Do Markets Discipline the Adoption of Dangerous Technology?*

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

September 2026

Abstract

We study technology adoption when disaster risk rises with the technology's capability and deployment. A survivable disaster that imposes the same proportional loss on adopting and non-adopting firms leaves relative payoffs unchanged while raising marginal utility. Disasters are more likely in high-capability states, so they tilt valuation toward states in which adopters are relatively productive. This raises adopter values and encourages adoption, which further raises disaster risk. Equilibrium deployment therefore exceeds the level that would prevail if the technology posed no disaster risk. Adopter-only disaster losses give adopters lower payouts in high-marginal-utility states and can significantly discipline adoption. A disaster that causes extinction eliminates financial claims in some high-capability states. This lowers the value of adoption and hence deployment, but only modestly, as the effect operates through physical, rather than risk-neutral, probabilities. Empirically, option-implied tail risk now moves with stock prices in a manner broadly consistent with a potential AI disaster.

Keywords: artificial intelligence, technology adoption, disaster risk, state prices, option-implied tails.
JEL codes: G12, O33, D62.

1 Introduction

The adoption of new technologies can increase productivity while also creating new dangers for society as a whole. Individual firms may therefore decide to adopt a new technology even when this adoption is not socially optimal, because they do not fully internalize the dangers their adoption creates (Acemoglu and Lensman, 2024; Jones, 2024). We show that financial markets can amplify this incentive rather than restrain it. In particular, if AI's productivity and its potential for causing a disaster are both increasing in its capability, then disaster states are more likely to be those where AI-adopting firms are relatively more productive. The high marginal utility of these states increases the value of adoption, potentially leading to more deployment than the case in which AI poses no disaster risk at all.

The nature of the disaster determines the direction and magnitude of this effect. If the disaster is survivable and poses equal proportional losses on both adopters and non-adopters, then the disaster creates no direct cash-flow loss for adopters and the valuation tilt associated with having

*Generative AI was used to produce this paper. All work was verified by the authors and any errors are our own. Correspondence can be directed to rready@uoregon.edu.

†Lundquist College of Business, University of Oregon.

more capable AI in high-marginal-utility disaster states unambiguously increases adopter values and overall adoption. If the disaster's damage is confined to adopters, then a direct relative cash-flow loss competes with the higher capability. In this case the sign is theoretically ambiguous, but for our illustrative calibrations we find that markets provide a substantial amount of discipline as adopters now have a lower relative payoff in high-marginal-utility states. This then lowers adopter values and discourages adoption. We also study an extinction-level disaster that removes all financial claims. In this case, adopters do not benefit from those high-capability states that end in disaster as they are effectively removed from the pricing distribution. This again lowers adopter values and disciplines adoption, but the effect is modest as the extinction adjustment uses physical survival probabilities rather than the much higher risk-neutral probabilities associated with a survivable disaster. Financial markets therefore provide little discipline against adopting a technology with the potential to cause extinction.

We first establish the mechanism in a fixed-date adoption model built on Pástor and Veronesi (2009), with individual firm adoption as in Gutierrez and Ready (2026). The simple setting lets us show analytically that a disaster imposing the same proportional loss on all firms encourages private adoption. We then allow firms to choose when to adopt as they learn about the technology, with consumption paid throughout. The numerical results show the same pattern: common survivable losses accelerate adoption, while adopter-only damage and extinction restrain it. Finally, we examine whether favorable technology-stock returns coincide with increases in both upside and downside option-implied tail probabilities, as the model suggests. The recent co-movement provides suggestive evidence of the mechanism.

The paper builds on the asset pricing of technological revolutions. Gutierrez and Ready (2026) show that when adoption changes only systematic risk, the private and planner conditions coincide because that risk is capitalized in firm values. We introduce a deployment externality and identify a state-price selection effect that increases the private return to the risky action. Acemoglu and Lensman (2024) study learning, disaster risk, and the regulation of transformative technologies; their deployment externality is the basis for our planner wedge, and our contribution is the private hedging channel. Jones (2024) studies the trade-off between growth and existential risk, while Chow et al. (2026) distinguish extinction from survivable aggregate disasters. That distinction matters here because the relative productivity payoff survives only in the latter case.

Our pricing logic connects to the climate-finance literature. In Pástor et al. (2021), green assets hedge climate risk, lowering their cost of capital and encouraging green investment. Baker et al. (2022) show that hedging demand can instead favor polluters, increasing investment in activities that generate environmental damage. We find a similar amplification when the damage directly reduces dividends and aggregate consumption, including adopters' dividends, rather than entering utility separately from financial wealth. A common proportional loss leaves relative payoffs unchanged, but capability-dependent disaster arrival shifts valuation toward states in which adoption is most productive. This separates the deployment externality from the private pricing incentive and reveals how that incentive changes under adopter-only losses or extinction. A different channel appears in

Baker et al. (2020), where disagreement and private insurance reduce support for disaster-prevention policy; our mechanism operates under common beliefs.

Related models connect innovation to displacement risk and asset returns (Gârleanu et al., 2012; Kogan et al., 2020). Chen (2026) studies AI stocks as hedges against consumption displacement and shows how incomplete insurance can lead households to block socially beneficial AI development. In our setting, hedging demand instead encourages deployment that raises disaster risk, even though disasters also harm adopters.

A growing literature uses asset prices and investment to assess expectations about AI (Eisfeldt et al., 2023; Andrews and Farboodi, 2025; Wachter and Wachter, 2026). Related work explains how AI can lower yields through stronger fiscal backing of government debt (Kung et al., 2026) or labor displacement with limited capital-market participation (Maresca, 2026). Shen (2026) examines the pricing of AI-risk news factors. Our question concerns how the pricing of AI-linked danger feeds back into technology deployment. Our empirical analysis examines the co-movement of technology-stock returns and option-implied tail probabilities.

Section 2 presents the simple model, incidence comparisons, and risk-neutral distributions. Section 3 develops the dynamic model and simulations. Section 4 examines technology stocks and tail risks in the data.

2 Simple Model

We begin with an illustrative model in which firms adopt at an exogenous review date, following Pástor and Veronesi (2009). As in Gutierrez and Ready (2026), individual firms can adopt at different costs, allowing partial adoption rather than an economy-wide switch. Fixing the adoption date isolates the private adoption incentive and the aggregate deployment externality. Section 3 relaxes this assumption, allowing firms to choose when to adopt as they learn about the technology.

2.1 Environment and adoption

A representative household owns a continuum of firms and consumes terminal wealth at T. Firms learn about a fixed but unknown technology quality ψ from a small pilot sector, then make an irreversible conversion choice at the review date t^*. Before T, output is reinvested. The pilot reveals a noisy public signal without contributing materially to aggregate wealth. Let $\mathcal{F}_t$ denote public information available at date t. The posterior is Gaussian, with mean $\widehat{\psi} = \widehat{\psi}_{t^*}$ and variance $\widehat{s}^2 = \widehat{s}^2_{t^*}$: $\psi \mid \mathcal{F}_{t^*} \sim \mathcal{N}(\widehat{\psi}, \widehat{s}^2)$. The mean responds to news and the variance falls as information accumulates. Appendix B gives the productivity processes and filtering equations, following Pástor and Veronesi (2009).

Let $\Delta = T - t^*$ be the remaining horizon after review. Conditional on quality, conversion multiplies terminal capital by $e^{A_2(\Delta)\psi}$ relative to remaining with the old technology. With productivity mean-reversion speed ϕ, the loading is

$$A_2(\Delta) = \Delta - \frac{1 - e^{-\phi\Delta}}{\phi}. \tag{1}$$

This loading reflects a gain that builds as productivity moves toward its higher long-run mean. We condition on current capital and productivity and normalize the common old-technology terminal payoff to one.

We decentralize adoption as in Gutierrez and Ready (2026). Firms have uniformly distributed fixed types $u \in [0, 1]$. Let κ be the midpoint conversion cost and ϵ its half-range, with $0 < \kappa - \epsilon < \kappa + \epsilon < 1$. Proportional conversion costs are

$$\kappa_\epsilon(u) = \kappa + \epsilon(2u - 1). \tag{2}$$

Conversion destroys that fraction of capital. Low-cost firms adopt first. If the adopted fraction is q, the marginal firm has type $u = q$. Total cost-adjusted adopted capital is $K(q)$, aggregate terminal wealth is $H(q, \psi)$, and the payoff gain from converting the marginal firm is $H_q(q, \psi) = \partial H(q, \psi)/\partial q$:

$$K(q) = \int_0^q [1 - \kappa_\epsilon(u)]\, du, \tag{3}$$

$$H(q, \psi) = 1 - q + K(q)e^{A_2(\Delta)\psi}, \tag{4}$$

$$H_q(q, \psi) = [1 - \kappa_\epsilon(q)]e^{A_2(\Delta)\psi} - 1. \tag{5}$$

The adopter gives up an old-technology payoff of one and receives the uncertain technology payoff, reduced by its conversion cost.

The household has constant relative risk aversion (CRRA) utility $U(C) = C^{1-\gamma}/(1 - \gamma)$ over consumption C, with risk-aversion coefficient $\gamma > 1$. Expectations $\mathbb{E}$ are conditional on the review-date posterior; we suppress its fixed variance in the arguments below. A dollar paid when consumption is low has greater market value because marginal utility is high. Without technological disasters, those state-price weights are proportional to $H^{-\gamma}$. Write G^0 for its marginal market-value gain in the no-disaster economy (superscript 0). The marginal firm is indifferent when

$$G^0(q, \widehat{\psi}) = \mathbb{E}[H(q, \psi)^{-\gamma} H_q(q, \psi)] = 0. \tag{6}$$

This is a market-value comparison, not a comparison of average physical payoffs. The planner has the same marginal condition: the effect of adoption on ordinary consumption risk is already priced. Disaster risk will break that equivalence by making aggregate deployment affect an outcome that each individual firm takes as given.

2.2 Common disaster risk and private adoption

We now add our central departure. Let N denote the number of technological disasters between t^* and T. Let $\lambda(\psi) > 0$ denote the annual disaster intensity at full deployment, with $\lambda'(\psi) > 0$, and let $\Lambda(q, \psi)$ denote the integrated intensity conditional on deployment and capability. After the adoption decision,

$$N \mid (q,\psi) \sim \text{Poisson}(\Lambda(q,\psi)),$$
$$\Lambda(q,\psi) = \underbrace{q^2}_{\text{deployment}} \underbrace{\lambda(\psi)}_{\text{capability}} \Delta, \qquad \lambda'(\psi) > 0. \tag{7}$$

Both capability and use increase danger. The first makes a system more able to cause harm; the second puts more activity at risk. The quadratic deployment term makes the current hazard and its marginal increase zero at zero adoption.[1] Risk then rises as adoption spreads. An individual firm takes q and hence $\Lambda(q,\psi)$ as given when it adopts. The planner instead recognizes that increasing q raises the common hazard. This is the new externality.

We consider three disaster structures: a common survivable disaster, an AI-sector survivable disaster, and extinction. The three cases differ in which cash flows survive the event.

Assume the valuation expectations below are finite. Suppose every disaster multiplies every firm's terminal payoff by the same survival factor $\omega \in (0,1)$. Let the superscript C label this common-disaster economy. Aggregate terminal consumption is

$$C_T^C = H(q,\psi)\omega^N. \tag{8}$$

Adopters are not physically safer than non-adopters. Let χ measure how a single disaster changes marginal-utility-weighted consumption, and define $R_D(q,\psi)$ as the expected state-price multiplier generated by the Poisson disaster process:

$$\chi = \omega^{1-\gamma} - 1 > 0, \qquad R_D(q,\psi) \equiv \mathbb{E}\Big[\omega^{(1-\gamma)N} \mid q,\psi\Big] = \exp\{\Lambda(q,\psi)\chi\}. \tag{9}$$

The multiplier combines a lower physical payoff with higher marginal utility. With $\gamma > 1$, the marginal-utility increase dominates, so $R_D > 1$. Higher capability increases this weight by making disasters more likely. This does not mean every high-productivity state has low consumption: it means common disasters shift valuation, relative to the no-disaster comparison, toward the high-capability realizations in which adopters do relatively well. After integrating over the number of disasters, write G_D^C for the decentralized marginal adoption gain; the subscript D denotes the decentralized allocation and P below denotes the planner. The adoption condition becomes

$$G_D^C(q,\widehat{\psi}) = \mathbb{E}\big[R_D(q,\psi)H(q,\psi)^{-\gamma}H_q(q,\psi)\big] = 0. \tag{10}$$

The loss cancels from the adopter–non-adopter payoff ratio, but a low-consumption payment still receives more weight in the market-value comparison. A disaster is a high-marginal-utility state. Because $\lambda'(\psi) > 0$, disaster states disproportionately select high-ψ realizations, precisely the states in which adoption is most productive.

[1] This choice aligns the market and planner fixed-date adoption thresholds at $q = 0$. The qualitative private hedging result does not require the quadratic form: at positive deployment, the covariance argument applies to $g(q)\lambda(\psi)$ whenever $g(q) > 0$.

Proposition 1 (Common disaster risk lowers the private threshold, more so when risk is greater). *Fix the deployed fraction of firms at $q > 0$. Technology quality is ψ, and its posterior mean $\widehat{\psi}$ summarizes current beliefs. Let $\lambda(\psi)$ be disaster intensity at full deployment and $\omega \in (0, 1)$ the fraction of each firm's payoff retained after a common disaster. Suppose risk aversion exceeds one ($\gamma > 1$), intensity rises with quality ($\lambda'(\psi) > 0$), posterior uncertainty is nondegenerate, and valuation expectations are finite. In each economy compared, assume the marginal private adoption gain crosses zero once from below as $\widehat{\psi}$ rises.*

The following conclusions hold:

(i) The belief threshold for private adoption is strictly lower with common survivable disaster risk than without it.

(ii) The threshold falls further when the common loss becomes more severe (lower ω) or intensity increases proportionally, holding other primitives fixed.

(iii) At zero deployment, $q = 0$, these fixed-date thresholds coincide.

Proof sketch: the covariance effect. At fixed deployment $q > 0$, recall three payoff objects: H is aggregate wealth without a disaster, $H_q = \partial H/\partial q$ is the payoff gain from marginal adoption, and R_D is the disaster-induced multiplier on valuation weights. Risk aversion is γ, so the no-disaster marginal-utility weight is $H^{-\gamma}$. Evaluate these objects at the no-disaster adoption threshold. Let P be the posterior probability measure and define its marginal-utility-weighted counterpart Q_0 by

$$\frac{dQ_0}{dP} = \frac{H^{-\gamma}}{\mathbb{E}[H^{-\gamma}]}, \qquad \mathbb{E}^{Q_0}[H_q] = 0,$$

Here $\mathbb{E}^{Q_0}$ denotes expectation under Q_0. The private marginal gain under common disasters, G_D^C, then satisfies

$$\frac{G_D^C}{\mathbb{E}[H^{-\gamma}]} = \mathbb{E}^{Q_0}[R_D H_q] = \mathrm{Cov}^{Q_0}(R_D, H_q) > 0. \tag{11}$$

The covariance is computed under Q_0. Both the disaster valuation multiplier R_D and the marginal adoption payoff H_q increase strictly with capability ψ. Disaster risk therefore places more valuation weight on the realizations in which adoption pays most. The firm prefers adoption at the old threshold, so single crossing implies a lower threshold. Greater severity or a proportional intensity increase adds another increasing tilt to the valuation weights; the same covariance argument gives the further threshold reduction. Appendix A provides the full proof and comparative statics.

Remark (capability versus deployment). If intensity is capability independent, write its constant full-deployment value as $\bar{\lambda}$. Then $R_D(q, \psi) = \exp\{q^2 \bar{\lambda} \Delta\chi\}$ factors out of the private condition, leaving the no-disaster threshold unchanged. A capability-independent additive increase in intensity likewise cancels. The planner wedge remains because deployment raises the hazard. The private threshold result requires capability-dependent risk.

The restriction $\lambda'(\psi) > 0$ captures the idea that more capable systems are both more useful and able to cause harm on a larger scale. Capability may instead arrive together with better monitoring, alignment, or control. If those improvements are strong enough to make disaster risk independent of capability, the private hedging channel disappears; if they make risk decline with capability, the channel reverses. The deployment externality remains whenever aggregate use raises the hazard.

2.3 Incidence and the planner

The planner sees the same state-price selection effect but also internalizes the dependence of disaster intensity on aggregate deployment. Recall that H is aggregate terminal wealth, H_q is the marginal conversion payoff, and R_D is the disaster valuation multiplier. Write V_P^C for the planner's expected utility in the common-loss economy:

$$V_P^C(q) = \frac{1}{1-\gamma}\mathbb{E}\left[H(q,\psi)^{1-\gamma} R_D(q,\psi)\right]. \tag{12}$$

With $\chi = \omega^{1-\gamma} - 1$ as defined above, the planner marginal gain $G_P^C = \partial V_P^C / \partial q$ satisfies

$$G_P^C(q,\widehat{\psi}) = \mathbb{E}\left[R_D H^{-\gamma}\left(H_q - H\frac{2q\lambda(\psi)\Delta\chi}{\gamma-1}\right)\right] = 0. \tag{13}$$

The first term is exactly the decentralized condition. The second is the marginal social cost of the additional disaster hazard created by deployment. Hence the same disaster risk that makes private adoption more attractive makes the planner more reluctant to deploy.

This deployment externality is the planner mechanism of Acemoglu and Lensman (2024). Proposition 1 adds the state-price selection result.

Proposition 2 (The planner always values marginal adoption less)**.** *Let G_P^C and G_D^C be the planner and private marginal gains under a common loss. For every adopted share $q > 0$ with $\lambda(\psi) > 0$, $\omega \in (0,1)$, and $\gamma > 1$,*

$$G_P^C(q,\widehat{\psi}) < G_D^C(q,\widehat{\psi}). \tag{14}$$

At any interior decentralized allocation satisfying $G_D^C = 0$, the planner strictly prefers less adoption. If each marginal condition crosses zero once from above as q rises, then $q_P < q_D$, where q_P and q_D are planner and decentralized adoption, respectively.

Proofs are in Appendix A.

Adopter-only losses. Now put the loss on adopters alone. Their productivity advantage is no longer protected in relative terms: a disaster damages the adopted claim while leaving the old claim intact. State-price selection still favors capable technology, but it must overcome this direct loss. A lower private threshold is therefore possible, not guaranteed.

Suppose a survivable disaster damages only capital operated with AI. We use the superscript A for this AI-sector disaster economy. Recall that N counts disasters, ω is the fraction of affected payoffs retained per event, and $K(q)$ is cost-adjusted adopted capital. Conditional on $N = n$, let C_n^A denote aggregate consumption and $H_{q,n}^A$ the payoff gain from marginal adoption:

$$C_n^A = 1 - q + K(q) e^{A_2(\Delta)\psi} \omega^n, \tag{15}$$

$$H_{q,n}^A = [1 - \kappa_\epsilon(q)] e^{A_2(\Delta)\psi} \omega^n - 1. \tag{16}$$

Let G_D^A denote the private marginal gain in this adopter-only economy. Its adoption condition is

$$G_D^A(q, \widehat{\psi}) = \mathbb{E}_{\psi,N}\left[(C_N^A)^{-\gamma} H_{q,N}^A\right] = 0. \tag{17}$$

Disaster selection still raises the state-price weight on high-ψ states, but the adopting firm now suffers direct damage relative to a non-adopter. Let Q_0 denote the no-disaster measure with density proportional to $H^{-\gamma}$, and let Q_A be the joint measure over (ψ, N) proportional to $f_P(\psi)\Pr(N = n \mid q, \psi)(C_n^A)^{-\gamma}$, where f_P is the review-date posterior density of quality. At the no-disaster threshold, AI-sector disaster risk encourages private adoption if and only if

$$\mathbb{E}^{Q_A}\left[e^{A_2(\Delta)\psi} \omega^N\right] > \mathbb{E}^{Q_0}\left[e^{A_2(\Delta)\psi}\right]. \tag{18}$$

The sign is therefore ambiguous: the disaster-state selection gain must exceed the direct physical damage to adopters.

The planner also accounts for the extra damage that deployment imposes on existing adopters. If an additional event lowers consumption, this term makes the planner value marginal adoption less than the firm. The derivative is given in Appendix C.

Extinction. Extinction changes which payouts exist, rather than their relative size after a survivable event. Both technologies pay zero after termination, so a firm can earn its technology advantage only in surviving states. If survival is less likely at high capability, precisely those valuable adopter payouts are less likely to be received.

Use the superscript E for the extinction economy. The financial-claim mechanism follows the extinction logic in Chow et al. (2026): extinction ends consumption and the payoffs of long-lived financial claims. Let the indicator S equal one if civilization survives and zero if extinction occurs, and define the survival probability $s(q, \psi)$ by

$$s(q, \psi) = \Pr(S = 1 \mid q, \psi) = e^{-\Lambda(q,\psi)}. \tag{19}$$

For welfare, we assign a separate reservation-consumption level $\underline{C} > 0$ after extinction. It keeps CRRA utility finite and is not corporate property. All firm claims pay zero after termination. Appendix C states these payoffs explicitly.

Recall that H is surviving aggregate wealth and H_q is the marginal adoption payoff. Since both alternatives pay zero after extinction, the private marginal gain G_D^E weights that payoff by survival probability:

$$G_D^E(q,\widehat{\psi}) = \mathbb{E}\big[s(q,\psi)H(q,\psi)^{-\gamma}H_q(q,\psi)\big] = 0. \tag{20}$$

Survival probability falls with ψ, while the payoff advantage of adoption rises with ψ. Extinction therefore removes precisely the high-ψ states in which adoption is most valuable, raising the private adoption threshold relative to the no-disaster benchmark.

The planner additionally accounts for the loss of survival probability caused by deployment. This is a negative marginal effect when surviving utility exceeds reservation utility.

Extinction works through the survival probability, rather than an adopter payoff multiplier. It is not the $\omega \to 0$ limit of a survivable disaster with a claim that continues to pay.

Figure 1 summarizes the central result by comparing market and planner adoption thresholds with the no-disaster benchmark. In the common-loss panel, disaster risk lowers the market threshold while the planner internalizes the additional hazard and raises its threshold. Adopter-only damage instead raises the market threshold in this calibration. In the extinction panel, the market threshold lies only slightly above the no-disaster benchmark. Appendix D derives the corrective deployment charge that closes the market–planner gap.

Disaster states are rare at the reference deployment level, helping explain the modest private extinction effect. Conditional on capability, extinction enters through the physical survival probability $e^{-\Lambda(q,\psi)}$. No financial claim pays in extinction, so there is no marginal-utility amplification of extinction-state payouts. Common survivable losses, by contrast, receive the high state-price weight associated with low consumption. Surviving payoffs still carry $H^{-\gamma}$; it is the survival factor that is physical rather than risk-neutral. A small private response therefore need not imply a small welfare cost.

2.4 Risk-neutral payoff distributions

The adoption result concerns the relative value of two technologies. Options concern the whole distribution of a claim's payoff. We connect the two by comparing a low-cost firm, a high-cost firm, and an aggregate consumption claim at two review-date beliefs. Deployment is zero at the lower belief and 30% at the higher belief. The low-cost firm adopts only in the higher state; the high-cost firm remains a non-adopter in both. The adoption decisions are then held fixed while subsequent payoffs are priced.

Figure 2 shows risk-neutral densities, meaning physical probabilities weighted by the household's marginal utility and normalized. Denote this normalized risk-neutral measure by Q; every claim uses it, and $\mathbb{E}^Q$ denotes expectation under it. A low-consumption outcome receives more weight, even if its physical probability is small. The illustration restores continuous old-economy risk, which cancels from the adoption comparison but makes non-adopter payoffs continuous as well. Appendix G gives the payoff equations, pricing weights, and distinct fixed-date calibration.

Both tails broaden for all three claims as deployment rises. The adopter gains uncertainty about the technology payoff as well as disaster exposure. Those risks are related: common disasters are

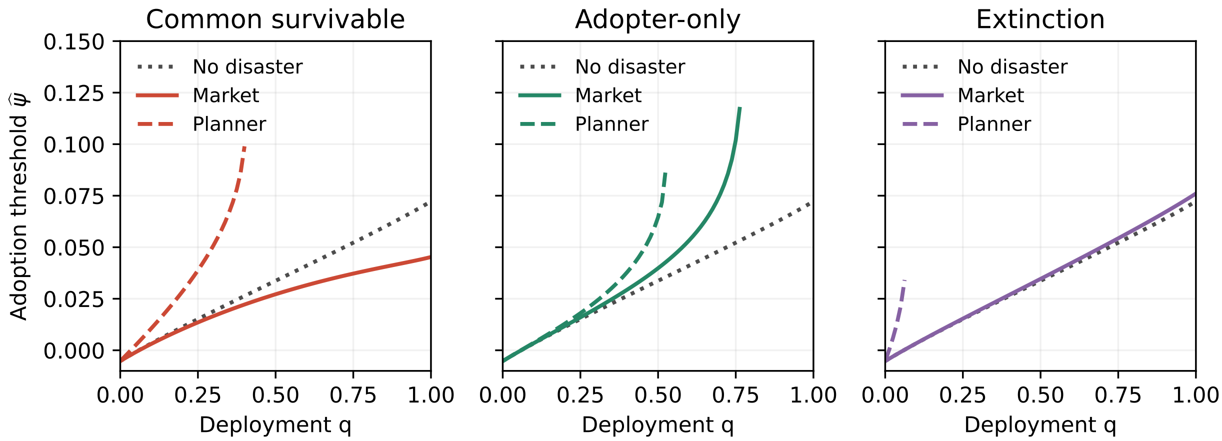

Figure 1: Fixed-date marginal adoption thresholds. The panels show common survivable losses, adopter-only losses, and extinction, each with the dotted no-disaster benchmark. Solid curves are market thresholds and dashed curves are planner thresholds. The hazard schedule is held fixed across cases. Survivable losses are 50%. Extinction reservation consumption is 0.30 in units of the normalized old-technology terminal payoff; it affects welfare only. Curves report the first upward zero crossing of marginal adoption gains and end where none exists in the searched belief range $[-0.08, 0.15]$.

more likely when the technology is more productive. Its productivity advantage can therefore partly offset the common loss. The non-adopter receives no such gain and remains exposed to the common damage.

That difference does not imply that only the non-adopter's left tail expands. Each density is centered at its own risk-neutral mean of log payoff. Fewer disasters than anticipated in that mean can give a non-adopter a relatively high outcome without any technology gain. Both of its centered tails can therefore widen. Nor can the adopter's productivity and disaster variances simply be added: the correlation between them matters.

The lower row decomposes the higher-belief density into probability-weighted no-disaster and disaster components. Their sum is the red curve in the upper row. The physical probability of at least one disaster is 2.98%, compared with a risk-neutral probability of 31.00%. State prices emphasize the low-consumption disaster outcomes. The residual peaks in the non-adopter's disaster component reflect successive proportional losses blurred by continuous old-economy risk; they are not discrete payoff probabilities.

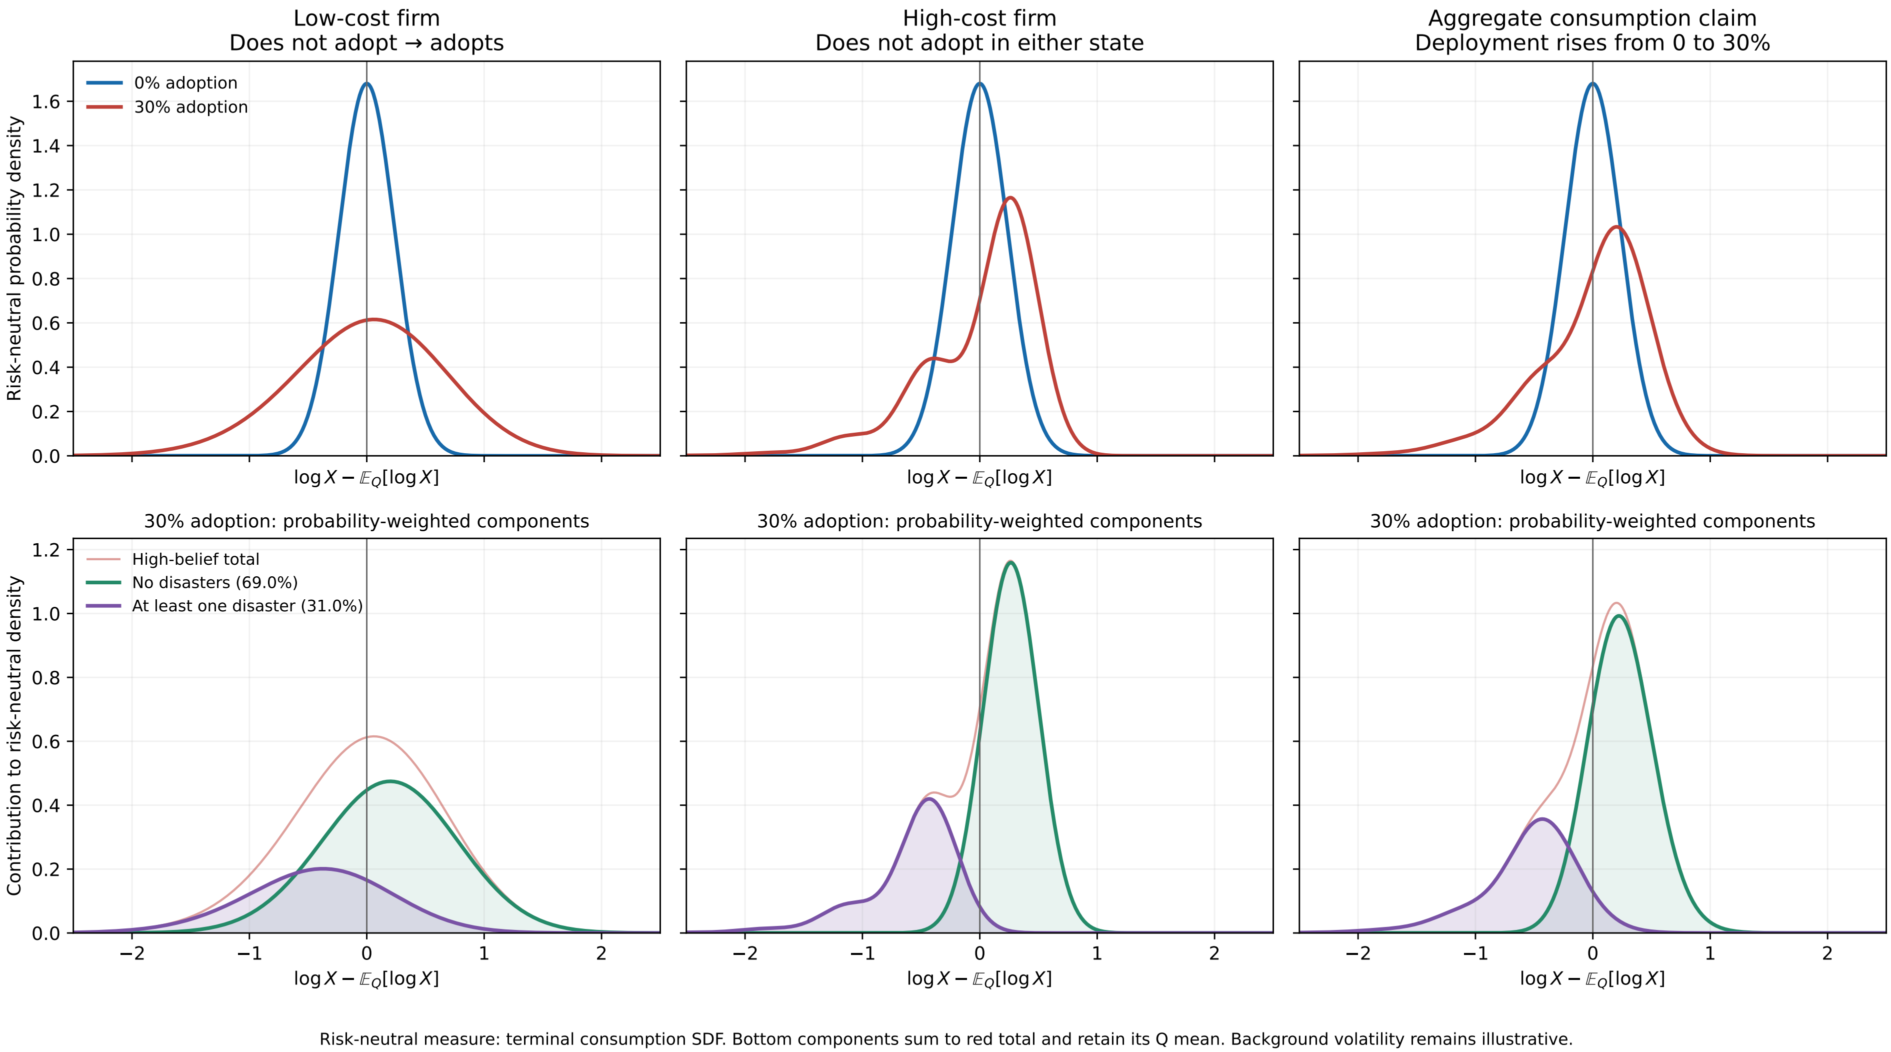

Figure 2: **Risk-neutral payoff densities as adoption rises from zero to** 30%**.** Columns show the low-cost firm, high-cost firm, and aggregate consumption claim. Each top-row density is centered at its own Q mean of log payoff. The bottom row decomposes the higher-belief density into probability-weighted no-disaster and disaster components. The components share the total density's center and integrate to their event probabilities. They sum to the red curve. The horizon is 22 years after review.

Appendix Table 4 provides a second comparison using each claim's terminal payoff X_i and its own forward payoff $F_i = \mathbb{E}^Q[X_i]$, where i indexes the three claims. Both $Q(X_i < 0.8F_i)$ and $Q(X_i > 1.2F_i)$ increase. These forward-relative tails use a different center from the plotted log densities. Their asymmetry does not quantitatively reproduce the stronger IWM downside exposure found in some empirical windows. The model comparison is also long horizon, whereas the empirical options are one year. Its role is to explain qualitative tail expansion rather than fit those option probabilities.

Damage confined to adopters changes both their physical losses and the common pricing measure. Extinction instead puts physical payoff mass at zero and removes high-capability payouts through survival. Their distributions require separate calculations; neither is obtained by relabeling a common-disaster curve. Their adoption conditions are given above.

The appendix also separates expected cash flows from state-price adjustments in stock prices.

3 Dynamic Model and Simulations

The fixed-date model compares conversion with non-adoption. We now let each firm choose when to convert, keeping its own fixed cost and its option to wait for more information. This is a separate dividend-flow implementation of the mechanism. Consumption is the sum of firms' dividends, so the planner and investors value the same resources.

3.1 Dividends, information, and installed adoption

An unconverted firm's dividend flow is D_t^O, where O denotes the old technology. Let g be its log-growth rate, σ_O its volatility, and W_t a standard Brownian motion shared by all firms. Away from disasters, its log growth is $g\,dt + \sigma_O dW_t$. Write y_t for cumulative log productivity relative to the old technology, φ for quality's loading into its growth, and σ_N for frontier volatility. The frontier Brownian motion Z_t is independent of W_t and shared by all adopters. After conversion, a firm also receives the increment

$$dy_t = \varphi\psi\,dt + \sigma_N dZ_t. \tag{21}$$

Quality ψ is fixed but unknown. The growth loading φ belongs to this dividend-flow model; Section 2's ϕ is the mean-reversion speed in the fixed-date model. The pilot reveals the same frontier signal that adopters experience. Let $\pi_t(\psi)$ denote the public posterior density and $\widehat{\psi}_t = \int \psi\pi_t(\psi)d\psi$ its mean. Investors observe productivity news, disaster arrivals, and survival. The pilot signal is independent of deployment; disaster observations reflect the intensity specified below.

As in Section 2, $u \in [0,1]$ ranks firms by conversion cost. We abbreviate $\kappa_\epsilon(u)$ as $\kappa(u) = \kappa + \epsilon(2u-1)$, with midpoint κ and half-range ϵ. At conversion, its dividend scale falls by the fraction $\kappa(u)$; subsequent dividends receive the frontier increment directly. Existing adopters have different adoption dates and accumulated productivity gains. Let q_t be their mass and let a_t be their aggregate dividends divided by the current old-technology dividend per firm. Writing $D_t(u)$ for firm u's dividend, aggregate consumption is

$$C_t = \int_0^1 D_t(u)\,du = D_t^O\big[1 - q_t + a_t\big]. \tag{22}$$

The two components are actual non-adopter and adopter dividends, not a log-linear approximation based on adopted mass. The additional state a_t records the relevant adoption history. A fraction q_t of firms need not supply that same fraction of dividends.

Suppress time subscripts for a conversion decision, and let $r \geq q$ be the adopted mass immediately afterward. A superscript + denotes the post-conversion value before subsequent growth or disasters. For this cost-ordered conversion, define $K(q) = \int_0^q [1 - \kappa(u)]du$. The installed adopter flow becomes

$$a^+ = a + K(r) - K(q), \qquad C^+ - C = -D^O \int_q^r \kappa(u)\, du. \tag{23}$$

New adopters enter at their own cost-reduced scale. Existing adopters' dividends are unchanged by others' conversion. This accounts for the resource loss once, through productive scale and hence dividends and consumption.

3.2 Individual exercise and aggregate adoption

Each firm compares the market value of conversion with old-technology dividends plus its own future option to convert. Both alternatives use marginal utility from (22). A firm takes the aggregate state and deployment path as given, while the mass of exercising firms must equal aggregate adoption. The state includes installed mass q, installed relative dividends a, the public posterior, and remaining time. Conversion type stays fixed when a firm waits.

We solve the marginal firm's indifference condition and update all exercising firms simultaneously. There is no cohort rule, deployment throttle, or capacity constraint. Cost ordering and exercise inequalities are checked numerically. Appendix E gives the stopping recursion, the installed-dividend transition, and numerical diagnostics.

3.3 Disaster incidence, the planner, and the terminal date

Let $\lambda(\psi)$ be annual full-deployment disaster intensity at true capability ψ, with $\lambda'(\psi) > 0$. With adopted mass q_t, the current annual intensity is

$$\lambda_t = q_t^2 \lambda(\psi). \tag{24}$$

It depends on deployed firm mass and capability, rather than the dividend share. Beliefs condition on the joint history of productivity and disaster observations. Appendix E gives the posterior and numerical recursion.

As in the simple model, $\omega \in (0, 1)$ is the fraction of affected dividends retained after an event. A common event multiplies all dividends and consumption by ω. An adopter-only event multiplies installed adopter dividends alone by ω. Recalling that a is adopter dividends divided by an old firm's dividend and q is adopted mass, define the adopter-only consumption multiplier Ω_A by

$$\Omega_A(q, a) = \frac{1 - q + \omega a}{1 - q + a}. \tag{25}$$

The loss is weighted by actual adopter dividends. Repeated losses compound on the affected dividends. Extinction terminates all financial payouts. For welfare only, we retain a reservation-consumption fraction $\underline{c}$ of counterfactual consumption in the termination interval, held constant through T. That utility convention is not surviving corporate property.

The planner maximizes expected utility from the same aggregate dividends and conversion losses. It additionally internalizes how deployment changes disaster risk and future information.

Final flow dividends and consumption occur at $T = 30$. Continuation values then vanish, with no terminal lump sum or liquidation payment.

3.4 Illustrative parameters and the conditional experiment

We use frontier-growth loading $\varphi = 0.3551$, $\sigma_N = 0.07$, prior standard deviation 0.04, risk aversion $\gamma = 4$, and annual utility discount rate $\delta = 0.01$. Old-technology log growth and volatility are each 0.02. Conversion costs have $\kappa = 0.10$ and $\epsilon = 0.025$, so retained scale reflects costs from 7.5% to 12.5%. Survivable losses are 30% and 50%, with an unchanged quadratic deployment hazard across severities. Full-deployment intensity is 0.017 at reference capability $\psi_\star = 0.029742$, capped at 0.10 by the bounded schedule in Appendix E. The historical frequency scale is illustrative, not an estimate of AI disaster probability (Barro and Ursúa, 2008). Extinction reservation consumption is 30% of counterfactual consumption at termination; it affects planner welfare, not financial payouts. The static example uses the same conversion-cost range, risk aversion, and hazard schedule, with a 50% survivable loss matching the more severe dynamic case. It retains a fixed review date and terminal consumption. Its extinction reservation level is 0.30 in units of the normalized old-technology terminal payoff, whereas here the reservation is a fraction of counterfactual consumption at termination.

All displayed paths start with a zero-mean prior, adopted mass $q_0 = 0$, and installed relative adopter dividends $a_0 = 0$. We condition on a favorable realization of a transformative technology, with fixed true quality $\psi = 0.29527$, zero realized frontier shocks, and no realized disaster. The same quality and productivity news are used across cases; firms nevertheless account for uncertain future signals and disasters when choosing adoption.[2] Welfare calculations integrate future uncertainty from the initial prior rather than evaluating utility along this selected path.

3.5 Simulation results

Figure 3 compares incidence and severity under identical news. The top row shows market adoption alongside the no-disaster benchmark. The bottom row reports perceived one-year risk $\int[1 - \exp\{-q_t^2\lambda(\psi)\}]\pi_t(\psi)d\psi$, holding current deployment fixed. It is not cumulative disaster probability.

At year six, common-loss market adoption is 31.55% under a 30% loss and 91.92% under a 50% loss, compared with 29.19% without disasters. Common disaster risk therefore accelerates deployment on this news path, with a stronger response under greater severity. Adopter-only losses yield 26.42% and 24.93% adoption, while extinction yields 27.42%. These incidence comparisons are numerical outcomes; the adopter-only sign depends on direct damage relative to the valuation benefit.

The severe-loss scenario illustrates a potentially sizeable adoption response, rather than an estimate of its empirical magnitude. The experiment does not jointly fit adoption, interest rates, or equity returns. Magnitudes depend on disaster severity, preferences, the hazard schedule, and numerical resolution; Appendix F reports the numerical checks.

[2] At year six the productivity-only posterior mean is about 3.3 prior-predictive standard deviations above its initial expectation. The experiment is conditional on favorable news, not a forecast of typical adoption.

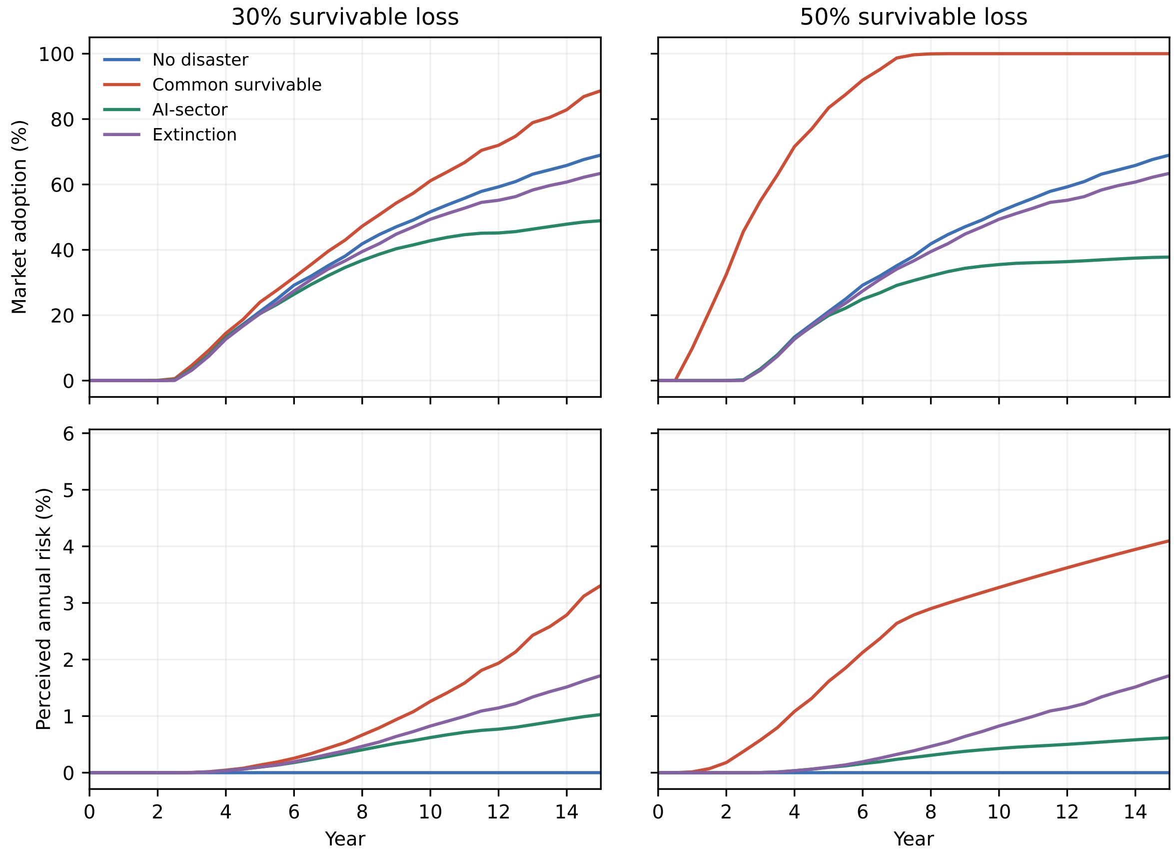

Figure 3: **Adoption and disaster risk with dynamic adoption.** Columns compare 30% and 50% survivable losses under identical favorable productivity news. The top row reports market adoption, including the no-disaster benchmark. The bottom row reports perceived one-year disaster probabilities at current deployment; extinction is conditional on survival.

Current quadratic hazard vanishes at zero adoption. Future deployment nevertheless affects both the adopted claim and the waiting option, so the dynamic onset need not coincide across incidence cases. Exact first-positive dates are sensitive to interpolation and are not a quantitative result of the exercise.

Figure 4 compares market and planner paths for the 30% loss benchmark. Both use the same dividend accounting and information. Appendix F reports stochastic welfare and numerical diagnostics.

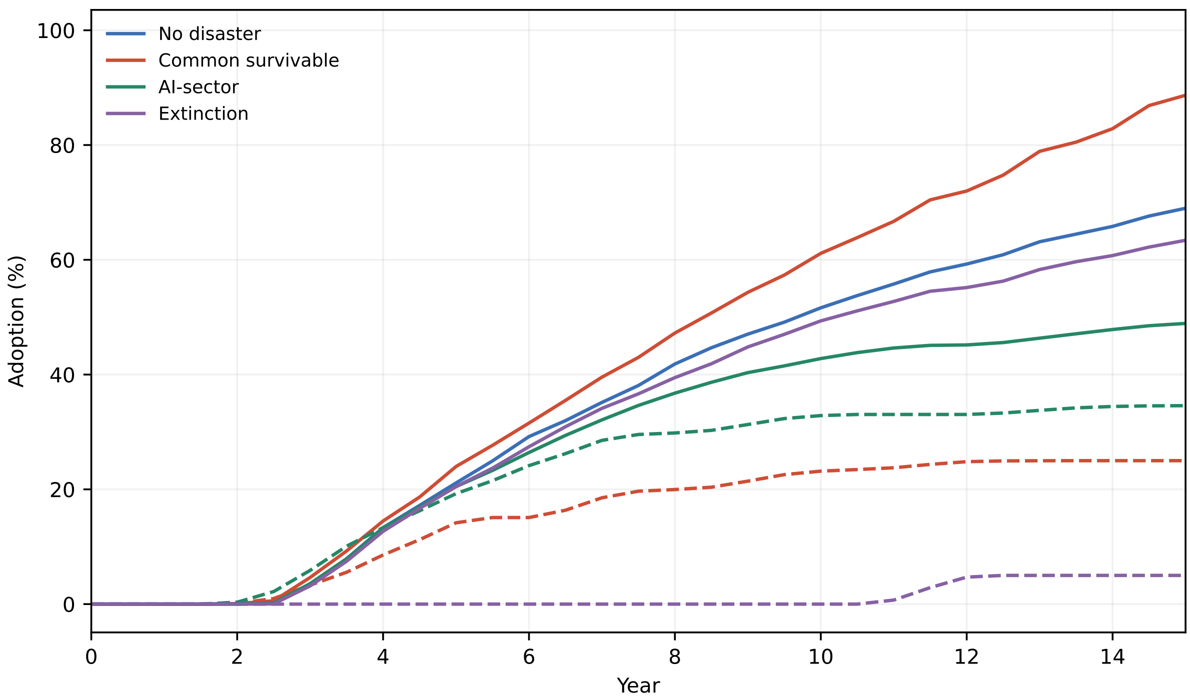

Figure 4: **Market and planner under the same conditional productivity-news path.** Solid curves are market adoption and dashed curves are planner adoption. Consumption aggregates firm dividends exactly in both problems.

4 Technology Stocks and Tail Risks in the Data

The simple model connects technology adoption to broader risk-neutral payoff distributions. We examine one-year option-implied tail probabilities for three equity portfolios represented by exchange-traded funds (ETFs): Invesco QQQ (QQQ), which tracks the technology-heavy Nasdaq-100 Index; the iShares Russell 2000 ETF (IWM), which tracks small-cap U.S. stocks; and the SPDR S&P 500 ETF Trust (SPY), which tracks the large-cap U.S. market. We select these ETFs for their highly liquid options markets: SPY, QQQ, and IWM rank first, second, and third in ETF options trading volume.[3] QQQ represents technology-exposed, relatively low-cost adopters, while IWM represents

[3]The ranking uses average daily contract volume across the Cboe options exchanges in 2026 through August 28, as compiled by Macroption. For the S&P 500 tail probabilities, we use SPX index options and pair them with SPY returns.

relatively high-cost adopters. QQQ returns proxy for changes in the value of technology-intensive firms.

The mapping draws on the greater AI use among large firms documented by Bonney et al. (2026) and the labor-task exposure measure of Eisfeldt et al. (2023). Focusing on firms in the two extreme exposure groups—artificial (A) and human (H)—41.1% of IWM firms are H and 58.9% are A, compared with 25.7% H and 74.3% A for QQQ. IWM therefore has a higher unweighted human-task share among firms in these extreme groups. These shares exclude the middle groups and unclassified firms. The classification concerns labor-task exposure, not other determinants of adoption, such as firm size, or conversion costs (Appendix H).

4.1 Tail probabilities and the rolling specification

Let j index the equity series and t index month-end dates. Write $S_{j,t}$ for the underlying index or ETF price, $F_{j,t}$ for its one-year forward price, and $\tau = 365/365$ years for maturity. Let Q_t denote the risk-neutral probability measure conditional on date-t information. The downside and upside objects are approximately

$$p^{-}_{j,t} = Q_t(S_{j,t+\tau} < 0.8F_{j,t}), \qquad p^{+}_{j,t} = Q_t(S_{j,t+\tau} > 1.2F_{j,t}). \tag{26}$$

We estimate the probabilities from option spreads centered on these thresholds. Write $F = F_{j,t}$ and suppress the index and maturity arguments. If $\mathcal{P}(k)$ and $\mathcal{C}(k)$ denote put and call prices at strike k, and d_t is the discount factor to option expiry, the maturity-specific estimators are

$$\widetilde{p}^{-}_{j,t} = \frac{\mathcal{P}(0.85F) - \mathcal{P}(0.75F)}{d_t\, 0.10F}, \tag{27}$$

$$\widetilde{p}^{+}_{j,t} = \frac{\mathcal{C}(1.15F) - \mathcal{C}(1.25F)}{d_t\, 0.10F}. \tag{28}$$

These are strike-band averages of digital probabilities. Available expirations are interpolated to one year; month-end observations are selected separately by wing and daily returns compounded monthly. Changes and returns are in percentage points, so a coefficient measures the probability change per percentage-point QQQ return, controlling for SPY.

For index j and tail $s \in \{-,+\}$ (downside or upside), let w be a window of month-end observations. Let r^{QQQ}_t and r^{SPY}_t denote monthly QQQ and SPY returns. In the regression, $\Delta\widetilde{p}^{s}_{j,t}$ is the monthly change in the estimated probability, $\alpha^{s}_{j,w}$ is the intercept, $\beta^{s}_{j,w}$ and $\theta^{s}_{j,w}$ are the two return slopes, and $\varepsilon^{s}_{j,t}$ is the residual. Each window estimates

$$\Delta\widetilde{p}^{s}_{j,t} = \alpha^{s}_{j,w} + \beta^{s}_{j,w} r^{QQQ}_t + \theta^{s}_{j,w} r^{SPY}_t + \varepsilon^{s}_{j,t}, \qquad t \in w. \tag{29}$$

Each window uses 13 consecutive months with the same intercept and QQQ/SPY return controls. Figure 5 plots $\beta^{s}_{j,w}$ and pointwise normal 95% Newey–West intervals with three lags. QQQ's coefficient includes an own-return relation. SPY-labelled probabilities come from SPX options.

The history combines OptionMetrics and a Databento extension with the established vendor adjustments. Appendix H describes the splice and its measurement limitations. The displayed windows end from February 2023 through August 2026. The dashed Ramp adoption series uses a separate right-hand axis as context; it is neither a regressor nor a treatment.

4.2 Findings and interpretation

Appendix Table 4 provides the model counterpart to these empirical tail probabilities. Moving to a higher posterior mean $\widehat{\psi}$ of technology quality raises equilibrium adoption and disaster risk, increasing both downside and upside probabilities measured relative to each claim's forward payoff. We therefore examine whether favorable technology-stock returns coincide with increases in both option-implied tails. This interpretation treats technology returns as a proxy for favorable capability news, not as an identified shock to $\widehat{\psi}$.

Figure 5 presents separate downside and upside coefficients. Broad-market and QQQ tails generally load positively in later windows, consistent with the both-tail expansion in Table 4 and Figure 2. QQQ's upside response is more precisely estimated than its downside response. IWM's later downside response is more persistent and more precisely estimated, while its upside response is mostly insignificant, in contrast to the earlier episode of significant upside exposure. The pattern is not uniform: in the latest window, both IWM wings are significant and its upside coefficient exceeds its downside coefficient, while QQQ's downside interval includes zero. Differences in significance do not establish differences between coefficients; overlapping windows and pointwise bands also limit inference across dates and indices.

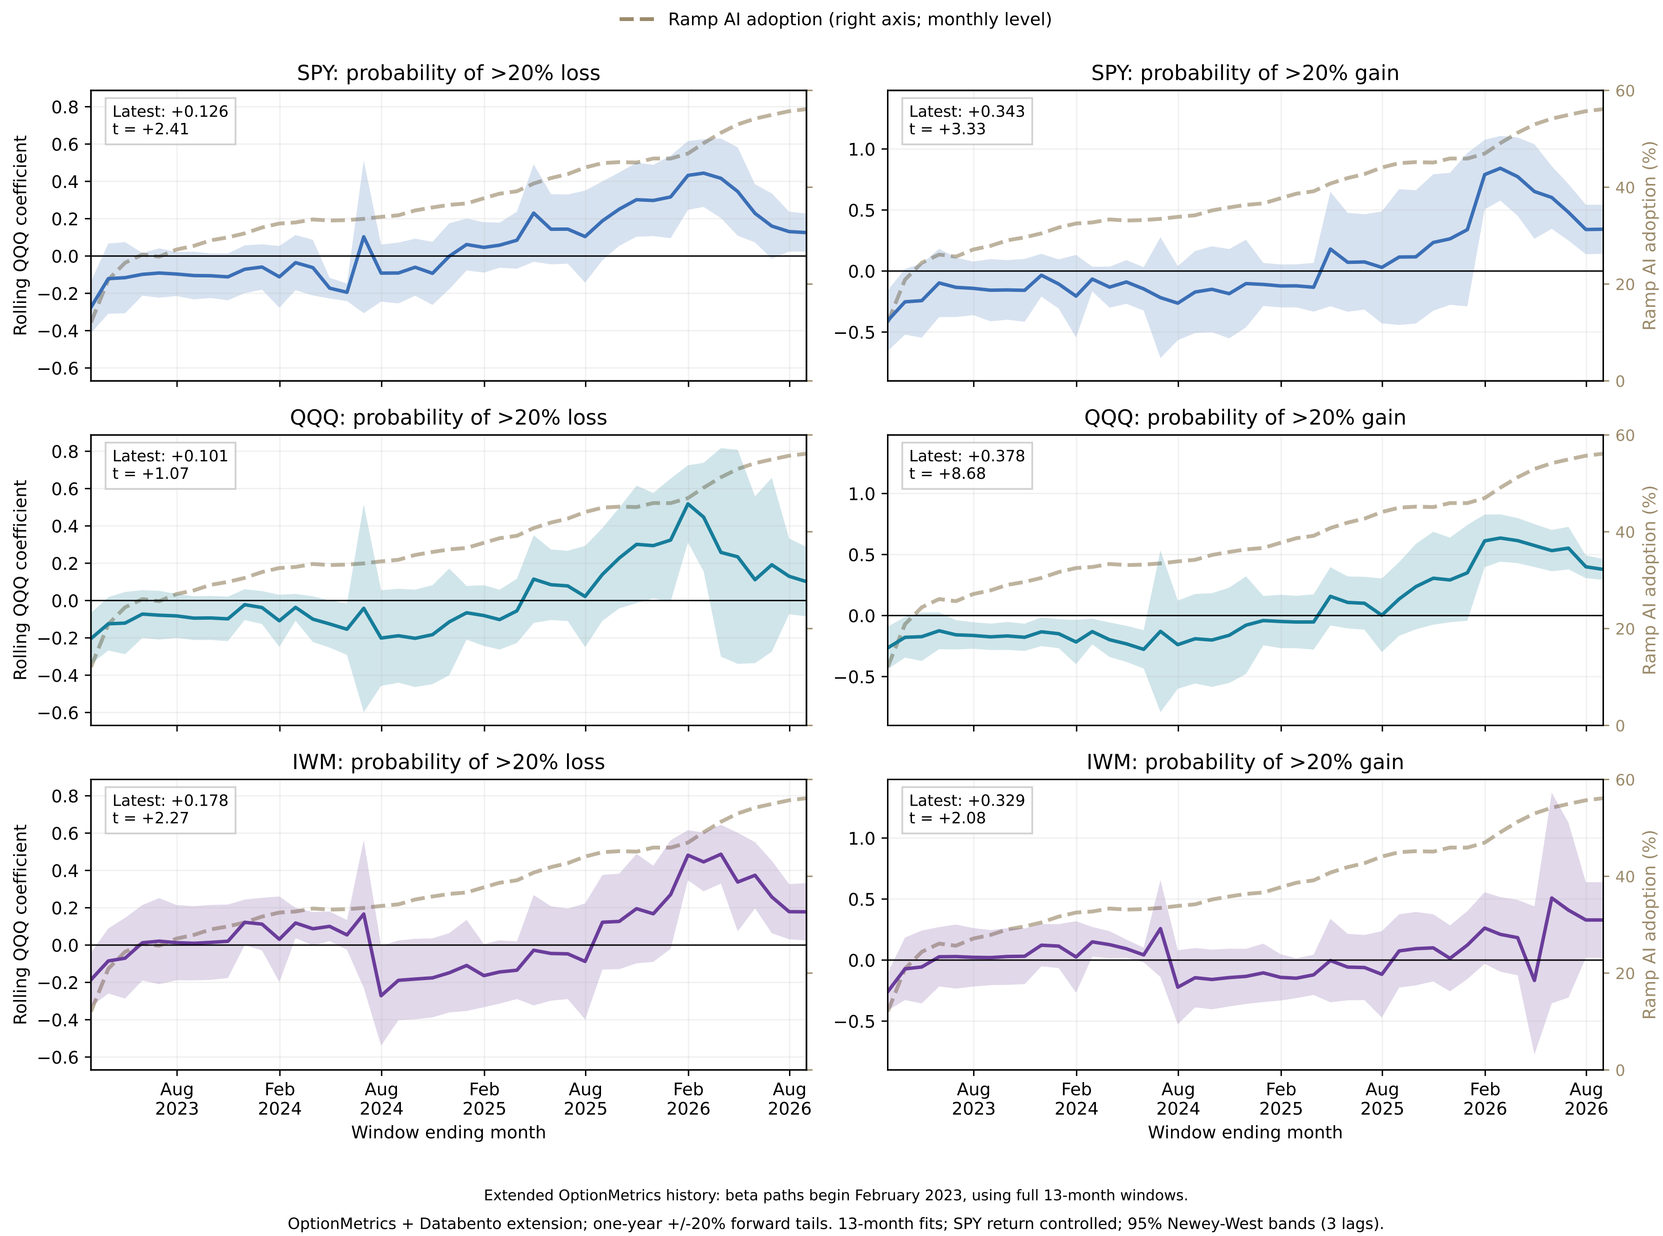

Figure 5: **Rolling QQQ-return coefficients for one-year downside and upside probabilities.** Rows show the broad market (SPX probabilities, labeled SPY), QQQ, and IWM. Columns show loss and gain probabilities beyond 20% of the one-year forward. Each fit includes an intercept and SPY and QQQ monthly returns, uses 13 months, and ends in the displayed month. Shading gives pointwise 95% Newey–West intervals with three lags. The dashed Ramp adoption level uses the right axis and is not included in the regressions. February 2023–August 2026.

The model connection is qualitative: the long-horizon illustration in Table 4 is not a quantitative prediction for the relative responses of one-year option tails.

Many other mechanisms could produce these patterns. QQQ and IWM mix adopters and suppliers, and changes in volatility, leverage, monetary news, growth expectations, or risk appetite can jointly move technology returns and option-implied tails. The estimates also combine physical beliefs with risk prices. Widespread generative-AI adoption is recent, leaving a short history with few independent observations of changes in capability and deployment; reliable inference will therefore remain difficult even with careful measurement. We therefore view the evidence as suggestive of the proposed mechanism, without identifying a causal effect or measuring AI disaster probabilities.

5 Conclusion

Markets can reward the adoption of a technology that makes a disaster more likely even when adopting and non-adopting firms suffer the same proportional loss. When capability raises both productivity and danger, disasters arrive disproportionately in the states where adopters are relatively most productive, and those are states of high marginal utility. Private firms price that payoff pattern while taking the aggregate hazard as given, so disaster risk raises the private return to adoption above what it would be if the technology were safe. A planner, internalizing the hazard that deployment creates, moves the other way. Decentralized adoption is therefore excessive for two reinforcing reasons. Firms ignore the externality, and prices reward the payoff pattern the externality produces.

The nature of the calamity determines whether this happens. A common survivable loss cancels from the adopter–non-adopter payoff ratio and survives only in state prices, where its dependence on capability encourages adoption. Damage confined to adopters is a direct relative cash-flow loss that competes with the selection benefit and restrains adoption in our benchmark. Extinction removes the financial claims that would carry the hedge and discounts the high-capability states in which adoption is most valuable. Markets reward deployment only when harm is widely shared and claims survive.

The mechanism relies on danger rising with capability. If advances in control and monitoring arrive alongside advances in capability, the selection effect weakens, and if danger falls with capability it reverses. Measures that reduce the hazard itself therefore weaken the hedging incentive along with the danger, while the incentive is strongest exactly when the technology is most dangerous.

These results suggest several policy responses. A corrective deployment charge can make firms internalize the additional hazard they create. An industry-wide penalty conditional on a survivable disaster could instead reduce the relative payoff to adoption in disaster states, provided liability is linked to each firm's deployment and the proceeds are transferred rather than destroyed. Such penalties cannot be collected after extinction. Measures that reduce disaster likelihood or severity offer another route, lowering expected damage and, under the conditions of Proposition 1, weakening the private hedging incentive.

References

Acemoglu, D. and Lensman, T. (2024). Regulating transformative technologies. *American Economic Review: Insights*, 6(3):359–376.

Andrews, I. and Farboodi, M. (2025). Do markets believe in transformative ai? Technical Report 34243, National Bureau of Economic Research. Revised July 2026.

Baker, S. D., Hollifield, B., and Osambela, E. (2020). Preventing controversial catastrophes. *The Review of Asset Pricing Studies*, 10(1):1–60.

Baker, S. D., Hollifield, B., and Osambela, E. (2022). Asset prices and portfolios with externalities. *Review of Finance*, 26(6):1433–1468.

Barro, R. J. and Ursúa, J. F. (2008). Macroeconomic crises since 1870. *Brookings Papers on Economic Activity*, 2008(1):255–350.

Bonney, K., Breaux, C., Dinlersoz, E., Foster, L., Haltiwanger, J., and Pande, A. (2026). The microstructure of ai diffusion: Evidence from firms, business functions, and worker tasks. Technical Report CES-26-25, U.S. Census Bureau.

Chen, A. Y. (2026). Hedging the singularity. Technical report, Working paper.

Chow, T., Halperin, B., and Mazlish, J. Z. (2026). Transformative ai, existential risk, and real interest rates. Technical report, Working paper.

Eisfeldt, A. L., Schubert, G., and Zhang, M. B. (2023). Generative ai and firm values. Technical Report 31222, National Bureau of Economic Research. Revised January 2026.

Gârleanu, N., Kogan, L., and Panageas, S. (2012). Displacement risk and asset returns. *Journal of Financial Economics*, 105(3):491–510.

Gutierrez, R. C. J. and Ready, R. C. (2026). Technological revolutions and stock prices revisited. Technical report, University of Oregon. Working paper.

Jones, C. I. (2024). The ai dilemma: Growth versus existential risk. *American Economic Review: Insights*, 6(4):575–590.

Kogan, L., Papanikolaou, D., and Stoffman, N. (2020). Left behind: Creative destruction, inequality, and the stock market. *Journal of Political Economy*, 128(3):855–906.

Kung, H., Lustig, H., and Paron, J. D. (2026). U.s. treasury investors are long in ai. Technical report, Stanford University Graduate School of Business.

Maresca, C. (2026). Agi could lower interest rates. Technical report, New York University.

Pástor, Ľ., Stambaugh, R. F., and Taylor, L. A. (2021). Sustainable investing in equilibrium. *Journal of Financial Economics*, 142(2):550–571.

Pástor, Ľ. and Veronesi, P. (2009). Technological revolutions and stock prices. *American Economic Review*, 99(4):1451–1483.

Shen, Y. (2026). Are ai risks priced in the u.s. stock market? evidence from financial news factors. Technical Report 2609.05485, arXiv.

Wachter, J. A. and Wachter, J. (2026). What investment data implies about the ai transition. Technical Report 35290, National Bureau of Economic Research.

A Fixed-date threshold proofs

A.1 Private threshold

Proof. Recall that H is aggregate terminal wealth without disasters, $H_q = \partial H/\partial q$ is the marginal conversion payoff, γ is risk aversion, and R_D is the common-disaster valuation multiplier. To see the sign, let P denote the physical probability measure and Q_0 the no-disaster state-price measure. Evaluated at the no-disaster threshold, its density is

$$\frac{dQ_0}{dP} = \frac{H^{-\gamma}}{\mathbb{E}[H^{-\gamma}]}. \tag{30}$$

At that threshold $\mathbb{E}^{Q_0}[H_q] = 0$. Both $R_D(q, \psi)$ and $H_q(q, \psi)$ are increasing in ψ, so for $q > 0$,

$$\mathbb{E}^{Q_0}[R_D H_q] = \mathrm{Cov}^{Q_0}(R_D, H_q) > 0. \tag{31}$$

The marginal firm therefore prefers adoption at the old threshold, and single crossing implies a lower threshold. At $q = 0$, $\Lambda(0, \psi) = 0$ and $R_D(0, \psi) = 1$.

For the comparative statics, compare two common-disaster economies with multipliers R_1 and R_2, holding q, the posterior variance, and all payoff primitives fixed. A lower retained fraction raises $\chi = \omega^{1-\gamma} - 1$. A proportional increase in intensity has the same effect on the coefficient of $\lambda(\psi)$ in the exponent. In either case,

$$\frac{R_2(q, \psi)}{R_1(q, \psi)} = \exp\{c\lambda(\psi)\}, \qquad c > 0. \tag{32}$$

At economy 1's threshold, define its normalized valuation measure by

$$\frac{dQ_1}{dP} = \frac{R_1 H^{-\gamma}}{\mathbb{E}[R_1 H^{-\gamma}]}. \tag{33}$$

Then $\mathbb{E}^{Q_1}[H_q] = 0$, while

$$\mathbb{E}^{Q_1}\left[\frac{R_2}{R_1} H_q\right] = \mathrm{Cov}^{Q_1}\left(\frac{R_2}{R_1}, H_q\right) > 0. \tag{34}$$

Strictness follows from nondegenerate capability uncertainty and strict monotonicity of both factors. Economy 2 therefore has a positive adoption gain at economy 1's threshold; its single-crossing condition implies a strictly lower threshold. At $q = 0$ both multipliers equal one, so the comparative effect vanishes. An additive capability-independent intensity increase instead makes R_2/R_1 constant in capability and leaves the threshold unchanged. All comparisons require finite valuation expectations.

□

A.2 Planner wedge

Proof. Recall that G_D^C and G_P^C are the decentralized and planner marginal gains, $\chi = \omega^{1-\gamma} - 1$ is the disaster valuation increment, and Δ is the remaining horizon. Subtracting the private condition from the planner condition gives

$$G_P^C - G_D^C = -\mathbb{E}\left[R_D H^{1-\gamma} \frac{2q\lambda(\psi)\Delta\chi}{\gamma - 1}\right] < 0. \tag{35}$$

Every factor inside the expectation is positive. The wedge is the marginal increase in the common disaster hazard created by additional deployment, which the firm takes as given and the planner internalizes. □

B Learning and pricing details

B.1 Productivity processes

The fixed-date economy has horizon $[0, T]$. A representative household owns firms, reinvests output before T, and consumes terminal wealth. Old-technology capital B_t^O grows at productivity ρ_t^O. Its long-run mean is $\bar{\rho}$, its mean-reversion speed is ϕ, and its innovation volatility is σ_O. Let $Z_{0,t}$ be the common standard Brownian motion:

$$dB_t^O = \rho_t^O B_t^O \, dt, \tag{36}$$

$$d\rho_t^O = \phi(\bar{\rho} - \rho_t^O) \, dt + \sigma_O \, dZ_{0,t}. \tag{37}$$

Productivity returns toward its long-run mean $\bar{\rho}$ at speed ϕ. The common old-economy shock is $Z_{0,t}$.

At t_0, investors begin observing a pilot sector using the new technology. Its long-run productivity is $\bar{\rho} + \psi$, with prior quality $\psi \sim \mathcal{N}(\mu_J, \sigma_J^2)$, where μ_J and σ_J are the prior mean and standard deviation. The pilot is small enough to reveal information without contributing materially to aggregate wealth. Write ρ_t^N for its productivity and $Z_{1,t}$ for an independent standard Brownian motion. Its loadings on the common and pilot-specific shocks are $\sigma_{N,0}$ and $\sigma_{N,1}$, respectively:

$$d\rho_t^N = \phi(\bar{\rho} + \psi - \rho_t^N) \, dt + \sigma_{N,0} \, dZ_{0,t} + \sigma_{N,1} \, dZ_{1,t}. \tag{38}$$

The second shock, $Z_{1,t}$, is independent of the old-economy shock. After subtracting observed common shocks and known mean reversion, investors obtain a noisy signal about ψ. The resulting posterior is Gaussian. Its mean $\widehat{\psi}_t$ responds to news, and its variance $\widehat{s}_t^2$ falls as observations accumulate. The signal and filter are given below.

Firms review adoption at t^*. We write $\widehat{\psi} = \widehat{\psi}_{t^*}$ and $\widehat{s}^2 = \widehat{s}_{t^*}^2$, so

$$\psi \mid \mathcal{F}_{t^*} \sim \mathcal{N}(\widehat{\psi}, \widehat{s}^2). \tag{39}$$

The fixed-date comparison isolates the uncertain long-run productivity gain. Write ρ_t^A for an adopter's productivity. It retains the old-economy shock but raises its long-run mean to $\bar{\rho} + \psi$:

$$d\rho_t^A = \phi(\bar{\rho} + \psi - \rho_t^A) \, dt + \sigma_O \, dZ_{0,t}, \qquad t \in [t^*, T]. \tag{40}$$

The dynamic extension also exposes adopters to the shared frontier shock; it is a separate dividend-flow specification rather than a numerical solution of these mean-reverting productivity equations.

B.2 The fixed-date productivity signal

The coefficient $\sigma_{N,0}$ is exposure to the observable old-economy shock $Z_{0,t}$; $\sigma_{N,1} > 0$ is exposure to the new-technology-specific Brownian motion $Z_{1,t}$, which is independent of $Z_{0,t}$. Thus observed old- and new-sector productivity reveal a noisy signal about ψ. After removing the common shock and known mean-reversion terms, the signal can be written

$$dY_t = \psi\, dt + \frac{\sigma_{N,1}}{\phi}\, dZ_{1,t}, \tag{41}$$

where Y_t is the cumulative orthogonalized productivity signal.

Let $\mathcal{F}_t$ denote the public information generated by observed old- and new-sector productivity through date t. Gaussian filtering implies

$$\psi \mid \mathcal{F}_t \sim \mathcal{N}(\widehat{\psi}_t, \widehat{s}_t^2), \qquad \widehat{s}_t^2 = \left[\frac{1}{\sigma_J^2} + \frac{\phi^2}{\sigma_{N,1}^2}(t - t_0)\right]^{-1}. \tag{42}$$

Here $\widehat{\psi}_t$ is the posterior mean and $\widehat{s}_t^2$ is the posterior variance. The posterior mean is a martingale that responds to unexpected new-technology productivity, whereas the posterior variance falls deterministically as signals accumulate. Equivalently, if $\widetilde{Z}_{1,t}$ denotes the innovation Brownian motion,

$$d\widehat{\psi}_t = \widehat{s}_t^2 \frac{\phi}{\sigma_{N,1}}\, d\widetilde{Z}_{1,t}. \tag{43}$$

Favorable AI news is therefore a positive innovation $d\widetilde{Z}_{1,t}$ and an upward revision in expected technology quality.

C Fixed-date incidence and welfare derivatives

C.1 Adopter-only damage

Recall that C_n^A is aggregate consumption after n adopter-only disasters and $U(C) = C^{1-\gamma}/(1-\gamma)$. Write $\Lambda_q = \partial\Lambda/\partial q$ for the marginal integrated hazard. In addition to the private gain G_D^A, the planner marginal gain G_P^A includes damage imposed on existing adopters:

$$G_P^A = G_D^A + \mathbb{E}_\psi\big[\Lambda_q \mathbb{E}_N\{U(C_{N+1}^A) - U(C_N^A)\}\big] = 0, \qquad \Lambda_q = 2q\lambda(\psi)\Delta. \tag{44}$$

Here $\mathbb{E}_N$ conditions on quality ψ, and $\mathbb{E}_\psi$ averages over its posterior. Because $C_{N+1}^A < C_N^A$, the additional planner term is negative.

C.2 Extinction payoffs and the planner

Recall that S is the survival indicator and $H(q,\psi)$ is surviving terminal wealth. Utility is evaluated at an exogenous reservation-consumption floor $\underline{C} > 0$ after extinction,

$$C_T^E = SH(q,\psi) + (1-S)\underline{C}. \tag{45}$$

For $q > 0$, define $\bar{k}(q) = K(q)/q$ as the average cost-adjusted installed scale per adopter. Let $X_{O,T}^E$ and $X_{N,T}^E$ denote the terminal payoffs of an old-technology firm and the average new-technology adopter, respectively, in the extinction economy. Corporate payoffs are zero after extinction:

$$X_{O,T}^E = S, \qquad X_{N,T}^E = S\bar{k}(q)e^{A_2(\Delta)\psi}. \tag{46}$$

The reservation endowment is not corporate property. It is only a reduced-form convention that keeps welfare finite in the extinction state.

Let $s(q,\psi) = e^{-\Lambda(q,\psi)}$ be survival probability and $s_q = \partial s/\partial q$ its deployment derivative. The planner gain G_P^E adds the effect on survival to the private gain G_D^E:

$$G_P^E = G_D^E + \mathbb{E}[s_q(q,\psi)\{U(H(q,\psi)) - U(\underline{C})\}] = 0, \tag{47}$$

where

$$s_q(q,\psi) = -2q\lambda(\psi)\Delta\, s(q,\psi) < 0. \tag{48}$$

If surviving consumption exceeds the reservation floor, this additional term is negative.

D Fixed-date corrective charge

Let τ_C^* be the corrective charge per marginal conversion under common survivable damage, in units of terminal consumption. Recall that R_D is the disaster valuation multiplier, H is surviving aggregate wealth, $\chi = \omega^{1-\gamma} - 1$, and $\Lambda = q^2\lambda(\psi)\Delta$ is integrated hazard. Then

$$\tau_C^* = \frac{\mathbb{E}[R_D H^{1-\gamma}\, 2q\lambda(\psi)\Delta\chi/(\gamma-1)]}{\mathbb{E}[H^{-\gamma}\exp\{\Lambda(\omega^{-\gamma}-1)\}]}. \tag{49}$$

The numerator is the hazard externality, not a premium for the firm's own physical exposure. For adopter-only damage it is the negative expected utility change from an additional event, weighted by Λ_q. For extinction it is $-\mathbb{E}[s_q\{U(H) - U(\underline{C})\}]$. The charge depends on beliefs, deployment, and the horizon over which that deployment creates exposure. A charge on a period of exposure must use that period's integrated hazard.

D.1 Illustrative corrective charges

The policy implication follows from the static deployment externality, as in Acemoglu and Lensman (2024). Let $Z(q) = \mathbb{E}[C_T^{-\gamma}]$ and consider a marginal charge τ in units of terminal consumption, returned to the representative household as a lump-sum transfer. The firm takes the rebate as given. Suppressing the disaster-incidence superscript, write G_D and G_P for the private and planner marginal gains. The firm's adoption condition becomes $G_D - Z\tau = 0$. At an interior planner allocation, the corrective charge is

$$\tau^*(q,\widehat{\psi}) = \frac{G_D(q,\widehat{\psi}) - G_P(q,\widehat{\psi})}{Z(q)}. \tag{50}$$

In the extinction case, a financial charge can be collected only in surviving states, so its denominator is $\mathbb{E}[sH^{-\gamma}]$. All quantities here use the fixed-date model's normalized old-technology payoff. The charge prices the additional hazard created by deployment. It is not compensation for a firm's own exposure or a tax on its productivity advantage. Its magnitude depends on deployment, beliefs, and how long the new hazard is present; Appendix D gives the incidence-specific formulas.

Table 1 evaluates the common reference belief used in the density exhibit, holding the hazard function fixed across incidence cases. We solve the decentralized condition and maximize expected utility globally over $q \in [0,1]$, including the endpoints. If V_D and V_P are expected utilities at those allocations, define certainty equivalents $CE_\ell = [(1-\gamma)V_\ell]^{1/(1-\gamma)}$ for allocation $\ell \in \{D,P\}$, with D denoting the market and P the planner. The loss is $1 - CE_D/CE_P$, the proportional consumption reduction from inefficient deployment in this normalized benchmark. Extinction welfare uses $\underline{C} = 0.30$ only as a reservation endowment; it is not a financial payoff.

Table 1: Static corrective charges and welfare losses at the common reference belief

Case	Market q (%)	Planner q (%)	Charge (%)	CE loss (%)
No disaster	26.92	26.92	0.00	0.000
Common survivable	30.00	13.50	18.57	3.286
AI-sector	25.73	23.64	2.70	0.033
Extinction	26.58	4.33	37.78	17.147

The charge is expressed as a percentage of one unit of normalized old-technology terminal payoff, evaluated at the planner allocation. Fixed-date severity is $\omega = 0.50$ (a 50% loss); extinction reservation consumption is 0.30 in the same normalized payoff units.

Common survivable risk raises market deployment while the planner internalizes losses across all consumption. Extinction produces a larger market–planner gap under the reservation level used here, illustrating the sensitivity of its welfare cost to that convention. The numerical magnitudes depend on severity, the capability–hazard relation, and the reservation convention under extinction. They are not calibrated policy rates for AI firms. Section 3 solves a dynamic planner and evaluates welfare separately. Its resource accounting and reservation-consumption convention differ from this fixed-date benchmark, so these static charges are not asserted to implement its policy.

How much extra adopter loss offsets the common hedge? Recall that $K(q)$ is cost-adjusted adopted capital and $A_2 = A_2(\Delta)$ is the fixed-date productivity loading. To separate selection from direct damage, fix the common event multiplier at $\omega = 0.5$ and add a proportional loss d to adopters at every event. Old claims receive ω^N and adopter claims receive $[\omega(1-d)]^N$. After removing the common multiplier, consumption is $1 - q + K(q)e^{A_2\psi}(1-d)^N$. At $q = 0.30$, holding the hazard function fixed, we evaluate the private gain at the no-disaster threshold and solve for the extra loss that makes it zero. The crossover is $d = 19.82\%$. This is extra damage on top of a common disaster, not a universal threshold for the pure AI-sector case. The calculation illustrates the competition in (18) without asserting that adopter damage always reverses the sign.

D.2 Threshold comparisons and additional adopter loss

Table 2: Fixed-date threshold comparisons at positive deployment

Case	Market versus no disaster	Planner versus market
No disaster	Benchmark	Equal
Common survivable	Lower under Proposition 1	Higher under single crossing
Adopter-only survivable	Ambiguous; direct loss competes with selection	Higher when additional events lower consumption
Extinction	Higher under single crossing	Higher when surviving utility exceeds reservation utility

The planner–market comparisons concern marginal conditions at a given state. They require the stated crossing conditions to translate into threshold or quantity orderings. A higher planner threshold than the market threshold does not by itself establish a comparison with the no-disaster threshold in every parameterization.

E Dynamic aggregation, valuation, and computation

E.1 Public beliefs and the installed-dividend state

Quality ψ is fixed but unknown. The relative-productivity signal is $dy_t = \varphi\psi\, dt + \sigma_N dZ_t$, as in Section 3. Let m_t and v_t denote the mean and variance of the Gaussian posterior based on this signal alone. For a time step dt,

$$v_{t+dt}^{-1} = v_t^{-1} + \varphi^2 dt/\sigma_N^2, \qquad m_{t+dt} = v_{t+dt}(m_t/v_t + \varphi\, dy_t/\sigma_N^2). \tag{51}$$

Let N_t be the cumulative number of observed disasters and $E_t = \int_0^t q_s^2 ds$ cumulative deployment exposure. Conditional on the productivity and disaster histories, the full posterior is

$$\pi_t(\psi) = \frac{\varphi_N(\psi; m_t, v_t)\lambda(\psi)^{N_t} e^{-E_t\lambda(\psi)}}{\int \varphi_N(z; m_t, v_t)\lambda(z)^{N_t} e^{-E_t\lambda(z)} dz}, \tag{52}$$

where $\varphi_N(\cdot; m, v)$ denotes the normal density. The public posterior mean $\widehat{\psi}_t$ is the mean of this density, not generally m_t. For extinction, financial decisions occur only before termination, so $N_t = 0$ and survival enters through E_t. These sufficient statistics preserve the non-Gaussian posterior without imposing a Gaussian approximation. Deployment affects disaster exposure but not the independent productivity signal.

On the displayed zero-shock productivity path from $m_0 = 0$,

$$m_t = \psi(1 - v_t/v_0), \qquad v_t^{-1} = v_0^{-1} + \varphi^2 t/\sigma_N^2. \tag{53}$$

The full posterior additionally conditions on survival along each allocation's deployment path. Thus identical productivity news need not imply identical full posteriors.

Let $B_t = D_t^O$ be the old per-firm dividend and $h_t = 1 - q_t + a_t$ consumption divided by B_t, so $C_t = B_t h_t$. Here q is deployed firm mass and a is installed adopter dividends divided by the old per-firm dividend. Conversion to mass $r \geq q$ changes installed dividends to $a^+ = a + K(r) - K(q)$, where K integrates retained conversion scales. Write R_O^0 for old-dividend growth before disaster losses and n for the number of disasters during the next interval. Then

$$B' = BR_O^0\omega^n, \quad a' = a^+e^{dy}, \qquad \text{common loss,} \tag{54}$$

$$B' = BR_O^0, \quad a' = a^+e^{dy}\omega^n, \qquad \text{adopter-only loss,} \tag{55}$$

$$C' = B'(1 - r + a'), \quad E' = E + r^2 dt, \quad N' = N + n. \tag{56}$$

Conversion-resource losses and repeated disasters therefore enter aggregate consumption through the sum of firm dividends. The numerical coordinate $x = \log[a/K(q)]$ records installed productivity relative to conversion scale, with $x = 0$ at $q = 0$. At conversion it becomes $\log\{[a + K(r) - K(q)]/K(r)\}$ before subsequent growth and losses.

E.2 Stopping values and the planner

Let $s = (q, a, m, v, E, N)$ collect the state excluding the old-dividend scale B, which is removed by homogeneity. The CRRA stochastic discount factor is

$$M_{t,t+dt} = e^{-\delta dt}(C'/C)^{-\gamma}. \tag{57}$$

For a fixed aggregate policy, let $\mathcal{A}_t(s)$ be the value per unit of an incumbent adopter's dividend and $\mathcal{W}_t(u;s)$ the value per unit of an old firm's dividend, including its option to convert at fixed cost type u. At proposed aggregate conversion r, the value of immediate adoption is $J_t(u;r)$ and the value of waiting is $L_t(u;r)$, both per unit of pre-conversion dividend:

$$J_t(u;r) = [1-\kappa(u)]\mathbb{E}_t[M_{t,t+dt}R_A\{dt + \mathcal{A}_{t+dt}(s')\}], \tag{58}$$

$$L_t(u;r) = \mathbb{E}_t[M_{t,t+dt}R_O\{dt + \mathcal{W}_{t+dt}(u;s')\}]. \tag{59}$$

The expectation uses the full posterior and the joint distribution of productivity news and disaster arrivals. R_A and R_O are adopter and old-firm dividend growth, including incidence-specific losses; both financial claims terminate at extinction. Aggregate consistency requires

$$r = q + \int_q^1 \mathbf{1}\{J_t(u;r) \geq L_t(u;r)\}\, du. \tag{60}$$

The solver finds marginal indifference or a corner, checks conversion-cost ordering, and selects the first downward crossing if several exist. No deployment throttle or cohort rule is imposed.

Let $V_t(B,s)$ denote planner welfare. Before extinction, the planner chooses $r \geq q$ to maximize discounted utility from the same aggregate dividends, including the effect of its choice on risk and information:

$$V_t(B,s) = \max_{r\geq q} \mathbb{E}_t[e^{-\delta dt}\{dt\, U(C') + V_{t+dt}(B',s')\}], \qquad U(C) = \frac{C^{1-\gamma}}{1-\gamma}. \tag{61}$$

For extinction, the terminal branch replaces financial consumption with the stated reservation-consumption convention. We set $\underline{c} = 0.30$ times counterfactual consumption in the termination interval and hold that level fixed through T. This is a welfare convention, not surviving corporate payouts. The market-policy welfare recursion holds r at the decentralized policy. Final flow payments occur at $T = 30$ and continuation values thereafter are zero.

E.3 Hazard and numerical approximation

Full-deployment intensity is bounded and increasing in true capability:

$$\lambda(\psi) = \lambda_{\max}\left[1 + \exp\left\{-\left(\alpha_D + \beta\frac{\psi - \psi_\star}{\bar{s}}\right)\right\}\right]^{-1}. \tag{62}$$

We use $\lambda_{\max} = 0.10$, $\psi_\star = 0.02974194$, $\bar{s} = 0.03118$, $\beta = 0.48990471$, and $\alpha_D = \log[0.017/(0.10 - 0.017)]$. The fixed scale $\bar{s}$ is not posterior uncertainty. The bounded intensity ensures finite disaster valuation moments; its upper bound is not an estimate of maximum AI risk.

The reported implementation uses half-year decisions, 61 signal-mean nodes on $[-0.12, 0.36]$, 17 installed-productivity nodes on $[-2, 5]$, adoption/type spacing 0.05, seven exposure nodes, and disaster counts from zero through eight. Extinction requires only the surviving count state. Gaussian quadrature integrates the signal and latent quality, and the interval disaster sum retains counts zero through seven. Continuation values interpolate in continuous states, while disaster counts remain discrete. Count-grid expansion, exposure-grid refinement, and quarterly decisions are checked separately; these checks do not constitute a joint convergence or equilibrium-uniqueness proof.

F Dynamic welfare and numerical diagnostics

Table 3 separates conditional year-six adoption, q_6, from ex-ante welfare at the initial prior. With initial expected utilities V_D under market adoption and V_P under planner adoption, the consumption-equivalent loss is $1-(V_D/V_P)^{1/(1-\gamma)}$. It integrates stochastic future outcomes, rather than evaluating utility at the displayed path.

Table 3: Conditional adoption and ex-ante welfare: 30% survivable loss

Case	Market q_6 (%)	Planner q_6 (%)	CE loss (%)
Common survivable	31.55	15.09	0.039
AI-sector	26.42	24.13	0.003
Extinction	27.42	0.00	1.497

Independent constant-intensity checks reproduce the previous market policies and welfare recursions and eliminate any dependence on disaster-count or exposure history. Across the reported grids, marginal-indifference residuals are below 10^{-6}, transition-probability errors are below 10^{-7}, and conversion-cost ordering holds. Displayed market and planner paths remain within the state grids. The planner's initial welfare weakly exceeds welfare under the market policy.

Some grid states have multiple downward crossings of the private marginal condition. The reported solution selects the first downward crossing; the calculations do not establish global equilibrium uniqueness.

Separate common-loss market calculations expand the count and exposure grids, refine the adoption and productivity grids, or halve the decision interval. Year-six market adoption with arrival information remains about 31–32% under a 30% loss and 91–92% under a 50% loss. These separate checks support the qualitative comparison but do not establish joint convergence for all incidence cases or the planner. The figures are quantitative illustrations, not estimated magnitudes.

G Risk-neutral density construction

Recall that ψ is quality, q is adopted mass, $H(q,\psi)$ is aggregate wealth before common disaster losses, and $\Delta = T - t^*$ is the remaining horizon. The loading $A_2 = A_2(\Delta)$ comes from the fixed-date productivity process with mean-reversion speed ϕ and volatility σ_O. The comparison restores the continuous common old-technology payoff that was normalized out of the adoption condition. Write this common factor as e^{ξ_O}, where ξ_O is its log payoff and v_O its variance, independently of ψ and the disaster process. Conditional on current old-technology productivity, the Gaussian loading in (37) gives

$$\xi_O \sim \mathcal{N}(0, v_O), \quad v_O = \frac{\sigma_O^2}{\phi^2}\left[\Delta - \frac{2(1-e^{-\phi\Delta})}{\phi} + \frac{1-e^{-2\phi\Delta}}{2\phi}\right]. \tag{63}$$

The mean is a scale normalization. This independent factor cancels from the adoption conditions but gives all three claims continuous payoff distributions, including the non-adopter. At the endpoints of the conversion-cost distribution, $\kappa_L = \kappa - \epsilon$ and $\kappa_H = \kappa + \epsilon$. For firm claim $i \in \{L, H\}$, write X_i for its terminal payoff and $\mathbf{1}\{\cdot\}$ for an indicator. With disaster count N and retained fraction ω, it pays

$$X_i = e^{\xi_O}\left[\mathbf{1}\{\text{adopt}_i\}(1-\kappa_i)e^{A_2\psi} + \mathbf{1}\{\text{do not adopt}_i\}\right]\omega^N, \qquad C_T = e^{\xi_O}H(q,\psi)\omega^N. \tag{64}$$

Let P be the physical posterior distribution of payoffs and Q its normalized marginal-utility-weighted counterpart. With risk aversion γ, all claims share the same pricing measure:

$$\frac{dQ}{dP} = \frac{C_T^{-\gamma}}{\mathbb{E}^P[C_T^{-\gamma}]}. \tag{65}$$

Under Q, ξ_O is Gaussian with mean $-\gamma v_O$ and variance v_O. Writing $f_P(\psi)$ for the physical posterior density of quality, its marginal density under Q is proportional to

$$f_P(\psi)H(q,\psi)^{-\gamma}\exp\{\Lambda(q,\psi)(\omega^{-\gamma}-1)\}, \tag{66}$$

and $N \mid \psi$ is Poisson with mean $\Lambda(q,\psi)\omega^{-\gamma}$. The exponent differs from the adoption multiplier $\chi = \omega^{1-\gamma} - 1$ because a probability receives the marginal-utility weight $\omega^{-\gamma N}$, whereas the payoff gain in that condition supplies an additional ω^N. Applying a different tilt to each claim would be inconsistent with a common household stochastic discount factor.

The fixed-date illustration uses $\gamma = 4$, $A_2 = 19.185$, $\widehat{s} = 0.03118$, $\Delta = 22$, $\kappa = 0.10$, $\epsilon = 0.025$, and $\omega = 0.50$. Risk aversion, conversion costs, and the capability-dependent hazard match the dynamic benchmark; survivable severity matches its 50% loss case. Background productivity volatility is $\sigma_O = 0.02$ in both models. The fixed-date horizon, productivity loading, and Gaussian review-date posterior remain specific to this terminal-consumption illustration. The annual intensity uses the same bounded logistic schedule as Appendix E, with a cap of 0.10, reference quality $\psi_\star = 0.029742$, and reference intensity 0.017. For deployment q, the Poisson mean is $q^2\lambda(\psi)\Delta$. Holding this hazard schedule fixed, we solve for the belief that produces private deployment $q = 0.30$. The higher belief is 0.01659; the lower belief, -0.02085, lies half a posterior standard deviation below the low-cost firm's entry threshold and produces $q = 0$.

Stock-price channels. Let X_i be the terminal payoff, Π_i its price, and B^{unit} the unit-payoff price. With physical and risk-neutral expectations $\mathbb{E}^P$ and $\mathbb{E}^Q$, respectively, $\Pi_i = B^{\text{unit}}\mathbb{E}^Q[X_i]$. Here $\Delta \log$ denotes a finite change between beliefs:

$$\Delta \log \Pi_i = \underbrace{\Delta \log B^{\text{unit}}}_{\text{unit-payoff price}} + \underbrace{\Delta \log \mathbb{E}^P[X_i]}_{\text{expected cash flow}} + \underbrace{\Delta \log \frac{\mathbb{E}^Q[X_i]}{\mathbb{E}^P[X_i]}}_{\text{risk adjustment}}. \tag{67}$$

Capability affects expected cash flows and their state-price valuation. The distribution exhibit isolates the latter through a common Q. Under extinction, the unit-payoff comparison claim is survival contingent.

Table 4: Risk-neutral probabilities beyond ±20% of own forward payoff, percent

Claim	0%: down	0%: up	30%: down	30%: up
Low-cost firm	20.6	18.8	47.8	28.1
High-cost firm	20.6	18.8	32.3	31.9
Aggregate claim	20.6	18.8	33.4	29.6

H Empirical data construction

Option probabilities are constant-maturity, strike-band digital approximations constructed from OptionMetrics and a Databento extension. The main-text equations state the put and call spreads, discounting, and forward-relative thresholds. For QQQ and IWM, each wing's Databento probability is adjusted by subtracting the median Databento-minus-OptionMetrics level difference on overlapping dates. OptionMetrics observations are retained through their final date; adjusted Databento observations are appended afterward. The adjustment is applied separately to downside and upside probabilities. SPY-labelled probabilities use the saved SPX series splice and are paired with SPY returns.

Early-2022 OptionMetrics observations extend the usable probability changes back to February 2022, permitting the first full 13-month regression window to end in February 2023. Saved overlap checks reproduce existing probabilities and monthly observations. Available expirations are interpolated to a one-year maturity. Month-end probabilities use the last available observation separately for each wing; monthly returns compound daily returns. Probability changes and returns are in percentage points. The final complete month is August 2026. Every regression includes an intercept and QQQ and SPY returns and uses Newey–West covariance estimates with three lags.

A vendor change can affect probability changes near the join even after a level adjustment. ETF American exercise, forward estimation, strike interpolation, and finite spread widths are additional approximations. These objects average digital probabilities over strike bands rather than measure exact pointwise tails. Ramp's monthly adoption level is displayed only over its available dates on a separate axis.

The portfolio comparison matches December 31, 2025 CRSP Mutual Fund Holdings to the total workforce-task exposure scores of Eisfeldt et al. (2023), using date-valid security identifiers and CRSP–Compustat links. The October 2025 score release reflects March 2022 workforces and March 2023 task assessments. We exclude non-equity positions and combine share classes. Reconstructed November 2022 NYSE exposure quintiles define H as the bottom quintile and A as the top; these approximate the authors' sorts rather than reproduce exact portfolio memberships. The main-text shares count each firm once and condition on membership in A or H, excluding the middle quintiles and unclassified firms. IWM has 232 A and 162 H firms, giving a denominator of 394; QQQ has 26 A and 9 H firms, giving a denominator of 35. These extremes represent 37.3% of IWM's 1,055 classifiable firms and 41.7% of QQQ's 84. Scores cover 54.4% of IWM's 1,941 firms and 84.0% of QQQ's 100 firms. Unclassified firms are excluded, not assigned to H. The comparison is therefore conditional on coverage and describes labor-task exposure, not firm size or other determinants of adoption.