

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

¹The 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 Eislefeldt 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

²For 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

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

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:

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

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

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

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 $\hat{\sigma}_t^2 = 0$, 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,”

³This ratio is an endogenous function of the time- t state variables: $(\rho_t^N, \rho_t, \hat{\psi}_t, \text{ and } t)$. 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.

⁴These 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.¹²

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.

⁵A 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:

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

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

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

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:

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

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.

⁶We 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 ($\frac{\partial \log(M_t^N)}{dZ_{1,t}}$). 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

⁷The 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

⁸See Online Appendix.

⁹For 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.

¹⁰See 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 $E_t[t^{**}]$. 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 $\log(M_t^N) \approx \log(M_t^{IA}) + \left(\frac{E_t[t^{**}] - t}{T - t}\right) A_2(\tau)^2 \gamma \hat{\sigma}_t^2$. Panel D shows the conditional exposure of the new-economy stock price to new-economy productivity shocks $\left(\frac{\partial \log(M_t^N)}{\partial \bar{Z}_{1,t}}\right)$.

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,

¹¹For 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

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

as the penalty needed to deter a firm from adopting unilaterally while the aggregate economy

is still old ($r = 0$), and the lag cost

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

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,

$$V^{a|r} = E_{\bar{t}^{**}} [\pi_{\bar{t}^{**}, T}^r B_{i, T}^a],$$

where π^r is the regime- r pricing kernel and $B_{i, T}^a$ 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 $\hat{\sigma}^2 \equiv \hat{\sigma}_{\bar{t}^{**}}^2$. The planner, join (J), and unilateral (U) thresholds for $\hat{\psi}_{\bar{t}^{**}}$ are*

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

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

$$(c_J^{lead}, c_J^{lag}) = (\gamma A_2(\tau)^2 \hat{\sigma}^2, 0),$$

$$(c_U^{lead}, c_U^{lag}) = (0, \gamma A_2(\tau)^2 \hat{\sigma}^2),$$

and

$$(c_{planner}^{lead}, c_{planner}^{lag}) = (\frac{1}{2}\gamma A_2(\tau)^2 \hat{\sigma}^2, \frac{1}{2}\gamma A_2(\tau)^2 \hat{\sigma}^2).$$

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

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

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

¹²We 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.

¹³We 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 .

¹⁴Recall 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

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

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

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

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:

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

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 family¹⁵

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

¹⁵This 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

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

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

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

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

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

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

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

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

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

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

To see why, define

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

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 ,

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

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 distinc-

tion 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 $B_{i,T}^0$ 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 $B_{i,T}^1 = (1-\kappa)B_{i,T}^0 \exp(A_2(\tau)\psi)$, 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 $W_T = (1-\kappa)^r B_T^0 \exp(rA_2(\tau)\psi)$, so $\pi_{\bar{t}^{**},T}^r \propto (B_T^0)^{-\gamma} \exp(-\gamma r A_2(\tau)\psi)$, 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 $V^{a|r}/V^{0|r} = E_{\bar{t}^{**}}[\pi_{\bar{t}^{**},T}^r B_{i,T}^a]/E_{\bar{t}^{**}}[\pi_{\bar{t}^{**},T}^r B_{i,T}^0]$, 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 $\psi \mid \mathcal{F}_{\bar{t}^{**}} \sim N(\hat{\psi}_{\bar{t}^{**}}, \hat{\sigma}_{\bar{t}^{**}}^2)$ then gives, for $a, r \in \{0, 1\}$,

$$\frac{V^{a|r}}{V^{0|r}} = \frac{E_{\bar{t}^{**}}[\pi_{\bar{t}^{**},T}^r B_{i,T}^a]}{E_{\bar{t}^{**}}[\pi_{\bar{t}^{**},T}^r B_{i,T}^0]} = (1-\kappa)^a \exp\left(aA_2(\tau)\hat{\psi}_{\bar{t}^{**}} + \frac{a}{2}(a-2\gamma r)A_2(\tau)^2\hat{\sigma}_{\bar{t}^{**}}^2\right). \quad (*)$$

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

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

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

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,

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

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

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. $E_{\bar{t}^{**}}[\pi_{\bar{t}^{**}, T}^1 B_{i, T}^1] = E_{\bar{t}^{**}}[\pi_{\bar{t}^{**}, T}^1 B_{i, T}^0]$, or $D^1 = 0$. By $(**)$ this is

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

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. $E_{\bar{t}^{**}}[\pi_{\bar{t}^{**}, T}^0 B_{i, T}^1] = E_{\bar{t}^{**}}[\pi_{\bar{t}^{**}, T}^0 B_{i, T}^0]$, or $D^0 = 0$. By $(**)$ this is

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

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

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

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

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

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

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

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 (**),

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

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 (**),

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

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

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

$$c_{planner}^{lead} = c_{planner}^{lag} = \frac{1}{2}\gamma A_2(\tau)^2\hat{\sigma}^2,$$

each equal to half of $c_J^{lead} = c_U^{lag}$. This completes the proof.

B.3 Proof of Proposition 2

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

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

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

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

Define the marginal gain

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

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

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

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

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

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,

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

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

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

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.