernel level falls, and a claim on the new technology is repriced upward as the wave proceeds.

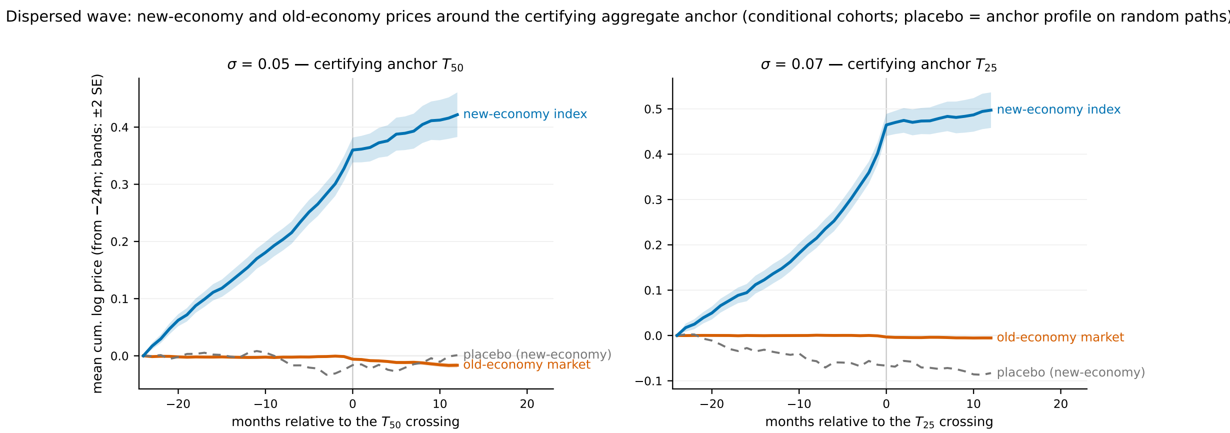

Figure D2: **New-economy prices around aggregate adoption milestones.**

Mean cumulative log price of the new-economy claim relative to old-economy book (solid), of the old-economy market claim (dashed) and of the placebo (dotted), in months relative to the milestone, with bands of ± 2 standard errors of the cohort mean. Left panel, $\sigma = 0.05$ and anchor T_{50}. The new-economy price is $+0.36$ at the crossing and $+0.42$ twelve months after, the old-economy market -0.01, the placebo -0.02. Right panel, $\sigma = 0.07$ and anchor T_{25}. The new-economy price is $+0.46$ at the crossing and $+0.50$ twelve months after, the old-economy market 0.00, the placebo -0.07. Conditioning on a milestone selects paths on which beliefs have climbed, so a rise into the anchor is expected under any persistent-signal economy. The content of the figure is the absence of any decline before or after the anchor, and the flat placebo. $\varepsilon = 0.05$, sorted schedule, derived filter.

What dispersion leaves of the crash is a firm-level phenomenon. Aligning each adopter on its own conversion date rather than on an aggregate milestone, the adopter's *own* claim falls into that date in every (ε, σ) cell. The fall is -0.18 per year at $\sigma = 0.05$ and -0.16 at 0.07 over the final two months (standard errors below 0.01), after a mild two-year erosion of about -0.02 per year, and there is no run-up in the firm's own price. This is the result of firms surrendering their waiting option as crossing arrives. It is the same object the lag cost of the previous subsection measures from the deviation side, and it occurs while the index the firm belongs to is rising. It is not a bubble in the sense of Section 2, but it is the piece of that mechanism that remains. The statement the results support is therefore that cost dispersion preserves the new-economy boom and removes the index crash.

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). We solve the corresponding free-boundary PDE on a grid that moves with the static threshold; no additional state variable is required.

The unilateral aggregate-preemption boundary is a fixed point of a best-response recursion. At each grid point, a firm compares adoption today with waiting under its off-equilibrium continuation payoff. The public state is $\hat{\psi}_t$, but the adoption payoff also depends on the unobserved true productivity ψ_{true}. We therefore integrate over $\psi_{true} \mid \hat{\psi}_t \sim N(\hat{\psi}_t, \hat{\sigma}_t^2)$ at every step. Exponential-affine terms are averaged in closed form; terms containing the exercise-boundary maximum are evaluated by Gauss–Hermite quadrature. The continuation value is weighted by the pricing kernel induced by the future aggregate adoption rule, so the boundary jointly reflects the adoption rule and the systematic-risk channel it creates. The equilibrium transitions to aggregate adoption when this fixed-point boundary is reached.

References

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.