Will Systematic Risk Burst the AI Bubble? Technological Revolutions and Stock Prices Revisited*

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

January 20, 2026

Abstract

Pástor and Veronesi (2009) provide a rational explanation for stock price bubbles observed during technological revolutions. We argue that the proposed mechanism, based on sharply rising systematic risk as the new technology nears wider adoption, is unlikely to explain bubbles. We show that this mechanism is only present in an illustrative model with the simplifying assumption that a one-time all-or-nothing adoption decision for the new technology occurs at an exogenously fixed date. In contrast, there is no stock price bubble in the baseline calibration of the more realistic model that allows for optimal adoption of the technology at any time. More generally, we argue that the bubble pattern is not a natural feature of the more realistic model. We also argue that the behavior of earnings and stock prices in the DotCom bubble, as well as the current AI stock boom, provide no evidence of the necessary mechanism. Our results suggest that the bursting of an AI bubble is unlikely to arise from increased systematic risk accompanying further positive news about AI technology.

*We thank Ľuboš Pástor and Pietro Veronesi for making their replication code available. Generative AI tools were used in both the analysis and the writing for this paper. All AI output used in the paper was verified by the authors. All 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) provide a rational explanation for bubble patterns observed in the stock prices of firms with transformative technologies. *Ex post*, technological revolutions are characterized by a long series of unexpectedly positive productivity shocks leading to the technology's broad adoption in the economy. These shocks lead to an increase in expected cash-flow that increases stock prices, but they also increase the probability of wider adoption, which leads to an increase in systematic risk that decreases stock prices. Early in the revolution, the probability of adoption is insignificant, so the first channel dominates and prices rise. Later in the revolution, when the probability of adoption becomes high, the second effect dominates and prices fall, creating the characteristic bubble pattern.

PV has been highly influential, and its predictions are often at the center of recent discussions regarding a potential stock price bubble for companies involved in the development of generative artificial intelligence (AI).[1] In particular, the PV mechanism suggests that further upward revisions in beliefs regarding AI's potential to revolutionize the economy may generate a sharp decline in AI stock prices.

In this paper, we argue that this outcome is unlikely to occur. Our first argument is theoretical. There are two versions of the model in PV: an illustrative version under the simplifying assumption that the representative agent makes a one-time all-or-nothing decision whether to adopt the new technology at an exogenously preset date, and a more realistic version where an endogenous adoption time is chosen optimally. We show that in the baseline calibrations of PV, the *ex post* bubble is only present in the illustrative case. While the endogenous adoption time model generates a bubble pattern in market-to-book (M/B) ratios, we show that this is primarily due to expected mean-reversion in the M/B ratio. In the baseline calibration of the more realistic model, there is no period of negative average returns, and hence no price bubble in a typical technological revolution.

[1] For example, see the article "This is how the AI stock boom plays out". https://www.bloomberg.com/opinion/articles/2025-10-28/this-is-how-the-ai-stock-boom-plays-out

We also argue that the bubble highlighted in PV is a *result* of the simplifying assumption of an exogenously specified adoption date, rather than a natural feature of a more realistic model. To do this, we derive a simple approximation that links the exogenous and endogenous specifications by expressing prices as a function of cash-flow expectations and a systematic risk term that scales linearly with the expected time until adoption. Using this approximation, we demonstrate that a price crash requires a region in which small cash-flow shocks generate large reductions in the expected time remaining until the new technology is adopted.

The assumption of an all-or-nothing decision at a preset time creates exactly such a region. Approaching the preset time, if beliefs about the technology's improvements to productivity are near the level required for adoption, small productivity shocks can sharply increase adoption probabilities. The expected adoption time then jumps from very far in the future (effectively the terminal time T if the technology is not adopted) to the immediate present. The resulting increase in systematic risk raises discount rates enough to overwhelm the positive cash-flow effects of the shock, generating large negative returns.

In contrast, when adoption timing is endogenous, adoption occurs when productivity crosses a continuous boundary. This boundary is generally downward sloping in time due to the gradual learning about productivity improvements that reduces both the uncertainty of the new technology's productivity and the value of the continuation option to keep learning. This implies that productivity shocks translate into smooth, incremental changes in expected adoption times, even far from the boundary. As a result, the increase in systematic risk is spread evenly over the course of the revolution, preventing it from dominating cash-flow effects and thereby eliminating any declines in prices.

We believe that this core feature of a learning model, the reduction of uncertainty as time passes, will make it very difficult to generate a bubble pattern in the endogenous case with realistic parameters. However, we note that even the endogenous adoption time model can be parameterized to generate a bubble, notably with an up-front adoption cost that is much higher than in the baseline calibration.[2] We therefore present a second, empirical, examination of the mechanism. We argue that the empirical signature of the discount-rate mechanism in the PV model is that a series of *positive* cash-flow shocks generate realized returns that are first increasing, and then decreasing, as the impact of these shocks on the discount rate is changing over time. In other words, conditional on a series of positive cash-flow shocks strong enough to create a revolution, the latter stage of the revolution is characterized by expected cash-flows which are continuing to rise while stock prices begin to fall.

To empirically examine this mechanism, we construct measures of forecasted aggregate earnings for both NASDAQ companies in the late 1990s and early 2000s around the DotCom Bubble, and in the more recent boom in AI stocks. We find that the boom and bust period of the DotCom Bubble corresponded to rising, and then falling, expected cash-flows. Likewise in the current AI Boom, stock prices are closely tracking forecasts of future earnings despite the fact that adoption probabilities are currently likely to be high, though extremely uncertain, which is precisely the time when the discount-rate mechanism should be most active. In short, we do not find positive earnings shocks to be paired with price falls.

Taken together, our theoretical and empirical arguments do not suggest that observed bubbles are inherently evidence of irrational investors.[3] Our argument here is simply that increasing systematic risk is unlikely to create a sudden drop in stock prices for new technologies. In the context of AI technologies, an increase in discount rates from positive news about improvements in AI productivity is unlikely to cause a crash.

[2]This is an artifact of the finite-horizon model specification. See Section 3 for a discussion.

[3]For instance, Pástor and Veronesi (2006) show that a highly uncertain new technology should have high valuations through a simple Gordon Growth argument: $\frac{1}{r-E[g]} < E[\frac{1}{r-g}]$ when g is uncertain. This means that, for the median uncertain technology, the return will be negative (i.e. if uncertainty about g resolves to E[g]).

2 Related Literature

While the paper most closely related is 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.

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 in Pástor and Veronesi (2009)

We first briefly review the model of Pástor and Veronesi (2009).[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 are infinitesimal in size relative to the old

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

economy. Hence, the new economy only impacts terminal wealth through its impact on the productivity process of the old economy p_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

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

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 $\hat{\psi}_t$. Following PV, we consider these orthogonalized unanticipated shocks to new economy productivity (controlling for shocks to old economy productivity) as $d\tilde{Z}_{1,t}$.

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 show, 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

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

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

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. The PV model also 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 M_t/B_t and M_t^N/B_t^N to normalize the market values. 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 by PV 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. We note that this figure is a replication of some of the results in Figures 3 and 4 in PV which present results for the exogenous adoption time case, and Figure 6 which presents results for the endogenous adoption time case. Following PV, here we consider adoptions in the endogenous case that occur within one year of the exogenous threshold $E[t^{**}] = 8$.

The left-hand panels represent the exogenous adoption time case, and the right-hand panels represent the endogenous adoption time case. As Panels A and B show, a technology adoption is characterized, *ex post*, by a long string of unexpectedly positive productivity shocks. These raise the subjective belief $\hat{\psi}_t$ enough that it becomes optimal to adopt. 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. The intuition is that early in the revolution, adoption probability is low, so the increased cash-flow effect leads to positive returns and increasing valuations. Later in the revolution, those same positive cash-flow shocks also increase the discount rate through the increased probability of adoption, which overwhelms the cash-flow effect and leads to falling prices. This intuition is precisely the cause of the bubble in the exogenous case. Panel E plots cumulative returns to the new economy in this case, which inherit the bubble pattern of M/B.

For the endogenous adoption time case, the pattern is very different. As Panel F shows, cumulative returns rise steadily throughout the revolution in the baseline calibration, with no bubble in stock prices. This may seem surprising given that 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 will argue that the difference is not due to a particular calibration, but instead due to a fundamental difference in discount-rate dynamics across the two models.

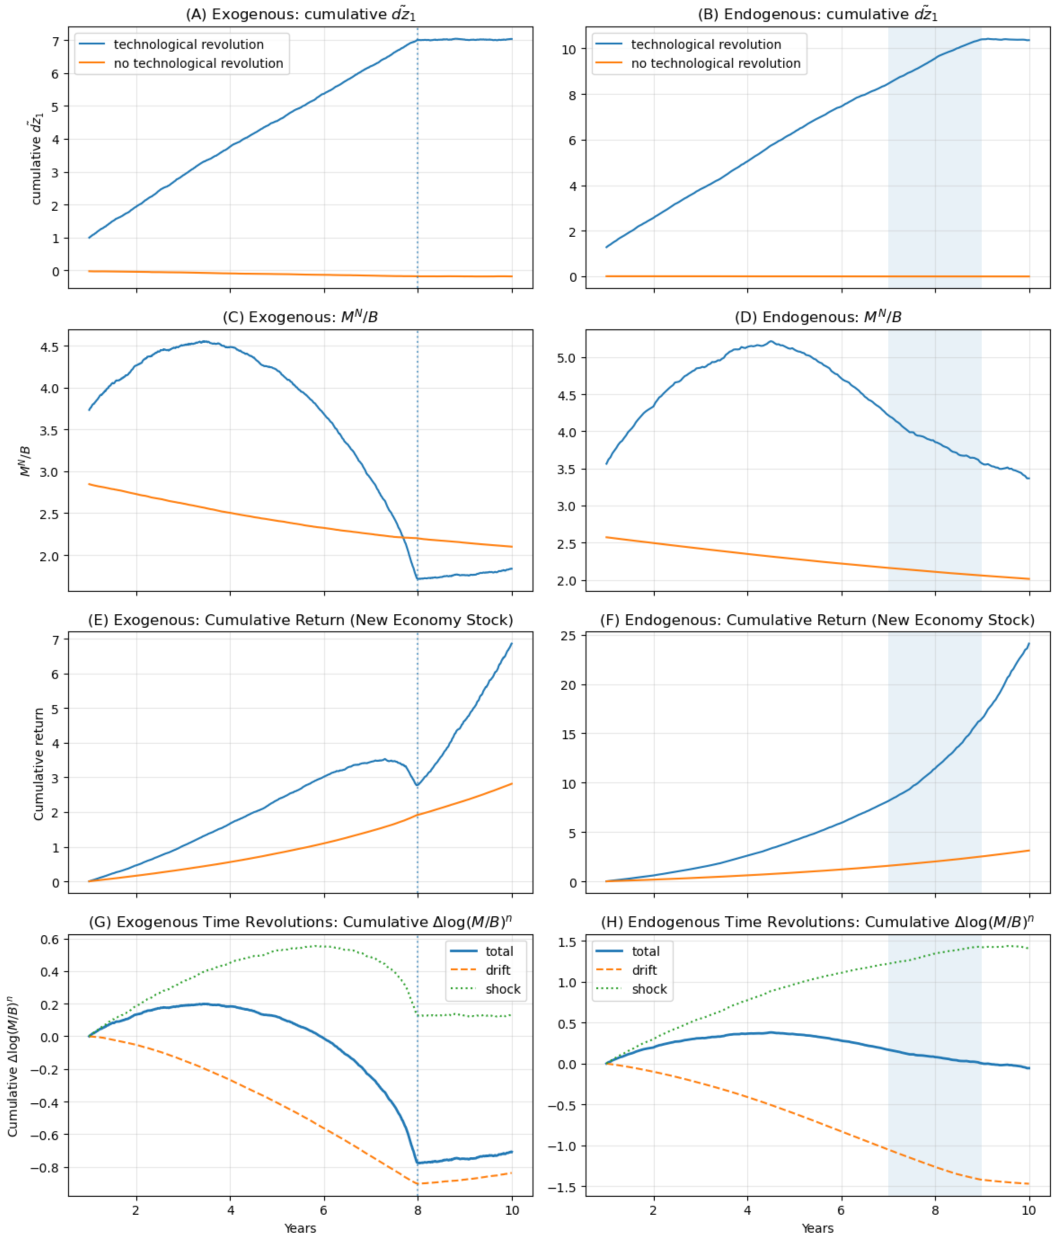

Figure 1: **New economy bubbles in baseline calibration**

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 average conditional on adoption (blue line) and on no adoption (orange line). Following PV in the endogenous time case we focus on 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 adoption M/B ratios in Panels C and D into unexpected innovations and conditionally expected drift.

3.1 Discount-rate effects in technological adoptions

For the purpose of understanding the model, and in particular the relation between the exogenous and endogenous adoption models, we introduce a simple approximation to help understand the mechanism.

Extreme cases: Never Adopt vs. Immediately Adopt Consider the new economy at any time t, and define $\frac{M_t^{N,NA}}{B_t^N}$ ("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 $\frac{M_t^{N,IA}}{B_t^N}$ ("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

$$\frac{M_t^{N,NA}}{B_t^N} = e^{\bar{C}_0(\tau)+A_1(\tau)\rho_t^N+A_2(\tau)\hat{\psi}_t+\frac{1}{2}A_2(\tau)^2\hat{\sigma}_t^2}$$

$$\frac{M_t^{N,IA}}{B_t^N} = e^{\bar{C}_0(\tau)+A_1(\tau)\rho_t^N+A_2(\tau)\hat{\psi}_t+\frac{1}{2}A_2(\tau)^2(1-2\gamma)\hat{\sigma}_t^2}$$

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 is in the $\hat{\sigma}_t^2$ term in the exponent. The immediately-adopt scenario has an additional $-\gamma A_2(\tau)^2\hat{\sigma}_t^2$ 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 $M_{IA}^N \leq M_{NA}^N$. Note that this difference 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 new 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

[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 $t \in [t^*, t^{**}]$.

case for systematic risk. If you wait to adopt, there will be less time remaining for shocks to the new economy to impact expectations about the old economy. Likewise, committing to never adopt is the best possible case for systematic risk. 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 $M_t^{N,IA} \leq M_t^N \leq M_t^{N,NA}$.

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:

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

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, when the decision no longer impacts new-economy value. In the endogenous case, $\mathbb{E}_t[t^{**}]$ is the expected time when beliefs cross the adoption threshold conditional on current subjective beliefs.

To help provide intuition, we derive an approximation for the current market value M_t^N that is a function of $\mathbb{E}_t[t^{**}]$. To do so, we start with an equation to define the two extreme values in terms of expected book value and an expected continuous discount rate, or yield, of $r_t^{N,NA}$ or $r_t^{N,IA}$

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

Prior to adoption t^{**}, when the adoption is still uncertain, we define r_t^N as the constant rate of expected return that equates expected discounted future book value to the current market price so that $M_t^N = \mathbb{E}_t\left[B_T^N\right] \exp\left(-\int_t^T r_t^N\, ds\right)$. We then assume as an approximation

$$\int_t^T r_t^N\, ds \;\approx\; \int_t^{\mathbb{E}_t[t^{**}]} r_t^{N,NA}\, ds + \int_{\mathbb{E}_t[t^{**}]}^T r_t^{N,IA}\, ds \tag{3}$$

This approximates the uncertain adoption case by assuming that the current expected return is equivalent to earning the "never adopt" expected return until the expected adoption time, and then the "immediately adopt" expected return after that point. We then further approximate this by

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

We then have that market to book in the uncertain adoption case is simply

$$\log\left(\frac{\hat{M}_t^N}{B_t^N}\right) \approx \log\left(\frac{M_t^{NA}}{B_t^N}\right) \frac{\mathbb{E}_t[t^{**}] - t}{T - t} + \log\left(\frac{M_t^{IA}}{B_t^N}\right) \frac{T - \mathbb{E}_t[t^{**}]}{T - t}. \tag{5}$$

This equation, which we confirm works extremely well in approximating the true market value in both the endogenous and exogenous case, 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 extreme M/B ratio cases is

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

And rewrite this approximation as

$$\log\left(\frac{M_t^N}{B_t^N}\right) \approx log\left(\frac{M_t^{N,IA}}{B_t^N}\right) + \left(\frac{\mathbb{E}_t[t^{**}] - t}{T - t}\right) A_2(\tau)^2 \gamma \hat{\sigma}_t^2 \tag{6}$$

The insight here is that the market price of the new economy has a discount-rate term that is 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 remaining 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 $M_t^{N,IA}$ by increasing the expected terminal wealth of the new economy

$$\frac{\partial \log(M_t^{N,IA})}{d\tilde{Z}_{1,t}} = A_2(\tau)\frac{\partial\hat{\psi}_t}{d\tilde{Z}_{1,t}} + A_1(\tau)\frac{\partial\rho_t^N}{d\tilde{Z}_{1,t}} \geq 0 \tag{7}$$

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

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

$$\frac{\partial\left(\frac{\mathbb{E}_t[t^{**}]-t}{T-t}A_2(\tau)^2\gamma\hat{\sigma}_t^2\right)}{d\tilde{Z}_{1,t}} = \frac{\partial\mathbb{E}_t[t^{**}]}{d\hat{\psi}_t}\frac{\partial\hat{\psi}_t}{d\tilde{Z}_{1,t}}\frac{1}{T-t}A_2(\tau)^2\gamma\hat{\sigma}_t^2 \leq 0 \tag{8}$$

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 through the increased discount rate driven by an earlier adoption time. The relative size of this negative effect across the two model cases is entirely determined by $\frac{\partial\mathbb{E}_t[t^{**}]}{d\hat{\psi}_t}$. In the endogenous case, this derivative is governed by the local slope of the adoption boundary $\bar{\psi}(t)$ given in Equation 2.

To understand why, suppose that $\bar{\psi}(t)$ is locally differentiable at time t. Linearizing the boundary around t,

$$\bar{\psi}(t+s) \approx \bar{\psi}(t) + \bar{\psi}'(t)\,s,$$

and using that $\hat{\psi}_t$ is locally a martingale, the expected remaining time to adoption admits the first–order approximation

$$\mathbb{E}_t[t^{**}] - t \;\approx\; \frac{\bar{\psi}(t) - \hat{\psi}_t}{-\bar{\psi}'(t)}, \qquad \bar{\psi}'(t) < 0.$$

Differentiating with respect to the current belief then yields

$$\frac{\partial \mathbb{E}_t[t^{**}]}{\partial \hat{\psi}_t} \;\approx\; \frac{1}{\bar{\psi}'(t)}.$$

Thus, the sensitivity of expected adoption time to changes in $\hat{\psi}_t$ is inversely proportional to the local time slope of the adoption threshold. A flat or nearly flat threshold implies a large (in magnitude) response of $\mathbb{E}_t[t^{**}]$ to belief innovations, while a downward-sloping threshold bounds this sensitivity and hence the discount-rate effect.

Note also that this approximation is effectively an upper bound on the sensitivity of the expected adoption time to shocks to productivity shocks. In regions far from the boundary, the expected adoption time will be nearly constant at T, and in those regions, the derivative will be near zero. As you rise into regions where adoption is possible, even with very low probabilities, the discount rate effect will smoothly increase, and spread across the revolution without generating a sudden decrease in prices.

This behavior stands in sharp contrast to the exogenous case. The simplifying assumption of the exogenous stopping time effectively evaluates $\hat{\psi}_t$ against a fixed threshold $\bar{\psi}$ when close to the fixed time $\bar{t^{**}}$. When adoption is evaluated against a fixed hurdle, the expected adoption time exhibits a highly nonlinear dependence on beliefs: Far from the threshold, small innovations to $\hat{\psi}_t$ have essentially no effect on $\mathbb{E}_t[t^{**}]$, as adoption is unlikely to occur within the relevant horizon. However, as $\hat{\psi}_t$ approaches the hurdle, the sensitivity of $\mathbb{E}_t[t^{**}]$ to belief shocks rises sharply. In the limit, there is an arbitrarily large response to infinitesimal belief innovations near the threshold just prior to $\bar{t}^{**}$.

This pattern, where discount rates are insensitive to productivity far from the threshold, when adoption is unlikely, and then extremely sensitive near the boundary, is precisely what generates the bubble pattern in the exogenous case.

In the endogenous case, the downward sloping adoption boundary comes from the structure of Equation 2, where there are three terms. The latter two terms tend to generate a downward slope as both the continuation option and the variance penalty are decreasing in time ($\frac{\partial \chi(t)}{dt} \leq 0$ and $\frac{\partial}{dt}\left(\frac{1}{2}(\gamma-1)A_2(\tau)\hat{\sigma}_t^2\right) \leq 0$). This is both due to the terminal period nearing and $\hat{\sigma}_t^2$ falling with the passage of time.

The first term generates an upward slope in the threshold with $\frac{\partial}{dt}\left(-\frac{1-\kappa}{A_2(\tau)}\right) \geq 0$. This is due to the "use it or lose it" nature of adoption due to the destruction of capital when $\kappa > 0$. You must adopt early enough to reap the benefits of increased production before the terminal period to offset the high up-front cost. However, note that this is an artifact of the finite-horizon setup with terminal wealth, as it operates through the $A_2(\tau)$ term. In contrast, the two channels through which falling $\hat{\sigma}_t^2$ generates the negative slope, 1) the reduction of uncertainty associated with adoption, and 2) the reduction of the value of the continuation option, are fundamental to the learning aspect of the model.[7]

Since learning induces a downward-sloping adoption boundary through the gradual reduction of uncertainty, the endogenous model spreads the discount rate effect smoothly over time and eliminates the bubble pattern present in the exogenous case.

Figure 2 visualizes how this effect plays out across the two model cases in the baseline calibration of the model in PV. The four plots of Figure 2 each correspond to a different time t in the baseline calibration presented in PV.

In each plot, old economy productivity is at its unconditional mean. The x-axis is the difference between the productivity of the new economy at and the mean productivity of the old economy at time t. Moving right on the x-axis increases new economy productivity, which directly impacts the M/B ratio and also determine the subjective $\hat{\psi}_t$.

The gray band in the plots show the range between $log\left(\frac{M_t^{N,IA}}{B_t^N}\right)$ on the bottom and $log\left(\frac{M_t^{N,NA}}{B_t^N}\right)$ on the top, which the true M/B for both specifications will lie between. The

[7] We verify that the endogenous case with a high κ can generate regions in which rising productivity generates falling prices.

true M/B for the endogenous adoption time specification is given by the blue line and the exogenous specification by the orange line. In the smaller plots beneath, the expected time to adoption is given for both cases.[8] Finally, the dashed lines in the main plot show the approximation of true M/B given by Equation 6, which works very well for both cases. The orange vertical line on the two plots shows the median level of productivity at time t conditional on an adoption occurring in the exogenous specification.

In Panel A, it is early in the revolution at $t = 1.5$, and for most revolutions, $\rho_{N,t}$ is only slightly positive due to the short history. In both specifications, the expected adoption time is close to the terminal period $T = 30$, so small shocks to $\rho_{N,t}$ move M/B along the upper boundary of the gray band, with only the cash-flow effect operating. At this stage, expected adoption time responds smoothly to changes in productivity in both cases.

Further along in the revolution at $t = 5$, shown in Panel B, typical productivity conditional on a revolution lies in a region where, in the exogenous case, $\mathbb{E}_t[t^{**}]$ responds sharply to productivity changes. Increases to productivity, and the corresponding increase to subjective beliefs $\hat{\psi}_t$, near the adoption threshold substantially increase the probability of adopting soon (in roughly three years) rather than never adopting (effectively in 25 years at $T = 30$). As a result, rising productivity lowers $\mathbb{E}_t[t^{**}]$ rapidly enough that the discount-rate effect dominates, and returns respond negatively to positive cash-flow shocks. This effect intensifies in the final panels, when $t = 7$ and $t = 7.5$. As the exogenous threshold approaches, the positive productivity shocks required for adoption cause stock prices fall sharply, and the bubble bursts.

In contrast, for the endogenous case (blue lines), $\mathbb{E}_t[t^{**}]$ responds smoothly to productivity shocks throughout the revolution. Expected adoption times move steadily closer with successive positive shocks, spreading the discount-rate effect over time. Without the exogo-

[8] Expected adoption times are computed by approximating posterior beliefs as a Gaussian diffusion with deterministically declining variance due to learning, and integrating the implied first-passage probabilities to the PDE-implied adoption boundary (with non-adoption mapped to the terminal horizon). The true M/B ratios for the new economy are computed by numerically solving PDEs in a manner similar to the replication code provided by PV.

neous adoption date to force a concentrated discount-rate shock, a bubble pattern does not arise.

To show that this effect is more general across reasonable parameters. Figure 3 shows the cumulative unexpected returns for parameter permutations of the most important discount rate parameters $(\sigma_{N,1}, \gamma, \kappa, \phi)$. This is analogous to the unexpected M/B paths for the revolutions in the bottom two panels of 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.[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.

While we argue above that the lack of a bubble in the endogenous adoption time case is a general feature of learning models, it may be the case that a bubble conditional on adoption could be obtained through other specifications. We therefore also present an empirical argument regarding the mechanism of PV.

4 Earnings and stock prices during the DotCom Bubble and AI Boom

We now turn to an empirical test of the mechanism in PV. The central prediction of the model is not simply that stock prices rise and then fall during technological revolutions, but that at some point positive cash-flow news becomes negative stock-return news. Intuitively, sufficiently strong cash-flow realizations accelerate the expected adoption of the new technology, raising risk and discount rates enough to offset higher expected profitability. As a result, the model predicts that prices fall even as earnings expectations improve: positive

[9] For parameterizations with vary 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.

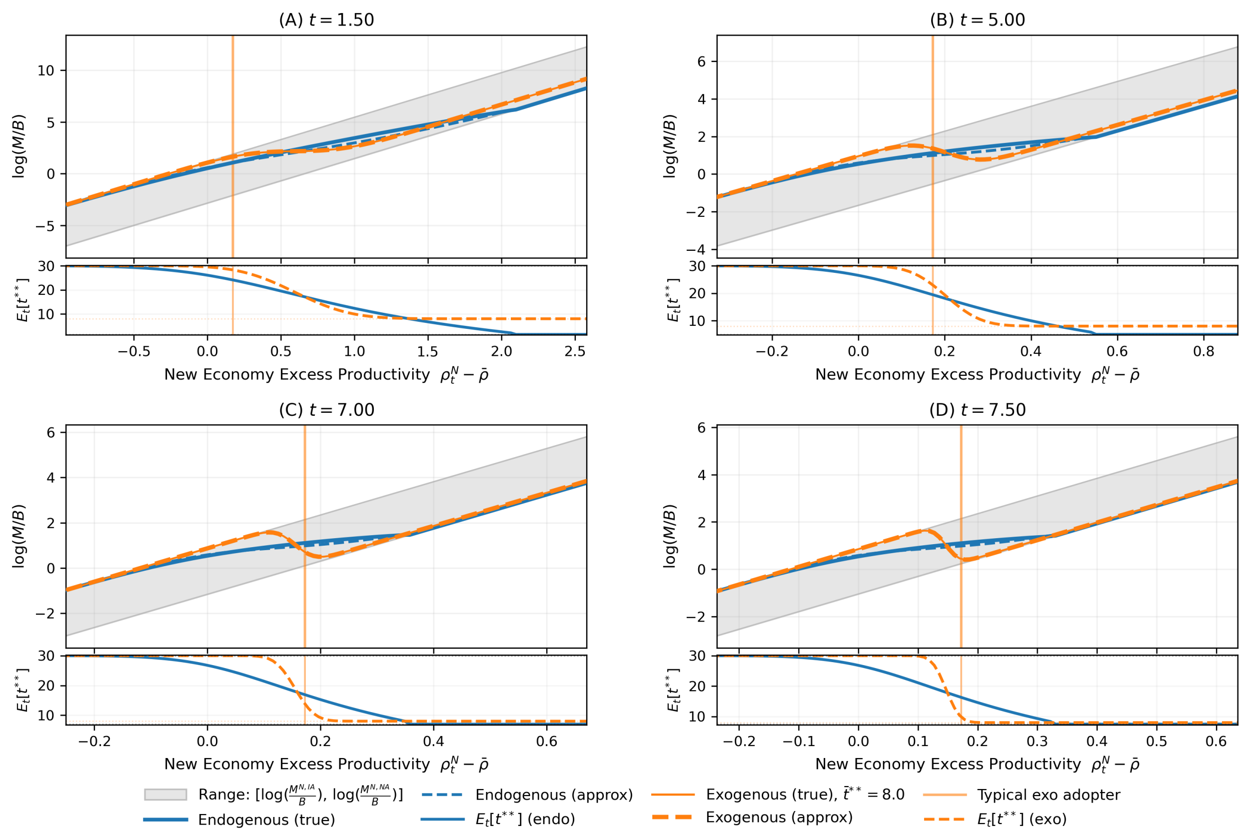

Figure 2: **Expected adoption times and new economy stock market prices.** Each panel shows new economy M/B (larger plots) and expected time to adoption $E[t^{**}]$ (smaller plots underneath) as a function of the productivity of the new economy relative to the old economy at a given time t in the model. Old economy productivity is assumed to be at its unconditional mean. The gray band shows the range between two extremes of M/B, where the upper boundary of the band is $log\left(\frac{M_t^{N,NA}}{B_t^N}\right)$, or the M/B ratio of the new economy if it can never be adopted (even optimally), and the lower boundary is $log\left(\frac{M_t^{N,IA}}{B_t^N}\right)$ is the M/B ratio if it were immediately adopted (even sub-optimally). The blue lines show the true (solid) and approximate (dashed) M/B ratio, and the orange lines show the true and approximate M/B when adoption is an all-or-nothing decision that occurs at the exogenous time $E[\bar{t}^{**}] = 8$. The approximation is $\log\left(\frac{M_t^N}{B_t^N}\right) \approx log\left(\frac{M_t^{N,IA}}{B_t^N}\right) + \left(\frac{\mathbb{E}_t[t^{**}]-t}{T-t}\right) A_2(\tau)^2\gamma\hat{\sigma}_t^2$. The vertical orange line shows the median relative productivity at a given time conditional on an adoption in the exogenous case.

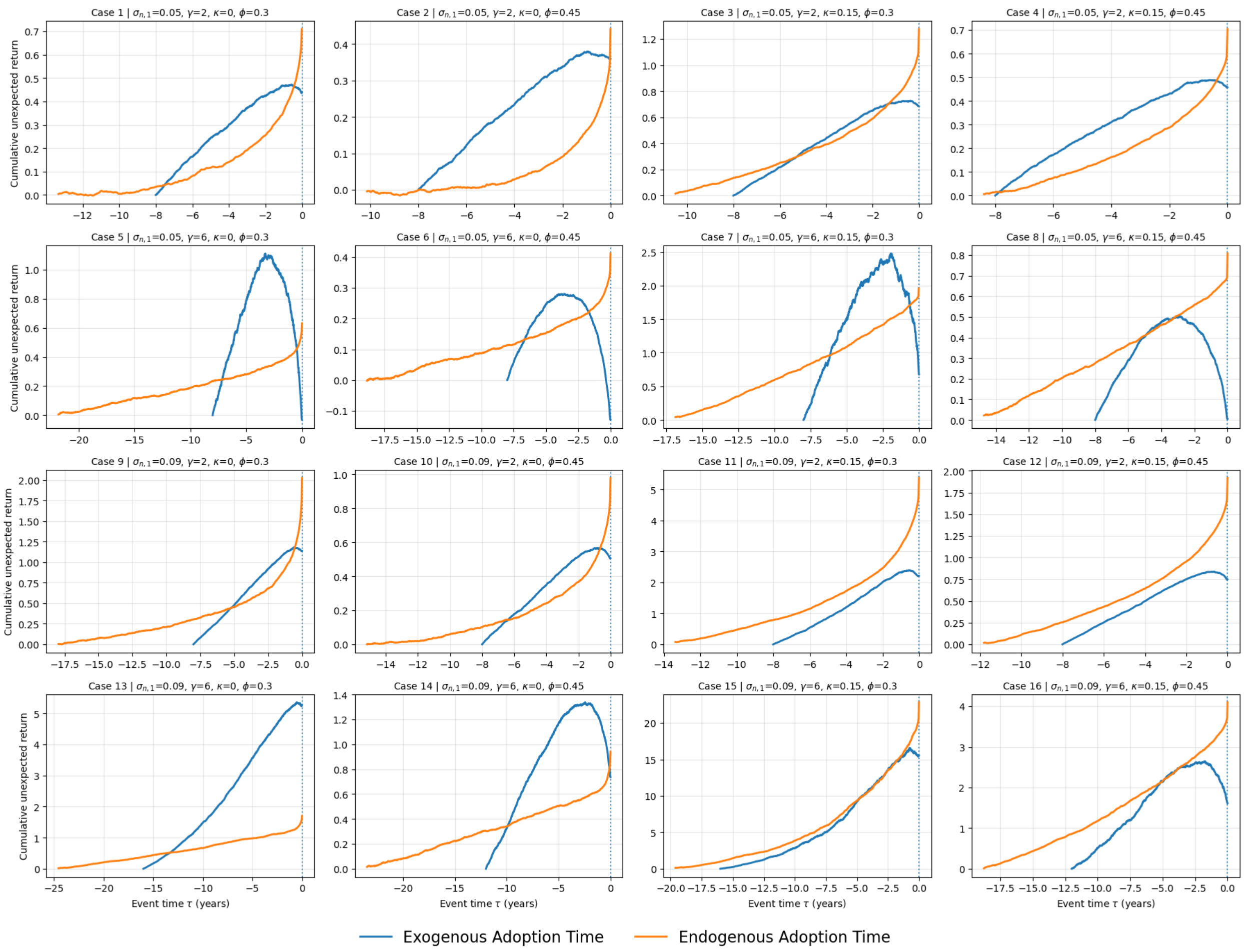

Figure 3: **Cumulative unexpected returns in revolutions across parameterizations** Each panel shows new economy cumulative unexpected returns conditional on a technical revolution in the endogenous and exogenous specification for a given set of parameters. Exogenous adoptions occur at time $t^{**} = 8$. For the endogenous adoption time specification we show adoptions in the two year band around the median adoption occurrence. The 16 panels represent the 16 possible combinations of parameters ($\kappa \in (0, 0.15), \gamma \in (2, 6), \sigma_{N,1} \in (0.5, 0.9), \phi \in (0.3, 0.45)$), 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.

cash-flow shocks themselves pop the bubble.

To evaluate this implication, Figure 4 compares the joint evolution of earnings measures and stock prices during the DotCom Bubble of the late 1990s and the recent AI boom beginning in late 2022.

For the DotCom episode, we follow PV and compare the NASDAQ and the Dow Jones Industrial Average (DJIA), focusing on the top 100 NASDAQ firms by market capitalization at the start of the period and the 30 DJIA constituents.[10] For the AI boom, we compare the DJIA to firms in the VanEck Semiconductor ETF (SMH), which we use as a proxy for exposure to the build-out of AI infrastructure. Stock price indices for the NASDAQ, DJIA, and SMH are plotted alongside earnings measures and normalized to 100.

To proxy for productivity and expected cash flows in the model, we use two measures. The first is forecast earnings scaled by book equity (BE). As PV note, this measure of return-on-equity (ROE) corresponds directly to productivity in the model, and ROE is also used as a proxy for cash-flow shocks in the broader literature (e.g., Vuolteenaho (2002)). To proxy for the growth parameter $\hat{\psi}_t$, we use forecast earnings growth.

Both measures are constructed using IBES consensus mean earnings-per-share forecasts for the next two fiscal years (FY1 and FY2). Aggregate forecast earnings levels are computed as a weighted average of FY1 and FY2 forecasts, scaled by shares outstanding and total book equity. Forecast earnings growth is defined as FY2/FY1 – 1 and aggregated across firms using FY1 earnings weights. All series are normalized to 100 at the start of each sample window.

The DotCom Bubble is shown in the top row. Stock prices exhibit the familiar boom–bust pattern, with NASDAQ prices rising sharply through 1999 and early 2000 before collapsing. The left panel show that this collapse was accompanied by declining earnings expectations, rather than rising cash-flow prospects, a pattern inconsistent with the discount-rate mechanism emphasized in PV.[11] The right panel shows little evidence that either the boom or the

[10]We use the top 100 firms to create a stable panel; the results are qualitatively unchanged using the full NASDAQ universe.

[11]We verify that realized aggregate earnings closely track one-year-ahead IBES forecasts. This pattern is also highlighted in recent work by Gómez-Cram and Lawrence (2025) who shows that realized earnings of

collapse was driven by changing earnings growth expectations.

The bottom row shows the recent AI boom. Beginning in 2023, AI stock prices rise alongside expected earnings and, unlike the DotCom episode, alongside rising growth expectations. Earnings growth expectations in particular closely track AI stock prices.

Given that adoption of AI appears likely but still highly uncertain, this positive association between cash flows and prices is at odds with the mechanism in PV, as this is precisely the period when the discount-rate channel should be most active.[12]

Taken together, neither episode displays the negative relation between stock prices and cash-flow expectations that is central to the bubble mechanism in the exogenous adoption model of PV.

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 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.

This result has 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

software firms relative to expectations were positive and then negative during the rise and fall of the Nasdaq.

[12] For instance, 90% of S&P 500 firms mentioned AI in their 2024 10-K filings, compared to only 25% in 2023.

time that the discount-rate channel should be most active according to the exogenous model, with cash-flow shocks and prices moving in opposite directions. However, we find empirically that beliefs about future cash-flows, as proxied by analyst forecasts of earnings, appear to be comoving positively with prices of AI stocks.

This pattern is consistent with each bit of news about the usefulness of AI producing incremental changes in the discount rate that attenuate, but do not overwhelm, the price impacts from cash-flow shocks, as the exogenous adoption time model predicts.Good cash-flow news is good stock price news for AI firms. If the bubble does in fact burst, it will be caused by some other mechanism than that proposed by Pástor and Veronesi (2009).

References

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

Garleanu, Nicolae, Stavros Panageas, and Jianfeng Yu, "Technological growth and asset pricing," *The 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.

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.

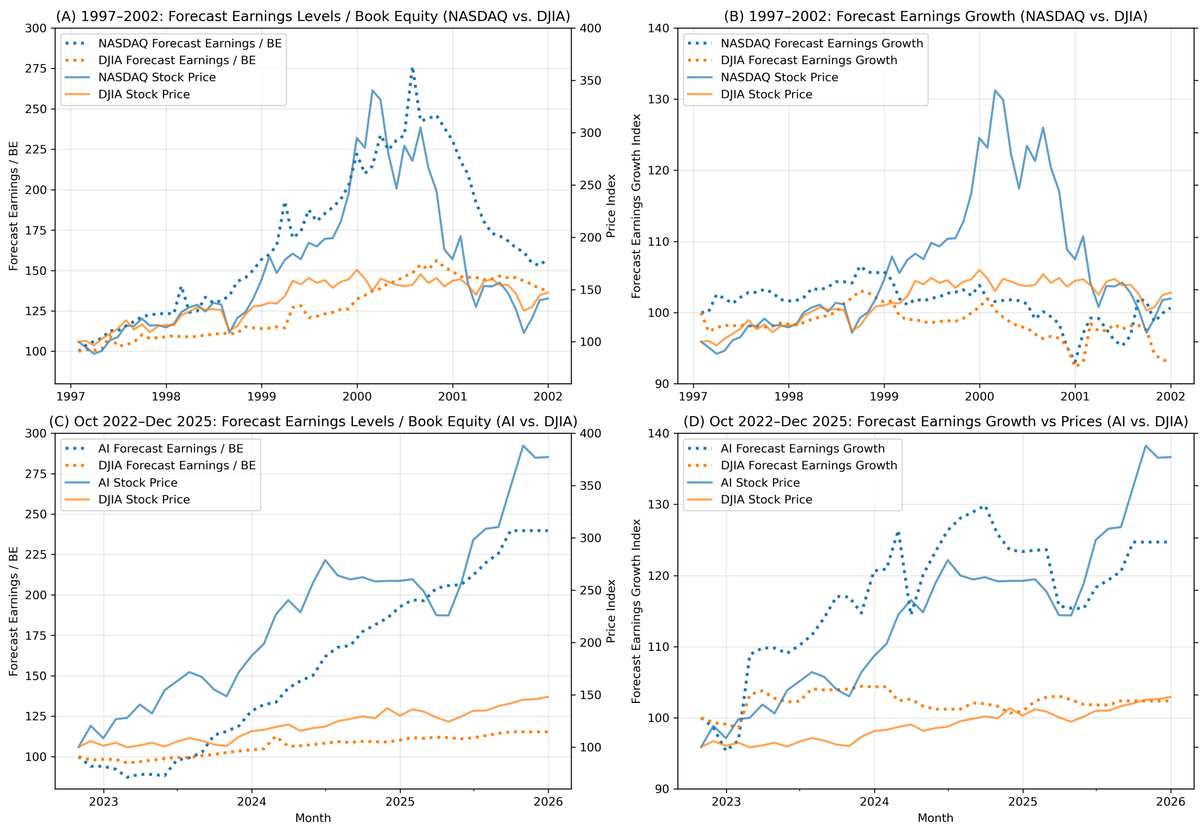

Figure 4: **Earnings and stock prices during the DotCom Bubble and AI Boom.** The top row shows the DotCom Bubble (1997–2002) and the bottom row shows the AI boom (October 2022–December 2025), where "AI" is proxied by the constituents of the VanEck Semiconductor ETF (SMH). In each row the left column plots one year forecast earnings / book equity built from a weighted average of IBES FY1 and FY2 consensus EPS forecasts; both are converted to dollars using shares outstanding, summed across firms in the universe, and scaled by total book equity. The right column plots a forecast earnings growth index, defined as FY2/FY1–1 aggregated using FY1 earnings dollars as weights. Each panel plots the relevant stock price indices (NASDAQ or SMH vs. DJIA), and all series are normalized to 100 at the start of each window.

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

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.

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