Technological Revolutions and Stock Prices Revisited*

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

March 27, 2026

Abstract

Pástor and Veronesi (2009) proposes that stock price bubbles can arise during technological revolutions due to sharply rising systematic risk as economy-wide technology adoption approaches. This mechanism is unlikely to explain bubbles. We show that, in the baseline calibration, a bubble arises only in an illustrative model that assumes the adoption decision occurs at an exogenously fixed date. More generally, bubbles are not a natural feature when firms may adopt optimally at any time. We also find no evidence of the unique empirical signature of the mechanism within the Dotcom bubble. We discuss implications for the recent AI stock boom.

*We thank Ľuboš Pástor and Pietro Veronesi for making their replication code available. 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) develops a rational explanation for stock price bubbles during technological revolutions. In the PV framework, positive shocks to the productivity of a new technology raise expected cash flows for new-technology firms, pushing their stock prices up, but also increase the likelihood of widespread adoption, raising systematic risk and pushing prices down through higher discount rates. Early in a revolution, the cash-flow effect dominates; later, as economy-wide adoption becomes more likely, the discount-rate effect can dominate. Conditional on a successful revolution, the stock prices of new-technology firms can therefore display a bubble pattern.

The mechanism, whereby exposure to new-technology productivity shocks is largely idiosyncratic early in a revolution but becomes increasingly systematic as widespread adoption approaches, is intuitive and has been central to recent discussions of potential bubbles in stocks related to generative artificial intelligence (AI).[1] The key implication is that further positive news about the new technology can become bad news for prices if it sufficiently increases systematic risk.

In this paper, we argue that this outcome is unlikely to occur. PV presents two models of technological adoptions: an illustrative case with an all-or-nothing decision to adopt the new technology available only once at an exogenously preset date, and a more realistic version with adoption endogenously and optimally timed. PV states that the dynamics in the more realistic model are "very similar" to those in the illustrative model, but we show that this is not the case. In the baseline calibration, a price bubble arises only in the illustrative model. When adoption timing is endogenous, stock prices instead rise steadily throughout the entire revolution.[2]

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

[2] A source of confusion is that, while PV emphasize stock price bubbles in the exposition, they report only market-to-book ratios for the endogenous model. The bubble pattern in market-to-book ratios reflects mean reversion and does not imply a bubble in prices. See Section 3.

To understand the disparity, we derive a simple approximation for new-technology stock prices that applies to both versions of the model. In this approximation, the impact of systematic risk on the log of new-technology stock price is linear in the expected time remaining until adoption. A stock price bubble therefore requires a region late in the revolution where small positive productivity shocks generate large and immediate reductions in expected adoption time, so that the resulting discount-rate increase overwhelms the positive cash-flow effect on price.

The assumption of an unnatural, all-or-nothing decision at an exogenous time creates exactly such a region. Approaching the preset time, if beliefs about the technology's productivity are just below the level required for adoption, small positive productivity shocks can sharply increase adoption probabilities. A small shock near the knife-edge threshold can make adoption jump from "never" to "now", causing a spike in systematic risk that overwhelms the positive cash-flow effect.

In contrast, in the more realistic endogenous model, adoption occurs when beliefs about new-technology productivity rise above a continuous boundary. Because of the agent's gradual learning about the productivity of the new technology — the central feature of the PV model — the adoption boundary in the endogenous model generally slopes downward with time. This downward slope means that productivity shocks translate into smooth, incremental changes in expected adoption time, even early in the revolution when beliefs are far from the boundary. Systematic risk of the new technology therefore rises smoothly throughout the revolution, preventing it from ever resulting in price declines. In short, optimal adoption timing removes the knife-edge decision that generates a price bubble.[3]

Having established the theoretical argument, we complement it with an empirical evaluation of the mechanism during the Dotcom bubble of the early 2000s, a main focus of PV. Rather than conducting a formal statistical test, which would necessitate measuring latent

[3]We find this lack of a bubble pattern holds for a broad range of reasonable parameterizations. However, under extreme (we argue implausible) parameterizations, the endogenous model can generate a bubble for some productivity paths. See Online Appendix for a discussion.

economy-wide adoption probabilities, our analysis focuses on whether the dynamics of stock prices and cash flows align qualitatively with the PV model. The empirical signature of the discount-rate mechanism is that just prior to economy-wide adoption, stock prices should be falling while expected cash flows are still rising. We construct forecasts and realizations of profitability for NASDAQ companies during the bubble period and find this signature is absent: the price collapse corresponds with strongly falling profitability, consistent with other papers' findings (e.g., Pástor and Veronesi (2006)).

Lastly, we comment on the recent AI boom. Since economy-wide AI adoption appears likely (e.g., Bick et al. (2026) and Yotzov et al. (2026)), but still uncertain, exposure to AI-productivity shocks should already be transitioning from an idiosyncratic to a systematic risk. If technological revolutions generate bubbles through the PV mechanism, this is the stage at which AI stock prices should already be falling as AI firms' profitability continues to rise. Instead, we find that profitability and prices are steadily rising together, consistent with the endogenous model's smooth change in systematic risk throughout a revolution.

2 Related Literature

While the paper focuses on the model of Pástor and Veronesi (2009), the analysis also relates to Pástor and Veronesi (2006), who emphasize that uncertainty about future growth rates can produce high valuations even in a fully rational setting because valuation ratios are convex in growth expectations. Note that this cash-flow mechanism is an alternative driver of prices rising more and then falling more than asset-pricing models without learning would predict.

More broadly, our work connects to the literature on bubbles driven by disagreement, speculation, or investor sentiment (e.g. Harrison and Kreps (1978), Scheinkman and Xiong (2003), and Barberis et al. (1998)) and to more recent papers that study technological innovation and asset prices (e.g Garleanu et al. (2012), Kung and Schmid (2015), and Kogan et al. (2020)). Finally, our empirical discussion relates to evidence on the Dotcom bubble (e.g. Ofek and Richardson (2003) and Greenwood and Nagel (2009)).

3 Model predictions and mechanisms

We briefly review the model of Pástor and Veronesi (2009) and then focus on the underlying mechanism that generates bubble dynamics in stock prices.[4] The economy has a finite horizon $[0, T]$ with a representative agent 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 $Y = \rho_t B_t$. This output is used to grow the capital stock so that $dB_t = Y_t d_t = \rho_t B_t d_t$. 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 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

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

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

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 $N(0, \sigma_J^2)$ with known variance. After t^*, the new-economy capital stock (B_t^N) and productivity (ρ_t^N) are observable and evolve according to

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

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 agent 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 agent decides at a pre-specified time $\bar{t}^{**}$ whether to adopt the new technology on a large scale. Adoption occurs if and only if the posterior mean exceeds a threshold:

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

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

In the *endogenous* adoption time scenario, adoption occurs optimally at the time when adoption maximizes the agent's expected utility. As PV shows, this problem is akin to the optimal exercise of an American option, whereby the agent 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 agent adopts if $\hat{\psi}_t \geq \bar{\psi}(t)$. This threshold can be written as the sum of two terms of the static adoption threshold for a given time t, and a continuation value, so that

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

Here $\chi(t) \geq 0$ is the continuation option value of waiting to adopt. A closed form for this term is not available and the endogenous adoption boundary is therefore obtained through the numerical solution of the PDEs laid out in PV. We note that $\chi(t) = 0$ when $t = T$ or $\hat{\sigma}_t^2 = 0$, 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 to 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.

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 representative agent'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. 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 characteristic bubble pattern. PV reports the plot for the endogenous case (its Figure 6), and after observing the bubble pattern in $\frac{M_t^N}{B_t^N}$, states that the "conclusions (...) are unaffected by endogenizing t^{**}". However, the economically relevant object, and the stated focus of PV, is a price bubble.

To examine whether the stock price inherits the bubble pattern of M/B, Panel E plots cumulative returns to the new economy under an exogenous adoption time. Stock prices in this case do inherit the bubble pattern of M/B, 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 inherit the 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 due to 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 $\bar{\rho}$, 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 the bubble pattern, while in the endogenous case the downward slope 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.

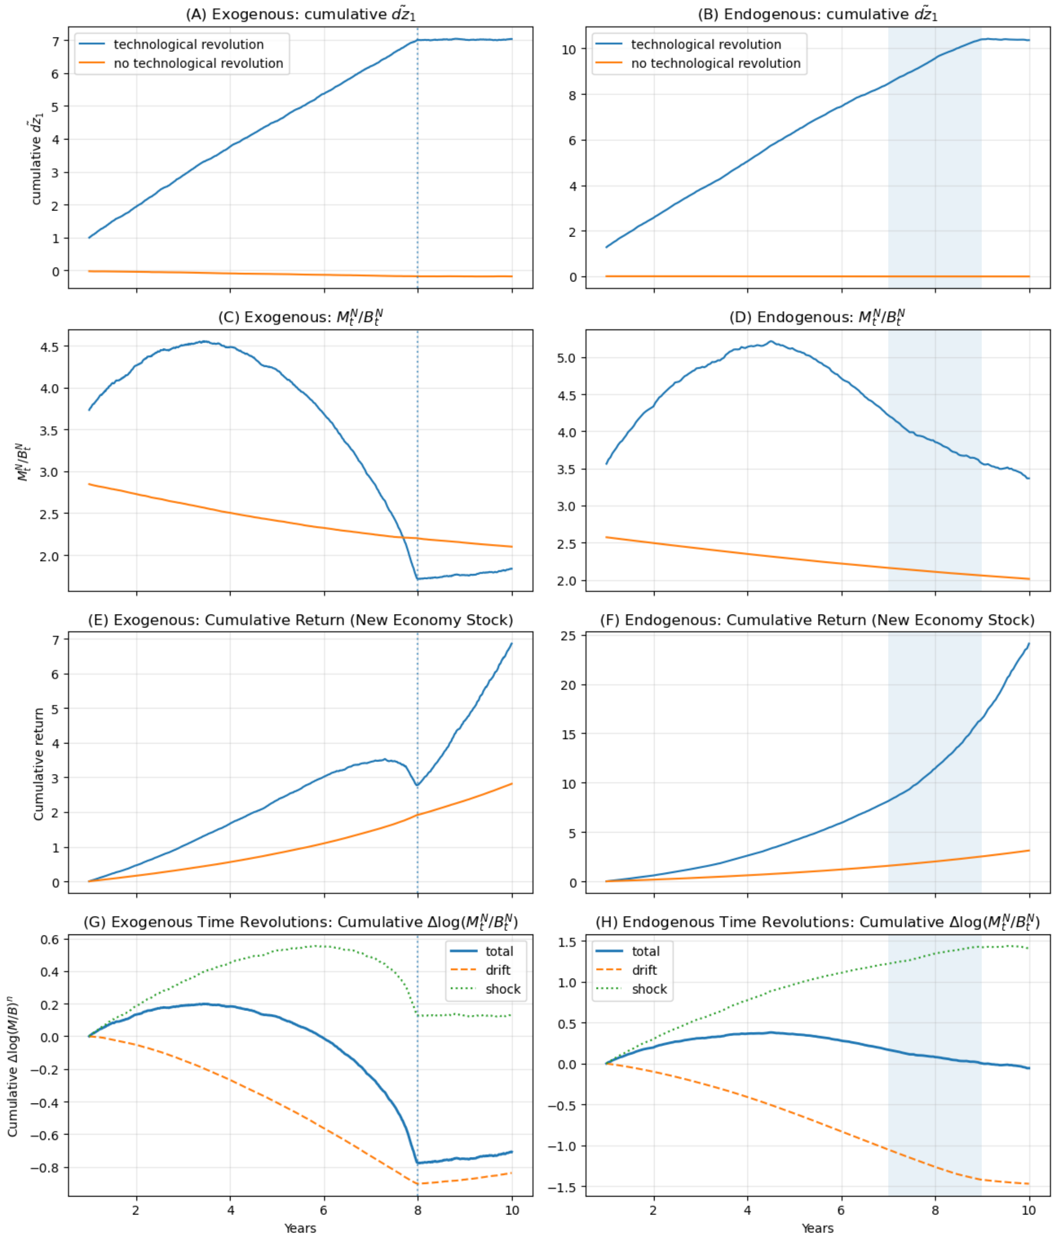

Figure 1: **New-economy bubbles in baseline calibration**

The figure shows simulated time series for the new economy from the baseline calibration of the model in Pástor and Veronesi (2009): $\gamma = 4$, $\phi = 0.3551$, $\bar{\rho} = 0.1217$, $\mu_J = 0$, $\sigma_J = 0.04$, $\sigma_0 = \sigma_{N,0} = \sigma_{N,1} = 0.07$, and $\kappa = 0.1$. The left-hand side shows results from the model where adoption decisions occur at an exogenous time $\bar{t}^{**} = 8$, while the right-hand side shows the model where endogenous adoption time is chosen optimally. Panels A to F show the time-series means conditional on adoption (blue line) and on no adoption (orange line). 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.

3.1 Discount-rate effects in technological revolutions

To illuminate the discount-rate mechanism of the PV model, we introduce a simple approximation of stock prices that reveals the disparities between the exogenous and endogenous models.

Extreme cases: Never Adopt vs. Immediately Adopt 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.[5] 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.[6] Note that the only difference between these extremes is that the immediately-adopt scenario has an additional term $-\gamma A_2(\tau)^2\hat{\sigma}_t^2$ in the exponent, representing the lower price due to the systematic risk incurred when the new technology is 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 difference between these extremes 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 when adoption no longer has any effect on the economy. Therefore, the "never adopt" case is equivalent to adopting at time T.

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

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

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 $M_t^{IA} \leq M_t^N \leq M_t^{NA}$.

Approximating Market Value in the PV model. Now consider the case where adoption is uncertain and may occur at some unknown time t^{**}. Let $\mathbb{E}_t[t^{**}]$ denote the expected adoption time. In the exogenous case, this is:

$$\mathbb{E}_t[t^{**}] = p_t \bar{t}^{**} + (1 - p_t)T, \tag{3}$$

where p_t is the probability at time t that adoption will occur at the pre-specified $\bar{t}^{**}$ (i.e. the posterior probability at time t that $\hat{\psi}_{\bar{t}^{**}} \geq \bar{\psi}$). If the adoption does not occur at $\bar{t}^{**}$, it will never occur, which is equivalent to adopting at time T. In the endogenous case, $\mathbb{E}_t[t^{**}]$ is the expected time when beliefs cross the adoption threshold conditional on current subjective beliefs.

To provide the fundamental intuition, we derive an approximation for the current market value M_t^N that is a function of $\mathbb{E}_t[t^{**}]$. We start by defining r_t^j to be the constant expected continuous discount rate, or yield, that equates the expected discounted future book value to the current market price. The yields for the two extreme market values as well as the actual new-economy stock price therefore satisfy

$$M_t^j = \mathbb{E}_t\left[B_T^N\right] \exp\left(-\int_t^T r_t^j \, ds\right), \qquad j \in \{NA, IA, N\}.$$

Our approximation then assumes that the yield associated with the market value of the new economy at any time prior to t^{**}, when the adoption is still uncertain, is

$$\int_t^T r_t^N \, ds \;\approx\; \int_t^{\mathbb{E}_t[t^{**}]} r_t^{NA} \, ds + \int_{\mathbb{E}_t[t^{**}]}^T r_t^{IA} \, ds \tag{4}$$

This approximates the true current constant yield of M_t^N as equivalent to earning the "never adopt" yield until the expected adoption time, and then the "immediately adopt" yield after that point.[7] This then implies

$$\int_t^T r_t^N \, ds \approx \frac{\mathbb{E}_t[t^{**}] - t}{T - t} \int_t^T r_t^{NA} \, ds + \frac{T - \mathbb{E}_t[t^{**}]}{T - t} \int_t^T r_t^{IA} \, ds \tag{5}$$

Using the definitions of the three yield terms, the market price of the new economy can be approximated as

$$\log\left(M_t^N\right) \approx \log\left(M_t^{NA}\right) \frac{\mathbb{E}_t[t^{**}] - t}{T - t} + \log\left(M_t^{IA}\right) \frac{T - \mathbb{E}_t[t^{**}]}{T - t}. \tag{6}$$

This equation, which we confirm works extremely well in approximating the true market value in both the endogenous and exogenous cases for most parameterizations, gives an intuitive result: The current log price of the new-economy stock is a weighted average of the immediate-adopt and never-adopt scenarios, where the weight on the immediate-adopt scenario is the portion of the remaining time to T that is expected to occur after adoption.

We then use the fact that the difference in logs between the two extremes is

$$\log(M_t^{NA}) - \log(M_t^{IA}) = A_2(\tau)^2 \gamma \hat{\sigma}_t^2$$

and rewrite the approximation as

$$\log\left(M_t^N\right) \approx log\left(M_t^{IA}\right) + \left(\frac{\mathbb{E}_t[t^{**}] - t}{T - t}\right) A_2(\tau)^2 \gamma \hat{\sigma}_t^2 \tag{7}$$

While PV highlight that systematic risk is linked to the probability of adoption, the insight

[7] This approximation implicitly assumes that the difference between expected yields in the two cases is constant, but it actually falls as uncertainty resolves. In cases where this drop is extremely fast (ϕ is very high), the approximation performs poorly. See Online Appendix for a discussion.

here is that the market price of the new economy has a discount-rate term that is approximately linear in the expected adoption time. This in turn means that, to create a large discount-rate effect, a shock must create a large change in the expected time until a technology is adopted. This approximation also allows for a pure decomposition of the cash-flow and discount-rate effects of a shock to new-economy productivity.

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:

Cash-Flow Effect (positive): Shocks to productivity raise the value of M_t^{IA} by increasing the expected terminal wealth of the new economy

$$\frac{\partial \log(M_t^{IA})}{d\tilde{Z}_{1,t}} = A_2(\tau)\frac{\partial \hat{\psi}_t}{d\tilde{Z}_{1,t}} + A_1(\tau)\frac{\partial \rho_t^N}{d\tilde{Z}_{1,t}} \geq 0 \tag{8}$$

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 M_t^{IA} (and M_t^{NA} equally as well).

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

$$\frac{\partial\left(\frac{\mathbb{E}_t[t^{**}]-t}{T-t}A_2(\tau)^2\gamma\hat{\sigma}_t^2\right)}{d\tilde{Z}_{1,t}} = \frac{\partial \mathbb{E}_t[t^{**}]}{d\hat{\psi}_t}\frac{\partial \hat{\psi}_t}{d\tilde{Z}_{1,t}}\frac{1}{T-t}A_2(\tau)^2\gamma\hat{\sigma}_t^2 \leq 0 \tag{9}$$

Here the inequality holds since $\mathbb{E}_t[t^{**}]$ decreases with an increase in the subjective belief about productivity. Therefore a positive shock to productivity 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 is 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 $dZ_{t,1}$ are equal to zero. This leads to positive values of $d\tilde{Z}_{t,1}$ 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 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 continuation option.[8]

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 7 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.

The exogenous specification behaves differently because the adoption threshold is not

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

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 described by Equation 3 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. Expected adoption time 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 7 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.[9] In contrast, the concentrated decrease in expected adoption time near the revolution in the exogenous case leads to a large discount-rate effect late in the revolution, which overwhelms the cash-flow effect and creates a bubble pattern.

Panel D provides yet another perspective on the two models. The panel graphs the exposure of new-economy stock prices to new-economy productivity shocks ($\frac{\partial \log(M_t^N)}{d\tilde{Z}_{1,t}}$). In the exogenous case, this value is strongly negative just prior to adoption, meaning the *positive* productivity shocks that are necessary to create an adoption are the actual triggers of *negative* stock returns. This is the core empirical signature of the systematic-risk mechanism. The existence of this region, which is proved in Corollary 2 of PV for the exogenous model (but not the endogenous model), is necessary to produce a bubble pattern when conditioning on an adoption. Since this region does not exist in this simulation for the endogenous model, no bubble occurs.

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

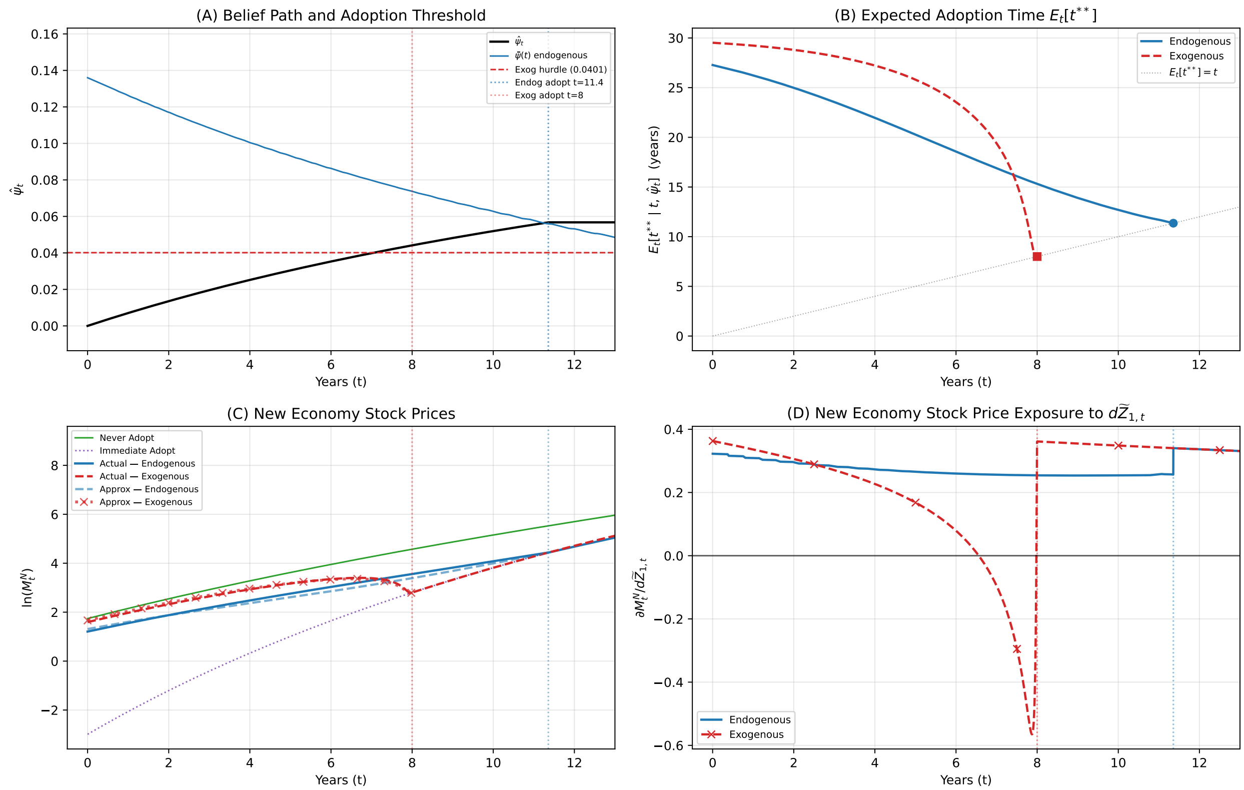

Figure 2: **Expected adoption times and new-economy stock market prices.**

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 times to adoptions $\mathbb{E}_t[t^{**}]$. Panel C shows the actual prices $\log(M_t^N)$ of the new economy in the two model specifications which fall between the two extremes of $\log(M_t^{NA})$ and $\log(M_t^{IA})$, as well as the approximated price using $\log\left(M_t^N\right) \approx \log\left(M_t^{IA}\right) + \left(\frac{\mathbb{E}_t[t^{**}]-t}{T-t}\right) A_2(\tau)^2 \gamma \hat{\sigma}_t^2$. Panel D shows the conditional exposure of the new-economy stock price to new-economy productivity shocks ($\frac{\partial \log(M_t^N)}{d\tilde{Z}_{1,t}}$).

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.[10] 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.

We do 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 is performing well across all these parameterizations.[11]

These results show that the lack of a bubble in the endogenous case holds for a wide set of reasonable parameters. While it is possible to choose parameters to generate a bubble pattern in some subset of adoptions, we find that these parameterizations are likely to be empirically implausible.[12] In short, the systematic-risk channel for price bubbles is not reasonably attainable within the PV model with endogenous adoption.

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

[11] See Figure A1 in the Online Appendix.

[12] For instance, in the Online Appendix we show that with implausibly fast mean reversion of productivity (ϕ), bubble patterns can be generated conditional on extremely fast adoption.

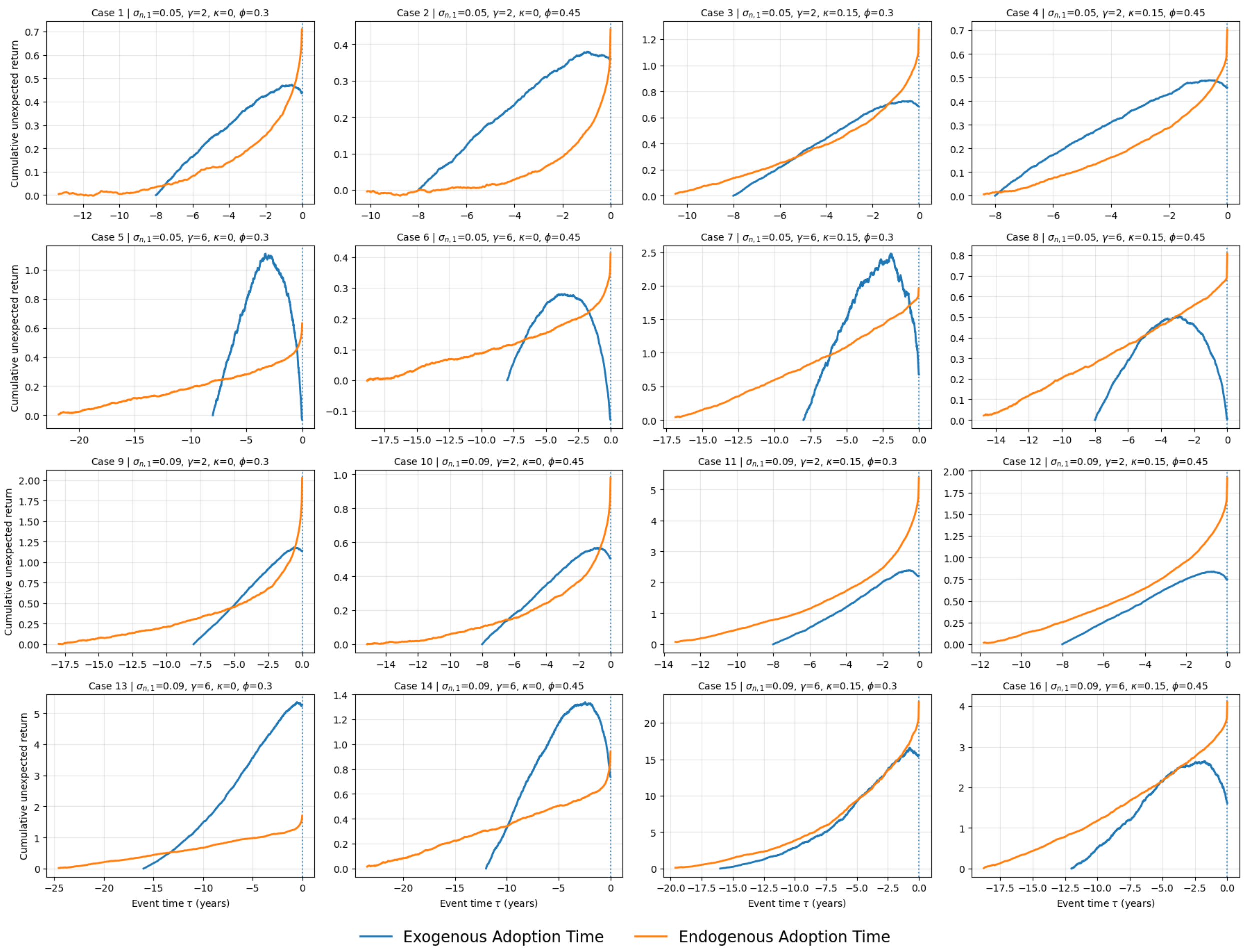

Figure 3: **Cumulative unexpected returns in revolutions across parameterizations** Each panel shows new-economy cumulative unexpected returns conditional on a technical revolution in the endogenous and exogenous cases for a given set of parameters. Exogenous adoptions occur at time $t^{**} = 8$. For the endogenous adoption time case, we show adoptions in the two-year band around the median adoption occurrence. 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.

4 Revisiting the Dotcom Bubble and Interpreting the AI Boom

The core prediction of the PV mechanism is that, at some point prior to economy-wide adoption of a new technology, the impact of positive cash-flow news on stock prices transitions from boosting prices up to pulling prices down (Panel D of Figure 2). Near the end of a revolution, new-economy prices fall despite profitability continuing to improve. In short, positive cash-flow shocks themselves pop the bubble.

We revisit the Dotcom bubble of the early 2000s with this core prediction in mind. PV cites this episode as evidence of technological revolutions inducing bubble patterns in stock prices. To proxy for productivity and expected cash flows, we focus on earnings scaled by book equity (BE). As PV notes, this return-on-equity (ROE) corresponds directly to productivity (ρ_t^N) in the model, and shocks to current ROE provide information about the long-run level of productivity (ψ). In other words, changes in current profitability represent cash-flow shocks in the model.[13]

We create a monthly series of one-year-ahead consensus forecasts of earnings from LSEG (2026).[14] As a robustness check, we also examine actual earnings realized for the fiscal quarter in which the current month resides. To aggregate up, we sum total earnings across companies and divide by their cumulative book equity obtained from S&P Global Market Intelligence (2026).

For the Dotcom episode, we follow Brunnermeier and Nagel (2004) and use the top quintile of price-to-sales firms listed on NASDAQ (at end of 1996) as our new-economy stocks and NYSE/AMEX as the old-economy stocks. The top two panels of Figure 4 show the realizations of these cash flow measures, along with the relevant stock indices created

[13] Pástor and Veronesi (2006) use ROE as a proxy for cash flows in a similar model setup as PV, and ROE is also used as a proxy for cash-flow shocks in the broader literature (e.g., Vuolteenaho (2002)).

[14] This measure is constructed as a weighted average of fiscal-year-one (FY1) and fiscal-year-two (FY2) forecasts, where the weights are chosen so that the weighted average of the upcoming fiscal year-end date and the next fiscal year-end date is one year from the current date.

using data from Center for Research in Security Prices (CRSP).

Prices of Dotcom stocks exhibit the familiar boom–bust pattern, with prices rising sharply relative to the NYSE/AMEX through 1999 and early 2000 before collapsing. Using either ROE measure, this price collapse was accompanied by a period of declining ROE rather than rising cash-flow prospects. This finding is not new to the literature (e.g., Pástor and Veronesi (2006) and Gómez-Cram and Lawrence (2025)).[15] Moreover, Dotcom ROE declines far more than old-economy ROE before the end of the Internet revolution, estimated in PV to be mid-2002. Hence, the Dotcom bubble is missing the good profitability news required to trigger the discount-rate effects of the PV model.

Lastly, we consider the recent AI boom because it has notably been discussed in connection with the PV model.[16] The bottom two panels of Figure 4 show ROE for the union of stocks held by the five largest AI-focused ETFs according to ETF.com (2026) as of January 2026. Again the NYSE/AMEX serves as the old economy.[17] As the figure shows, AI stock prices rise sharply relative to NYSE/AMEX stocks, with ROE for AI firms rising much more than old-economy ROE.

Given that economy-wide adoption of AI appears likely but still uncertain (see, e.g., Bick et al. (2026) and Yotzov et al. (2026)), this is the time when AI-productivity shocks should have the largest impact on expectations regarding broader adoption. In other words, the transition of positive cash-flow news from boosting to sinking stock prices should likely already have begun according to PV. The caveat here is that it is difficult, if not impossible, to precisely calculate an economy-wide adoption probability, so any conclusions drawn from the AI plots will be somewhat subjective. However, the finding of Yotzov et al. (2026) that approximately 70% of all firms are actively incorporating AI suggests that adoption is well underway, with no sign of the price drops predicted by PV.

[15]For instance, Pástor and Veronesi (2006) states "there was a fundamental reason for Nasdaq prices to come down after the 1990s, namely, an unprecedented decline in the profitability of Nasdaq-traded firms in 2000 and 2001".

[16]See footnote 1.

[17]The AI ETFs are AIQ, ARTY, ARKQ, BAI, and BOTZ. In unreported results, we examine other indices and sets of stocks and find that results are not sensitive to the choice of index.

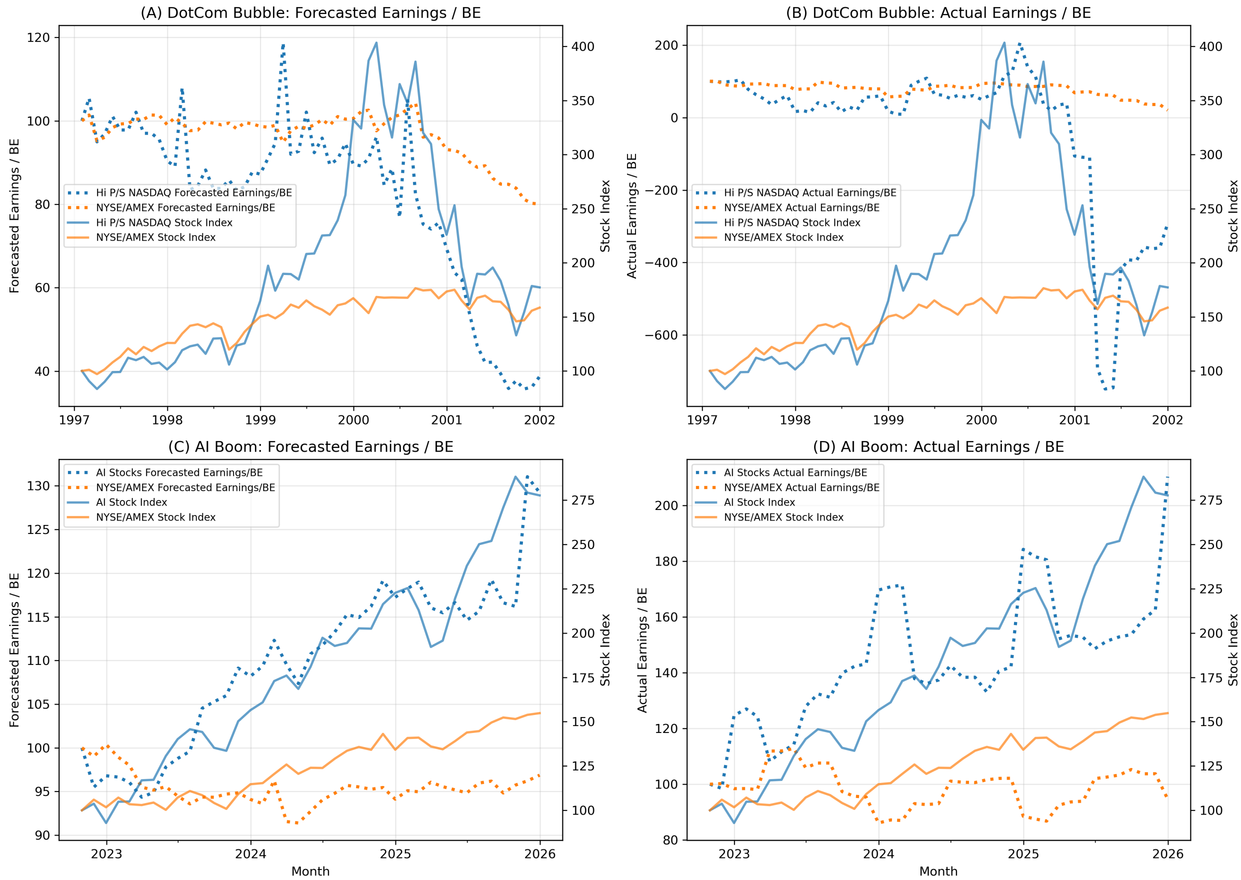

Figure 4: **Earnings and stock prices during the Dotcom Bubble and AI Boom.** The top two panels show the Dotcom bubble (1997–2002) and plots values for the top price/sales quintile of NASDAQ stocks. The bottom two panels show the AI boom (October 2022–December 2025), where the AI portfolio consists of the union of firms held by AI-focused ETFs (AIQ, ARTY, ARKQ, BAI, and BOTZ). In all plots, values for all NYSE/AMEX stocks are shown as comparison. The left column plots monthly one-year-ahead forecasts of earnings divided by book equity where the forecasts are weighted averages of FY1 and FY2 from the LSEG (2026) consensus file with the weights chosen so that the weighted average of FY1 and FY2 end dates is one year from the current date. Forecasts are converted to dollars using shares outstanding, summed across firms in the universe, and scaled by total book equity across all firms. The right column uses actual realized earnings, calculated for the fiscal quarter in which a month falls and obtained from S&P Global Market Intelligence (2026). All panels include value-weighted stock indices constructed with data from Center for Research in Security Prices (CRSP). All series are normalized to 100 at the start of the period.

5 Conclusion

In this paper, we examine the mechanism proposed in Pástor and Veronesi (2009) for explaining *ex post* stock price bubbles associated with transformative technologies. We show that the model calibration in Pástor and Veronesi (2009) only produces a stock price bubble in the illustrative case with the simplifying assumption of an all-or-nothing adoption decision at an exogenous point in time. In the more realistic case of an endogenous optimal adoption time, the baseline calibration yields no stock price bubble.

We show that this difference is due to the structure of the two specifications. To do so, we derive a simple approximation for stock prices in the model that applies to both, and use it to show that the bubble pattern arises as a result of the simplifying assumption, and is not a natural feature of a learning model with an endogenous adoption time.

In addition, we revisit the Dotcom bubble of the late 1990s and early 2000s. The Pástor and Veronesi (2009) model requires that positive productivity shocks in the late stage of a successful revolution drive prices down. Dotcom stocks display no evidence of this requirement.

Finally, our findings have implications for the current AI stock price boom. Given that the adoption of AI into the broader economy is well underway, it is likely the case that beliefs about systematic risk are changing as expectations for AI's usefulness adjust. This is precisely the time that the discount-rate channel should be most active. Thus far however, the AI stock prices have continued to rise with profitability.

To the extent that discount rates are now a function of the economy-wide use of AI technology, we argue that changes to the discount rate will be incrementally small and incapable of overwhelming the positive cash-flow shocks required to justify a technological revolution as modeled in Pástor and Veronesi (2009). Our findings suggest that, if an AI "bubble" does in fact burst, rising discount rates driven by an increasing probability of an economy-wide AI adoption is unlikely to be the reason.

References

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

Bick, Alexander, Adam Blandin, and David J Deming, "The rapid adoption of generative AI," *Management Science*, 2026.

Brunnermeier, Markus K. and Stefan Nagel, "Hedge funds and the technology bubble," *Journal of Finance*, 2004, *59* (5), 2013–2040.

Center for Research in Security Prices (CRSP), "CRSP US Stock Database," Dataset 2026. Chicago Booth School of Business / WRDS. Accessed March 2026.

ETF.com, "Artificial Intelligence ETFs," ETF.com topic page 2026. Accessed January [day], 2026.

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.

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

Gómez-Cram, Roberto and Alastair Lawrence, "The value of software," *American Economic Review*, 2025, *115* (10), 3514–3558.

Greenwood, Robin and Stefan Nagel, "Inexperienced investors and bubbles," *Journal of Financial Economics*, 2009, *93* (2), 239–258.

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.

LSEG, "I/B/E/S Summary History and Detail History," Dataset, accessed via Wharton Research Data Services (WRDS) 2026. Accessed March 2026.

Ofek, Eli and Matthew Richardson, "DotCom Mania: The Rise and Fall of Internet Stock Prices," *Journal of Finance*, 2003, *58* (3), 1113–1137.

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.

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

S&P Global Market Intelligence, "Compustat North America," Dataset, accessed via Wharton Research Data Services (WRDS) 2026. Accessed March 2026.

Vuolteenaho, Tuomo, "What drives firm-level stock returns?," *Journal of Finance*, 2002, *57* (1), 233–264.

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

Yotzov, Ivan, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Steven J Davis, Kevin M Foster, Aaron Jalca, Brent H Meyer, Paul Mizen, Michael A Navarrete et al., "Firm Data on AI," Technical Report, National Bureau of Economic Research 2026.

Internet Appendix

Technological Revolutions and Stock Prices Revisited

Roberto C. Gutierrez Jr. Robert C. Ready

This appendix provides additional detail on the endogenous adoption model of Pástor and Veronesi (2009). We first provide an additional perspective on Figure 3 in the main text, and then highlight a situation in which the endogenous model can produce a bubble pattern in some adoptions, but argue that this calibration is not a realistic possibility for observed bubble patterns.

A1 Adoptions at various times across parameterizations

Figure 3 in the main text shows average price patterns for adoptions in different parameterizations of the model. Figure A1 replicates that figure but shows price patterns for adoptions occurring at different times. To create these figures we draw individual values of ψ and generate single adoption paths in the manner of Figure 2 in the main text. We draw values of $\psi = k \cdot \sigma_J$ where k takes integer values from one to six (inclusive). Higher values of ψ generate quicker adoptions as subjective beliefs rise more quickly. For all parameterizations, all of the observed adoptions do not have the bubble pattern. Even though expected adoption times fall quickly for high draws of ψ, they still fall smoothly and no bubble patterns arise.

The figure also shows the performance of the approximation derived in the main text. For most parameterizations the values calculated using the approximate price track very closely with the true prices, indicating the approximation is performing well. The notable exception is that the approximation is less accurate, though still generates qualitatively similar behavior, when values of ϕ and γ are high. We therefore explore more extreme parameterizations in this direction.

A2 Producing bubbles with extreme parameterizations of the endogenous model

Our goal with theoretical argument in the main paper is not to show that the endogenous model is incapable of generating bubble-like price patterns. Rather, our argument is that under the benchmark specification in Pástor and Veronesi (2009), and for a wide range of nearby parameterizations, such patterns do not arise along equilibrium adoption paths, suggesting that they are not a typical feature of the model.

In the main text, we argue that reason for this is that learning generates a downward-sloping endogenous adoption threshold, which implies that the expected adoption time $\mathbb{E}_t[t^{**}]$ responds smoothly to productivity shocks, even when adoption is expected far in the future. When you combine this fact with the approximation for market prices, in which the discount-rate impact of the rising systematic risk associated with adoption is linear in $\mathbb{E}_t[t^{**}]$, this then implies that increases in systematic risk are spread over the entire revolution, rather than concentrated near adoption as in the exogenous-timing model.

Despite this general result, there exist parameterizations under which the model can generate bubble-like price dynamics. In a broad search over parameter values, we find that such patterns arise only under extreme values of ϕ, corresponding to extremely transitory productivity shocks to the economy, which in turn imply very rapid learning. While we were unable to derive a formal characterization of this region, we present one such parameterization and explain why it generates a bubble pattern, and also why we view it as empirically implausible.

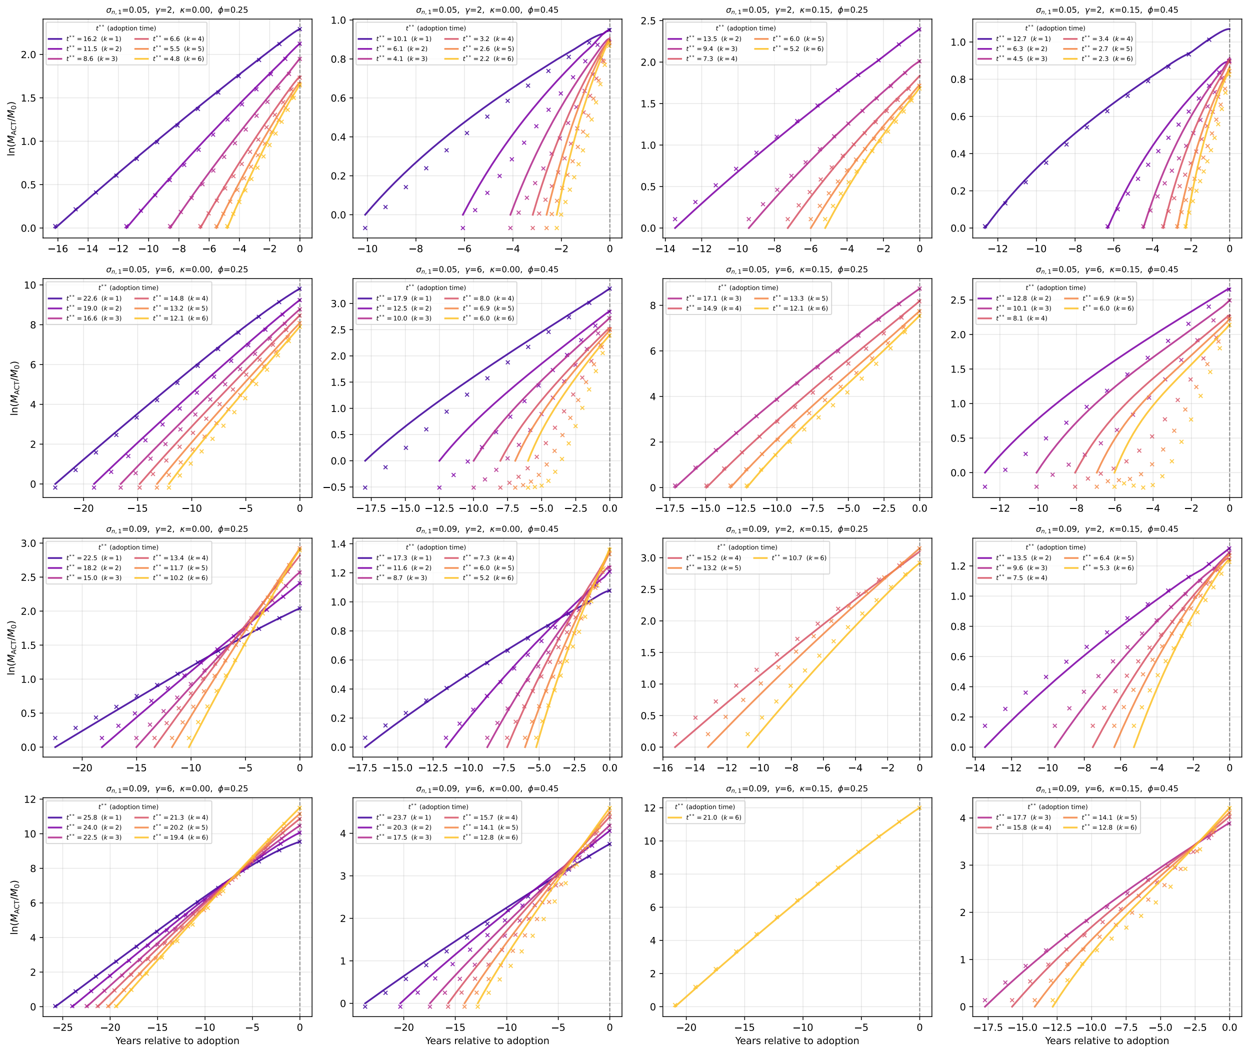

Figure A1: **Cumulative unexpected returns in revolutions across parameterizations: Various adoption times** Each colored line shows the log price path $\log(M_t^N/M_0^N)$ in the endogenous adoption-time model for a different value of the true technology quality $\psi_{\text{true}} = k \cdot \sigma_J$, $k = 1, 2, \ldots, 6$, under the baseline calibration ($\phi = 0.3551$, $\gamma = 4$, $\kappa = 0.1$). Paths are deterministic (Kalman filter with no idiosyncratic noise). Each panel shows a different parameterization corresponding to the parameterizations in Figure 3 of the main text. The (x) symbols plot the same values calculated using the approximation for the log of new-economy stock prices derived in the main text.

The emergence of a bubble pattern is closely related to the accuracy of the approximation derived in the main text. That approximation expresses the current yield of the new economy r_t^N, defined as the constant continuous discount rate that equates the current market price to the discounted value of the expected terminal cashflow, as a time-weighted average of the current value of the never-adopt (NA) and immediate-adopt (IA) yields. Under this assumption, we have

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

which produces the weighted log-price formula

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

The key step is treating the current IA expected return over the remaining horizon as a constant yield. In reality, however, the IA expected return is time-varying. At a future date $s > t$, the difference in expected returns for the NA and IA cases is given by $\gamma \, A_2(T-s)^2 \, \hat{\sigma}_s^2/(T-s)$, which declines over time as learning reduces $\hat{\sigma}_s^2$. As a result, an adoption that occurs later is associated with a lower IA expected return than one that occurs earlier. Using an implied constant value of r_t^{IA} for all future expected returns ignores this variation and treats the remaining expected return of the IA case as independent of the timing of adoption.

This distinction becomes important when learning is very rapid. Since subjective volatility evolves according to

$$\hat{\sigma}_t^2 = \frac{\sigma_0^2}{1 + \frac{\phi^2}{\sigma_1^2} \, \sigma_0^2 \, t},$$

large values of ϕ cause uncertainty to collapse quickly early in the revolution. In this case, the difference between IA and NA expected returns is sharply declining over time: it is large initially and close to zero by the middle of the horizon.

As a result, changes in $\mathbb{E}_t[t^{**}]$ have different effects depending on their timing. A decrease in $\mathbb{E}_t[t^{**}]$ shifts potential adoption into an earlier period when $\hat{\sigma}_t$ is still high, implying that the relevant IA yield at adoption is higher than the average yield over the entire interval. Subsequent increases in productivity therefore have a larger discount rate effect. The constant-yield approximation therefore understates the increase in discount rates associated with earlier adoption. For sufficiently early adoption paths, this amplified discount-rate effect can dominate the cash-flow component, producing the boom-and-bust patterns observed in the simulations.

Figure A2 illustrates this mechanism. The first panel shows the benchmark calibration. The second increases ϕ, which accelerates learning but does not generate a bubble. The third increases risk aversion γ, again without producing bubble-like dynamics. The fourth panel combines high ϕ and high γ, generating bubble patterns for paths with very early adoption.[1]

Figure A3 provides further context for this parameterization. The first two panels repeat the price paths from the previous figure, and then show the performance of the approximation for the extremely fast adoption case with the most pronounced bubbles. As the figure shows, the approximation performs very poorly due to the argument above, as it does not capture the differences in the level of the systematic risk between an early and late adoption arising from the large differences in subjective volatility created by very fast learning.[2]

However, we think that there are three features of this mechanism that limit its empirical relevance. First, it requires extreme realizations of ψ_{true}: only the paths with $\psi_{\text{true}} \geq 4\sigma_J$ exhibit a bubble, which is far out in the right tail of the prior. As panel (d) of Figure A3 shows, the average adoption path across Monte Carlo draws does not have a noticeable bubble, because such extreme early adoptions are extremely rare (they occur with 4, 5, and 6 standard-deviation draws of ψ).

[1]In unreported results, we find that increasing γ or κ alone does not generate bubble patterns. With high ϕ and γ, increasing κ can amplify the observed bubble patterns, but this reflects the finite-horizon "use it or lose it" feature of κ arising as an artifact of the finite-time horizon.

[2]These differences are apparent in the graph via the initially large, but quickly narrowing, gap between the IA and NA prices.

Second, the rapid collapse of $\hat{\sigma}_t^2$ produces a pronounced U-shape in return volatility, visible in panel (c) of Figure A3. Annualized return volatility begins near 35%, falls to roughly 5% as the posterior variance collapses, and then rises back to near 35% as the remaining gap Δ_t is traversed near the adoption event. This pattern arises because return volatility through the learning about cash-flows channel is proportional to $\hat{\sigma}_t$, which is quickly falling due to rapid learning. Additionally, the positive productivity shocks $d\tilde{Z}_{1,t}$ that push $\hat{\psi}_t$ toward the adoption generate large negative price effects near the boundary through the rising adoption discount when $\hat{\sigma}_t$ is still large further attenuating the cash-flow effect. This channel is eliminated after adoption allowing volatility to rebound quickly (it then falls quickly again with learning). This extreme volatility path, and in particular the large drop in volatility prior to the price drop, is not characteristic of observed technology booms.

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

References

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

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

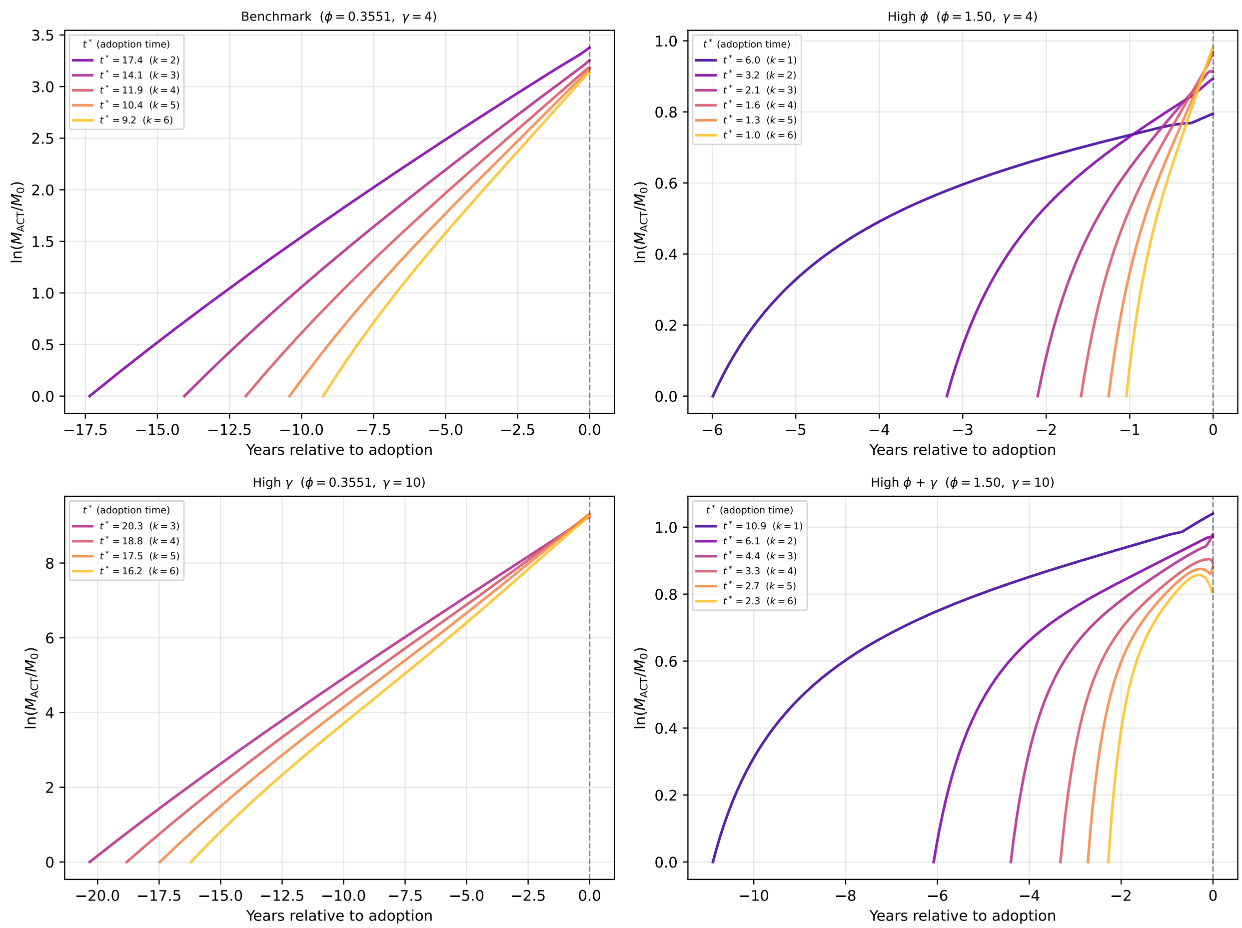

Figure A2: **Deterministic price paths: 2×2 parameterization grid (endogenous model).** Each panel plots $\log(M_t^N/M_0^N)$ paths for $\psi_{\text{true}} = k \cdot \sigma_J$, $k = 1, \ldots, 6$ (colored lines), as in Figure A1. The four panels vary the speed of mean reversion ϕ and the coefficient of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

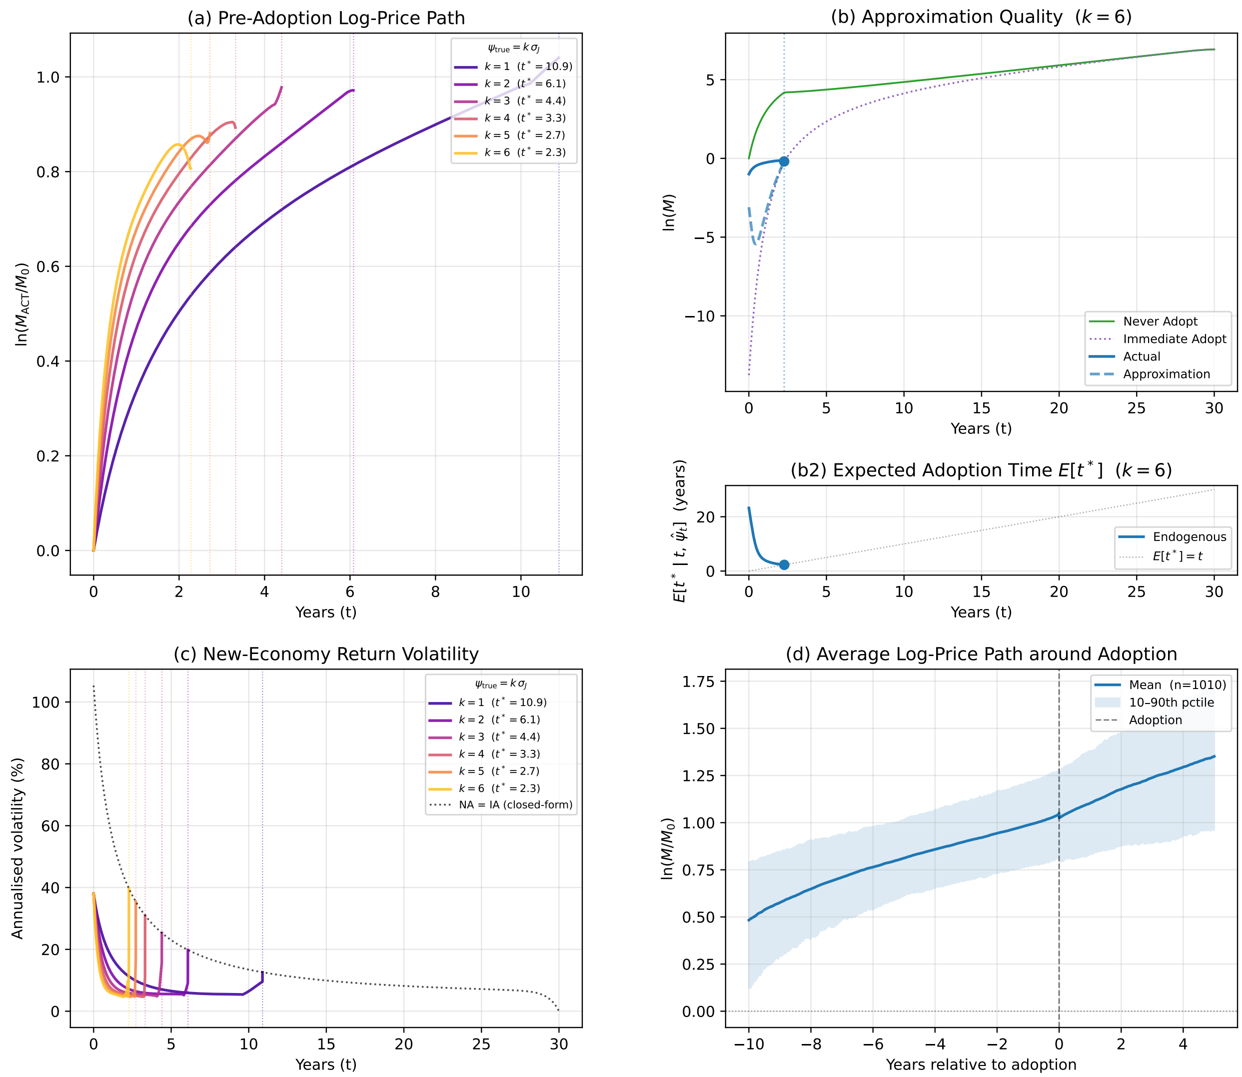

Figure A3: **High-ϕ, high-γ parameterization: price paths, beliefs, volatility, and Monte Carlo average (endogenous model, $\phi = 1.50$, $\gamma = 10$).** Panel (a): $\log(M_t^N/M_0^N)$ for $\psi_{\text{true}} = k \cdot \sigma_J$, $k = 1, \ldots, 6$ (colored lines). Panel (b): Plots of actual M_t^N/M_0^N along with the approximated value for the fast adoption ($k = 6$) simulation. Panel (b2) shows the expected time to adoption. Panel (c): annualized return volatility for each path. Panel (d): average $\log(M_t^N/M_0^N)$ path aligned to adoption from a Monte Carlo simulation with $N = 5{,}000$ draws (shaded band shows the 10th–90th percentile range). Adoption cost $\kappa = 0.1$.