

NBER WORKING PAPER SERIES

TECHNOLOGICAL REVOLUTIONS AND STOCK PRICES

Lubos Pastor
Pietro Veronesi

Working Paper 11876
<http://www.nber.org/papers/w11876>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2005
Revised February 2008

We thank Sreedhar Bharath, Markus Brunnermeier, Leonid Kogan, Lars Lochstoer, Lu Zhang, and especially Judy Chevalier and two anonymous referees for very helpful comments. We also thank Malcolm Baker, Robert Barro, Efraim Benmelech, Gene Fama, Bob Fogel, Boyan Jovanovic, John Heaton, Ali Lazrak, Robert Novy-Marx, Rob Stambaugh, Dmitriy Stolyarov, and the audiences at the 2008 AEA meeting, 2007 AFA meeting, 2006 EFA meeting, 2006 UBC Summer Conference, Fall 2005 NBER Asset Pricing meeting, CERGE-EI Prague, Dartmouth College, Ente Einaudi, Harvard University, Indiana University, London Business School, London School of Economics, New York University, Stockholm Institute for Financial Research, Stockholm School of Economics, University of Chicago, University of Pennsylvania, and University of Vienna. Shastri Sandy has provided valuable research assistance. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2005 by Lubos Pastor and Pietro Veronesi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Technological Revolutions and Stock Prices
Lubos Pastor and Pietro Veronesi
NBER Working Paper No. 11876
December 2005, Revised February 2008
JEL No. G1

ABSTRACT

We develop a general equilibrium model in which stock prices of innovative firms exhibit "bubbles" during technological revolutions. In the model, the average productivity of a new technology is uncertain and subject to learning. During technological revolutions, the nature of this uncertainty changes from idiosyncratic to systematic. The resulting "bubbles" in stock prices are observable ex post but unpredictable ex ante, and they are most pronounced for technologies characterized by high uncertainty and fast adoption. We find empirical support for the model's predictions in 1830-1861 and 1992-2005 when the railroad and Internet technologies spread in the United States.

Lubos Pastor
Graduate School of Business
University of Chicago
5807 South Woodlawn Ave
Chicago, IL 60637
and NBER
lubos.pastor@chicagogsb.edu

Pietro Veronesi
Graduate School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
pietro.veronesi@gsb.uchicago.edu

1. Introduction

Technological revolutions tend to be accompanied by bubble-like patterns in the stock prices of firms that employ the new technology. After an initial surge, stock prices of innovative firms usually fall in the presence of high volatility. Recent examples of such price patterns include the “Internet craze” of the late 1990s, the “biotech revolution” of the early 1980s, and the “tronics boom” of the early 1960s, as characterized by Malkiel (1999).¹ Other examples include the 1920s and the turn of the 20th century; in both periods, technological innovation spread rapidly while the stock market boomed and then faltered (e.g., Shiller, 2000).²

The bubble-like stock price behavior during technological revolutions is frequently attributed to market irrationality (e.g., Shiller, 2000, Perez, 2002). We propose another possible explanation that does not involve irrationality. We argue that new technologies are characterized by high uncertainty about their future productivity, and that the time-varying nature of this uncertainty can also produce the observed stock price patterns.

We build a general equilibrium model of a finite-horizon representative-agent economy with two sectors: the “new economy” and the “old economy.” The old economy implements the existing technologies in large-scale production whose output determines the representative agent’s wealth. The new economy, which is created when a new technology is invented, implements the new technology in small-scale production that does not affect the agent’s wealth. It is optimal for the new technology to be initially employed on a small scale because its future productivity is uncertain. By observing the new economy, the representative agent learns about the average productivity of the new technology before deciding whether to adopt the technology on a large scale. We show that this irreversible adoption takes place if the agent learns that the new technology is sufficiently productive. We define a technological revolution as a period concluded by a large-scale adoption of a new technology.

We show that the nature of the risk associated with new technologies changes over time. Initially, this risk is mostly idiosyncratic due to the small scale of production and a low probability of a large-scale adoption. The risk remains idiosyncratic for those technologies

¹According to Malkiel (1999), “What electronics was to the 1960s, biotechnology became to the 1980s... Valuation levels of biotechnology stocks reached levels previously unknown to investors... From the mid-1980s to the late 1980s, most biotechnology stocks lost three-quarters of their market value.”

²“Every previous technological revolution has created a speculative bubble... With each wave of technology, share prices soared and later fell... The inventions of the late 19th century drove p-e ratios to a peak in 1901, the year of the first transatlantic radio transmission. By 1920 shares prices had dropped by 70% in real terms. The roaring twenties were also seen as a “new era”: share prices soared as electricity boosted efficiency and car ownership spread. After peaking in 1929, real share prices tumbled by 80% over the next three years.” (The Economist, September 21, 2000, Bubble.com)

that are never adopted on a large scale. For the technologies that are ultimately adopted, however, the risk must gradually change from idiosyncratic to systematic. As the probability of adoption increases, the new technology becomes more likely to affect the old economy and with it the representative agent's wealth, so systematic risk in the economy increases.

This time-varying nature of risk has interesting implications for stock prices. Initially, while uncertainty about the new technology is mostly idiosyncratic, the new economy stocks command high market values. As the adoption probability increases, the resulting increase in systematic risk pushes up the discount rates and thus depresses stock prices in both the new and old economies. The new economy stock prices fall deeper because their discount rates rise higher due to an increase in the new economy's market beta.

Stock prices are affected not only by discount rates but also by expected cash flows. The technologies that are ultimately adopted must turn out to be sufficiently productive before the adoption. This positive cash flow news pushes stock prices up, countervailing the effect of the higher discount rate. The cash flow effect prevails initially, pushing the new economy stock prices up, but the discount rate effect prevails eventually, pushing the stock prices down. The resulting pattern in the new economy stock prices looks like a bubble but it obtains under rational expectations through a general equilibrium effect.

The bubble-like pattern in stock prices arises in part due to an ex post selection bias. Researchers study technological revolutions with the ex post knowledge that the revolutions took place, but investors living through those periods did not know whether the new technologies would eventually be adopted on a large scale. The representative agent in our model never expects stock prices to fall; she always expects to earn positive stock returns commensurate to the stocks' riskiness, and she subsequently earns those fair returns, on average. However, in those rare periods that are recognized as technological revolutions ex post, the agent's realized returns tend to be initially positive due to good news about productivity and eventually negative due to bad news about systematic risk.

Uncertainty about new technologies affects not only the level but also the volatility of stock prices. Due to this uncertainty, the new economy stocks are more volatile than the old economy stocks. After an initial decline, the new economy's volatility rises sharply when the stochastic discount factor becomes more volatile as a result of a higher probability of a large-scale adoption. The same effect also pushes up the new economy's market beta and the old economy's volatility, two different aspects of systematic risk in the economy.

Our model makes many empirical predictions for technological revolutions: The "bubble"

in stock prices should be much stronger in the new economy than in the old economy; stock prices in both economies should reach the bottom at the end of the revolution; the new economy's market beta should increase sharply before the end of the revolution; the new economy's volatility should also rise sharply and it should exceed the old economy's volatility; the old economy's volatility should rise but less than the new economy's volatility; the new economy's beta and both volatilities should all peak at the end of the revolution; and the old economy's productivity should begin rising at the end of the revolution.

All of these predictions are supported by empirical evidence from the recent Internet revolution. According to the model, this revolution ended (i.e., the probability of a large-scale adoption of the Internet technology reached one) in 2002. The “bubble” pattern was much stronger in the NASDAQ index (our proxy for the new economy) than in the NYSE/AMEX index (the old economy); both stock price indexes reached the bottom in 2002; NASDAQ's beta doubled between 1997 and 2002; NYSE/AMEX's return volatility also doubled and NASDAQ's volatility tripled over the same period; NASDAQ's volatility always exceeded NYSE/AMEX's volatility; NASDAQ's beta and both volatilities peaked in 2002; and the productivity growth of the U.S. economy accelerated sharply after 2002.

We also examine stock prices during the first major technological revolution in the U.S. since the opening of the U.S. stock market – the introduction of steam-powered railroads. In the 1830s and 40s, there was substantial uncertainty about whether the railroad technology would be adopted on a large scale. We analyze stock prices before the Civil War, and find that they fell before and during year 1857, with railroad stocks falling more than non-railroad stocks. The railroad stock volatility and price-dividend ratios consistently exceeded their non-railroad counterparts. The volatility of all stocks rose in 1857. The railroad stock beta increased sharply in the 1850s, before falling right after 1857. In the context of our model, all of this evidence is consistent with a large-scale adoption of the railroad technology around 1857, soon after railroads began expanding west of the Mississippi River.

Much of the literature on technological innovation analyzes issues different from those addressed here. Unlike Romer (1990), Aghion and Howitt (1992), and others, we take technological inventions to be exogenous. We do not examine the links between technological revolutions and human capital (e.g., Chari and Hopenhayn, 1991, Caselli, 1999, Manuelli, 2003). Different but related models of learning are presented in Jovanovic (1982), Jovanovic and Nyarko (1996), and Atkeson and Kehoe (2007). We empirically examine the Internet and railroad revolutions, while other technological revolutions are examined by Jovanovic and Rousseau (2003, 2005), Mazzucato (2002), and others. Mokyr (1990) argues that tech-

nological progress is discontinuous, as assumed in our model, and that occasional seminal inventions (“macroinventions”) are the key sources of economic growth.

A small but growing literature explores the links between technological innovation and stock prices (e.g., Jovanovic and MacDonald, 1994, and Laitner and Stolyarov, 2003, 2004a,b). According to Greenwood and Jovanovic (1999) and Hobijn and Jovanovic (2001), innovation causes the stock market to drop because the incumbent firms are unable or unwilling to implement the new technology. Similar initial stock market drops are obtained in the models of Laitner and Stolyarov (2003) and Manuelli (2003). In our model, the stock market value of the old economy also drops after the new technology is invented, mostly because of the costs and risks associated with a large-scale adoption of the new technology, but our focus is on the subsequent bubble-like stock price pattern in the new economy.

The paper is organized as follows. Section 2 presents the model. Section 3 solves for stock prices and analyzes their dynamics. Section 4 investigates the model’s empirical predictions for stock prices during technological revolutions. Section 5 empirically examines the behavior of stock prices in 1830–1861 and 1992–2005 when the railroad technology and the Internet technologies, respectively, spread in the United States. Section 6 concludes.

2. The Economy

We consider an economy with a finite horizon $[0, T]$. A representative agent has preferences defined by power utility over terminal wealth W_T , with risk aversion $\gamma > 1$:

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

At time $t = 0$, the agent is endowed with capital B_0 . Subsequently, capital is invested in a linear technology producing output (net of depreciation) at the rate of $Y_t = \rho_t B_t$. Since there is no intermediate consumption, all output is reinvested, and capital B_t follows

$$dB_t = Y_t dt = \rho_t B_t dt. \quad (2)$$

Productivity ρ_t follows a mean-reverting process whose mean is determined by the available technology. There are two technologies: “old” and “new.” Initially, only the old technology is available, and the long-run mean of ρ_t is equal to $\bar{\rho}$. At time t^* , the new technology becomes available. If the representative agent adopts the new technology at time $t^{**} \geq t^*$, the long-run mean of ρ_t changes from $\bar{\rho}$ to $\bar{\rho} + \psi$. Thus, the dynamics of ρ_t are given by

$$d\rho_t = \phi(\bar{\rho} - \rho_t) dt + \sigma dZ_{0,t}, \quad 0 < t < t^{**} \quad (3)$$

$$d\rho_t = \phi(\bar{\rho} + \psi - \rho_t) dt + \sigma dZ_{0,t}, \quad t^{**} \leq t < T, \quad (4)$$

where ϕ is the speed of mean reversion, $\bar{\rho}$ is the mean productivity of the old technology, ψ is the “productivity gain” brought by the new technology, and σ^2 is the variance of productivity shocks, represented by the Brownian increments $dZ_{0,t}$. That is, the adoption of the new technology is equivalent to a shift in the economy’s average productivity.

The representative agent chooses whether and when to adopt the new technology to maximize utility in equation (1) under the market-clearing condition $W_T = B_T$. In equilibrium, the agent’s final wealth must equal the amount of capital accumulated by time T .

Our key assumption is that the productivity gain ψ is unobservable. When the new technology appears at time t^* , ψ is drawn from a normal distribution with known variance:

$$\psi \sim N(0, \hat{\sigma}_{t^*}^2). \quad (5)$$

All other parameters are known. The adoption of the new technology is irreversible and costly. Converting capital to the new technology incurs a proportional conversion cost $\kappa \geq 0$.

The agent has three choices at time t^* when the new technology becomes available:

- (i) Adopt the new technology
- (ii) Begin learning about the new technology (i.e., about ψ)
- (iii) Discard the new technology

We show below that the agent optimally chooses option (ii), so he begins learning at time t^* . The agent learns about ψ by “experimenting” with the new technology – i.e., by implementing it on a small scale. After time t^* , the economy consists of two sectors: the small-scale “new economy,” which employs the new technology, and the large-scale “old economy,” whose productivity ρ_t follows equation (3). The capital B_t^N used in the new economy is infinitely smaller than B_t , so the agent’s wealth W_T is affected by the new technology only if this technology is adopted on a large scale (i.e., by the old economy). Denoting the new economy’s productivity by ρ_t^N , the processes of B_t^N and ρ_t^N for $t > t^*$ are given by

$$dB_t^N = \rho_t^N B_t^N dt \quad (6)$$

$$d\rho_t^N = \phi(\bar{\rho} + \psi - \rho_t^N) dt + \sigma_{N,0} dZ_{0,t} + \sigma_{N,1} dZ_{1,t}, \quad (7)$$

where $Z_{1,t}$ is a Brownian motion uncorrelated with $Z_{0,t}$. The agent learns about ψ by observing ρ_t^N and ρ_t . The learning process is characterized by Lemma A1 in the Appendix. The posterior distribution of ψ conditional on $\mathcal{F}_t = \{(\rho_\tau^N, \rho_\tau) : t^* \leq \tau \leq t\}$ is normal,

$$\psi | \mathcal{F}_t \sim N(\hat{\psi}_t, \hat{\sigma}_t^2),$$

where the posterior mean $\widehat{\psi}_t$ is a martingale (see equation (17)) and the posterior variance $\widehat{\sigma}_t^2$ declines deterministically over time due to learning (see equation (18)). If the new technology is adopted at time t^{**} , the agent continues to learn about ψ by observing ρ_t^N and ρ_t , but the old economy's productivity follows equation (4) rather than equation (3).

We define a *technological revolution* as the period $[t^*, t^{**}]$ concluded by a large-scale adoption of the new technology. We treat the invention of the new technology as given, and study the conditions under which the invention leads to a technological revolution.

2.1. Optimal Adoption of the New Technology

The agent can adopt the new technology anytime between times t^* and T (or never). We solve for the optimal adoption time t^{**} numerically in Section 4.2. Until then, we focus on a simpler problem in which t^{**} denotes an exogenously given time at which the agent decides whether or not to adopt the new technology. This simpler problem admits a closed-form solution for stock prices, which improves our understanding of the stock price dynamics. Our numerical results in Section 4.2. show that the dynamics obtained when t^{**} is endogenously chosen are very similar to those obtained here with an exogenous t^{**} .

The sequence of events in the model is summarized in Figure 1. We assume that if a new technology is not adopted at time t^{**} , it continues to operate on a small scale until time T . Our history is full of examples of technologies that have not been adopted on a large scale but still survive on a small scale (e.g., direct-current electric motors, airships, etc.)

Proposition 1: It is never optimal to adopt the new technology immediately at time t^* .

Adopting the new technology is risky – it may increase or decrease average productivity, depending on the sign of ψ . The prior for ψ in equation (5) is centered at zero, making the increases and decreases in productivity equally likely as of time t^* .³ Since the agent is risk averse, immediate adoption of the new technology is suboptimal. This intuition is formalized in the Appendix, which shows that the adoption of the new technology at time t^* yields lower expected utility than no adoption. Proposition 1 holds for any κ , including $\kappa = 0$, as it is driven by the increase in risk resulting from the adoption of the new technology.

Proposition 2: The new technology is adopted at time t^{**} if and only if

$$\widehat{\psi}_{t^{**}} \geq \underline{\psi} = -\frac{\log(1 - \kappa)}{A_2(\tau^{**})} + \frac{1}{2}(\gamma - 1) A_2(\tau^{**}) \widehat{\sigma}_{t^{**}}^2, \quad (8)$$

³If the prior is centered at $\widehat{\psi}_{t^*} \neq 0$, Proposition 1 is modified so that it is not optimal to adopt the new technology at time t^* unless $\widehat{\psi}_{t^*}$ is sufficiently high. See Proposition 2 for an analogous relation.

where $\tau^{**} = T - t^{**}$ and $A_2(\tau) = \tau - (1 - \exp(-\phi\tau)) / \phi > 0$.

The new technology is adopted at time t^{**} if the expected productivity gain, $\widehat{\psi}_{t^{**}}$, is positive and sufficiently large. The threshold $\underline{\psi}$ is always positive, and it increases in the conversion cost κ , uncertainty $\widehat{\sigma}_{t^{**}}$, and risk aversion γ , which is intuitive. Note that the agent makes the adoption decision without knowing the true value of ψ . Regardless of the outcome of the adoption decision, learning about ψ continues after time t^{**} .

Proposition 3: It is optimal to begin experimenting with the new technology at time t^* .

This proposition, proved in the Appendix, shows that the agent chooses to set up the new economy to begin learning about the new technology immediately after this technology becomes available at time t^* . The intuition is simple. Experimenting allows the agent to learn about the productivity gain ψ . If this learning leads the agent to believe at time t^{**} that ψ is sufficiently high, then it becomes optimal to adopt the new technology (Proposition 2). Otherwise, the status quo will prevail. Since experimenting is costless and there is no downside to it, it gives the agent a valuable option for free.⁴

Since option value generally increases with uncertainty, high uncertainty $\widehat{\sigma}_{t^*}$ makes a new technology desirable for experimentation. If it were costly to experiment with new technologies, or if the agent had to choose from a subset of technologies at time t^* , then the technologies with the highest $\widehat{\sigma}_{t^*}$ would be selected for experimentation, *ceteris paribus*. Uncertainty about productivity gains is thus a natural feature of innovative technologies.

3. Stock Prices

The stocks of the old and new economies pay liquidating dividends B_T and B_T^N , respectively, at time T . There is also a riskless bond in zero net supply, whose yield we normalize to zero, for simplicity. Since the two shocks in the model are spanned by the two stocks, markets are complete. Standard arguments then imply that the state price density is uniquely given by

$$\pi_t = \frac{1}{\lambda} E_t [W_T^{-\gamma}], \quad (9)$$

where λ is the Lagrange multiplier from the utility maximization problem of the representative agent. The market values (shadow prices) of the old and new economy stocks, denoted

⁴The problem we solve resembles the problem of making an irreversible marriage decision. It is generally suboptimal to marry a new acquaintance immediately because of substantial uncertainty regarding the quality of the personality match (cf. Proposition 1). Instead, it seems advisable to first develop the relationship on a small scale, by dating without any commitment (cf. Proposition 3), and then to marry if we learn that the relationship is likely to work in the long run (cf. Proposition 2).

by M_t and M_t^N , respectively, are given by the standard pricing formulas

$$M_t = E_t \left[\frac{\pi_T B_T}{\pi_t} \right] \quad \text{and} \quad M_t^N = E_t \left[\frac{\pi_T B_T^N}{\pi_t} \right]. \quad (10)$$

To normalize the market values, we form “market-to-book” (M/B) ratios M_t/B_t and M_t^N/B_t^N . It seems reasonable to interpret capital as the book value of equity, and this interpretation is exact for B_t and B_t^N in equations (2) and (6) if we also interpret output and productivity as earnings and profitability, respectively (Pástor and Veronesi, 2003).

Let p_t denote the probability at time t , $t^* \leq t < t^{**}$, that the new technology will be adopted at time t^{**} . Lemma A3 in the Appendix shows that $p_t = 1 - \mathcal{N} \left(\underline{\psi}; \widehat{\psi}_t, \widehat{\sigma}_t^2 - \widehat{\sigma}_{t^{**}}^2 \right)$, where $\mathcal{N}(\cdot; a, s^2)$ denotes the c.d.f. of the normal distribution with mean a and variance s^2 .

Proposition 4: For any $t \in [t^*, t^{**})$, the state price density is given by

$$\pi_t = \lambda^{-1} B_t^{-\gamma} \left\{ (1 - p_t) \widetilde{G}_t^{no} + p_t \widetilde{G}_t^{yes} \right\}, \quad (11)$$

where \widetilde{G}_t^{no} and \widetilde{G}_t^{yes} are expectations of the marginal utility of wealth conditional on whether or not the new technology is adopted at time t^{**} . Both values are given in the Appendix.

Corollary 1. For any $t \in [t^*, t^{**})$, the dynamics of π_t are given by

$$\frac{d\pi_t}{\pi_t} = -\gamma A_1(\tau) \sigma d\widetilde{Z}_{0,t} - S_{\pi,t} \widehat{\sigma}_t^2 \frac{\phi}{\sigma_{N,1}} d\widetilde{Z}_{1,t}, \quad (12)$$

where $\tau = T - t$, $A_1(\tau) = (1 - e^{-\phi\tau})/\phi$, $S_{\pi,t}$ is given in the Appendix, and so are the orthogonalized Brownian motions $(\widetilde{Z}_{0,t}, \widetilde{Z}_{1,t})$, which capture the agent’s expectation errors.

This corollary illustrates the time-varying nature of risk during technological revolutions. When a new technology arrives at time t^* , the adoption probability p_{t^*} is generally small, which makes S_{π,t^*} small as well ($p_t = 0$ implies $S_{\pi,t} = 0$). The volatility of the stochastic discount factor in equation (12) then depends only slightly on $\widehat{\sigma}_t^2$, making uncertainty about ψ mostly idiosyncratic. During a technological revolution, the adoption probability increases, which makes $S_{\pi,t}$ larger.⁵ As a result, the volatility of the stochastic discount factor becomes more closely tied to $\widehat{\sigma}_t^2$, making uncertainty about ψ increasingly systematic.

Proposition 5: For any $t \in [t^*, t^{**})$, the market-to-book ratios are given by

$$\frac{M_t}{B_t} = \frac{(1 - p_t) G_t^{no} + p_t G_t^{yes}}{(1 - p_t) \widetilde{G}_t^{no} + p_t \widetilde{G}_t^{yes}} \quad \text{and} \quad \frac{M_t^N}{B_t^N} = \frac{(1 - p_t) K_t^{no} + p_t K_t^{yes}}{(1 - p_t) \widetilde{G}_t^{no} + p_t \widetilde{G}_t^{yes}}, \quad (13)$$

⁵In a technological revolution, p_t rises from $p_{t^*} \approx 0$ to $p_{t^{**}} = 1$, and $S_{\pi,t}$ rises from $S_{\pi,t^*} \approx 0$ to $S_{\pi,t^{**}} = \gamma A_2(\tau^{**}) > 0$. That is, as p_t increases, $S_{\pi,t}$ increases from approximately zero to a positive number.

where \tilde{G}_t^{no} , \tilde{G}_t^{yes} , G_t^{no} , G_t^{yes} , K_t^{no} , and K_t^{yes} are given in the Appendix.

In the special case $p_t = 0$, the market-to-book ratio of the new economy simplifies into

$$\frac{M_t^N}{B_t^N} = e^{C_0(\tau) + A_1(\tau)\rho_t^N + A_2(\tau)\hat{\psi}_t + \frac{1}{2}A_2(\tau)^2\hat{\sigma}_t^2}, \quad (14)$$

where $A_1(\tau)$ is defined in Corollary 1, $A_2(\tau)$ in Proposition 2, and $C_0(\tau)$ is in the Appendix. Note that M^N/B^N increases when uncertainty about ψ , $\hat{\sigma}_t^2$, increases. This relation, first pointed out by Pástor and Veronesi (2003) in a simpler framework, is due to the idiosyncratic nature of uncertainty. When $p_t = 0$, the state price density does not depend on uncertainty about ψ , but when $p_t > 0$, it does. When p_t is sufficiently large, uncertainty is mostly systematic, and the associated risk reverses the positive relation between M^N/B^N and $\hat{\sigma}_t^2$.

The return processes for both stocks are given in Corollary A1 in the Appendix. Not surprisingly, the expected stock returns are given by the return covariances with $d\pi_t/\pi_t$, and the return volatilities of both stocks increase with uncertainty $\hat{\sigma}_t^2$.

3.1. The Dynamics of Prices during a Technological Revolution

In a technological revolution, the adoption probability p_t rises from a small value at time t^* to the value of one at time t^{**} . The effect of p_t on stock prices is analyzed next.

Proposition 6: The new (old) economy's M/B ratio is increasing in p_t if and only if $h_{new} > 0$ ($h_{old} > 0$), where h_{new} and h_{old} are functions of $\hat{\psi}_t$ given in the Appendix.

For plausible parameter values, $h_{new} > 0$ when $\hat{\psi}_t$ is close to zero, but $h_{new} < 0$ when $\hat{\psi}_t$ approaches the threshold $\underline{\psi}$. That is, the condition $h_{new} > 0$ holds shortly after time t^* , but it becomes violated as the adoption at time t^{**} becomes more likely. Proposition 6 then implies that the new economy's M/B is initially increasing but ultimately decreasing in p_t during a technological revolution. The condition $h_{old} > 0$ is never satisfied for the baseline parameter values, so the old economy's M/B is always decreasing in p_t .

While analyzing M/B as a function of p_t seems informative, p_t is driven primarily by $\hat{\psi}_t$. Stock prices depend on $\hat{\psi}_t$ through two opposing effects. On one hand, an increase in $\hat{\psi}_t$ is good news for prices because it increases expected cash flows ($E_t[B_T]$ and $E_t[B_T^N]$) in both economies. This *cash flow effect* is stronger for the new economy whose perceived productivity is immediately affected; the old economy's productivity is not affected until time t^{**} , if at all. On the other hand, an increase in $\hat{\psi}_t$ is bad news for prices because the higher adoption probability makes the risk embedded in the new technology increasingly

systematic, thereby raising the discount rate. This *discount rate effect* is also stronger for the new economy because π_t covaries more with ρ_t^N than with ρ_t (since both π_t and ρ_t^N correlate with revisions in $\widehat{\psi}_t$ but ρ_t does not). Moreover, the discount rate effect has a growing impact on the new economy's M/B because the dependence of π_t on revisions in $\widehat{\psi}_t$ increases as p_t increases. For the old economy, the discount rate effect generally outweighs the cash flow effect from the very beginning, leading to a gradual decline in M/B during a revolution. For the new economy, the cash flow effect tends to dominate at first, but the discount rate effect dominates in the end, producing a “bubble”.

Although the dependence of M^N/B^N on $\widehat{\psi}_t$ is complicated, its key features can be established locally at times t^* and t^{**} . We show below that M^N/B^N is increasing (decreasing) in $\widehat{\psi}$ around time t^* (t^{**}), under certain assumptions.

Proposition 7: For any $t \geq t^*$ there exists $\bar{p} > 0$ such that if $p_t < \bar{p}$ then $\frac{\partial(M_t^N/B_t^N)}{\partial\widehat{\psi}_t} > 0$.

In words, if the probability of adoption p_t is sufficiently small, then M^N/B^N is increasing in $\widehat{\psi}_t$. When p_t is close to zero, so is its sensitivity to changes in $\widehat{\psi}_t$; thus an increase in $\widehat{\psi}_t$ does not produce a large discount rate effect.⁶ The cash flow effect is large, though, because M^N/B^N in equation (14) is strongly increasing in $\widehat{\psi}_t$. Proposition 7 follows.

When a new technology arrives at time t^* , its probability of eventual adoption is typically small because only a small fraction of new technologies are adopted by the whole economy. Proposition 7 then implies that, for most new technologies, the cash flow effect initially prevails over the discount rate effect and M^N/B^N is increasing in $\widehat{\psi}_t$ shortly after time t^* .

We also have some local results at time t^{**} . Below, we compare the M/B ratio of the new economy under two scenarios: $\widehat{\psi}_{t^{**}} = \underline{\psi} \pm \varepsilon$, where $\varepsilon > 0$ is small.

Corollary 2:

(a) If $\widehat{\psi}_{t^{**}} = \underline{\psi} + \varepsilon$, then the new technology is adopted at time t^{**} , and

$$\frac{M_{t^{**}}^N}{B_{t^{**}}^N} = e^{\bar{C}_0(\tau^{**}) + A_1(\tau^{**})\rho_{t^{**}}^N + A_2(\tau^{**})\widehat{\psi}_{t^{**}} + \frac{1}{2}A_2(\tau^{**})^2(1-2\gamma)\widehat{\sigma}_{t^{**}}^2}. \quad (15)$$

(b) If $\widehat{\psi}_{t^{**}} = \underline{\psi} - \varepsilon$, then the new technology is not adopted at time t^{**} , and

$$\frac{M_{t^{**}}^N}{B_{t^{**}}^N} = e^{\bar{C}_0(\tau^{**}) + A_1(\tau^{**})\rho_{t^{**}}^N + A_2(\tau^{**})\widehat{\psi}_{t^{**}} + \frac{1}{2}A_2(\tau^{**})^2\widehat{\sigma}_{t^{**}}^2}. \quad (16)$$

⁶Analogously, if a stock option is deep out of the money, a small increase in the stock price does not change the option value by much since its delta is small and the option remains deep out of the money.

The new economy's M/B is clearly lower when the technological revolution takes place. The reason is the uncertainty term $\widehat{\sigma}_t^2$, whose coefficient is negative in part (a) and positive in part (b). In part (a), $\widehat{\sigma}_t^2$ is systematic (it affects π_t), whereas in part (b), it is idiosyncratic (it does not affect π_t). Since $\widehat{\psi}_t$ (expected cash flow) is essentially the same in both scenarios, the difference between M/B in parts (a) and (b) is due to the discount rate effect. This knife-edge case shows that M^N/B^N is likely decreasing in $\widehat{\psi}_t$ close to time t^{**} .

In summary, the cash flow effect usually dominates close to time t^* , leading to an initial positive relation between M^N/B^N and $\widehat{\psi}_t$, but the discount rate effect usually dominates close to time t^{**} , leading to an eventual negative relation. During a technological revolution, $\widehat{\psi}_t$ generally increases, leading to a bubble-like pattern in M^N/B^N .

3.2. Discussion

Corollary 2 shows that the adoption reduces the new economy's M/B , holding $\widehat{\psi}_t$ constant. Intuitively, the adoption does not bring any benefit to the new economy, which already uses the new technology. On the contrary, it increases systematic risk and thus reduces the new economy's market value. The model features only one shareholder, the representative agent, who employs infinitely more capital in the old economy than in the new economy. This agent wants the adoption to take place because the utility gain from making the old economy more productive outweighs the negligible loss of market value in the new economy.

Analogous to Corollary 2, we can show that the old economy's market value also decreases at time t^{**} if the adoption takes place when $\widehat{\psi}_{t^{**}}$ is close to $\underline{\psi}$. Interestingly, the representative agent chooses to adopt the new technology even if doing so reduces the market value of her stocks. There is a difference between maximizing utility and maximizing market value. The adoption occurs only if it increases the agent's expected utility. This adoption changes the economic environment by installing (what the agent perceives to be) a more productive technology and by increasing expected stock returns. In this new environment, stock prices are lower (due to higher discount rates) but expected utility is higher (due to higher expected wealth). Expected utility and stock prices need not move in the same direction because stock prices are related to the agent's marginal utility rather than to the level of utility.

We solve the social planner's problem in which a utility-maximizing representative agent owns all output by holding the stocks of the old and new economies. When a new technology is invented, it becomes property of the social planner. The social planner finds it optimal to set up a small-scale new economy to learn about the new technology before deciding whether

to adopt this technology in the large-scale old economy. Upon adoption, there is no transfer from the old economy to the new economy because the new economy does not own the new technology (the social planner does). As an example of a new economy firm, Amazon was an early user of the Internet but it did not own the Internet technology.

As an alternative to the social planner’s problem, we analyze a competitive economy in which firms independently decide whether to adopt the new technology while maximizing their own market values. We present this alternative decentralized model in the Appendix, and find that it produces the same stock price dynamics as the social planner’s problem. The alternative model features “network externalities,” in that the average productivity of a technology increases as the fraction of firms using this technology increases. Each firm makes its own adoption decision independently, taking the decisions of all other firms as given. Adopting the same technology as other firms has two opposing effects. On one hand, it hurts the firm, because the technology adopted by all other firms carries more systematic risk. On the other hand, it benefits the firm through network externality gains. We show that it is possible to choose the magnitude of the network externality gains such that the solution is identical to that in the social planner’s problem. Specifically, the Nash equilibrium at time t^{**} is such that all firms adopt the new technology if $\hat{\psi}_{t^{**}} \geq \underline{\psi}$, but none of them do if $\hat{\psi}_{t^{**}} < \underline{\psi}$, analogous to our Proposition 2. As a result, all pricing formulas are the same as in the social planner’s problem, and the same “bubbles” in stock prices obtain.

The alternative model highlights the lack of coordination among firms in a competitive economy. Although each firm maximizes its own market value, the aggregate effect of the firms’ adoptions is to reduce market values. The reason is that firms adopting the new technology do not fully internalize the resulting increases in the volatility of the stochastic discount factor. Each adopting firm imposes a negative externality on the market values of other firms by increasing systematic risk in the economy. We see that the stock price patterns obtained in our simple model with a utility-maximizing social planner hold also in a more complicated model featuring value-maximizing competitive firms.

Other ways of decentralizing the model could also lead to similar stock price dynamics. For example, suppose that firms facing different conversion costs observe signals about ψ . As $\hat{\psi}_t$ rises during a technological revolution, the proportion of firms that adopt the new technology also rises. This proportion might play the same role as the adoption probability in our model: As the proportion rises from about zero to one, the volatility of the stochastic discount factor also rises, making the uncertainty about ψ increasingly systematic.

In our simple model, all output represents firm profits, so productivity and profitability

coincide. In reality, technological advances lead to permanent increases in productivity but only temporary increases in profitability. In the long run, new technology tends to benefit workers and consumers, not producers. Therefore, we also analyze a richer model in which labor income drives a wedge between productivity and profitability.⁷ In this model, productivity gains from new technology last until time T , but profitability gains last only until time $t^{***} < T$, after which all productivity gains go to labor. Profitability affects the cash flow to stocks, whereas productivity affects the discount rates. Systematic risk depends on uncertainty about productivity because the agent's total wealth depends on productivity. As a result, our basic mechanism is unaffected by the shorter profitability horizon. Indeed, we find that this richer model produces stock price dynamics very similar to those reported here. For the same parameter values, the bubble pattern is somewhat less pronounced, but more dramatic patterns can be easily obtained after plausible parameter changes.

4. Empirical Implications

The purpose of this section is to analyze the model-implied paths of the key variables during technological revolutions. We simulate 50,000 samples of shocks in our economy and compute the paths of quantities such as the M/B ratios and volatilities in each simulated sample. We split the 50,000 samples into two groups, depending on whether or not the new technology is adopted at time t^{**} , and plot the average paths of prices and volatilities across all samples within each group. Our objective is to understand how these paths differ depending on whether or not the new technology leads to a technological revolution.

Table 1 shows the parameters used in our simulations. For the productivity processes, we choose parameter values close to those estimated by Pástor and Veronesi (2006) for the dynamics of profitability. The relation between productivity and profitability in our model is explained in Section 3.2. The parameter values for the conversion cost, time horizon, risk aversion, and prior beliefs about ψ are varied later in our sensitivity analysis.

Figure 2 plots the average paths of $\widehat{\psi}_t$, p_t , and $\sigma_\pi \equiv \text{Std}(d\pi_t/\pi_t)$. Panel A shows that the average drift in $\widehat{\psi}_t$ during technological revolutions is positive, due to conditioning on the ex post event that $\widehat{\psi}_{t^{**}} \geq \underline{\psi}$ (without such conditioning, $\widehat{\psi}_t$ is a martingale; see equation (17)).⁸ Analogously, conditional on $\widehat{\psi}_{t^{**}} < \underline{\psi}$, $\widehat{\psi}_t$ in Panel B (no revolution) drifts downward. The drift is less pronounced in Panel B than in Panel A because $\widehat{\psi}_{t^*} = 0$ and $\underline{\psi} > 0$. The

⁷This model is presented in the Technical Appendix, which is downloadable from the authors' websites.

⁸Brown, Goetzmann and Ross (1995) provide a mathematical proof of a related statement in their analysis of stock returns conditional on the stock's survival through the end of the sample.

average probability of adoption, p_t , drifts up in Panel C (revolution) and down in Panel D (no revolution), as expected. The volatility of the stochastic discount factor, σ_π , is almost flat while p_t is low, but it increases as p_t increases (Panel E).

Figure 3 plots the average paths of M/B and volatility for the new economy (solid line) and the old economy (dashed line). The panels on the left are based on the samples in which $p_{t^{**}} = 1$ (revolution); the panels on the right condition on $p_{t^{**}} = 0$ (no revolution).⁹ The dotted vertical lines mark the time when the new technology arrives, $t^* = 1$, and the time at which the agent decides whether to adopt the technology, $t^{**} = 9$.

Panel A of Figure 3 plots the average paths of M/B across all technological revolutions. The new economy’s M/B rises and then falls, as predicted in Section 3.1. Since we are conditioning on the adoption of the new technology at time t^{**} , $\hat{\psi}_t$ must go up between t^* and t^{**} (Figure 2). This increase in $\hat{\psi}_t$ has two countervailing effects on prices. First, it increases expected future cash flow from the new technology, pushing M/B up. Second, it increases the adoption probability, which makes the risks embedded in the new technology ever more systematic (affecting W_T), which then increases the discount rate applied to future cash flow, pushing M/B down. For the old economy, the discount rate effect outweighs the cash flow effect from the outset, leading to a slow decline in M/B. For the new economy, the cash flow effect is stronger at first, but the discount rate effect prevails in the end, producing a “bubble.” Since the path plotted in Panel A is an *average* across all revolutions, it shows that apparent bubbles in high-tech stock prices are not merely possible in a rational world; they should in fact be *expected* during technological revolutions.

Different technological revolutions produce different paths of M/B, depending on the path of realized productivity. These individual paths look mostly like bubbles that peak at different times, and they are less smooth than the average path plotted in Panel A of Figure 3. On this average path, the peak-to-bottom drop in the new economy’s M/B lasts 5 years, but for some revolutions, the price drop is much more abrupt. For example, for 10% of all revolutions, the peak-to-bottom drop lasts less than 2.56 years, and for 5% of all revolutions, it lasts less than 1.68 years. The magnitude of the price drop also exhibits substantial dispersion across revolutions. On the average path, M/B falls by 2.9 from the peak to the bottom, but for 5% of all revolutions, it falls by more than 7.8.

Panel B of Figure 3 plots the average paths of M/B across all samples in which $p_{t^{**}} = 0$

⁹The fraction of the simulated samples in which $p_{t^{**}} = 1$ is approximately equal to the ex ante probability of adoption implied by our parameter choices, $p_{t^*} \approx 2\%$, as expected. In principle, any product innovation could potentially lead to a technological revolution, but very few do, both in reality and in our model.

(no revolution). In these samples, $\widehat{\psi}_t$ declines slightly between t^* and t^{**} , nudging the M/Bs down as well. The decline is larger in the new economy, for two reasons. One, the new economy's M/B is more sensitive to $\widehat{\psi}_t$, as discussed earlier. Two, uncertainty about ψ gradually declines due to learning, which reduces M/B for the new economy but not for the old economy (see equation (14)). Thanks in part to this uncertainty, the level of M/B is higher in the new economy than in the old economy, in both Panels A and B. Higher productivity is another reason why the new economy's M/B is higher in Panel A, even after time t^{**} . Although the adoption makes the long-run means of productivity equal in both economies, the productivity at time t^{**} is higher in the new economy ($\rho_{t^{**}}^N$ is likely to be high to make $\widehat{\psi}_{t^{**}} > \underline{\psi}$), lifting the M/B of the new economy above that of the old economy.

Panel C of Figure 3 plots the average paths of stock return volatility across all technological revolutions. Volatility is higher in the new economy than in the old economy, partly due to higher volatility of the fundamentals, but mostly due to uncertainty about ψ . To understand the U-shape in the new economy's volatility, recall that shocks to $\widehat{\psi}_t$ affect stock prices via the discount rate and cash flow effects, which work in opposite directions. Around time t^* (t^{**}), the cash flow (discount rate) effect dominates, so the two effects do not offset each other much and the volatility is high. The volatility is lowest when the two effects cancel each other, which happens at some point between times t^* and t^{**} ; hence the U-shape. For the old economy, the discount rate effect dominates from the outset, so the old economy's volatility slowly increases as the rising adoption probability makes the stochastic discount factor more volatile. The spike in volatility at time t^{**} is caused by those simulated paths for which $\widehat{\psi}_{t^{**}}$ is close to $\underline{\psi}$ because then p_t swings a lot shortly before time t^{**} , making returns highly volatile (Corollary 2). We show later that the volatility spike disappears (but all other effects remain) when t^{**} is chosen optimally instead of being fixed exogenously. Panel D plots the average return volatility across all no-revolution samples. In these samples, the discount rate effect is weak and volatility is roughly constant over time.

Panels A and B of Figure 4 plot the market beta of the new economy, β , defined as the slope from the regression of the new economy stock returns on the old economy stock returns. In Panel A, where we condition on $p_{t^{**}} = 1$ (revolution), β exhibits an asymmetric U-shape pattern, for the following reason. Positive shocks to $\widehat{\psi}_t$ always reduce the market value of the old economy stocks, but they increase the value of the new economy stocks initially while the cash flow effect prevails over the discount rate effect, leading to an initial decrease in β . Only after the discount rate effect overcomes the cash flow effect, shocks to $\widehat{\psi}_t$ begin affecting the market values of both economies in the same direction, leading to an increase in β . Since the effect of $\widehat{\psi}_t$ on the old economy stocks increases with the adoption

probability, the rise in β is more dramatic than the initial fall. After a mild decline in the first half of the revolution, β doubles in the second half, from 0.75 to 1.5. The average beta in the no-revolution samples, plotted in Panel B, is almost flat over time.

As explained above, two aspects of systematic risk increase during technological revolutions: the old economy's volatility and the new economy's beta. The increase in the old economy's volatility raises the discount rates for both economies, old and new, holding β constant. The increase in β gives an additional boost to the discount rate of the new economy, which is why stock prices fall by more in the new economy than in the old economy.

The remaining panels of Figure 4 plot the average realized returns (solid line) and expected returns (dashed line).¹⁰ In technological revolutions, realized stock returns are first positive and then negative for both economies, due to an *ex post selection bias*. Ex post, we know that a technological revolution took place at time t^{**} , but ex ante, we only have a probability assessment of this event. Before time t^{**} , stock prices are not expected to rise and fall; expected returns are given simply by the covariances with the stochastic discount factor. However, conditioning on a technological revolution means that the adoption probability p_t must be revised upward between times t^* and t^{**} , causing a bubble-like pattern in prices through the cash flow and discount rate effects discussed earlier. The bias of realized returns relative to expected returns is due solely to ex post conditioning on $p_{t^{**}} = 1$; when this conditioning is removed, the bias disappears. (Across all 50,000 simulations, average realized returns are equal to average expected returns.) The rise and fall in stock prices during technological revolutions are observable ex post but not predictable ex ante.

The unexpected arrival of the new technology causes the old economy's market value to drop immediately at time t^* (Panel E of Figure 4). This drop is driven by two forces. The possibility of eventual adoption means that conversion costs might be paid at time t^{**} , and it also increases systematic risk and so drives up the discount rate.

4.1. Sensitivity Analysis

This section examines the sensitivity of the price dynamics to our parameter choices. Figure 5 is the counterpart of Panel A of Figure 3 (revolution), with various parameter changes.

In Panel A of Figure 5, risk aversion $\gamma = 3$, as opposed to $\gamma = 4$ in Figure 3. Lower risk aversion increases M/B in both economies, as expected, but the pattern of M/B is otherwise the same as that in Figure 3. A hump-shaped pattern in M^N/B^N obtains for any $\gamma > 1$.

¹⁰All returns are annualized by multiplying each interval- dt return by $1/dt$.

In Panel B of Figure 5, the conversion cost is $\kappa = 0$, as opposed to $\kappa = 0.1$ in Figure 3. The main effect of the lower κ is to decrease M^N/B^N . The lower conversion cost makes it more likely that the new technology will be adopted, which increases discount rates and thus depresses prices. For the old economy, there is also a counterbalancing effect, as the lower conversion cost increases the old economy's post-conversion capital $B_{t^{**}} = B_{t^{**}}(1 - \kappa)$. The two effects approximately cancel out, so the old economy's M/B is almost unaffected by the change in κ . Most important, the price patterns look just like those in Figure 3.

In Panel C, prior uncertainty about ψ is $\hat{\sigma}_{t^*} = 8\%$, compared to $\hat{\sigma}_{t^*} = 4\%$ in Figure 3. The higher uncertainty increases M^N/B^N , especially close to time t^* when p_t is small (equation (14)). However, as p_t increases during a revolution, uncertainty becomes increasingly systematic, pushing M^N/B^N down, and this discount rate effect is stronger when systematic uncertainty is higher. Therefore, in technological revolutions characterized by high uncertainty, the new economy firms tend to start out with high valuations that exhibit a large decline. High uncertainty amplifies the bubble-like pattern in stock prices.

In Panel D of Figure 5, the time until the adoption decision is shortened to $t^{**} - t^* = 6$, compared to $t^{**} - t^* = 8$ in Figure 3. Faster adoption increases M^N/B^N . To understand this effect, we note two facts. First, faster adoption implies higher uncertainty about ψ at time t^{**} because there is less time to learn (equation (18)). Second, faster adoption implies a higher adoption threshold $\underline{\psi}$ because t^{**} is lower and $\hat{\sigma}_{t^{**}}$ is higher (equation (8)). Since $\hat{\psi}_t$ has less time to reach a higher threshold, the adoption probability $p_{t^{**}}$ is lower, which implies that systematic risk is initially lower and M^N/B^N starts higher than in Figure 3. M^N/B^N then rises higher and falls deeper than in Figure 3, conditional on $p_{t^{**}} = 1$, because both the cash flow effect and the discount rate effect are stronger when adoption is faster. The cash flow effect is stronger because in order for $\hat{\psi}_t$ to reach a higher threshold in shorter time, the increase in $\hat{\psi}_t$ must be sharper. The discount rate effect is stronger because uncertainty at time t^{**} is higher, and conditional on $p_{t^{**}} = 1$, this uncertainty is systematic. Since both effects are stronger, the rise and fall in M^N/B^N are more striking than in Figure 3. Faster adoption of the new technology magnifies the bubble-like pattern in stock prices.

4.2. Optimal Adoption Time

In this section, we relax the assumption that t^{**} is exogenously given. Without this assumption, no closed-form solutions are available, but we can solve the problem numerically. The agent is choosing the optimal time t^{**} , $t^* \leq t^{**} \leq T$ to adopt the new technology (no adoption is a possibility). This is essentially a problem of solving for the best time to exercise an

American real option. The details of our solution are in the Appendix.

Figure 6 plots the average paths of M/B and volatility when t^{**} is chosen optimally. Depending on the path of productivity, the adoption can occur anytime between t^* and T , but averaging across very different t^{**} 's would not be meaningful. For comparison with Figure 3 in which $t^{**} = 9$ years, the left panels of Figure 6 report averages across those simulations in which the optimal t^{**} is between years 8 and 10. Our main results are unaffected by endogenizing t^{**} . During revolutions, the new economy's M/B exhibits a rise and fall similar to that in Figure 3, albeit slightly weaker (a stronger “bubble” pattern is obtained for $\gamma = 3$, as we show in an earlier draft). The path of volatility in Panel C is also similar, except that the volatility spike observed in Figure 3 disappears, as explained earlier.

5. Empirical Evidence

In this section, we empirically examine the behavior of stock prices during two technological revolutions, one recent and one distant. For both revolutions, we consider the key quantities in our model, such as the new economy's market beta and the level and volatility of stock prices, and compare their empirical dynamics with their model-implied dynamics.

5.1. The Internet Revolution

The Internet's predecessor, Arpanet, was created in 1969 with funding from the U.S. Department of Defense. Arpanet ceased to exist in 1990, roughly when the team of Tim Berners-Lee at CERN released the World Wide Web. The first Web site, info.cern.ch, appeared in 1991. The first graphics-based web browser, Mosaic, was launched in 1993 by Marc Andreessen at the National Center for Supercomputing Applications. In 1994, Andreessen co-founded Netscape Communications, which went public in August 1995 in the first Internet IPO. The first big pioneer of e-commerce was the online bookseller Amazon.com, which was launched by Jeff Bezos in 1995 and went public in May 1997. The Internet gradually became mainstream. The number of web servers grew from about 23,000 in mid-1995 to about 30 million in mid-2001 and 65 million in mid-2005 (see www.zakon.org/robert/internet/timeline/). A prominent example of the Internet's integration into traditional business models was the creation of the first “clicks-and-mortar” company through the merger of AOL and Time Warner.¹¹ Today, the Internet technology is an indelible part of the economic landscape.

¹¹AOL announced its plan to acquire Time Warner (for some \$182bn in stock) in January 2000, and the FTC approved the deal in January 2001. “The merger, the largest deal in history, combines the nation's top internet service provider with the world's top media conglomerate. The deal also validates the Internet's

To provide a benchmark for our empirical analysis, we plot the model-implied dynamics of some key variables in Figure 7. These are the expected dynamics during a revolution, in that we average the model-implied paths across many simulations in which the new technology is adopted at time t^{**} . We keep all parameters from the baseline case (Table 1) except that we shorten the duration of the revolution from eight to six years because the Internet revolution was relatively fast. Panel A of Figure 7 shows that the new economy’s market beta decreases slightly (from 0.9 to 0.7) in the first half of the revolution, but then it increases sharply in the second half, reaching 1.65 at time t^{**} before falling to one. This increase in beta is even steeper than in the baseline case (Panel A of Figure 4). Panel B shows that the increase in stock return volatility is also steeper than in the baseline case (Panel C of Figure 3), with the old economy’s volatility doubling to 38% and the new economy’s volatility rising to over 65%. Panel C plots the market values of both economies. There is a clear “bubble” in the new economy, whose market value quintuples and then falls by half. The old economy’s market value also rises and falls, but this pattern is much weaker than in the new economy. Panel D shows that the old economy’s productivity begins rising immediately after the adoption of the new technology, when it begins mean-reverting toward a higher mean.

Figure 8 is an empirical counterpart of Figure 7 for the period 1992–2005. For simplicity, we assume that the technology-loaded NASDAQ index represents the new economy and the NYSE/AMEX index is the old economy. We obtain daily index returns from CRSP.

Panel A of Figure 8 plots the market beta of the NASDAQ index, along with its two-standard-error confidence bands. The beta is computed daily as the slope coefficient from the regression of the NASDAQ returns on the NYSE/AMEX returns over the most recent 500 trading days (i.e., about two years). After a slight decrease from about 1.2, NASDAQ’s beta doubles from 1.0 to 2.0 between 1997 and mid-2002, and this increase is highly statistically significant. This empirical pattern is strikingly similar to the model-implied pattern in Panel A of Figure 7, in which the beta also decreases by about 0.2 before rising 2.3-fold by the end of the revolution. According to the model, the time when the beta peaks is the time of the large-scale adoption; hence the evidence on NASDAQ’s beta is consistent with the probability of the Internet’s large-scale adoption reaching one by mid-2002.

Panel B of Figure 8 plots the standard deviations of returns on the NASDAQ and NYSE/AMEX indices, computed daily over the most recent 500 trading days. NASDAQ’s volatility falls from 17% in 1992 to 11% in 1995, before rising to 47% at the beginning of

role as a leader in the new world economy, while redefining what the next generation of digital-based leaders will look like.” (CNN Money, Jan 10, 2000).

2002. NYSE/AMEX’s volatility falls from 13% in 1992 to 8% in 1995, before rising to 21% by the end of 2002. These patterns are similar to the model-implied patterns in Panel B of Figure 7 in several ways: (i) the new economy’s volatility always exceeds the old economy’s volatility; (ii) both volatilities generally rise over time, with a bit of a U-shape pattern; (iii) the new economy’s volatility rises much faster; and (iv) both volatilities peak at about the same time. In the model, both volatilities peak at the time of the adoption; the volatility evidence is thus consistent with the Internet revolution ending sometime in 2002.

Panel C of Figure 8 plots the index levels for NASDAQ and NYSE/AMEX, namely, the value of \$1 invested in these indices in January 1992, with dividend reinvestment. The NASDAQ index quadruples between 1996 and March 2000, but then it falls back to the 1996 level by October 2002, after which it rises again. In contrast, NYSE/AMEX exhibits a much smaller rise and fall over the same period. This pattern is similar to the model-implied pattern in Panel C of Figure 7, in which the new economy’s market value also exhibits a “bubble” but the old economy’s rise and fall are much less pronounced.¹² According to the model, the time when both indices stop falling is the time of the large-scale adoption; hence this evidence is consistent with the Internet’s adoption by October 2002.

Panel D of Figure 8 plots a three-year moving average of multifactor productivity growth in the private business sector of the U.S. economy. (This is the most commonly used multifactor productivity measure, according to the Bureau of Labor Statistics, which is the source of the data.) In year t , we plot the average annual productivity growth in years $t - 2$, $t - 1$, and t . Multifactor productivity growth averaged about 1% per year in the 1990s, but it increased sharply after year 2002: from 1% per year in 2002 to 1.5% in 2003 and 2.5% in 2004 and 2005. A similar pattern is observed for labor productivity.¹³ The observed productivity pattern is similar to the model-implied pattern in Panel D of Figure 7, except that

¹²Pástor and Veronesi (2006) show that NASDAQ’s M/B dropped by 5.3 (from 8.5 to 3.2) from the peak to the bottom in 2.7 years. Both the duration and the magnitude of this drop in M/B correspond closely to their counterparts in our model. The corresponding model-implied average pattern in M/B is plotted in Panel D of Figure 5 (in which the revolution lasts 6 years, as it does in Figure 7). This average M/B falls by 4.5 from the peak to the bottom, also in 2.7 years, matching the observed values remarkably closely.

¹³In his Remarks Before Leadership South Carolina on August 31, 2006, Ben Bernanke argued that “One of the most important economic developments in the United States in the past decade or so has been a sustained increase in the growth rate of labor productivity... From the early 1970s until about 1995, productivity growth in the U.S. nonfarm business sector averaged about 1.5% per year... Between 1995 and 2000, however, the rate of productivity growth picked up significantly, to about 2.5% per year... Talk of the “new economy” faded with the sharp declines in the stock valuations of high-tech firms at the turn of the millennium. Yet, remarkably, productivity accelerated further in the early part of this decade. From the end of 2000 to the end of 2003, productivity rose at a 3.5% annual rate and it is estimated to have increased at an average annual rate of 2.25% since the end of 2003. These advances were achieved despite adverse developments that included the 2001 recession, the terrorist attacks of September 11, [etc.]”

that figure plots the level of productivity as opposed to its growth rate.¹⁴ In the model, the economy’s productivity begins rising at the time of the adoption; hence the productivity evidence is consistent with a large-scale adoption of the Internet by 2002.

Overall, we find Figure 8 remarkably similar to Figure 7. The patterns of NASDAQ’s beta and NYSE/AMEX’s volatility show that both sectors experienced large increases in systematic risk in 1997–2002, supporting the key prediction of the model. To summarize, the empirical evidence seems consistent with the joint hypothesis that our model holds and that the Internet technology was adopted on a large scale by 2002.

5.2. American Railroads Before the Civil War

Our paper is motivated by the technological revolutions, listed in the introduction, that were accompanied by apparent bubbles in stock prices. In this section, we conduct an “out-of-sample” analysis of a revolution whose stock prices do not seem to have been analyzed before. We analyze the first major technological revolution that took place in the U.S. since the New York Stock Exchange was organized in 1792 – the introduction of steam-powered railroads (RRs). In the early days of the RR, there was substantial uncertainty about whether the RR technology would be ultimately adopted on a large scale. After examining the historical milestones of American RRs in Section 5.2.1., we argue that the probability of a large-scale adoption rose gradually, and that it approached one in the late 1850s after the RR expansion west of the Mississippi River. We then empirically examine the behavior of the RR stock prices in 1830–1861 in Section 5.2.2. In the context of our model, our evidence is consistent with a large-scale adoption of the RR technology around year 1857.

5.2.1. Brief History

The steam engine, an 18th-century invention, was first used for rail-based transportation in the early 19th century in Britain. The United States followed shortly afterwards. The first RR act in the U.S. was passed in 1815 when the New Jersey legislature awarded a charter to Colonel John Stevens to build a RR between the Delaware and Raritan rivers.¹⁵ In 1825, Stevens operated the first locomotive in America – his 16-foot “Steam Waggon” ran around a circular rail track in Hoboken at 12 miles per hour. The construction of the first RR, the Baltimore & Ohio, began in July 1828. The Baltimore & Ohio initially used

¹⁴Our comparison seems reasonable because in the model, average productivity can grow only via technological revolutions, whereas in reality, there are also many non-revolutionary improvements in productivity. Therefore, in the data, it is the growth rate of productivity that sets a technological revolution apart.

¹⁵The discussion in this section draws especially on Stover (1961), Fogel (1964), and Klein (1994).

horses to draw its cars, but it replaced them in 1830 by a steam locomotive, Peter Cooper's "Tom Thumb." In 1830, both passenger and freight service commenced on the Baltimore & Ohio. RRs spread quickly. On Christmas Day in 1830, the "Best Friend of Charleston," the first locomotive built for sale in the U.S., made the first scheduled steam-RR train run in America. Between 1830 and 1840, the RR mileage in the U.S. grew from 23 to 2,808 miles. In 1840, only four of the 26 states had not completed their first mile of track.

The new RR technology competed with the existing modes of transportation such as wagons, stagecoaches, steamboats, and canals. Those were not without problems – wagons were slow and expensive, stagecoaches were uncomfortable, steamboats were dangerous and limited in scope, and canals froze over in winter. However, it was far from obvious in the 1830s and 1840s that the RRs would later come to dominate the transportation industry. For example, waterways were much less expensive than RRs, and wagons were not restricted to rails. While the RR mileage caught up with the canal mileage in the early 1840s, waterways still carried the great bulk of the nation's freight in the late 1840s. Writes Fogel (1964): "Far from being viewed as essential to economic development, the first RRs were widely regarded as having only limited commercial application. Extreme skeptics argued that RRs were too crude to insure regular service, that the sparks thrown off by belching engines would set fire to buildings and fields, and that speeds of 20 to 30 miles per hour could be "fatal to wagons, road and loading, as well as to human life." More sober critics questioned the ability of RRs to provide low cost transportation, especially for heavy freight. [Some] placed "a RR between a good turnpike and a canal" in transportation efficiency."

Nearly all RRs organized as corporations funded by private investors. More than half of the more than \$300 million invested in American RRs in 1850 was represented by capital stock, the remainder being in bonds. The freight business was economically more important than passenger traffic, which typically produced around 30% of the total revenue.

While most early RRs were built with local capital to provide local transportation, RR building became more ambitious in the 1850s. This decade "was one of the most dynamic periods in the history of American RRs" (Stover, 1961). RR mileage expanded from 9,021 in 1850 to 30,626 in 1860, and total investment in the industry increased from about \$300m to about \$1,150m over the same period. This growth was spurred by land grants to RRs by the federal government. The first land-granting act was passed by the Congress in 1850, aiding the Illinois Central and the Mobile & Ohio RRs. The RR growth in the 1850s was also stimulated by the discovery of gold in California and the lure of the trans-Pacific trade. In the 1850s, New York, Philadelphia, and Baltimore all achieved their rail connections with

the west. In 1853, an all-rail route opened from the East to Chicago, and Chicago quickly became the rail capital of the nation. The RR technology also advanced in the 1850s – telegraph was first used to dispatch trains, T-rails became the general rule, and so did the standard track gauge, at least in the North.¹⁶ “Instead of merely serving as connectors between navigable bodies of water as originally conceived, RRs were replacing them as the preferred way of transport” (Klein, 1994).

The dramatic RR growth in the 1850s is also evident in Figure 9, which plots the total rail consumption in the U.S., measured by the number of track-miles of rails laid each year (Fogel, 1964). Rail consumption grew fast in the 1830s, but especially fast during the decade leading up to 1856. After 1856, rail consumption slowed down and even declined in 1861 when the Civil War began, but it accelerated again after the war.

The diffusion of the RR technology made a leap in 1856 when two milestone RRs were completed: the Illinois Central, the longest RR in the world (705 miles), and the Sacramento Valley, the first RR in California. Also in 1856, the first RR bridge across the Mississippi was built near Davenport, Iowa, heralding future westward expansion into the region then known as the “Great American Desert.” This westward expansion was the defining feature of the RR growth in the decades to come. The RRs shaped the economy of the West, creating new national markets and fostering unprecedented economic specialization across the nation.

By the late 1850s, it seemed clear that the RR had become a dominant form of transportation. According to Stover (1961), “By 1860 the canal packets and river steamers had lost much of their passenger traffic” to the RR. In 1860, every state save Minnesota and Oregon had RR mileage, and 29 of the 33 states had more than 100 miles of line. Klein (1994) argues that “By 1860... [the RR] had emerged not only as the preferred form of transportation but also as the chief weapon of commercial rivalry.” This evidence suggests that a large-scale adoption of the RR technology took place by the end of the 1850s.

5.2.2. Railroad Stock Prices

To examine the behavior of RR stock prices in the early days of the RR (1830–1861), we use the data compiled by Goetzmann, Ibbotson, and Peng (2001). These data contain monthly individual stock prices for NYSE stocks from 1815 to 1925, as well as annual dividends for a

¹⁶The Northern RRs were using 11 different track gauges in the 1850s, but the standard gauge, 4’8.5”, became by far the most common by 1860, according to Stover (1961). The South was still mostly on the 5’ gauge. Benmelech (2007) exploits the diversity of track gauges in 19th century American railroads to examine the effect of asset liquidation value on capital structure.

subset of stocks from 1825 to 1870. The data are provided at <http://icf.som.yale.edu/nyse/>.

To focus on common stocks, we exclude stocks classified as “preferred” or “scrip” in the database. (Scrips are certificates convertible into shares when fully paid-in.) If such classification is not provided, we examine the stock name and exclude stocks whose name contains an indication of non-common status such as “pref,” “pr.,” “pf,” or “scrip.” Among the 671 stocks in the database, we identify and exclude 85 preferred stocks and 29 scrips.

We identify 284 RR stocks (42.32% of the whole sample) by examining the stock names. The first RRs that appear in our price index (discussed below) in 1831 are Camden & Amboy, Canajoharie & Catskill, Harlem, and Ithaca & Oswego. All RRs that have at least one valid monthly common stock return between 1830 and 1861 are listed in Table 2.

We clean the monthly price file to remove apparent data errors. To proceed in a systematic fashion, we exclude all prices that imply implausibly large return reversals. Specifically, we exclude prices that more than tripled compared to the most recent available price and then fell to less than a third at the nearest future observation, as well as prices that experienced the same reversals in reverse order (first down, then up). We eliminate 34 such prices in our 1830–1861 sample. We also examine all price sequences in which the price increased or decreased at least tenfold without reversal, and eliminate six suspicious price entries between 1830 and 1861. We retain the price entries that imply returns below -90% at the very end of a stock’s price series because these could be stocks heading for bankruptcy. Altogether, we delete 40 of the 15,276 price entries between 1830 and 1861, or 0.26% of the sample.

Before the price coverage in the database improves in 1848, uninterrupted price sequences for RR stocks are rare. In no month before 1848 are there more than five RR stocks with valid monthly returns, and there are months with zero RR returns. An important part of the problem are gaps in the price series, in which one or several missing values are sandwiched between two valid prices for a given stock. To alleviate the data shortage, we fill in such gaps by linear interpolation, but only for gaps that are no more than three months long. This procedure substantially increases the price coverage early in the sample. For example, without interpolating, the RR year-end price-dividend ratio discussed below would have only three valid observations prior to 1847; with interpolation, the number of valid observations increases to eight. Without interpolating, our results would be noisier, with more missing values, but they would lead to the same basic conclusions.

We compute monthly RR (non-RR) index returns as price-weighted averages of monthly

capital gains across all RR (non-RR) stocks.¹⁷ We use capital gains rather than total returns because the dividend data available to us are annual, not monthly, and because these data are spotty, especially early in the sample (Goetzmann et al. (2001) suggest that their dividend sample is incomplete). The resulting return series can be viewed as approximations to the actual returns earned by investors.

Panel A of Figure 10 plots the beta of the RR index, with a two-standard-error confidence band. The beta in month t is the slope coefficient from the regression of the most recent 36 monthly RR returns (in months $t-35$ through t) on the non-RR returns. Not surprisingly, the beta estimates computed from only 36 observations have wide confidence bands. Nonetheless, it appears that the largest increase in beta took place in the 1850s: The beta estimate rose from about 0.2 in 1850 to about 1.8 in 1856, before falling to about 1.0 right after 1857. This empirical pattern is quite similar to the model-implied pattern in Panel A of Figure 4 if we assume that the RR technology was adopted on a large scale in 1857.

Panel B plots the volatility of returns in the RR and non-RR industries, computed annually as the standard deviation of monthly industry returns within the year. Two facts seem noteworthy. First, the RR volatility exceeds the non-RR volatility in every year except 1841, consistent with the presence of uncertainty about the RR technology. The volatility difference is also due in part to the fact that the RR portfolio is less diversified than the non-RR portfolio, but it persists also after the number of RRs with valid monthly stock returns increases sharply (from 6 in December 1847 to 15 in January 1848, to 25 in July 1850). The second interesting fact in Panel B is that return volatility increases sharply in 1857, to 33.5% per year for RRs and to 23.1% for non-RRs. Comparison with Panel C of Figure 3 shows that both facts are consistent with a large-scale adoption of the RR technology in 1857.

Panel C plots the stock price index levels for the RR and non-RR industries, obtained by cumulating monthly returns in each industry. The general downward trend in the price indexes is partly due to the absence of dividends and partly due to the absence of inflation in the economy. The biggest price declines occur in the mid-1850s. For example, between June 1853 and October 1857, the RR price index falls by 58.3%, whereas the non-RR index falls by 33.9%. Both the sharp price decline for RRs and the milder decline for non-RRs are consistent with the RR technology being adopted on a large scale around 1857. Recall that our model predicts that the new economy (RR) stock prices fall by more than the old economy (non-RR) stock prices shortly before the adoption of the new technology.

¹⁷Goetzmann et al. (2001) argue that price-weighting best approximates the return on a buy-and-hold portfolio, given the absence of information about market capitalization and book value in their database.

Various events played a role in the stock price decline in 1857. Investor confidence was shaken by embezzlement at the Ohio Life Insurance and Trust Company in August, as well as by the government’s loss of a large amount of gold at sea in September. Other commonly cited negative influences include falling grain prices, British withdrawals of capital from U.S. banks, and manufacturing surpluses. The stock market bottomed in October 1857 amidst a number of bank failures. However, the stock price decline cannot be fully attributed to the banking panic. According to Mishkin (1991), “Rather than starting with the banking panic in October 1857, the disturbance to the financial markets seems to arise several months earlier with the rise in interest rates, the stock market decline... and the widening of the interest rate spread.” Mishkin’s last observation is particularly interesting. He shows that the spread between the yields of low- and high-quality corporate bonds was unusually high in 1857–1859, higher than at any future time before the 1930s. These high yield spreads indicate that the risk premia in the late 1850s were high, consistent with our story. Mishkin also opines that the decline in stock prices in the late 1850s “might be linked to the general rise in interest rates which lowers the present discounted value of future income streams.” This is precisely our story - stock prices fall shortly before the adoption of the new technology because discount rates increase due to an increase in systematic risk.

In Panel D of Figure 10, we do not plot productivity as we did in Figure 8 because, to our knowledge, productivity in this period has been computed only at a ten-year frequency from census data. We note, however, that the evidence points to a large increase in productivity after the late 1850s. For example, Cochrane (1979) argues that productivity advanced sharply just before the Civil War, and Craig and Weiss (1993) conclude that “the 1860s saw the greatest increase in output per farm worker of any decade in the 19th century.” This evidence on productivity further strengthens the case for a large-scale adoption of the RR technology in the late 1850s in the context of our model (cf. Panel D of Figure 7).

Panel D plots the aggregate price-to-dividend ratio (P/D) for the RR and non-RR industries. Each year, we compute P/D as the sum of year-end prices divided by the sum of dividends paid in that year, summing across all RR (or non-RR) stocks with valid price and dividend data. Note three main results. First, the P/D of RRs almost invariably exceeds the P/D of non-RRs before the mid-1850s. Second, the RR P/D falls from 24.9 in 1846 to 15.8 in 1852, to 6.5 in 1857. Third, the non-RR P/D falls as well, but less dramatically: from 14.0 to 12.8 to 9.1 over the same period. While interpreting the noisy data requires caution, all three results in Panel D are consistent with the joint hypothesis that our model is true and that the new RR technology was widely adopted around 1857.

5.3. Other Evidence

Three recent papers explicitly test some predictions of our model. Bharath and Viswanathan (2006) empirically analyze the model’s risk implications at the firm level. They examine 252 brick-and-mortar firms that launched commercial websites (i.e., adopted the Internet technology as a way of doing business) in 1995–2004. The authors find that adopting the new technology is associated with an increase in firm risk, with differences between the early and late adopters: Firms that adopted the Internet before March 2000 (while stock prices were rising) experienced significant increases in idiosyncratic risk, whereas firms that adopted after March 2000 had significant increases in systematic risk. The authors conclude that their evidence provides strong support for our model.

Mazzucato and Tancioni (2006) analyze stock prices and patent-related measures of innovation in a sample of firms in the pharmaceutical and biotechnology industries between 1975 and 1999. They find that the firms’ price-earnings ratios are positively related to innovation as well as to idiosyncratic risk, and argue that this evidence supports our model.

Hoberg and Phillips (2006) empirically examine the real and financial outcomes following industry booms in 1972–2004. They test the risk predictions of our model at the industry level. They find that market betas increase and idiosyncratic risk declines after booms, consistent with our model. The authors find strong support for our model in competitive industries but not in concentrated industries. It would be useful to extend our simple model to analyze the effect of product market competition theoretically.

In earlier work, Mazzucato (2002) studies the early phases of the life-cycles of the automobile and PC industries in the U.S. She finds that in both industries, stock prices were the most volatile when technological change was the most radical. She also finds idiosyncratic risk to be higher in the early stage of industry evolution, consistent with our model.

6. Conclusions

We develop a general equilibrium model that produces stock price “bubbles” during technological revolutions. The model features Bayesian learning about the average productivity of the new technology. Stock prices of innovative firms initially rise due to good news about this productivity, but they ultimately fall as the risk of the technology changes from idiosyncratic to systematic. The rise and fall in stock prices are observable only in hindsight – this pattern is unexpected by investors in real time but we observe it ex post when we focus on

technologies that eventually led to technological revolutions. The “bubbles” should be most pronounced in revolutions characterized by high uncertainty and fast adoption.

The model makes many empirical predictions. We find support for these predictions in the evidence from 1830–1861 and 1992–2005 when the railroad and Internet technologies spread in the United States. In the context of our model, the empirical evidence is consistent with large-scale adoptions of railroads by the late 1850s and the Internet by 2002.

We focus on stock prices but our model also has implications for productivity. The new technology does not bring productivity gains immediately upon arrival because the agent finds it optimal to learn about a new technology before adopting it. Since the agent chooses the adoption time optimally depending on what she learns, the time it takes for the productivity gains to begin emerging is endogenous in the model. The implication that productivity gains arrive with a lag seems reasonable; for example, although electric power first appeared around 1880, it was not until the 1920s that the productivity of the U.S. economy increased as a result of a large-scale adoption of electricity (David, 1991).

Our model has no implications for investment because the agent invests only a negligible amount in the new technology for learning purposes. As a result, the model fails to predict the large amount of investment that often accompanies technological revolutions. For example, it seems clear ex post that the 1990s witnessed significant overinvestment in the Internet infrastructure.¹⁸ Billions of dollars worth of optical fibers laid in the 1990s remain unused to this day as the market appears to have overestimated the extent to which the Internet would revolutionize the delivery of bandwidth-intensive content. It is possible that market irrationality, which is absent from our model, contributes to the overinvestment and stock price “bubbles” observed during technological revolutions. Future research can test our model against alternatives that involve behavioral biases. Some predictions of our model, such as those involving market beta, are unlikely to follow from behavioral models, in which there is typically no role for systematic risk. Since we find empirically that systematic risk increased sharply during two prominent revolutions, it seems unlikely that behavioral biases can fully explain the observed stock price patterns. Such biases could certainly be a part of the story, though, and quantifying their relative importance would be interesting.

¹⁸Johnson (2007) and DeMarzo, Kaniel, and Kremer (2007) develop models that can generate overinvestment in new technologies. These models have a different focus and they rely on different mechanisms – learning about the curvature of the production function and relative wealth concerns, respectively.

Table 1
Parameters used in Simulations.

$\bar{\rho}$	$\hat{\psi}_{t^*}$	$\hat{\sigma}_{t^*}$	
0.1217	0	0.04	
ϕ	σ	$\sigma_{N,0}$	$\sigma_{N,1}$
0.3551	0.07	0.07	0.07
κ	$t^{**} - t^*$	T	γ
0.1	8	30	4

Table 2
Railroads Appearing in our Price Index.

This table lists all railroads in our sample that have at least one valid monthly common stock return between 1830 and 1861. The railroads are sorted by the year of appearance of their first valid monthly return.

Year	Railroad
1831	Camden & Amboy; Canajoharie & Catskill; Harlem; Ithaca & Oswego
1832	Boston & Providence
1833	Boston & Worcester; Brooklyn & Jamaica
1835	Hudson & Berkshire; Long Island
1839	Auburn & Syracuse
1841	Auburn & Rochester
1844	Housatonic
1847	Hudson River; Macon & West
1848	Hartford & New Haven; New York & Erie
1849	Erie
1850	Albany & Schenectady; Baltimore & Ohio; Michigan Central; New York & Harlem
1851	Chemung
1852	Michigan & Southern
1853	Cincinnati, Hamilton & Dayton; Cleveland, Columbus & Cincinnati; Cleveland & Pittsburg; Cleveland & Toledo; Galena & Chicago; Illinois Central; Little Miami
1854	Chicago & Rock Island
1855	Michigan Southern & Northern Indiana
1856	Eighth Avenue; Lacrosse & Milwaukee; Macon & Western
1857	Chicago, Burlington & Quincy; Delaware, Lackawanna & Western; Indianapolis & Cincinnati
1858	Brooklyn City; Buffalo & State Line; Cleveland, Painesville & Ashtabula




Figure 1. The Sequence of Events. In this chart, t^{**} , the time when the agent decides whether to adopt the new technology, is taken as given. We initially take t^{**} as given for the purpose of obtaining closed-form solutions for prices, but later we solve for the optimal time t^{**} to adopt the new technology.

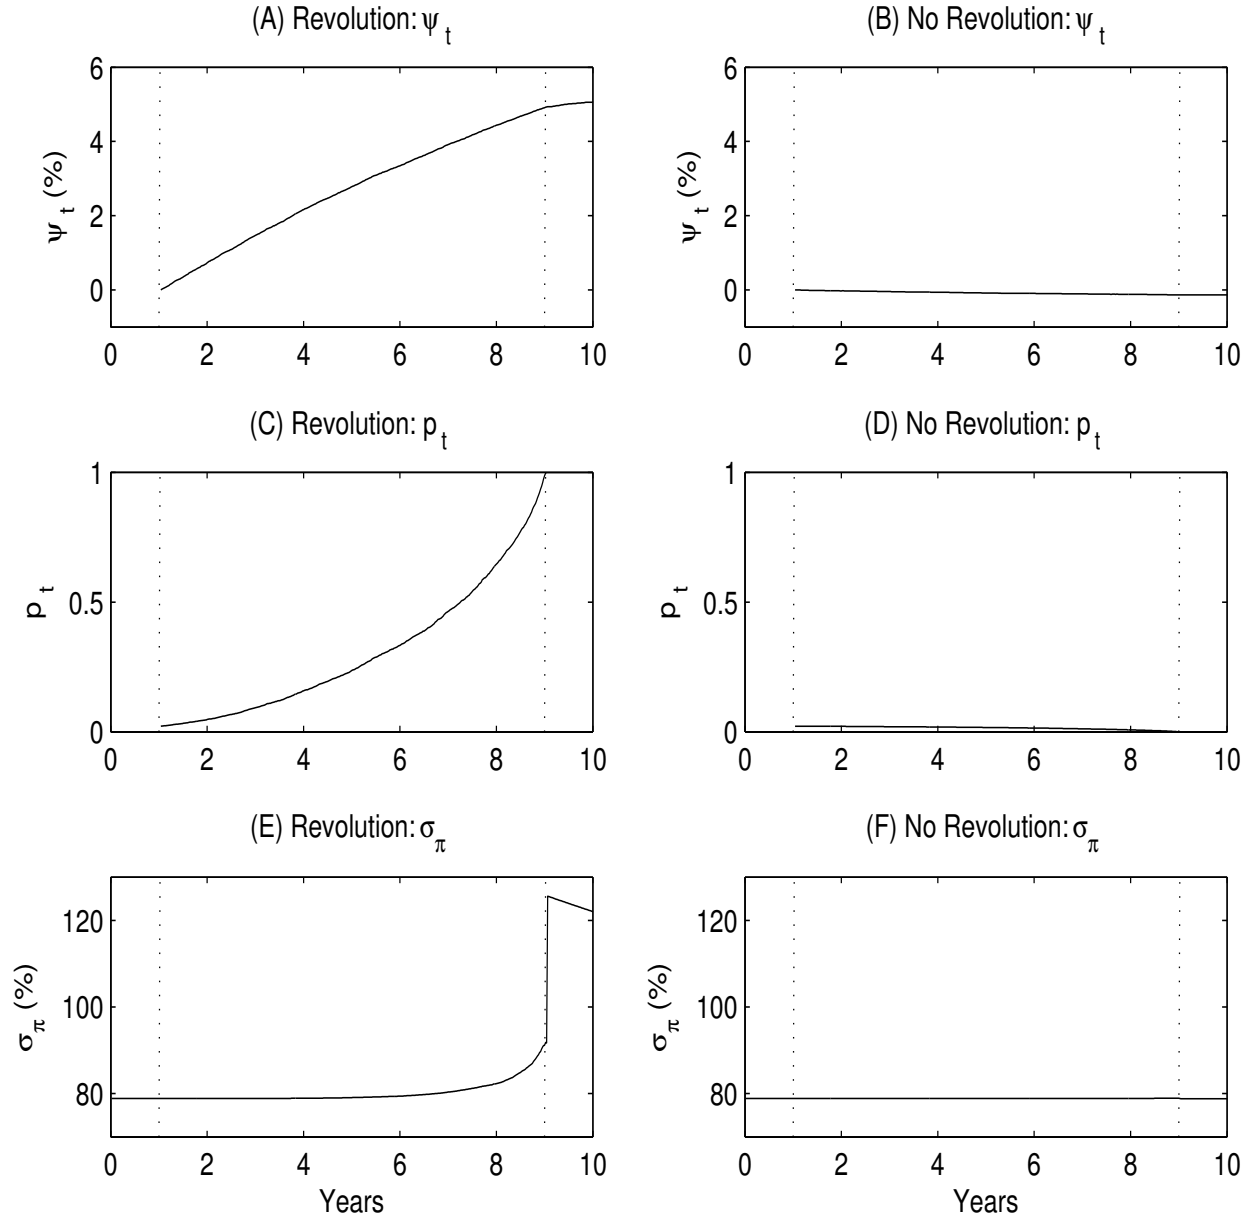


Figure 2. Average $\widehat{\psi}_t$, p_t , $\sigma_{\pi,t}$ in Simulations. The left panels plot the perceived productivity gain $\widehat{\psi}_t$ (Panel A), the adoption probability p_t (Panel C), and the volatility of the stochastic discount factor $\sigma_{\pi,t}$ (Panel E), averaged across all simulations in which the new technology was adopted at time t^{**} ($p_{t^{**}} = 1$). The right panels (B, D, and F) plot the same quantities but the average is taken across all simulations in which the new technology was not adopted at time t^{**} ($p_{t^{**}} = 0$). In each panel, the first vertical line denotes $t^* = 1$, the time when the new technology becomes available, and the second vertical line denotes $t^{**} = 9$, the time at which the agent decides whether to adopt the technology on a large scale. All parameters are given in Table 1.

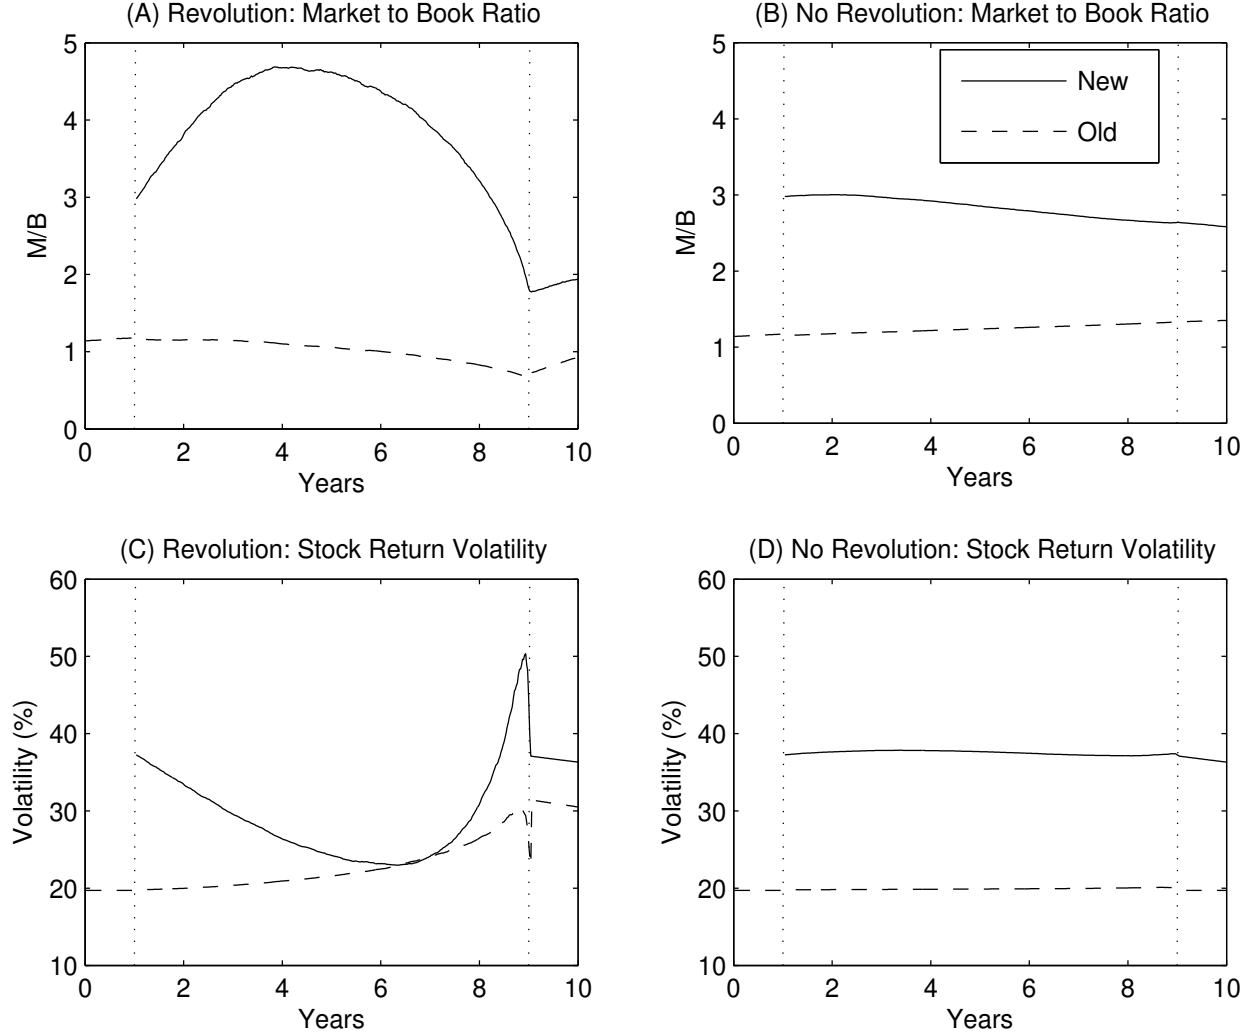


Figure 3. Average M/B and Volatility in Simulations. Panel A plots the path of the market-to-book ratio of the new economy (solid line) and old economy (dashed line) averaged across all simulations in which the new technology was adopted at time t^{**} ($p_{t^{**}} = 1$). Panel B is an equivalent of Panel A, except that the averages are computed across all simulations in which the new technology was not adopted at time t^{**} ($p_{t^{**}} = 0$). Panels C and D are equivalents of Panels A and B, respectively, with M/B replaced by the volatility of stock returns. In each panel, the first vertical line denotes $t^* = 1$, the time when the new technology becomes available, and the second vertical line denotes $t^{**} = 9$, the time at which the agent decides whether to adopt the technology on a large scale. All parameters are given in Table 1.

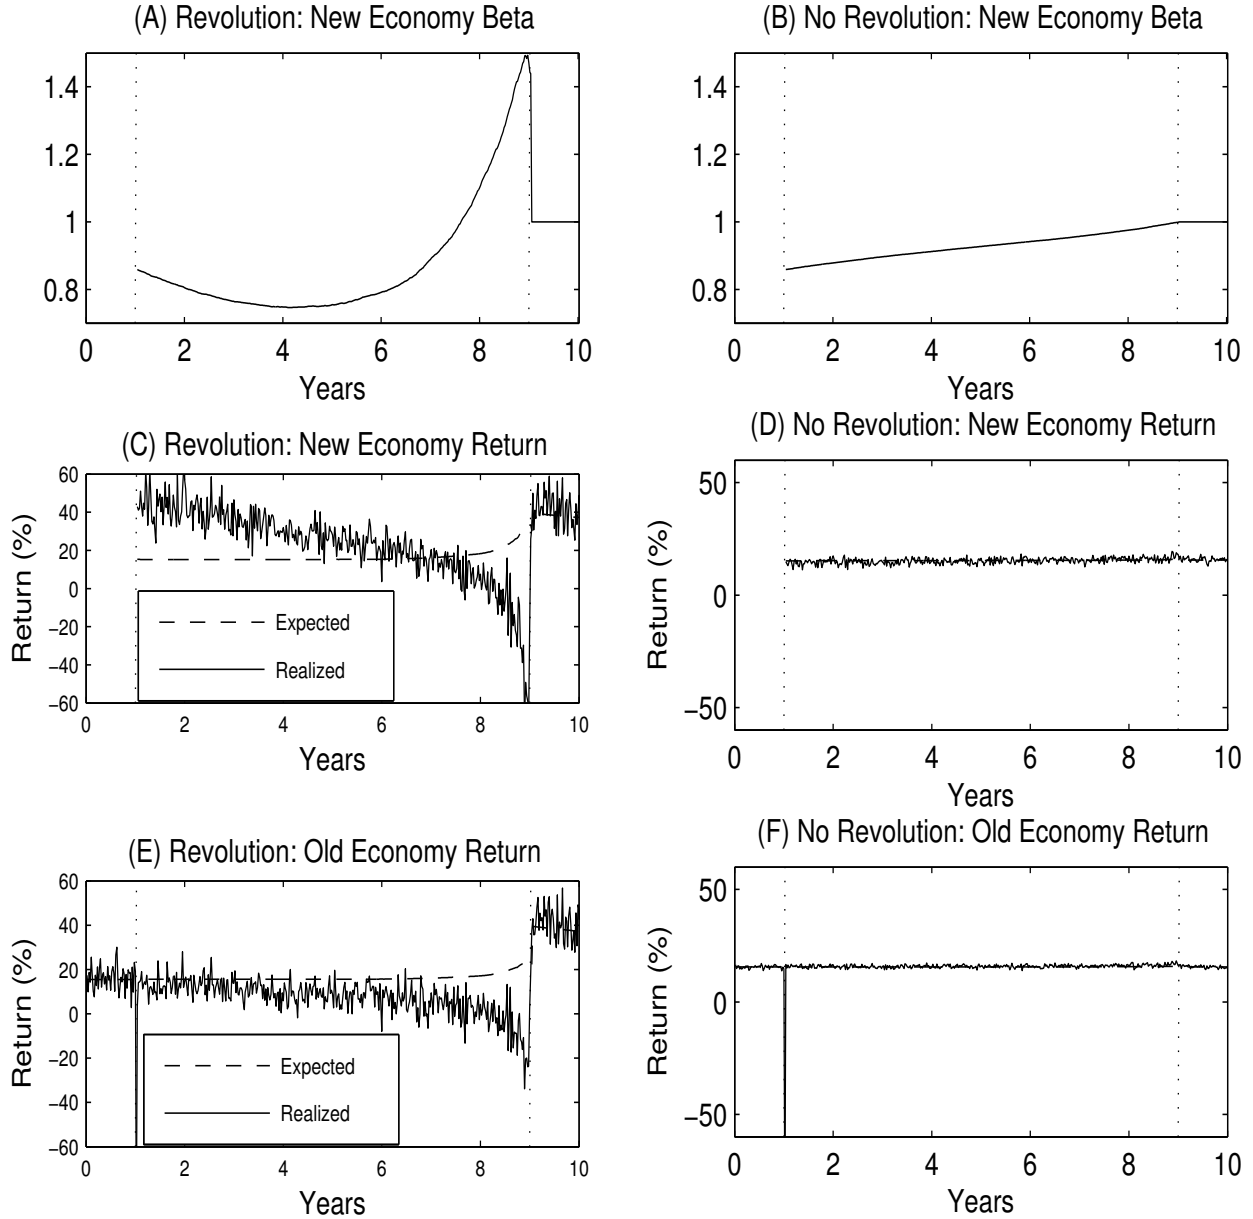


Figure 4. Beta and Average Stock Return in Simulations. The left panels plot the beta of the new economy (Panel A), the realized stock return (solid line) and the expected stock return (dashed line) for the new economy (Panel C) and the old economy (Panel E), averaged across all simulations in which the new technology was adopted at time t^{**} ($p_{t^{**}} = 1$). The right panels (B, D and F) plot the same quantities but the average is taken across all simulations in which the new technology was not adopted at time t^{**} ($p_{t^{**}} = 0$). In each panel, the first vertical line denotes $t^* = 1$, the time when the new technology becomes available, and the second vertical line denotes $t^{**} = 9$, the time at which the agent decides whether to adopt the technology on a large scale. All parameters are given in Table 1.




Figure 5. Average M/B in Simulated Revolutions: Sensitivity Analysis. All four panels plot the paths of the market-to-book ratio of the new economy (solid line) and old economy (dashed line) averaged across all simulations in which the new technology was adopted at time t^{**} ($p_{t^{**}} = 1$). All parameters are given in Table 1, except for one change that varies across the panels. In Panel A, the risk aversion $\gamma = 3$ instead of the benchmark case $\gamma = 4$. In Panel B, the conversion cost $\kappa = 0$ instead of the benchmark case $\kappa = 0.1$. In Panel C, the uncertainty $\sigma_{t^*} = 0.08$ instead of the benchmark case $\sigma_{t^*} = 0.04$. In Panel D, the time until the adoption $t^{**} - t^* = 6$ instead of the benchmark case $t^{**} - t^* = 8$ years. In each panel, the first vertical line denotes t^* , the time when the new technology becomes available, and the second vertical line denotes t^{**} , the time at which the agent decides whether to adopt the technology on a large scale.

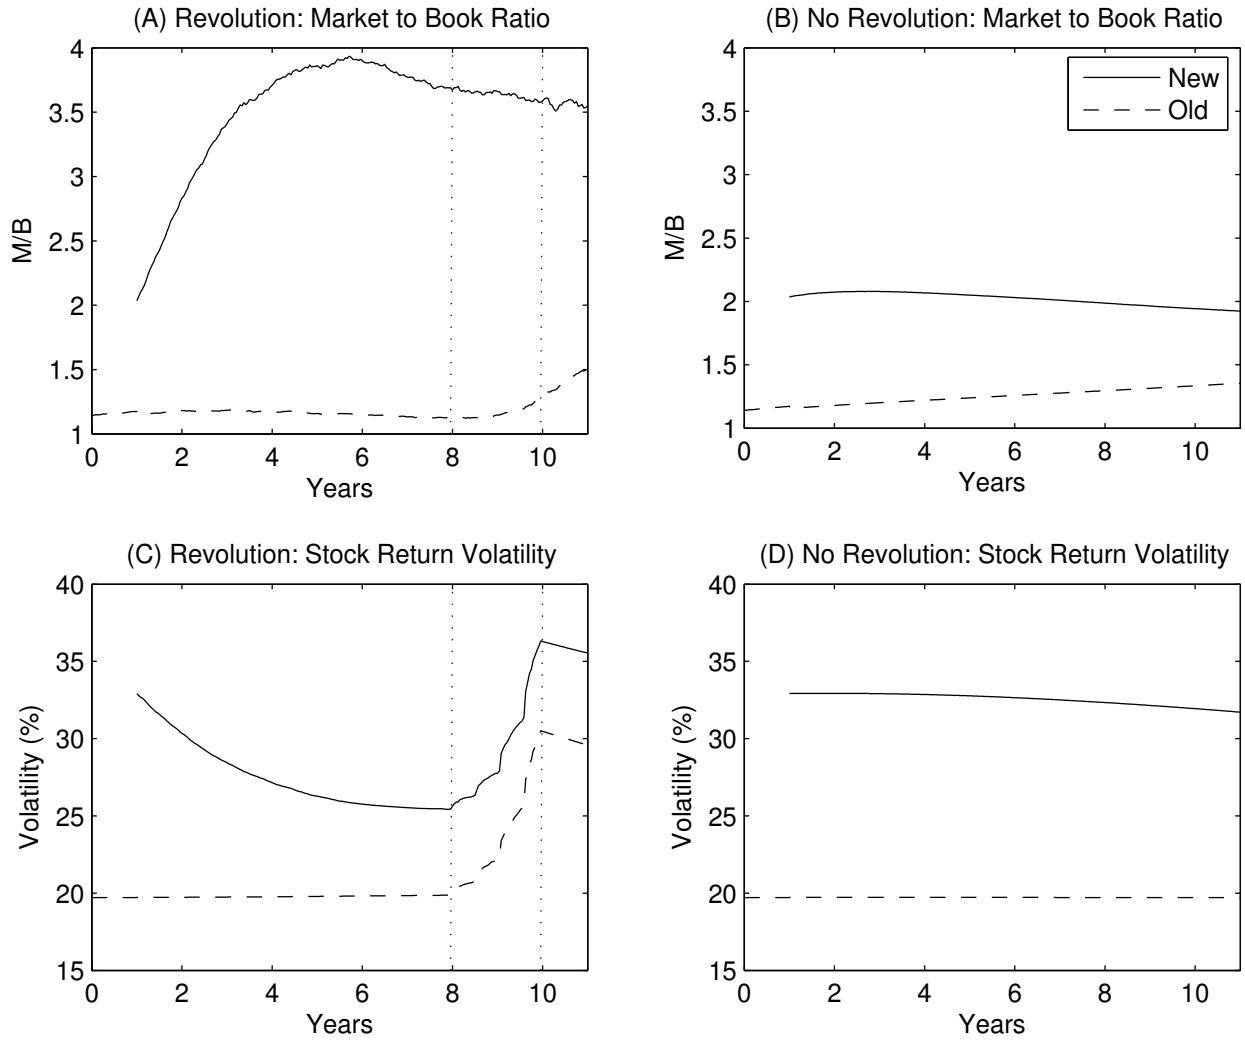


Figure 6. Average M/B and Volatility in Simulations with Optimal Adoption Time. Panel A plots the path of the market-to-book ratio of the new economy (solid line) and old economy (dashed line) averaged across all simulations in which the new technology was adopted at an optimally chosen time t^{**} between years 8 and 10. Panel B is an equivalent of Panel A, except that the averages are computed across all simulations in which the new technology was never adopted. Panels C and D are equivalents of Panels A and B, respectively, with M/B replaced by the volatility of stock returns. All parameters are in Table 1.




Figure 7. The Internet Revolution: Theory. This figure plots the model-implied dynamics of selected quantities in a revolution characterized by fast adoption of the new technology (such as the internet revolution). All quantities are averages computed across all simulations that led to a revolution (i.e., adoption at time $t^{**} = 6$). Panel A reports the market beta of the new economy. Panel B plots the market volatility of the new economy (solid) and old economy (dashed). Panel C plots the market values of the new economy (solid) and old economy (dashed). Panel D plots the productivity of the old economy.




Figure 8. The Internet Revolution: Data. Panel A plots the market beta of the NASDAQ index, with a two-standard-error confidence band. The beta in day t is the slope coefficient from the regression of daily NASDAQ returns on NYSE/AMEX returns over the 500 trading days (i.e., about two years) immediately preceding day t (i.e., days $t - 499$ through t). Panel B plots the standard deviations of returns on the NASDAQ and NYSE/AMEX indices, also computed from daily data over the 500 trading days immediately preceding day t . Panel C plots the value of \$1 invested in January 1992 in the NASDAQ and NYSE/AMEX indices, with dividend reinvestment. Panel D plots a three-year moving average of multifactor productivity growth in the private business sector (in year t , we plot the average annual productivity growth in years $t - 2$, $t - 1$, and t). The vertical dotted line marks October 2002.




Figure 9. Total Rail Consumption in the United States. The figure plots the number of track-miles of rails laid each year in the U.S., as estimated by Fogel (1964, p.174). A track-mile of rails is defined as one half of the length of the rails in a mile of single track. The total includes rails used in the construction of new track as well as in the replacement of worn-out rails.

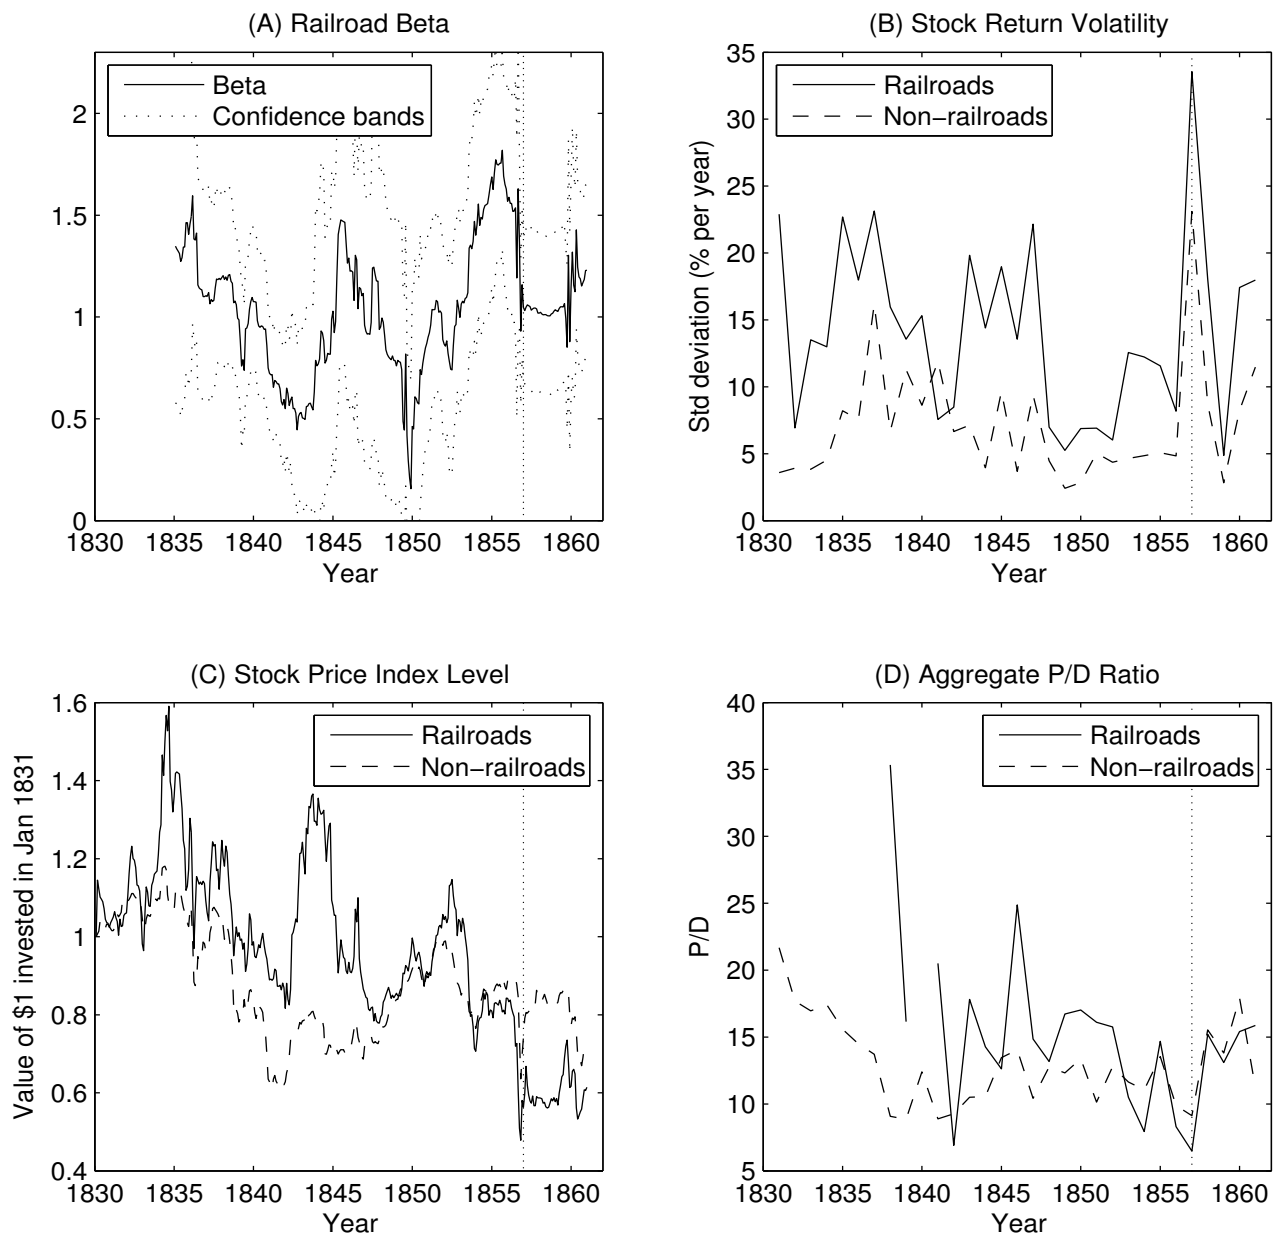


Figure 10. The Railroad Revolution: Data. Panel A plots the beta of the railroad index, with a two-standard-error confidence band. The beta in month t is the slope coefficient from the regression of the most recent 36 monthly railroad returns (in months $t - 35$ through t) on non-railroad returns. The monthly railroad (non-railroad) index returns are computed as price-weighted averages of monthly capital gains across all railroad (non-railroad) stocks. Panel B plots the standard deviation of returns in the railroad and non-railroad industries. Each year, this standard deviation is computed across all monthly price-weighted average industry returns in the given year. Panel C plots the stock price index level, obtained by cumulating monthly capital gains to \$1 invested in January 1831 in the railroad and non-railroad indices. Panel D plots the aggregate price-to-dividend ratio for the railroad and non-railroad industries. Each year, this ratio is computed as the sum of year-end prices divided by the sum of dividends paid in that year, summing across all railroad (non-railroad) stocks with valid price and dividend data. The vertical dotted line marks 1857.

Appendix.

The Appendix contains technical material omitted from the text. The proofs of all results are available in the Technical Appendix, which is downloadable from the authors' websites.

Lemma A1: Suppose the prior distribution of ψ at time t^* is normal, $\psi \sim N(0, \hat{\sigma}_{t^*}^2)$. Then the posterior distribution of ψ at time t , $t^* < t < t^{**}$, conditional on $\mathcal{F}_t = \{(\rho_\tau^N, \rho_\tau) : t^* \leq \tau \leq t\}$ is also normal, $\psi | \mathcal{F}_t \sim N(\hat{\psi}_t, \hat{\sigma}_t^2)$, where the posterior mean $\hat{\psi}_t$ follows the process

$$d\hat{\psi}_t = \hat{\sigma}_t^2 \frac{\phi}{\sigma_{N,1}} d\tilde{Z}_{1,t}, \quad (17)$$

and the posterior variance $\hat{\sigma}_t^2$ is given by

$$\hat{\sigma}_t^2 = \frac{1}{(\hat{\sigma}_{t^*}^2)^{-2} + \left(\frac{\phi}{\sigma_{N,1}}\right)^2 (t - t^*)}. \quad (18)$$

Moreover, the productivity processes can be rewritten as

$$d\rho_t = \phi(\bar{\rho} - \rho_t) dt + \sigma d\tilde{Z}_{0,t} \quad (19)$$

$$d\rho_t^N = \phi(\bar{\rho} + \hat{\psi}_t - \rho_t^N) dt + \sigma_{N,0} d\tilde{Z}_{0,t} + \sigma_{N,1} d\tilde{Z}_{1,t}, \quad (20)$$

where the Brownian motions $(\tilde{Z}_{0,t}, \tilde{Z}_{1,t})$, which capture the agent's expectation errors, follow

$$\begin{pmatrix} d\tilde{Z}_{0,t} \\ d\tilde{Z}_{1,t} \end{pmatrix} = \begin{pmatrix} \sigma & 0 \\ \sigma_{N,0} & \sigma_{N,1} \end{pmatrix}^{-1} \begin{pmatrix} d\rho_t \\ d\rho_t^N - E_t \left[\begin{matrix} d\rho_t \\ d\rho_t^N \end{matrix} \right] \end{pmatrix}. \quad (21)$$

Lemma A1 follows from Theorem 10.3 in Liptser and Shiryaev (1977).

Re Proposition 1: Define the value function at time t as

$$V(B_t, \rho_t, \hat{\psi}_t, \hat{\sigma}_t, t; T) = E_t \left[\frac{W_T^{1-\gamma}}{1-\gamma} \right], \quad (22)$$

where ρ_t follows the process in equation (4). Lemma A2 in the Technical Appendix shows:

$$V(B_t, \rho_t, \hat{\psi}_t, \hat{\sigma}_t, t; T) = E_t \left[\frac{B_T^{1-\gamma}}{1-\gamma} \right] = \frac{B_t^{1-\gamma}}{1-\gamma} e^{A_0(\tau) + (1-\gamma)A_1(\tau)\rho_t + (1-\gamma)A_2(\tau)\hat{\psi}_t + \frac{1}{2}(1-\gamma)^2 A_2(\tau)^2 \hat{\sigma}_t^2}, \quad (23)$$

where $\tau = T - t$, $A_1(\tau)$ and $A_2(\tau)$ are given in the text, and

$$A_0(\tau) = (1-\gamma)\bar{\rho}(\tau - A_1(\tau)) + \frac{\sigma^2(1-\gamma)^2}{2\phi^2} \left\{ \tau + \frac{1 - e^{-2\phi\tau}}{2\phi} - 2\frac{1 - e^{-\phi\tau}}{\phi} \right\}.$$

Since $\gamma > 1$, $V(\cdot)$ is negative, decreasing in $\hat{\sigma}_t^2$, and increasing in B_t . As a result,

$$V(B_{t^*}(1-\kappa), \rho_{t^*}, 0, \hat{\sigma}_{t^*}^2, t^*; T) < V(B_{t^*}, \rho_{t^*}, 0, 0, t^*; T). \quad (24)$$

Proposition 1 follows immediately from equation (24). The left-hand side is the expected utility conditional on adopting the technology at time t^* , and the right-hand side is the expected utility conditional on not adopting the technology at time t^* or any time afterwards.¹⁹ The expected utility from no adoption at time t^* exceeds the right-hand side (and hence also the left-hand side) because it includes the value of the option to adopt after time t^* .

Re Proposition 2: Using (23), equation (8) follows from the optimality condition

$$V\left(B_{t^{**}}(1-\kappa), \rho_{t^{**}}, \widehat{\psi}_{t^{**}}, \widehat{\sigma}_{t^{**}}^2, t^{**}; T\right) \geq V\left(B_{t^{**}}, \rho_{t^{**}}, 0, 0, t^{**}; T\right). \quad (25)$$

Lemma A3: The distribution of $\widehat{\psi}_{t^{**}}$ conditional on $\widehat{\psi}_t$ is normal:

$$\widehat{\psi}_{t^{**}}|\widehat{\psi}_t \sim N\left(\widehat{\psi}_t, \widehat{\sigma}_t^2 - \widehat{\sigma}_{t^{**}}^2\right).$$

In addition, $p_t \equiv \text{Prob}\left(\widehat{\psi}_{t^{**}} > \underline{\psi}|\widehat{\psi}_t\right) = 1 - \mathcal{N}\left(\underline{\psi}; \widehat{\psi}_t, \widehat{\sigma}_t^2 - \widehat{\sigma}_{t^{**}}^2\right)$, where $\mathcal{N}(x; a, s^2) \equiv \int_{-\infty}^x (2\pi s^2)^{-1/2} \exp(y-a)^2/(2s^2) dy$ is the c.d.f. of a normal distribution, $N(a, s^2)$.

Re Proposition 3: Define the value function at time t , $t^* \leq t < t^{**}$, as

$$\mathcal{V}\left(B_t, \rho_t, \widehat{\psi}_t, \widehat{\sigma}_t^2, t; T\right) = E_t \left\{ \max_{\{\text{yes, no}\}} E_{t^{**}} \left[\frac{W_T^{1-\gamma}}{1-\gamma} \right] \right\}, \quad (26)$$

where the maximization involves choosing whether or not to adopt the new technology at time t^{**} , following Proposition 2.²⁰ Lemma A4 in the Technical Appendix shows:

$$\mathcal{V}\left(B_t, \rho_t, \widehat{\psi}_t, \widehat{\sigma}_t^2, t; T\right) = \frac{B_t^{1-\gamma}}{1-\gamma} \{(1-p_t) G_t^{no} + p_t G_t^{yes}\}, \quad (27)$$

where

$$G_t^{no} = e^{A_0(\tau) + (1-\gamma)A_1(\tau)\rho_t} \quad (28)$$

$$G_t^{yes} = G_t^{no} (1-\kappa)^{1-\gamma} R_t e^{(1-\gamma)A_2(\tau^{**})\widehat{\psi}_t + \frac{1}{2}(1-\gamma)^2 A_2(\tau^{**})^2 \widehat{\sigma}_t^2} \quad (29)$$

$$R_t = \frac{1 - \mathcal{N}\left(\underline{\psi}; \widehat{\psi}_t + (1-\gamma)A_2(\tau^{**})(\widehat{\sigma}_t^2 - \widehat{\sigma}_{t^{**}}^2), \widehat{\sigma}_t^2 - \widehat{\sigma}_{t^{**}}^2\right)}{1 - \mathcal{N}\left(\underline{\psi}; \widehat{\psi}_t, \widehat{\sigma}_t^2 - \widehat{\sigma}_{t^{**}}^2\right)} < 1. \quad (30)$$

Using (23) and (27), we prove Proposition 3 in the Technical Appendix by showing that expected utility is higher when experimentation takes place:

$$\mathcal{V}\left(B_{t^*}, \rho_{t^*}, 0, \widehat{\sigma}_{t^*}^2, t^*; T\right) > V\left(B_{t^*}, \rho_{t^*}, 0, 0, t^*; T\right). \quad (31)$$

¹⁹On the right hand side, V is evaluated at $\widehat{\psi}_{t^*} = \widehat{\sigma}_{t^*}^2 = 0$. If the agent decides not to adopt the new technology, ρ_t follows the process in equation (3), which is equivalent to equation (4) when $\psi = 0$.

²⁰Note the difference between the value functions \mathcal{V} in equation (26) and V in equation (23). Whereas \mathcal{V} includes the value of the option to adopt the new technology at the future time t^{**} , V does not include such option value because it applies to settings in which the adoption decision has already been made.

Re Proposition 4: The functions \tilde{G}_t^{no} and \tilde{G}_t^{yes} are given by

$$\tilde{G}_t^{no} = E_t \left[\left(\frac{B_T}{B_t} \right)^{-\gamma} \mathbb{1}_{\{\hat{\psi}_{t^{**}} < \underline{\psi}\}} \right] = e^{\bar{A}_0(\tau) - \gamma A_1(\tau) \rho_t} \quad (32)$$

$$\tilde{G}_t^{yes} = E \left[\left(\frac{B_T}{B_t} \right)^{-\gamma} \mathbb{1}_{\{\hat{\psi}_{t^{**}} \geq \underline{\psi}\}} \right] = \tilde{G}_t^{no} (1 - \kappa)^{-\gamma} \tilde{R}_t e^{-\gamma A_2(\tau^{**}) \hat{\psi}_t + \frac{1}{2} \gamma^2 A_2(\tau^{**})^2 \hat{\sigma}_t^2}, \quad (33)$$

where

$$\tilde{R}_t = \frac{1 - \mathcal{N} \left(\underline{\psi}; \hat{\psi}_t - \gamma A_2(\tau^{**}) (\hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2), \hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2 \right)}{1 - \mathcal{N} \left(\underline{\psi}; \hat{\psi}_t, \hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2 \right)} < 1 \quad (34)$$

$$\bar{A}_0(\tau) = -\gamma \bar{\rho}(\tau - A_1(\tau)) + \frac{\sigma^2 \gamma^2}{2 \phi^2} \left\{ \tau + \frac{1 - e^{-2\phi\tau}}{2\phi} - 2 \frac{1 - e^{-\phi\tau}}{\phi} \right\}. \quad (35)$$

Re Corollary 1: Let $\tilde{p}_t = 1 - \mathcal{N} \left(\underline{\psi}; \hat{\psi}_t - \gamma A_2(\tau^{**}) (\hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2), \hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2 \right)$. Then

$$S_{\pi,t} = \frac{\left(\gamma A_2(\tau^{**}) - \frac{1}{\tilde{p}_t} \frac{\partial \tilde{p}_t}{\partial \hat{\psi}_t} \right) \tilde{G}_t^{yes} + \frac{\partial \tilde{p}_t}{\partial \hat{\psi}_t} \tilde{G}_t^{no}}{(1 - p_t) \tilde{G}_t^{no} + p_t \tilde{G}_t^{yes}}. \quad (36)$$

Re Proposition 5: The functions \tilde{G}_t^{no} , \tilde{G}_t^{yes} , G_t^{no} and G_t^{yes} are presented above. The functions K_t^{no} , and K_t^{yes} are given by

$$K_t^{no} \equiv E_t \left[\left(\frac{B_T}{B_t} \right)^{-\gamma} \frac{B_T^N}{B_t^N} \mathbb{1}_{\{\hat{\psi}_{t^{**}} < \underline{\psi}\}} \right] = K_t R_{L,t}^N$$

$$K_t^{yes} \equiv E_t \left[\left(\frac{B_T}{B_t} \right)^{-\gamma} \frac{B_T^N}{B_t^N} \mathbb{1}_{\{\hat{\psi}_{t^{**}} \geq \underline{\psi}\}} \right] = (1 - \kappa)^{-\gamma} K_t^N R_{H,t}^N$$

where

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

$$K_t^N = K_t e^{-\gamma A_2(\tau^{**}) \hat{\psi}_t + \frac{1}{2} \gamma A_2(\tau^{**}) (\gamma A_2(\tau^{**}) - 2A_2(\tau)) \hat{\sigma}_t^2}$$

$$R_{L,t}^N = \frac{\mathcal{N} \left(\underline{\psi}; \hat{\psi}_t + \sigma_{y\hat{\psi}}^L, \hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2 \right)}{\mathcal{N} \left(\underline{\psi}; \hat{\psi}_t, \hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2 \right)} \quad \text{with} \quad \sigma_{y\hat{\psi}}^L = A_2(\tau) \hat{\sigma}_t^2 - A_2(\tau^{**}) \hat{\sigma}_{t^{**}}^2$$

$$R_{H,t}^N = \frac{1 - \mathcal{N} \left(\underline{\psi}; \hat{\psi}_t + \sigma_{y\hat{\psi}}^H, \hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2 \right)}{1 - \mathcal{N} \left(\underline{\psi}; \hat{\psi}_t, \hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2 \right)} \quad \text{with} \quad \sigma_{y\hat{\psi}}^H = \sigma_{y\hat{\psi}}^L - \gamma A_2(\tau^{**}) (\hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2)$$

$$C_0(\tau) = (1 - \gamma) \bar{\rho}(\tau - A_1(\tau)) + \frac{1}{2\phi^2} \left\{ \tau + \frac{1 - e^{-2\phi\tau}}{2\phi} - 2 \frac{1 - e^{-\phi\tau}}{\phi} \right\} (\gamma^2 \sigma^2 - 2\gamma \sigma_{N,0} \sigma + (\sigma_{N,0}^2 + \sigma_{N,1}^2)).$$

Corollary A1: For any $t \in [t^*, t^{**})$, the stock return processes are given by

$$\frac{dM_t}{M_t} = \mu_{M,t} dt + \sigma_{M,t}^0 d\tilde{Z}_t^0 + \sigma_{M,t}^1 d\tilde{Z}_t^1 \quad \text{and} \quad \frac{dM_t^N}{M_t^N} = \mu_{M,t}^N dt + \sigma_{M,t}^{N,0} d\tilde{Z}_t^0 + \sigma_{M,t}^{N,1} d\tilde{Z}_t^1,$$

where expected returns are equal to the return covariances with $d\pi_t/\pi_t$,

$$\mu_{M,t} = -\sigma_{M,t}^0 \sigma_{\pi,t}^0 - \sigma_{M,t}^1 \sigma_{\pi,t}^1; \quad \mu_{M,t}^N = -\sigma_{M,t}^{N,0} \sigma_{\pi,t}^0 - \sigma_{M,t}^{N,1} \sigma_{\pi,t}^1, \quad (37)$$

and the components of the return volatilities are

$$\sigma_{M,t}^0 = A_1(\tau) \sigma; \quad \sigma_{M,t}^1 = (S_{M,t} + S_{\pi,t}) \hat{\sigma}_t^2 \frac{\phi}{\sigma_{N,1}} \quad (38)$$

$$\sigma_{M,t}^{N,0} = A_1(\tau) \sigma_{N,0}; \quad \sigma_{M,t}^{N,1} = A_1(\tau) \sigma_{N,1} + (S_{M,t}^N + S_{\pi,t}) \hat{\sigma}_t^2 \frac{\phi}{\sigma_{N,1}}, \quad (39)$$

with $S_{M,t}$ and $S_{M,t}^N$ given by

$$S_{M,t} = \frac{-\frac{\partial p_t}{\partial \psi_t} G_t^{no} + \left((1-\gamma) A_2(\tau^{**}) + \frac{1}{p_t} \frac{\partial \bar{p}_t}{\partial \psi_t} \right) G_t^{yes}}{(1-p_t) G_t^{no} + p_t G_t^{yes}} \quad (40)$$

$$S_{M,t}^N = \frac{\left(A_2(\tau) + \frac{1}{p_{L,t}^N} \frac{\partial p_{L,t}^N}{\partial \psi} \right) K_t^{no} + \left((A_2(\tau) - \gamma A_2(\tau^{**})) + \frac{1}{p_{H,t}^N} \frac{\partial p_{H,t}^N}{\partial \psi} \right) K_t^{yes}}{(1-p_t) K_t^{no} + p_t K_t^{yes}} \quad (41)$$

$$\bar{p}_t = 1 - \mathcal{N}\left(\underline{\psi}; \hat{\psi}_t + (1-\gamma) A_2(\tau^{**}) (\hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2), \hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2\right) \quad (42)$$

$$p_{L,t}^N = \mathcal{N}\left(\underline{\psi}; \hat{\psi}_t + \sigma_{y\hat{\psi}}^L, \hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2\right) \quad (43)$$

$$p_{H,t}^N = 1 - \mathcal{N}\left(\underline{\psi}; \hat{\psi}_t + \sigma_{y\hat{\psi}}^H, \hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2\right). \quad (44)$$

Re Proposition 6.: The functions h_{old} and h_{new} are given by

$$h_{old} = -\tilde{\kappa} + A_2(\tau^{**}) \hat{\psi}_t + \frac{1}{2} (1-2\gamma) A_2(\tau^{**})^2 \hat{\sigma}_t^2 \quad (45)$$

$$- \log \left(\frac{1 - \mathcal{N}\left(\underline{\psi}; \hat{\psi}_t - \gamma A_2(\tau^{**}) (\hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2), \hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2\right)}{1 - \mathcal{N}\left(\underline{\psi}; \hat{\psi}_t + (1-\gamma) A_2(\tau^{**}) (\hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2), \hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2\right)} \right) \quad (46)$$

$$h_{new} = -\gamma A_2(\tau^{**}) A_2(\tau) \hat{\sigma}_t^2 - \log \left(\frac{\mathcal{N}\left(\underline{\psi}; \hat{\psi}_t + \sigma_{y\hat{\psi}}^L, \hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2\right)}{\mathcal{N}\left(\underline{\psi}; \hat{\psi}_t, \hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2\right)} \right) \quad (47)$$

$$- \log \left(\frac{1 - \mathcal{N}\left(\underline{\psi}; \hat{\psi}_t - \gamma A_2(\tau^{**}) (\hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2), \hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2\right)}{1 - \mathcal{N}\left(\underline{\psi}; \hat{\psi}_t - \gamma A_2(\tau^{**}) (\hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2) + \sigma_{y\hat{\psi}}^L, \hat{\sigma}_t^2 - \hat{\sigma}_{t^{**}}^2\right)} \right). \quad (48)$$

Re Corollary 2: In the corollary, $\bar{C}_0(\tau^{**}) = C_0(\tau^{**}) - \bar{A}_0(\tau^{**})$.

Re Optimal Adoption Time: The value function is

$$\mathcal{V} \left(B_t, \rho_t, \widehat{\psi}_t, \widehat{\sigma}_t^2, t; T \right) = E_t \left\{ \max_{t^{**}} E_{t^{**}} \left[\frac{W_T^{1-\gamma}}{1-\gamma} \right] \right\} = B_t^{1-\gamma} e^{(1-\gamma)A_1(T-t)\rho_t} \mathcal{V}_2 \left(\widehat{\psi}_t, t; T \right), \quad (49)$$

where $\mathcal{V}_2 \left(\widehat{\psi}_t, t; T \right)$ satisfies the partial differential equation

$$0 = \frac{\partial \mathcal{V}_2}{\partial t} + \left((1-\gamma) A_1(T-t)\phi\bar{\rho} + \frac{1}{2} (1-\gamma)^2 A_1(T-t)^2 \sigma^2 \right) \mathcal{V}_2 + \frac{1}{2} \frac{\partial^2 \mathcal{V}_2}{\partial \widehat{\psi}^2} \left(\widehat{\sigma}_t^2 \frac{\phi}{\sigma_{N,1}} \right)^2,$$

with the boundary conditions $\mathcal{V}_2 \left(\widehat{\psi}_T, T \right) = \frac{1}{1-\gamma}$ if $t^{**} > T$ and

$$\mathcal{V}_2 \left(\widehat{\psi}_t, t; T \right) \geq \frac{(1-\kappa)^{1-\gamma}}{1-\gamma} e^{A_0(\tau) + (1-\gamma)A_2(\tau)\widehat{\psi}_t + \frac{1}{2}(1-\gamma)^2 A_2(\tau)^2 \widehat{\sigma}_t^2},$$

with an equality at $t = t^{**}$. We solve the PDE by using the finite difference method.

Alternative Decentralized Model.

Below, we develop a model of a competitive economy in which firms independently decide whether to adopt the new technology while maximizing their own market values. We show that, in equilibrium, this decentralized economy produces the same adoption rules and same stock price dynamics (including “bubbles”) as the social planner’s problem in Section 2.

Consider a finite-horizon economy with a continuum of identical firms and investors. All investors maximize utility as in Section 2. Before time t^* , all firms employ the same “old technology.” The productivity $\rho_{i,t}$ of any firm i that uses the old technology follows

$$d\rho_{i,t} = \phi \left(g \left(s_t^O \right) - \rho_{i,t} \right) dt + \sigma dZ_{0,t}, \quad (50)$$

where $s_t^O \in [0, 1]$ is the fraction of firms using the old technology at time t , and $g'(s) > 0$. The function $g(s)$ captures “network externalities”: the average productivity of a technology increases as the fraction of firms using this technology increases. Denote $g(1) = \bar{\rho}$ and $g(0) = \bar{\rho}_0$. We refer to $\bar{\rho} - \bar{\rho}_0 > 0$ as the network externality gain. Since $s_t^O = 1$ for $t \leq t^*$, the productivity of all firms before time t^* is identical to equation (3) in Section 2.

At time t^* a new unique firm N appears, equipped with a new technology. The productivity process of firm N is given by

$$d\rho_{N,t} = \phi \left(\bar{\rho} + \psi - \rho_{N,t} \right) dt + \sigma_{N,0} dZ_{0,t} + \sigma_{N,1} dZ_{1,t}, \quad (51)$$

where ψ is an unobservable productivity gain from using the new technology, as in Section 2. The prior distribution for ψ at time t^* is given in equation (5). After time t^* , all firms learn about ψ by observing $\rho_{N,t}$ and $\rho_{i,t}$ for all i .

Firms using the old technology can adopt the new technology either immediately at time t^* , or at a given later time t^{**} , or never. When they adopt, they incur a proportional

conversion cost κ . If the fraction $s_t^N = 1 - s_t^O$ of all firms use the new technology at time t , the productivity of each firm i that uses the new technology is given by

$$d\rho_{i,t} = \phi(g(s_t^N) + \psi - \rho_{i,t}) dt + \sigma dZ_{0,t}. \quad (52)$$

If all firms adopt the new technology ($s_t^N = 1$), then the process (52) is identical to equation (4). Aggregate productivity (obtained by aggregating across identical firms i) therefore follows the same process (4) as it does in the social planner's problem. If no firm adopts the new technology ($s_t^N = 0$), then the aggregate economy uses the old technology. Aggregate productivity is then given by the process (50) with $s_t^O = 1$, which is identical to equation (3). Once again, aggregate productivity follows the same process as in the planner's problem. We show below in Proposition D2 that, in equilibrium at time t^{**} , either $s_{t^{**}}^N = 1$ or $s_{t^{**}}^N = 0$. Therefore, the pricing kernel for $t \geq t^{**}$ is the same as in the planner's problem, and it depends on whether $s_{t^{**}}^N = 1$ or $s_{t^{**}}^N = 0$.

The market value of each firm at time t^{**} depends not only on whether the firm adopts the new technology, but also on the adoption decisions of all other firms, because those decisions affect the pricing kernel. We denote the market value of firm i at time t^{**} as $M_{i,t^{**}}^{s^N=1,\text{yes}}$, $M_{i,t^{**}}^{s^N=1,\text{no}}$, $M_{i,t^{**}}^{s^N=0,\text{yes}}$, or $M_{i,t^{**}}^{s^N=0,\text{no}}$, where "yes" or "no" indicates whether firm i adopts or not, and " $s^N = 0$ " or " $s^N = 1$ " indicates whether other firms adopt or not. Closed-form expressions for these four values are given in the Technical Appendix. Taking the choice of other firms as given ($s^N = 1$ or $s^N = 0$), each firm adopts the new technology if doing so maximizes its market value. That is, firm i adopts if and only if

$$M_{i,t^{**}}^{s^N=1,\text{yes}} > M_{i,t^{**}}^{s^N=1,\text{no}} \quad \text{if everybody else adopts} \quad (53)$$

$$M_{i,t^{**}}^{s^N=0,\text{yes}} > M_{i,t^{**}}^{s^N=0,\text{no}} \quad \text{if nobody else adopts} \quad (54)$$

The following proposition characterizes the value-maximizing adoption decision of any individual firm, conditional on the decisions of all other firms at time t^{**} .

Proposition D1: (a) If $s_{t^{**}}^N = 1$ (i.e., all firms adopt the new technology at time t^{**}), then firm i adopts the new technology at time t^{**} if and only if

$$\widehat{\psi}_{t^{**}} \geq \underline{\psi}^{s^N=1} = -\frac{\log(1-\kappa)}{A_2(\tau^{**})} + \frac{1}{2}(2\gamma-1)A_2(\tau^{**})\widehat{\sigma}_{t^{**}}^2 - (\bar{\rho} - \bar{\rho}_0) \quad (55)$$

(b) If $s_{t^{**}}^N = 0$ (i.e., no firm adopts the new technology at time t^{**}), then firm i adopts the new technology at time t^{**} if and only if

$$\widehat{\psi}_{t^{**}} \geq \underline{\psi}^{s^N=0} = -\frac{\log(1-\kappa)}{A_2(\tau^{**})} - \frac{1}{2}A_2(\tau^{**})\widehat{\sigma}_{t^{**}}^2 + (\bar{\rho} - \bar{\rho}_0) \quad (56)$$

In general, $\underline{\psi}^{s^N=1} \neq \underline{\psi}^{s^N=0}$. Both thresholds consist of three terms. The first term, which reflects the conversion cost κ , is the same in both cases. The second term, which reflects uncertainty $\widehat{\sigma}_{t^{**}}$, is positive in (55) but negative in (56). This difference is due to systematic risk. In case (a), all firms adopt the new technology, so the equilibrium stochastic discount

factor is heavily affected by the risk of the new technology. Hence, the new technology carries higher systematic risk than the old technology. If a firm adopts the new technology, its beta increases, pushing its discount rate up and market value down. This undesirable feature of the new technology increases the adoption threshold. In case (b), the new technology carries less systematic risk, so the $\widehat{\sigma}_{t^{**}}$ term reduces the adoption threshold.

The third term reflects the network externality gain, $\bar{\rho} - \bar{\rho}_0 > 0$. In case (a), all other firms have already adopted the new technology, so its average productivity is higher and, as a result, the adoption threshold is lower. In case (b), the new technology is less productive because no other firm has adopted it, so the adoption threshold is higher.

Intuitively, adopting the same technology as other firms has two effects. On one hand, it hurts the firm, because the technology adopted by all other firms carries more systematic risk. On the other hand, it benefits the firm through network externality gains.

We are interested in setting up a competitive environment that supports the social planner's solution. This can be done by a judicious choice of the magnitude of the network externality gains, $\bar{\rho} - \bar{\rho}_0$. When we choose $\bar{\rho}_0$ as

$$\bar{\rho}_0 = \bar{\rho} - \frac{1}{2}\gamma A_2(\tau^{**})\widehat{\sigma}_{t^{**}}^2, \quad (57)$$

the thresholds in Proposition D1 become equal to each other:

$$\underline{\psi}^{s^N=1} = \underline{\psi}^{s^N=0} = \underline{\psi} = -\frac{\log(1-\kappa)}{A_2(\tau^{**})} + \frac{1}{2}(\gamma-1)A_2(\tau^{**})\widehat{\sigma}_{t^{**}}^2.$$

Moreover, both thresholds are now equal to the threshold $\underline{\psi}$ in equation (8). In other words, the adoption thresholds in the competitive problem and the social planner's problem are identical. Still operating under condition (57), we obtain the following proposition.

Proposition D2: If $\widehat{\psi}_{t^{**}} \geq \underline{\psi}$, then $s_{t^{**}}^N = 1$ is a Nash equilibrium. If $\widehat{\psi}_{t^{**}} < \underline{\psi}$, then $s_{t^{**}}^N = 0$ is a Nash equilibrium.

Proposition D2 shows that, in equilibrium, either all firms adopt the new technology at time t^{**} or none of them do.²¹ Moreover, the adoption decision is analogous to that in Proposition 2, with an identical threshold. As a result, the equilibrium stochastic discount factor is the same, and all pricing formulas are identical to those in the planner's problem.

Note that when all firms choose to adopt the new technology at time t^{**} , they increase the equilibrium discount rates and thus decrease their market values. Although each firm maximizes its own market value, the aggregate effect of the firms' adoptions is to depress market values at time t^{**} . The reason is that firms that adopt the new technology do not fully internalize the resulting increases in systematic risk. To summarize, this decentralized model with value-maximizing competitive firms produces the same stock price dynamics as our simpler model with a utility-maximizing social planner.

²¹We can also prove a related result for time t^* : $s_{t^*}^N = 1$ is not a Nash equilibrium but $s_{t^*}^N = 0$ is, under plausible parametric conditions (see Technical Appendix). That is, in equilibrium, no firm adopts the new technology at time t^* , analogous to our Proposition 1 in the social planner's problem.

REFERENCES

- Aghion, Philippe, and Peter Howitt, 1992, A model of growth through creative destruction, *Econometrica* 60, 323–351.
- Atkeson, Andrew, and Patrick J. Kehoe, 2007, Modeling the Transition to a New Economy: Lessons from Two Technological Revolutions, *American Economic Review* 97, 64–88.
- Benmelech, Efraim, 2007, Asset salability and debt maturity: Evidence from 19th century American railroads, *Review of Financial Studies*, forthcoming .
- Bharath, Sreedhar T., and Siva Viswanathan, 2006, Is the Internet bubble consistent with rationality?, Working paper, University of Michigan.
- Brown, Stephen J., William N. Goetzmann, and Stephen A. Ross, 1995, Survival, *Journal of Finance* 50, 853–873.
- Caselli, Francesco, 1999, Technological revolutions, *American Economic Review* 89, 78–102.
- Chari, V.V., and Hugo Hopenhayn, 1991, Vintage human capital, growth, and the diffusion of new technology, *Journal of Political Economy* 99, 1142–1165.
- Cochrane, W., 1979, *The Development of American Agriculture*, Minneapolis.
- Craig, Lee A., and Thomas Weiss, 1993, Agricultural Productivity Growth During the Decade of the Civil War, *Journal of Economic History* 53, 527–548.
- David, Paul A., 1991, Computer and dynamo: The modern productivity paradox in a not-too-distant mirror, In *Technology and Productivity: The Challenge for Economic Policy* (OECD, Paris), 315–347.
- DeMarzo, Peter, Ron Kaniel, and Ilan Kremer, 2007, Technological Innovation and Real Investment Booms and Busts, *Journal of Financial Economics* 85, 735–754.
- Fogel, Robert W., 1964, *Railroads and American Economic Growth: Essays in Econometric History*, Johns Hopkins Press, Baltimore, MD.
- Goetzmann, William N., Roger G. Ibbotson, and Liang Peng, 2001, A new historical database for the NYSE 1815 to 1925, *Journal of Financial Markets* 4, 1–32.
- Greenwood, Jeremy, and Boyan Jovanovic, 1999, The information-technology revolution and the stock market, *American Economic Review Papers and Proceedings* 89, 116–122.
- Hoberg, Gerard, and Gordon Phillips, 2006, Real and financial industry booms and busts, Working paper, University of Maryland.
- Hobijn, Bart, and Boyan Jovanovic, 2001, The Information-technology revolution and the stock market: Evidence, *American Economic Review* 91, 1203–1220.
- Johnson, Timothy C., 2007, Optimal learning and new technology bubbles, *Journal of Monetary Economics* 54, 2486–2511.
- Jovanovic, Boyan, 1982, Selection and the evolution of industry, *Econometrica* 50, 649–670.
- Jovanovic, Boyan, and Glenn M. MacDonald, 1994, The life cycle of a competitive industry, *Journal of Political Economy* 102, 322–347.

- Jovanovic, Boyan, and Yaw Nyarko, 1996, Learning by doing and the choice of technology, *Econometrica* 64, 1299–1310.
- Jovanovic, Boyan, and Peter L. Rousseau, 2003, Two technological revolutions, *Journal of European Economic Association* 1, 419–428.
- Jovanovic, Boyan, and Peter L. Rousseau, 2005, General purpose technologies, NBER Working paper #11093.
- Klein, Maury, 1994, *Unfinished Business: The Railroad in American Life*, University Press of New England.
- Laitner, John, and Dmitriy Stolyarov, 2003, Technological change and the stock market, *American Economic Review* 93, 1240–1267.
- Laitner, John, and Dmitriy Stolyarov, 2004a, Aggregate returns to scale and embodied technical change: Theory and measurement using stock market data, *Journal of Monetary Economics* 51, 191–233.
- Laitner, John, and Dmitriy Stolyarov, 2004b, The role of owned ideas in stock market run-ups, Working paper, University of Michigan.
- Malkiel, Burton G., 1999, *A Random Walk Down Wall Street*, W. W. Norton & Co., NY.
- Manuelli, Rodolfo E., 2003, Technological change, the labor market and the stock market, Working paper, University of Wisconsin-Madison.
- Mazzucato, Mariana, 2002, The PC industry: New economy or early life-cycle?, *Review of Economic Dynamics* 5, 318–345.
- Mazzucato, Mariana, and Massimiliano Tancioni, 2006, Stock price volatility and patent citation dynamics, Working paper, Open University.
- Mishkin, Frederic S., 1991, Asymmetric information and financial crises: A historical perspective, In *Financial Markets and Financial Crises*, edited by R. Glenn Hubbard, University of Chicago Press, Chicago, 69–108.
- Mokyr, Joel, 1990, *The Level of Riches: Technological Creativity and Economic Progress*, Oxford University Press, New York.
- Pástor, Ľuboš, and Pietro Veronesi, 2003, “Stock valuation and learning about profitability,” *Journal of Finance* 58, 1749–1789.
- Pástor, Ľuboš, and Pietro Veronesi, 2006, “Was there a Nasdaq bubble in the late 1990s?” *Journal of Financial Economics* 81, 61–100.
- Perez, Carlota, 2002, *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*, Edward Elgar, Cheltenham, UK.
- Romer, Paul, 1990, Endogenous technological change, *Journal of Political Economy* 98, S71–S102.
- Shiller, Robert, 2000, *Irrational Exuberance*, Princeton University Press, Princeton, NJ.
- Stover, John F., 1961, *American Railroads*, University of Chicago Press.