Commercialization of GPTs and Stock Market Crashes*

Lukasz A. Drozd

Federal Reserve Bank of Philadelphia

lukasz.drozd@phil.frb.org

Preliminary draft — limited circulation

October 31, 2026

Abstract

We study how developing a viable general purpose technology (GPT) can generate a collapse in fundamentally priced equity. Such collapses are usually attributed to finance: they are abrupt, often follow financial triggers, and leave peak valuations looking unjustified. We show that these features can arise without financial frictions or mispricing. Firms accumulate proprietary capabilities that both advance commercialization and determine their share of its payoff. Research by others accelerates payment but dilutes ownership of future profits. In equilibrium, firms withdraw together at a common threshold, so modest adverse news triggers an abrupt collapse of research and equity, and capabilities lost in the pause raise the bar for resuming. Realized receipts then fall short of earlier prices, as they did for the railroad and dot-com leaders. This spiral can delay commercialization far beyond what a planner would choose.

Keywords: general purpose technology; innovation; coordination; commercialization; global games; stock-market crashes.

JEL: D82, E32, G12, O31, O33.

*The views expressed in this paper are those of the author and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

1 Introduction

The build-out of a major general purpose technology (GPT) has repeatedly coincided with a boom and bust in asset values, even when the technology ultimately proved viable and delivered a large, lasting gain in productivity (Perez, 2002). We focus on two defining episodes here: the American railroad expansion that ended in 1873 and the information technology boom around 2000. Railways transformed transportation and dominated it for nearly a century. Yet in both Britain and the United States, their early build-out brought a crash that shook the broader economy. In 1873, the failure of Jay Cooke & Co. set off a panic and the longest contraction in the NBER chronology. The dot-com crash of 2000 followed a strikingly similar pattern. In both cases, severe equity losses and a collapse in investment interrupted the build-out of a technology whose commercial use was still expanding. Similar episodes accompanied British canals and the automobile revolution after World War I.[1]

These crashes are usually attributed to finance, and for understandable reasons. The collapses break abruptly, often around financial events such as the failure of Jay Cooke & Co.; financial distress is visible in every episode; and after the fact, peak valuations look unjustified by fundamentals. The standard account is that a new technology breeds overoptimism, a bubble forms, and financial forces turn its correction into a crash: the financial accelerator (Kiyotaki and Moore, 1997; Bernanke et al., 1999; Brunnermeier and Sannikov, 2014), runs (Diamond and Dybvig, 1983; Goldstein and Pauzner, 2005), and related mechanisms. Fraud, misleading promotion, and leverage are part of these histories too, but in very different measure: British railways were built with private share capital, American railroads with heavy bond finance and land grants, and internet firms largely with equity (Odlyzko, 2010; White, 2011). The crashes nonetheless look alike, and the profitable leaders of the American railroad and technology booms later return much of their peak value to shareholders. In asset prices, a railway or dot-com crash can look much like the Dutch tulip mania of the 1630s, although the technology behind it was real. Why do viable GPTs, which eventually transform the economy, appear equally fragile? Can the real tradeoffs of commercialization alone produce similarly abrupt and nonlinear episodes?

[1] Aviation and radio followed the same path in the 1920s, but their crash came in 1929 and cannot be separated from the general market collapse. Section 2 documents the railroad, automobile, and technology episodes, and Appendix A.4 gives the sources. For the NBER dates, see National Bureau of Economic Research (2026c).

We show that they can. These tradeoffs alone can generate a large collapse in investment and in fundamentally priced equity, together with a delay that coordinated investment would avoid, in a technology that succeeds under every policy. Leverage and other financial forces would amplify such a cycle, but they are not needed to produce it; Appendix A.2.5 studies a funding wedge.

The source of a fundamental fragility is the commercialization process itself, with equity valuations and market dynamics reflecting these fundamental forces in prices and rational innovation decisions. Specifically, the return to developing a GPT depends on when complementary infrastructure, applications, and organizational knowledge make it commercially productive, that is, when the scale of use allows the investment to be recouped at reasonable margins. The model thus captures the productivity J-curve of GPT build-outs (Brynjolfsson et al., 2021): investment in intangible capabilities comes first, and the returns follow only once commercialization is reached. The crash-prone element is that the commercialization date depends on investment across the industry, and not individual investment. A developer therefore invests partly because it expects others to continue. Their withdrawal postpones the common payoff, reducing the return to its own research. A deterioration in development requirements can consequently induce a collective retreat even when the technology's eventual productive potential is unchanged.

The key implication is that fundamental asset pricing can coexist with inefficient investment. Firms may correctly value a delayed stream of receipts while failing to internalize how their research could bring those receipts forward. Coordination would then raise both investment and the value of existing claims. The resulting crash is not a correction of an earlier pricing mistake within the model. It reflects a deterioration in the equilibrium commercialization path.

An early online retailer illustrates the underlying complementarity. Its prospective market depended on investments in logistics, payments, and widespread use by buyers and sellers. Those investments could accelerate its profits without changing the eventual usefulness of online retail. The same distinction between eventual usefulness and the timing of returns arises in artificial intelligence, whose commercial use depends on complementary investments by model developers, infrastructure providers, and the firms that adopt it.

We connect this mechanism to historical evidence from US railroads around 1873 and technology firms around 2000, with supporting comparisons from British railways and automobiles. The evidence es-

tablishes three features the theory seeks to reconcile. Equity prices collapse during the development of technologies whose commercial use continues to expand. Investment retreats persistently: annual railroad construction falls 78.4 percent between 1872 and 1875 and does not regain its earlier flow until 1881. Around the dot-com contraction, equipment expenditure and venture financing fall sharply. Real US business-funded R&D falls 6.4 percent between 2000 and 2002 and remains below its 2000 level through 2005, although this last measure covers the whole economy rather than IT alone.

The third feature is substantial eventual value among firms that were already large and profitable near the valuation peak, even as failures and exits spread through the rest of the industry. We fix leader portfolios using dated profitability information and follow shareholder payouts, funding contributions, failures, acquisitions, and residual claims. At an 8 percent annual nominal discount rate, subsequent receipts cover approximately 71 percent of the initial value of the ten-firm portfolios in both episodes. Composition matters: the five-firm fractions are 87 percent for railroads and 34 percent for technology. For the technology industry as a whole, a dollar placed in the value-weighted Fama–French business-equipment portfolio in January 2000 delivers 125 percent of its value at the same discount rate, although the portfolio lost 80 percent between March 2000 and September 2002. These realized receipts document enduring value through a crash. They do not establish that initial prices were correct, measure investors' required returns, or identify a counterfactual commercialization date. Likewise, the investment and news chronology does not identify the historical contribution of coordination. The evidence motivates the joint occurrence of commercial success, severe valuation losses, and prolonged investment withdrawal.

Our model makes the payoff date an outcome of firms' research decisions. Firms own capabilities: proprietary knowledge that performs commercialization work, supports further research, and determines ownership of the eventual payoff. Research creates capabilities in proportion to the firm's inherited stock, adapting the portfolio scaling of Klette and Kortum (2004). Commercialization occurs when the industry has completed a common requirement, reflecting the complementary development emphasized by Helpman and Trajtenberg (1998). A fixed aggregate payoff is then divided according to capability ownership. Research by other firms brings payment forward but also dilutes existing claims. Strategic complementarity arises when acceleration dominates dilution. It follows from the development technology rather than from an imposed reward for matching others' actions.

Two reinforcing forces amplify adverse news. The first is static informational clustering: firms withdraw together at a common threshold. Each firm receives a private signal of the remaining requirement and must judge both how much work remains and whether others will continue, because their withdrawal postpones the common payoff and lowers the return to its own research. Dispersed information disciplines these expectations. A firm whose signal lies exactly at the withdrawal threshold cannot tell whether its signal is high or low relative to the signals of others, so it regards every participation rate as equally likely; the threshold is the requirement at which research just pays on average over that belief. As signals become precise, the thresholds of all firms converge, and a small change in news moves the industry from nearly full research to nearly full withdrawal. Differences in research costs do not spread this response out. Suppose the most expensive third of firms withdrew first, at news under which the others still research. At its own threshold it would then expect participation between two-thirds and one, more than the cheapest firms would expect at theirs, between zero and one-third. The expensive firms would therefore need worse news to withdraw than the cheap firms, the reverse of the order assumed. Separate thresholds unravel in this way, and firms with different costs withdraw together at a common boundary while holding different beliefs about participation. In the calibrated economy, news worse by 0.4 prior standard deviations lowers costly participation from 97.7 to 2.3 percent.

We establish the mechanism first in a tractable economy with one research decision. The technology succeeds under every feasible investment policy, and expected returns are smooth in requirements and aggregate research. Private signals select a unique withdrawal boundary. As signal noise vanishes, arbitrarily close fundamentals induce full research on one side and full withdrawal on the other. The associated delay and loss of net producer surplus remain strictly positive, while a planner continues research throughout. Introducing different research costs need not smooth this response. Applying the belief accounting of Sákovics and Steiner (2012) and the clustering argument of Drozd and Serrano-Padial (2018), we prove convergence to a common boundary over an explicit range of cost dispersion. Across that boundary, every costly type and the continuing core lose net equity value. The core nevertheless gains a larger eventual ownership share. These results jointly establish selection, inefficient postponement, and losses on fundamentally priced claims in a technology that remains viable.

The second force is the dynamic deterioration of the industry's knowledge base, its stock of capabilities, and the recursive model explains why the retreat persists. Research both performs commercialization work and replenishes the capabilities that make further research productive. When firms pause, the industry completes less work and creates no new capabilities, while part of its existing stock is lost as specialized teams disperse and know-how becomes obsolete. We call this loss depreciation. Its rate, the share of capabilities lost each period, determines how quickly a pause erodes the industry's ability to resume, and therefore how long a withdrawal persists.[2] The industry therefore enters the next period with more work remaining and fewer capabilities with which to perform it. Where acceleration still dominates dilution, this deterioration lowers the return to resuming, so restarting requires more favorable news than sustaining research did. While investment stays low, the stock keeps depreciating and the restart threshold stays high. This depreciation spiral feeds withdrawal back into subsequent investment long after the initiating news has passed, and even after participation recovers, research expands from a smaller base. The two forces reinforce each other. Because clustering takes nearly all costly firms out at once, a large part of the knowledge base depreciates together; the deteriorated base raises the bar for the next round of research, and the industry clusters out again at news it would previously have absorbed. The spiral works through the industry's productive state rather than its balance sheets, and it does not require permanent technological failure or permanently pessimistic beliefs: a core whose costless learning continues eventually completes commercialization. The analytical results state the sign condition under which the feedback tightens the restart boundary; the numerical analysis checks incentives in the recursive specialization.

The quantitative exercise measures the scope of the mechanism. We fit the initial core share to the railroad construction contraction and maintain the other technology and information parameters. An 8 percent annual real reference rate maps a model period into 5.85 years and yields a mean commercialization horizon of 41 years. This is an illustrative multidecade transition, not an estimated historical

[2]The integration of capability depreciation—where a pause in research physically destroys usable knowledge and raises the bar for restarting—is what turns a static coordination failure into a persistent slump that delays commercialization for many years. The calibrated rate, 0.3 per model period of 5.85 years or about 5 percent a year, is conservative relative to estimates for business R&D capital (Li and Hall, 2020). At half that rate (Appendix A.5.2), withdrawals become frequent and short rather than rare and persistent, and the average surplus shortfall rises to 35.3 percent: lower depreciation changes the form of the inefficiency rather than removing it.

adoption date. The recursion uses a common current research cost; the heterogeneous-cost result provides a separate analytical foundation for clustered withdrawal.

At the initial state, news worse by 0.4 prior standard deviations lowers costly participation from 97.7 to 2.3 percent. With identical subsequent innovations, commercialization is postponed by 15.96 years privately and by 0.56 years under a planner who observes current news but not future innovations. The two forces act together rather than as separable channels: after the adverse news, 45.2 percent of histories contain a later round in which costly firms again withdraw together, against virtually none after the favorable news, and each such round further erodes the knowledge base. Net core equity falls 40 percent despite continued core research and a larger eventual share. Across independent histories, private investment produces 21.4 percent less expected surplus than the planner. A coordinated rule using only the public physical state recovers 28.8 percent of that gap. A fresh numerical refinement leaves these results nearly unchanged. A shorter decision interval preserves large conditional losses, but makes withdrawal less frequent and reduces the average surplus shortfall to about 5 percent. These outcomes demonstrate amplification, rather than a complete historical fit: conditional core receipts recover about 54 percent of pre-news value, below the ten-firm historical portfolios, and rebuilding is slower than the railroad benchmark. The contribution is a disciplined account of how complementary development can generate large, persistent investment and valuation losses alongside eventual technological success. Financing frictions and changing risk premia can operate in addition to this timing channel.

Related literature

The conventional view of technology crashes relies either on financial accelerators, leverage, and runs (Kiyotaki and Moore, 1997; Bernanke et al., 1999; Brunnermeier and Sannikov, 2014; Diamond and Dybvig, 1983; Goldstein and Pauzner, 2005), or on learning models such as Pástor and Veronesi (2009), in which a successful technology's price falls because its risk turns from idiosyncratic to systematic upon mass adoption. We propose an alternative mechanism driven by the link between research, the productivity of research, and the timing of profitable adoption. The mechanism generates a large price collapse without mispricing, simply by pushing the commercialization date back through a coordinated withdrawal of complementary research.

A number of contributions study related and complementary aspects of the boom-bust cycle. We share with Pástor and Veronesi (2006) their focus on fundamental valuation. In their model, uncertainty about average future profitability raises firm value and can account for Nasdaq valuations at the 2000 peak without a bubble. Gârleanu et al. (2012) study adoption cycles in which technological growth options become assets in place, with implications for valuations and risk premia. Hirano et al. (2026) link a general purpose technology to a rational bubble. In their framework, prices exceed fundamental value while growth is unbalanced, and these prices are expected to collapse once the technology matures. We hold the aggregate commercialization payoff and discount rate fixed to isolate a different margin: complementary private investment can inefficiently postpone that payoff. The distinctive prediction concerns the joint response of research, the endogenous payment date, and net equity value to coordinated withdrawal.

The production structure combines the complementary innovation needed to prepare a GPT for productive use (Bresnahan and Trajtenberg, 1995; Helpman and Trajtenberg, 1998) with the dependence of firm innovation on inherited knowledge (Klette and Kortum, 2004). We adapt the latter's portfolio scaling to programs that expand proprietary capabilities; physical depreciation here destroys usable knowledge rather than transferring product leadership.

The selection argument builds on global games (Carlsson and van Damme, 1993; Morris and Shin, 2003; Frankel et al., 2003). Our heterogeneous-cost results apply the belief constraint of Sákovics and Steiner (2012) and the clustering and translation arguments of Drozd and Serrano-Padial (2018). The new application links selected withdrawal to a physically determined commercialization date, net equity losses, and inefficient delay. Selection is sensitive to information: public signals and revealing prices can restore multiplicity (Angeletos and Werning, 2006; Hellwig et al., 2006; Angeletos et al., 2006), and dynamic learning introduces further issues (Frankel and Pauzner, 2000; Angeletos et al., 2007). Firms here commit research before prices reveal current news. Our unconditional uniqueness results concern the tractable games; the recursive analysis states additional conditions and reports numerical incentive checks.

Our mechanism also relates to theories of speculative booms (Abreu and Brunnermeier, 2003; Scheinkman and Xiong, 2003; Caballero et al., 2006), to real fluctuations driven by news, noise,

and endogenous learning (Beaudry and Portier, 2004; Jaimovich and Rebelo, 2009; Lorenzoni, 2009; Fajgelbaum et al., 2017), and to feedback between asset prices and investment (Angeletos et al., 2023).[3]

2 Three historical facts

Three observations motivate the theory: crashes recur during the build-out of successful technologies; their apparent triggers are limited in scope relative to the ensuing, persistent contraction in valuations and investment; and profitable incumbents retain substantial eventual value. We develop these observations for US railroads around 1873 and technology firms around 2000, with supporting evidence from British railways and the automobile industry.

The evidence follows the sequence from equity losses to investment withdrawal and eventual investor receipts. These are distinct measurements. A successful technology can coexist with losses on its original equity; investment expenditure differs from financing; and later receipts need not validate earlier prices. The episodes motivate the model without identifying its causal contribution. In particular, neither event-window returns nor the subsequent recovery identify the fundamental news investors received or the commercialization path that would have occurred without a contraction. Figure 1 follows monthly stock prices around each peak. Figure 2 compares investment in the build-out with the growth of commercial use. Figure 3 reports how much of the leaders' peak valuation was later delivered to investors who held their shares.

2.1 Crashes recur during the build-out of successful technologies

We measure a crash by the fall in an inflation-adjusted stock-price index from its peak month to its subsequent low. Figure 1 plots two such indices for each episode, each divided by the consumer price index and set to 100 in the month of the sector peak. The dashed line is the sector: the Macaulay index of all railroad stocks and the Nasdaq Composite. The solid line is a buy-and-hold portfolio of the industry

[3]Several related empirical works document the investment build-outs that precede commercialization. Brynjolfsson et al. (2021) show that GPTs require large complementary investments in intangible capital. These investments produce a J-curve in measured productivity. Jovanovic and Rousseau (2005) document the gradual diffusion of electricity and information technology. Comprehensive historical accounts of technology booms and busts can be found in Perez (2002) and Quinn and Turner (2020). The analysis in Perez (2002) is particularly relevant for our narrative: a technological revolution's installation period ends in a crash, which precedes a deployment period in which the technology spreads across the economy. We share this sequence but trace the crash to the timing of commercialization rather than to the dynamics of financial capital.

leaders defined in Section 2.3, weighted by their market values at the peak. Neither line includes dividends; payouts enter the valuation measure of Section 2.3. Shaded bands mark the news events discussed in Section 2.2. Appendix A.4 describes the construction and sources of every series used in this section.

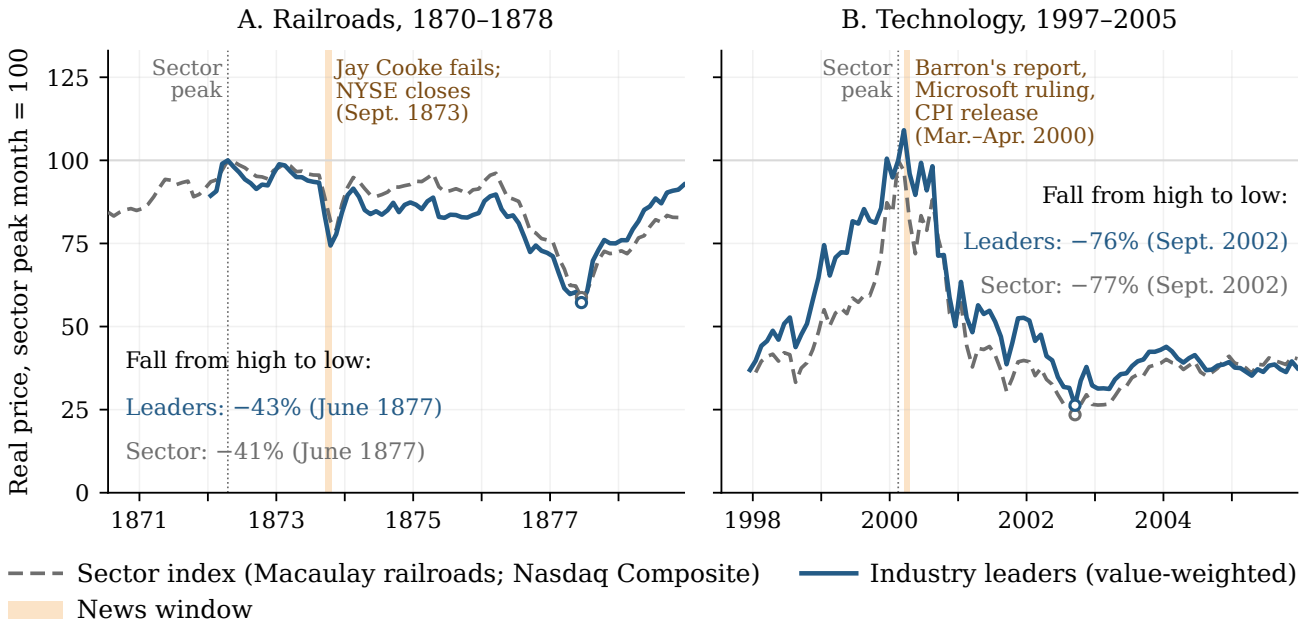

Figure 1: Stock prices around the railroad and technology peaks

Notes: Monthly prices without dividends, deflated by the consumer price index; the month of the sector peak (April 1872; February 2000) equals 100. Sector: Macaulay railroad stock-price index and Nasdaq Composite. Leaders: buy-and-hold portfolios of the ten railroad and ten technology leaders of Section 2.3, weighted by peak market value. Circles mark each line's low after its high near the peak. Railroad prices are deflated by the annual consumer price index interpolated to months; Appendix A.4.1 reports the declines with monthly wholesale prices. Sources and construction: Appendix A.4.

The railroad and internet crashes occurred while the commercial use of both technologies was still expanding. The Macaulay railroad index falls 41.3 percent in real terms between April 1872 and June 1877, and the ten railroad leaders fall 43 percent over the same months. Covering all ten leaders matters: the four roads with monthly quotations in the Yale price file, which hold 40 percent of the leaders' 1872 value and include the two that fell furthest, would show a 58 percent decline, whereas the complete portfolio falls about as much as the sector. The depth of the real decline also depends on the deflator. Wholesale prices fell twice as fast as consumer prices over these years; with the monthly Warren–Pearson wholesale price index in place of the interpolated annual CPI, the sector and the leaders fall 34 and 35 percent, and in nominal terms 48 and 49 percent (Appendix A.4.1). Panel A shows two stages: a sharp drop at the September 1873 panic, a partial recovery, and a second, deeper decline in 1876 and 1877, years after the

panic. Yet freight traffic rises 17 percent between 1873 and 1877 and 43 percent between 1873 and 1878, and the Northern Pacific completes its transcontinental line in 1883. Nasdaq loses 76.5 percent in real terms from February 2000 to September 2002, and the ten technology leaders lose 76 percent from their March 2000 high (Panel B). The fall is not specific to the Nasdaq: the Fama–French business-equipment portfolio loses 80 percent in nominal terms, dividends included, from March 2000 to September 2002. Retail e-commerce nevertheless expands from 0.925 percent of US retail sales in 2000 to 17.05 percent in the first half of 2026. At both peaks, use was still a small fraction of its eventual scale. Figure 8 in Section 5 places these expanding uses on a multidecade scale. It provides context for the model's calendar convention, rather than an estimated date of commercialization.

This sequence recurs in other technological transitions. British railway shares decline from their 1845 peak through 1849, during a substantial expansion of the network financed almost entirely with private share capital (Odlyzko, 2010, 2011). Britain's first railway boom had also ended in a crash, in 1837; nearly all the lines it authorized were nevertheless in service by 1843, and by 1845 investors regarded the episode as a success (Odlyzko, 2010). In the US automobile episode, listed prices fall 64.2 percent between November 1919 and January 1922; annual passenger-car production then more than doubles between 1921 and 1923. Appendix A.4.6 documents that episode. Across these cases, a crash interrupts the industry's development without marking the end of the technology's commercial expansion.

2.2 The collapse is large relative to its apparent triggers

The disparity concerns both the scope of the news and the persistence of the response. The events associated with these crashes concern an intermediary, firms' ability to finance continued operations, or the terms on which investors can trade their claims. The ensuing contraction extends across the industry and persists for years. Historical accounts differ over which events matter most and how they interact with existing vulnerabilities.[4] We describe the associated news, measure how much of the price decline occurs in selected event windows, and then document the investment contraction and failures that follow. This ordering separates the chronology from a causal interpretation.

[4]The comparison concerns the apparent scope of the events, not an identified change in fundamental value. Appendix A.4.4 reviews the proposed catalysts and reports event-window price changes.

In 1873, the conspicuous catalyst is the failure of Jay Cooke & Co. on September 18 over its financing of the Northern Pacific. The NYSE suspends trading two days later. Nitschke (2018) emphasizes how Cooke's reputation made this failure consequential for confidence beyond the firm itself. The boom had also produced notorious scandals, among them Crédit Mobilier, exposed in 1872 (White, 2011). Mixon (2008) finds that markets already perceived greater risks among more leveraged railroads before the panic.

For the internet episode, the March 20, 2000 *Barron's* report provides a concrete example of a public reassessment of firms' ability to continue. Using already public financial statements, it identifies 51 of 207 internet companies that would exhaust their cash within twelve months at unchanged late-1999 revenue and expenditure rates (Willoughby, 2000). Ofek and Richardson (2003) emphasize a different mechanism: expiring IPO lockups bring additional potential sellers into the market. The accounting frauds most associated with the episode, at Enron and WorldCom, surfaced in late 2001 and 2002, well into the decline (Healy and Palepu, 2003; U.S. Securities and Exchange Commission, 2002). These developments occur during the monetary tightening of 1999–2000. Neither account requires a discovery that the internet has lost its technological usefulness.

To compare these events with the decline that follows, we measure the share of the sector's eventual log price loss that occurs within the news window. A share of one would mean that the entire measured decline occurs within the selected window; a small share means that most of the loss occurs outside it. This accounting does not attribute the later decline to a separate shock or measure the information content of the initial event. The railroad index falls 15 percent from August to October 1873, which is 25 percent of its eventual loss. On the Nasdaq, the *Barron's* session and the session of April 14, 2000, when the March CPI was released, together account for 9 percent of the eventual loss. Adding the session of April 3, when the court issued its conclusions of law in the Microsoft case, raises the share to 15 percent. Even the entire fall from the March 10 peak through April 14 accounts for only 28 percent. These shares depend on how each window is dated. Moving each 2000 window by one or two trading sessions leaves the three sessions together with between −7 and 6 percent of the eventual loss, where a negative share means that prices rose in the shifted window; moving the railroad window by a month gives 13 or 19 percent (Appendix Table A3). The window shares describe the chronology of the decline, not the size

of the news. In both episodes, as Figure 1 shows, prices continue to fall for two and a half to four years after the news.

The valuation decline is accompanied by a sustained retreat in resources committed to commercialization. Figure 2 plots investment flows against measures of use, each set to 100 in the peak year and drawn on a logarithmic scale so that equal vertical distances represent equal proportional changes. In Panel A, annual first-track completions fall from 7,439 miles in 1872 to 1,606 in 1875, a decline of 78.4 percent, while freight traffic continues to grow. Construction does not regain its earlier annual flow until 1881, nine years after the peak. In Panel B, between 2000 and 2003, expenditure on computers and communications equipment falls 33.0 percent in CPI-deflated terms, real venture investment falls 82.4 percent, and technology IPO proceeds fall 95.1 percent, while the e-commerce share of retail sales nearly doubles. Research also contracted: real US business-funded R&D fell 6.4 percent between 2000 and 2002 and remained below its 2000 level through 2005 (National Center for Science and Engineering Statistics, 2025). This economy-wide series supplies context; it is not an estimate of listed IT firms' research. Equipment expenditure, venture investment, and IPO proceeds likewise measure different margins: purchases of equipment, venture financing, and public equity issuance. Their common retreat documents a contraction across these margins without treating them as interchangeable measures of physical investment.

The contraction also appears in failures and exits. From 1873 to 1879, 183 railroad receiverships covered 24,388 miles of track, and the mileage placed in receivership in 1875 alone was a tenth of the network. Only 48 percent of the internet start-ups that sought venture funding in 1998–2002 survived to 2004. In automobiles, 52 of the 88 producers operating in 1921 exited by 1926.

The object to explain is this broad, persistent withdrawal around news of limited apparent scope, a withdrawal that continues for years after the dated events. Financial distress, fraud, and leverage belong to the sequence; the chronology alone does not determine whether it initiates or reflects deteriorating commercialization prospects. The model examines how firms' responses to news can themselves worsen those prospects and prolong the investment contraction.

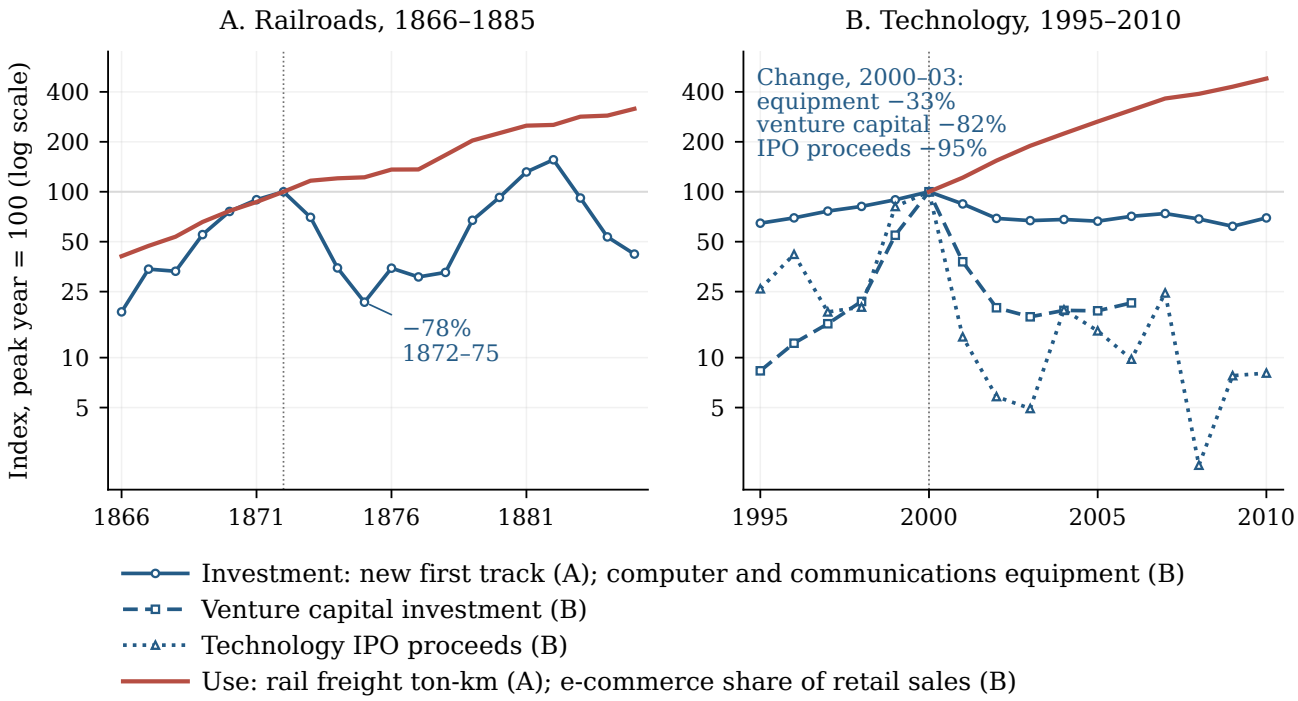

Figure 2: Investment in the build-out and growth in commercial use

Notes: Annual series, peak year (1872; 2000) equals 100, logarithmic scale. Panel A: new first-track mileage and freight ton-kilometers. Panel B: computer and communications equipment expenditure, venture capital investment, and technology IPO proceeds, each deflated by the CPI, and the e-commerce share of retail sales. Sources and construction: Appendix A.4.

2.3 Profitable incumbents retain substantial eventual value

A large interim price loss does not reveal how much value an original shareholder eventually receives. We therefore follow payouts and residual claims from fixed reference valuations near the peaks. The calculation measures realized recovery at stated discount rates; it is not a test of the correctness of ex ante prices. To assess how much initial value is eventually recovered, we select large firms and rank them by profitability, fixing portfolio membership before following payouts, failures, and acquisitions. A10 and A5 denote the ten and five highest-ranked firms. Reference valuations are the 1872 annual average for railroads and January 4, 2000 for technology firms.[5] Membership is held fixed when following outcomes, so failures, acquisitions, and firms implicated in misconduct, such as the Erie (Daggett, 1908) and Lucent and Computer Associates (U.S. Securities and Exchange Commission, 2004a,b), remain in the shareholder accounting. The railroad screen nevertheless uses accounts published after the reference valuation year;

[5]Railroad selection uses fiscal-1872 accounts published in 1873. The technology screen uses earnings reported before January 4, 2000 where original releases are available. Appendix A.4 gives the selection rules, information dates, and corporate-action treatment.

it is not a portfolio shown to be implementable with information available at every 1872 purchase date. The railroad screen also ranks firms on net earnings before interest, whereas the technology screen uses net income after interest; each identifies profitable firms within its own episode rather than imposing a common accounting measure. The appendix records this information lag and the historical inputs that require reconstruction.

The *justified fraction* is discounted payouts and residual value per initial dollar invested:

$$J(r) = \sum_i w_i \left[\sum_s \frac{c_{is}}{(1+r)^{t_s}} + \frac{v_{iH}}{(1+r)^H} \right], \qquad \sum_i w_i = 1. \tag{1}$$

Here w_i is the initial equity weight, c_{is} is firm i's payout net of shareholder contributions per initial dollar at date t_s, v_{iH} is its residual value per initial dollar at horizon H, and r is an annual effective nominal discount rate. The endpoints are 1895–98 for railroads and October 2026 for technology. A value of $J(r) = 1$ means that an investor who bought the portfolio at the reference valuation earned exactly the annual return r; a value below one measures the shortfall from that return. Unlike the price indices in Figure 1, $J(r)$ counts dividends and cash distributions and discounts each receipt by the time required to obtain it.

At an 8 percent discount rate, subsequent payouts and residual claims cover 70.7 percent of railroad A10's initial value and 70.6–71.0 percent of technology A10's value. Approximately seven-tenths of these portfolios' pre-crash valuations is therefore recovered after discounting receipts over more than two decades. Portfolio composition matters: railroad A5 recovers 87.3 percent and technology A5 recovers 33.7 percent. Table 1 and Figure 3 report these values together with those at a 10 percent discount rate. For the railroads, most of the value comes as dividends: at 8 percent, dividends account for 55 of A10's 71 points and the 1895–98 market value for 17, less one point of assessments. Composition and reconstructed inputs matter. Microsoft supplies 47.6 percentage points of technology A10's recovery; excluding it lowers the fraction to 35.1–35.6 percent. Setting estimated railroad dividends to zero lowers railroad A10's fraction to 48.3 percent; that calculation still uses estimated opening and terminal prices, so it is not a lower bound on recovery. Appendix A.4 reports these sensitivities. The similarity between the two A10 headline fractions should therefore not be read as a stable cross-episode parameter.

Automobiles provide a check on the same calculation, using industry price and dividend indices rather than a selected leader portfolio. The Cowles indices for General Motors and for the other listed makers deliver 102 and 105 percent of their May 1919–April 1920 value by 1924–26 at 8 percent; the stock market as a whole delivers 101 percent over the same window (Appendix A.4.6). The technology industry as a whole shows the same pattern. A dollar placed in the value-weighted Fama–French business-equipment portfolio at the start of January 2000 and held through August 2026, with distributions reinvested, delivers 125 percent of its value at 8 percent and 77 percent at 10 percent, after the 80 percent fall documented above (Appendix A.4.5).

Table 1: Justified fraction of the peak valuation

Portfolio, initial-value weights	$J(0.08)$	$J(0.10)$
Railroads, A10	70.7	57.1
Railroads, A5	87.3	71.8
Technology, A10	70.6–71.0	46.1–46.5
Technology, A5	33.7	24.1
Automobiles, General Motors (check)	102.4	93.1
Automobiles, other listed makers (check)	105.1	95.6
Memo: all listed stocks, 1919–20 reference	101.1	92.1

Notes: Percent of the reference valuation, equation (1). The technology A10 range reflects two unobserved inputs. Sources and construction: Appendix A.4; alternative railroad end window: Appendix A.4.5.

Eventual recovery coexists with substantial initial losses. The median worst real-price decline is 44.9 percent among the railroad leaders and 87.2 percent among the western comparison roads. Technology A5 loses 69.4 percent in real price through end-2002, close to Nasdaq's 68.1 percent over the same dates. Between 1873 and 1880, two of the ten railroad leaders entered receivership or default, against eight of ten comparison roads. Together with the exits documented in Section 2.2, where only 48 percent of the internet start-ups of 1998–2002 survived to 2004 and 52 of the 88 automobile producers of 1921 exited by 1926, this contrast sets a further requirement for any account of these crashes: the retreat removes a large part of the industry within a few years, while the profitable leaders survive and keep much of their value. The leaders' later receipts reveal value that a crash-depth comparison alone cannot measure. They also leave substantial shortfalls at the stated discount rates. The model supplies one interpretation of this combination: a firm can retain a claim on a valuable technology while losing present value as the

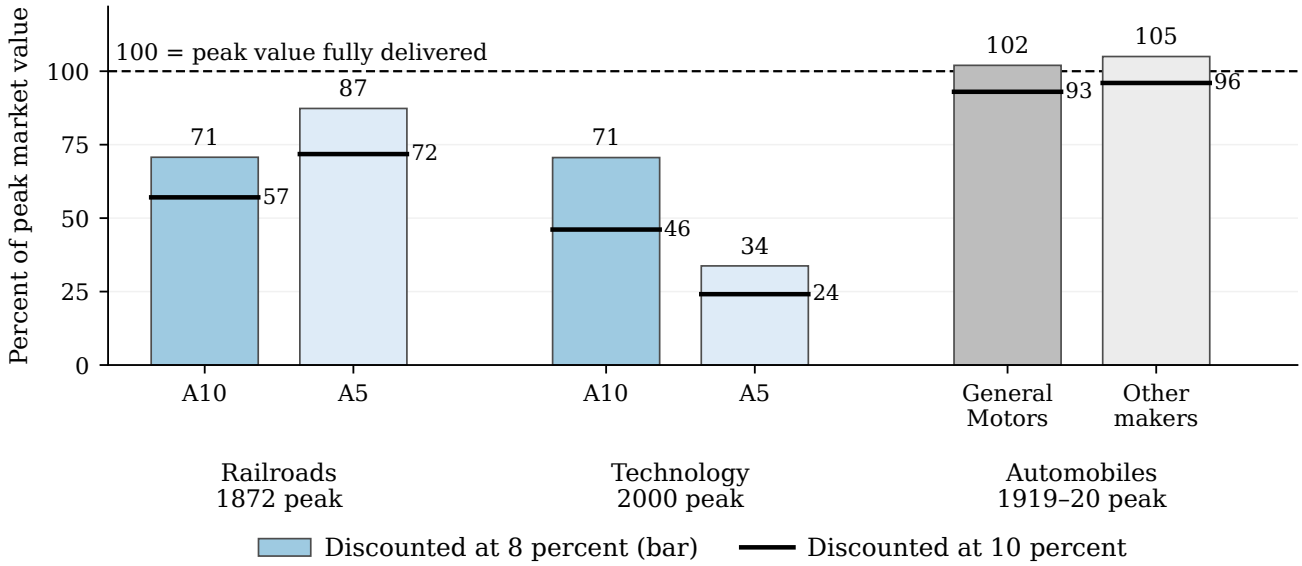

Figure 3: Share of the leaders' peak valuation later delivered to investors

Notes: Bars: justified fraction $J(0.08)$, the present value at the reference date of later payouts plus the market value of the remaining claims at the end date, in percent of the reference valuation. Black lines: $J(0.10)$. The dashed line at 100 corresponds to an investment that earned exactly the discount rate. Automobile bars use industry indices and serve as a check. Values as in Table 1; the technology A10 bar shows the lower end of its range. Sources: Appendix A.4.

shared payoff is postponed. The historical calculation does not identify how much of its shortfall reflects that delay, risk, changed profit expectations, or initial mispricing.

The three measurements define the empirical question taken to the model. Commercial use expands through a severe equity contraction; the commitment of new resources retreats persistently, with failures concentrated outside the profitable leaders; and some original claims retain substantial discounted value. The model asks whether endogenous commercialization timing can generate this combination, and how its private investment path differs from a coordinated allocation.

3 Model of dynamic commercialization

Firms accumulate capabilities that perform commercialization work and give their owners claims on a common payoff. Each period, firms decide whether to research using their current costs, private information about remaining requirements, and continuation values generated by future equilibrium play. The model allows them to pause and resume; the commercialization date is determined by their accumulated work. We first define the productive state, then the information and research decision, and

finally the equilibrium values and planning benchmark. Section 4 uses this structure to separate current coordination from dynamic propagation.

The fixed completion payoff isolates a timing question. It represents the capitalized profits from large-scale use, rather than a claim that historical railroads or technology firms earned no earlier revenue. Capabilities collect the proprietary knowledge needed to earn those profits; remaining work collects the common engineering and implementation requirements. Existing businesses can generate revenues outside the modeled GPT claim. Bringing the completion payoff forward is valuable even when its eventual size is unchanged.

3.1 Capabilities, firms, and commercialization

Time is discrete. The discount factor is $\beta = \exp(-r)$ with $r > 0$, and the public state at the start of a period is

$$S = (C, F, R). \tag{2}$$

The positive stocks C and F are capabilities owned by the core and costly firms. A capability is useful proprietary technological knowledge: it supplies commercialization work, supports further research, and gives its owner a proportional claim on the eventual payoff. The state $R > 0$ is remaining commercialization work. Completion delivers a fixed aggregate payoff $m > 0$, which can represent the capitalized value of subsequent operating profits. The GPT claim pays no operating dividend before commercialization. Investment changes when the common payoff arrives and who receives it.

A costly firm owns n capabilities. Running its research program creates $b\,n$ new capabilities at resource cost $b\,z\,n$, where $b > 0$ and z is its current cost per new capability. Pausing leaves its inherited stock unchanged before depreciation. This specification adapts the chained research mechanism of Klette and Kortum (2004): accumulated knowledge supports further innovation, and current expansion is proportional to the firm's inherited capabilities. The core creates $b\,C$ capabilities without an incremental resource bill, representing learning generated by other ongoing activities. Costless learning supplies continuing progress through a contraction. It is also individually optimal for an atomistic core owner, who acquires additional claims without a resource bill. The maintained assumption is that these oppor-

tunities remain available during a contraction; profitability alone does not make a historical incumbent part of this core.

Current research costs are drawn independently across costly firms and dates from a publicly known distribution H on $[\underline{z}, \overline{z}]$, with $0 < \underline{z} \leq \overline{z} < \infty$. The draws are independent of inherited capabilities, commercialization news, and signal errors. Each firm observes its cost before deciding whether to research. Because costs are drawn afresh, every costly capability has the same distribution of future research opportunities at a given public state. A point mass at z gives the common-cost case used in the numerical exercise.

Organizational capital and intangible know-how depreciate continuously within each period at rate $\delta \geq 0$: specialized teams disperse, routines deteriorate, and technical expertise becomes obsolete. This is a loss of productive capability, distinct from the transfer of product leadership to a competitor in Klette and Kortum (2004). For aggregate participation $a \in [0, 1]$ among costly capabilities, post-program capacity and work performed per capability during the period are

$$N(a) = (1 + b)C + (1 + b\,a)F, \qquad h = \int_0^1 \exp(-\delta x)\, dx. \tag{3}$$

Thus $h = (1 - \exp(-\delta))/\delta$ when $\delta > 0$ and $h = 1$ otherwise. Current news θ changes remaining requirements. A positive innovation raises the work needed and is adverse; a negative innovation reduces it. The work flow $hN(a)$ is supplied by the entire stock of usable capabilities after current programs, not just the newly created units. New research consequently expands both the productive base and the ownership claims on the payoff. Commercialization is certified at the period end if

$$R + \theta \leq hN(a). \tag{4}$$

Otherwise, next period's state $S' = (C', F', R')$ is

$$C' = \exp(-\delta)(1 + b)C, \qquad F' = \exp(-\delta)(1 + b\,a)F, \qquad R' = R + \theta - hN(a). \tag{5}$$

Requirements follow a random walk after subtracting completed work. They can include engineering, integration, infrastructure, compliance, and financing costs and frictions (without finance-implied propagation mechanisms). An adverse innovation raises remaining work once and does not predict further adverse innovations. Persistence after a crash therefore comes from withdrawal itself, through the depreciation spiral, rather than from the news process. At completion, an owner of n post-program capabilities receives $mn/N(a)$. Proportional depreciation within the final period cancels from ownership shares.

3.2 Information and the research decision

Let $\sigma, \nu > 0$. At each date, $\theta \sim N(0, \sigma^2)$ is independent of earlier innovations. Firm i observes its current cost z_i and a private signal

$$s_i = \theta + \nu \varepsilon_i, \qquad \varepsilon_i \sim N(0,1), \tag{6}$$

with conditionally independent errors. Firms then commit their programs. The common innovation is disclosed before work occurs, so next period's physical state is public. Current research decisions use the private signals. There is no intervening price signal. Figure 4 summarizes the information and payment sequence. This timing lets post-decision equity reflect current news without allowing that price to coordinate programs that have already been committed.

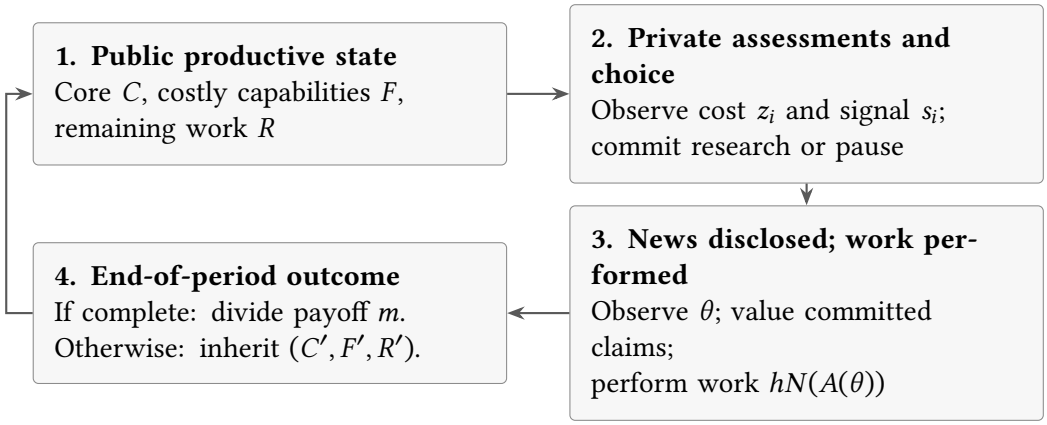

Figure 4: Information, research, and payment within a model period

Notes: Research costs are paid when programs are committed. The current innovation is then disclosed, and existing and new capabilities perform work. Completion pays at the period end; otherwise depreciated capabilities and unfinished work form the next public state. The return arrow applies only before completion. Prices formed after commitment cannot inform the current research choice.

Firms are atomistic. Each takes the aggregate research policy and its implied commercialization date as given. Every future capability stock, research bill, and terminal receipt scales with a firm's inherited n. Its value before the current cost draw and signal is therefore

$$\text{firm value} = nv(S), \tag{7}$$

where v is the value per costly capability. Write u for the value per core capability. Fresh cost draws make v common across costly firms, including firms with different costs today.

To evaluate the current choice, consider one capability after the program has been chosen. If completion occurs now, its discounted receipt is $\beta m/N(a)$. Otherwise it leaves $\exp(-\delta)$ capabilities next period. Define

$$Q(\theta, a) = \beta \begin{cases} m/N(a), & R + \theta \leq hN(a), \\ \exp(-\delta)v(S'), & R + \theta > hN(a). \end{cases} \tag{8}$$

Define $U(\theta, a)$ by replacing v with u in the continuation line. For a core owner, the gain from a costless program is $b\,n\mathbb{E}[U \mid s]$. It is strictly positive: Gaussian news gives a positive probability of current completion, and continuation receipts are nonnegative. Thus allowing the core the same binary program choice would support its maintained expansion rule.

Both returns depend on the current state and future equilibrium research; the state argument is suppressed, and S' is given by (5).

Let $a_z(\theta)$ be participation among costly capabilities with current cost z, conditional on the innovation. Aggregate participation and the total resource bill are

$$A(\theta) = \int a_z(\theta)\, dH(z), \qquad \text{research bill} = bF \int z a_z(\theta)\, dH(z). \tag{9}$$

A firm with signal s_i and cost z_i compares pausing, worth $n\mathbb{E}[Q(\theta, A(\theta)) \mid s_i]$, with researching, worth $(1 + b)n\mathbb{E}[Q(\theta, A(\theta)) \mid s_i] - b\,z_i n$. It researches when

$$\mathbb{E}[Q(\theta, A(\theta)) \mid s_i] > z_i, \tag{10}$$

and may choose either action at equality. Its cost is independent of the information about θ, so it changes the required return without changing the posterior. Inherited capabilities remain valuable under either action.

For example, with $b = 0.5$ and $z_i = 2.5$, an expected return of three per post-program capability makes pausing worth three per inherited capability. Research is worth $1.5 \times 3 - 0.5 \times 2.5 = 3.25$. The positive program gain is 0.25, while current expansion remains bounded by inherited knowledge. The firm is indifferent only when the expected return equals 2.5.

3.3 Recursive equilibrium and continuation boundaries

Before current costs and signals arrive, values satisfy

$$v(S) = \mathbb{E}\left[(1 + bA(\theta))Q(\theta, A(\theta)) - b\int z a_z(\theta)\, dH(z)\right], \tag{11}$$

$$u(S) = (1 + b)\mathbb{E}\big[U(\theta, A(\theta))\big]. \tag{12}$$

The actual cost-weighted participation appears in the first equation because low-cost programs can be more likely to run.

Definition 1 (Recursive equilibrium)**.** A recursive equilibrium consists of research choices as functions of the public state, current cost, and private signal; their conditional participation rates $a_z(\theta)$ and aggregate $A(\theta)$; and capability values v and u. Choices maximize each firm's value, participation aggregates those choices according to (9), and the value functions satisfy (11)–(12) at every public state. Values equal discounted commercialization receipts net of subsequent research costs, with no discounted residual value at an infinite horizon.

When equilibrium choices have a cutoff representation, a cost-z firm researches at $s_i < c(z; S)$. Suppressing the state argument,

$$a_z(\theta) = \Phi\left(\frac{c(z) - \theta}{\nu}\right), \qquad A(\theta) = \int \Phi\left(\frac{c(z) - \theta}{\nu}\right) dH(z), \tag{13}$$

where Φ is the standard normal cdf. Set $w = \sigma^2/(\sigma^2 + v^2)$. Bayes' rule gives $\theta \mid s_i \sim N(ws_i, wv^2)$. Each marginal type satisfies

$$\mathbb{E}[Q(\theta, A(\theta)) \mid s_i = c(z)] = z. \tag{14}$$

These conditions jointly determine the cutoff schedule against future continuation values. A candidate schedule must also make the prescribed choice optimal at every other signal.

In the common-cost specialization, the schedule reduces to one cutoff c:

$$A(\theta) = \Phi\left(\frac{c - \theta}{v}\right), \tag{15}$$

$$G(c; S) = \mathbb{E}[Q(\theta, A(\theta)) \mid s_i = c] - z = 0. \tag{16}$$

Appendix A.2 gives all-signal verification, dominance bounds that include future research options, and sufficient conditions for a unique current cutoff in this specialization. The next section explains how cost dispersion affects clustered withdrawal and how the physical state changes the continuation boundary over time.

3.4 Equity and the planning benchmark

Opening industry equity is

$$I(S) = Cu(S) + Fv(S). \tag{17}$$

After programs are committed and current news is disclosed, prices are

$$P_C = (1 + b)CU(\theta, A(\theta)), \qquad P_F = (1 + bA(\theta))FQ(\theta, A(\theta)). \tag{18}$$

These are date-of-decision values: the end-of-period receipt or continuation claim is discounted once by β. They are gross of the current research bill and net of future bills. The original owners' current net

value subtracts $bF \int z a_z(\theta)\, dH(z)$. Summing the Bellman equations gives

$$I(S) = \mathbb{E}\left[-bF \int z a_z(\theta)\, dH(z) + \beta \begin{cases} m, & \text{completion this period,} \\ I(S'), & \text{otherwise.} \end{cases}\right]. \tag{19}$$

Equity is expected discounted net payouts. Firms do not treat news as harmless: they know its distribution, and their values anticipate future adverse innovations and the withdrawals these would trigger. Opening equity therefore prices a crash at its expected value. Valuations change with the timing and allocation of those payouts and with future funding calls. As in Pástor and Veronesi (2006), pricing is convex: the payoff is discounted over an uncertain commercialization date, and a firm pays a funding call only when its expected return exceeds its cost. A peak valuation can therefore look unjustified after an adverse outcome without having been mispriced.

Define a realized payoff ratio J as an original owner's discounted receipts, less funding calls, divided by the opening price. Fundamental pricing implies $\mathbb{E}[J \mid S] = 1$. Conditioning later on adverse news, withdrawal, or survival can change the mean of this ratio when the denominator remains the opening price. A survivor's rising market share is not itself an excess-return calculation.

The causal planner observes current news and current research costs, but not future innovations or cost draws. Given total participation a, it assigns programs to the lowest-cost capabilities. Define

$$K(a) = \int_0^a H^{-1}(p)\, dp, \qquad 0 \le a \le 1, \tag{20}$$

where H^{-1} is the cost quantile function. The minimum resource bill is $bFK(a)$. The planner's value satisfies

$$W(S) = \mathbb{E} \max_{a \in [0,1]} \left\{ -bFK(a) + \beta \begin{cases} m, & R + \theta \le hN(a), \\ W(S'), & R + \theta > hN(a). \end{cases} \right\}. \tag{21}$$

The maximization is inside the expectation because current news is known; future uncertainty remains in the continuation value. For a common cost z, $K(a) = za$. The comparison combines coordinated

investment with pooled current information. The quantitative section also evaluates a coordinated rule that uses only the public physical state.

These are producer-surplus comparisons when m denotes profits. With m interpreted as total allocable surplus and $\beta(1 + b) < 1$, a minimal general-equilibrium closure has households own the firms, value numeraire consumption linearly with discount factor β, and receive an endowment sufficient to finance feasible programs, with finite discounted value. Consumption equals the endowment less research expenditures, plus the completion receipt. The planner's objective differs from aggregate welfare only by the discounted endowment. This closure fixes the discount rate and resource cost without introducing another demand channel.

4 Mechanisms of fragile commercialization

The model contains two distinct sources of amplification. At a given productive state, research by other firms can make a firm's own program profitable by advancing the common payoff. Dispersed information then determines where a collective withdrawal occurs. Across dates, withdrawal changes the productive state itself: the capabilities and unfinished work inherited by the next research game. The first mechanism explains an abrupt response; the second explains why its consequences outlast the initial news.

We develop the argument in three steps. A two-period version of the commercialization technology isolates selection and inefficient delay with a single research cost. Three costs then show why heterogeneity need not smooth the contraction. Finally, the recursive transition identifies the productive losses left by a pause and the condition under which they tighten subsequent research incentives. The first two exercises have closed-form structure; the dynamic condition is stated separately because acceleration need not dominate dilution at every recursive state.

4.1 Simplified model of fragile implementation

To isolate static amplification, suppress depreciation and subsequent research, leaving one program decision and two possible dates of payment. Research takes place at date 0. The technology can be imple-

mented at date 1 or date 2 and pays the same total amount at either date. We first give all costly firms the same research cost, then add two cost types in Section 4.2.

Capabilities, research, and ownership. An atomistic firm owns n capabilities. A capability is a unit of useful proprietary technological knowledge. The underlying knowledge performs work needed for commercialization and creates further research opportunities. Ownership gives a proportional claim on the technology's eventual payoff. Running a research program doubles the firm's capabilities, from n to $2n$, at cost $2n$. Pausing leaves n capabilities. Thus the program is linear in inherited capabilities, and an individual firm takes the aggregate investment rate and commercialization date as given.

Initially, the costless core owns $C = 1/4$ capabilities and costly firms own $F = 1/2$. The core doubles its stock through learning generated by other activities. If costly firms run programs on a fraction a of their inherited capabilities, the total stock after the decision is

$$N(a) = 2C + (1 + a)F = 1 + \frac{a}{2}. \tag{22}$$

No further capabilities are created in this example. The knowledge embodied in each capability performs one unit of commercialization work in each period. At implementation, the fixed prize $m = 16$ is divided in proportion to capability ownership. The discount factor is $\beta = 1/4$. These numbers make the calculation transparent.

What firms know. The public remaining-work baseline is $R = 3/4$. The total requirement is $R + \theta + \eta$, where θ is uniform on $[-1/4, 1/4]$ and η is independently uniform on $[0, 1]$. A higher θ is worse news. Before choosing its program, firm i observes

$$s_i = \theta + \nu\,\varepsilon_i, \qquad \varepsilon_i \sim \text{Uniform}[-1, 1], \qquad 0 < \nu < \frac{1}{6}. \tag{23}$$

The signal errors are independent across firms conditional on θ. Current news θ is disclosed after programs are committed. The remaining certification requirement η becomes known at date 1. Neither

firms nor the planner know it at the investment date. Uniform signals give a closed-form example; the recursive quantitative model uses Gaussian signals.

A common commercialization timeline. The first period completes $N(a)$ units of work. The probability of implementation at date 1, conditional on θ and a, is therefore

$$\Pr\{\tau = 1 \mid \theta, a\} = \Pr\{R + \theta + \eta \leq N(a)\} = \frac{1}{4} + \frac{a}{2} - \theta. \tag{24}$$

This expression lies in $[0, 1]$ over the entire state and action space. If the first period is insufficient, the second completes another $N(a)$ units. Total requirements are at most two and $N(a) \geq 1$, so *implementation by date 2 is certain even if every costly firm pauses.* Research changes the timing and division of the fixed prize. It does not determine whether that prize ultimately exists.

Why a firm wants others to continue. A capability after the current program is worth

$$Q(\theta, a) = \frac{m}{N(a)}\mathbb{E}[\beta^{\tau} \mid \theta, a] = 3 - \frac{5/4 + 3\theta}{1 + a/2}. \tag{25}$$

The numerator is the expected discounted common prize. The denominator captures dilution. Both effects come from the work and ownership technology; no additional payoff from matching others' actions is imposed. Here acceleration dominates dilution: Q decreases in θ and strictly increases in a everywhere on the prior support. The expected return is smooth in both arguments. Thus an abrupt equilibrium response need not come from a discontinuity in the return to a given aggregate research allocation. It comes from firms' responses to one another.

At signal s, pausing gives $n\mathbb{E}[Q \mid s]$, while researching gives $2n\mathbb{E}[Q \mid s] - 2n$. The firm researches exactly when

$$\mathbb{E}[Q \mid s] > 2. \tag{26}$$

Its inherited capabilities remain valuable in either case. If the fundamental were common knowledge, research would be strictly dominant for $\theta < -1/12$, and pausing would be strictly dominant for $\theta > 1/12$. Between those bounds, both universal research and universal withdrawal are equilibria. Noisy

information determines where withdrawal occurs inside that interval. The equilibrium contraction is therefore a response to a change in requirements under a specified information structure. It is not a switch between arbitrarily selected equilibria at the same realized fundamental.

The return calculation in numbers. Hold the fundamental at $\theta = 0$. If every costly firm pauses, there is one capability in the economy, the probability of early completion is $1/4$, and a capability is worth 1.75. If all costly firms research, the stock rises to 1.5, early completion has probability $3/4$, and each capability is worth $13/6 \simeq 2.167$. Research therefore loses 0.25 per inherited capability against universal withdrawal and gains about 0.167 against universal research. The larger stock dilutes each claim by one-third, but the earlier common receipt more than offsets that dilution. These are two aggregate allocations at the same fundamental; the information structure determines the equilibrium allocation.

The belief constraint and selected boundary. For a cutoff $c(\nu)$, firms research when $s < c(\nu)$. Aggregate participation is

$$A(\theta) = \left[\frac{c(\nu) - \theta + \nu}{2\nu}\right]_0^1, \qquad [x]_0^1 = \min\{1, \max\{0, x\}\}. \tag{27}$$

The relevant *belief constraint* concerns the marginal firm's rank in the population of private signals. Let F_ε denote the noise cdf. Conditional on θ, a firm's percentile rank is $F_\varepsilon(\varepsilon_i)$. At the marginal signal $s_i = c(\nu)$, other firms research precisely when their signals lie below its own signal. Aggregate participation therefore equals its rank:

$$A(\theta) = F_\varepsilon\left(\frac{c(\nu) - \theta}{\nu}\right) = F_\varepsilon(\varepsilon_i) \quad \text{when } s_i = c(\nu). \tag{28}$$

The flat prior implies that, conditional on this interior cutoff signal, the firm's own error retains its original distribution. The probability integral transform then makes its rank uniform on $[0, 1]$.[6]

[6]The dominance bounds give $c(\nu) \in [-1/12, 1/12]$. Since $\nu < 1/6$, Bayes' rule gives $f_{\varepsilon|s}(e \mid c) \propto f_\theta(c - \nu e) f_\varepsilon(e) = \text{constant} \times f_\varepsilon(e)$ throughout the noise support. For any continuous cdf F and $X \sim F$, $\Pr\{F(X) \le u\} = F(F^{-1}(u)) = u$, $0 < u < 1$, using the generalized inverse. Thus $F(X)$ is uniform.

Rank matters at the margin because it determines how many peers continue, and their participation determines the return against which the marginal firm is indifferent. This indifference pins down the equilibrium cutoff. Even as noise vanishes, uncertainty about the marginal firm's rank remains: uncertainty about the fundamental and the differences between signals shrink together. The belief constraint therefore averages the nonlinear return over the entire participation distribution, rather than evaluating it at its mean of one-half. Section 4.2 extends this argument to heterogeneous costs using the belief accounting of Sákovics and Steiner (2012) and the clustering logic of Drozd and Serrano-Padial (2018).

Substituting $\theta = c(\nu) + \nu - 2\nu a$ into the marginal firm's indifference condition gives

$$\int_0^1 [Q(c(\nu) + \nu - 2\nu a, a) - 2]\, da = 0. \tag{29}$$

The candidate is optimal at every signal because the realized return $Q(\theta, A(\theta))$ strictly decreases with the fundamental. Appendix A.1 gives the closed-form solution and proves uniqueness by iterated dominance, including among strategies that were not initially required to be cutoffs. For $\nu = 0.01$, the cutoff is -0.00629. News of -0.01629 induces full research, while news of 0.00371 induces complete withdrawal. The change in the requirement is only 0.02; all of the discrete change in the research set comes from firms' optimal responses to that news and to each other's responses.

Coordination and the planner. The causal planner knows θ but faces the same unresolved η. Its value, net of current research expenditures, is

$$W(\theta, a) = m\mathbb{E}[\beta^\tau \mid \theta, a] - a = \frac{7}{4} - 3\theta + \frac{a}{2}. \tag{30}$$

It always chooses $a = 1$. That policy does not depend on the information available about θ; its advantage here comes entirely from coordinating investment. Expected implementation time is

$$\mathbb{E}[\tau \mid \theta, a] = 2 - \Pr\{\tau = 1 \mid \theta, a\} = \frac{7}{4} + \theta - \frac{a}{2}. \tag{31}$$

At $\theta = 0$, for example, coordinated research raises the present value of the aggregate prize from 1.75 to 3.25. Paying the research bill of one leaves net surplus of 2.25, compared with 1.75 under withdrawal. Each firm takes the payment date as given, but their collective decision changes it.

More generally, at a fixed fundamental in this one-decision economy, aggregate net value is $W = N(a)Q(\theta, a) - bzFa$. Its marginal change is

$$W_a = bF(Q - z) + NQ_a.$$

The first term is the private return on the added capabilities; the second is the change in the value of all existing post-program claims. An atomistic firm omits that second term. It is positive here because the common acceleration benefit exceeds dilution. Fundamental pricing therefore coexists with inefficient research: correctly priced private claims do not make firms internalize their effect on the common date.

Proposition 1 (Inefficient postponement of a successful technology)**.** *The two-period economy has a unique Bayesian equilibrium, up to choices at indifferent signals. Its cutoff is determined by* (29)*, and the technology is implemented by date 2 under every feasible research policy.*

Compare the two current innovations $\theta_G = c(v) - v$ and $\theta_B = c(v) + v$. Private participation falls from one to zero. The associated increase in expected implementation time is $1/2 + 2v$ privately and $2v$ under the planner. At the adverse innovation, withdrawal therefore adds exactly one-half period of expected delay and lowers net surplus by exactly $1/2$, relative to coordinating all programs at the same innovation. These losses remain positive as the difference between the two innovations tends to zero.

The result combines equilibrium selection, successful eventual adoption, and an inefficient change in timing. For $v = 0.01$, the private response is 0.52 period and the coordinated response is 0.02, a ratio of 26. The excess response, 0.50, does not rely on a ratio with a small denominator. The appendix gives the full proof; the welfare and timing claims follow directly from (30) and (31).

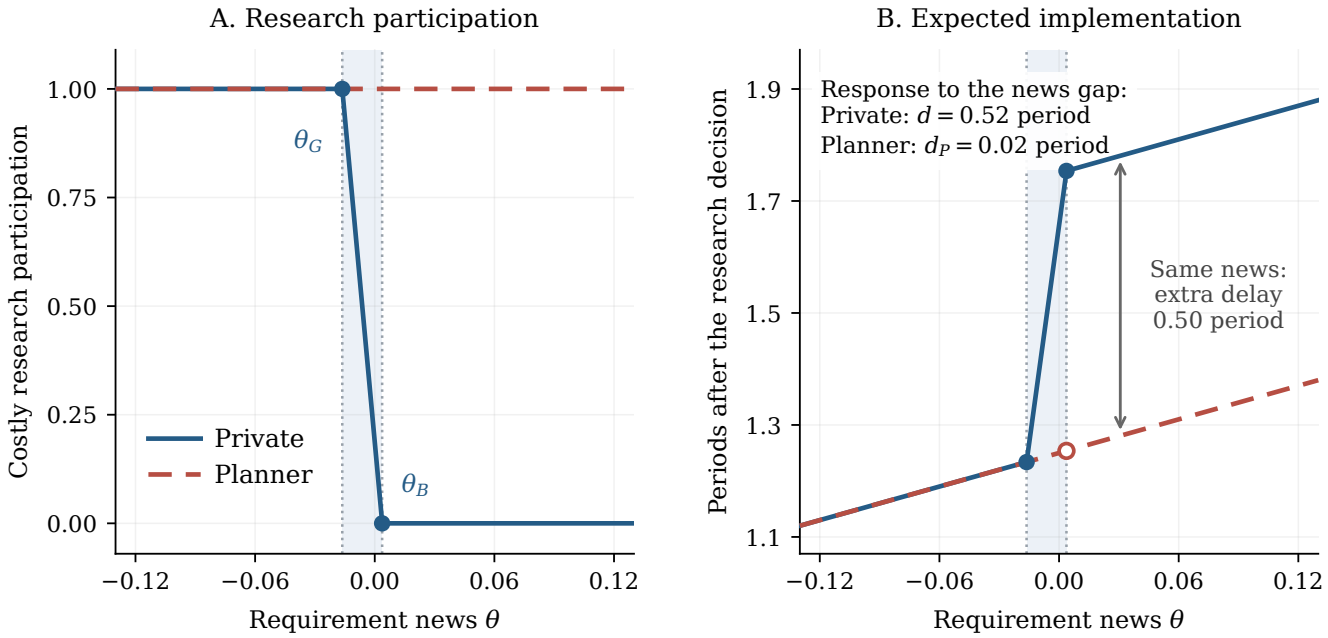

Figure 5: Coordinated withdrawal delays a technology that succeeds under every policy

Notes: Exact two-period example at $\nu = 0.01$. Panel A plots the unique private equilibrium and the planner's constant research policy. Panel B translates those policies into expected implementation dates. The shaded interval joins $\theta_G = c(\nu) - \nu$ and $\theta_B = c(\nu) + \nu$. Its width is 0.02, but private expected implementation moves back 0.52 period, compared with 0.02 for the planner. The same-fundamental policy gap is 0.50 period after withdrawal. There is one research decision at date 0 and a fixed prize paid at date 1 or 2.

Why survivors can gain share while their equity falls. The core owns 1/2 capability after its costless program. Its eventual market share rises from 1/3 when costly firms research to 1/2 when they pause. Yet its equity value falls because the common prize is paid later. At the limiting cutoff, its value is approximately 1.089 under universal research and 0.883 under universal withdrawal, a decline of 18.9 percent. Both prices are expected discounted receipts from the same fixed prize. The example therefore separates conquering a larger eventual share from preserving the value of the original claim.

Figure 6 puts the ownership and valuation effects side by side. The core creates the same capabilities under either policy. Its larger eventual share after withdrawal is therefore an ownership gain, while the fall in its current claim reflects the later common payment. The comparison holds the aggregate prize fixed and uses the same requirement news for the private and coordinated allocations.

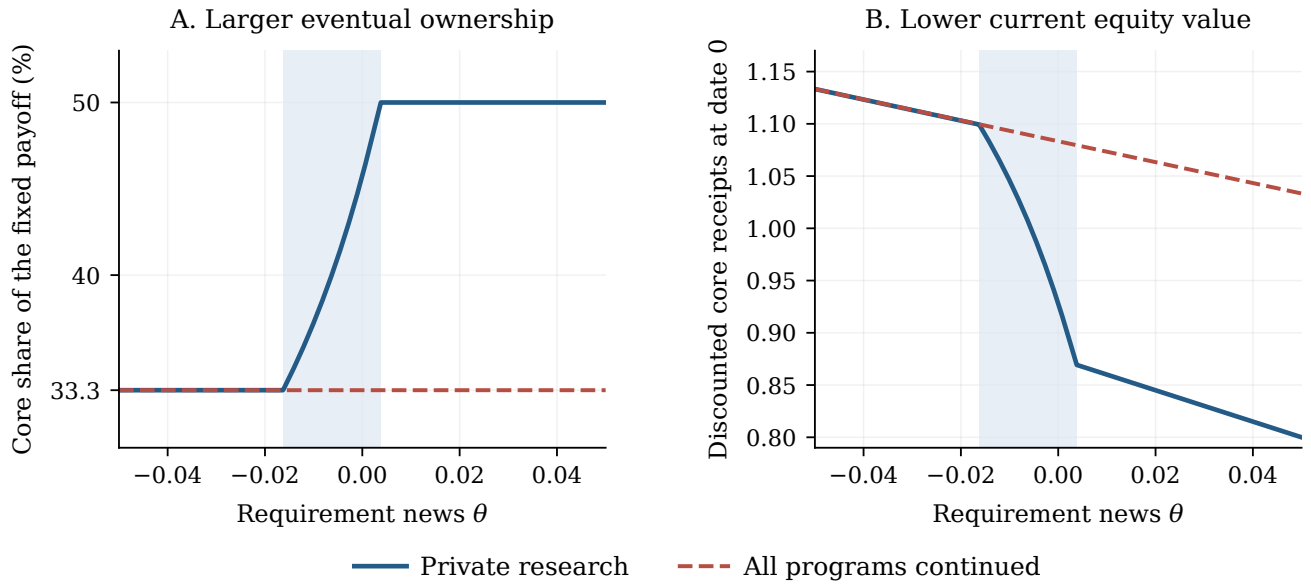

Figure 6: A larger eventual share can accompany a lower fundamental equity value

Notes: Exact two-period example, $\nu = 0.01$, with fixed aggregate payoff $m = 16$. Panel A: the core's terminal ownership share, $(1/2)/(1 + A(\theta)/2)$. Panel B: its date-0 expected discounted receipts, $(1/2)Q(\theta, A(\theta))$. The core researches throughout. Solid curves use private equilibrium participation; dashed curves continue all costly programs. Shading marks the same news interval as Figure 5. Both allocations deliver the technology by date 2. The dashed allocation is a coordination benchmark, not a subsidized equity return.

4.2 Clustering across heterogeneous research costs

Now divide the costly sector into three equal capability groups with research costs $z_1 = 1.9$, $z_2 = 2$, and $z_3 = 2.1$. Each owns one-third of F. The work technology, common payoff, and information structure remain those of the preceding example. A type-i firm researches when its expected capability return exceeds z_i. If c_i is its cutoff, aggregate participation is

$$A(\theta) = \frac{1}{3}\sum_{i=1}^{3}\left[\frac{c_i - \theta + \nu}{2\nu}\right]_0^1. \tag{32}$$

Lower costs imply $c_1 > c_2 > c_3$: cheaper firms continue under less favorable signals. The issue is whether these differences persist as information becomes precise, or whether the types withdraw at a common fundamental.

Cost differences suggest a gradual retreat: expensive programs close first, leaving cheaper programs in operation. That reasoning holds the remaining programs' returns fixed. Here each departure also postpones the common payoff. If the loss of complementary research is large enough, a cheaper program

cannot profitably continue after the expensive ones have stopped. To determine whether exits separate, we must therefore compare costs with returns in the different participation environments each marginal type expects.

Different marginal beliefs, one belief constraint. With distinct cutoffs, a firm's signal rank no longer equals total participation. At its own cutoff, a low-cost firm can expect many higher-cost firms to have already paused. A high-cost firm requires more participation to justify its program. Their marginal beliefs therefore differ, even when the cutoff signals become close in levels.

The belief constraint links these assessments. For every bounded continuous function f,

$$\frac{1}{3}\sum_{i=1}^{3}\mathbb{E}[f(A(\theta)) \mid s_i = c_i] = \int_0^1 f(a)\,da. \tag{33}$$

Thus the capability-weighted mixture of marginal participation beliefs is uniform. Each type is evaluated at its own cutoff signal. The identity follows because its posterior density is the rate at which its participation falls as the fundamental worsens; adding those densities gives $-A'(\theta)$, and changing variables from θ to $a = A(\theta)$ gives the right-hand side. Appendix A.1.1 supplies the derivation. Taking $f(a) = a$ makes the average marginal expectation one-half, but individual types need not expect one-half. The full distributional restriction is needed for the nonlinear return. This identity is the belief constraint of Sákovics and Steiner (2012). Drozd and Serrano-Padial (2018) extend it to subsets of types; that version, applied below, determines which types pool.

The cluster in numbers. Table 2 reports the three marginal firms' beliefs. At $\nu = 0.01$, the low-cost type is indifferent while expecting participation of 0.290; the high-cost type requires an expectation of 0.718. Both act optimally: their respective expected capability returns equal 1.9 and 2.1. As noise vanishes, their cutoff signals converge, while their marginal expectations remain different. The low-cost type expects participation of 0.258 and the high-cost type 0.753. Differences between cutoff signals shrink with the noise, preserving differences in how each marginal type ranks itself relative to the other types' withdrawal thresholds.

Table 2: Different costs and distinct marginal beliefs

Cost z_i	Cutoff $c_i(0.01)$	$\mathbb{E}_i[A]$, $\nu = 0.01$	$\lim_{\nu\to 0} \mathbb{E}_i[A]$
1.9	-0.001016	0.290	0.258
2.0	-0.006241	0.492	0.489
2.1	-0.012146	0.718	0.753

Notes: Each type owns one-third of strategic capabilities. Each expectation conditions on the corresponding type-specific cutoff signal. The last column reports the small-noise limit of the displayed equilibrium branch; all cutoffs converge to $\theta^* = -0.005616$.

Proposition 2 (Heterogeneous costs and a common withdrawal boundary)**.** *In the three-cost example with $0 < \nu < 7/60$, the equilibrium is unique up to choices at indifferent signals, and it is in monotone cutoffs. As ν tends to zero, the cutoffs of all three types converge to*

$$\theta^* = \frac{1}{6\log(3/2)} - \frac{5}{12} \simeq -0.005616. \tag{34}$$

Aggregate participation converges to one below this boundary and zero above it. The planner continues all three types at every fundamental. Full private withdrawal adds one-half period of expected delay and lowers surplus by one-half relative to the planner at the same fundamental, while implementation by date 2 remains certain.

At this boundary, a capability is worth 1.767 under universal withdrawal, below even the lowest cost, and 2.178 under universal research, above even the highest cost. Continued participation sustains every type; collective withdrawal makes every type prefer to pause. Away from this region, incentives are unambiguous: all types prefer research for $\theta < -7/60$ and pausing for $\theta > 2/15$, regardless of others' decisions.

To see why separate exits fail, suppose the types withdraw separately as news worsens. The expensive type must leave first, while both cheaper groups continue. At its marginal signal, it would therefore face participation between two-thirds and one. The cheapest type must leave last, after the other groups have paused, and would face participation between zero and one-third. The intermediate type would face participation between one-third and two-thirds. Each type is uncertain about its own group's participation because its signal rank is uniform, just as in the homogeneous example.

These different strategic environments can outweigh the differences in costs. Applying each type's indifference condition under the proposed separation gives boundaries of -0.020229, -0.000618, and 0.007843 for costs 1.9, 2, and 2.1. This would make the cheapest type leave first, contradicting the incentive ordering that required it to leave last. Pooling only two groups yields the same contradiction. There is consequently no range of fundamentals over which the proposed separate exits can be sustained. Withdrawal by expensive programs reduces the return to cheaper programs enough to bring their exits together. This is the ordering argument in the three-type example of Drozd and Serrano-Padial (2018), applied to commercialization rather than enforcement. It answers the question posed in Section 2.3: adverse news does not remove the most expensive programs first and the rest later, but takes out costly research together, while the core continues. The full proof, including existence, uniqueness at finite noise, and all possible partitions, is in Appendix A.1.1. Clustering is not a general consequence of heterogeneity alone: in this example complementarity is strong enough relative to the dispersion in costs to prevent separate exit regions. Greater cost dispersion or weaker acceleration can separate those regions. The proposition concerns the stated three-cost economy, while the quantitative recursion below maintains a common current cost.

The result persists over an open range of heterogeneity. With costs $2 - d$, 2, and $2 + d$, all three types cluster whenever $0 < d < 0.128534$ and private noise is sufficiently small. Moreover, withdrawal lowers every type's equity value even after subtracting its research bill; the core loses value despite gaining ownership share. Appendix A.1.2 proves these statements and gives the exact dispersion bound. The contraction therefore persists with strictly different costs and consistently measured net owner values.

At the common limiting fundamental, the belief constraint and the three indifference conditions give

$$\int_0^1 Q(\theta^*, a)\, da = \frac{z_1 + z_2 + z_3}{3} = 2. \tag{35}$$

The average research cost equals that in the common-cost example, so the limiting boundary is the same. Heterogeneity is accommodated through different marginal beliefs. Improved news reverses the withdrawal mechanism: expected participation by other types raises research returns and supports a clustered expansion. The shared commercialization timeline generates both responses.

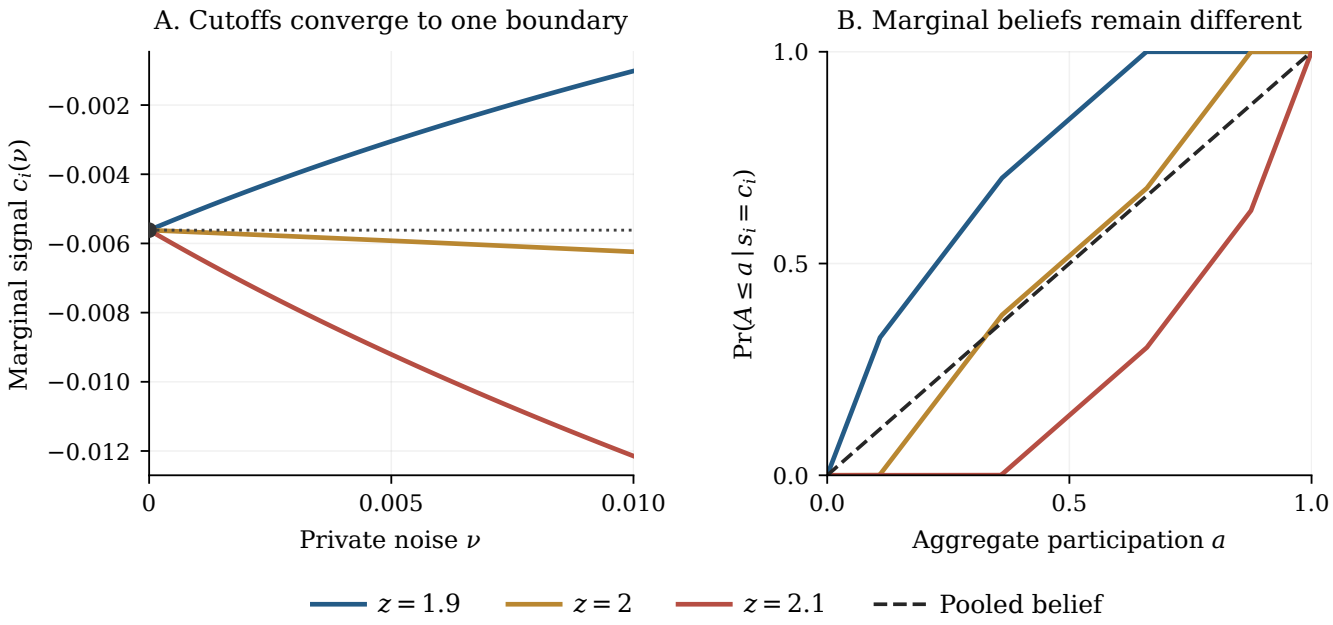

Figure 7: A common withdrawal boundary with different marginal beliefs

Notes: Three equal capability groups with costs 1.9, 2, and 2.1. Panel A shows their equilibrium cutoff signals as private noise varies; the dotted line is θ^*. Panel B shows limiting participation cdfs conditional on each type's own marginal signal. The cheaper type puts more probability on low participation, and the expensive type on high participation. Their equal-weight mixture is the uniform diagonal, as required by (33).

4.3 Dynamic propagation through capability depreciation

Static clustering makes the initial withdrawal abrupt. Its subsequent effects depend on the physical state inherited by the next research game. A pause reduces the creation of capabilities while the existing stock continues to depreciate. This endogenous deterioration can lower future continuation returns and make reentry require more favorable news: the depreciation spiral. Two losses need to be distinguished. Depreciation destroys part of the inherited stock under either action; pausing also forgoes the expansion that would have offset that loss. Remaining work can still decline during a pause because the surviving stock continues working. What deteriorates unambiguously relative to continued research is the industry's productive position: fewer capabilities survive and less work has been completed. The effect on research incentives additionally depends on dilution, as the condition below makes explicit.

Inaction is not neutral. A costly firm's pause stops its creation of new capabilities while its inherited stock continues to depreciate. This loss represents the decay of specialized teams, the scattering of supplier relationships, and the obsolescence of prototypes. Relative to continued research, withdrawal

leaves both fewer costly capabilities and more work unfinished. If completion has not occurred,

$$\frac{R'}{C' + F'} = \exp(\delta) \left[\frac{R + \theta}{N(a)} - h \right]. \tag{36}$$

Lower participation strictly raises this burden per unit of surviving capacity. The core may offset part of the costly sector's contraction, but it cannot undo the work forgone in that period. Moreover, the next program can expand a firm's stock only in proportion to what survives. Information can improve immediately; rebuilding lost productive capacity takes time.

For orientation, the numerical model below uses $b = 0.5$ and $\delta = 0.3$. Each inherited costly capability leaves $\exp(-0.3) = 0.741$ capabilities next period if its program pauses, and $1.5 \exp(-0.3) = 1.111$ if it runs. Resuming at the following date applies that same expansion factor to the smaller stock. After a pause followed by research, the stock is only two-thirds of what two consecutive research programs would have produced. The difference is a lost productive position, even when both firms subsequently choose the same action.

How complementarity evolves. The continuation return depends on the capabilities and unfinished work inherited from earlier decisions. In a differentiable continuation region, subscripts denote derivatives with respect to the named state variable, and

$$Q_a = \beta \exp(-\delta) b F \big[\exp(-\delta) v_F(S') - h v_R(S') \big]. \tag{37}$$

When $v_R < 0$, the term $-h v_R$ is the benefit of clearing more work. The term $\exp(-\delta) v_F$ incorporates the effect of a larger inherited costly sector, including dilution and future research. Their sum determines whether research is complementary. Once this period's completion is assured, further research cannot advance the payment date and only changes ownership shares:

$$Q_a = -\frac{\beta m b F}{N(a)^2} < 0. \tag{38}$$

Private research can still be profitable because it acquires ownership claims. Losing complementarity is therefore not a rule for shutting down investment. The constant-capacity calculation $r\tau = 1$ in the appendix gives the same acceleration–dilution tradeoff in a simpler setting.

Why the restart condition moves. A pause changes the physical state at which the next game is played. The interaction between G_F, which captures the continuation and dilution effects of inherited capabilities, and G_R, which captures unfinished work, determines how that change affects the marginal firm. When preserving capacity and clearing work raise its continuation return, withdrawal tightens the restart condition. This is the mechanism behind the depreciation spiral.

Proposition 3 (Withdrawal tightens the restart condition)**.** *In the common-cost specialization, fix the current state and news. Suppose participation choices $a_L < a_H$ both leave commercialization incomplete. Along the induced next states $S'(a)$, suppose the continuation game has a differentiable cutoff $c(S')$, with $G_c < 0$, and*

$$\exp(-\delta)G_F - hG_R > 0. \tag{39}$$

The derivatives concern the next physical state, holding the candidate signal cutoff fixed and using equilibrium continuation values. Then $c(S'(a_L)) < c(S'(a_H))$: after greater withdrawal, firms require more favorable news to restart. For any $k > 0$, if the cutoff gap is at least $2kv$, the same next innovation can generate participation at most $\Phi(-k)$ after a_L and at least $\Phi(k)$ after a_H.

Proof. An extra unit of current participation leaves $\exp(-\delta)bF$ more costly capabilities and hbF less unfinished work next period. It therefore raises the marginal firm's continuation surplus, at a fixed candidate cutoff, by $bF[\exp(-\delta)G_F - hG_R]$. Differentiating $G(c(S'(a)); S'(a)) = 0$ gives

$$\frac{dc(S'(a))}{da} = -\frac{bF[\exp(-\delta)G_F - hG_R]}{G_c} > 0. \tag{40}$$

A higher return permits a less favorable marginal signal because $G_c < 0$. Integrating over $[a_L, a_H]$ orders the two cutoffs. Writing them as $c_L < c_H$, any next innovation in $[c_L + kv, c_H - kv]$ gives the participation bounds through (15). □

The condition compares two opposing effects of preserving capacity. More capabilities and less unfinished work bring the common payoff closer; additional claims can also dilute each capability's receipts. When the first effect dominates, the productive position lost during a pause makes the restart threshold more demanding. An improvement in news then confronts a different industry. It neither recreates the missing organizational knowledge nor completes the work forgone. Even when research resumes, its scale is limited by the smaller inherited stock. The core keeps making progress and can help restore incentives to invest; no restart date is imposed. Appendix A.3.1 shows explicitly why reversing a temporary requirement shock leaves these physical losses in place.

A growing core, $g = \exp(-\delta)(1+b) > 1$, guarantees eventual completion almost surely despite strategic withdrawal. The depreciation spiral can therefore prolong commercialization without making technological failure permanent. Appendix A.2 develops the recursive construction and the conditions under which its current games select a unique continuation boundary.

5 Quantitative Analysis

The recursive game measures how coordinated withdrawal changes commercialization, continuation incentives, and equity values. The quantitative exercise uses the common-cost specialization, with H a point mass at $z = 2.5$. The 1870s railroad contraction disciplines the model's initial composition; a discount-rate convention gives it a multidecade calendar interpretation. Private decisions and coordinated allocations then face the same sequence of requirement innovations.

5.1 Parameters and historical scale

Initial capabilities sum to one. The core's share is chosen to match the 78.41 percent fall in first-track completions between 1872 and 1875. The model counterpart is the contraction in total new-capability creation, $b(C + FA)$, across nearby news realizations $\theta_G = c - 2v$ and $\theta_B = c + 2v$. Costly participation is $\Phi(2)$ and $\Phi(-2)$, respectively. With $q = 1606/7439$ denoting the historical trough-to-peak ratio, the fitted share is

$$C = \frac{q\Phi(2) - \Phi(-2)}{1 - \Phi(-2) - q[1 - \Phi(2)]} = 0.193582, \qquad F = 1 - C.$$

Physical buildout disciplines capability creation. The core captures the continued creation of capabilities through other activities during a research contraction. The costly resource bill, $bzFA$, is recorded separately. The model contains no financial accelerator, so matching this moment attributes the full 1872–75 contraction to the commercialization mechanism. This fit does not identify its historical contribution: the same investment contraction may reflect several forces. The remaining outcomes assess what the fitted mechanism implies, rather than provide independent estimates of its share of the crash.

We set $b = 0.5$, $\delta = 0.3$, $z = 2.5$, $m = 100$, $\sigma = 0.5$, and $\nu = 0.05$, with initial remaining work $R = 8.7051$. An annual real reference discount rate of 8 percent maps $\beta = \exp(-0.45)$ into $0.45/\log(1.08) = 5.847$ years per model period. Private firms and the planner share this discount factor. Mean commercialization occurs 41.22 years after the initial state, placing the transition on an approximately forty-year horizon. Figure 8 provides the historical scale: freight per unit of real GDP first crosses half its observed maximum twenty-eight years after 1872 and reaches that maximum forty-five years after the peak.

The e-commerce series supplies a second reference for the commercialization horizon. Its observed share first exceeds half of a fitted 24.52 percent long-run reference in 2020, twenty years after the technology valuation date. The fitted path reaches 90 percent in 2034. The railroad maximum is observed; the e-commerce reference and future crossing depend on the logistic fit. These paths place the build-out on a multidecade scale, while the model's completion distributions describe variation across simulated histories. An observed diffusion path and a distribution of discrete completion dates are different objects. We do not fit the model's initial work requirement to either crossing, and a logistic extrapolation is not evidence about a realized future adoption date.

Table 3 distinguishes the fitted investment contraction from the model's other implications. The initial composition fits the depth of the contraction; the calendar convention sets the scale of the transition. The completion-date 90th percentile is 64.32 years across simulated histories. Core capabilities grow at 1.82% per year and inherited capabilities depreciate at 5.00% per year under this calendar mapping.

The model's rate is real; the historical valuation calculations use nominal rates. Deflation makes the railroad 8 percent nominal rate approximately 9.65% in real terms over 1872–96. Applying that real rate to the same dimensionless model gives 4.883 years per period and mean completion of 34.43 years. Appendix A.4 reports the discount-rate comparison.

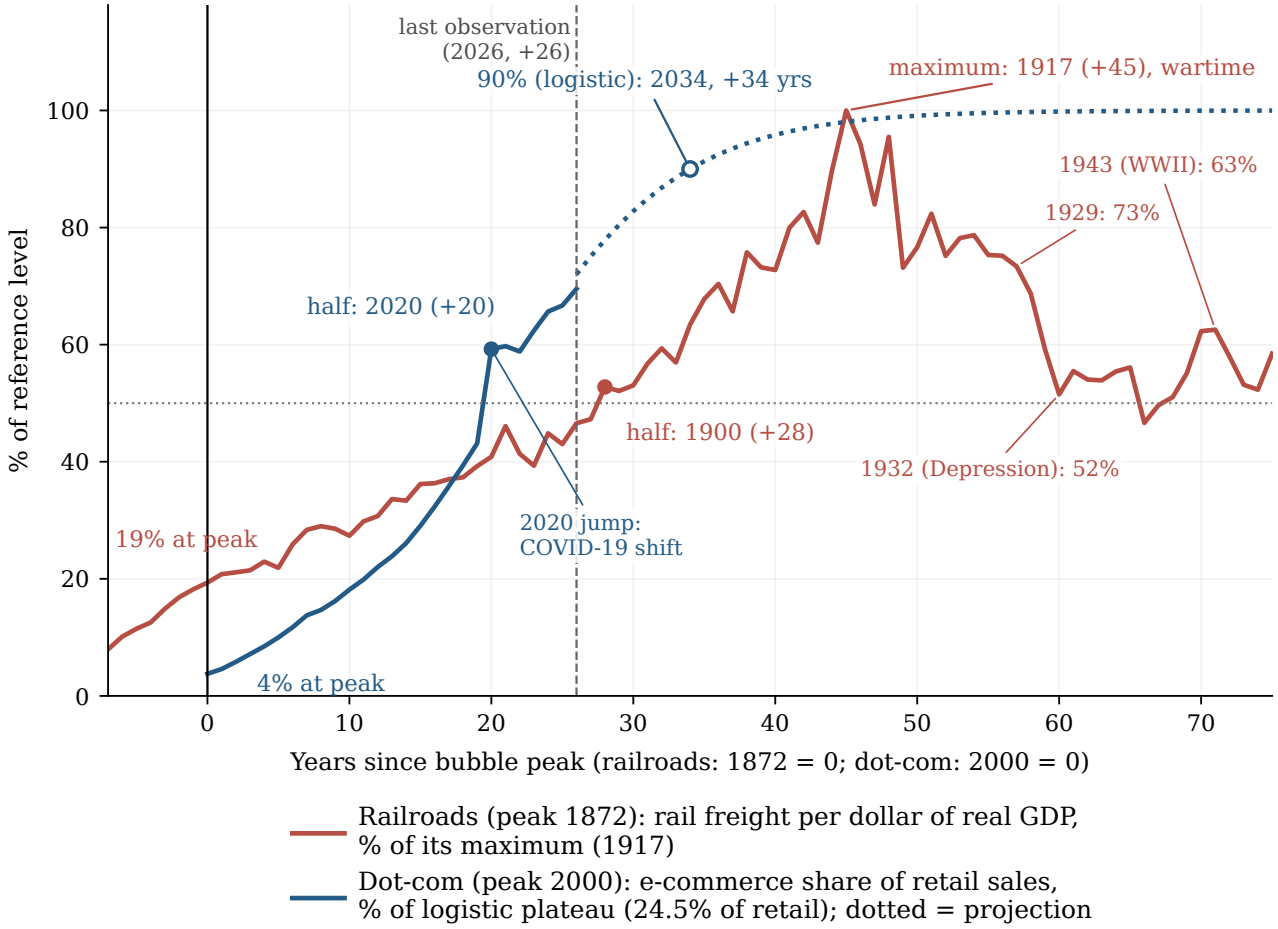

Figure 8: Commercial use since the reference valuation year

Notes: Rail: CHAT freight tonne-kilometers divided by MeasuringWorth real GDP, scaled by the observed 1917 maximum. E-commerce: Census retail share divided by the fitted logistic reference of 24.52 percent. The 2026 observation is the first-half average. The dotted continuation and hollow 90 percent marker are conditional projections. These within-episode paths measure commercial use; model completion dates instead vary across simulated histories.

Table 3: Calibration targets and model implications

Moment	Reference	Model	Role
Capability-flow contraction	78.41%	78.41%	Fits C
Unconditional mean completion	—	41.22	Model implication
Adverse-news mean completion	~ 30 years	50.26	Tentative scale check
Completion-date 90th percentile	—	64.32	Untargeted
Recovery, observed-history median	9 years	17.54	External comparison

Notes: Initial total capabilities are normalized to one. The core stock is fitted from the contraction in railroad construction. Initial work and the discount convention set an approximately forty-year scale. Technology and information parameters are maintained. Outcomes use 20,000 simulated histories; completion quantiles describe dispersion across histories. The recovery median conditions on recovery before completion.

5.2 The participation cliff and inefficient delay

A change in the remaining requirement of 0.2—0.4 prior standard deviations, or 2.30% of initial work—lowers costly participation from 97.72% to 2.28%. Total capability creation falls by 78.41%, and the initial costly research bill falls by 97.67%. Each economy then follows its own equilibrium policy and physical state under exactly the same future innovations.

Under private decisions, mean commercialization moves from 34.30 years after the initial state to 50.26. Under the causal planner, the corresponding means are 34.21 and 34.76. The planner observes each current innovation and its stochastic process, without knowing future realizations.

The adverse news therefore adds 15.96 years to private commercialization, compared with 0.56 years under the planner. The amplification ratio is

$$M = \frac{\text{additional private delay}}{\text{additional planner delay}} = 28.55. \tag{41}$$

The paired Monte Carlo standard error is 0.90. In levels, the private response exceeds the planner's response by about 15.40 years; reporting that difference makes the size of the effect visible independently of the small planner response in the denominator. The comparison holds future news fixed. The participation cliff is the static response at the initial state; the subsequent depreciation spiral changes the states and continuation returns governing later research. The delay ratio reflects their combined effect.

The surplus loss extends beyond this local comparison. Across 20,000 histories drawn from the initial prior, private expected surplus is 3.795, compared with 4.828 under the planner: a 21.40% shortfall. Decentralized research delivers a later prize and a different sequence of resource costs. Section 5.5 reports the coordinated allocation based only on the public physical state alongside these outcomes.

The initial cutoff lies 0.35 prior standard deviations above the prior mean. The prior probability of participation below ten percent is 31.520%. The conditional news comparison measures the size of the response at the boundary; this probability describes the mass of initial innovations on its withdrawal side.

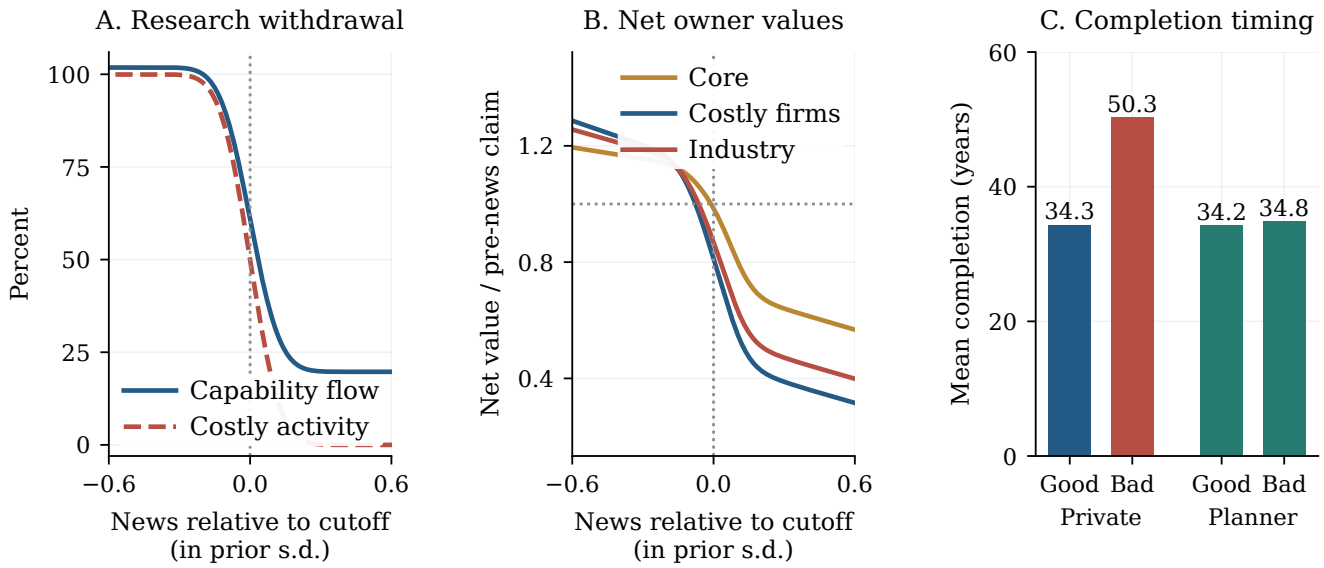

Figure 9: A small requirement change crosses the private continuation boundary

Notes: The horizontal news scale is centered at the initial equilibrium cutoff and measured in prior standard deviations. Panel A plots costly participation and total new-capability creation, the latter relative to its favorable-news level. Panel B reports net owner values, including the current research bill, on a common pre-news base. Panel C compares expected completion at $\theta_G = c - 2\nu$ and $\theta_B = c + 2\nu$ using 10,000 paired histories. Future innovations are identical across the histories being compared. The planner knows current news and its distribution, but not the future realizations.

5.3 Dynamic amplification through the depreciation spiral

The continuation boundary is a function of inherited capabilities and unfinished work. A firm that waits does not preserve its productive position: its next stock is $\exp(-\delta)n$ rather than $\exp(-\delta)(1+b)n$. At the retained parameters, maintaining the costly sector's stock requires participation of at least $[\exp(0.3) - 1]/0.5 = 0.700$. A partial restart below that rate still allows its capabilities to shrink.

After the adverse initial realization, costly capabilities fall from 0.806 to 0.604. Remaining work itself falls because the core and surviving capabilities keep working, but work per surviving capability rises from 8.71 to 9.80. The cutoff falls from 0.1765 to 0.0602. Since research occurs below the cutoff, this is a stricter condition for returning.

A same-news comparison isolates the effect of current withdrawal. Hold the adverse realization fixed and coordinate full research for this period only; then let subsequent private play determine the continuation. Net owner value rises from 1.956 to 4.136 even after deducting the larger current research bill. The next cutoff is 1.8399, compared with 0.0602 after equilibrium withdrawal. Thus the next research game differs

sharply even with identical requirement news. The change comes through capabilities and unfinished work, not an assumed persistence of pessimistic beliefs.

Repeating the original lower-hurdle realization $\theta_G = 0.0765$ after withdrawal induces only 37.2% participation next period; the same news induced 97.72% before the collapse. Returning to more favorable news need not restore the original investment response. Figure 10 shows this comparison directly: the same news intersects different participation curves because the industry has changed. With zero subsequent innovations, however, participation already reaches 88.6% at the next decision. Better news can restart research; restarting on a depleted capability base does not replace the work or productive capacity lost in the meantime.

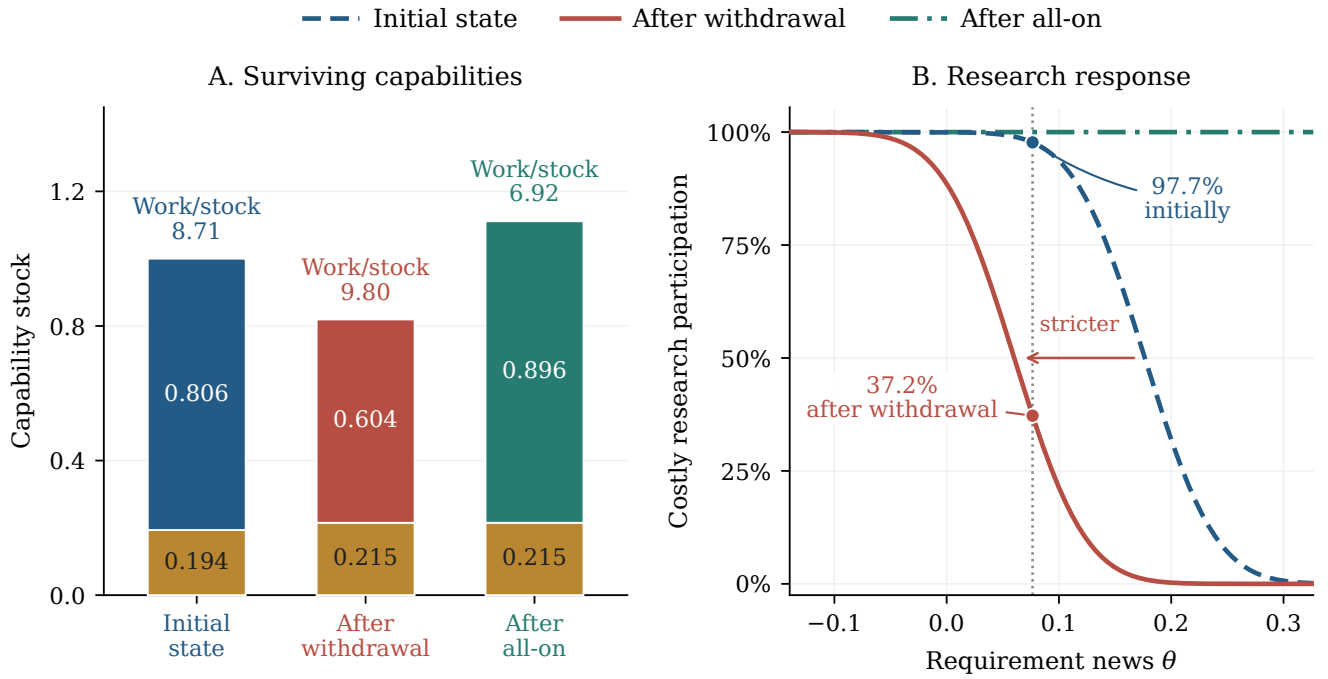

Figure 10: Same news, different industry

Notes: Panel A compares the initial stock with next-period stocks following identical current news, $\theta_B = 0.2765$. The selected equilibrium has costly participation of 2.275%; the coordinated comparison sets participation to one for the current period and retains subsequent private policies. Gold segments are core capabilities; annotations give remaining work divided by total surviving capabilities. Panel B plots $A(\theta) = \Phi((c - \theta)/v)$ at each public state. Research requires $s < c$, so the fall in the cutoff from 0.1765 to 0.0602 tightens the restart condition. The same realization $\theta_G = 0.0765$ (dotted line) produces 97.725% participation initially and 37.226% after withdrawal. Following coordinated research, the cutoff is 1.8399, outside the displayed range. These are conditional games at different inherited states.

Figure 11 follows one history selected by a stated rule, not by its crash depth. The initial participation cliff is followed by 18.5% participation at the next date, then almost full research at date 2, or 11.69 years.

The cutoff initially tightens and then relaxes as the full physical state evolves. Along this history, the burden $R/(C+F)$ rises slightly again between dates 1 and 2 even though the cutoff improves. That ratio helps explain the damage but is not a sufficient state: composition, absolute scale, and continuation opportunities also matter.

The core keeps contributing work during the pause. Together with subsequent research, it eventually supports implementation. The adverse-news economy completes at date 8, against date 6 with the lower initial hurdle and the same later innovations, a difference of 11.69 years. The long postponement includes a short initial pause and slow rebuilding after participation recovers. The model does not require every firm to remain inactive until the commercialization date.

The boundary as commercialization approaches. Remaining time is itself an equilibrium outcome. At each public state on the illustrated histories, we compute $\mathbb{E}[\tau - t \mid S_t]$ using fresh current and future innovations and the private continuation policy. The forecast does not use the history's subsequently realized news. Figure 12 plots the cutoff against both elapsed time and this expected remaining time. The initial forecast is about 40.94 years. After the adverse first period, 5.85 years have elapsed, yet the expected remaining wait rises to 44.45 years. After lower-hurdle initial news, the corresponding wait falls to 28.46 years. Withdrawal has changed the physical state enough to move expected arrival further away even as calendar time advances. Including the elapsed first period, the adverse-state forecast places completion 50.29 years after the initial withdrawal decision. This conditional wait is distinct from the unconditional horizon at the pre-news initial state.

The cutoff gives this change an immediate risk interpretation. Since $A(\theta) = \Phi((c-\theta)/\nu)$, the public-state probability of costly participation below ten percent is

$$\Pr\{A(\theta) < 0.1 \mid S\} = 1 - \Phi\left(\frac{c(S) + \nu\Phi^{-1}(0.9)}{\sigma}\right).$$

It rises from 31.52% initially to 40.18% after the adverse period. After lower-hurdle initial news it is only 0.00154%. Along the subsequent adverse history, renewed research and completed work eventually raise

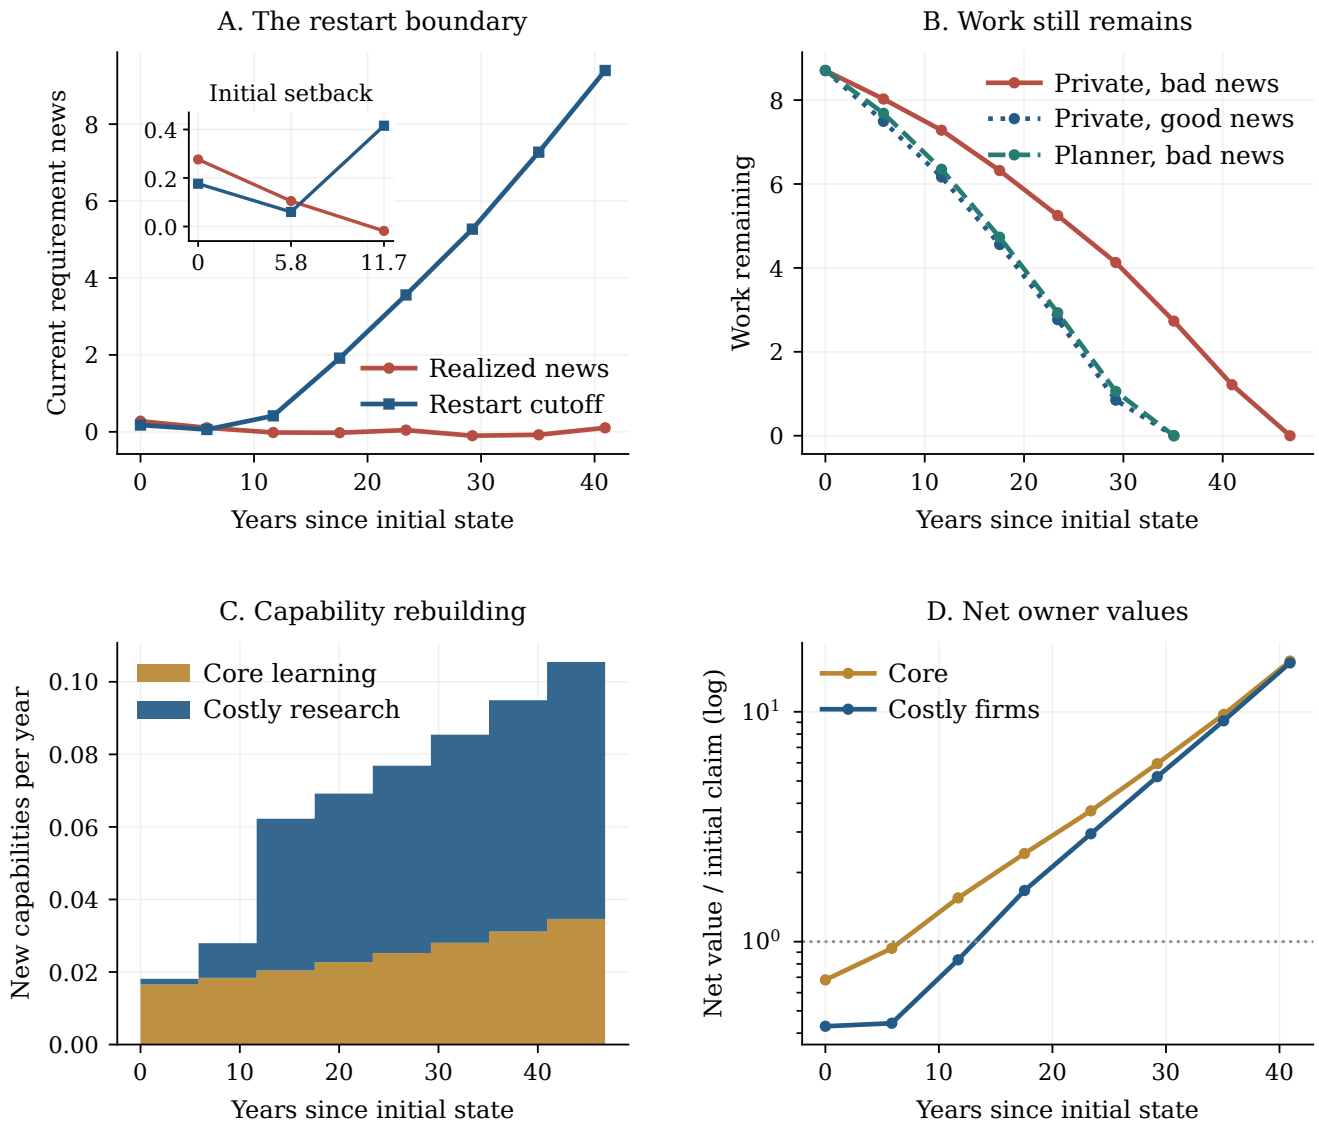

Figure 11: The restart boundary moves because the productive state changes

Notes: One paired history from the median positive additional-delay bin, chosen within that bin to minimize the largest absolute later innovation. Panel A plots news and the private cutoff; research is attractive when news lies below the cutoff. Its inset resolves the initial tightening. Panel B compares remaining work under the two initial news realizations and the causal planner. Panel C separates core capability creation from costly research after adverse news. Panel D follows core and costly net owner values relative to their initial pre-news claims, on a logarithmic scale. These are contemporaneous price ratios, not discounted realized returns. Each economy keeps its own state while receiving the same later innovations. The figure illustrates a history, not an average response.

the cutoff and reduce this risk. The two paths do not define a universal relation between age, expected remaining time, and fragility: capability composition and scale remain part of the state.

Late cutoffs also require care. A threshold deep in the adverse tail means ordinary news sustains research; it does not describe an observed late crash. Filled points in the figure additionally satisfy $Q(c, 0) < z <$

$Q(c, 1)$ at the displayed boundary, holding future private policies fixed. The last favorable point fails that endpoint test and is hollow. Even a positive endpoint test at a very unlikely requirement realization need not imply substantial current crash risk. Near completion, research can remain privately attractive through its ownership payoff even when advancing the payment date no longer supplies strategic complementarity, as (38) shows.

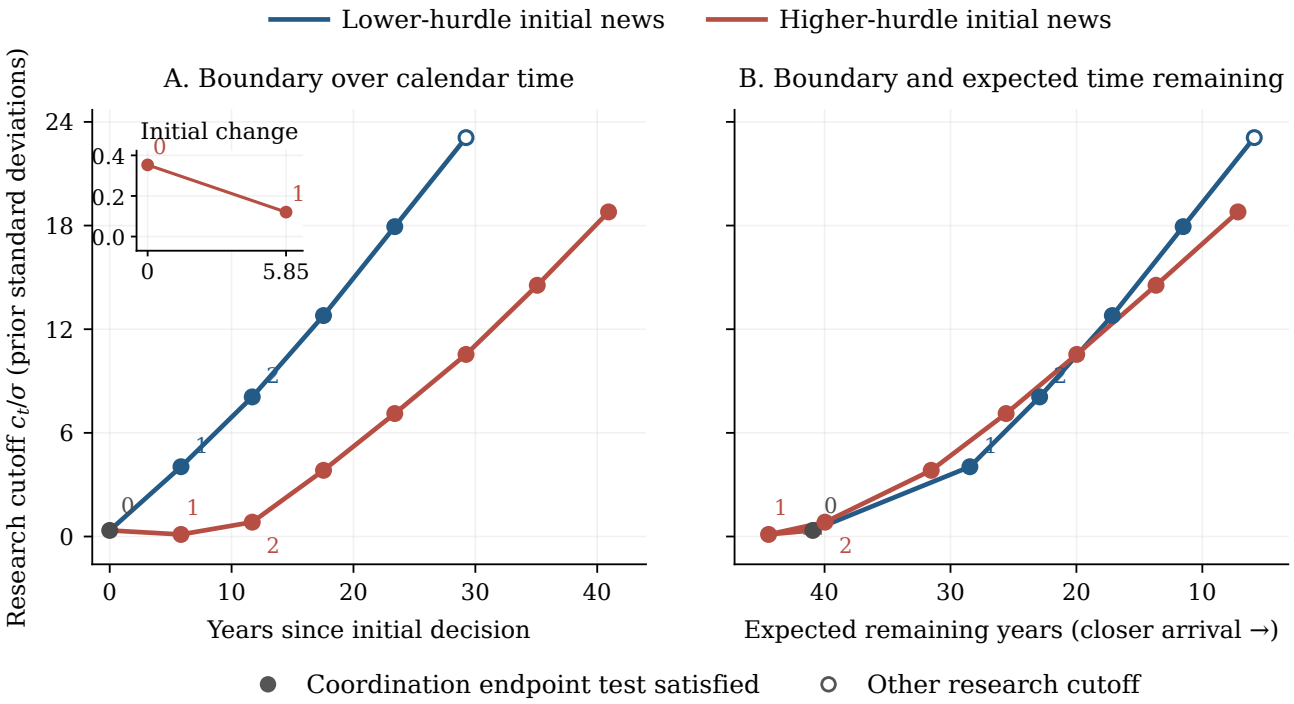

Figure 12: The withdrawal boundary changes with the expected commercialization horizon

Notes: Public states on the same paired history as Figure 11. Panel A plots the signal cutoff in prior-standard-deviation units. A higher cutoff permits research under worse news; the inset shows the initial tightening after withdrawal. Panel B uses expected remaining years, with closer arrival to the right. Each forecast integrates fresh current and future news over 1,024 common antithetic histories; horizontal bars show ±1.96 Monte Carlo standard errors, whose maximum is 0.318 year. Filled points satisfy $Q(c, 0) < z < Q(c, 1)$; the hollow point is a research cutoff outside that endpoint coordination region. All displayed current games and forecast decisions pass the recorded incentive checks. Labels 0, 1, and 2 denote decision dates. The lines stop at the last pre-completion decision; no research boundary is defined after commercialization. Expected remaining time alone is not a sufficient state variable.

Two reinforcing forces. The private delay reflects static informational clustering and the deterioration of the knowledge base acting together, and the paired histories show how each feeds the other. Clustering takes nearly all costly firms out at once, so a large part of their capabilities depreciates in the same period; the deteriorated base tightens the next continuation boundary; and the industry then clusters out again at news it would previously have absorbed. After the favorable news, costly firms

research at the next decision in 99.99 percent of the histories, and a later clustered withdrawal occurs in 0.01 percent. After the adverse news, research resumes at the next decision in only 54.9 percent, and 45.2 percent of histories contain at least one later decision at which costly participation falls below one-half; those histories average 2.5 such decisions, about 15 years. Later participation remains clustered: it lies between 10 and 90 percent at only 2.5 percent of later decisions. Because each force operates through the other, their separate contributions to the delay are not additively defined.

The role of depreciation. The depreciation rate is not estimated, so Appendix A.5.2 solves the model again at half the calibrated rate. Whether initial work is held fixed or raised to keep the mean horizon, the same news at the withdrawal boundary then postpones private commercialization by 5.3 to 5.5 years rather than 16.0, and a later clustered withdrawal follows in 3 to 7 percent of adverse-news histories rather than 45 percent. Depreciation is what makes a withdrawal persist: with less of it, a pause is cheap to reverse. At the same horizon, pauses also become more frequent, and the average surplus shortfall rises to 35.3 percent. Lower depreciation thus changes the form of the inefficiency, from rare and persistent crashes to frequent and short pauses, rather than removing it.

5.4 Equity, funding calls, and realized recovery

The cash-flow convention is common across all firms and all comparisons. The post-commitment prices P_C and P_F are gross of the current costly research bill. Net owner values are P_C for the core and $P_F - bzFA$ for costly firms; industry owner value is their sum. The pre-news claims Cu and Fv already deduct all expected current and future research bills. From favorable to adverse news, net core value falls by 40.12%, costly value by 63.12%, and industry value by 55.64%. These are comparisons between two news realizations at the same initial state. The announcement return from the pre-news expected industry claim to the adverse-news net owner value is instead a 48.81% decline. Comparing a gross post-commitment price with a net pre-news claim would mix the funding convention with the effect of news.

The costly-sector calculation makes the distinction concrete. Under the lower requirement, its gross post-commitment claim is 3.9592, but the current funding call is 0.9851, leaving owner value 2.9741.

Under the higher requirement the corresponding amounts are 1.1199, 0.0229, and 1.0970. Subtracting the bill in both states gives the reported 63.12% decline. A reduction in funding needs partly offsets the fall in the gross claim; it is not a dividend or a gain in the value of an existing capability.

Table 4: Equity measured consistently before and after news

Net owner claim	Pre-news expectation	Favorable news	Adverse news
Core	1.259	1.434	0.859
Costly firms	2.561	2.974	1.097
Industry	3.820	4.408	1.956

Notes: Values use the same date-0 numeraire. The pre-news expectation deducts all expected funding calls. The news-contingent costly claim deducts its current bill; future calls are already included in the continuation value. Favorable and adverse refer to the lower and higher requirement realizations, respectively; both can exceed the prior mean. These price levels are not realized investment returns.

The model also records realized value on the same basis. For every history, J_C divides discounted core receipts by the pre-news core claim. The industry ratio J subtracts all discounted research costs from discounted terminal receipts before dividing by the pre-news industry claim. Their unconditional simulated means are 0.996 and 0.993, consistent with fundamental pricing up to numerical and sampling error. Conditional on initial costly participation below ten percent, the means are 0.538 and 0.372. This conditioning describes the histories that experience a withdrawal; it does not convert the realized recovery statistic into a test of ex ante pricing. A later increase in contemporaneous equity prices can include the normal return for the passage of time; it is not itself a recovery of discounted initial value. The core's mean terminal share rises from 19.48% after favorable news to 36.51% after adverse news. Enduring capabilities can gain share even while their discounted value falls because commercialization is delayed.

The historical justified fraction and the model's realized payoff ratio both compare discounted receipts with an initial valuation. Historical receipts include dividends and residual market value, whereas the model pays a terminal prize net of research costs. The comparison concerns the recovery of an earlier valuation after withdrawal; the historical leader portfolios do not identify the model's costless core.

Discussion. The contraction target disciplines the initial composition, while recovery and ownership provide further comparisons. In 92.14% of adverse-news histories, capability creation regains its initial

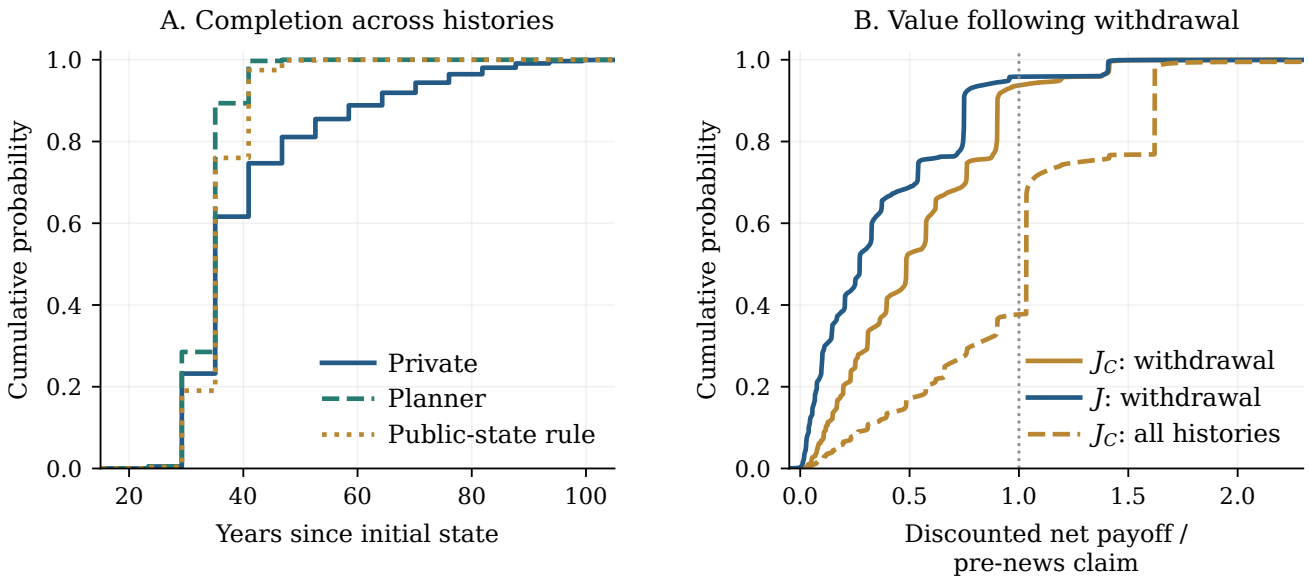

Figure 13: Uncertain completion and incomplete recovery of an earlier equity value

Notes: Panel A gives completion-date distributions across the independent simulated histories. It is not a within-economy adoption curve. Panel B compares recovery conditional on initial costly participation below ten percent with unconditional core recovery (dashed). Each ratio divides realized discounted net receipts by the pre-news price. The core is more resilient than the industry, but its conditional mean recovery is only 0.538. The corresponding unconditional pricing identities are checked separately.

favorable-news level before commercialization; the conditional median is 17.54 years, compared with nine years for railroad construction. The coarse decision interval and slow rebuilding lengthen this recovery. The model's core gains terminal ownership share, whereas the measured railroad leaders' freight share falls. Relative equity value and output shares describe different margins. These comparisons leave a richer mapping from capabilities to construction and commercial use as an important extension. The exercise is strongest as a demonstration of amplification and weaker as a joint fit of historical recovery and ownership dynamics.

The 5.85-year decision interval is coarse. A sensitivity exercise halves it, preserving annual discounting, depreciation, full-research capability growth, and the physical variance of requirement innovations. Firms then receive signals and revise programs twice as often, so the exercise changes the economic game as well as its time resolution. At the same participation endpoints, private delay remains 15.89 years and core equity falls 38.73 percent. At a fixed 0.4-standard-deviation news gap, the corresponding outcomes are 12.97 years and 33.15 percent. Initial withdrawal probability falls from 31.52 to 5.82 percent, however, and the unconditional surplus shortfall falls from about 21 to 4.78 percent. Large conditional

amplification is therefore more stable than its frequency or average surplus implications. This sensitivity uses a finite Bayesian signal approximation at some visited states; Appendix A.5.1 gives its construction and scope.

A separate numerical refinement keeps the economic model fixed. Re-solving the full recursion on a finer state grid with doubled quadrature leaves the initial valuation and matched simulation outcomes nearly unchanged. Independent integration and action-optimization checks also reconcile the sampled private and planner decisions. Some late nodes remain unresolved; the appendix reports their locations and a physical reachability bound. These checks strengthen the initial-state calculation without establishing global recursive uniqueness or uniform numerical convergence.

5.5 Coordinated investment and surplus

In the withdrawal region studied here, coordinating research preserves capabilities and brings expected commercialization forward. Continued research changes the physical state inherited by every firm: more capabilities survive and less work remains. When the acceleration benefit dominates dilution, this limits the depreciation spiral and raises the return to subsequent investment.

The relevant margin is a research program whose continuation changes that shared timeline. Support that induces such a program can also strengthen other firms' incentives to continue. The heterogeneous-agent analysis of Sákovics and Steiner (2012) relates targeting to these external effects. Near completion, the acceleration benefit recedes. Once the payment date is secured, additional research reallocates ownership and can absorb resources without advancing commercialization.

A coordinated rule based only on the public physical state recovers 28.81% of the decentralized surplus gap. Let $a_P(S, \theta)$ be the causal planner's research choice. The public-state allocation sets

$$a_{\text{pub}}(S) = a_P(S, 0).$$

It uses capabilities and remaining work, without observing current private signals or the current innovation. The realized innovation enters the subsequent physical transition, and the rule is applied again at the new public state.

Table 5 reports outcomes under common future shocks. The public-state rule raises expected surplus from 3.795 to 4.093, compared with the planner's 4.828, and reduces mean commercialization time from 41.22 to 35.51 years. Its gains come from coordinating the timing and scale of research using publicly observable industry conditions. All reported surplus measures deduct the resources consumed by research. The public-state rule shows that gains do not require observing the current private innovation. It is an allocation benchmark, rather than an implemented subsidy or an optimal policy within that restricted information set. When m represents profits, the comparison concerns producer surplus. A welfare interpretation requires the surplus interpretation and household closure stated in Section 3.

Table 5: Completion, research, and surplus under common shocks

	Private	Causal planner	Public-state rule
Mean completion (years)	41.224	34.024	35.510
Discounted research cost	2.154	2.749	2.751
Undiscounted research cost	6.387	5.810	5.808
Discounted net surplus	3.795	4.828	4.093

Notes: Means over 20,000 common-shock histories. Research costs include only the costly sector's resource bill, not costless core learning. Surplus is discounted terminal payoff less discounted research costs. Monte Carlo standard errors are reported in the replication tables; the private completion mean has a standard error of 0.10 years. The planner has no advance knowledge of future innovations.

6 Conclusion

We have shown that the build-out of a general purpose technology can generate an abrupt collapse in fundamentally priced equity and a persistent freeze in investment. This occurs entirely without financial frictions or irrational pricing. The crash arises from the strategic timing of commercialization. Firms accumulate capabilities and advance a common payoff date. Their coordinated withdrawal in response to modest adverse news postpones that date, destroys accumulated knowledge, and triggers a depreciation spiral. This mechanism reconciles the large interim equity losses of the railroad and internet eras with the eventual technological success and enduring value of their industry leaders.

Our framework offers a direct lens for the ongoing commercialization of artificial intelligence (AI). Like the internet and the railroads, AI is a general purpose technology that requires substantial complementary investments in infrastructure, software, and organizational redesign before its productivity gains

are realized. The current AI boom is defined by an intense accumulation of these proprietary capabilities. Our model suggests that the primary fragility of this build-out need not lie in financial leverage or a resale bubble. It lies in the strategic coordination of the investment itself and the uncertain timeline, regardless of how productive AI will ultimately turn out to be.

The commercial viability of AI depends on synchronized efforts across foundation model developers, infrastructure providers, and downstream enterprise adopters. The industry remains vulnerable to collective withdrawal because aggregate investment determines usability and profitable adoption. We define this threshold as the point of commercialization. A contraction erodes specialized organizational knowledge and inefficiently delays the technology's payoff. This generates significant dynamic propagation of the initial coordinated withdrawal. Our paper addresses the ubiquitous question of whether a boom is a bubble. It demonstrates that a technology can be fundamentally priced, ultimately successful, and still suffer a crash and inefficient delay due to the fragile coordination of its developers.

References

Abreu, Dilip and Markus K. Brunnermeier, "Bubbles and Crashes," *Econometrica*, 2003, *71* (1), 173–204.

Angeletos, George-Marios and Iván Werning, "Crises and Prices: Information Aggregation, Multiplicity, and Volatility," *American Economic Review*, 2006, *96* (5), 1720–1736.

_ **, Christian Hellwig, and Alessandro Pavan**, "Signaling in a Global Game: Coordination and Policy Traps," *Journal of Political Economy*, 2006, *114* (3), 452–484.

_ **,** _ **, and** _ , "Dynamic Global Games of Regime Change: Learning, Multiplicity, and the Timing of Attacks," *Econometrica*, 2007, *75* (3), 711–756.

_ **, Guido Lorenzoni, and Alessandro Pavan**, "Wall Street and Silicon Valley: A Delicate Interaction," *Review of Economic Studies*, 2023, *90* (3), 1041–1083.

Beaudry, Paul and Franck Portier, "An Exploration into Pigou's Theory of Cycles," *Journal of Monetary Economics*, 2004, *51* (6), 1183–1216.

Bernanke, Ben S., Mark Gertler, and Simon Gilchrist, "The Financial Accelerator in a Quantitative Business Cycle Framework," in John B. Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*, Vol. 1C, Amsterdam: Elsevier, 1999, chapter 21, pp. 1341–1393.

Board of Governors of the Federal Reserve System, "Report to the Congress on the Availability of Credit to Small Businesses," Board of Governors of the Federal Reserve System, Figure 7 (PwC/NVCA MoneyTree data) 2007.

Bresnahan, Timothy F. and Manuel Trajtenberg, "General Purpose Technologies: 'Engines of Growth'?," *Journal of Econometrics*, 1995, *65* (1), 83–108.

Brunnermeier, Markus K. and Yuliy Sannikov, "A Macroeconomic Model with a Financial Sector," *American Economic Review*, 2014, *104* (2), 379–421.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," *American Economic Journal: Macroeconomics*, 2021, *13* (1), 333–372.

Caballero, Ricardo J., Emmanuel Farhi, and Mohamad L. Hammour, "Speculative Growth: Hints from the U.S. Economy," *American Economic Review*, 2006, *96* (4), 1159–1192.

Campbell, Gareth, "Deriving the Railway Mania," *Financial History Review*, 2013, *20* (1), 1–27.

Carlsson, Hans and Eric van Damme, "Global Games and Equilibrium Selection," *Econometrica*, 1993, *61* (5), 989–1018.

Comin, Diego and Bart Hobijn, "The CHAT Dataset," NBER Working Paper 15319, National Bureau of Economic Research 2009.

_ **and** _ , "An Exploration of Technology Diffusion," *American Economic Review*, 2010, *100* (5), 2031–2059.

Commercial and Financial Chronicle, "Course of Prices of Railroad and Miscellaneous Stocks during the Year," *Commercial and Financial Chronicle*, issues of January 4, 1873; January 10, 1874; January 9, 1875; January 8, 1876; January 6, 1877; January 5, 1878; and January 4, 1879. Accessed through FRASER, Federal Reserve Bank of St. Louis 1873–1879.

Cowles III, Alfred and Associates, *Common-Stock Indexes, 1871–1937* Cowles Commission Monograph No. 3, Bloomington, IN: Principia Press, 1938.

CRSP US Stock Database, "CRSP Stock (Monthly) and CRSP Stock (Daily)," Provided by Wharton Research Data Services (WRDS) 2026. Available at: https://wrds-www.wharton.upenn.edu/p

ages/about/data-vendors/center-for-research-in-security-prices-crsp/, last accessed on October 22, 2026.

Daggett, Stuart, *Railroad Reorganization*, Vol. 4 of *Harvard Economic Studies*, Boston and New York: Houghton, Mifflin and Company, 1908.

Diamond, Douglas W. and Philip H. Dybvig, "Bank Runs, Deposit Insurance, and Liquidity," *Journal of Political Economy*, 1983, *91* (3), 401–419.

Drozd, Lukasz A. and Ricardo Serrano-Padial, "Financial Contracting with Enforcement Externalities," *Journal of Economic Theory*, 2018, *178*, 153–189.

Fajgelbaum, Pablo D., Edouard Schaal, and Mathieu Taschereau-Dumouchel, "Uncertainty Traps," *Quarterly Journal of Economics*, 2017, *132* (4), 1641–1692.

Federal Open Market Committee, "FOMC Statement," Board of Governors of the Federal Reserve System, press release, June 30 June 1999.

_ , "FOMC Statement and Board Discount Rate Action," Board of Governors of the Federal Reserve System, press release, May 16 May 2000.

Federal Reserve Bank of Minneapolis, "Consumer Price Index, 1800–," Federal Reserve Bank of Minneapolis 2026.

Frankel, David and Ady Pauzner, "Resolving Indeterminacy in Dynamic Settings: The Role of Shocks," *The Quarterly Journal of Economics*, 2000, *115* (1), 285–304.

Frankel, David M., Stephen Morris, and Ady Pauzner, "Equilibrium Selection in Global Games with Strategic Complementarities," *Journal of Economic Theory*, 2003, *108* (1), 1–44.

French, Kenneth R., "12 Industry Portfolios," Data library, Tuck School of Business 2026. CRSP data through August 2026.

Gârleanu, Nicolae, Stavros Panageas, and Jianfeng Yu, "Technological Growth and Asset Pricing," *Journal of Finance*, 2012, *67* (4), 1265–1292.

Goetzmann, William N., Roger G. Ibbotson, and Liang Peng, "A New Historical Database for the NYSE 1815 to 1925: Performance and Predictability," *Journal of Financial Markets*, 2001, *4* (1), 1–32.

Goldfarb, Brent, David Kirsch, and David A. Miller, "Was There Too Little Entry During the Dot Com Era?," *Journal of Financial Economics*, 2007, *86* (1), 100–144.

Goldstein, Itay and Ady Pauzner, "Demand–Deposit Contracts and the Probability of Bank Runs," *Journal of Finance*, 2005, *60* (3), 1293–1327.

Healy, Paul M. and Krishna G. Palepu, "The Fall of Enron," *Journal of Economic Perspectives*, 2003, *17* (2), 3–26.

Hellwig, Christian, Arijit Mukherji, and Aleh Tsyvinski, "Self-Fulfilling Currency Crises: The Role of Interest Rates," *American Economic Review*, 2006, *96* (5), 1769–1787.

Helpman, Elhanan and Manuel Trajtenberg, "A Time to Sow and a Time to Reap: Growth Based on General Purpose Technologies," in Elhanan Helpman, ed., *General Purpose Technologies and Economic Growth*, Cambridge, MA: MIT Press, 1998, pp. 55–83.

Hirano, Tomohiro, Keiichi Kishi, and Alexis Akira Toda, "General-Purpose Technologies and Stock Market Bubbles," 2026. arXiv:2501.08215v3, revised August 16, 2026. Original January 2025 version titled Bursting Bubbles in a Macroeconomic Model.

Jaimovich, Nir and Sergio Rebelo, "Can News about the Future Drive the Business Cycle?," *American Economic Review*, 2009, *99* (4), 1097–1118.

Johnston, Louis and Samuel H. Williamson, "What Was the U.S. GDP Then?," MeasuringWorth 2026. Accessed October 20, 2026.

Jovanovic, Boyan and Peter L. Rousseau, "General Purpose Technologies," in Philippe Aghion and Steven N. Durlauf, eds., *Handbook of Economic Growth*, Vol. 1B, Elsevier, 2005, chapter 18, pp. 1181–1224.

Kiyotaki, Nobuhiro and John Moore, "Credit Cycles," *Journal of Political Economy*, 1997, *105* (2), 211–248.

Klette, Tor Jakob and Samuel Kortum, "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, 2004, *112* (5), 986–1018.

Li, Wendy C. Y. and Bronwyn H. Hall, "Depreciation of Business R&D Capital," *Review of Income and Wealth*, 2020, *66* (1), 161–180.

Lorenzoni, Guido, "A Theory of Demand Shocks," *American Economic Review*, 2009, *99* (5), 2050–2084.

Mixon, Scott, "The Crisis of 1873: Perspectives from Multiple Asset Classes," *Journal of Economic History*, 2008, *68* (3), 722–757.

Morris, Stephen and Hyun Song Shin, "Global Games: Theory and Applications," in Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, eds., *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, Vol. 1, Cambridge University Press, 2003, chapter 3, pp. 56–114.

Nasdaq, Inc., "NASDAQ Composite [NASDAQCOM]," FRED, Federal Reserve Bank of St. Louis n.d. Daily closing index; accessed October 20, 2026.

National Bureau of Economic Research, "American Railroad Stock Prices for United States," NBER Macrohistory, series m11005, from F. R. Macaulay, pp. A142–A161; accessed through FRED 1938.

_ , "Index of Wholesale Prices, Variable Group Weights, 1850–1894, Series m04048a," NBER Macrohistory Database 2026. Source: Warren and Pearson (1933), pp. 12–13. Accessed October 31, 2026.

_ , "Miles of Railroad Built, Series a02082a," NBER Macrohistory Database 2026.

_ , "US Business Cycle Expansions and Contractions," 2026. Business cycle chronology, accessed October 20, 2026.

National Center for Science and Engineering Statistics, "U.S. R&D Totaled $892 Billion in 2022; Estimate for 2023 Indicates Further Increase to $940 Billion," Technical Report NSF 25-327, National Science Foundation 2025. Figure 4: business-funded R&D in constant 2017 dollars.

Nitschke, Christoph, "Theory and History of Financial Crises: Explaining the Panic of 1873," *Journal of the Gilded Age and Progressive Era*, 2018, *17* (2), 221–240.

Odlyzko, Andrew, "Collective Hallucinations and Inefficient Markets: The British Railway Mania of the 1840s," Working Paper, University of Minnesota 2010. SSRN 1537338, preliminary version of January 15.

_ , "The Collapse of the Railway Mania, the Development of Capital Markets, and the Forgotten Role of Robert Lucas Nash," *Accounting History Review*, 2011, *21* (3), 309–345.

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

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

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

Perez, Carlota, *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*, Edward Elgar Publishing, 2002.

Poor, Henry V., *Manual of the Railroads of the United States for 1880*, New York: H. V. and H. W. Poor, 1880. Scan: Internet Archive, poorsmanualofr1880newyuoft.

Quinn, William and John D. Turner, *Boom and Bust: A Global History of Financial Bubbles*, Cambridge: Cambridge University Press, 2020.

Richardson, Gary and Tim Sablik, "Banking Panics of the Gilded Age," Federal Reserve History December 2015. Written December 4, 2015; accessed October 20, 2026.

Ritter, Jay R., "IPO Statistics," University of Florida, Warrington College of Business 2026.

Sákovics, József and Jakub Steiner, "Who Matters in Coordination Problems?," *American Economic Review*, 2012, *102* (7), 3439–3461.

Scheinkman, José A. and Wei Xiong, "Overconfidence and Speculative Bubbles," *Journal of Political Economy*, 2003, *111* (6), 1183–1220.

Seltzer, Lawrence H., *A Financial History of the American Automobile Industry: A Study of the Ways in Which the Leading American Producers of Automobiles Have Met Their Capital Requirements*, Boston: Houghton Mifflin, 1928.

Swain, Henry H., "Economic Aspects of Railroad Receiverships," *Economic Studies*, 1898, *3* (2), 53–161. As tabulated in B. Hansen, "Bankruptcy Law in the United States," EH.Net Encyclopedia, Table 3.

U.S. Bureau of Economic Analysis, "Private Fixed Investment in Computers and Peripheral Equipment [B935RC1Q027SBEA] and in Communications Equipment [C275RC1A027NBEA]," FRED, Federal Reserve Bank of St. Louis 2026.

U.S. Bureau of Labor Statistics, "Consumer Price Index for All Urban Consumers: All Items in U.S. City Average [CPIAUCSL]," FRED, Federal Reserve Bank of St. Louis 2026.

US Census Bureau, "E-Commerce Retail Sales as a Percent of Total Sales (ECOMPCTSA)," Quarterly Retail E-Commerce Report, via FRED, Federal Reserve Bank of St. Louis 2026. https://fred.stlouisfed.org/series/ECOMPCTSA.

U.S. Securities and Exchange Commission, "SEC Charges WorldCom with $3.8 Billion Fraud," Litigation Release No. 17588 June 2002.

_ , "Lucent Settles SEC Enforcement Action Charging the Company with $1.1 Billion Accounting Fraud," Press Release 2004-67 May 2004.

_ , "SEC Files Securities Fraud Charges Against Computer Associates International, Inc., Former CEO Sanjay Kumar, and Two Other Former Company Executives," Press Release 2004-134 September 2004.

Warren, George F. and Frank A. Pearson, *Prices*, New York: John Wiley & Sons, 1933.

White, Richard, *Railroaded: The Transcontinentals and the Making of Modern America*, New York: W. W. Norton, 2011. ISBN 978-0-393-06126-0.

Willoughby, Jack, "Burning Up," Barron's, March 20 March 2000. Contemporary article reproduced on Aswath Damodaran's NYU website; original publication date confirmed by the author's follow-up.

Online Appendix

Commercialization of GPTs and Stock Market Crashes

Lukasz A. Drozd

Preliminary draft — limited circulation

Contents

A.1 Proof of the simplified model **A-2**

A.2 Recursive equilibrium: statements, proofs, and scope **A-10**

A.3 Fundamental prices and the acceleration–dilution boundary **A-18**

A.4 Historical data and measurement **A-21**

A.5 Computation and numerical scope **A-41**

A.1 Proof of the simplified model

Proof of Proposition 1. The commercialization timeline and private return. Requirements are at most $3/4 + 1/4 + 1 = 2$. Since $N(a) \geq 1$, two periods always suffice. The first-period completion probability is $N(a) - 3/4 - \theta$, which belongs to $[0, 1]$ throughout the feasible domain. Expected discounted aggregate receipts are therefore

$$16\left\{\frac{1}{16} + \left(\frac{1}{4} - \frac{1}{16}\right)[N(a) - 3/4 - \theta]\right\} = 3N(a) - 5/4 - 3\theta.$$

Dividing by $N(a)$ gives (25). In particular,

$$Q_\theta = -\frac{3}{N(a)} < 0, \qquad Q_a = \frac{5/4 + 3\theta}{2N(a)^2} > 0.$$

The second inequality is strict even at the lowest fundamental in the prior support. The investment incentive is consequently increasing in others' research at every information set.

Dominance bounds and best responses. With everybody else pausing, $Q(\theta, 0) = 7/4 - 3\theta$. With everybody else researching, $Q(\theta, 1) = 13/6 - 2\theta$. These give the complete-information dominance bounds $-1/12$ and $1/12$. For the private-information game, the posterior is uniform on the intersection of $[s - \nu, s + \nu]$ and $[-1/4, 1/4]$. When $s \in [-1/12, 1/12]$ this intersection is the full interval $[s - \nu, s + \nu]$, because $\nu < 1/6$. Its mean is s. The best response to universal pausing is therefore the cutoff $s < -1/12$, and the best response to universal research is $s < 1/12$.

If other firms use a cutoff, their aggregate participation is nonincreasing in θ. The realized return $Q(\theta, A(\theta))$ is then strictly decreasing in θ. Posterior distributions shift upward with s, including where

their support is truncated by the prior. Expected return is thus strictly decreasing in s, and the unique best response is a cutoff. Its boundary lies in $[-1/12, 1/12]$ by the dominance bounds just established.

Selection and uniqueness without imposing cutoff strategies. Start with universal pausing as a lower bound on every player's strategy and universal research as an upper bound. Update each bound by best response. Strategic complementarity makes the lower sequence weakly increase and the upper sequence weakly decrease. After their first update both bounds are cutoff strategies. Their cutoffs remain in the compact interval $[-1/12, 1/12]$ and hence converge. The best-response cutoff is continuous in the opponents' cutoff: aggregate participation is continuous, returns are bounded and continuous, and conditional return strictly decreases through its indifference point. Dominated convergence therefore makes both limiting bounds cutoff equilibria.

At any such fixed point, the marginal posterior is the untruncated uniform distribution on $[c - v, c + v]$. Substitution $\theta = c + v - 2va$ gives

$$\begin{aligned} G(c) &= \int_0^1 \left[1 - \frac{5/4 + 3c + 3v - 6va}{1 + a/2}\right] da \\ &= 1 - (5/4 + 3c)L + v(12 - 15L), \qquad L = 2\log(3/2), \end{aligned}$$

where $\int_0^1 (1 + a/2)^{-1} da = L$ and $\int_0^1 a(1 + a/2)^{-1} da = 2 - 2L$. This expression has slope $-3L < 0$, so its only zero is

$$c(v) = \frac{1/L - 5/4}{3} + v\left(\frac{4}{L} - 5\right), \qquad L = 2\log(3/2). \tag{A.1}$$

As noise vanishes, $c(v) \to (1/L - 5/4)/3 \simeq -0.00562$. The limiting lower and upper best-response cutoffs coincide. Every Bayesian equilibrium, including an asymmetric or mixed strategy profile, lies between every pair of iterated bounds: if its opponents' strategies lie between two bounds, monotonicity puts its best responses between the updated bounds. Consequently every equilibrium uses the same strategy except at the indifferent signal, a null set. This proves the stated unrestricted uniqueness. There is no independently selected confidence state.

Planner, delay, and survivor shares. The resource bill is $2Fa = a$. Subtracting it from expected discounted aggregate receipts gives $W(\theta, a) = 7/4 - 3\theta + a/2$, so $a = 1$ is the unique planner policy at every in-

novation. This remains its optimum without observing θ because the derivative is always $1/2$. Since completion is at date 1 or 2, expected completion time is $2 - \Pr\{\tau = 1\} = 7/4 + \theta - a/2$. The innovations $\theta_G = c(\nu) - \nu$ and $\theta_B = c(\nu) + \nu$ give participation one and zero, respectively. Substitution establishes the private increase $1/2 + 2\nu$, the planner increase 2ν, and the same-fundamental loss $1/2$ in both expected delay and net surplus. These comparisons use the same distribution of unresolved certification requirements. With the same realization of that requirement, additional research can only advance completion in this example.

Finally, the core owns $1/2$ capability after its initial program, so its terminal share is $1/3$ at $a = 1$ and $1/2$ at $a = 0$. At the limiting cutoff, $5/4 + 3c = 1/L$. Its values are consequently $(1/2)[3 - 2/(3L)]$ and $(1/2)[3 - 1/L]$, respectively. They are the values reported in the text. A larger terminal share need not compensate for a slower common timeline. □

A.1.1 Three research costs in the common-timeline example

Proof of Proposition 2. Index the costs in increasing order, $z_1 = 1.9$, $z_2 = 2$, and $z_3 = 2.1$, with capability weights $w_i = 1/3$. Write c_i for type i's signal cutoff and

$$A_c(\theta) = \sum_{i=1}^{3} w_i \left[\frac{c_i - \theta + \nu}{2\nu} \right]_0^1.$$

The return remains $Q(\theta, a) = 3 - (5/4 + 3\theta)/(1 + a/2)$.

Existence and ordered cutoffs. For sufficiently small noise, a monotone cutoff equilibrium exists. To see this directly, set

$$\ell_i = \frac{7/4 - z_i}{3}, \qquad u_i = \frac{13/6 - z_i}{2}, \qquad D = \prod_{i=1}^{3} [\ell_i, u_i].$$

For $0 < \nu < 7/60$, every candidate cutoff in D has its entire posterior interval inside $[-1/4, 1/4]$. Against any $c \in D$, the conditional expected return is continuous and strictly decreasing in the firm's signal. Indeed, $Q_\theta < 0$, $Q_a > 0$, and A_c is nonincreasing, so $Q(\theta, A_c(\theta))$ strictly decreases in θ. The conditional uniform distribution of θ shifts upward in first-order stochastic dominance as the signal rises, including where the prior truncates its support.

At $s = \ell_i$, expected return is at least $\mathbb{E}[Q(\theta, 0) \mid s = \ell_i] = z_i$; at $s = u_i$, it is at most $\mathbb{E}[Q(\theta, 1) \mid s = u_i] = z_i$. Each type therefore has a unique best-response cutoff in $[\ell_i, u_i]$. The resulting cutoff map is continuous and maps the compact convex set D into itself. Brouwer's theorem supplies a fixed point. Strict monotonicity of conditional returns makes these cutoff choices optimal at every signal. Since all types face the same conditional return function but $z_1 < z_2 < z_3$, every such equilibrium satisfies $c_1 > c_2 > c_3$. This argument establishes existence in monotone cutoffs.

Uniqueness at finite noise. For $0 < \nu < 7/60$, every equilibrium, including one in non-cutoff or mixed strategies, coincides with a monotone cutoff equilibrium except at indifferent signals, and that cutoff equilibrium is unique. The argument follows the translation argument of Drozd and Serrano-Padial (2018). Write

$$G(s; c) = \mathbb{E}[Q(\theta, A_c(\theta)) \mid s]$$

for the expected capability return at signal s when other firms use the cutoff profile c. It is common to all types, which differ only in cost.

First, the iterated best-response bounds of Appendix A.1 apply type by type. Starting from universal pausing and universal research, the first updates are the cutoff profiles (ℓ_i) and (u_i), because every signal in $[\ell_i, u_i]$ has an untruncated posterior with mean s. Strategic complementarity makes the lower sequence of cutoff profiles increase and the upper sequence decrease within D, and continuity of the best-response cutoffs makes their limits $\underline{c} \leq \bar{c}$ cutoff equilibria. Every equilibrium lies between them. It therefore suffices to show that $\underline{c} = \bar{c}$.

Two properties of G deliver this. Since $Q_a > 0$ and a higher cutoff raises participation at every fundamental, $G(s; c)$ is nondecreasing in each component of c. Translating every cutoff by Δ translates participation, $A_{c+\Delta}(\theta + \Delta) = A_c(\theta)$. For a signal whose posterior is the untruncated uniform distribution on $[s - \nu, s + \nu]$, the change of variables $\theta = \theta' + \Delta$ therefore gives

$$G(s + \Delta; c + \Delta) = G(s; c) - 3\Delta\, \mathbb{E}\left[\frac{1}{1 + A_c(\theta)/2} \,\middle|\, s\right] < G(s; c) \qquad \text{for } \Delta > 0,$$

because $Q_\theta = -3/(1 + a/2) < 0$. This is the counterpart of the common-translation step in the working-paper version of Drozd and Serrano-Padial (2018) (WP 18-21, Appendix B.1.1, Lemma 6); the sign reverses because a higher requirement is adverse whereas greater enforcement capacity is favorable. Translation raises the requirement while leaving the participation profile faced by the marginal firm unchanged.

Suppose that $\underline{c} \neq \bar{c}$, and let $\Delta = \max_i(\bar{c}_i - \underline{c}_i) > 0$, attained by type j. Then $\underline{c} + \Delta \geq \bar{c}$ componentwise and $\bar{c}_j = \underline{c}_j + \Delta$. Both cutoffs of type j lie in $[\ell_j, u_j]$, so their posteriors are untruncated, and both satisfy type j's indifference condition. Hence

$$z_j = G(\bar{c}_j; \bar{c}) \leq G(\underline{c}_j + \Delta; \underline{c} + \Delta) < G(\underline{c}_j; \underline{c}) = z_j,$$

a contradiction. The smallest and largest equilibria coincide, which proves uniqueness. The limit characterization below applies to this unique equilibrium.

The full belief constraint. At type i's marginal signal, the posterior density of the fundamental is

$$f_i(\theta) = \frac{1}{2\nu} \mathbf{1}\{|\theta - c_i| < \nu\}.$$

Except at finitely many kinks,

$$-A_c'(\theta) = \sum_{i=1}^{3} w_i f_i(\theta).$$

Consequently, for every bounded measurable function f,

$$\sum_{i=1}^{3} w_i \mathbb{E}[f(A_c(\theta)) \mid s_i = c_i] = \int f(A_c(\theta))[-A_c'(\theta)]\, d\theta = \int_0^1 f(a)\, da. \tag{A.2}$$

The last equality uses the fact that A_c is continuous and nonincreasing from one to zero. Flat portions carry no mass under $-A_c'$. Thus the capability-weighted mixture of marginal participation beliefs is uniform. Individual types' marginal beliefs need not be uniform. In particular, taking $f(a) = a$ gives a weighted mean of one-half; taking $f(a) = a^2$ gives one-third. The full identity, rather than just its first moment, is needed to average the nonlinear return.

Why separate limiting thresholds are impossible. Consider any sequence of monotone cutoff equilibria with v tending to zero. The cutoffs lie in D, so a subsequence converges. Its distinct limiting cutoffs partition the ordered cost types into contiguous blocks. Let B be such a block, w_B its capability mass, $\bar{z}_B$ its mean cost, and a_0 the mass of cheaper types outside B. In a shrinking neighborhood of the block's limiting threshold θ_B, cheaper outside types always research and more expensive outside types always pause. Applying the preceding density identity to the block alone, and then taking the limit in its marginal indifference conditions, gives

$$\bar{z}_B = \frac{1}{w_B} \int_{a_0}^{a_0+w_B} Q(\theta_B, a)\, da. \tag{A.3}$$

All posterior supports within the block converge to θ_B; continuity and boundedness of Q justify this step even if their cutoff differences vanish more slowly than v.

Define

$$L_B = \frac{2}{w_B} \log\left(\frac{1 + (a_0 + w_B)/2}{1 + a_0/2} \right).$$

Equation (A.3) determines the candidate boundary

$$\theta_B = \frac{(3 - \bar{z}_B)/L_B - 5/4}{3}. \tag{A.4}$$

There are only three proper contiguous partitions. For singleton blocks, the candidate boundaries are

$$\begin{aligned} t_1 &= \frac{11}{180 \log(7/6)} - \frac{5}{12} \simeq -0.020229, \\ t_2 &= \frac{1}{18 \log(8/7)} - \frac{5}{12} \simeq -0.000618, \\ t_3 &= \frac{1}{20 \log(9/8)} - \frac{5}{12} \simeq 0.007843. \end{aligned}$$

They satisfy $t_1 < t_2 < t_3$, contrary to the required ordering of cutoffs. For a pooled pair and a singleton, the additional candidates are

$$t_{12} = \frac{7}{60 \log(4/3)} - \frac{5}{12} \simeq -0.011126, \qquad t_{23} = \frac{19}{180 \log(9/7)} - \frac{5}{12} \simeq 0.003347.$$

The partition $\{1,2\} \mid \{3\}$ would require $t_{12} > t_3$, while $\{1\} \mid \{2,3\}$ would require $t_1 > t_{23}$; both inequalities are reversed. Hence every limiting equilibrium pools all three types.

For that full block, $w_B = 1$, $a_0 = 0$, and $\bar{z}_B = 2$, so

$$\theta^* = \frac{1}{6\log(3/2)} - \frac{5}{12} \simeq -0.00561609. \tag{A.5}$$

Every convergent subsequence has this same limit. Compactness therefore implies $c_i(\nu) \to \theta^*$ for every type, along any selection of monotone cutoff equilibria. Aggregate participation converges to one for $\theta < \theta^*$ and to zero for $\theta > \theta^*$.

Timing and efficiency. Let a_i be the fraction of type i researching and $a = \sum_i w_i a_i$. The planner's value is

$$W(\theta, a_1, a_2, a_3) = \frac{7}{4} - 3\theta + \sum_{i=1}^{3} w_i \left(\frac{3}{2} - \frac{z_i}{2}\right) a_i.$$

Every coefficient of a_i is positive, so the planner selects all programs at every fundamental. Full withdrawal lowers surplus by $1/2$ and adds $1/2$ period to expected implementation, relative to full research at the same fundamental. Implementation by date 2 remains certain under every policy. The full-participation and full-withdrawal fundamentals $\min_i c_i - \nu$ and $\max_i c_i + \nu$ converge to the same θ^*; thus arbitrarily small news can induce these positive private losses. □

A.1.2 An open range of cost dispersion and net equity losses

The three-cost result is not confined to the numerical costs in the text. Keep the same work technology, prior, and payoff, and let $z_1 = 2 - d$, $z_2 = 2$, and $z_3 = 2 + d$, with equal capability weights. Define

$$L_1 = 6\log(7/6), \quad L_2 = 6\log(8/7), \quad L_3 = 6\log(9/8),$$
$$L_{12} = 3\log(4/3), \qquad L_{23} = 3\log(9/7),$$

and

$$d_* = \frac{L_{12} - L_3}{L_{12} + L_3/2} \simeq 0.128534.$$

Proposition 4 (Robust clustering and net equity losses). *For every $0 < d < d_*$ and $0 < v < 1/6 - d/2$, the structural three-cost game has an essentially unique Bayesian equilibrium. All three cutoffs converge to the boundary in Proposition 2 as noise vanishes. Set $\theta_G(v) = \min_i c_i(v) - v$ and $\theta_B(v) = \max_i c_i(v) + v$. These fundamentals induce full research and full withdrawal and converge to the same boundary. Between them, the limiting net equity loss is strictly positive for every costly type and for the core. Aggregate net owner value falls by $1/2$, and expected implementation is postponed by $1/2$ period. The core's eventual share nevertheless rises from $1/3$ to $1/2$.*

Proof. The dominance rectangle has bounds $\ell_i = (7/4 - z_i)/3$ and $u_i = (13/6 - z_i)/2$. The noise restriction keeps every marginal posterior untruncated. The existence, extremal-best-response, and translation arguments in Appendix A.1.1 therefore apply without change. They give uniqueness, ordered cutoffs, and the block indifference equation.

For separate singleton blocks the proposed boundaries, in increasing cost order, are

$$t_1 = \frac{(1+d)/L_1 - 5/4}{3}, \quad t_2 = \frac{1/L_2 - 5/4}{3}, \quad t_3 = \frac{(1-d)/L_3 - 5/4}{3}.$$

Their required order fails since $t_1 < t_2$ for $d < L_1/L_2 - 1 \simeq 0.154415$. For the partition $\{1\} \mid \{2,3\}$, the proposed cheap-type boundary is below the pooled boundary when

$$d < \frac{L_1 - L_{23}}{L_{23} + L_1/2} \simeq 0.140547.$$

For $\{1,2\} \mid \{3\}$ the same ordering contradiction is equivalent to $d < d_*$. Both inequalities, as well as the singleton contradiction, hold on the stated range. All proper partitions are excluded. The mean cost remains two, so the pooled boundary is unchanged. Compactness and the block belief constraint establish convergence as in the preceding proof.

Let $L = 2\log(3/2)$. At the common boundary, returns per post-program capability under full research and withdrawal are

$$Q_1 = 3 - \frac{2}{3L}, \qquad Q_0 = 3 - \frac{1}{L}.$$

The core's value falls by $(Q_1 - Q_0)/2 > 0$. A costly owner with one inherited capability receives net value $2Q_1 - z_i$ under research and Q_0 under withdrawal. The difference is

$$2Q_1 - z_i - Q_0 = 3 - z_i - \frac{1}{3L} \geq 1 - d_* - \frac{1}{3L} > 0.$$

Thus the result concerns values net of the current funding call, not only gross prices of additional claims. Aggregating with $C' = 1/2$, $F = 1/2$, and mean cost two gives an exact net owner-value difference of $1/2$. The expected-delay and ownership comparisons depend only on aggregate participation and follow from the work technology. □

These are strict inequalities on an open parameter range. They establish robustness to cost dispersion within this structural example; they are not a claim of robustness to every technology, information structure, or recursive continuation. In particular, the quantitative recursion still uses a common current research cost.

A.2 Recursive equilibrium: statements, proofs, and scope

This appendix collects the sufficient results used to interpret the recursive economy. The state S and continuation functions are held fixed in each current-game result unless explicitly stated otherwise. The linear-value calculation applies to the general cost distribution H. The subsequent dominance, scalar-cutoff, and dynamic uniqueness results specialize to a point mass at a common cost $z > 0$. The unrestricted uniqueness established in the simplified model is not assumed for the recurring game.

A.2.1 Linear values and current program decisions

Fix an aggregate policy and its induced public-state process. Under any feasible individual research history, future capability stocks, research expenditures, and terminal receipts are proportional to the initial n. Dividing by n leaves an optimization problem independent of firm size. Fresh cost draws give every costly capability the same distribution of future opportunities at a given public state. Its value before current costs and signals arrive is therefore $nv(S)$, proving (7). At current cost z and signal s, pausing gives $n\mathbb{E}[Q \mid s]$, while researching gives $(1 + b)n\mathbb{E}[Q \mid s] - bzn$. Their difference is $bn(\mathbb{E}[Q \mid s] - z)$,

proving (10). Conditional on the true innovation, averaging over current costs and private signals gives expansion $bA(\theta)$ and cost $b\int za_z(\theta)\,dH(z)$ per inherited capability. Averaging over the innovation gives (11). The core's costless program gives (12). Linearity permits positive rents because inherited capabilities constrain each firm's expansion.

A.2.2 Dominance, including the option to research later

Lemma 1 (Bounds on capability values)**.** *Suppose $0 < \beta < 1$, values equal fundamental present values, research costs are nonnegative, and $C > 0$. Then*

$$0 \le Q(\theta, a), U(\theta, a) \le \frac{\beta m}{(1+b)C}. \tag{A.6}$$

If also $\beta(1+b) < 1$, then

$$v(S), u(S) \le \frac{\beta(1+b)m}{C+F}, \qquad Q(\theta, a), U(\theta, a) \le \frac{\beta m}{N(a)}. \tag{A.7}$$

Under $\beta(1+b) < 1$, research is strictly dominated at every signal when

$$\frac{\beta m}{(1+b)C+F} < z. \tag{A.8}$$

If $\beta m/[(1+b)(C+F)] > z$, sufficiently favorable signals make research strictly optimal uniformly over other firms' current and future policies. Sufficiently adverse signals make research strictly suboptimal uniformly over those policies.

Proof. Start with one capability after the current program. Its lineage can grow no faster than the post-program core stock $(1+b)C$, because both have the same maximum research expansion and depreciation. Its terminal ownership share is at most $1/[(1+b)C]$ on every path. The earliest receipt is one period away. Discounting and dropping nonnegative costs gives the upper bound in (A.6). Pausing forever is feasible and has nonnegative receipts, giving the lower bound for a costly capability. Core programs are costless, so the same lower bound holds for a core capability.

If completion occurs $j \geq 1$ periods from a pre-program state, legacy aggregate capabilities alone leave at least $(C+F)\exp(-j\delta)$ survivors. One initial capability has at most $(1+b)^j \exp(-j\delta)$ descendants. Its discounted gross receipt is thus bounded by $m[\beta(1+b)]^j/(C+F)$, which is at most $\beta(1+b)m/(C+F)$ when $\beta(1+b) < 1$. A capability after the current program has only $j-1$ remaining expansion opportunities before a receipt j periods away. Its corresponding bound is $\beta m[\beta(1+b)]^{j-1}/N(a) \leq \beta m/N(a)$. This proves (A.7) and the all-signal dominance condition because $N(a) \geq (1+b)C + F$.

For the favorable tail, every participation rate completes currently when $\theta \leq h[(1+b)C+F] - R$. On that event $Q \geq \beta m/[(1+b)(C+F)]$. Its posterior probability tends to one as $s \to -\infty$. Nonnegativity of returns proves favorable-signal dominance under the stated strict inequality.

For the adverse tail, set $g = \exp(-\delta)(1+b)$. Under maximal current and future research, completed work over j periods is at most

$$W_j^{\max} = h(1+b)(C+F)\sum_{\ell=0}^{j-1} g^\ell.$$

Conditional on current signal s, the sum of the current and next $j-1$ innovations is normal with mean ws and variance $wv^2 + (j-1)\sigma^2$. A union bound therefore gives, for every policy,

$$\Pr(\tau \leq T \mid s) \leq B_T(s) := \min\left\{1, \sum_{j=1}^{T} \Phi\left(\frac{W_j^{\max} - R - ws}{\sqrt{wv^2 + (j-1)\sigma^2}}\right)\right\}.$$

The core ownership bound applies to every future program history, so

$$\mathbb{E}[Q \mid s] \leq \frac{m}{(1+b)C}\left[\beta B_T(s) + \beta^{T+1}\right]. \tag{A.9}$$

Choose T large enough that the discounted-tail term, including its prefactor, is below $z/2$. For this fixed T, $B_T(s) \to 0$ as $s \to \infty$, making the remaining term less than $z/2$ as well. This bound allows every feasible future resumption or expansion program. □

Eventual completion and a numerical tail bound. If $g > 1$, core-only cumulative work grows geometrically. Cumulative mean-zero Gaussian innovations divided by elapsed periods converge to zero almost surely. Hence the growing core completes every finite initial requirement almost surely under

any costly research policy. For a finite horizon T, retaining passive costly capabilities gives the bound

$$W_T^{\min} = h\left[(1+b)C\sum_{j=0}^{T-1} g^j + F\sum_{j=0}^{T-1}\exp(-j\delta)\right], \tag{A.10}$$

$$\Pr(\tau > T \mid S) \le \Phi\left(\frac{R - W_T^{\min}}{\sigma\sqrt{T}}\right). \tag{A.11}$$

These are policy-uniform completion and fixed-policy payoff bounds. They do not by themselves prove that equilibrium thresholds are insensitive to a computational boundary; that comparison is separate.

A.2.3 Verifying a cutoff at every signal

Lemma 2 (All-signal verification and cutoff uniqueness)**.** *Fix a state and continuation functions. If $Q(\theta, A(\theta)) - z$ crosses zero once, from positive to negative as θ rises, any solution of (16) is an optimal cutoff at every signal. For a piecewise continuously differentiable return, sufficient conditions are the favorable and adverse signs, downward jumps, and*

$$Q_\theta - \frac{Q_a}{v}\phi((c-\theta)/v) < 0$$

at each continuous cost-level crossing.

For uniqueness within the cutoff class, couple marginal posterior draws as $\theta = wc + v\sqrt{w}Z$, with Z standard normal. Suppose every smooth piece satisfies

$$Q_\theta + \frac{v}{\sigma^2}\phi((c-\theta)/v)Q_a \le 0, \tag{A.12}$$

and every crossing between pieces is downward along this coupling. If the integrated inequality is strict on every nonempty cutoff interval, G is strictly decreasing and has at most one zero. On continuation pieces, the stronger condition

$$-Q_\theta > \frac{v}{\sigma^2\sqrt{2\pi}}\max\{Q_a, 0\} \tag{A.13}$$

is sufficient. In the completion region, the smooth inequality is strict because $Q_\theta = 0$ and $Q_a < 0$.

Proof. For $s_2 > s_1$, Gaussian posterior densities satisfy a strict monotone likelihood ratio: $p(\theta \mid s_2)/p(\theta \mid s_1)$ increases in θ. Let θ_0 be the sign crossing. If expected net return is zero at s_1, expected net return at s_2 is

$$\int_{-\infty}^{\infty} \left[\frac{p(\theta \mid s_2)}{p(\theta \mid s_1)} - \frac{p(\theta_0 \mid s_2)}{p(\theta_0 \mid s_1)} \right] [Q(\theta, A(\theta)) - z] p(\theta \mid s_1)\, d\theta < 0.$$

Both sides of the crossing contribute nonpositively, and the strict signs have positive posterior probability. Reversing the argument gives positive expected net return at every lower signal. This proves all-signal optimality. The displayed crossing derivative follows from $A'(\theta) = -\phi((c-\theta)/v)/v$. The stated regularity and downward jumps exclude a second sign crossing.

For uniqueness, keep Z fixed across candidate marginal signals. Along the coupling,

$$\frac{d\theta}{dc} = w, \qquad \frac{dA(\theta)}{dc} = \frac{1-w}{v}\phi((c-\theta)/v).$$

The realized return's derivative on a smooth piece is $w[Q_\theta + (v/\sigma^2)\phi((c-\theta)/v)Q_a]$ because $(1-w)/(wv) = v/\sigma^2$. Downward crossings preserve this order through nonsmooth pieces. Taking expectations and using the strictness condition makes G strictly decreasing. □

Completion boundaries and endogenous continuation. On continuation pieces,

$$Q_\theta = \beta \exp(-\delta) v_R(S'). \tag{A.14}$$

The coupling crosses from completion to continuation as c increases if

$$\frac{hbFv}{\sigma^2\sqrt{2\pi}} < 1, \tag{A.15}$$

because the derivative of $R + \theta - hN(A(\theta))$ is $w[1 - (hbFv/\sigma^2)\phi((c-\theta)/v)]$. When $\beta(1+b) < 1$, the continuation side satisfies

$$\beta \exp(-\delta) v(S') \le \beta(1+b)\frac{\beta m}{N(a)} < \frac{\beta m}{N(a)},$$

so that crossing is downward. The same conclusion holds for U. The conditions are sufficient certificates for the continuation functions being considered. They do not turn the full recursive equilibrium operator into a contraction. An indifference residual, an all-signal best-response check, and a cutoff-uniqueness certificate are three different requirements.

A.2.4 A finite strategic horizon and conditional dynamic uniqueness

Proposition 5 (Finite strategic decisions inside an infinite-lived project)**.** *Suppose* $g = \exp(-\delta)(1+b) > 1$, $C_0 > 0$, *and* $z > 0$. *Define*

$$C_D = \frac{\beta m}{(1+b)z}, \qquad t_D = \min\{t \geq 0 \,:\, C_0 g^t > C_D\}. \tag{A.16}$$

Then t_D *is finite, and costly research is strictly dominated at every uncompleted state on this core-stock orbit from* t_D *onward. The no-costly-research tail has unique fundamental values.*

Fix an admissible class of current Bayesian strategies. At each preceding state reachable under any feasible policy, suppose the stage game induced by its already determined future values has a unique equilibrium in that class. Then the recurring private game has a unique recursive equilibrium in that class on the invariant core-stock orbit.

Proof. Once $C_0 g^t > C_D$, the core bound gives $Q < z$ for every innovation, current participation, remaining requirement, and future policy. The same strict inequality holds thereafter. Every equilibrium must therefore have zero costly research in the tail; no assumption about expecting others to pause is needed.

Starting at a tail state, completion is the stopping time

$$\tau = \inf\left\{ j \geq 1 \,:\, R + \sum_{\ell=0}^{j-1} \theta_\ell \leq h\left[(1+b)C \sum_{\ell=0}^{j-1} g^\ell + F \sum_{\ell=0}^{j-1} \exp(-\ell\delta)\right]\right\}. \tag{A.17}$$

It is finite almost surely. Fundamental tail values are uniquely

$$v(S) = m\mathbb{E}\left[\frac{\beta^{\tau}}{C(1+b)^{\tau}+F}\right], \tag{A.18}$$

$$u(S) = m\mathbb{E}\left[\frac{\beta^{\tau}(1+b)^{\tau}}{C(1+b)^{\tau}+F}\right]. \tag{A.19}$$

Depreciation cancels from ownership shares. These expectations value a stochastic completion time; they do not force completion at t_D.

At $t_D - 1$, future values are unique at all possible next states. The stage hypothesis makes the current Bayesian policy unique, and its Bellman expectations determine unique current values. Induction back through $t_D - 2, \ldots, 0$ proves the result. A convenient invariant domain at date t consists of $C = C_0 g^t$, $F_0 \exp(-t\delta) \leq F \leq F_0 g^t$, and all $R > 0$. Every feasible transition lies in the next such domain. It suffices to impose the stage hypothesis there, or on the smaller domain reachable under any feasible policy. The strictly increasing core stock identifies the date along this orbit, so the construction is Markov in the physical state. Any other equilibrium in the same class agrees in the dominant tail and, inductively, at every preceding state. □

The hypothesis is substantive. If $\beta m/N(1) < z < \beta m/N(0)$, immediate completion creates an interior dilution problem in the favorable-news tail. A lower-signal cutoff that forces universal research in that tail cannot be an all-signal equilibrium. Such states require a richer strategy class. Audits of selected numerical states do not prove the stage hypothesis over every off-path state or establish unrestricted dynamic uniqueness.

A.2.5 Small-noise beliefs, temporary funding costs, and fragility

Lemma 3 (The marginal firm's small-noise belief)**.** *Fix a bounded continuation return. Suppose cutoff equilibria satisfy $c(v) \to c$ as $v \downarrow 0$ and the return is continuous at $(\theta, a) = (c, a)$ for almost every $a \in (0,1)$. Then*

$$\int_0^1 Q(c,a)\,da = z. \tag{A.20}$$

The result also holds for converging continuation functions under the same dominated-convergence condition. It is a necessary limiting belief restriction, not a stand-alone existence or uniqueness theorem.

Proof. At the marginal signal, write $\theta = wc(v) + v\sqrt{w}Z$. Because $c(v)$ has a finite limit,

$$\theta \longrightarrow c, \qquad \frac{c(v) - \theta}{v} = \frac{(1-w)c(v)}{v} - \sqrt{w}Z \longrightarrow -Z.$$

The marginal firm's participation belief converges to $\Phi(-Z)$, which is uniform on $[0,1]$ by the probability integral transform. Dominated convergence in marginal indifference gives the uniform integral. Isolated completion boundaries in a have zero limiting rank probability, so almost-everywhere continuity suffices. □

A temporary cost or financing wedge. Let a publicly known, one-period increment f raise the current private cost from z to $z + f$, holding future costs and continuation functions fixed. The current condition becomes $G(c) - f = 0$. At a differentiable interior cutoff with $G'(c) < 0$, implicit differentiation and the Gaussian aggregation rule give

$$\frac{dc}{df} = \frac{1}{G'(c)} < 0, \qquad \frac{\partial A(\theta)}{\partial f} = \frac{\phi((c-\theta)/v)}{v\,G'(c)} < 0. \tag{A.21}$$

The participation derivative holds realized news fixed. Since the future regime is fixed, the wedge does not directly change the return conditional on current participation. Calling the wedge financial does not settle its social accounting: a resource cost belongs in welfare, while a transfer requires an identified counterparty.

The local response to small requirement news. Fix a state and verified private cutoffs with $c(v) \to c$. For a fixed $k > 0$, define

$$\theta_G = c(v) - kv, \qquad \theta_B = c(v) + kv. \tag{A.22}$$

The innovation gap is $2kv$ and tends to zero. Substitution into the Gaussian aggregation rule gives, exactly at every positive v,

$$A(\theta_G) - A(\theta_B) = 2\Phi(k) - 1 > 0. \tag{A.23}$$

Suppose returns converge at these innovations and their associated participation rates. Gross equity converges to (18) at the same limiting innovation c but at participation rates $\Phi(k)$ and $\Phi(-k)$. If the limiting Q and U are strictly increasing in participation between those rates, the adverse innovation causes positive limiting declines in both populations' gross equity with the ultimate prize unchanged. Increasing U raises the core price; nonnegative, strictly increasing Q makes $(1 + b\,a)Q$ strictly increasing as well. If the causal planner's current objective is strictly increasing on $[0, 1]$ for all nearby innovations, its unique action is one on both sides.

This sufficient local statement does not rank private and planner research at every state. It also does not infer a causal delay from a valuation decline. Completion times require following the state transitions.

A.3 Fundamental prices and the acceleration–dilution boundary

Aggregate equity and normal returns. Multiply (11) by F and (12) by C. Upon current completion, gross discounted claims sum to $\beta m[(1 + b)C + (1 + b\,A)F]/N(A) = \beta m$. Otherwise they sum to

$$\beta \exp(-\delta)[(1 + b)Cu(S') + (1 + b\,A)Fv(S')] = \beta I(S').$$

Subtracting the actual resource bill $bF \int z a_z(\theta)\, dH(z)$ gives (19) for the general cost distribution. The after-commitment prices are therefore gross of the current bill and net of future bills. Iterating to a finite horizon gives discounted receipts minus research bills, plus discounted remaining equity. The fundamental-value restriction removes that residual in the limit, so opening equity equals expected discounted net payouts. Equivalently, next-period gross receipts net of the current bill carried forward at rate $1/\beta$ have conditional expectation $I(S)/\beta$. This is the normal-return identity; positive research rents or a realized news-driven price jump do not violate it.

The constant-capacity boundary $r\tau = 1$. In a final research round followed by frozen capacity N and deterministic remaining work R, completion takes $\tau = R/N$. The discounted value per capability is $(m/N)\exp(-rR/N)$, and

$$\frac{\partial \log[(m/N)\exp(-rR/N)]}{\partial \log N} = \frac{rR}{N} - 1 = r\tau - 1.$$

The timing gain exceeds dilution when $r\tau > 1$ and is smaller when $r\tau < 1$. In the latter case other firms' research reduces a capability's value, but that value can still exceed its creation cost. Strategic substitution is therefore compatible with continued private research. This derivative is a transparent benchmark, not the derivative of the recurring return Q, which includes future program choices and changing ownership. In the discrete recurring economy the current-completion region gives the exact substitution result (38). Neither calculation is a general investment-stopping rule.

A.3.1 Physical memory: transition details and scope

Proposition 3 in the main text gives the conditional restart comparison for a common research cost. The physical accounting below applies to any cost distribution H. Fix S, current news, and two participation choices $a_H > a_L$. If both choices leave the project incomplete, then

$$F'(a_H) - F'(a_L) = \exp(-\delta)bF(a_H - a_L), \tag{A.24}$$

$$R'(a_L) - R'(a_H) = hbF(a_H - a_L), \tag{A.25}$$

while C' is the same. These identities follow from (5). Because $C' + F' = \exp(-\delta)N(a)$, differentiating (36) yields

$$\frac{d}{da}\left(\frac{R'}{C' + F'}\right) = -\frac{\exp(\delta)bF(R+\theta)}{N(a)^2} < 0.$$

Current continuation implies $R + \theta > hN(a) > 0$, establishing the sign. Hence lower participation leaves strictly more unfinished work per unit of surviving capacity. Translating that physical deterioration into a lower restart cutoff additionally requires (39).

A temporary shock can leave physical damage after it reverses. For an exact two-date comparison, start two histories at the same state. In the first date, one continues fully while the other pauses and receives an additional requirement innovation of size $\Delta\theta$. At the next date, reverse exactly that extra innovation and give both histories the same research participation a. Provided both remain incomplete, the paused history still has more work left after that second date, by

$$hbF\left[1+\exp(-\delta)(1+b\,a)\right]>0,$$

and fewer costly capabilities, by

$$\exp(-2\delta)(1+b\,a)bF>0.$$

The first-period news difference cancels, but the lost first-period work and the smaller capacity available for second-period work do not. If the two histories instead have the same subsequent participation at every date, their costly-stock ratio remains $1/(1+b)$ after the initial pause versus full research. These are physical comparisons holding subsequent actions equal; endogenous later actions can narrow or amplify the gap, so they are not an unconditional ranking of equilibrium paths.

Limits of the interpretation. Under complete pausing, total inherited capacity changes by

$$(C'+F')-(C+F)=\left[\exp(-\delta)(1+b)-1\right]C-\left[1-\exp(-\delta)\right]F.$$

Thus total capacity can grow even as costly capacity contracts. The core's next-period share also rises when costly research falls. Neither feature eliminates the strict deterioration relative to continued research, but both can offset its effect on the cutoff. Sufficiently favorable news can still induce research wherever the good-signal dominance condition holds. The model supports failure to reverse a pause with an ordinary improvement in news; it does not generally support the claim that no favorable news could ever reverse it.

A.4 Historical data and measurement

This appendix documents every series behind Section 2: what it measures, where it comes from, how it is constructed, and where it enters the analysis. It follows the order of Section 2. Table A1 summarizes the series. Subsection A.4.1 states the conventions common to all calculations. The next four subsections describe, in turn, stock prices around the peaks, commercial use, the news events and the contraction that followed, and the leader portfolios and their valuation recovery. The automobile episode is documented separately, and the last subsection states how each source was accessed and the terms of its use. Every number computed in Section 2 and this appendix is produced by the replication package from stored inputs; numbers quoted from published studies are cited where they appear.

These measures describe valuation recovery, investment contraction, and the later scale of use. Their reference dates do not identify the commercialization date that would have prevailed without a crash. Mapping them to a common work requirement or an inefficient delay therefore requires the model's maintained assumptions.

Table A1: Historical series: source, construction, and use

Series	Source	Construction	Enters
Stock prices			
Railroad stock prices	Macaulay index, NBER Macrohistory m11005	Monthly; deflated by CPI; April 1872 = 100	Fig. 1A; trigger shares
Railroad leaders, 4 roads	Yale ICF NYSE monthly prices, 1815–1925	Month-end prices; internal gaps interpolated	Fig. 1A
Railroad leaders, 6 roads	*Commercial and Financial Chronicle*, annual price tables, 1872–78	Midpoint of monthly high and low, transcribed from page images	Fig. 1A
Technology stock prices	Nasdaq Composite, FRED NASDAQCOM	Month-end close deflated by CPI-U; daily closes for news windows	Fig. 1B; trigger shares
Technology leaders	CRSP US Stock Database, via WRDS	Chained monthly returns excluding distributions	Fig. 1B
Business equipment	Fama–French 12 industry portfolios	Value-weighted total return	Sec. 2 text; Subsec. A.4.5
Commercial use			

Series	Source	Construction	Enters
Railroad freight	CHAT ton-kilometers	Annual; also divided by real GDP	Fig. 2A; Fig. 8
US real GDP	MeasuringWorth (Johnston and Williamson, 2026)	Millions of 2017 dollars	Fig. 8
E-commerce share	Census, FRED ECOMPCTSA	Annual mean of quarterly shares; logistic fit	Fig. 2B; Fig. 8
Investment, research, and exits			
New railroad track	NBER Macrohistory a02082a	Miles of first track per year	Fig. 2A; calibration target
IT equipment	BEA private fixed investment	Computers plus communications equipment, deflated by CPI-U	Fig. 2B
Venture capital	PwC/NVCA MoneyTree, via Federal Reserve Board (2007)	Annual, deflated by CPI-U	Fig. 2B
Technology IPO proceeds	Ritter, *IPO Statistics*	Annual, deflated by CPI-U	Fig. 2B
Business-funded R&D	NCSES, constant 2017 dollars	Annual	Sec. 2 text
Railroad receiverships	Swain (1898)	Annual count and mileage	Sec. 2 text
Start-up survival	Goldfarb, Kirsch, and Miller (2007)	Share surviving to 2004	Sec. 2 text
Automobile producers	Seltzer (1928), Table 5	Entries, exits, number operating	Text; Fig. A1B
Valuation recovery			
Railroad selection	*Poor's Manual* 1873	Size, then fiscal-1872 profitability	Tables 1, A5
Railroad dividends and prices	*Poor's Manual* (1873, 1880), *Financial Review*, *Chronicle*, Yale ICF; Daggett (1908)	Dividends 1873–96; 1895–98 terminal price; Erie assessments	$J(r)$: Table 1, Fig. 3
Technology selection	*FT Global 500*; annual filings	Size, then profitability at January 4, 2000	Tables 1, A4
Technology shareholder ledger	Price histories, issuer filings, depositary records	Payouts and claims per initial share	$J(r)$: Table 1, Fig. 3
Automobile indices	Cowles Commission indexes, groups 1, 41–43	Price and dividend-reinvested indices	Table A8; Fig. A1A
Passenger-car production	NBER Macrohistory m01107a	Monthly units	Sec. 2 text

Series	Source	Construction	Enters
Deflators			
Annual CPI	Federal Reserve Bank of Minneapolis	Interpolated to months	Railroad real values
Wholesale prices	NBER Macrohistory m04048a (Warren and Pearson)	Monthly, 1910–14 = 100	Railroad real values, alternative
CPI-U	BLS, FRED CPIAUCSL and CPIAUCNS	Monthly	Technology and automobile real values

A.4.1 Conventions

Each episode is dated from its sector peak: April 1872 for railroad stock prices, February 2000 (month-end) for the Nasdaq, and November 1919 for automobile stocks. Valuation calculations use reference dates fixed by the selection rule: the 1872 average price for railroads, January 4, 2000 for technology firms, and the May 1919–April 1920 average for automobiles.

Real values divide nominal values by the consumer price index of the same month or year. For railroads, the annual CPI of the Federal Reserve Bank of Minneapolis is placed at July of each year and interpolated log-linearly to months (Federal Reserve Bank of Minneapolis, 2026). For technology firms, the monthly CPI-U for all items, seasonally adjusted (FRED CPIAUCSL), deflates prices; the annual CPI-U deflates investment flows (U.S. Bureau of Labor Statistics, 2026). Automobile prices use the unadjusted CPI-U (FRED CPIAUCNS).

For the railroads we also report declines deflated by the monthly Warren–Pearson index of wholesale prices, NBER Macrohistory series m04048a (National Bureau of Economic Research, 2026a; Warren and Pearson, 1933). It is a different deflator, not a more finely observed consumer price index: it measures wholesale commodity prices, which fell twice as fast as consumer prices over these years (23.2% against 11.1% between April 1872 and June 1877). Deflated by wholesale prices, the Macaulay index falls 34.0% from its real high in August 1871 to June 1877, and the leader portfolio falls 34.7% from March 1872 to May 1877. Between the fixed dates of the CPI-deflated extremes, April 1872 and June 1877, the wholesale-price declines are 32.1% and 33.8%. With the interpolated CPI, the declines are 41.3% and 42.8%; in nominal

terms they are 47.8% and 49.1%. An interpolated annual index also smooths the monthly real series. Figure 1 uses the CPI so that both episodes share a consumer-price deflator.

The justified fraction $J(r)$ in (1) follows the claims purchased at the reference date: c_{is} is a cash receipt per initial dollar, net of any shareholder contribution, v_{iH} is the value of the remaining claim at the end date, and w_i is the initial portfolio weight. Throughout, r is an annual effective nominal discount rate, set to 8%, with 10% as the only alternative. A terminal market value capitalizes receipts after the end date; those receipts are not added again. The selection rule, weights, discount rates, and end dates are fixed before any outcome is computed.

A.4.2 Stock prices around the peaks

Figure 1 plots, for each episode, a sector price index and a buy-and-hold portfolio of the industry leaders defined in Subsection A.4.5. Both are monthly, exclude dividends, are deflated by the CPI, and equal 100 in the month of the sector peak. Each portfolio holds its securities in proportion to their market values at the reference date; each security's price is divided by its reference-date price.

Railroad sector. Macaulay's index of American railroad stock prices, monthly, NBER Macrohistory series m11005 (National Bureau of Economic Research, 1938), is a chain index weighted by shares outstanding.[7] Its real value peaks in April 1872. It falls 41.3% in real terms to June 1877 and 47.8% in nominal terms over the same months.

Railroad leaders. The portfolio holds the ten railroad leaders, eleven securities because Chicago & North Western common and preferred are both held, weighted by 1872 market value (par value times the 1872 average price). Monthly prices come from two sources.

- Lake Shore & Michigan Southern; Chicago & North Western, common and preferred; Chicago & Rock Island; and Erie: month-end prices from the Yale International Center for Finance NYSE monthly price file, 1815–1925 (Goetzmann et al., 2001). These securities account for 40% of the

[7] NBER source documentation: https://data.nber.org/databases/macrohistory/rectdata/11/docs/m11005.txt.

leaders' 1872 value. Internal gaps are filled by log-linear interpolation: over 1872–78, 12 of the 420 monthly observations of these five securities are interpolated (Chicago & North Western common 3 and preferred 5, Rock Island 3, Erie 1), including Rock Island in October 1873.

- New York Central & Hudson River; Delaware, Lackawanna & Western; Chicago, Burlington & Quincy; Central Railroad of New Jersey; Michigan Central; and Illinois Central: the midpoint of each month's highest and lowest sale price at the New York Stock Exchange, from the *Commercial and Financial Chronicle*'s annual table "Course of Prices of Railroad and Miscellaneous Stocks" (Commercial and Financial Chronicle, 1873–1879). These securities account for 60% of the leaders' 1872 value. The Yale file has no quotations for Chicago, Burlington & Quincy or Central Railroad of New Jersey in these years and quotes the other four only from 1873 or later.

The *Chronicle* tables used are, for prices in year y, those published in the first or second January issue of year $y + 1$: 1872, January 4, 1873 (vol. 16, p. 11); 1873, January 10, 1874 (vol. 18, p. 34); 1874, January 9, 1875 (vol. 20, p. 37); 1875, January 8, 1876 (vol. 22, p. 35); 1876, January 6, 1877 (vol. 24, p. 9); 1877, January 5, 1878 (vol. 26, p. 10); and 1878, January 4, 1879 (vol. 28, p. 11). All were accessed as page images through FRASER, Federal Reserve Bank of St. Louis. The tables report each month's highest and lowest prices at the New York Stock Exchange; from 1876 they state that they are compiled from all sales and print each month's low before its high. Values were read from the page images, because the scanned text layer is unreliable, and recorded with every hard-to-read entry flagged and explained. To check the transcription, we compare it with Yale month-end prices for all securities quoted in both sources: in 483 of 492 overlapping months (98%), the Yale price lies within the transcribed monthly range, up to half a dollar. The transcription flags 223 entries as hard to read, 222 of which carry a value and 153 of which enter the leader portfolio. Omitting the flagged values, moving each by half a dollar in either direction, or replacing it with the alternative reading changes the real decline of the leader portfolio by less than 0.02 percentage point, and no single cell changes it by more than 0.016 point; these are effects on the crash depth, not on every plotted month. The page images of thirteen January issues, 1873–1879 (the seven listed and the adjacent issues consulted to locate the tables), were downloaded from FRASER in October 2026. They are retained by the author and not redistributed (Subsection A.4.7).

Chicago, Burlington & Quincy distributed $25 of bonds per share to holders of record on December 10, 1872, in its consolidation with the Burlington & Missouri River Railroad (Commercial and Financial Chronicle, 1873–1879); its 1872 prices are reduced by $25 so that the series is continuous. With all ten roads, the leaders' real price falls 43% from April 1872 to June 1877. Alternative constructions give 42% (all eleven securities from *Chronicle* midpoints, 1872–1877) and 44% (without the Burlington adjustment). With only the four roads in the Yale file, the decline would be 58%; those roads include Erie and Chicago & North Western, which fell furthest.

Technology sector. The Nasdaq Composite, daily close, FRED series NASDAQCOM (Nasdaq, Inc., n.d.). Figure 1 uses month-end closes; its real month-end maximum is February 2000. The index falls 76.5% in real terms to September 2002. In daily nominal closes, its peak is 5,048.62 on March 10, 2000 and its trough 1,114.11 on October 9, 2002, a decline of 77.9%.

Technology leaders. The portfolio holds all ten technology leaders (Alcatel, Hewlett-Packard, IBM, Intel, Lucent, Applied Materials, Computer Associates, Microsoft, Ericsson, and CMGI), weighted by market value on January 4, 2000 from the *Financial Times* Global 500. Each leader's price is chained from monthly returns excluding distributions in the CRSP US Stock Database, accessed through Wharton Research Data Services (WRDS) (CRSP US Stock Database, 2026), starting from its January 4, 2000 close. The series is therefore a split-adjusted price without dividends. CRSP permanent numbers are 77604, 27828, 12490, 59328, 83332, 14702, 25778, 10107, 33452, and 80114, respectively. As a check, for the six leaders with public price histories, the CRSP index differs from one built from Yahoo Finance closing prices by a median of 1.7% over 1998–2006. The ten leaders fall 76% in real terms from March 2000 to September 2002.

Business-equipment portfolio. The value-weighted total return of the Business Equipment portfolio in the Fama–French 12 industry portfolios (French, 2026), compounded from April 2000 through September 2002, is −80% in nominal terms.

A.4.3 Commercial use

Railroad freight. US freight ton-kilometers come from the Cross-country Historical Adoption of Technology data set (Comin and Hobijn, 2009, 2010). Freight rose from 26.444 billion ton-kilometers in 1873 to 30.924 billion in 1877 and 37.753 billion in 1878, increases of 16.94% and 42.77%. Figure 2 plots the series relative to 1872.

For the calendar of Section 5, freight is divided by Johnston–Williamson real GDP from MeasuringWorth, in millions of 2017 dollars, with historical estimates linked to BEA data from 1929. For freight F_t in millions of ton-kilometers and GDP Y_t, the index plotted in Figure 8 is

$$I_t = 100 \frac{F_t / Y_t}{\max_{s \in \{1865,\ldots,1993\}} F_s / Y_s}.$$

Freight intensity reaches 751.8 ton-kilometers per $1,000 of real GDP in 1917. The index is 19.36 in 1872, 47.27 in 1899, and 52.77 in 1900. The first half-maximum crossing thus occurs 28 years after 1872, and the maximum 45 years after it. Subsequent levels of 73.37 in 1929, 51.51 in 1932, and 62.56 in 1943 show that 1917 is a historical maximum rather than a permanent plateau. Freight intensity also responds to production composition, haul distance, and transport mode.[8]

E-commerce. E-commerce use is the Census Bureau's estimated share of e-commerce in total retail sales, seasonally adjusted, distributed through FRED as ECOMPCTSA (US Census Bureau, 2026). Annual observations average the four quarterly shares, rather than dividing annual dollar totals; the 2026 point averages the first two quarters. The share rises from 0.925% in 2000 to 16.35% in 2025 and 17.05% in the first half of 2026. A logistic reference fitted with equal annual weights to 2000–2025 and that half-year observation has a ceiling of 24.52%. The observed share first exceeds half this reference in 2020. The projected 90% crossing around 2034 depends on the fitted ceiling, which has not been observed.

[8]GDP definition and construction: https://www.measuringworth.com/datasets/usgdp/. The normalization uses the 1865–1993 overlap with CHAT.

A.4.4 News and the subsequent contraction

Railroads, 1873. Historical explanations distinguish the event around which the panic crystallized from the vulnerabilities that preceded it. Mixon (2008) finds that equity, bond, and option prices reflected growing concern about riskier and more leveraged railroads before the panic. Nitschke (2018) emphasizes the informational consequences of Jay Cooke's unexpected failure: his reputation gave the bankruptcy significance beyond the balance sheet of one intermediary. The failure on September 18 and exchange closure on September 20 provide a clear chronology (Richardson and Sablik, 2015). They do not separately identify a primitive financing shock and the industry's endogenous response to deteriorating prospects.

British railways, 1845–50. The British sequence illustrates why proposed catalysts and subsequent contraction must be distinguished. Campbell (2013) finds that calls for further payments on partly paid shares put downward pressure on prices, but dividend changes have substantially greater explanatory power. Odlyzko (2011) examines the role of Robert Lucas Nash's published financial analyses, which contemporaries blamed for the 1848 crash, and finds their influence smaller than those accounts suggested. George Hudson's accounting irregularities, the episode's best-known scandal, came to light in 1849, at the trough in prices. Odlyzko (2010) argues that the shortfall in aggregate revenue was foreseeable by 1846; the losses he documents come from construction costs about half above projections and revenues 30 to 40 percent below them, and he notes that corruption was also widespread in the 1830s boom, which investors later judged a success. Equity prices had already peaked in 1845. Construction expenditure nevertheless rose from £13 million in 1845 to £44 million in 1847 before returning to £13 million in 1850. Previously committed construction can therefore continue after valuations turn down; the two contractions need not begin together.

Internet firms, 2000. Several proposed catalysts concern the conditions under which firms could remain active before becoming profitable. Willoughby (2000) reassesses cash needs using public late-1999 accounts. Ofek and Richardson (2003) instead link the collapse to expiring IPO lockups, additional potential sellers, and constraints on short selling. Those expiration dates were prescheduled. Monetary policy also tightened: the federal funds target rose from 4.75 to 6.5 percent between June 1999 and May

2000, a sequence spanning the March 2000 market peak (Federal Open Market Committee, 1999, 2000). These explanations concern different changes in information, market participation, and discounting. Their chronology does not establish a single initiating shock or measure its fundamental magnitude.

Price changes in the news windows. Table A2 compares sector-price changes around dated events with the full decline. For a window from a to b, its share of the eventual log loss is

$$\frac{\log(P_a/P_b)}{\log(P_{\text{peak}}/P_{\text{trough}})},$$

computed on nominal prices. The railroad denominator is the nominal Macaulay-index decline of 47.81% between April 1872 and June 1877. The Nasdaq denominator is the nominal decline of 77.93% between the daily close of March 10, 2000 and that of October 9, 2002. These are price-index measures, separate from the shareholder valuation calculations.

Table A2: Price changes around dated news events

Index and event window	Price change	Share of log loss
Railroads, Aug.–Oct. 1873: September panic	−15.26%	25.47%
Nasdaq, Mar. 17–20, 2000: *Barron's* issue date	−3.92%	2.65%
Nasdaq, Mar. 31–Apr. 3, 2000: Microsoft conclusions of law	−7.64%	5.26%
Nasdaq, Apr. 13–14, 2000: March CPI release	−9.67%	6.73%

Notes: Macaulay railroad prices, NBER Macrohistory m11005, monthly; Nasdaq Composite daily closing prices, FRED NASDAQCOM. Each Nasdaq row contains one close-to-close trading session.

The railroad window brackets the failure of Jay Cooke & Co. and the exchange closure (Richardson and Sablik, 2015). The March 20 issue of *Barron's* discussed internet firms' cash needs (Willoughby, 2000). On April 3, the district court issued its conclusions of law in *United States v. Microsoft*. On April 14, the BLS released March CPI at 8:30 a.m.; the seasonally adjusted increases were 0.7% for the headline index and 0.4% excluding food and energy.[9] Together, the *Barron's* and CPI sessions account for 9.38% of the full log decline; adding the Microsoft session raises the share to 14.63%. The entire decline from the March 10 peak through April 14 accounts for 27.71%. The windows describe when price losses occurred; other

[9]Bureau of Labor Statistics, *Consumer Price Index: March 2000*, April 14, 2000, USDL-00-102: https://www.bls.gov/news.release/history/cpi_04142000.txt.

news and changing expectations also affect those prices. They do not measure the size of an exogenous technology shock.

Dating and sensitivity of the windows. Each 2000 window is the close-to-close change over the first trading session in which the news could trade: the session ending on the issue date of *Barron's*, on the date of the court's conclusions, and on the CPI release date. The railroad window runs from the month before the event month to the month after it, because the Macaulay index is monthly. These rules make the dating transparent, but they were applied after the episodes, in a retrospective exercise. Table A3 moves each window. The three 2000 sessions together account for 14.6% of the eventual loss at the stated dates, but between −7.0% and 5.6% when shifted by one or two sessions; a negative share means that prices rose in the shifted window. The railroad share is 25.5% for August–October 1873 and 12.9% or 19.1% when the window is moved by a month in either direction. The window shares are therefore fragile measures of chronology, not measurements of the news.

Table A3: Share of the eventual log price loss in shifted news windows

	Shift of the window				
Window	−2	−1	0	+1	+2
Nasdaq Composite, 2000; shifts in trading sessions					
Barron's issue, March 20	−1.92	−1.12	2.65	−1.44	−2.12
Microsoft conclusions of law, April 3	2.72	−1.68	5.26	1.18	−0.32
March CPI release, April 14	4.84	1.65	6.73	−4.20	−4.59
Barron's and CPI sessions	2.93	0.53	9.38	−5.65	−6.71
Three sessions combined	5.64	−1.16	14.63	−4.47	−7.03
March 10 peak through April 14 (end shifted)	19.33	20.98	27.71	23.51	18.91
Macaulay railroad index, 1873; shifts in months					
August–October (baseline window)		12.87	25.47	19.05	
September–October			13.42		
August–September			12.05		
July–November			31.92		
August–November			31.10		

Notes: Percent of the eventual nominal log price loss (Appendix Table A2 gives the denominators). Column 0 uses the stated dates. Each 2000 window is the close-to-close change over the first session in which the news could trade; a shift moves the whole window by the stated number of sessions. A negative share means that prices rose in the shifted window. Railroad alternatives change the window's span rather than shifting it. Computed by the Section 2 data reviewer from the series of Appendix A.4; see `data/historian_update2/`.

Railroad construction. Miles of new first track are NBER Macrohistory series a02082a, drawn from Census series Q43 and *Railway Age* (National Bureau of Economic Research, 2026b). Completions fall from 7,439 miles in 1872 to 1,606 in 1875, a decline of 78.4%, and first regain the earlier annual flow in 1881. The pre-1880 history covers roads operating in 1880, so the series measures a defined construction flow rather than all railroad expenditure or the entire original population of projects. The 1872–75 decline is the target used to fit the initial core share in Section 5.

Information-technology investment. Computer and communications equipment is the sum of BEA private fixed investment in computers and peripheral equipment (series B935RC1Q027SBEA, the annual average of quarterly annual rates) and in communications equipment (series C275RC1A027NBEA), in current dollars (U.S. Bureau of Economic Analysis, 2026). A common CPI-U deflator converts the sum to consumption units; it is not a chained equipment-volume aggregate. Venture capital investment is PwC/NVCA MoneyTree totals as reproduced in the Federal Reserve Board's 2007 report to Congress on small-business credit, Figure 7 (Board of Governors of the Federal Reserve System, 2007); the series ends in 2006. Technology IPO gross proceeds are from Ritter (2026). Venture investment and IPO proceeds measure financing flows and are reported separately from expenditure. Between 2000 and 2003, in CPI-deflated terms, equipment expenditure falls 33.0%, venture investment 82.4%, and technology IPO proceeds 95.1%.

Research and development. Business-funded research and development is from Figure 4 of National Center for Science and Engineering Statistics (2025), in billions of constant 2017 dollars. It equals 255.73 in 2000, 253.37 in 2001, 239.29 in 2002, 241.76 in 2003, 241.98 in 2004, 254.76 in 2005, and 270.22 in 2006. The 2000–2002 decline is therefore 6.43%, and 2006 is the first year that regains the 2000 level. The series measures business funding of domestic R&D across performing sectors, rather than technology-sector R&D alone.

Failures and exits. Railroad receiverships established each year and the mileage placed in receivership come from Swain (1898), as tabulated in Hansen's "Bankruptcy Law in the United States" (EH.Net Encyclopedia, Table 3). From 1873 to 1879, 183 receiverships covered 24,388 miles, and the mileage placed

in receivership in 1875 was 9.9% of the network. The survival rate of internet start-ups that sought venture funding in 1998–2002 is from Goldfarb et al. (2007). Entries, exits, and the number of passenger-car producers are from Seltzer (1928), Table 5; Subsection A.4.6 discusses their interpretation.

A.4.5 Leader portfolios and valuation recovery

Railroad selection. The initial size universe consists of the thirty largest roads by stock and funded debt in *Poor's Manual* for 1873. Requiring an NYSE common-stock quotation and excluding guaranteed-rental, leased, and closely held claims leaves sixteen eligible firms. The profitability ranking divides fiscal-1872 net earnings by stock and funded debt. The ten highest-ranked firms, A10, are Delaware, Lackawanna & Western; Rock Island; Central Railroad of New Jersey; Chicago, Burlington & Quincy; New York Central & Hudson River; Chicago & North Western; Lake Shore & Michigan Southern; Michigan Central; Illinois Central; and Erie. A5 comprises the first five. Chicago & North Western has both common and preferred shares in the portfolio, so A10 contains eleven securities, each weighted by its 1872 market value.[10]

The railroad cohort is selected from historical accounts rather than from firms surviving today. Its information date nevertheless matters: fiscal-1872 accounts were published in 1873, after the average-price baseline. It is therefore a retrospective profitability screen, not a trading strategy implementable throughout 1872. Some earnings and debt entries are historical estimates. Railroad net earnings are before interest, whereas the technology screen below uses net income after interest; the rankings identify profitable firms within each episode without imposing identical accounting measures across episodes.

Technology selection. The technology size universe comprises thirty large US-listed firms from the *FT Global 500*, including foreign issuers' US-listed claims. Initial equity capitalizations are dated January 4, 2000. The ranking statistic is

$$\text{score}_i = \frac{\text{reported annual net income available at the cutoff}}{\text{January 4, 2000 equity capitalization} + \text{financial debt}}.$$

[10] *Poor's Manual of Railroads of the United States*, 1873: https://archive.org/details/poorsmanualofrai00unse.

Calendar-year firms generally enter with fiscal-1998 accounts. Firms whose later fiscal years had already been reported, including Microsoft, Lucent, Computer Associates, and CMGI, enter with those accounts. Annual filings and issuer statements provide the income and debt observations; Applied Materials' November 17, 1999 results announcement supplies its fiscal-1999 figures.

Table A4: Technology profitability ranking near the portfolio boundary

Firm	Initial cap.	Net income	Debt	Score (%)	Fiscal year	Report date
ALA	45.6	2.6217	4.5992	5.2227	1998	See note
HP	108.6	3.4910	4.8690	3.0766	Oct. 1999	See note
IBM	201.0	6.3280	29.4130	2.7464	1998	Mar. 29, 1999
INTC	277.1	6.0680	0.8610	2.1830	Dec. 1998	Mar. 26, 1999
LU	237.7	4.7660	7.0260	1.9475	Sep. 1999	Dec. 21, 1999
AMAT	46.5	0.7466	0.6266	1.5843	Oct. 1999	Nov. 17, 1999
CA	37.0	0.6260	2.5240	1.5838	Mar. 1999	May 28, 1999
MSFT	586.2	7.7850	0.0000	1.3280	June 1999	Sep. 28, 1999
ERIC	129.6	1.6404	2.4181	1.2425	1998	See note
CMGI	39.0	0.4762	0.0403	1.2199	July 1999	Oct. 29, 1999
DELL	119.6	1.4600	0.5120	1.2155	Jan. 1999	Apr. 27, 1999

Notes: Dollar amounts are billions. Calculations use unrounded income and debt. Report dates refer to original filings except for Applied Materials' results announcement. HP announced its fiscal-1999 results in November 1999; its annual report was filed January 27, 2000. Using its fiscal-1998 accounts leaves A5 and A10 unchanged. Alcatel's French-GAAP figures are reported in later issuer historical statements. Ericsson's 1998 board report is dated January 28, 1999; that date need not be its public-release date. Foreign-currency amounts are converted using annual-average exchange rates.

The first ten firms in Table A4 form A10, and the first five form A5. The closest membership comparison is CMGI versus Dell. CMGI reported net income of $476.240 million and financial debt of $40.318 million. At its $39.0 billion initial capitalization, its score slightly exceeds Dell's. Using income attributable to common shareholders, $474.578 million, preserves this ordering, but the gap is small relative to the rounding of capitalizations. Including CMGI's $411.283 million of redeemable preferred capital in the denominator puts Dell ahead. The rule uses reported net income and financial debt, so CMGI enters A10.[11] Other accounting choices change ranks without changing membership. HP's fiscal-1998 income of $2.945 billion and debt of $3.308 billion leave the selected sets intact. Replacing Alcatel's French-GAAP income of EUR2.340 billion with US-GAAP income of EUR1.179 billion does the same. Applied Materials and Computer Associates are nearly tied at ranks six and seven, and Microsoft is eighth.

Two qualifications apply to the technology screen. First, it conditions on the FT size universe and the available historical accounts. The predecessor firm's fiscal-1998 observation for Veritas is unavailable,

[11] See CMGI's October 1999 filing and Applied Materials' results announcement.

and some foreign issuers' original publication dates and accounting entries are incomplete. Second, reported net income includes nonoperating gains. CMGI had an operating loss and substantial investment gains; Alcatel's income included disposals. A screen based on operating income or internal financing would identify a different economic concept from the reported-income rule used here.

Railroad dividends, prices, and assessments. Railroad justified fractions use annual dividends from 1873 through 1896 and the 1895–1898 average price. Dividends and prices are from *Poor's Manual of the Railroads of the United States* (1873, 1880), the *Financial Review* (1889, 1893), the *Commercial and Financial Chronicle*, and the Yale NYSE file; Erie's shareholder assessments are from Daggett (1908). Of 264 annual dividend entries, including zeros, 147 have cited observations and 117 are estimates. The seven-year statement for the Chicago, Rock Island & Pacific in the 1880 *Manual* (Poor, 1880) reports dividends of \$1,659,172 in fiscal 1873–74, \$1,678,384 in each of 1874–75 and 1875–76, and \$2,097,980 in 1876–77, and prints the capital stock as \$20,500,000 in 1873 and \$20,980,000 in 1874–79. The payouts are therefore about 8% of the stock through 1875–76 and 10% in 1876–77. The 1876 entry, previously estimated at 8%, is 10%, and the 1873–75 entries are confirmed at 8%; a fiscal year beginning in April of year t is assigned to year t. The correction raises railroad $J(0.08)$ by 0.08 point for A10 and 0.16 point for A5. Nine of the eleven terminal prices equal the mean of the Yale month-end quotations for 1895–98. The other two, Rock Island and the reorganized Erie, are estimates from *Chronicle* ranges; replacing them with their single Yale quotations raises $J(0.08)$ by at most 0.3 point. Annual dividends are dated $t_s = y - 1872$. Stock distributions multiply subsequent claims, and Erie assessments enter as negative cash flows at their annual or fractional-year dates. The 1895–1898 average terminal quotation is dated 24.5 years after the baseline.

Technology shareholder claims and corporate actions. Each technology investment is followed on a consistent share-unit basis. Ordinary dividends and cash acquisition proceeds enter c_{is}; stock distributions alter the quantity of subsequent claims. A cash subscription or reorganization assessment is a negative cash flow. Repurchases are not added to the dividends of a shareholder who retains the position. Their effect, together with new issuance, is reflected in per-share values and ownership. Opening equity capitalizations determine w_i rather than replacing this shareholder accounting.

The historical price series embed some spin-off adjustments that reinvest the distribution in the parent. Other distributions are valued at the event date and held as cash. Lucent's stock acquisition and Alcatel's exchange into Nokia are treated in the latter way. Annual observations for delisted Alcatel, Lucent, and Computer Associates contain estimated prices, as do some acquisition-event valuations. Alcatel's 2013 rights issue is not fully resolved.

Ericsson's closing-price series adjusts for stock splits but not for the bonus element of the 2002 rights issue. The baseline investor does not subscribe. We consider rights-sale proceeds ranging from zero to the highest underlying-share price during the rights-trading window, $4.80 per modern normalized unit, discounted from August 15, 2002. The upper case is a price-based scenario for an unobserved sale rather than an observed rights quotation. Depositary distribution records determine the 2005–2013 cash payments, including the 2012 payment, with pre-2008 amounts converted for the ADS ratio. Duplicate or inconsistent provider cash entries are excluded. Other dividends retain their price-provider observations. These conventions give Ericsson $J(0.08) = 5.09–10.08\%$ and $J(0.10) = 3.88–8.63\%$.[12]

CMGI's subsequent claims are followed through ModusLink and Steel Connect. The cash-flow calculation incorporates the 2007 one-for-ten reverse split, the March 31, 2011 special dividend of $0.9134 per then-current share, the 2023 reverse/forward split, and the January 2, 2025 payment of $11.45 plus a contingent-value right. The position is assumed large enough to survive the 2023 small-shareholder cash-out; fractional rounding is immaterial at portfolio scale. On the post-January-2000-split basis, the identified 2011 dividend is $0.09134 and the 2025 cash payment is $0.1226786. The exact January 4, 2000 quotation is unavailable. The annual report gives split-adjusted high and low prices of $163.50 and $48.09 for the containing quarter. We vary the opening price across those bounds and the contingent right between zero and the issuer's $1.15 estimate per 2025 share. The latter is a valuation assumption, not a contractual maximum. CMGI therefore contributes identified cash to $J(r)$, although its complete intervening price path is unavailable.[13]

[12]Citi's depositary distributions: https://depositaryreceipts.citi.com/adr/guides/pgm_dispaDivd.aspx?cusip=294821608&pageId=15&subpageID=113. Ericsson's rights-offer terms: https://news.cision.com/ericsson/r/ericsson-sets-terms-for-rights-offering,c2247752.

[13]CMGI, *2000 Annual Report*, p. 64: https://www.annualreports.com/HostedData/AnnualReportArchive/m/NASDAQ_MLNK_2000.pdf; ModusLink's 2011 dividend announcement: https://www.sec.gov/Archives/edgar/data/914712/000119312511057745/dex992.htm; Steel Connect's 2023 Form 10-K: https://www.sec.gov/

For technology, elapsed time is actual days divided by 365.25 from January 4, 2000. The terminal date is October 19, 2026, so $H = 26.7899$ years. Annual dividend totals are dated at the last quoted trading day of their year unless transaction-level dates are available. Microsoft's series uses ex-dividend dates as proxies for payment dates; Ericsson combines depositary payment dates with the remaining provider dates. This convention preserves the timing of identified transactions while avoiding false daily precision for annual totals.

Results. Table A5 reports $J(r)$ for the four portfolios with initial-value weights; Table 1 in the main text rounds the same values. The technology A10 range varies Ericsson rights proceeds and CMGI's opening-price and contingent-right assumptions. The alternatives reported are the 10% discount rate and, for railroads, a later end window (Table A6) and a calculation that sets every estimated dividend to zero. The last gives railroad A10 48.3% at 8% and 38.0% at 10%. It still uses estimated opening and terminal prices, so it measures sensitivity to the dividend estimates rather than a lower bound on shareholder recovery.

Table A5: Valuation recovery at the two discount rates

Portfolio, initial-value weights	$J(0.08)$	$J(0.10)$
Railroads A10	70.72	57.07
Railroads A5	87.32	71.80
Technology A10	70.61–71.00	46.11–46.47
Technology A5	33.74	24.12

Notes: Percentages of initial valuation. Railroad accounts include estimated dividend observations.

Table A6: Railroad valuation recovery under alternative end windows

Calculation, initial-value weights	$J(0.08)$	$J(0.10)$
A10, 1895–98 terminal prices (baseline)	70.72	57.07
A10, 1899–1902 terminal prices and dividend carry	74.11	58.44
A10, 1899–1902 terminal prices without added dividends	71.99	57.13

Notes: Percentages of initial valuation. Later terminal prices average annual quotations in 1899–1902 where at least three monthly observations are available, and are dated at 28.5 years. The dividend-carry case extends each road's 1896 dividend rate through 1900. Rock Island retains its original terminal observation and date, without additional dividends.

Archives/edgar/data/914712/000162828023037522/stcn-20230731.htm; and its January 2025 Form 8-K: https://www.sec.gov/Archives/edgar/data/914712/000121390025000385/ea022628901-8k_steel.htm.

Annual nominal discount rates need not correspond to the same real required return across the two episodes. The alternative calculation that equalizes realized average real rates uses a railroad nominal rate of approximately 3.65%, giving A10 a justified fraction of about 125%. This is an inflation-adjusted check of the discount-rate rule rather than an estimate of investors' original expected inflation or risk compensation. The model's real, per-period discounting and its calendar conversion are defined separately in Section 5.

Industry-level context. For technology, the Fama–French business-equipment portfolio provides an industry counterpart to the automobile indices of Subsection A.4.6. A dollar placed in its value-weighted portfolio at the start of January 2000 and held through August 2026, with distributions reinvested, has $J(0.08) = 125\%$ and $J(0.10) = 77\%$; starting in February or March 2000 gives $J(0.08)$ of 133% and 113%. The portfolio follows the industry as a whole, including firms that entered after 2000, rather than a fixed cohort of claims selected at the peak.

A company's aggregate dividends, repurchases, and terminal market capitalization do not measure the receipts of its original shareholders unless subsequent issuance and changes in ownership are also followed. Table A7 illustrates the distinction. Subtracting cash issuance from company-wide distributions moves the issuer measure toward the shareholder measure, but it does not account for all stock-financed acquisitions, noncash issuance, or successor ownership fractions.

Table A7: Shareholder and issuer measures for technology A5

Measure, initial-value weights	8%	10%
Initial shareholders' $J(r)$	33.74	24.12
Issuer gross distributions and terminal capitalization	48.38	38.53
Issuer distributions net of cash issuance and terminal capitalization	43.53	34.40

Notes: Percentages of initial equity capitalization. Issuer distributions include dividends, gross repurchases, and specified distributions. Fiscal-year cash totals are prorated for the portion after January 4, 2000 and discounted from that portion's midpoint; dated transactions retain their event dates.

Losses, failures, and relative value of the leaders. Sector-index drawdowns, fixed-portfolio returns, and firm-level worst losses answer different questions. The railroad loss comparison in Section 2 uses the eleven A10 securities and four western comparison securities (Union Pacific; Columbus, Chicago

& Indiana Central; Ohio & Mississippi; and Toledo, Wabash & Western). For each security, the worst decline is the lowest annual real price in 1873–1880 relative to its 1872 average; the text reports the median. A separate 1873–1880 receivership comparison records two of ten leader roads and eight of ten comparison roads; restricting the entries to primary-source-confirmed events gives one and six. These are descriptive comparisons of historically selected groups. Financial relative value also differs from product-market share. A thirteen-road freight sample, which is not A10, has a share that falls from 41.3% in 1882 to 31.9% in 1890; its flat earlier share partly reflects how pre-1882 totals were assembled.

A.4.6 Automobiles, 1919–1926

The American automobile episode provides a shorter interval between the price collapse and subsequent valuation recovery. Cowles' monthly automobile-and-truck price index falls 64.22% from its November 1919 peak to January 1922. We use its price and dividend-reinvested indices for all automobiles (group 41), the General Motors group (group 42), other listed makers (group 43), and the broad stock market (group 1) (Cowles III and Associates, 1938). The GM group includes Fisher Body until July 1926. Ford was privately held and is absent from these listed-stock comparisons.

Cowles' indices are value-weighted and chain-linked, with changing constituents. Their valuation recovery therefore describes the corresponding evolving index portfolio, rather than a fixed cohort of firms selected in 1919. Let P_{is} and T_{is} be its price and dividend-reinvested indices in month s. Cash per initial index investment is inferred from

$$c_{is} = \frac{P_{i,s-1}(T_{is}/T_{i,s-1}) - P_{is}}{\bar{P}_{i0}}, \qquad v_{iH} = \frac{\bar{P}_{iH}}{\bar{P}_{i0}}.$$

The baseline $\bar{P}_{i0}$ is the May 1919–April 1920 average. The central terminal value $\bar{P}_{iH}$ averages 1924–1926 quotations and is assigned to that window's midpoint, 5.67 years after the initial window midpoint. Cash payments after April 1920 and through the terminal midpoint are discounted monthly. This window convention reduces dependence on a single initial or terminal month and ends before the 1927–1929 stock-market boom.

Table A8: Automobile index valuation recovery

Window and rate	GM group	Other makers	All autos	All stocks
1924–1926 endpoint, $r = 8\%$	102.40	105.14	99.24	101.09
1924–1926 endpoint, $r = 10\%$	93.07	95.63	90.26	92.13
1923–1925 endpoint, $r = 8\%$	72.75	88.25	78.22	93.33
Maximum-month initial price, $r = 8\%$	86.41	92.91	86.31	96.75

Notes: Percentages. Cowles groups 42, 43, 41, and 1, respectively. Unless otherwise indicated, the initial reference is the May 1919–April 1920 average; the maximum-month variant retains the 1924–1926 terminal window.

At 8%, discounted dividends and terminal values cover approximately the initial valuation for both the GM group and other listed makers. The broad market has a similar result. An earlier endpoint or a single peak-month initial price lowers these fractions, as Table A8 shows.

The automobile inputs are transcriptions. In all 120 rows of the transcribed Cowles tables, the annual figure equals the mean of the transcribed months to print precision. The all-automobile group lies outside the range of the GM and other-maker groups in 16 of 360 month-cells, all in one run from November 1919 to February 1921. Because the indices are chain-linked with changing constituents, a group need not lie between its subgroups, so these cells are flagged for comparison with the printed page rather than treated as errors; no Cowles cell has yet been compared with a scan of the original. Seltzer's counts agree with the public scan of his Table 5 in 29 of 30 cells for 1917–26. For 1918 the scan's text layer reads 73 producers remaining; the transcribed 71 satisfies the stock-flow identity in both 1918 and 1919 and is used.

Production and industry structure also changed. Monthly passenger-car output fell from 193,600 in March 1920 to 45,400 in January 1921. Annual output was 1.905 million in 1920, 1.442 million in 1921, 2.272 million in 1922, and 3.624 million in 1923.[14] Seltzer's count of producer enterprises falls from 88 in 1921 to 44 in 1926, with eight gross entries and 52 gross exits during 1922–1926 (Seltzer, 1928). The count treats producing units within GM separately; these are enterprise-level turnover totals, not exits from an identified cohort of 88 independent corporations. Figure A1 places that turnover beside listed prices. Panel A plots real prices without dividends, as in Figure 1, so that the depth and timing of the decline are visible: General Motors falls 59 percent and the other listed makers 61 percent from the reference

[14]NBER Macrohistory passenger-car production, M0107AUSM543NNBR: https://fred.stlouisfed.org/series/M0107AUSM543NNBR. The series is not seasonally adjusted and covers approximately 90% of production through 1921, with full coverage from 1922.

average to their lows. Dividends are essential for valuation, and they enter through $J(r)$: with dividends discounted at 8%, the same indices deliver 102 and 105 percent of their reference value (Table A8).

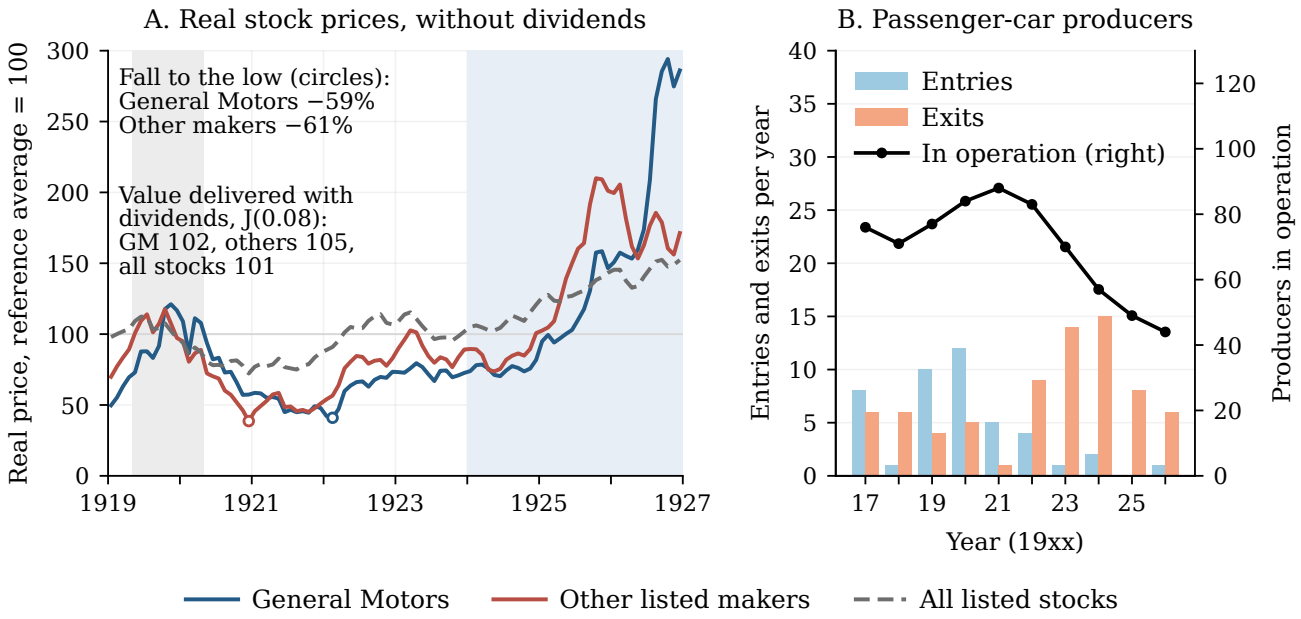

Figure A1: Automobile stock prices and producer turnover

Notes: Panel A: Cowles price indices (series P) for General Motors, other listed automobile makers, and all listed stocks, deflated by the CPI and set to 100 at the May 1919–April 1920 average (gray band). Dividends are excluded from the plotted lines; they enter the justified fraction $J(0.08)$ printed in the panel, computed from the price and dividend-reinvested indices (series P and C) with the 1924–1926 end window (blue band). The indices have changing constituents. Panel B: Seltzer's passenger-car producing enterprises, including entries, exits, and the number in operation. Sources: Cowles III and Associates (1938); Seltzer (1928).

The automobile evidence thus combines a deep listed-price decline, subsequent recovery of much of the initial valuation, and substantial turnover among producers. It complements the longer railroad and technology episodes without treating the index portfolios as equivalent to their fixed initial-firm samples.

A.4.7 Data access, citation, and terms of use

The historical series combine public, licensed, and permission-restricted sources. Access for this paper does not imply permission to redistribute every series, and the replication files keep the two apart.

- *CRSP through WRDS.* The technology leader index in Figure 1B is calculated based on data from the CRSP US Stock Database, Center for Research in Security Prices, LLC, an affiliate of the University of Chicago Booth School of Business (CRSP US Stock Database, 2026). The data were provided by Wharton Research Data Services (WRDS), which was used in preparing this paper. The monthly

stock file was extracted on October 22, 2026 for the ten permanent numbers listed above, June 1997–December 2006, with the fields RETX, RET, PRC, CFACPR, and SHROUT, and the daily file supplied the January 3–4, 2000 returns. The firm-level records are licensed and are not part of the replication files, which contain only the value-weighted aggregate index. This version uses no Compustat data.

- *Commercial and Financial Chronicle.* The page images were obtained from FRASER, Federal Reserve Bank of St. Louis (Commercial and Financial Chronicle, 1873–1879), whose terms require attribution and leave rights with the owners of the material. The transcriptions, which give the issue, volume, and page of every table, are part of the replication files; the images are not redistributed.

- *MeasuringWorth.* Real GDP is from Johnston and Williamson (2026). MeasuringWorth permits use for non-profit educational purposes with proper credit and asks that other uses be cleared with it. The paper prints only ratios derived from the series (Figure 8); permission to republish the series will be sought before the replication files are released publicly.

- *Other sources.* The NBER Macrohistory Database (Macaulay railroad prices, new track, passenger-car production, and Warren–Pearson wholesale prices), the Yale International Center for Finance NYSE price file, the Kenneth French data library, FRED, BEA, BLS, the Census Bureau, NCSES, Jay Ritter's IPO statistics, and the Federal Reserve Bank of Minneapolis CPI are publicly available and cited where used. *Poor's Manual* (1873, 1880) and Seltzer (1928) were read from public scans at the Internet Archive, and the Cowles indexes were transcribed from Cowles III and Associates (1938).

A.5 Computation and numerical scope

The quantitative exercise uses the common-cost specialization: H is a point mass at $z = 2.5$. The core stock is deterministic, $C_t = Cg^t$ with $g = \exp(-\delta)(1 + b)$. We therefore solve each date's continuation values on a two-dimensional grid in $(R, F/C)$: 241 equally spaced work nodes on $[0, 30]$ and 65 ratio nodes on $[0, 4.1657589]$. The full recursion underlying this example was solved backward from computational date 40. These value functions are reused when the historical timing interpretation changes; no calendar

interpretation enters their per-period economic parameters. The continuation at date 40 and outside the work grid is zero. This is a numerical truncation, not a forced commercialization date. Interpolation is bilinear, while an actual completed transition uses the terminal payoff directly.

For the private game, a candidate signal cutoff solves marginal indifference by bisection. Normal-density integrals are split at the cutoff neighborhood, the completion threshold, and continuation-grid crossings. The prior integral uses 12 Gauss–Legendre nodes per segment; the posterior uses 16. Candidate policies are then checked for profitable deviations at 50 signal points, including points around the cutoff and in the prior tails. When this check fails, the code uses a finite Bayesian signal game and a potential-based local search with a recorded iteration limit. This is not a proof of global dynamic uniqueness.

The full grid contains 626600 state–date nodes. Dates 0 through 10 have continuous-cutoff candidates passing the recorded signal checks at every node, with maximum checked regret 1.1×10^{-11}. At later dates, 125146 nodes use the finite approximation, and 35435 of these remain flagged as unresolved under the iteration limit. Their locations and residuals are retained in the numerical files; the calculation is not presented as a uniformly certified solution over the whole state grid. All states visited in the reported main simulation have a continuous cutoff passing the recorded checks, with maximum checked regret 3×10^{-12}. Current games are re-solved against their continuation functions at every visited state; the simulation does not interpolate a cutoff or an action across states.

The planner uses the same full state grid and date-40 truncation. It integrates current innovations with 48 Gauss–Legendre nodes per segment in Gaussian probability, splitting known completion thresholds. The continuous action optimum is found by checking the interpolation-cell boundaries, the exact completion trigger, and interior vertices of the piecewise quadratic objective. The planner observes current news but integrates future news. The public-state rule uses that planner action at zero current news; its realized transitions still use the actual innovation.

Positive core growth also gives a physical tail check. Even if costly research never resumes, completion by date j receives cumulative work of at least

$$hC(1+b)\frac{g^j-1}{g-1} + hF\frac{1-\exp(-\delta j)}{1-\exp(-\delta)}.$$

The sum of the first j requirement innovations is normal with standard deviation $\sigma\sqrt{j}$. At the initial state, the resulting upper bound on the probability of noncompletion by date 40 is smaller than 10^{-458}. The discounted aggregate terminal-prize remainder under a fixed policy is at most 1.52×10^{-6}. This probability calculation does not establish uniqueness or stability of an equilibrium selection. The same physical bound, applied to the work coordinates of the unresolved regions and their adjacent interpolation cells, bounds the probability of entering any such region by 3.2×10^{-9} under any research policy. The core-only return bound makes costly research uniformly dominated from date 43 onward, independently of news; the simulation ends earlier through completion, not through that dominance date.

Initial remaining work is maintained at $R \simeq 8.7$. The core share alone is fitted to the conditional capability-flow contraction. The unconditional and adverse-news conditional completion means are reported separately. The independent validation uses 20,000 Gaussian histories and seed 20261020. The good/bad comparison uses 10,000 antithetic histories and seed 8675309. Future innovations are identical across the two initial-news economies and across private, planner, and public-rule policies; each economy keeps its own endogenous state. Standard errors in the antithetic experiments use independent pair means. The standard error of $M = d/d_P$ uses the paired influence quantity $(d_i - M d_{P,i})/d_P$. Pricing checks compare simulated discounted receipts, net of all funding calls, with the recursive net owner claims.

Research-flow recovery means the first later decision date at which $b(C_t + F_t A_t)$ regains its favorable-news initial level. If the flow has not recovered before completion, the observation is censored; it is not assigned a recovery date. The reported conditional recovery dates therefore come with the fraction observed. For the illustrated path, we select the rounded median positive extra-delay bin and then the smallest maximum absolute future shock. The chosen path index is 149 and the selection record is saved. Neither the path nor a simulation seed is selected to match an asset-price or amplification target.

The replication package records the full configuration, fixed seeds, numerical statuses, and funding-call conventions.

A.5.1 Independent checks and sensitivity

A separate implementation recomputes physical transitions, research bills, terminal ownership, and discounted net receipts from 120,000 saved policy histories, containing 764,671 decision records. The iden-

tities reconcile to numerical precision. This establishes internal accounting consistency, not historical accuracy or equilibrium optimality.

We also recheck 52 visited private states across dates 0–17 using independent Python interpolation and adaptive posterior integration, splitting at every continuation-grid crossing. Each state is checked at 742 signal values. The largest marginal indifference residual is 2.98×10^{-12}; the largest profitable-deviation residual after adaptive rechecking is 1.50×10^{-12}. These are sampled incentive checks against the supplied continuation arrays. Independently optimizing the planner's scalar action at 40 visited states changes its reported objective by at most 7.2×10^{-15}. Neither check proves global recursive uniqueness.

A fresh grid refinement. We solve the entire private recursion again on 361 work nodes and 97 ratio nodes, doubling the prior and posterior quadrature orders to 24 and 32 nodes per segment. The planner is also re-solved on that grid with 96 quadrature nodes. No original continuation arrays are used to initialize this calculation. The economic parameters, computational horizon, and initial state are unchanged. The initial net industry claim changes from 3.819983 to 3.818228, and the cutoff changes from 0.176497 to 0.176188. With 2,000 identical unconditional histories, the private completion mean changes by 0.0029 years and the surplus-shortfall percentage by 0.0100 percentage points. The corresponding 2,000-history conditional comparisons appear in Table A9.

The refined full grid still includes unresolved late states. The physical upper bound on entering their work regions, including adjacent interpolation cells, is 3.75×10^{-9}. No visited state in the matched baseline or refined simulations is unresolved. The agreement supports the accuracy of the reported initial-state exercise. One refinement and sampled incentive checks do not certify uniform convergence or bound changes in equilibrium selection.

A shorter decision interval. Let $\Delta = 1/2$ in units of the original period. We set

$$b_\Delta = (1+b)^\Delta - 1, \quad r_\Delta = \Delta r, \quad \delta_\Delta = \Delta\delta, \qquad R_\Delta = R/\Delta,$$

and $\sigma_\Delta = \sigma/\sqrt{\Delta}$, $\nu_\Delta = \nu/\Delta$. Work is measured in units that make the unit flow per capability per calendar year unchanged. These transformations preserve full-research stock growth, depreciation, dis-

Table A9: Numerical refinement and decision-interval sensitivity

	Years per decision	Cutoff / prior s.d.	Withdrawal probability (%)	Extra delay (years) Private	Planner	Core value fall (%)
Original grid	5.85	0.353	31.52	15.91	0.55	40.12
Refined grid	5.85	0.352	31.54	15.91	0.56	40.13
Half interval	2.92	1.389	5.82	15.89	0.61	38.73

Notes: Each row uses its own equilibrium cutoff and compares participation $\Phi(2)$ with $\Phi(-2)$ under 2,000 antithetic histories (1,000 independent pairs). Withdrawal probability means initial participation below ten percent. The first two rows use identical innovations and the same economic model. The third changes decision and information frequency as described below; its news gap is 0.566 prior standard deviations, compared with 0.4 in the first two rows. Values are net of current funding calls. Headline results elsewhere use the larger original simulations.

counting, and physical requirement variance per calendar year, as well as physical signal-error variance at each decision. They do not preserve the information structure over a year: firms decide and receive signals twice as often. Research costs are incurred at different dates and the within-period capability path changes. This is an economic sensitivity exercise, not a numerical refinement of the same game. The horizon is 80 shorter periods and the work grid is $[0, 60]$, with 241 work nodes and 65 ratio nodes; no parameter is refitted.

At the same participation endpoints, the extra private delay is 15.89 years and the core-value fall is 38.73 percent. Initial withdrawal probability is 5.82 percent, compared with 31.52 percent in the original model. The unconditional private completion mean is 36.37 years and the surplus shortfall is 4.78 percent. Holding the news gap instead at 0.4 prior standard deviations moves participation from 92.1 to 7.9 percent; the additional private delay is 12.97 years and core value falls 33.15 percent. All visited states in these shorter-interval simulations pass the recorded current-game checks. The finite Bayesian signal approximation is used at 22.2 percent of unconditional visited decisions, however, so this sensitivity is not certified as a converged continuous-signal solution. The comparison distinguishes robustness of a large conditional loss from sensitivity of its probability and aggregate consequences to the decision and information interval.

Initial work and event frequency. Changing initial work is a state sensitivity, so the existing continuation solution applies. With all parameters unchanged, moving R from 8.0 to 9.0 raises the probability of initial participation below ten percent from 2.13 to 56.56 percent. The core-value fall across each state's

own $c \pm 2\nu$ comparison remains between 39.31 and 42.17 percent. Thus a large conditional response is compatible with very different withdrawal frequencies. The maintained initial work requirement should not be interpreted as an estimate of a historical crash hazard.

A.5.2 Sensitivity to the depreciation rate

The depreciation rate is not estimated. The calibrated $\delta = 0.3$ per period corresponds to about 5 percent a year. Table A10 solves the private and planner problems again at $\delta = 0.15$, half the calibrated rate, and repeats the headline experiments with the same seeds. The core share is unchanged, because its fit to the construction contraction does not involve δ. Lower depreciation also raises the growth of capabilities under continued research, from 1.11 to 1.29 per period, and so shortens the horizon at the calibrated initial work to 26.8 years. The second column holds initial work fixed; the third raises it to 13.87 so that the mean private horizon again equals 41.2 years, the scale for which initial work was chosen.

Both comparisons show that depreciation governs the persistence of a crash. News at the withdrawal boundary worse by 4ν postpones private commercialization by 5.5 and 5.3 years at the lower rate, against 16.0 years in the calibration, and a later clustered withdrawal follows in 7 and 3 percent of adverse-news histories, against 45 percent. A pause that destroys fewer capabilities is cheaper to reverse, so the industry resumes instead of clustering out again. How often withdrawal occurs depends on where the boundary lies. At the calibrated initial work, the faster-growing industry is close enough to completion that the boundary lies 9.4 prior standard deviations above the mean, and no plausible news triggers withdrawal. At the matched horizon, the boundary lies 1.3 prior standard deviations below the mean: initial withdrawal occurs with probability 87 percent, against 32 percent in the calibration, and the average surplus shortfall is 35.3 percent rather than 21.4 percent. Lower depreciation therefore replaces rare, persistent crashes with frequent, short pauses rather than removing the inefficiency.

Turning depreciation off entirely moves the boundary at the calibrated initial work to about 22 prior standard deviations above the mean. At $\delta = 0$, however, many states visited by the paired and horizon-matched experiments lack a verified equilibrium, so those experiments are not reported. The solutions at $\delta = 0.15$ leave a few unresolved decisions on simulated paths, reported in the table; the calibrated

solution leaves none. Rates above $\log(1+b) \approx 0.41$ per period would make even continuously researched capabilities shrink and are outside this comparison.

Table A10: Sensitivity to the capability depreciation rate

	Calibration	Half the depreciation rate	
		Same initial work	Same mean horizon
Depreciation per period, δ	0.30	0.15	0.15
Annual depreciation rate (percent)	5.0	2.5	2.5
Initial work, R_0	8.71	8.71	13.87
Mean private horizon (years)	41.2	26.8	41.3
Withdrawal boundary (prior s.d.)	0.35	9.43	−1.28
Initial withdrawal probability (percent)	31.6	0.0	87.5
News at the boundary, worse by $4v$			
Private delay (years)	15.96	5.55	5.30
Planner delay (years)	0.56	0.32	0.37
Amplification ratio, M	28.6	17.3	14.4
Later clustered withdrawal (percent)	45.2	7.4	3.3
Private surplus shortfall (percent)	21.4	6.7	35.3
Unresolved decisions on simulated paths	0	8	5

Notes: Each column solves the private and planner problems at the stated δ with all other parameters at their calibrated values; the core share does not depend on δ. The third column raises initial work so that the mean private horizon over the 20,000 independent histories equals the calibration's. News rows use the 10,000 paired histories at $\theta = c \mp 2v$, where c is each economy's own initial cutoff. The withdrawal boundary is c in prior standard deviations. Later clustered withdrawal is the percent of adverse-news histories with a later decision at which costly participation is below one-half. The surplus shortfall is relative to the planner over the independent histories. Unresolved decisions count paired and independent histories together. Generated by `revision_work/sensitivity/delta_sensitivity.py`.