

Forcing Technologies: Why Firms Adopt and Who Profits

Ren (Electric Capital), GPT 6 Astra and Claude 5.1 Fable

October 20, 2026

Working draft, October 20, 2026.

The argument in one page

In June 2010, Spread Networks and Infinera announced a New York–Chicago fiber route below 13.33 milliseconds round trip. By February 2013, McKay Brothers reported 8.27 milliseconds on its Illinois–New Jersey microwave route (Infinera, 2010; McKay Brothers, 2013). Traders competing to arrive first faced a choice: buy speed or lose position. Their suppliers faced the next generation of the same race. Zayo agreed to buy Spread for \$127 million in November 2017 (Zayo Group Holdings, 2017). An urgent purchase had not guaranteed its first supplier a durable advantage.

Technology becomes economically forcing when its spread makes the old way of operating less viable for holdouts. That pressure creates urgency; market structure determines who captures the spending, and for how long. Rivals change prices, service expectations and access to customers. Attacks make defenses inadequate. Regulatory and private requirements supply another mechanism, with distinct enforcement risks.

The two-curve model separates the incentive to adopt from the buyer’s gain or loss and the supplier’s price ceiling. A tool can pay its first user and become necessary as rivals change the market. The paper follows that transition through competition, fulfillment costs, networks and defenses; prices, deployment and supplier profits require additional assumptions. Its test is an observable mechanism, an exposed buyer and a credible timing window: a margin crossing or deployment schedule can establish one without an announced deadline.

AI and crypto show how adoption pressure and supplier returns can differ across applications, platforms and infrastructure. NVIDIA earned \$193.7 billion of Data Center revenue in fiscal 2026; Ethereum’s 2022 Merge removed mining from its network while Bitcoin’s hardware race continued (NVIDIA Corporation, 2026; Ethereum.org Contributors, 2026b). Section 4 compares four crypto networks and five foundation-model strategies, separating useful services, platform advantages and repeated infrastructure spending.

The investment screen asks which companies can retain revenue after adoption becomes common. A comparison of sixteen investment opportunities identifies the following priorities (Section 6); these are conditional diligence judgments, not completed investment decisions:

1. **AI-operated specialist services.** Own accepted work and recurring customer contracts; prove that automation reduces cash cost after review and repricing.
2. **Repeatable industrial cells.** Own an application, reliable field service and financeable equipment; deployment capital and guarantees determine the return.
3. **Flexible power access, subject to enforceable contracts.** Accepted load flexibility can change usable capacity; a scheduler must possess enforceable operating and payment rights.
4. **Cross-border payout operations, subject to access to customers and local partners.** Stablecoin rails can improve a difficult endpoint; all-in costs, local liquidity and retained fees constrain the entrant.
5. **Fraud signals and seller negotiation, as options.** These opportunities require evidence of additional fraud prevented in live use and higher seller profit after costs. Complex certificate operations remain a requirement-led benchmark.

What would prove us wrong

The economic thesis fails in a proposed market if holdouts retain their margins, customers and access as adoption spreads. The investment thesis fails if an equivalent substitute captures the spending, delivery costs consume the gain or financing outlasts the useful window. Public deployments show that these technologies can work. The scenarios examine the adoption, pricing and cash requirements a new supplier would face. Neither a useful technology nor an urgent purchase guarantees an attractive company.

1 The model

Take a fixed field of rivals: the merchants selling into a state, the trading firms on an exchange, the swimmers in a final. Let x be the share of them that has adopted a technology, from none to all. Each rival is small enough that its own choice leaves x where it is. Two payoffs describe its position, both measured against the world before the technology existed: $U_A(x)$ if it adopts and $U_N(x)$ if it holds out. Zero means nothing has changed. Both curves can sit below zero at once, which is what an arms race looks like from the inside. The cost of adopting, c , is a flow per period: the subscription, the hardware amortized, the staff who run it.

Economic forcing is the deterioration of the holdout's position as the technology spreads: rivals take its customers, lower the price it can charge, raise the service customers expect, or leave it outside a useful network. A positive return to adoption is insufficient evidence of that mechanism. A tool can pay its first user before it changes anyone else's circumstances, then become necessary as competitors pass its benefits through to customers. Requirements imposed by a platform, a buyer or a regulator can produce similar urgency, but their enforcement is a separate source of pressure. We first identify the mechanism and the exposed buyer, then ask what a supplier can capture.

1.1 Four channels

Four sources of value or pressure determine the payoff gap.

Direct value, γ .

The technology does something for the adopter whether or not anyone else has it. An AI assistant increased issues resolved per hour by 15% in a study of 5,172 support agents, without requiring rival firms to adopt (Brynjolfsson et al., 2025). Net value also depends on its cost.

Network value, βx .

The technology is worth more as counterparties join. A bank connection is worth nothing until both the bank and the app are on it; Plaid's product got better with every bank it added and every app that integrated it.

Competitive edge, α and r .

The adopter takes something from rivals that have not adopted: business, a ranking, the next trade. The prize is fixed, so one side's gain is the other's loss, and the edge erodes as rivals catch up. Spread Networks' fiber gave its customers a speed advantage that competing routes could erode. The strength ratio r determines how the edge changes with adoption; it does not set the speed of diffusion.

Exclusion, $\varepsilon\Phi(x)$.

A gatekeeper penalizes holdouts, and nobody in the field collects the penalty. After *Wayfair*, a remote seller that did not collect sales tax answered to the state rather than to a rival. The gatekeeper can as easily be a browser, a card network, a dominant customer or an insurer, and Section 2 says what changes when it is. The pressure curve Φ rises from zero to one; θ locates

its rise in adoption share and k controls its steepness. A calendar deadline requires a separate time-dependent enforcement path.

Let p be the shared belief that the technology works. Subtracting the holdout's payoff from the adopter's gives the gain from adopting,

$$\underbrace{\Delta(x)}_{\text{adopt minus hold out}} = \underbrace{p\gamma}_{\text{direct}} + \underbrace{p\beta x}_{\text{network}} + \underbrace{\frac{p\alpha}{1 + (r-1)x}}_{\text{contest}} + \underbrace{\varepsilon\Phi(x)}_{\text{exclusion}} - c.$$

Beliefs scale the three technology channels and leave the penalty alone, which is why a company can doubt a technology and still buy it before a deadline. A rival acts when the gain exceeds the value of waiting, W , and once it has acted it cannot undo the decision. The four channels are the model's accounting; the economics that fill them, rivals, counterparties, cost curves, customers and gatekeepers, are the subject of Section 2. Appendix A derives the two curves, adds an outside option, and walks through each parameter one at a time.

The field and the supplier must be named separately. An application developer choosing a chain or model, a miner replacing machines and a laboratory buying compute face different payoff comparisons (Section 4). Network value means that other participants or complements improve the adopter's service. Spreading a fixed training bill over more customers is instead an economy of scale; migration and re-evaluation costs can retain customers without a network effect. Those distinctions determine what evidence can support β , and which business might charge for it.

1.2 Three tests and four patterns


With the technology proven ($p = 1$), three numbers summarize a technology for a rival of average cost:

$$\begin{aligned} \Delta(0) &= \gamma + \alpha - c && \text{Does adopting first pay?} \\ \Delta(1) &= \gamma + \beta + \alpha/r + \varepsilon - c && \text{Would dropping it hurt once everyone has it?} \\ y_R &= \gamma + \beta - c && \text{Is the field better off with everyone in?} \end{aligned} \tag{1}$$

The third line adds up everyone's payoff, so the contest cancels and the penalty disappears from the total: at full adoption there are no holdouts left to pay it, although a defector would still face it. The formal taxonomy calls a technology *forced* when $\Delta(1) > 0$, *competitively compelled* when $\Delta(0) \leq 0$, and *wasteful for the field* when $y_R < 0$. These are endpoint labels, not causal findings: a useful standalone product can satisfy the full-adoption test with no harm to holdouts. Economic forcing additionally requires evidence about how the holdout's position changes. Table 1 combines the first two tests with the third, and Figure 1 draws one technology of each kind.

Table 1: Four endpoint patterns. The sign of y_R gives the labels in the last two columns.

Pattern	Adopt first / stay at full adoption	$y_R > 0$: full adoption pays	$y_R \leq 0$: no net gain at full adoption
Dominant	yes / yes	productive	arms race
Tipping	no / yes	coordination	wasteful bandwagon
Selective	yes / no	selective	wasteful selective
Never	no / no	suppressed	inert



All panels: known to work ($p = 1$); identical costs ($\sigma = 0$); no outside option or waiting value ($Q = -\infty, W = 0$); exclusion only in the mandate panel. Illustrative payoff units.

Figure 1: One technology of each kind, with its parameters above each panel. The vertical gap between the curves is the incentive to adopt; the height of each curve relative to zero is the change from the world before the technology. In the arms race, everyone adopts and everyone ends below zero. In the coordination case, nobody moves until half the field has, and then everyone gains. Under selective forcing the edge wears away and adoption stops at 42%, where adopters and holdouts are equally badly off. Under the mandate, the adopter's curve never rises above -0.2 ; the penalty on holdouts does all the work, and the field as a whole ends worse off than it started.


Figure 2 places 14 historical cases on the last two tests. To the right of zero, staying adopted pays when everyone else has adopted; above zero, full adoption leaves the rival group better off than no adoption. In the fitted estimates, hybrid corn and containers have positive group payoffs across their displayed intervals. Dreadnoughts and clap skates sit in the lower-right area: *privately worth keeping, collectively costly*. Frequent-flyer programs and LZR suits lie near the welfare boundary, where the intervals matter more than the point estimates. These are comparisons at full adoption, including cases whose observed adoption remains partial; their interpretation depends on the model and its priors.

The welfare columns and historical map are the buyers' accounting; nothing in them says what a company earns.

Historical cases: adoption pressure and collective payoff

Dots: posterior medians; bars: 90% marginal credible intervals

Rivals' net benefit at full adoption, y_R



Normalized within each case; conditional on priors and the observation model.

Source: stored full-model posterior summaries, 14 historical cases.

Figure 2: Historical cases on the adoption-pressure and welfare map. Dots pair posterior medians of $\Delta(1)$ and y_R ; bars show separate 5th–95th percentile intervals, not joint regions. Estimates are conditional on priors and observation assumptions, use units normalized within each case, and exclude outside options, counters and wider spillovers. They describe full-adoption incentives and rival-group payoffs, not causal holdout losses or dollar comparisons across industries. Data and estimation: Appendices J and K.

1.3 Productivity can change the market around the buyer

The fixed-prize contest is one source of that deterioration, not a general model of price competition. Suppose automation saves a firm γ per period and competitors pass some of those savings through in lower market prices. Let $h(x)$ be the resulting revenue loss at a fixed output, common to adopters and holdouts, with $h(0) = 0$. Their changes from the old operating profit are

$$U_A^{\text{prod}}(x) = \gamma - c - h(x), \quad U_N^{\text{prod}}(x) = -h(x).$$

The switching gain remains $\gamma - c$, yet the holdout's operating profit, its original profit less $h(x)$, can cross zero. Adoption was profitable at the start and can become necessary to keep operating later. That crossing supplies a timing window even when no rule supplies a date; installation time must be subtracted from the window.

This common downward shift is an explicit extension to the base model's payoff levels. It leaves the direct-value adoption gap unchanged but changes producer welfare and viability: at full adoption the field's change is $\gamma - c - h(1)$, and consumers receive the price reduction. It is not a zero-sum transfer among the firms, and the base model's outside-option and welfare results cannot simply be applied to it. The case in Section 7.2 works through the extension with stated assumptions. In a network, likewise, a rising adoption gain alone establishes attraction; an exclusion or lost-demand mechanism must be observed before claiming that holdouts are being harmed.

1.4 The price ceiling moves with adoption

A buyer facing exposed rivals will pay at most its gross benefit from adopting,

$$B(x, p) = p \left[\gamma + \beta x + \frac{\alpha}{1 + (r - 1)x} \right] + \varepsilon \Phi(x), \quad (2)$$

less its own implementation cost κ and the value of waiting W . The company's price q has to fit inside $B - \kappa - W$, and the slope of B in x says what happens to that room as the market fills:

$$\frac{\partial B}{\partial x} = p \left[\beta - \frac{\alpha(r - 1)}{(1 + (r - 1)x)^2} \right] + \varepsilon \Phi'(x).$$

Under a contest the ceiling falls: the tenth buyer of a speed advantage gets less of it than the first, and the company either cuts price or ships the next generation. Under network value and exclusion the ceiling rises: holdouts face a larger exclusion penalty, and new participants gain access to more counterparties. Under direct value alone the incremental ceiling is flat. Price competition can still make the purchase urgent while reducing the buyer's ability to pay; the productivity case therefore checks continued operation and competing suppliers as well as the incremental ceiling. The 14 historical fits in Appendix J show both slopes: the fitted budget for a dreadnought falls from about 1.9 to 1.0 as navies catch up, and the budget for a container berth rises from 0.5 to 1.0 as ports join.

1.5 Two exits that bound the pressure

Holdouts have two ways out, and each caps what a company can charge. One is to leave the contested segment: ally, specialize, or exit, for a payoff $Q \leq 0$. Once holdouts take the outside option, adopters lose their competitive edge. The adoption gain falls by the same amount and can stop further adoption, under the shakeout conditions of Proposition 3. The other is a counter, a defense that shields a share φ of a holdout's stake at cost c_K . In this model, demand for a counter develops as the technology it defends against spreads (Appendix F). The two exits disarm different

things. An exit can end pressure from a contest or a gatekeeper; a counter can end pressure from a contest only. Neither disarms a technology whose direct and network gains alone cover its cost (Corollaries 1 and 4). For a company, an exit is a lost customer and a counter may be a second product line.

1.6 What each channel means for a company

Table 2 sets the five channels side by side as company business models.

Table 2: Supplier business models associated with each channel. “Ceiling” is the buyer’s maximum price $B - \kappa - W$ as adoption spreads.

Channel	Who pays	Why now	Ceiling as x rises	Durability	Main commercial risk
Direct value γ	The adopter, from its own P&L	ROI first; a margin or service threshold can later make delay costly	Incremental ceiling flat; viability and substitutes can cap it	While savings persist and the company retains a cost or workflow advantage	The incumbent bundles; competitors pass the saving to customers
Network βx	Adopters, once counterparties are on	After a threshold, if standalone value is insufficient	Rises; supplier capture depends on control of access	Strong where the operator retains the network	The standard opens; a gatekeeper reprices access
Contest α, r	Exposed rivals whose gains cover costs	Early buyers have the largest edge	Falls as the edge erodes, unless the company sells the next generation	Through re-buys, while the race runs	A ban; a counter; the venue sells the race itself
Exclusion $\varepsilon \Phi$	Holdouts, on the gatekeeper’s clock	Expected enforcement; often a deadline	Rises with expected penalty; need not fall after compliance	Recurring service or a temporary migration wave	Installed demand saturates; the capability is bundled; enforcement changes
Counter φ, c_K	The exposed side of someone else’s race	Only once the attack has spread	Rises with the attack	While the attack keeps evolving	The platform defends natively; the attack is banned

For investment, the rows answer two separate questions: what makes waiting costly for the buyer, and what lets a company retain part of the resulting budget? A contest can renew the purchase with each generation; productivity gains can make the old cost structure unviable; a network can make disconnection costly; a mandate can announce an enforcement date. None guarantees supplier profit. The next section traces these pressures before the paper applies a narrower screen for a new company’s entry, capture and timing.

2 Where the pressure comes from

The question is what happens to a firm that keeps operating as before while others adopt. It may lose customers to a faster rival, face lower market prices, fall short of a service expectation, or lose access to counterparties. Those are economic mechanisms. A procurement requirement, platform policy or statute can also make holding out costly, through an imposed condition rather than the market response alone. Both kinds can operate together: fraud creates a loss, for example, while a liability rule changes who bears it. The examples below establish mechanisms, not their frequency across the economy. For each, an investor needs evidence of the buyer’s exposure and a timing window, followed by a separate account of how a supplier can retain the budget.

2.1 Rivals force each other when the prize is fixed

The simplest forcing has no gatekeeper in it. When customers, rankings or fills are allocated by relative position, a technology that improves position is worth having only because rivals lack it, and a rival that lacks it loses ground whether or not it ever buys. Hirsch called the goods positional; Lazear and Rosen showed that when pay depends on rank, the effort each contestant puts in is set by the spread between the prizes rather than by what the effort produces; Frank and Cook described the markets in which small differences in position carry most of the reward, and why the spending they induce is wasted from the group’s point of view (Hirsch, 1976; Lazear and Rosen, 1981; Frank and Cook, 1995). In the model this is the contest term, $\alpha/(1+(r-1)x)$: a lone adopter takes ground from every holdout, each later adopter has less ground to take, and a holdout’s loss, $\alpha x/(1+(r-1)x)$, grows with every rival that moves. Adoption here is a strategic substitute in the sense of Bulow, Geanakoplos and Klemperer: rivals’ adoption lowers the gain from adopting while raising the cost of abstaining (Bulow et al., 1985). Two things follow. There is no tipping point, so adoption can begin without a critical mass of users; and a static equilibrium can lie well short of everyone, with the holdouts worse off than before the race began. Irreversible adoption can overshoot that equilibrium; the simulated field in Case 1, Section 7.1, reaches 72% at month 48. The field as a whole ends poorer whenever the direct value is below the cost: the arms race of Table 1, and the Red Queen of Barnett and Hansen, in which rivals’ improvements offset one another while their costs remain (Barnett and Hansen, 1996).

Markets with this structure are easy to list, because a venue ranks, matches or fills. Trading (Spread Networks), search and AI answers (Section 6.5), applicant screening and programmatic advertising all allocate by position. Retail media shows how the venue can profit from competition for placement. A retailer’s shelf, search box and app are positional: placement can move business between brands, giving exposed brands a reason to buy while the gain covers their costs. Retailers have turned that competition into a product line. Amazon’s advertising revenue went from \$37.7 billion in 2022 to \$56.2 billion in 2024 and about \$68.6 billion in 2025; Walmart’s grew 46% to nearly \$6.4 billion in its fiscal 2026; Instacart’s reached \$1.07 billion, and it sells the capability on to more than 310 other grocers (Amazon.com, Inc., 2025, 2026d; Marketing Dive, 2026; Walmart Inc., 2026; Maplebear Inc., 2026). More than 200 retailers ran a network by the end of 2024, up from 150 a year earlier, because, as one company put it, “pretty much every retailer feels like they need to be doing this” (Mewborn et al., 2023; Modern Retail, 2024), and 36% of the consumer-goods companies surveyed in late 2023 were buying from ten or more networks, against 6% in 2022 (Path to Purchase Institute, 2024). The software intermediary grew more slowly: Criteo’s retail-media revenue rose 2% in 2025 and fell 17% in the fourth quarter, which the company attributed partly to scope changes at two clients (Criteo S.A., 2026). The contrast illustrates different revenue paths; it does not measure how much profit each layer captured. Exposed rivals buy while the competitive edge covers their costs. Early adopters value it most, and new generations can sustain demand.

Venues can also sell the advantage directly, as exchanges do with colocation and search engines with advertising. Regulation can end this demand. A contest that imposes costs on the field may attract a ban (Section 3.1).

2.2 Network adoption can accelerate without an announced deadline

The opposite structure is the strategic complement: a rival’s adoption raises your gain from adopting. Katz and Shapiro and Farrell and Saloner worked out what follows when a technology is worth more the more others use it: an installed base nobody wants to be first to leave, inertia below a threshold and a bandwagon above it, and a winner chosen by expectations as much as by quality (Katz and Shapiro, 1985, 1986; Farrell and Saloner, 1985, 1986; Arthur, 1989). Under their strategic-complementarity conditions, Milgrom and Roberts established largest and smallest equilibria that bound adaptive play; convergence to one of those extremes is not a general result (Milgrom and Roberts, 1990). In the model this is βx . With identical costs, known effectiveness and no waiting value, a network-led tipping point occurs where $\Delta(x_T) = 0$ crosses from negative to positive. Below it no unassisted actor adopts; above it the gain from connecting grows. A falling holdout payoff additionally requires that unconnected firms lose access or business; the network term alone does not establish that loss. Someone has to pay for the crossing, because the pioneers adopt at a loss. Sometimes a regulator pays with a mandate. A large actor can instead fund the crossing: by subsidizing one side, which is Rochet and Tirole’s platform, pricing the two sides so as to internalize the externality between them (Rochet and Tirole, 2003); by switching the default on for an installed base it already owns, as Stripe and Shopify did with stablecoin acceptance (Section 4); or by being large enough that its own adoption is the tip. Switching costs then hold the field where it tipped (Klemperer, 1987; Farrell and Klemperer, 2007).

When the participants are competitors, these incentives can support coordinated pricing. The Justice Department’s complaint against RealPage, filed on August 23, 2024, alleges that rivals encouraged one another to use a tool: the software set prices for about three million apartments on data from sixteen million, with about 80% of its market, and the complaint alleges that landlords “know that others are using the same software,” asked whether competitors nearby were already on it before licensing it themselves, and said they would rather their competitors used it too, since widespread use would benefit them all (U.S. Department of Justice, Antitrust Division, 2024; U.S. Department of Justice, 2024). Those allegations describe a possible complementarity among landlords, with renters bearing the cost; they do not establish a measured β or a causal adoption path. Thoma Bravo paid \$10.2 billion for RealPage in April 2021 (RealPage, Inc., 2021). The Justice Department’s proposed settlement of November 24, 2025 would restrict competitors’ nonpublic data in pricing, require changes to price-alignment features and impose a monitor (U.S. Department of Justice, 2025). It was proposed without a trial or admission of wrongdoing. This is evidence of legal risk to a coordination mechanism, not evidence that RealPage retained its customers or lost all the reasons they bought.

Section 3.3 examines the commercial consequence: ownership of access can matter more than having funded the network’s growth. Shapiro and Varian’s advice to give the technology away until the field tips and charge for the lock-in afterwards works only if the pioneer owns the lock-in (Shapiro and Varian, 1999). Plaid built the connection layer but remained dependent on banks for data access; Sea-Land pioneered containerization without owning the resulting standard. The investor should ask whether the startup will control an asset that remains valuable after the network has formed.

2.3 A falling cost curve can make the old cost structure unviable

Wright found in 1936 that the labor cost of an airframe fell by a fixed fraction with every doubling of cumulative output, and Arrow built the economics of learning by doing on the same regularity (Wright, 1936; Arrow, 1962). A learning curve can make adoption profitable for a class of buyers in a similar period. Griliches showed that the timing and spread of hybrid corn tracked its profitability in different regions (Griliches, 1957). That first crossing establishes an adoption incentive. Economic forcing requires the next step: lower costs change the prices or service levels rivals can sustain, and the firms that keep the old process lose margin or demand.

The two steps need separate measurements. A pilot can establish a saving; price quotes, contract renewals, conversion rates and holdout margins show whether the saving is changing competition. Amazon’s description of its next-generation fulfillment center makes the proposed connection explicit: it targeted a 25% improvement in cost to serve during peak seasons and identified savings it could pass to customers (Amazon, 2024). That is a company target, not a measured industry effect. The fulfillment case in Section 7.2 assumes a pass-through path and shows exactly when it would exhaust a holdout’s margin, and how lower pass-through would remove that crossing. Its point is that useful automation and competitive pressure can coexist; neither a saving nor a deployment count establishes the second on its own. Published price and adoption data show how quickly these economics can change. BloombergNEF’s survey put the average lithium-ion pack at \$115 a kilowatt-hour in 2024, down 20% in a year, and at \$108 in 2025, with packs for battery-electric cars below \$100 and Chinese packs at \$84 (BloombergNEF, 2024, 2025); new-energy vehicles were 47.9% of China’s vehicle sales in 2025, 16.5 million of them (CnEVPost, 2026). IRENA put the global levelized generation cost of utility-scale solar at 4.3 cents a kilowatt-hour in 2024, about 90% below 2010, against 7.3 cents for coal and 8.5 for gas, and 91% of the renewable capacity commissioned that year was cheaper than the cheapest new fossil plant (International Renewable Energy Agency (IRENA), 2025); renewables were 85.6% of the generating capacity the world added in 2025 (TaiyangNews, 2026). A 60-watt-equivalent LED bulb that cost more than \$50 in 2007 cost \$5.90 in 2017, and LEDs went from 8% of installed US lighting in 2015 to about 48% in 2020 (U.S. Department of Energy, Solid-State Lighting Program, 2020; inside.lighting, 2024). These price and adoption series establish changing economics; they do not by themselves identify damage to comparable holdouts. The question for companies is who owns the curve and how much of its saving customers and competing suppliers take. Sutton’s argument about endogenous sunk costs supplies one route to concentration when maintaining quality requires continuing fixed investment (Sutton, 1991). When many producers drive a commodity’s curve together, margins can instead migrate to installation, finance and integration. In either case, an urgent purchase can coexist with a flat or falling supplier price ceiling.

2.4 Expectations ratchet to what most rivals offer

A fourth source of pressure is the customer. Kano separated features that delight when present from features that are expected and punish when absent, and features move from the first group to the second as rivals supply them (Kano et al., 1984); Kahneman and Tversky’s reference dependence is the mechanism, since customers judge an offer against the reference the market has set and weigh losses against it more than gains (Kahneman and Tversky, 1979). Appendix H.3 adds an expectation that follows adoption upward and can persist if adoption recedes. Business transferred to rivals belongs in the contest channel; demand lost to the whole field requires a separate loss term. Neither is established by the base network term alone. If matching the expectation adds cost without enough producer surplus, full adoption can leave firms worse off even though dropping the feature hurts. A wasteful bandwagon additionally requires that adopting alone did not pay; an initially profitable feature can instead produce the arms-race pattern. Customers may gain in either

case. Delivery speed is the running example. Amazon reported more than 13 billion Prime items delivered the same or next day worldwide in 2025; its US volume grew 30% from 2024 ([Amazon.com, Inc., 2026e](#)); Walmart’s chief executive promised in April 2025 to reach 95% of the US population in under three hours by the end of the year, and its February 2026 results reported sales through its expedited delivery channels up more than 50% ([TheStreet, via AOL, 2025](#); [Walmart Inc., 2026](#)); Target extended next-day delivery from 20% to 54% of the population in September 2025, an effort, in the wire report’s words, to “narrow gap with rivals” ([D’Innocenzio, 2025](#)). The surveys discussed in Section 3.2 support the role of delivery expectations, although they do not identify the causal loss suffered by slower retailers. Another rapid response took three months. Lyft let its drivers cash out daily in December 2015; Uber announced the same in March 2016, “a little behind the times” in Fortune’s words, and had paid out \$1.3 billion that way within a year ([Reisinger, 2016](#); [Yeung, 2017](#)). For companies, the consequences are timing and price. A discretionary budget can become difficult to defer when the feature becomes expected, often after platforms have begun to bundle it. The relevant window is the interval before lost conversion, retention or renewals exhausts the buyer’s tolerance; a calendar announcement is unnecessary. An adoption wave can also reflect an informational cascade. Bikhchandani, Hirshleifer and Welch showed that rational actors who watch each other’s choices can adopt in a wave that carries no information about payoffs, and such waves end as abruptly as they start ([Bikhchandani et al., 1992](#)); evidence for an economic ratchet requires holdout losses attributable to the changed expectation, which is the row of Table 10 on holdouts losing share.

2.5 Gatekeepers are not only regulators

Private requirements remain requirements: changing the identity of the gatekeeper does not establish competitive forcing. Their distinct evidence is enforcement and the gatekeeper’s incentive to continue it, rather than a measured market-price or demand response.

Exclusion, $\varepsilon\Phi(x)$, is a penalty on holdouts that no rival collects, and nothing in the term says the penalty comes from the state. A browser (Chrome 68), a card network (the EMV liability shift), a dominant customer (Walmart’s supplier mandates), a platform (Apple’s tracking prompt), an insurer (controls required at renewal), a standards body (the CA/Browser Forum’s certificate lifetimes) and the customers who send security questionnaires (SOC 2) each set a penalty, a date and a steepness, and more than half the mandates in Section 3 are of this kind. Walmart’s two tag mandates provide a useful comparison. On June 11, 2003 it told its 100 largest suppliers to put radio tags on every pallet and case by January 1, 2005, with the rest to follow by 2006; a tag cost 40 cents to a dollar, the supplier paid for it, and the expected savings, which Bernstein put at \$8.4 billion a year, would accrue to Walmart ([Violino, 2003](#); [Roberti, 2003](#)). The January 2006 deadline passed with few suppliers compliant, the program was “fizzling” by 2007, and the penalty Walmart eventually set, \$2 a pallet at Sam’s Club from February 2008, was cut to 12 cents within a year ([Supply Chain Digest, 2009](#)). The January 2022 expansion of the item-level program, already used for apparel, required tags on additional home, electronics and toy categories arriving in stores by September 2, 2022, and on nine more categories by February 1, 2024; by then a tag cost three to five cents, and the inventory accuracy it bought, from about 60% to 99% by SKU in Auburn University’s measurements, was the supplier’s gain as well as Walmart’s ([Swedberg, 2022](#); [CYBRA Corporation, 2026](#); [R.R. Donnelley & Sons, 2024](#)). The later mandate coincided with substantial growth for tag suppliers. The industry shipped 52.9 billion tag chips in 2024, 77% of them for apparel; Impinj’s revenue went from \$190 million in 2021 to \$366 million in 2024; Avery Dennison’s label business reached about \$850 million growing 18% a year, and the two companies announced a fresh-food expansion in October 2025 ([Taylor, 2025](#); [Impinj, Inc., 2026](#); [edgar.tools, 2026](#); [Avery Dennison Corporation, 2024](#); [Hickey, 2025](#)). Apple’s prompt is the platform version. From April

26, 2021 every iOS app had to ask permission to track; Meta told investors the change would cost it about \$10 billion in 2022, and AppLovin had agreed to pay about \$1 billion for Adjust, a measurement company, in the February before the rule took effect (9to5Mac, 2021; Meta Platforms, Inc., 2022; AppLovin Corporation, 2021). The gatekeeper paid nothing, its rivals paid a great deal, and fines of €150 million in France and €98.6 million in Italy have since been levied without the rule changing (Autorité de la concurrence, 2025; PPC Land, 2025). A private gatekeeper’s incentive to enforce matters for timing. Its own gain from compliance can strengthen the credibility of a penalty, although implementation costs, contracts and legal constraints also matter. The 2003 mandate asked suppliers to pay for a benefit that was Walmart’s, and Walmart did not enforce it; the 2022 mandate came with cheap tags, a measured gain on both sides and enforcement. The signal to read is the gatekeeper’s own gain from compliance, and the risk to price is the one the platform row of Table 11 names: a gatekeeper that supplies the capability itself can displace vendors unless they offer additional value.

2.6 Who keeps the budget: appropriability

Teece’s theory of appropriability explains why the innovator may not capture the returns. An innovator profits when imitation is hard, or when it owns the complementary assets the innovation needs to reach the customer: distribution, the workflow, the installed base, the venue (Teece, 1986). When a specified capability is easy to imitate or supplied through an open standard, the complementary asset may retain more value: the accounting suite above the free exchange (TeamSystem), the operating system above the age check, the rails above the stablecoin, the exchange above the fiber, the bank above the data connection. Sutton identifies a further route to concentration: continuing fixed investment in quality can prevent a growing market from fragmenting among small suppliers (Sutton, 1991). That mechanism is consistent with some consolidation in defense and certificate operations, but does not establish its cause. The opportunities in Section 6 therefore require either a capability that remains difficult to copy or complementary assets the entrant can retain: customer contracts, distribution, data, field service or financing capacity. Without either advantage, a startup risks selling a capability that competitors or platforms will provide at no additional charge.

Table 3: Six sources of pressure. Each requires its own evidence about exposure, timing and capture; the productivity row uses the explicit extension in Section 1.3.

Source	In the model	Timing evidence	Possible capture	Running examples
Rivals, over a fixed prize	$\alpha/(1 + (r - 1)x)$; holdouts' loss grows with x	The first rival's move; renewed by each generation; ended by a ban	The venue that hosts the race; the company selling the next increment	Spread Networks and the exchanges; LZR; AI answers
Counter-parties, through a network	βx ; a threshold if early gains do not cover costs	Whoever crosses the tipping point: a subsidized side, a switched default, a large joiner	The operator that retains control of access	Plaid; containers; stablecoin rails; rent-setting software
A cost curve and price response	$\gamma - c$ pays; common loss $h(x)$ erodes operating profit	Saving, observed pass-through and the holdout's margin crossing, less installation time	The curve's owner, complementary suppliers or customers through lower prices	Fulfillment automation; hybrid corn; batteries; solar; LEDs
Customers' expectations	History-dependent expectations; transfers and lost demand must be separated	The share of rivals already offering the feature	The platform that bundles it for many sellers at once	Free and same-day delivery; pay-later at checkout; mobile deposit
A private gatekeeper (platform, dominant customer, card network, insurer, standards body)	$\varepsilon\Phi(x)$; credibility depends on enforcement incentives	The policy's date, credible when compliance pays the gatekeeper	The gatekeeper, if it can supply the capability; otherwise the companies it tolerates	Chrome 68 and Let's Encrypt; EMV and Verifone; SOC 2 and Vanta; Walmart's tags; Apple's prompt
A regulator or court	$\varepsilon\Phi(x)$ is a proxy; a calendar rule needs time-dependent enforcement	The final text and enforcement; distinguish migration from recurring work	Required service providers; bundles can limit commodity fees	<i>Wayfair</i> and Avalara; SdI and TeamSystem; age assurance and the OS age signals

3 Economic pressure and the company’s business

The cases begin with rivals competing over position, then follow productive tools into changing customer expectations, networks and defenses. They ask what happens to the holdout and who retains the spending. Private requirements and public mandates supply comparisons with a more explicit enforcement mechanism. None of these channels establishes an attractive entry point for a new supplier on its own. Table 4, at the end of the section, records the company outcomes.

3.1 Arms races create repeated demand for an edge

In a pure positional contest, rivals spend to protect a relative advantage that disappears from the field’s total payoff. Purchases can recur with each generation, but supplier profit depends on who controls the next improvement and whether buyers can exit or defend. Bitcoin mining and frontier AI combine this pressure with security or productive demand; Section 4 separates those channels. A positional supplier faces four hazards: the edge erodes with x , the company selling the edge is itself in a race, the venue that hosts the contest can sell the race directly, and a regulator can end it.

Spread Networks lit a private fiber route from Chicago to the New Jersey data centers in June 2010, 825 miles at a round trip of 13.33 milliseconds, at a cost Forbes estimated at \$300 million (Matsumoto, 2010; Steiner, 2010). Trading firms paid because three milliseconds was the difference between the fill and the miss; the edge was pure α . McKay Brothers reported 8.27 milliseconds on its microwave route in February 2013 (McKay Brothers, 2013); when Zayo agreed to buy Spread in November 2017 for \$127 million fewer than a quarter of its 432 fibers were lit (Zayo Group Holdings, Inc., 2018; Hardy, 2017). Those disclosures do not establish customer migration or Spread’s lifetime return. They show competing generations of speed alongside a different position in the market: exchanges sell colocation and proprietary data to both sides of the race, revenue that Budish, Lee and Shim put at \$675 to 790 million in 2015 and projected at \$1.3 to 1.5 billion for 2022 (Budish et al., 2019); ICE’s exchange data and connectivity line alone was \$1.03 billion in 2025 (Intercontinental Exchange, Inc., 2026). IEX’s speed bump, approved by the SEC in June 2016, is the counter to this race, and it illustrates the limit of counters: it protected IEX’s own customers and left the race running everywhere else (U.S. Securities and Exchange Commission, 2016).

The LZR Racer shows how a regulator ends one. Speedo launched the suit in February 2008 at \$550; swimmers wearing it set 62 world records by the Beijing Games, and 108 fell in 2008 (Sports Illustrated, 2008; Associated Press, 2009). Arena’s polyurethane X-Glide beat it at the 2009 world championships, and on July 31, 2009 FINA’s Bureau confirmed January 1, 2010 as the date for the new textile-only requirements (FINA, 2009). Speedo’s parent, Pentland, is private, so the company’s revenue through the episode is not public; TYR’s antitrust suit against Speedo was dismissed in 2010, and the non-textile category ended by rule less than two years after the launch (Katten Muchin Rosenman LLP, 2010). An investment in this market depends on the supplier’s ability to keep improving, its relationship with the venue and the risk of a ban.

Some suppliers sustain revenue by combining a competitive advantage with services that are costly to replace. The Bloomberg terminal reached the market in December 1982 with twenty units at Merrill Lynch (McCracken, 2015); its combination of data, tools and communication has sustained paid demand for more than forty years. Subscribers passed 325,000 in 2015 and Bloomberg now says more than 350,000 (McCracken, 2015; BlueGamma, 2026); Burton-Taylor put Bloomberg’s share of a \$27.5 billion market-data industry at 33% in 2016, and the industry’s revenue at a record \$49.2 billion in 2025 (Burton-Taylor International Consulting, 2017, 2026); and the reported single-terminal price rose to \$31,980 a year from January 2025 (NeuGroup, 2024), with the customers

saying in 2000 what they say now: “We are upset about the increase, but what can we do about it?” (WatersTechnology, 2000). The terminal’s communication network and embedded workflows help explain its persistence: a departing desk must replace more than a faster data feed. That is a capture mechanism to test, not proof that customers cannot switch. Semrush kept the search-placement race for seventeen years. Its revenue grew from \$92 million in 2019 to \$444 million in 2025 and Adobe paid \$1.9 billion for it in April 2026, four times ARR (SEMrush Holdings, Inc., 2021; Semrush Holdings, Inc., 2026a; Adobe Inc. and Semrush Holdings, Inc., 2025; Adobe, 2026); net revenue retention had fallen from 120% to 104% over the same years, indicating slower expansion within the customer base without identifying its cause, and the race itself was moving to AI answers, the opportunity discussed in Section 6.5. The Trade Desk kept the programmatic bidding race for a decade, with customer retention above 95% every year and \$2.9 billion of revenue on \$13.4 billion of spend in 2025, in a market where programmatic is now more than 90% of US display advertising (The Trade Desk, Inc., 2026c,b; eMarketer, 2026), while competing with venues that also sell buying tools: its filing names Google and Amazon among its competitors. Revenue growth fell from 26% in 2024 to 3% in the second quarter of 2026, and the shares are down about 80% from their August 2025 level (The Trade Desk, Inc., 2026a; 24/7 Wall St., 2026). Those outcomes do not isolate the effect of venue competition. Retail media, in Section 2.1, provides another example of venues selling placement directly.

The dreadnought race, whose falling ceiling Section 1 uses, brought orders to the shipyards without insulating them from the later contraction. Britain’s annual naval expenditure rose from £31.5 million to £50 million, and the yards sold the race abroad as well: Armstrong Whitworth sold the dreadnought it was building for Brazil to the Ottoman navy for £2.75 million in December 1913 (Cartwright, 2025; Wikipedia contributors, 2026a). After the Armistice Vickers reported losses of more than £1.25 million, Armstrong lost £626,000 in the first eleven months of 1926, and the two merged their armaments businesses in 1927 (Wikipedia contributors, 2026c; Grace’s Guide to British Industrial History, 2026). Major League Baseball centralized the purchase of tracking systems for all thirty clubs (TrackMan’s radar in 2015 and Hawk-Eye’s cameras in 2020) and distributes the data to every club, so the sensor companies sold one contract each and the edge moved to the analysts (Major League Baseball, 2026; Ballpark Digest, 2020). What separates Bloomberg from Spread Networks is what the company owns besides the edge: the network, the workflow, or an exit to an incumbent before the venue moves.

3.2 A productive tool can become necessary as the market changes

When standalone value exceeds adoption cost and the value of waiting, the gain from adopting is already positive at $x = 0$. With direct value alone, that gain stays positive: the dominant pattern of Table 1. Diffusion then depends on budgets, information and installation capacity rather than a tipping point. Customer-support AI illustrates the direct benefit. In the staggered-rollout field study cited in Section 1, 5,172 agents at one software firm resolved 15% more issues per hour with an assistant, the least experienced gaining most (Brynjolfsson et al., 2025), and the gain needed no other firm to adopt. Economic forcing requires a further link: rivals’ lower prices, improved service or faster delivery must damage the firms that keep the old process. The fulfillment case in Section 7.2 makes the price channel explicit.

Broad demand for a useful product does not establish an advantage for any particular supplier. The budget that direct value creates is contestable by every company with access to the same models, and by the incumbent that already owns the workflow. Intercom’s Fin charges \$0.99 per resolution (Intercom, n.d.); Google folded Gemini into Workspace in January 2025 and Microsoft put Copilot into consumer Microsoft 365 the same week (Google Workspace, 2025; Rognier, 2025); and the standalone leaders are priced as if the market were already theirs: Sierra reported more

than \$150 million of ARR entering its third year, and Decagon raised \$250 million at a \$4.5 billion valuation in January 2026, triple its price of six months earlier (Sierra, 2026; Decagon, 2026). A differentiated product, distribution or installed workflow may protect it. The additional forcing question is whether adopters change prices, response times or service levels enough to damage comparable holdouts. Section 6 sets aside the broad support category on entry and capture grounds, while retaining that economic test for narrower markets.

Customer expectations supply a second route: rivals raise what customers expect, and a holdout loses demand to a standard it did not set. Section 2.4 gives the mechanism and Appendix H.3 models it as a ratchet, an expectation that follows adoption upward and stays there if adoption recedes. Amazon Prime is the case, and it has run for twenty years. It launched on February 2, 2005 at \$79 a year for unlimited two-day shipping (Amazon.com, Inc., 2005); in April 2019 Amazon put about \$800 million into a single quarter to make one-day delivery the default (The Motley Fool, 2019a); the same-day volumes of Section 2.4 are where that has reached, with 30-minute delivery following in May 2026 (Amazon.com, Inc., 2026c,b). The customers moved with it: 62% of US holiday shoppers defined fast shipping as two days or less in 2018, up from 54% a year earlier (Chain Store Age, 2018), and in 2025 62% said faster shipping decides between retailers at the same price (Radial, Inc., 2025). Retailers serving customers with those expectations faced pressure to improve fulfillment; existing logistics networks supplied one way to respond. Walmart’s own fulfillment service handled more than 40% of its marketplace sales by late 2024 and had 52% of its sellers by February 2026 (Walmart Inc., 2024; Investing.com, 2026); Target bought Shipt for \$550 million in December 2017 (Target Corporation, 2017); and Shopify, which paid \$2.1 billion for Deliverr in May 2022 because fast shipping was “now table stakes,” sold the business to Flexport thirteen months later for a 13% stake and booked \$1.7 billion of one-time charges (Shopify Inc., 2022, 2023a,b). The independents that sell Prime-like fulfillment to everyone else raised at \$1 billion (ShipBob, 2021) and \$3 billion (Stord, May 2026) (ShipBob, 2021; TechCrunch, 2026): a real company market, made by a rival’s service level rather than by a rule.

AI features in SaaS provide a possible parallel, with less evidence yet of a common customer requirement. Bessemer found AI features on the roadmaps of 96% of its portfolio companies in 2023 and wrote in August 2025 that “there is no cloud without AI anymore” (Crunchbase News, 2023; Bessemer Venture Partners, 2025); the bundling above is what meeting the expectation cost the incumbents; and some existing businesses reported weaker demand as customers turned to AI tools. Chegg told investors on May 1, 2023 that students’ interest in ChatGPT was cutting its new-customer growth, and lost 48% of its value the next day (Chegg, Inc., 2023; Associated Press, 2023); new questions on Stack Overflow fell from 1.34 million in 2022 to 109,000 in 2025 (PPC Land, 2026). These observations show changing offers and demand, but do not establish that every SaaS buyer must purchase an AI feature. Even where a requirement develops, suppliers of the capability remain in competition. What the ratchet adds is a signal, in the customer surveys and holdout revenues of Section 5, that a budget has become compulsory without any gatekeeper saying so.

3.3 Coordination goods can pay the operator

A network good with little standalone value can be worth little to the first adopter and much more to later ones, leaving early adopters below the profitability threshold shown in Figure 1: below the tipping point nobody moves, and whoever funds the crossing does so at a loss. A tipping point establishes a growing gain from connection. Economic pressure on holdouts additionally requires lost access or demand, such as customers leaving a merchant that cannot accept the standard. The question for companies is who owns the network and who can reprice its access.

Plaid built the connection layer between US banks and the apps that read their data, a business with network effects represented by βx : each bank added made every app more valuable and each

app made every bank connection more worth having. Visa agreed to buy it for \$5.3 billion in January 2020, the Justice Department sued, and the deal was abandoned in January 2021; Plaid then raised at \$13.4 billion in April 2021 and at \$6.1 billion in April 2025 ([Visa Inc., 2020](#); [U.S. Department of Justice, 2021](#); [Beltran, 2025](#)). That valuation decline preceded the bank-fee announcements below and cannot be attributed to them. JPMorgan told aggregators in July 2025 that they would pay for access to customer data, fees it said would total hundreds of millions of dollars, and the companies announced a renewed access agreement that September, without disclosing its financial terms ([Bloomberg Law, 2025a](#); [JPMorganChase, 2025](#)); the Consumer Financial Protection Bureau’s 2024 open-banking rule, which would have made that access a right, was enjoined in October 2025 while the Bureau reconsidered it ([ABA Banking Journal, 2025](#)). Demand for bank connections did not give Plaid control over access fees. The banks’ ability to set terms creates a capture risk; the cited announcements do not quantify its effect on Plaid’s margins.

Containerization tells the same story over sixty years. Malcom McLean’s container service began with the *Ideal X* voyage on April 26, 1956 and sold out to R.J. Reynolds in 1969; the box became an ISO standard from 1968; and the lines that adopted it destroyed more than \$100 billion of shareholder value over the twenty years to 2017, earning less than their cost of capital in most of them ([Rodrigue, 2026](#); [Wackett, 2018](#)). The standard created benefits beyond the pioneering shipping line, including lower transport costs for shippers. Those gains need not appear as durable profits for carriers. For a supplier, the investment case depends on control over access to the network or a credible acquisition route.

3.4 Counters are the arms race’s second market

Some competitive technologies create losses that make protection valuable. The model treats a counter as a third choice within a defined rival field ([Appendix F](#)); external threats require a separate exposure model. Greater loss can raise the willingness-to-pay ceiling for effective protection, but does not guarantee a higher realized price. It also has a failure mode the attack side lacks: the venue can defend natively, and then the counter is a feature of the platform rather than a company.

Distributed denial-of-service protection is the historical case. Attack sizes reset their record almost every year, to 31.4 terabits per second in late 2025 ([Cloudflare, 2026](#)); Akamai bought Prolexic, the largest standalone defender, for about \$370 million in February 2014, and Cloudflare, with mitigation at its core, grew from \$287 million of revenue in 2019 to \$2.17 billion in 2025 ([Akamai Technologies, 2014](#); [Cloudflare, Inc., 2026](#)). E-commerce fraud prevention followed the same path: Riskified listed in July 2021 at \$21 a share, and five years later reports \$345 million of revenue and names generative AI and agentic commerce as the next things changing fraud ([Riskified Ltd., 2021, 2026](#)). Cybersecurity as a whole is the largest counter market in the taxonomy. Gartner expects \$244 billion of information-security spending in 2026 ([Gartner, Inc., 2026](#)), and CrowdStrike, which sells endpoint defense, grew its recurring revenue from \$313 million in January 2019 to \$5.84 billion in July 2026 ([CrowdStrike Holdings, Inc., 2021, 2026](#)). Its July 2024 update crashed about 8.5 million Windows machines ([TechCrunch, 2024](#)), and the year still ended with recurring revenue up 23% and gross retention of 97% ([CrowdStrike Holdings, Inc., 2025](#)): persistent threat exposure can sustain demand through a service failure. Customers can still switch suppliers, and retention alone does not identify pricing power.

Identity verification remains a live benchmark in [Section 6](#): attacks can create renewed demand for defense, while requirements sustain some of the verification flows. Attackers have industrialized the deepfake: in Entrust’s 2026 Identity Fraud Report, published in November 2025 on more than a billion verification events, deepfakes were one in five biometric fraud attempts, deepfaked selfies rose 58% in a year and injection attacks 40% ([Entrust, 2025](#)). At the same time regulators mandate the check: the United Kingdom’s age-assurance rules took effect on July 25, 2025 and produced 69

million age checks in the second half of the year across a 32-service sample, 23 times the prior six months (Ofcom, 2026); Australia’s under-16 ban took effect on December 10, 2025 and 4.7 million accounts were removed within days (The Register, 2026); and Texas’s app-store law came into force after the Supreme Court declined to block it in July 2026 (Wiley Rein LLP, 2026). A counter that leaves the compliance requirement unmet cannot disarm it (Corollary 4): a platform that would rather not verify anyone still has to run a check, and the check is the flow a detection company sells into. Operating-system signals and platforms’ own estimation can reduce demand for paid checks. An evolving attack can preserve a separate detection budget where those controls remain inadequate. Socure reported \$364 million of recurring revenue growing 63% and raised at \$5.2 billion in August 2026; Persona raised \$200 million at \$2 billion in April 2025; Yoti’s revenue grew 62% to £29 million in 2025 (Biometric Update, 2026c,d; Persona, 2025; Borak, 2026). The risk is the one every counter carries: Apple and Google now ship age signals at the operating-system level, and Discord said at its September 2026 rollout that more than 90% of its users would never face an active check (Apple Inc., 2025; Biometric Update, 2026a).

3.5 Low supply costs can eliminate revenue from the required capability

A mandate can achieve widespread adoption without creating substantial revenue for suppliers. HTTPS is the case (Appendix J fits it as a mix of dominant and tipping adoption). Google made encryption a ranking signal in August 2014 and announced on February 8, 2018 that Chrome 68 would mark every HTTP page “Not secure” from July; by then 68% of Chrome traffic on Android and Windows was already encrypted (Bahajji and Illyes, 2014; Schechter, 2018). Chrome 68 shipped on July 24, 2018 (Google, 2018): a gatekeeper’s penalty with a date, applied to every website.

Free certificates limited the revenue available from basic issuance. Let’s Encrypt opened its public beta on December 3, 2015 and issued domain-validated certificates free, automatically, to anyone; by 2025 it served 762 million websites and had issued more than seven billion certificates in its first decade (Aas, 2015; Internet Security Research Group, 2025). Free issuance constrained prices for basic certificates; it did not eliminate every paid certificate service. Symantec agreed on August 2, 2017 to sell its certificate authority to DigiCert for \$950 million in cash and a 30% stake, as Chrome prepared to stop trusting Symantec’s certificates altogether (Symantec Corporation, 2017; O’Brien et al., 2017). A distinct paid market developed above basic issuance. Enterprises with more than 100,000 certificates pay for inventory, rotation and outage prevention: 86% of 450 organizations surveyed in 2025 had suffered a certificate-related outage in the past year, one in ten every week (Keyfactor, 2025). CyberArk bought Venafi, the largest certificate-lifecycle company, for \$1.54 billion in 2024 on about \$150 million of recurring revenue, and Palo Alto Networks announced its \$25 billion CyberArk deal in July 2025 citing machine identities (CyberArk Software Ltd., 2024; Palo Alto Networks, Inc., 2025). When a nonprofit or platform supplies an equivalent capability free, exclusion need not create a separately chargeable budget for it. Complex certificate operations can still support paid services. Shorter browser certificate lifetimes increase demand for that work (Section 6.5).

3.6 Mandates create migration waves and recurring work

A mandate often supplies an announced enforcement date, making its purchase window easier to observe than a competitive threshold. Expected enforcement can raise a holdout’s willingness to pay and compress the sales window. Once most firms comply, new-installation demand can fall even though the penalty for dropping compliance remains. The base model captures that continuing incentive; a migration sales wave additionally depends on the installed base and replacement cycle. Recurring filings, monitoring or new transactions can sustain paid work after the initial deadline.

Sales tax recurs every month. *Wayfair* let states require collection from sellers with no physical presence, and 33 states and the District of Columbia had adopted economic-nexus rules within eight months ([The Motley Fool, 2019b](#)). Avalara’s revenue grew from \$213.2 million in 2017 to \$272.1 million in 2018, \$382.4 million in 2019, \$500.6 million in 2020 and \$699.0 million in 2021 ([Avalara, Inc., 2018, 2020, 2022](#)); Vista agreed to pay \$8.4 billion for it in August 2022, and the company filed confidentially for a second listing in July 2025 ([Vista Equity Partners, 2022](#); [Srivastava, 2025](#)). Its chief executive was careful about the cause: “We filed for our IPO before the ruling, and we never counted on it. However, it has given us a tailwind” ([Prete and Baltz, 2019](#)). SOC 2 recurs every year, and it is enforced by customers rather than by a regulator: a security questionnaire that asks for the report is a gatekeeper with a budget attached. Vanta, founded in 2018, crossed \$100 million of recurring revenue in the fiscal year to January 2024 and \$300 million by April 2026, with 16,000 customers and a last valuation of \$4.15 billion ([Vanta, 2024](#); [Lazarus, 2026](#); [Bloomberg Law, 2025b](#)). Drata, founded in 2020, passed \$100 million in February 2025 ([Drata, 2025](#)).

EMV illustrates a migration wave under a continuing liability rule. From October 1, 2015, US merchants accepting chip cards without chip-capable terminals could bear counterfeit-fraud losses. Upgrading avoided that exposure; it did not require every merchant to buy a new terminal, nor did the liability incentive expire after the upgrade ([Visa, 2016](#)). Verifone’s North American revenue rose from \$526 million in fiscal 2014 to \$792 million in 2015 and \$804 million in 2016, then fell to \$630 million in 2017 ([VeriFone Systems, Inc., 2015, 2016, 2017a](#)). Its 2017 annual report says why: “revenue increases due to the adoption of standards such as EMV in the United States have in the past slowed and declined” ([VeriFone Systems, Inc., 2017b](#)). Francisco Partners took the company private in April 2018 for \$23.04 a share, about \$3.4 billion with debt, a 54% premium to a stock that had fallen to \$15 ([VeriFone Systems, Inc., 2018](#)). Y2K is the same shape at a larger scale. Gartner put worldwide remediation spending at \$300 to \$600 billion ([Edward W. Kelley, 1998](#)); Keane’s Y2K revenue went from \$369.5 million in 1998 to \$206.1 million in 1999 to \$5.4 million in 2000, and Cognizant’s from 44% of revenue to 0.4% over the same three years ([Keane, Inc., 2001](#); [Cognizant Technology Solutions Corp., 2001](#)). Post-quantum migration (Section 6.5) also creates migration work, but its deadlines are less uniform.

Private requirements also depend on the gatekeeper’s enforcement incentives. Walmart’s two RFID programs and Apple’s tracking prompt, discussed in Section 2.5, show why an announced policy, its enforcement and its effect on supplier revenue need separate evidence.

Italy is the precedent for the e-invoicing case in Section 7, and its warning is about bundling rather than dates. Italy introduced mandatory business-to-business e-invoicing for in-scope firms through the state’s own exchange, the Sistema di Interscambio, from January 1, 2019 ([European Commission, 2026](#)). The exchange is free. Free public transmission limited that fee opportunity; accounting software above the exchange remained a paid market. TeamSystem, the largest Italian company selling that software, reported €1.06 billion of revenue, €476 million of adjusted EBITDA and a €37.7 million consolidated net loss in 2025 ([TeamSystem Holdco S.p.A., 2026](#)). A September 2026 press report put the valuation in minority-stake transactions with Francisco Partners and KKR at €8 to 10 billion ([Private Equity Wire, 2026](#)). A French transmission startup would compete with the free and bundled offers described in Section 6.5.

Table 4: Who kept the budget. One row per running example, with a demand mechanism and a reported company outcome. Receipts and transaction values are not estimates of profits caused by the mechanism.

Example	Channel	Demand mechanism	Who paid	Who kept the budget, and what ended it
AI support assistants	Direct	ChatGPT demonstration, Nov. 2022	Support operations, from their own labor line	Model providers and workflow incumbents; bundling
HTTPS certificates	Exclusion	Chrome 68 “Not secure,” July 2018	Website operators	Free issuance constrains basic certificate fees; lifecycle companies a layer up (Venafi, \$1.54bn)
Bank data connections	Network	App demand; the 2024 open-banking rule, enjoined in 2025	Fintech apps	Plaid as operator; banks can charge for access, but its Sept. 2025 JPMorgan agreement’s financial terms were undisclosed
Sales-tax compliance	Exclusion, recurring	<i>Wayfair</i> , June 2018	Remote sellers	Avalara (\$8.4bn take-private, 2022); recurring compliance demand continues
SOC 2	Exclusion, recurring	Customer questionnaires	SaaS companies	Vanta (\$300m ARR, 2026), Drata; the risk is a free tier
EMV terminals	Exclusion; migration wave	Liability shift, Oct. 2015	US merchants	Verifone, during the upgrade wave (\$3.4bn take-private, 2018)
Italian e-invoicing	Exclusion	SdI mandate, Jan. 2019	In-scope VAT-registered firms	The accounting suite (TeamSystem, reported €8–10bn valuation), above the free public exchange
Low-latency fiber	Contest	Spread Networks, 2010	Trading firms	The next company (microwave), then the venue (exchange colocation); Spread sold for \$127m
Polyurethane suits	Contest	LZR Racer, Feb. 2008	Elite swimmers	Speedo for one season, rivals the next; FINA’s ban ended it, Jan. 2010
Market-data terminals	Contest and network	Bloomberg terminal, Dec. 1982; periodic price increases	Subscribing trading desks	Bloomberg, for forty years (more than 350,000 seats at about \$32,000); the terminal is also the network
Search placement	Contest	Google’s rankings; AI answers from 2024	Sites competing for organic traffic	Semrush (\$444m revenue, 2025; sold to Adobe for \$1.9bn, 2026), Ahrefs; the race moved to AI answers
Programmatic ad bidding	Contest	Real-time bidding; 90% of US display by 2025	Programmatic advertisers	The Trade Desk for a decade (\$2.9bn revenue, 2025), facing venue competitors Google and Amazon; later growth slowed
Dreadnoughts	Contest	HMS <i>Dreadnought</i> , 1906; the 1909 eight-ship programme	Navies, from estimates rising from £31.5m to £50m	Vickers and Armstrong while the race ran; losses after the Armistice and a merger in 1927

continued on the next page

Table 4: Who kept the budget. One row per running example, with the demand mechanism and reported company outcome (continued).

Example	Channel	Demand mechanism	Who paid	Who kept the budget, and what ended it
AI compute	Productive demand and contest	Repeated model and hardware generations	Labs, clouds and application providers	NVIDIA: \$193.7bn Data Center revenue in FY2026; corporate gross margin 71.1%, not a segment margin. Utilization and replacement risk remain downstream (Section 4)
Bitcoin mining	Reward contest and security	Relative hashpower and equipment efficiency	Miners, through hardware and power bills	Suppliers receive spending; difficulty, energy costs and the dollar reward determine miners' returns. Ethereum's Merge ended a different network's mining market (Section 4)
Endpoint security	Counter	Ransomware and breaches; Gartner forecasts \$244bn spending in 2026	Enterprises buying protection	CrowdStrike (\$5.8bn ARR, July 2026; 97% gross retention in FY2025); the risk is the platform's own defense
Identity verification	Counter and mandate	Deepfakes; age laws 2025–26	Platforms, from fines and fraud avoided	Socure, Persona, Yoti; the risk is OS-level age signals
RFID tags at Walmart's suppliers	Exclusion, customer as gatekeeper	Walmart's letters: pallets by Jan. 2005 (lapsed); items by Sept. 2, 2022	Suppliers in the named categories	Impinj (\$366m revenue, 2024), Avery Dennison; the 2003 mandate failed to achieve broad compliance
Mobile ad tracking	Exclusion, platform as gatekeeper	Apple's prompt, April 26, 2021	Advertisers and ad networks (Meta: about \$10bn in 2022)	The platform; measurement re-tooled (AppLovin bought Adjust, \$1bn); EU fines have not reversed the tracking rule
Rent-setting software	Coordination among rivals	DOJ alleges shared pricing data; suit Aug. 23, 2024	Landlords, then renters	RealPage (\$10.2bn, 2021) faced restrictions on its pricing mechanism in the proposed Nov. 2025 settlement; customer retention is unverified
Retail media	Contest; the venue sells the race	Amazon's ad line; more than 200 networks by 2024	Consumer brands	The retailers (Amazon \$56bn, 2024; Walmart \$6.4bn, FY2026); Criteo's 2025 retail-media revenue grew 2%
Battery packs and solar	Cost curve	Packs at \$115/kWh (2024) and \$108 (2025); solar at 4.3 cents/kWh	Automakers, utilities	Low-cost producers and customers may capture the saving; holdout losses require evidence of substitution and price response
Delivery speed	Expectations	Prime, Feb. 2005; one-day by default, 2019; 13bn same- or next-day items, 2025	Rival retailers (Walmart, Target), from fulfillment capex	Amazon and Walmart as operators; Shopify paid \$2.1bn for Deliverr and sold for a 13% stake; Stord and ShipBob sell to the rest

4 AI and crypto: useful products, competing platforms, expensive races

AI and crypto can create substantial value while competition reduces suppliers’ returns. A support worker using an assistant, a trader using Hyperliquid, a developer choosing Ethereum and a laboratory buying another GPU cluster are making different investments. Their gains cannot be inferred from one industry’s adoption curve. The comparison has to follow the buyer’s alternative, the source of pressure and the recipient of the payment.

The strongest common mechanism is competition to supply a base on which others build. Bitcoin, Ethereum, Solana and Hyperliquid offer different combinations of settlement, execution and liquidity; OpenAI, Anthropic, Google, Meta and DeepSeek offer models, distribution and tools. Applications can strengthen those bases while also acquiring the customers and retaining the margin. Meanwhile, miners and model developers buy infrastructure whose suppliers can earn before the ultimate economics are settled. The cases below test those mechanisms rather than treating either sector as a single regime.

4.1 Payoffs for users, platforms and infrastructure buyers

For a user, γ is the value of completing a task or accessing a service relative to the existing alternative, after holding other users’ adoption fixed. A network term belongs where additional counterparties or complementary applications improve that service. For a platform, the question is whether another application attracts users who then attract other applications, and whether the platform can charge for that connection. For an infrastructure buyer, the question is whether more capacity serves paying demand, protects a competitive position, or does both. Table 5 keeps these decisions separate.

Table 5: The paired comparison. Channels describe a specified decision, not an entire industry.

Decision	AI	Crypto	Evidence that identifies the mechanism
Use a service	Worker or firm buys a better task outcome; rivalry can later lower prices or raise expectations	Holder, trader, merchant or recipient buys asset exposure, execution, settlement or dollar access	Task quality and total cost against the actual alternative; losses incurred by comparable holdouts
Build on a base	Developer chooses a model, API, host and distribution channel	Developer chooses a chain, liquidity pool, settlement layer and distribution channel	Retained paid activity after credits or incentives end; complement use; cost of switching or using several bases
Supply capacity	Lab trains models; cloud serves training and inference	Miner competes for rewards; validator supplies consensus and execution	Realized utilization, replacement spending and returns after capital costs; changes in competitors’ outcomes

A larger training run is not automatically a network effect: spreading its cost over more customers is an economy of scale. Nor is a larger validator set interchangeable with more paying applications. The adoption share x must refer to a defined field making a comparable choice. Dollars of capital expenditure, hashrate and model benchmark scores are useful observations, but none is x . Months of technical leadership do not directly measure the model’s erosion parameter r .

4.2 Adoption: direct value can precede competitive pressure

The evidence for useful AI is strongest where the outcome is specified. The staggered introduction studied by [Brynjolfsson et al. \(2025\)](#) covered 5,172 support agents and raised issues resolved per hour by 15% on average, with larger gains for less experienced workers. This is evidence of a productive channel in that workplace, not a universal return on an AI subscription. Verification, integration, error correction and organizational change still enter the adopting firm's cost. A public demonstration can raise belief p without proving a particular production workflow.

Results differ across tasks and users. METR's early-2025 randomized study found experienced open-source developers took 19% longer with the available AI tools. Its February 2026 update reported apparent improvements in a later study but warned that developers and tasks were selecting out of the no-AI condition, making the new estimate unreliable ([Becker et al., 2026](#)). Neither result licenses a claim about all software work. Together with the support study, they make task-level testing necessary and show why self-reported enthusiasm, benchmark gains and realized productivity must remain separate.

Productivity can nevertheless become pressure. Once rivals pass savings to customers, a firm can lose margin while the technology continues to create value; [Section 1.3](#) models that distinction. Competition for a fixed position in search or an AI answer is a different mechanism, closer to α . AI-generated attacks create a third: the buyer pays to prevent losses imposed by attackers. [Section 7.3](#) evaluates that defense at the verification workflow, where catching an additional fraud attempt matters more than general model capability.

Crypto demand needs equally specific treatment. In the Federal Reserve's October 2025 household survey, 9% of US adults reported buying or holding cryptocurrency as an investment during the prior year, while 2% reported transactional use ([Board of Governors of the Federal Reserve System, 2026](#)). Those overlapping categories measure different purposes, not the success and failure of one payment product. A holder's demand for an asset, a trader's demand for liquidity and a merchant's demand for acceptance have different alternatives and network requirements.

Dollar access and cross-border payouts can have direct value before general merchant acceptance. In Stripe's customer account, Remote's payments executive estimates that optional USDC payouts reach contractors two to three days faster and avoid currency-conversion costs exceeding 1% of payout value ([Stripe, 2026](#)). These are company-published estimates, not a controlled comparison or a complete cost calculation: later redemption, local spending, custody and compliance can reduce the gain. But they identify a buyer and a plausible saving that a blanket coordination story misses. Conversely, Visa's filtering of stablecoin transfers removes substantial automated and other non-economic activity; even its adjusted series is not a census of purchases ([Sheffield, 2025](#)). Repeat use after all costs is the demand test.

4.3 Four Layer 1 cases

Bitcoin: monetary coordination and a separate mining contest. Bitcoin's original design combines peer-to-peer transfer with proof-of-work to establish a transaction history ([Nakamoto, 2008](#)). Its potential network advantage comes from other holders, counterparties, liquidity and infrastructure accepting the same asset and history. That differs from attracting developers to a general-purpose application runtime. Holding bitcoin does not entitle its owner to a corporate revenue stream. Users pay transaction fees; miners receive those fees and protocol issuance in exchange for block production. Custodians, exchanges, hardware makers and power suppliers have separate businesses. The monetary network may retain demand while competition erodes an individual miner's return. That separation is the central Bitcoin lesson for the model-provider comparison.

Ethereum: shared applications, with value divided across layers. Ethereum makes applications and assets available within a shared execution environment. Contracts can call other contracts, allowing a new application to use existing services rather than recruit every complement afresh ([Ethereum.org Contributors, 2026a](#)). That is a concrete network mechanism; developer counts alone are only a proxy. It does not ensure that all value accrues to ETH holders. Under EIP-1559 the base transaction fee is burned, while the priority fee goes to the block producer; issuance and other validator receipts require separate accounting ([Buterin et al., 2019](#)). Rollups can retain their own execution relationships while buying settlement and data availability from Ethereum. EIP-4844 introduced a separate blob transaction and fee mechanism intended to lower rollup data costs ([Buterin et al., 2022](#)). Greater ecosystem use can therefore coexist with a lower fee per operation at the base. The empirical question is whether demand growth and settlement dependence compensate for that compression.

Solana: execution quality must turn into retained demand. Solana’s case is a shared execution platform competing on the cost and performance of using applications. Its fee rules distinguish base fees, split between burning and the validator, from prioritization fees paid to the validator ([Solana Documentation, 2026](#)). Cheap execution can make new activities viable, but low fees and high transaction counts do not by themselves establish pricing power. Reliability is part of the product: the February 6, 2024 halt lasted approximately five hours, according to the developers’ incident report ([Anza, 2024](#)). The outage illustrates the importance of reliability, although it does not establish the network’s current performance. The durable-advantage test is whether applications and users remain after incentives fall and competing chains improve, paying enough to support the relevant suppliers. A trading boom, validator votes and repeated automated transactions cannot be treated as interchangeable measures of customer demand.

Hyperliquid: an application can build the base around itself. Hyperliquid reverses the assumption that a broad developer ecosystem must come first. HyperCore includes the order books and matching engine; HyperEVM supplies general-purpose execution under the same HyperBFT consensus and can connect applications to that trading state ([Hyperliquid, 2026a,b](#)). The proposed network loop is specific: trading liquidity attracts traders, and applications can reuse that liquidity. Fee routing gives a closer connection to monetization than token market capitalization does. The documentation directs trading fees among HLP, deployers and an assistance fund that converts fees into HYPE; its current description says that fund’s HYPE is burned ([Hyperliquid, 2026c](#)). These are protocol rules, not audited company earnings. The unresolved question is how much independent application demand persists beyond the core trading venue. Liquidity concentration, dependence on the venue’s risk controls and changing fee rules limit the inference from one successful application to a defensible general platform.

4.4 Five foundation-model strategies

A foundation model is a reusable model adapted to many downstream tasks ([Bommasani et al., 2021](#)). Reuse can support a platform, but the economics depend on what the provider controls besides the model. Table 6 summarizes the comparison; the following examples are dated demonstrations of different strategies, not a ranking of today’s model quality.

OpenAI: model access plus a direct customer relationship. OpenAI supplies models through an API and through ChatGPT, so a developer can buy a component from a firm that also serves end users. Its July 2024 GPT-4o mini release priced input tokens at \$0.15 and output tokens at \$0.60 per million, over 60% below GPT-3.5 Turbo ([OpenAI, 2024a](#)). This is evidence

of price competition and product segmentation, not a quality-adjusted industry price index. Distribution, integrations and accumulated workflow context may retain users when a capability lead narrows. Automatic learning from every customer is a weaker claim: OpenAI’s business and API policy excludes training on customer data by default, subject to explicit opt-in ([OpenAI, 2026](#)). A data advantage needs lawful access and evidence of improvement, not just a usage count.

Anthropic: enterprise distribution and a first-party application. The February 2025 Claude 3.7 release offered access through Anthropic, Amazon Bedrock and Google’s Vertex AI, alongside the introduction of Claude Code ([Anthropic, 2025](#)). The dated release illustrates two routes to customers: distribution through several clouds and a direct application built on the model. Neither requires all developers to assemble around one exclusive host. Cloud relationships can expand reach while giving the distributor bargaining power; a first-party coding product can retain workflow value while competing with developers that buy the underlying model. Renewals and paid use after a competing release are better evidence of durability than a launch benchmark.

Google: integration changes where the return is earned. Google included Gemini capabilities in Workspace Business and Enterprise plans in January 2025, removing the need for a separate AI add-on ([Google Workspace, 2025](#)). That gives the model an existing route to users and changes the price a standalone application must beat. Bundling does not make development or inference costless. It means the return may be earned through suite retention, cloud demand or other products, rather than a separately reported Gemini margin. Comparing a standalone model company’s revenue with the whole of Google’s business would obscure that distinction.

Meta’s Llama: access to weights changes the supplier. Llama 3.1 permits deployment and modification under its community license, including a special permission requirement for licensees above its 700-million-monthly-user threshold ([Meta, 2024](#)). The license makes the weights available subject to these restrictions. A user can purchase hosting, adaptation and operations from another supplier rather than pay the model originator for every call. That can shift the budget toward infrastructure and applications. It leaves inference, evaluation, support and subsequent upgrades to fund. Downloads reveal distribution; they do not establish Meta’s model revenue or the profitability of a third-party host. Zuckerberg’s stated rationale links model release to Meta’s own products and reduced dependence on closed suppliers ([Zuckerberg, 2024](#)); that is a strategy for earning elsewhere in the business, whose incremental return still needs evidence.

DeepSeek: efficiency can reset the cost of competing. DeepSeek-V3’s technical report describes 2.788 million H800 GPU-hours for its training and an estimated \$5.576 million at an assumed \$2 per GPU-hour, explicitly excluding prior research and ablation experiments ([DeepSeek-AI, 2024](#)). The reported result challenges a simple equation between spending and capability. It does not measure the cost of building the company, reproducing its research or serving customers. More efficient training may reduce entry costs; cheaper inference may also expand demand enough to increase total compute purchases. Which effect dominates is an empirical question, and neither follows from the training bill alone.

Both types of platform can subsidize entry and attract complements without retaining their eventual value. Electric Capital’s 2024 analysis found one in three crypto developers working across multiple chains; its current methodology counts such developers in individual ecosystem totals and covers open-source repositories ([Electric Capital, 2024, 2026](#)). Those totals cannot be added as exclusive market shares. In AI, availability through several hosts similarly creates options, although moving a workflow still requires testing. Neither multi-chain development nor model routing implies frictionless switching.

Table 6: Layer 1s and foundation models: where the platform analogy holds and breaks. Interpretations follow the case evidence above.

Economic question	Layer 1 networks	Foundation-model providers
What do users share?	An asset and transaction history; on application chains, state, contracts and liquidity	A reusable capability; tools and distribution may be shared, but customer workflows need not share state
Where can the network effect arise?	Counterparties, liquidity and composable applications; the mix differs across the four cases	Complementary tools and distribution; a data feedback loop requires access rights and useful learning
What makes moving costly?	Moving assets and liquidity, bridging risk, contract dependencies and counterparty coordination	Re-evaluation, integrations, model-specific behavior, data policies and committed cloud spend
Who receives the payment?	Miners or validators, application operators and intermediaries; burns change supply without creating corporate profit	Model API, application or hosting company; bundled returns may appear elsewhere in the provider’s business
What weakens capture?	Competing chains, cheaper blockspace, applications retaining fees and developers using several chains	Competing models, available weights, routing, cheaper inference and incumbent bundles

The “fat protocols” thesis proposed that shared state and tokens would concentrate value at the protocol layer (Monegro, 2016). Here it is a hypothesis to test. Ethereum’s rollups and Hyperliquid’s integration show why the application/base boundary matters; model providers moving into applications pose the same capture question. A rollup may still purchase Ethereum settlement, whereas a self-hosted model need not buy an inference call from its originator. They are different ways of changing the supplier, with different residual dependencies.

4.5 The arms race: mining supplies the mechanism, AI changes the prize

Bitcoin mining makes relative investment unusually visible. A miner’s expected share of rewards depends on its share of network hashpower. Difficulty adjusts, and issuance follows a protocol schedule, so adding machines does not proportionately expand the long-run number of blocks. The dollar value of the prize still changes with bitcoin’s price and fees (Bitcoin Developer Documentation, 2026). More efficient equipment can benefit its buyer while inducing rivals to replace equipment too. Budish’s model links this rent competition to the cost of deterring attacks (Budish, 2018). The security service prevents a conclusion that all mining expenditure is socially wasted merely because miners compete for rewards.

MARA’s reported results expose the pressure without requiring a welfare estimate. Year-end energized hashrate rose from 53.2 to 66.4 EH/s in 2025, yet fourth-quarter blocks won fell from 703 to 595 compared with the prior-year quarter. The company attributed higher purchased energy cost per bitcoin partly to network difficulty outrunning its hashrate growth (MARA Holdings, Inc., 2026). Energized capacity is a point-in-time capability, while blocks are a period outcome; uptime, fees and chance also matter. The evidence is consistent with a race and does not isolate the causal return to each machine. Its reported net loss also included large asset revaluations, which cannot be read as mining’s operating loss.

AI laboratories face an analogous incentive when an improvement wins customers, developers or distribution and rivals respond with another training run, release or price cut. But AI can expand

the prize by enabling work that was previously uneconomic. It also involves learning about uncertain demand. Compute-optimal training research demonstrates that how compute is allocated across model size and data matters to performance (Hoffmann et al., 2022). Capital expenditure alone cannot distinguish productive expansion, experimentation and defensive investment. Evidence of competition requires a sequence connecting rivals’ advances to purchases and to losses experienced by those that fall behind.

NVIDIA’s sales show the scale of spending on AI infrastructure: its fiscal-2026 Data Center revenue was \$193.7 billion. The 71.1% GAAP gross margin reported for that year belongs to the entire company, not specifically to Data Center (NVIDIA Corporation, 2026). Those sales combine several workloads and do not identify the fraction attributable to a positional contest. Nor do they establish that every infrastructure supplier earns equally durable returns. Financed capacity carries utilization, replacement and customer-credit risks after the equipment seller books revenue.

The two races also meet in physical infrastructure. Core Scientific’s 2025 filing reports \$65.4 million of colocation revenue, against \$24.4 million in 2024, as it adapted capacity for high-density computing (Core Scientific, Inc., 2026). Power and sites can have alternative uses; bitcoin ASICs cannot simply become AI accelerators. Conversion requires investment and a customer. The distinction separates an asset owner’s option from a claim that every miner can escape deteriorating economics by switching to AI.

4.6 A worked race and six possible outcomes

Consider one round of the model with $p = 1$, $r = 2$, $c = 0.4$ and $\beta = \varepsilon = 0$, in normalized payoff units. If $\gamma = 0.2$ and $\alpha = 1$, adopting first pays 0.8 and adopting when everyone else has pays 0.3, but the field’s saturated payoff is $y_R = -0.2$. Set $\gamma = 0.6$ and $\alpha = 0.2$ instead: the saturated adoption gain is still 0.3, while $y_R = +0.2$. Both can sustain adoption; one destroys value for the modeled field and the other creates it. This is an illustration, not a calibration to mining or AI, and the two examples need not have the same adoption path.

A sequence of new generations can restart the contest: a useful advance gives an early lead, rivals match it, prices or rankings adjust, and the next purchase restores an advantage. That sequence is an extension in interpretation; the base model makes a single irreversible adoption decision. It does not prove perpetual expenditure. Table 7 states conditional branches that may overlap, with no assigned probabilities.

Table 7: Conditional arms-race outcomes and observations that would distinguish them.

Outcome	Evidence supporting it	Contrary evidence; who bears the downside
Productive expansion	Paid demand and task quality grow enough to cover inference, development and replacement costs	Idle capacity and weak renewals undermine it; capacity owners and financiers bear excess investment
Returns consumed by competition	Shorter-lived leads, matching expenditure and falling returns after capital costs despite continued adoption	Durable differentiation or expanding surplus weakens it; exposed producers absorb the replacement bill
Consolidation or specialization	Exits from the frontier, durable niche customers or outsourcing at acceptable margins	Continued profitable entry weakens it; exiting firms bear stranded assets and creditors may lose
Several bases remain viable	Customers use multiple chains or models, switch after price changes and retain portable complements	Persistent exclusive retention weakens it; platforms fund recruitment without corresponding pricing power

Table 7: Conditional arms-race outcomes (continued).

Outcome	Evidence supporting it	Contrary evidence; who bears the downside
A technical or protocol change	Comparable service delivered with less resource use, or consensus changes remove a hardware requirement	Rebound demand or new bottlenecks sustain spending; owners of obsolete equipment bear write-downs
Rules partition the market	Export restrictions, access rules or customer procurement divide suppliers and demand	Substitution and continued trade weaken the partition; excluded suppliers lose sales and restricted buyers pay more

These markets can change when their underlying technical or regulatory rules change. Ethereum’s September 15, 2022 Merge ended proof-of-work on that network while preserving its transaction history and applications; it did not end Bitcoin mining ([Ethereum.org Contributors, 2026b](#)). Cambridge’s October 2021 account of China’s mining restrictions reports geographic displacement and a partial recovery in global hashrate by August ([Rauchs, 2021](#)). The race continued elsewhere. Available model weights and efficient inference can lower the cost of participating without shielding a rival from competitive loss. A counter in the formal model does the latter; a cheaper input need not.

4.7 Forecasts for 2031 and 2036

Our five-year forecast is that AI creates the broadest pressure through competition in repeatable services, with intense defensive spending in businesses exposed to automated attacks. Compute replacement creates a more concentrated race. In crypto, the strongest candidate beyond trading is settlement in costly payment corridors; within trading, liquidity can create stronger local dependence. By 2036, the leading additional candidates are agent-compatible customer access, shared collateral and physical automation. Their wider reach requires operational milestones that have not yet been demonstrated across the relevant populations.

These are conditional judgments from October 20, 2026, for October 20, 2031 and October 20, 2036. “Leading” identifies the strongest supported mechanism within a stated field; “narrow” separates intensity from breadth; “conditional” identifies a necessary milestone; “weak” means the proposed holdout-loss channel lacks support. The tiers are not estimated probabilities. The public evidence identifies useful capabilities and some operating outcomes, but does not identify comparable losses among holdouts, exits and adopters across these fields.

We distinguish pressure already present from additional deterioration, and an event reached *by* a horizon from one still active *at* it. Our provisional materiality test is a loss of at least 10% of positive baseline annual operating profit for at least half a defined commercial field, lasting two years after feasible alternatives, with adoption able to improve the actor’s position. The threshold is a research convention. Appendix [O.6](#) specifies sensitivity tests and identifies outcomes for which probabilities remain unestimated.

Table 8: Conditional forcing forecasts from October 20, 2026. Tiers compare mechanisms within the stated fields; they are not probability bins.

Capability / field	2031	2036
Specialist workflows	Leading: pressure in measurable, repeatable services.	Leading: pressure on staffing and service delivery if complete workflows become reliable.
Compute and power	Narrow but intense: capacity and replacement pressure.	Conditional: specialization and efficiency may relocate or reduce the bottleneck.
Agent interfaces	Conditional: important customer channels require compatibility.	Leading conditional candidate: machine-readable access becomes routine procurement.
Attack defenses	Leading within exposed workflows: repeated defense purchases.	Leading within exposed workflows if attacks keep adapting.
Physical automation	Narrow: repeatable tasks and favorable sites.	Conditional expansion: deployment and maintenance become repeatable across sites.
Stablecoin operations	Leading crypto candidate outside trading, in high-friction corridors.	Leading capability candidate; particular coins and chains remain substitutable.
Shared collateral	Conditional: specialist financial networks.	Conditional expansion: collateral mobility and integration become economically important.
Trading liquidity	Narrow but intense where execution quality concentrates.	Narrow and persistent if liquidity survives adverse cycles.
Bitcoin mining	Narrow and intense among continuing miners.	Conditional persistence after reward, energy and exit changes.
Bitcoin monetary use	Weak general forcing; possible specialist niches.	Conditional niches; no forecast of universal business dependence.
Agent payments	Conditional: selected production commerce.	Conditional expansion: authorization and settlement become access requirements.

AI: cost competition first, wider coordination later. For support, coding and back-office suppliers selling comparable output, the 2031 mechanism is verified savings becoming lower prices or lost orders for holdouts. The support study’s 15% productivity improvement establishes a task benefit, not that competitive transmission. METR’s later measurement difficulties and Anthropic’s conversation-based estimates also caution against equating apparent speed with accepted output or firm returns (Brynjolfsson et al., 2025; Becker et al., 2026; Anthropic, 2026a). Census measured AI use in any business function at 19.8% in early May 2026, following a broadened question; that is not the adoption share for a particular production workflow (U.S. Census Bureau, 2026). We would downgrade this leading forecast if comparable holdouts retain prices, orders and margins despite sustained rival savings.

By 2036, reliable performance across complete workflows could extend pressure from individual tasks to staffing, service speed and the organization of firms. The likely recipients are operators

selling accepted outcomes and companies controlling customer distribution. A specialist must retain value after model charges, review and integration. Physical automation could transmit similar pressure to logistics and factory operators, but its ranking remains below software workflows until multisite deployment and maintenance become repeatable. Existing industrial robot counts do not isolate AI capability or establish those returns ([International Federation of Robotics, 2026a](#)). Downtime or supervision that erases the cost advantage reverses the forecast.

Frontier laboratories and AI hosts face a different five-year problem: delivering competitive capacity and replacing it. NVIDIA’s realized sales and the IEA’s electricity analysis support the scale of infrastructure demand; neither separates productive expansion from positional expenditure. The IEA’s 2030 electricity figure is itself a projection ([NVIDIA, 2026](#); [International Energy Agency, 2026](#)). Delivered power, differentiated compute and infrastructure can receive this spending. Ten-year capture is less secure because efficiency, custom hardware, outsourcing and exits can move the bottleneck. Buyers that defer upgrades without losing capability or customers would weaken the forcing thesis; supplier revenue alone would not establish competitive pressure on buyers.

For financial verification and enterprise security, the forecast concerns additional expected losses after existing controls. Selected documented misuse supports the mechanism, without measuring its population incidence ([Anthropic, 2026b](#)). Through 2031, effective defenses can attract repeated spending in exposed workflows; through 2036, adaptation can sustain it. Owners of identity, payment and security signals have an advantage, while specialists need measured protection beyond bundled controls. If incremental losses stop rising or existing defenses absorb them, the independent budget shrinks. Attackers and defenders are separate populations here.

Crypto: particular financial relationships, with competing rails. For payout and treasury providers, begin with a named corridor, such as US-to-Mexico contractor payments. Remote’s reported USDC payout gains make this a concrete candidate, but complete redemption costs and competitor losses remain unmeasured ([Stripe, 2026](#)). By 2031, sustained delivered savings could force competing providers to match cost and availability. Distribution, redemption and local liquidity are plausible recipients. By 2036, continuous settlement could become ordinary service infrastructure, while the coin and chain remain replaceable. Project Agora’s tokenized bank-deposit prototype provides a plausible alternative to stablecoin settlement, although a prototype does not establish realized savings ([Bank for International Settlements, 2026](#)). Bank services matching the full economics would reverse the crypto-specific forecast.

Shared collateral ranks below corridor operations for 2031. The exposed actors are institutions and applications financing the same assets: useful composability becomes forcing only if outsiders suffer worse financing or lose counterparties. By 2036, usable rights, liquidity and integration could widen that dependence, benefiting custody, financing and application providers. Equivalent financing outside the network would defeat the stronger claim. Ethereum settlement, L2 execution and bank-led infrastructure therefore remain competing or complementary implementations, rather than a single adoption decision ([Ethereum.org Contributors, 2026a](#); [Buterin et al., 2022](#); [Bank for International Settlements, 2026](#)).

Trading venues can generate intense pressure on traders and market makers in particular instruments by 2031. For 2036, retained liquidity through adverse conditions matters more than current activity. Venues and order-flow distributors can charge only while execution or customer access remains difficult to replace. Cheap migration after incentives end would reverse the persistence forecast. Bitcoin mining supplies an existing replacement-pressure benchmark, with rewards, energy and exit determining its continuation; it does not establish a new economy-wide forcing event. General business dependence on holding bitcoin remains a weak forecast unless refusing it causes durable commercial losses despite alternatives.

The intersection: access can become necessary while implementations compete. MCP connects applications to tools and context; A2A coordinates agent systems. AP2 addresses payment authority, while x402 supplies a payment interface with multiple possible rails (Model Context Protocol Contributors, 2026; A2A Project, 2026; Parikh and Surapaneni, 2025; x402, 2025). These layers do not imply that agent commerce requires cryptocurrency. A priced x402 API endpoint establishes an offer, not repeat paid demand (Coinbase, 2026). For 2031, software suppliers and merchants face conditional pressure where customers actually delegate purchases. For 2036, interoperable access, authority and recourse could become procurement requirements. Workflow distributors and trusted payment operations are stronger capture candidates than generic connectors. Existing channels preserving orders cheaply would defeat the forcing claim; interchangeable connectors would defeat a supplier’s margin claim.

Table 9: The nine existing cases under the two horizons. These are conditional strategy judgments, not provider rankings or asset-return forecasts.

Case	2031 dependence	2036 structural change	Capture and challenge
OpenAI	Paid workflows and distribution can make the capability important; provider dependence remains unproved.	Renewal depends on workflow ownership, reliability and customers’ ability to route elsewhere.	Distribution and accepted outcomes could support returns; cheap switching and weak paid retention would undermine them.
Anthropic	Enterprise workflows can support demand where measured reliability pays for review and integration.	Interoperability can widen use while moving bargaining power toward workflow owners.	Reliable paid workloads could support returns; equivalent alternatives and cheap switching would weaken the case.
Google	Existing workspace and cloud distribution can bundle AI into ordinary purchases.	Capture depends on retained customer access and infrastructure returns as model supply changes.	Distribution and infrastructure could support returns; unprofitable bundling or customers bypassing those channels would weaken the case.
Meta / Llama	Available weights can reduce dependence on a single API and complement existing distribution.	Cheaper substitutes can spread capability while shifting revenue to hosts and applications.	Complementary businesses and deployment services; licenses and hosting still matter.
DeepSeek	Efficient models challenge the cost of matching frontier capability; full costs remain unobserved.	Efficiency may lower entry costs or expand demand; neither guarantees a lasting model margin.	Efficient operation with retained customers; a training-run estimate cannot establish profit.
Bitcoin	Strong existing replacement pressure for miners; weak general business need to hold the asset.	Mining depends on rewards, costs and exit; monetary forcing still needs counterparty loss evidence.	Power, equipment and paid services; token appreciation is a separate outcome.

Continued on next page

Continued from previous page

Case	2031 dependence	2036 structural change	Capture and challenge
Ethereum / L2s	Shared assets and collateral can create dependence within connected financial applications.	Wider collateral use is conditional; execution may move while base settlement persists or changes.	Liquidity and application access; base fees require separately measured settlement dependence.
Solana	Local dependence can form around retained applications, liquidity and execution quality.	Durability requires paid use through competing improvements and adverse cycles.	Applications, distribution and validation could support returns; cheap migration of useful activity would weaken the case.
Hyperliquid	The strongest case is concentrated trading liquidity; broader application dependence is conditional.	A wider platform requires external applications that retain paying users through stress.	Trading and application receipts; fee routing and burns alone do not establish holder profit.

Who can retain the spending. Table 9 carries the nine existing cases into both horizons without assuming today’s suppliers remain dominant. For incumbents, distribution, scarce delivered capacity and control of useful financial relationships offer the strongest capture mechanisms. For new specialists, the better candidates sell verified outcomes, incremental protection or difficult settlement operations across platforms. A new base needs retained customers and complements after subsidies; a protocol participant needs separately specified economic rights. Neither widespread use nor token appreciation establishes company profit.

Circle illustrates the payment boundary: its second-quarter 2026 distribution, transaction and other costs were \$412 million against \$701 million of revenue and reserve income, approximately 58.8%; the remainder is before operating expenses ([Circle Internet Group, 2026](#)). A strong settlement forecast therefore does not locate all the value at the issuer. We would revisit these rankings as independent cohorts establish accepted-output savings, changed competing offers and retained paid demand during 2027–2029. The stronger ten-year forecasts depend on those observations and on the specified deployment or interoperability milestones.

4.8 What would justify a new supplier

For another base platform, the burden is to identify demand that persists after recruitment subsidies and technical novelty fade, then show a defensible way to charge for it. For a supplier serving several bases, the test is whether portability, reliability or a specialized workflow saves customers enough to support a margin after infrastructure and support. Multi-platform demand is useful only if the platforms cannot cheaply absorb the function.

The stablecoin investment case also depends on the source of demand: issuer rules can create compliance budgets, but settlement can also earn through direct savings and counterparty demand. A new company needs corridor-level retention and net savings, not merely aggregate stablecoin supply. Likewise, useful AI does not guarantee a profitable model provider, and model-provider losses do not establish negative social value. Appendix O records the accounting boundaries and remaining evidence gaps. Section 5 turns these cases into observable signals; the investment memos then ask which of those signals leaves room for a new company.

5 Spotting a forcing event early

The useful signal is a change in the buyer’s economics that makes delay costly. Establish the mechanism first: a rival’s cost or performance improvement, a change in customer demand, a network access loss, a rising attack, or an enforced requirement. Then identify the exposed buyer and estimate the interval before the loss exceeds its tolerance, allowing for installation time. The market signals in Table 10 come first; Table 11 records the distinct evidence left by gatekeepers. Published enforcement dates can make timing easier to document.

A cost-saving pilot establishes what adoption can earn. Comparable holdouts’ prices, margins, win rates and retention show what abstention costs as adoption spreads. Use observations before and after a rival’s rollout and an unexposed comparison where possible; control for input prices, customer mix and the exit of weak firms. A capital budget, benchmark or expectation survey is an early hypothesis about pressure, not proof of it. The identification argument of Appendix I becomes practical here: adoption counts alone cannot separate productive value from competitive transfer.

A timing window can be expressed as a threshold and a range rather than a date. In the illustrative fulfillment case, an assumed rival deployment path reaches the holdout’s zero-profit threshold in about sixteen months; a nine-month installation leaves about seven months to begin. Lower pass-through removes the crossing entirely (Section 7.2). A buyer’s tolerance may bind before zero profit, but it has to be stated. The equivalent evidence in other channels is a renewal cohort losing customers, counterparties moving transactions, or expected attack losses overtaking the cost of defense.

Table 10: Leading indicators from the market, by channel. Lead times are illustrative diligence windows, not estimated predictive lags. Dated examples establish chronology, not causation.

Signal	Channel	Lead	What it looked like	False positive
Replacement spending rising while realized returns fall	Contest hypothesis	Observe each equipment or model generation	Compare mining efficiency with rewards per unit; AI capacity with paid utilization and replacement costs (Section 4)	Demand growth, experimentation or input-price shocks can explain spending; capacity is not an adoption share
Paid cohorts retained after incentives or a rival release	Platform durability	Cohort renewals; no universal lead	Measure fee-paying applications and paid model workloads, switching and use of several platforms (Appendix O)	Overlapping developers, airdrop activity, credits and account counts mistaken for exclusive retained customers
Capacity delivery, cash receipts and financing tracked separately	Supplier capture	Contract and delivery milestones	Distinguish orders, energized capacity, use and payment; trace supplier-funded purchases and residual asset risk	Backlog or a letter of intent counted as current sales; one dollar counted at several layers

continued on the next page

Table 10: Leading indicators from the market, by channel (continued).

Signal	Channel	Lead	What it looked like	False positive
A rival’s capital-expenditure or procurement guidance	Contest	Concurrent; 1–3 quarters before the company’s revenue	Meta raised its 2024 capex guidance on April 24, 2024 “to support our AI roadmap,” then guided \$60–65bn for 2025 and \$115–135bn for 2026; Alphabet guided \$175–185bn and Amazon about \$200bn for 2026 (Meta Platforms, Inc., 2024, 2026; Alphabet Inc., 2026; GeekWire, 2026); Nvidia’s data-center quarter went from \$22.6bn to \$62.3bn behind them (NVIDIA Corporation, 2024c, 2026)	A capacity signal read as demand: TD Cowen’s February 2025 report of Microsoft lease cancellations, and CoreWeave’s November 2025 capex cut for a late developer, while the company’s revenue kept rising (Fortune, 2025; TIKR, 2025)
A company’s generation cadence or price list	Contest	About 3 quarters to the company’s revenue	Blackwell announced March 18, 2024, samples in August, “in full production” in November, a \$35.6bn data-center quarter in February 2025; a “one-year rhythm” dated through 2028 (NVIDIA Corporation, 2024a,b, 2025); OpenAI’s price per token down 99% from 2022 to July 2024 (OpenAI, 2024a)	A slipped generation the last one covers (the August 2024 Blackwell delay report; Hopper filled the gap and revenue rose 16% a quarter (SiliconANGLE, 2024)); a price list falling because the capability is commoditizing, the last row of this table
A public benchmark or leaderboard result	Capability signal; can revise p	Weeks when the benchmark is a product; 5–15 months when a company can sell against it; years for a new industry	ImageNet 2012: 15.3% error against 26.2% (ImageNet, 2012); ChatGPT to Microsoft’s multibillion-dollar investment in 54 days (OpenAI, 2022; Microsoft, 2023); SWE-bench Verified (August 2024) to Cursor’s \$100m of ARR in five months and \$1bn in fifteen (OpenAI, 2024b; Anyosphere, Inc., 2025a,b)	A result with no buyer behind it: Inflection-2, “the second most capable LLM in the world,” November 2023, and the team hired away four months later (Inflection AI, 2023); a leaderboard that is gamed (27 private Llama-4 variants tested on Chatbot Arena) (Singh et al., 2025)
Competitors passing cost savings into prices or service, with holdout margins or retention falling	Cost and price response	The buyer’s tolerance threshold, less installation time	Fulfillment case: monitor realized cost per order, comparable contract quotes and rival deployment; the numerical path in Section 7.2 is illustrative	Savings retained as profit; unrelated price cuts; changing customer mix; holdouts able to specialize or exit cheaply
A learning curve on course to cross the incumbent’s cost for a class of buyers	Direct value, synchronized	12–36 months; computable from the curve	Battery packs at \$115 then \$108 a kilowatt-hour (BNEF, 2024–25) with Chinese EVs at parity; solar at 4.3 cents against coal at 7.3 (IRENA)	A curve that flattens before the crossing; a crossing in a segment too small to matter; a commodity whose producers’ margins go to zero as it crosses
Rivals naming the technology on earnings calls, or a CEO memo making it a condition of headcount	Contest or direct value	0–6 months (late)	Shopify’s April 2025 memo requiring teams to show AI could not do a job before hiring for it	Talk without a budget line; the mention is itself the competitive signal and the capability is bundled

continued on the next page

Table 10: Leading indicators from the market, by channel (continued).

Signal	Channel	Lead	What it looked like	False positive
Holdouts losing measurable share, fills, traffic or applicants, or saying so in a filing	Contest	Concurrent; the filing 1–2 quarters later	Pew’s observed 8% click-through with an AI summary against 15% without one, before establishing causation; slower trading firms losing fills in the latency race; Chegg’s May 2023 call, with ChatGPT “having an impact on our new customer growth” and the shares down 48% (Chegg, Inc., 2023; Associated Press, 2023); Nokia’s smartphone share from 50.9% to 40.8% in the year after the iPhone (Fierce Wireless, 2009)	Share lost for unrelated reasons, or an adoption share rising because holdouts exited (Appendix H.2); a holdout’s price hit with no share loss (Pearson fell 15% in sympathy with Chegg and called itself ChatGPT-proof (The Motley Fool, 2023)); an analyst’s forecast of holdout loss (Gartner’s projected 25% fall in search volume by 2026 (Gartner, Inc., 2024); revenue growth is not a direct test of a volume forecast)
A service level that most rivals match within a few quarters	Expectations	0–12 months	Same-day delivery (Amazon, Walmart and Target in 2025); daily driver payouts (Lyft to Uber in three months, 2015–16)	A cascade with no measured customer loss; a platform that bundles the feature before any company budget forms
Customer-expectation surveys	Expectations	2–5 years, once the supply side moves	PwC, 2018: 32% of customers would leave a brand after one bad experience (PwC, 2018); Salesforce, 2023: 77% expect to interact with someone immediately (Salesforce, Inc., 2023); the delivery surveys of Section 3.2; the support-AI budgets of 2023–26 followed (Fin, Sierra, Decagon)	A stated preference without a budget: openness to AI fell in the 2023 survey while support-AI budgets rose fastest
Attack statistics rising faster than defenses	Counter	6–18 months	Deepfakes in one of five biometric fraud attempts; DDoS record sizes reset yearly	The platform absorbs the defense (mail providers against AI-written outbound; job boards against AI-written applications)
A counter company appearing, or a venue adding a native defense	Counter	3–7 years from founding to an enterprise budget; its statistics lead by 2–3 years	DDoS: the attacks on Yahoo and eBay in February 2000, Prolexic founded in 2003, bought by Akamai for \$370m in 2014 (Wikipedia contributors, 2026b); deepfakes: Deeptrace’s 2019 census, Sumsu’s tenfold rise in 2023, a Gartner designation and a KPMG stake for Reality Defender in 2026 (Deeptrace, 2019; Sumsu, 2023; Reality Defender, 2026)	A counter founded before the attack is commercial (96% of the deepfakes Deeptrace counted in 2019 were pornography); the venue absorbs the defense, and the counter becomes a feature
The venue launching its own product for the race	Contest	Negative signal	Ads in ChatGPT and in Google’s AI Mode (2026); retailers’ own media networks, more than 200 by 2024, with Criteo’s retail-media line flat	Possible exception: customers retain independent tools that improve on the venue’s product

continued on the next page

Table 10: Leading indicators from the market, by channel (continued).

Signal	Channel	Lead	What it looked like	False positive
The price of the forced capability falling toward zero	Any	Negative signal	Let’s Encrypt, December 2015; a free e-invoicing tier inside an accounting suite	Possible exception: customers still pay for complementary integration or management services

Table 11: Leading indicators from a gatekeeper. Lead times are illustrative windows from the signal to the first budget, not estimated predictive lags.

Signal	Channel	Lead	What it looked like	False positive
A final rule with a dated obligation, in the Federal Register, the EU Official Journal or a ministry circular	Exclusion	12–36 months	CMMC’s acquisition rule, published September 10, 2025, with Level 2 certification required in solicitations from November 10, 2026; France’s September 1, 2026 and 2027 e-invoicing phases	A proposed rule; a rule with no enforcement (France levies no sanctions in 2026); a date that slips (the EU AI Act’s Annex III moved 16 months in July 2026) or is suspended (CMMC’s second phase, July 13, 2026)
A platform policy post with a version number and a date	Exclusion, steep Φ	3–18 months	Chrome’s February 8, 2018 announcement that version 68 would mark HTTP “Not secure”; the CA/Browser Forum’s 2025 ballot cutting certificate lifetimes; Apple’s age-range API	The platform supplies the capability itself, free, and the company market never forms
A new line item in security questionnaires, RFPs or subcontract flow-downs	Exclusion, buyer as gate-keeper	6–24 months	SOC 2 reports as a condition of enterprise sales, which built Vanta; CMMC clauses flowing down to subcontractors; AI-governance questions appearing in company reviews	Boilerplate questions nobody scores; a requirement an existing supplier satisfies by default
An insurer or lender condition	Exclusion	6–12 months	Cyber insurers tightening underwriting and stipulating controls such as multi-factor authentication and endpoint detection in the 2021 hard market, when US pricing rose 96% in a year (Marsh, 2021)	Conditions waived at renewal once the market softens
A dominant customer’s supplier letter with a tagging, labelling or data requirement and a date	Exclusion, customer as gate-keeper	6–24 months	Walmart’s January 2022 letter (tags in the specified additional categories by September 2, 2022) and its February 2024 expansion; Impinj’s revenue up 93% over 2021–24	A mandate the customer does not enforce, or whose gain it keeps for itself: Walmart’s 2003 pallet mandate had fizzled by 2007

Three signals warrant closer examination: announced deadlines, platform competition and pilot results. The first is the deadline that moves. The model separates belief in technical effectiveness from expected enforcement. A rule’s legal status does not remove uncertainty about enforcement or timing; ε must reflect the expected loss. Regulation 2026/1744 moved the AI Act’s Annex III high-risk requirements from August 2026 to December 2, 2027, with a separate August 2028 date for product-related high-risk systems ([European Parliament and Council of the European Union](#),

2026). France’s e-invoicing mandate was postponed from 2024 before it started in 2026. The SEC ended its defense of the stayed climate-disclosure rules in March 2025 and proposed rescission in May 2026; neither action was itself a completed repeal (U.S. Securities and Exchange Commission, 2025, 2026). A supplier planning around the original date could incur sales costs before delayed or cancelled demand. For a thesis that depends on a mandate, the strongest timing evidence is final text with credible enforcement. A proposal requires a probability and delay allowance; economic losses already borne by the buyer should be assessed separately.

The second is the platform that owns both sides of a race. AI-written outbound against inbox filtering is a real arms race, and Google and Microsoft include the defense in their mailbox products, limiting demand for a separate supplier. AI-written job applications against screening has the same shape, with LinkedIn, Workday and Greenhouse holding the defense. An independent supplier needs protection or information beyond the bundled controls. Identity verification is one such possibility, which is why these races inform that memo rather than establish separate investable markets here.

The third is the pilot. Its result may include a temporary positional advantage and omit the network value of later adoption. Appendix E.2’s $\alpha - \beta$ comparison is between the first adopter’s payoff and the field’s payoff at full adoption; it is not a general formula for a company’s price ceiling. The memos therefore estimate the buyer’s incremental ceiling $B(x, p) - \kappa - W$ along an adoption path, test continued operation and alternatives, and record what observable evidence would change that path. In a productive market, urgency can increase while the incremental ceiling stays flat.

6 Where we would invest now

We would prioritize companies that can operate at the new cost structure and suppliers of scarce complements that make adoption useful. The strongest new research priorities are narrowly operated AI services and repeatable industrial automation, followed by contract-backed power flexibility and corridor-specific payment operations. Security and compliance provide useful comparisons. Productive operators also merit consideration when they can demonstrate a customer advantage.

The distinction is between an attractive direction and an investable company. Public sources establish deployments, competing products and some company outcomes. They do not establish private unit economics, an entry valuation or a causal estimate of holdout losses in the proposed fields. The decisions below allocate diligence, not capital. A software-only seed investor should treat equipment financing, acquisitions and power contracts as separate capabilities to acquire or partner for, rather than hide them inside a software revenue multiple.

The analysis compares sixteen opportunities across crypto, AI, fintech and other markets, with detailed assessments of four and two-year operating scenarios for six. Appendix N gives their parameters and the full decision map. The repository note `paper/investment_directions_findings.md` contains the source ledger, public organization map, competing alternatives and calculations. The assessment uses public sources; interviews and paid experiments remain to be conducted.

6.1 Own the delivered service, then test the competitive pressure

For AI-operated specialist services, start with a defined document-completion or exception-resolution service in logistics. The investee sells accepted work and owns delivery quality; it is not merely selling seats in an assistant. C.H. Robinson reports deployed agents contributing to productivity, while its H1 2026 disclosure shows headcount declining alongside higher personnel expense, including restructuring (C.H. Robinson, 2025, 2026). That combination supports feasibility and warns against turning time saved directly into cash saved. It does not isolate an AI effect or identify non-adopting rivals.

Table 12: Investment research priorities as of October 20, 2026. “Advance” means investigate an identifiable paid position; it is not a completed investment decision.

	Direction	Position worth owning	Decision and binding condition
1	AI-operated specialist services	Recurring delivery contracts and dependable exception handling	Advance a narrow workflow; cash labor reduction and retained contribution must survive repricing.
2	Repeatable industrial cells	Application expertise, field service and financeable equipment	Advance one job family; uptime, changeovers and debt service must fit the fee.
3	Flexible power access	Accepted operating obligations and access/payment rights	Requires enforceable operating and payment rights.
4	Corridor payout operations	Embedded distribution and a difficult local endpoint	Conditional; a matched all-in saving and attainable retained flow are required.
5	Shared fraud signals	Useful receiving-account labels and operational access	Requires a data partner and evidence of losses prevented in live use.
6	Seller negotiation	Authority over a vertical commercial workflow	Run a small experiment; the scenario assumes no bargaining advantage until one is measured.

The constructed mature operator has \$20m revenue and \$12m of addressable labor. A realizable 20% saving gives $\gamma = \$2.4m$; \$400k upstream capability purchases and \$600k other costs leave $F = \$1.4m$. If competition passes through 80% of the saving, $h(1) = \$1.92m$ and the producer payoff change is $-\$520k$. With \$1.2m initial profit, the saturated adopter still earns \$680k while the holdout loses \$720k. The forcing mechanism is a change in market prices or service, not a mandate. The pass-through rate and adoption trajectory remain assumptions.

Value can persist in distribution, contracts and reliable completion as models improve. It need not persist in the first tool used. The investee pays models, staff and working capital before retaining the residual. The illustrative new-business funnel reaches \$3.27m annualized service receipts at month 24 and needs about \$4.06m of equity including a 25% contingency, before any acquisition. It remains cash-consuming at that point. The established operator used in the payoff calculation is much larger than the startup projected at month 24.

Convoy illustrates the risk that a logistics platform may change owners without delivering durable returns to its original investors: Flexport describes buying Convoy’s platform and later selling it to DAT (Flexport, 2025). That is not a controlled failure of AI. It is a reminder that distribution, finance and operating execution are part of the product.

6.2 Make the industrial application repeatable and financeable

For industrial automation, the initial product is a machine-tending or packaging cell designed for a consistent set of tasks. Formic’s 3D Machine case reports operating hours and a capacity gain under a service arrangement, without disclosing the price or cash saving (Formic, 2023). Greater capacity counts in γ only when it eliminates expense or earns contribution on sold output. The performance guarantee transfers some risk from buyer to operator.

The scenario assumes a \$150k reduction in labor expenditure and \$50k of scrap/rework savings, with a \$100k annual fee and \$60k of other buyer cost. Thus $F = \$40k$. With the assumed

effectiveness haircut and waiting value, the current margin is $-\$30k$: an application with slightly worse economics should wait or fail. The $\$100k$ fee is an assumed $\$25$ per paid hour at 4,000 annual hours, not Formic’s quoted price.

Supplier economics are a separate calculation. A $\$180k$ installed asset depreciated over five years, $\$35k$ annual cash delivery expense and $\$100k$ fee produce $\$29k$ economic gross profit before overhead. At half the fee, that becomes a $\$21k$ loss. The assumed equipment loan does not eliminate capital need: the 24-month path reaches about 29 active cells, requires $\$3.93m$ peak equity before contingency, and carries $\$3.51m$ equipment debt at month 24. Without that facility, peak cash need is $\$7.10m$ before contingency.

The industry comparison reinforces the distinction. Symbotic disclosed non-billable deployment overruns in its FY2024 reporting and a net loss alongside substantial FY2025 revenue ([Symbotic / SEC, 2024](#); [Symbotic / SEC Exhibit 99.1, 2025](#)). Its warehouse systems are not retrofit cells, but backlog growth does not establish that contract prices cover installation and maintenance risks. The attractive asset is a repeatable application with local service density and financing, rather than a generic robotics label.

6.3 Own a scarce complement: usable power or a payout endpoint

Flexible power access. The hypothesis is that accepted load flexibility lets a site obtain usable power sooner or on better terms. A Phoenix field study demonstrated a 25% power reduction on a 256-GPU cluster for three hours under its selected workload commitments; Google announced utility agreements; Indiana’s case index lists a Google capacity and demand-response arrangement as pending ([Coskun et al., 2025](#); [Google, 2025](#); [Indiana Office of Utility Consumer Counselor, 2026](#)). None is a disclosed startup fee or a general right to connect a new site.

For a 20MW site, assume 5MW becomes usable earlier and contributes $\$200k$ per MW-year after ordinary compute, power and workload opportunity costs, before the separately listed flexibility costs. Then $\gamma = \$1m$; a $\$250k$ fee and $\$350k$ other cost leave $F = \$400k$. The current haircut and waiting value require about 5.83MW to clear the purchase hurdle, so the assumed 5MW does not yet do so. If the utility recognizes no incremental capacity or access, set that benefit to zero.

The investee must control useful workload flexibility and possess enforceable payment/access rights. A cloud, site owner, utility or incumbent aggregator may own them instead. Duke’s modeled system headroom is not a site inventory or a software market ([Norris et al., Duke University, 2025](#)). The coordination-only funnel reaches four active sites and $\$1m$ annualized receipts at month 24, requiring roughly $\$7m$ equity with contingency. Generation, storage, buildings and transmission are excluded and require separate project finance.

Corridor payout operations. For payments, the scarce complement can be local redemption, liquidity and distribution. Stripe’s Remote case supports practical stablecoin use, while Wise’s falling aggregate take rate and fast transfers show that the competing rails are improving too ([Stripe, 2026](#); [Wise, 2026](#)). Compare completed recipient value in a named corridor, including conversion, failures, support and prefunding; blockchain speed alone is not the advantage.

The payout scenario assumes $\$200m$ annual eligible flow, a 20bp operational saving, a 6bp fee and 6bp other cost, with no unmeasured network premium. It gives $F = \$160k$ and a current decision margin of zero. At 80% pass-through, producer payoff falls $\$160k$ despite the positive adoption incentive. These are required economics to test, not a disclosed corridor quote.

The scale check is restrictive. McKinsey and Artemis estimate $\$226bn$ of annualized B2B stablecoin payment activity from December 2025, not total on-chain transfers or a full-year forecast ([McKinsey and Artemis Analytics, 2026](#)). At an assumed retained 6bp, that is $\$135.6m$ of global category receipts; 10% share gives $\$13.56m$ before corridor restrictions. $\$100m$ would require about

\$167bn flow, 74% of the dated global run rate. A narrow independent entrant needs exceptional distribution, a more valuable endpoint or additional paid services.

Reserve-income distribution is a different pool: Circle’s Q2 2026 filing records substantial Coinbase-related distribution expense ([Circle / SEC, 2026](#)). It illustrates the value of a distribution relationship, not a fee available to any payment startup. The 24-month customer acquisition scenario for payouts produces only \$0.55m annualized receipts against a \$4.69m equity envelope including contingency, even with a hypothetical liquidity facility. An existing operator improving its rail can be a better investment than building a low-fee corridor company from zero.

6.4 Negotiation and fraud services require further evidence

Seller negotiation. Automated procurement may create demand for seller-side negotiation tools. Buyers and rival suppliers must be modeled as separate groups. Rival suppliers can compete over contribution; buyer automation is an external state. Pactum documents procurement deployments and supplier satisfaction, while Vendavo already supplies seller pricing and commercial workflows ([Pactum, n.d.](#); [Vendavo, 2026](#)). These establish activity and substitutes, not the seller’s incremental bargaining gain.

With no measured bargaining advantage, the scenario counts only \$100k in administrative savings against \$200k in annual costs, giving $F = -\$100k$. At $r = 2$, at least \$200k of lone-adopter contest edge would be needed merely to restore $F = 0$, with no added cost. That edge must be measured in contribution after volume, payment timing and concessions. The interesting product may be a system that prices terms, finances a concession or refuses an unprofitable agreement. Conversational skill alone is not the investment.

Shared fraud signals. Retain the bank’s residual-loss problem but narrow the product to receiving-account information an entrant can legally and operationally use. Pay.UK’s reported detection uplift came from historical transaction data; Mastercard describes live customer deployments; BIS Hertha used synthetic data ([Pay.UK, 2024](#); [Mastercard, 2023](#); [Bank for International Settlements, 2025](#)). Those are three different evidence levels. None directly estimates an entrant’s incremental net prevented loss.

The fraud scenario assumes 15% prevention of \$5m residual annual loss and no separately counted network benefit: $\gamma = \$750k$, $c = \$600k$, $F = \$150k$. With the current haircut and waiting value, the buyer needs 25% net prevention to buy now. Network operators and established suppliers already have access, so a new company needs a specific data right, useful incremental signal and paying counterparties. A rising expected loss can raise a willingness-to-pay ceiling; competition and false positives still determine the realized price. The fraudsters and banks are not interchangeable shares in the paper’s same-field counter theorem.

6.5 Other investment opportunities and benchmarks

Certificate lifecycle and cryptographic agility. This remains the clearest requirement-led benchmark for heterogeneous machine estates. The adopted public-certificate schedule moves to 200-day validity in March 2026, 100 days in March 2027 and 47 days in March 2029 ([CA/Browser Forum, 2025](#)). Inventory, rotation and outage prevention can remain valuable above a cheap credential. Free automation and incumbent integration constrain the independent company’s fee. Five hundred customers at \$50–200k means \$25–100m revenue, reaching the top only at the highest fee and, in a 2,000-account pool, 25% share. Post-quantum migration and recurring certificate operation have different renewal economics.

Deepfake identity and age verification. Retain incremental defense inside an existing paid workflow. Separate the requirement to check age from fraud prevention. Counts of checks and requests from different platforms are not automatically deduplicated, independently payable events. Require incremental prevented loss or improved acceptance after false rejects, review expense and a credible bundled substitute. The counter model does not predict that realized per-check prices must rise.

Civil counterdrone. Keep it for an investor with the appropriate hardware, installation and customer channel. Legal authority and a funding line can open a market without forcing a recurring purchase. A grant, procurement order or site count is not software ARR. Recurring maintenance, detection quality and the platform's ownership of the customer determine the paid position.

Placement in AI answers. Keep the contest as a model benchmark; downgrade a new horizontal measurement entrant. The venue controls access and can sell the optimization workflow. A transaction feed, regulated workflow or proprietary outcome measurement may justify a narrower company. The category spending estimate and the worked supplier revenue scenario measure different things. Neither establishes that a new entrant can reach \$100m in revenue.

Invoice exceptions. Retain a vertical exception-resolution opportunity only where it beats an already compliant system. A mandate for the base platform does not make the add-on compulsory. The worked example in Section 7 compares it against suite capabilities. Five country opportunities of €18m yield €90m before proving comparability, adoption and share; an eventual category pool is not the startup's sales plan.

Horizontal compliance and CMMC. Monitor enforcement dates before prioritizing an investment. Underwrite the current enforceable condition, who owns the workflow, recurring paid effort and alternatives. A multiyear assessment contract is not annual recurring revenue. A deadline can synchronize shopping while most spending remains with suites, auditors and established service firms.

What would reverse this ranking

Lower the priority of AI-operated services if delivery savings disappear after all costs or accrue entirely to customers. Lower the priority of industrial cells if failures, changeovers or financing eliminate the operator's margin. Do not advance flexible power access without accepted rights, or payout operations without matched corridor economics and distribution. Seller negotiation rises in priority only with measured seller contribution; shared fraud signals rise only with live incremental prevention and usable data rights. Across all sixteen, persistent profitable holdouts or a compatible lower-priced incumbent can invalidate the forcing or capture thesis. Absence of public AI claims is not evidence that a firm is a holdout.

7 Four worked cases

The cases begin with two economic mechanisms: brands competing for answer placement, then fulfillment operators whose productivity gains change market prices. They are followed by a hybrid defense sold into existing verification flows and an invoice add-on sold above a compliant platform. The first case's field path uses the contest model; the second is the explicit price-response extension in Section 1.3, with its deployment path assumed rather than estimated. The final two compare the startup with an alternative that already satisfies the mandate, so the penalty cancels. Detection earns avoided fraud losses; invoice exceptions earn labor savings. Three cases trace monthly company revenue paths, while fulfillment isolates the timing and capture consequences of a changing

cost structure. Inputs and generation commands are recorded in Appendix S; planning assumptions are not forecasts of named companies.

7.1 Case 1: answer placement for a consumer brand

No imposed adoption requirement drives this case; brands compete for placement on the answer platforms. The buyer is a consumer brand with \$10 million of digital marketing spend, of which search takes 22.8%, the share Gartner reports for SEM and SEO (Semrush, 2026): \$2.28 million a year. Pew observed clicks on conventional results in 8% of visits with an AI summary, against 15% without one (Chapekis and Lieb, 2025). That comparison motivates exposure to changing search behavior; it does not show that named brands receive all the lost traffic. This case assumes a fixed prize. A category’s answers name a handful of brands, so a brand that optimizes while its rivals do not takes placement from them, and gives it back as they catch up. We assume answer placement decides \$100,000 a year of the brand’s gross profit, about 4% of its search budget, and that an optimizing brand is named five times as often as one that is not ($r = 5$); together those put the lone adopter’s edge at $\alpha = \$400,000$ a year, the number a pilot would show while no rival had the tool. With every rival on it the edge is $\alpha/r = \$80,000$, plus \$20,000 of measurement value that does not depend on rivals. Against a \$60,000 price, just under Semrush’s \$64,000 enterprise average (Semrush Holdings, Inc., 2026b), \$60,000 of the brand’s own content and feed work and \$10,000 of waiting value, the margin to buy first is \$290,000 and the margin at full adoption is minus \$30,000 (Table 13). The launch price clears until 66% of the brand’s rivals have a tool (Figure 4). These are purchase margins at the launch price, not price ceilings. The full-adoption ceiling is $B(1, 1) - \kappa - W = \$30,000$; leaving the buyer a further \$20,000 of margin reduces the supplier price to \$10,000.

The field’s own path says when that happens. Simulating the brand’s category from the same parameters (costs dispersed with $\sigma = 0.3$ around the mean, 3% of brands on a tool at launch, eligible non-adopters each having a 4.5% monthly adoption probability; Appendix S) puts the field at 50% after 18 months, 58% after 24 and 72% after 48; it passes 66% after 34. This is an irreversible path, not the static equilibrium: adoption continues beyond month 48 and can strand brands that bought while the edge was larger. Holdouts lose ground even though full adoption does not pay the mean-cost brand. Under these assumptions, the supplier can sustain its launch price for 34 months.

The supplier’s early revenue and acquisition costs appear attractive. Six hundred qualified brands over 24 months, a 15% win rate against Profound, Peec, Adobe’s Brand Visibility and the enterprise SEO platforms, four installations a month, \$6,000 of variable cost a customer (a 90% gross margin), \$35,000 of acquisition and onboarding (a 7.8-month payback) and \$250,000 a month of fixed cost give \$2.1 million of ARR at month 12 and \$3.3 million at month 18, with \$2.5 million of net new ARR in the trailing year and a burn multiple of 1.0 (Table 14). Those early results could support a Series A, but the launch price becomes unsustainable in month 34, within the financing period considered here. The path table compares two policies against the same simulated field at the launch price. It does not solve the feedback from this supplier’s repricing to industry adoption; customers can leave the supplier while retaining the capability elsewhere. Hold the price, and from month 35 customers do not renew: ARR peaks at \$11.3 million in month 34 and is \$3.5 million at month 48, and the second generation, placement in agent-driven purchases launched at month 20 with a smaller edge at the same price, is past its own break-even by month 39, because its edge erodes on the same timetable. Cut the price to the ceiling, keeping \$20,000 of margin for the buyer, and the customers stay: ARR is \$8.2 million at month 48, with the first generation at \$33,000 a year by month 48. Revenue depends on successive product generations. The two-generation illustration does not establish a sustainable release cadence. Both the next product’s value and its timing must

Table 13: Case 1 scenarios: an answer-placement tool sold to a consumer brand whose rivals are buying the same thing. Monetary amounts are USD thousands, annual except acquisition, onboarding and cumulative cash. Peak ARR is measured over the first 24 months. The field’s adoption share is simulated from the same parameters with the model’s dynamics; “launch price held” means customers whose margin turns negative do not renew.

Input or result (USD thousands per year unless stated)	Downside	Central	Upside
Direct value, γ	10	20	40
Lone adopter’s edge, α	280	400	500
Strength ratio, r	5.0	5.0	3.0
Gross buyer benefit B with no rivals on the tool ($x = 0$)	290	420	540
Gross buyer benefit B with every rival on it ($x = 1$)	66	100	207
Launch price, q	60	60	60
Other annual buyer cost, κ	60	60	60
Waiting value, W	20	10	0
Margin to buy first, $M(0)$	150	290	420
Margin at full adoption, $M(1)$	-74	-30	87
Share of rivals on the tool at which the launch price stops clearing	28.8%	65.9%	100%
Month the field reaches that share	9	34	never
Field adoption at month 18, 24 and 48	40 / 44 / 52%	50 / 58 / 72%	50 / 60 / 82%
Qualified accounts over 24 months	300	600	1,000
Win fraction	10%	15%	20%
Installations per month after ramp	2	4	6
Annual churn while the price clears	15%	10%	8%
Second generation: launch month	never	20	16
Active customers at month 24, launch price held	2	74	120
Peak ARR in first 24 months, launch price held (month)	545 (9)	6,897 (24)	12,196 (24)
ARR at month 24, launch price held	118	6,897	12,196
ARR at month 24, price cut to the ceiling	272	6,231	12,196
Gross margin at the launch price	90%	90%	90%
CAC payback (months)	7.8	7.8	7.8
ARR at month 18 (Series A), launch price held	218	3,305	7,405
Burn multiple, trailing 12 months	n/a	1.0	0.5
Cumulative cash low point	-5,853	-4,701	-5,130

be tested against the platform-controlled contest described in Section 6.5.

The size of the prize and the rate at which rivals adopt limit attainable revenue. Five thousand enterprise brands at a 10% share is \$30 million of ARR at the launch price, \$16.7 million at the price the field supports at month 48 and \$5 million at the price full adoption supports; a category leader at a 40% share clears \$100 million at the launch price and not at the month-48 price. The downside is the small prize: at \$70,000 the launch price clears for nine months, ARR peaks at \$545,000 and the company has spent \$5.9 million. The upside is the large one: at \$250,000, with a holdout one third as likely to be named, the price clears at any adoption and ARR reaches \$12.2 million at month 24 on a burn multiple of 0.5, which is the category leader’s case. What separates the three is the prize and the strength ratio, and a pilot measures neither. The two numbers a seed memo in a contest has to carry are the ceiling at full adoption and the share of rivals already on a

Table 14: Case 1 central path under two price policies, and the ceiling at three prices.

Central case (USD thousands)	Launch price held	Price cut to the ceiling
Field adoption at month 12		39%
Field adoption at month 24		58%
Field adoption at month 36		67%
Field adoption at month 48		72%
ARR at month 12	2,092	2,092
ARR at month 24	6,897	6,231
ARR at month 36	10,186	8,244
ARR at month 48	3,499	8,192
of which the second generation, month 48	1,884	3,453
Price at month 48 (gen. 1, USD thousands)	60	33
Series A snapshot, month 18: ARR	3,305	3,305
Net new ARR, prior 12 months	2,493	2,493
Burn multiple	1.0	1.0
Gross margin at the launch price		90%
Eligible accounts	5,000	
Achievable share	10%	
Ceiling ARR at the launch price (60)	30,000	
Ceiling ARR at the price the field at month 48 supports (33)	16,727	
Ceiling ARR at the price full adoption supports (10)	5,000	

tool.

7.2 Case 2: fulfillment savings become competitive pressure

An illustrative fulfillment operator processes one million orders a year. Automation lowers its operating cost; as rivals adopt, competition passes part of the saving to customers through lower prices. The mechanism has observable precursors: Amazon has described cost savings it aims to pass to customers, and Symbotic’s January 2025 agreement envisaged systems at 400 Walmart pickup and delivery centers, subject to performance criteria (Amazon, 2024; Symbotic Inc., 2025). Neither source establishes the price response or holdout losses assumed here. Every numerical input below is a planning assumption, recorded in `data/investor/productivity.json`; none estimates a named company.

At the initial price of \$5 an order, the operator’s \$4.50 of operating cost leaves \$0.50, or \$500,000 a year. Costs include the allocated operating costs of the site at this volume. Automation saves \$1.20 per order before a supplier charge of \$0.35 and the buyer’s own \$0.15 of integration and operating costs. The supplier charge includes annualized equipment and support. At unchanged prices the adopter earns \$1.20 per order, or \$1.2 million a year. Adoption already pays: the \$700,000 improvement exceeds the assumed \$50,000 annual equivalent of waiting value. Economic pressure arrives when rivals change the price the old process must earn.

Let $s = \$1.20$ be the gross saving and let $\rho = 0.8$ be the assumed share of that saving passed through when the rival field is fully automated. At rival adoption share x , suppose the market price is

$$P(x) = P_0 - \rho s x = 5 - 0.96x.$$

With annual volume N , incumbent cost C_0 , supplier price q and other buyer cost κ , annual operating profits are

$$\pi_N(x) = N[P(x) - C_0], \quad \pi_A(x) = N[P(x) - (C_0 - s) - q - \kappa].$$

These are operating profits, not the changes from the old world used in the base model. Subtracting the original \$500,000 from both yields the extension in Section 1.3: $\gamma = Ns$, $c = N(q + \kappa)$ and the common loss $h(x) = N\rho sx$. A price reduction is received by customers; it is not a fixed prize redistributed among the operators.

Table 15: Illustrative fulfillment economics at fixed annual volume. Supplier price is \$0.35 per order throughout. The combined ceiling is the lower of the incremental-return and continued-operation bounds, before competing suppliers. All values follow from stated assumptions.

Rivals automated	0%	50%	75%	100%
Market price per order (USD)	5.00	4.52	4.28	4.04
Holdout operating profit (USD thousands)	500	20	-220	-460
Adopter operating profit (USD thousands)	1,200	720	480	240
Incremental margin after waiting value (USD thousands)	650	650	650	650
Incremental price ceiling per order (USD)	1.00	1.00	1.00	1.00
Viability price ceiling per order (USD)	1.55	1.07	0.83	0.59
Combined price ceiling per order (USD)	1.00	1.00	0.83	0.59

The holdout breaks even at $x^* = (P_0 - C_0)/(\rho s) = 52.1\%$. At full adoption it loses \$460,000 a year; the adopter still earns \$240,000. The old process has become unviable even though the incremental gain from switching remains \$700,000, or \$650,000 after waiting value, at every adoption share in Table 15. This is why a positive-ROI product can also become forcing, and why urgency alone does not imply a rising supplier price.

The threshold gives a conditional sales window. Suppose rival adoption is currently 20% and rises by two percentage points a month. Zero holdout profit is then about sixteen months away. With nine months to install, implementation must start within roughly seven months to finish before that crossing. This is an assumed deployment path across firms with different constraints, not a simulated equilibrium or a forecast. At 40% pass-through, even full adoption leaves the holdout \$20,000 of annual profit and this particular deadline disappears; at 100%, the installation window shrinks to about two months. The records to collect are realized cost per order, matched contract prices, rival deployments and the buyer's installation lead time. Lost volume or a minimum acceptable margin would move the deadline earlier; customer differentiation or an inexpensive exit could weaken the pressure.

The supplier faces two buyer-side bounds. Incremental return permits at most $s - \kappa - w = \$1.00$ per order, with $w = \$0.05$ of waiting value. Continued operation also requires $q \leq P(x) - (C_0 - s) - \kappa$, assuming exit yields zero and no extra financing buffer is needed. The combined ceiling therefore falls from \$1.00 at launch to \$0.59 at full adoption. At the assumed \$0.35 price, the supplier earns \$350,000 of annual revenue per site. An assumed \$0.18 per order of delivery cost, including annualized equipment, maintenance and support, leaves \$170,000 of gross profit, a 49% gross margin before acquisition costs and central overhead. An equivalent competing offer at \$0.28 would cap the price below either buyer-side bound.

For scale, assume 1,000 technically suitable sites and a 10% attainable share: annual supplier revenue is \$35 million, or \$28 million at that competing price. These are illustrative site counts, not a market estimate or an acquisition forecast. Reaching \$100 million at the original price needs at least 286 sites with the same volume and economics. The investment question is whether a supplier can win and service that many sites while retaining a cost, deployment or integration advantage.

Economic pressure establishes the buyer’s urgency; it does not establish that advantage. A seed test would measure both the saving and the price response on comparable contracts, and a subsequent financing would need repeat deployments with retained gross margin.

7.3 Case 3: a detection layer sold into verification flows

This is the attack–defense branch of Section 4: more capable generative tools can raise losses at an existing verification flow, creating demand for a counter. The calculation values incremental fraud prevented after the base check; a general model benchmark cannot establish the detector’s catch rate or the buyer’s avoided loss.

The hypothetical buyer is a consumer platform that already runs a compliant age/identity check and considers a separate detection layer. Its requirement is unchanged, so $\varepsilon = 0$. For the numerical illustration, assume external expected loss $L = \alpha x / (1 - \varphi y)$ and incremental avoided loss $z_K = \varphi L$. Here x is the deepfake share of fraud attempts, y is the defended share of platforms, and α is a loss normalizer. These are different populations: this is a stipulated exposure/displacement scenario, not an application of the common-population equilibrium in Appendix F, and α is not an identified contest stake $V(r - 1)$. Holding L fixed would instead give a simple defense-ROI case.

Entrust’s one-in-five share of biometric fraud attempts and Pindrop’s reported deployment statistic provide context, not a matched estimate for this hypothetical platform (Entrust, 2025; Biometric Update, 2026b). The central scenario sets $x = 0.2$, $y = 0.1$, $\varphi = 0.8$ and $\alpha = \$3.45m$, which produces \$600k avoided loss by construction. At two million checks per year, an assumed \$0.15 price per check, \$150k other cost and \$50k waiting value, the margin is \$100k and the price ceiling is \$0.20 per check. These are annual scenarios, not a disclosed detection quote or measured prevention effect. False rejects, review, displaced attacks and residual losses must be reconciled to the customer’s actual records. The downside scenario’s smaller attack share and catch rate make adoption fail (Table 16).

The company sells the layer at a 73% gross margin, after \$0.04 of inference and review cost per check (\$0.05 and \$0.03 in the downside and upside). The funnel is deliberately slow: 100 qualified accounts over 24 months, a 15% win rate because Socure, Persona, Incode, Entrust and Veriff sell or build detection and are in every deal, one installation a month after a six-month ramp, 10% annual churn. At the flat price that produces 13 active customers and \$3.9 million of ARR at month 24, \$2.9 million at month 18 with \$2.1 million of net new ARR in the trailing year and a burn multiple of 1.9, and a cumulative cash need of \$7.6 million through month 24 at \$350,000 a month of fixed cost. The financing case depends heavily on the assumed win rate against established providers that could build the same capability.

The flat price is one policy; pricing at a fixed fraction of a changing ceiling is another explicit assumption. Entrust reported 58% growth in the count of deepfaked selfies (Entrust, 2025). Applying that growth rate to the share of attempts is a sensitivity exercise, not a valid inference about the denominator. Under that assumed share path, fixed catch rate and displacement rule, avoided loss rises from \$600k to \$948k and \$1.5m, and the price ceiling rises from \$0.20 to \$0.37 and \$0.65 per check. The calculation does not imply a company can charge those amounts.

A company that could retain 75% of the ceiling would charge about \$0.49 at month 24 and reach the higher revenue path in Table 17. The \$12.5m ARR and 92% gross margin in that path additionally assume unchanged per-check inference/review costs, no competitive fee response and successful repricing. The flat-price path remains valid even if expected loss rises. Likewise, the higher avoided loss when more platforms defend is a stipulated displacement stress, not a measured network effect. Incremental prevented loss, operational costs, supplier alternatives and renewal prices decide which path is attainable. A rising ceiling alone cannot establish pricing power.

Table 16: Case 3 scenarios: a detection layer sold into a platform’s verification flow. Monetary amounts are USD thousands per year except prices per check and the one-time acquisition and onboarding costs; run rates are at the stated month. The burn multiple is net cash consumed over net new ARR in the twelve months to month 18; “n/a” where no ARR was added, “cash positive” where the business funds itself.

Input or result (USD thousands per year unless stated)	Downside	Central	Upside
Verifications per year	1,000,000	2,000,000	3,000,000
Price per verification (USD)	0.10	0.15	0.20
Deepfake share of fraud attempts, x	10%	20%	30%
Platforms already running detection, y	10%	10%	10%
Catch rate of the layer, φ	50%	80%	90%
Assumed annual loss normalizer, α	3,450	3,450	3,450
Fraud avoided by the layer, $z_K = \varphi\alpha\tilde{x}$	182	600	1,024
Mandate penalty avoided beyond the existing check, ε	0	0	0
Annual contract value, q	100	300	600
Other annual buyer cost, κ	150	150	150
Waiting value, W	100	50	0
Buyer’s margin to purchase now, M	-168	100	274
Qualified accounts over 24 months	60	100	150
Win fraction	8%	15%	20%
Installations per month after ramp	1	1	2
Annual churn	20%	10%	8%
Active customers at month 24	0	13	26
ARR at month 24	0	3,860	15,738
Gross margin	50%	73.3%	85%
CAC payback (months)	45.6	8.2	3.4
ARR at month 18 (Series A)	0	2,908	11,797
Burn multiple, trailing 12 months	n/a	1.9	cash positive
Cumulative cash low point	-8,400	-7,641	-3,497

A simple scale calculation would add Incode’s 4.1 billion checks and Socure’s 2.7 billion requests (Burt, 2025; Biometric Update, 2025). Multiplying 10% of that sum by \$0.15 gives \$102m arithmetically, but does not establish an addressable independent detection market. The units, overlap, eligibility and suppliers’ rights to attach another paid layer have not been reconciled. A higher assumed price leaves the uncertainty about eligible transaction volume unresolved. A qualified flow inventory and the revenue the incumbent leaves to an independent company are required before using this as venture-scale evidence.

7.4 Case 4: invoice exceptions for a French mid-sized firm

The buyer is a French mid-sized firm with several finance systems, no invoice-automation tool today, and 12,000 invoice exceptions a year, which at 20 minutes each and €60 a loaded hour (Eurostat puts France’s average at €44.30 and its finance sector at €70.50 (Eurostat, 2026a,b)) cost it €240,000 of staff time; Billentis’s benchmarks put exceptions at 20 to 30% of paper invoices (Koch, 2017). It already has a compliant platform for the September 2026 mandate, because the accounting suites include one (Section 6.5). The startup sells an add-on that removes exceptions. The mandate’s penalty therefore cancels in the comparison, as it did in Italy, and the add-on is worth only γ : 40% of the exception work removed is valued at €96,000 a year. This assumes the released staff capacity can be used or converted to cash savings; it is not a measured payroll reduction. At €48,000 of price, €18,000 of other cost and €6,000 of waiting value, the margin to buy is €24,000, and the

Table 17: Case 3 under the stipulated external-exposure scenario: the attack, the exposure, the ceiling, and the company’s path if the price tracks it. Central sensitivity: the deepfake share of fraud attempts starts at one in five and is assumed to rise 58% a year, the layer catches 80% of the loss, and 10% of platforms are already defended. USD thousands unless stated.

Central case	Month 0	Month 12	Month 24
Deepfake share of fraud attempts, x (growing 58% a year)	20%	31.6%	49.9%
Exposure of an undefended platform, $\tilde{x} = x/(1 - \varphi y)$	0.22	0.34	0.54
Fraud cost exposed, $\alpha \tilde{x}$ (USD thousands)	750	1,185	1,872
Fraud avoided by the layer, $z_K = \varphi \alpha \tilde{x}$ (USD thousands)	600	948	1,498
Buyer’s margin at the flat price of 0.15 a check (USD thousands)	100	448	998
Price ceiling per check, $(z_K - \kappa - W)/\text{checks}$ (USD)	0.20	0.37	0.65
Price per check if the company keeps its month-0 share of the ceiling (USD)	0.15	0.28	0.49
Fraud avoided at month 0 if half of platforms run detection ($y = 0.5$)	920		
Deepfake share at which the flat price stops paying the buyer	16.7%		
Company ARR (USD thousands)	Month 12	Month 24	Month 48
Price tracks the ceiling	3,563	12,523	46,091
Flat price	1,905	3,860	7,202
<i>Indexed-price operating measures (USD thousands unless stated)</i>			
Series A ARR, month 18			7,210
Net new ARR, prior 12 months			6,035
Burn multiple, prior 12 months			0.3
Gross margin at month 24			91.8%
Cumulative cash low point, first 24 months			-4,274

Table 18: Case 3 central company path and unverified market arithmetic. The final four rows reproduce an unreconciled check/request sum; they do not estimate an eligible market.

Central case (USD thousands)	
ARR at month 12	1,905
ARR at month 24	3,860
ARR at month 36	5,619
ARR at month 48	7,202
Series A snapshot, month 18: ARR	2,908
Net new ARR, prior 12 months	2,060
Burn multiple	1.9
Gross margin	73%
Unverified raw sum of checks and requests (millions)	6,800
Assumed fraction of that raw count	10%
Illustrative unit price (USD)	0.15
Revenue arithmetic, not an estimated market	102,000

Detection layer: the price ceiling rises with the deepfake share of fraud attempts

$$\text{Ceiling per check} = (0.8 \times \$3.45\text{m} \times x / (1 - 0.8 \times 0.1) - \$200\text{k}) / 2\text{m checks}$$




Figure 3: Case 3 under the stipulated external-exposure scenario: the price ceiling per check against the deepfake share of fraud attempts, with the central scenario's other inputs fixed. The ceiling is $(z_K - \kappa - W)$ per check, where $z_K = \varphi \alpha \tilde{x}$ is the loss the layer shields; the dots mark Entrust's one in five and where a 58% annual rise takes it by months 12 and 24, and the dashed line is the price of a company that keeps the share of the ceiling it held at signing. Run with $p = 1$, $r = 1$, $\varphi = 0.8$, $y = 0.1$, $\beta = 0$, $\sigma = 0$ and no sheltering; the formula is printed on the figure.

product must remove more than 30% of the work to be bought at all (Figure 4); a ten-point miss on the 40% ends the purchase. If the suite's next release, or the AP-automation companies already selling exception workflows (Basware, Escher, Yooz, Medius, Coupa, SAP), halves the exceptions first, the add-on's 40% is worth €48,000 and the margin is minus €24,000: the mandate goes on driving compliance while the startup's product stops paying.

The scenario produces attractive unit economics but limited revenue scale. A 75% gross margin, a nine-month payback on €28,000 of acquisition and onboarding, and a lean fixed base of €100,000 a month give a burn multiple of 0.8 at month 18 on €1.9 million of ARR, with €2.1 million of cumulative cash needed. The €28,000 assumes distribution through the accounting firms that Pennylane and Cegid already sell through; a field sales motion into mid-sized firms costs about twice that and doubles the payback. ARR reaches €2.5 million at month 24 and €4.6 million at month 48. Assume 2,500 eligible French firms have the multi-system exception problem, an unverified customer count. Then a 15% share against the suites at €48,000 is €18 million of ARR, and at the five-times-revenue multiple Vertex paid for ecosio in 2024 (Vertex, Inc., 2024), a €90 million exit. The upside scenario reaches €5.2 million at month 24, but does not establish a larger eligible market. Under the central scale assumptions, five comparable countries yield €90 million of ARR; at least six would be needed to exceed €100 million. Comparable demand, localization costs and achievable share must be established country by country.

What would change the answer is the sequence of mandates. Domestic B2B requirements began in Belgium in January 2026 and were phased in across Poland in February and April 2026. Germany's issuance transition generally ends in January 2027 for firms above €800,000 of prior-year turnover and in January 2028 for the remaining in-scope firms, with an EDI transition through 2027. ViDA's cross-border B2B reporting starts in July 2030 (Belgian Federal Government, 2024; Polish Ministry of Finance, 2026; German Federal Ministry of Finance, 2026; European Commission, 2025). These are different populations and requirements; they offer possible expansion windows

Table 19: Case 4 scenarios: invoice exceptions relative to a compliant bundle. Monetary amounts are EUR thousands per year except the one-time acquisition and onboarding costs.

Input or result (EUR thousands per year unless stated)	Downside	Central	Upside
Exception work removed, e	20%	40%	60%
Gross annual benefit, γ	48	96	144
Annual contract value, q	36	48	60
Other annual buyer cost, κ	24	18	12
Waiting value, W	12	6	0
Buyer's margin to purchase now, M	-24	24	72
Qualified accounts over 24 months	100	300	500
Win fraction	10%	20%	20%
Installations per month after ramp	1	4	6
Annual churn	20%	10%	8%
Active customers at month 24	0	51	87
ARR at month 24	0	2,470	5,203
Gross margin	66.7%	75%	75%
CAC payback (months)	15.0	9.3	8.3
ARR at month 18 (Series A)	0	1,861	3,887
Burn multiple, trailing 12 months	n/a	0.8	0.4
Cumulative cash low point	-2,400	-2,099	-2,116

Table 20: Case 4 central path and ceiling.

Central case (EUR thousands)	
ARR at month 12	1,219
ARR at month 24	2,470
ARR at month 36	3,596
ARR at month 48	4,609
Series A snapshot, month 18: ARR	1,861
Net new ARR, prior 12 months	1,318
Burn multiple	0.8
Gross margin	75%
Eligible accounts	2,500
Achievable share	15%
Mature ACV	48
Ceiling ARR	18,000

rather than interchangeable copies of France, and established providers such as Pagero and Sovos already serve multiple jurisdictions. We would advance a French exceptions startup as a venture case only if it has paying customers in a second mandate country before that country's deadline, an expansion milestone that should be tested before further financing. Italy's precedent holds either way: in the mandate's first year 3.9 million Italian firms sent 2.09 billion invoices through the state's exchange ([Osservatori Politecnico di Milano, 2026](#)); TeamSystem sold the suite above that exchange, and the exchange itself was free to use.




Figure 4: One input at a time, with the central scenario's other inputs fixed. The vertical value is the buyer's margin to buy now, $M = B - q - \kappa - W$, in local currency; dots mark the central case and the break-even. Case 3 varies the price per check, with assumed avoided loss at today's attack share as the benefit (Figure 3 varies the attack); Case 4 varies the share of exception work removed. In the two add-on cases the mandate penalty cancels against the compliant alternative, with $p = 1$, $\beta = 0$, $r = 1$ and $\sigma = 0$. Case 3's α is the stipulated external-loss normalizer, not the contest stake of the common-population counter theorem; Case 4 sets $\alpha = 0$. Case 1 varies rival adoption, which moves the contest's ceiling; the orange dot is the simulated share at the Series A. The settings and formulas are printed on the figure. The productivity case uses Table 15 to separate switching gains from operating profits.

8 Where the framework misleads

The model organizes the reasons firms adopt, but adoption curves alone cannot identify those reasons. We fitted it to 14 historical adoption series, from hybrid corn to generative AI, with every observation traced to a source (Appendix J). The fits describe the paths well and say almost nothing about where the value came from. In every case the split between productive value and competitive transfer stays close to its prior. In the linear restriction with no outside option and no waiting value, a direct benefit of 0.3 with an edge of 0.2 traces the same curve as a benefit of 0.1 with an edge of 0.4 (Proposition 12); that exact equivalence does not hold for every nonlinear specification. With adoption data alone, the classification remains sensitive to prior assumptions. The classification has to come from holdout outcomes, bans, beliefs and buyer records, which is why every memo in Section 6 needs a buyer, an observable mechanism and a timing window; an S-curve alone supplies none of them.

The base model holds the market environment fixed except for its specified contest, network and exclusion terms. It does not solve for output prices, demand or firms' deployment schedules. The common revenue loss in Section 1.3 is an explicit extension, and the fulfillment case's price response and adoption path are assumptions. They establish a conditional threshold, not an estimated competitive effect or forecast. Volume, financing constraints, specialization and exit can move that threshold. Stronger evidence would connect a measured cost advantage to a market response and then to comparable holdouts' outcomes.

The investment screen adds another limitation. Its emphasis on identifiable buyers and accessible sales windows favors some requirement-led opportunities. That is a preference in underwriting, not evidence that economic forcing is rare or that useful products cannot create it. A supplier's weak capture can justify passing on an investment while leaving the forcing mechanism intact.

The model has four further limitations. It treats a private gatekeeper's penalty like a statute's, though enforcement depends on the gatekeeper's incentives and constraints; Walmart's failed 2003 tag mandate illustrates the risk (Section 2.5). It treats a ban as outside the game, although FINA's ban eliminated the polyurethane-suit market. It treats a deadline as fixed, although Section 5 lists three that moved. And it bounds the company's price by the buyer's gain without determining how many companies share that gain. The certificate case shows that a mandate can induce adoption without generating supplier profits. Assessing those profits requires a separate analysis of supply and competition.

The AI and crypto comparison adds a boundary on the unit of analysis. Repeated hardware replacement, new model generations, changing token prices, multi-platform use and entry change the field and the prize; a single irreversible adoption decision cannot determine their long-run equilibrium. A miner's security service and a model's benefits to downstream users also lie outside the profits of competing producers. Private-company losses, token burns, supplier revenue and capital expenditure measure different objects. Section 4.7 adds conditional 2031 and 2036 forecasts with separate judgments about capability pressure and company capture. Appendix O.6 distinguishes those qualitative tiers from constructed stress tests: neither supplies fitted event probabilities, validated long-horizon accuracy or a ranking by estimated social welfare.

Finally, the historical refits in Appendix J are hindsight. They use today's priors on cases chosen because they mattered, and the early-window estimates would not have been available on the day an investor needed them. We do not claim the framework would have picked historical winners. The investment implication is to pair evidence of a deteriorating holdout position with an asset the supplier can retain: accepted work, customer access, effective protection or difficult operating capacity. The former establishes urgency; the latter must survive competition, delivery costs and financing.

References

- 24/7 Wall St. Live: Will trade desk’s q2 earnings tonight reset its 80% selloff? 24/7 Wall St., 2026. URL <https://247wallst.com/investing/2026/08/06/live-will-trade-desks-q2-earnings-tonight-reset-its-80-selloff/>. Published 2026-08-06. Accessed 2026-10-20.
- 9to5Mac. Apple reminds developers about ATT requirements ahead of iOS 14.5 release, April 2021. URL <https://9to5mac.com/2021/04/20/apple-reminds-developers-about-att-requirements/>. April 20, 2021; quotes Apple developer notice. Accessed 2026-10-20.
- A2A Project. Agent2Agent Protocol Specification, 2026. URL <https://a2a-protocol.org/v1.0.1/specification/>. Accessed October 20, 2026.
- Josh Aas. Entering public beta. Let’s Encrypt, December 3, 2015, 2015. URL <https://letsencrypt.org/2015/12/03/entering-public-beta.html>. Accessed 2026-10-19.
- ABA Banking Journal. Kentucky federal court enjoins CFPB from enforcing current 1033 final rule. ABA Banking Journal, November 3, 2025, 2025. URL <https://bankingjournal.aba.com/2025/11/kentucky-federal-court-enjoins-cfpb-from-enforcing-current-1033-final-rule/>. Accessed 2026-10-19.
- Eric Abrahamson and Lori Rosenkopf. Institutional and competitive bandwagons: Using mathematical modeling as a tool to explore innovation diffusion. *Academy of Management Review*, 18 (3):487–517, 1993.
- Adobe. Adobe Completes Semrush Acquisition, Strengthening CX Enterprise with Enhanced Brand Visibility Capabilities, 2026. URL <https://www.semrush.com/news/455953-adobe-completes-semrush-acquisition-strengthening-cx-enterprise-with-enhanced-brand-visibility-capabilities>. 2026-04-28. Accessed 2026-10-19.
- Adobe Inc. and Semrush Holdings, Inc. Adobe to acquire semrush. Semrush Newsroom, 2025. URL <https://www.semrush.com/news/435826-adobe-to-acquire-semrush/>. Published 2025-11-19. Accessed 2026-10-20.
- Akamai Technologies. Akamai completes acquisition of Prolexic. Press release, February 18, 2014; announced December 2, 2013 at about \$370 million, 2014. URL <https://www.ir.akamai.com/news-releases/news-release-details/akamai-completes-acquisition-prolexic>. Accessed 2026-10-19.
- Alphabet Inc. Alphabet announces fourth quarter and fiscal year 2025 results, february 4, 2026. Alphabet Investor Relations, 2026. URL https://s206.q4cdn.com/479360582/files/doc_financials/2025/q4/2025q4-alphabet-earnings-release.pdf. Accessed 2026-10-20.
- Amazon. Amazon unveils the next generation of fulfillment centers powered by AI and 10 times more robotics. About Amazon, 2024. URL <https://www.aboutamazon.com/news/operations/amazon-fulfillment-center-robotics-ai>. The 25 percent peak-season cost-to-serve improvement is a target, not a measured industry result. Accessed 2026-10-20.
- Amazon.com, Inc. Amazon.com announces record free cash flow fueled by lower prices and free shipping; introduces new express shipping program: Amazon prime, press release (exhibit 99.1 to form 8-k), february 2, 2005. SEC EDGAR, 2005. URL <https://www.sec.gov/Archives/edgar/data/1018724/000119312505017451/dex991.htm>. Accessed 2026-10-20.

Amazon.com, Inc. Amazon.com, Inc. Form 10-K for fiscal 2024 — Financial Report (XBRL), Segment Information, February 2025. URL https://www.sec.gov/Archives/edgar/data/1018724/000101872425000004/Financial_Report.xlsx. Advertising services net sales 2022–2024. Accessed 2026-10-20.

Amazon.com, Inc. Amazon.com Announces Fourth Quarter Results, 2026a. URL <https://ir.aboutamazon.com/news-release/news-release-details/2026/Amazon-com-Announces-Fourth-Quarter-Results/default.aspx>. February 5; year ended December 31, 2025; accessed 2026-10-20.

Amazon.com, Inc. Amazon rolls out amazon now to dozens of cities across the u.s., providing 30-minute delivery to millions of customers, press release, may 12, 2026. Amazon US Press Center / Business Wire, 2026b. URL <https://press.aboutamazon.com/2026/5/amazon-rolls-out-amazon-now-to-dozens-of-cities-across-the-u-s-providing-30-minute-delivery-to-millions-of-customers>. Accessed 2026-10-20.

Amazon.com, Inc. Amazon sets new prime delivery speed record in 2025 with over 13 billion items arriving the same or next day around the world, press release, february 3, 2026. Amazon US Press Center / Business Wire, 2026c. URL <https://press.aboutamazon.com/2026/2/amazon-sets-new-prime-delivery-speed-record-in-2025-with-over-13-billion-items-arriving-the-same-or-next-day-around-the-world>. Accessed 2026-10-20.

Amazon.com, Inc. Amazon.com Announces Fourth Quarter Results (Exhibit 99.1 to Form 8-K), February 2026d. URL <https://www.sec.gov/Archives/edgar/data/1018724/000101872426000002/amzn-20251231xex991.htm>. February 5, 2026. Accessed 2026-10-20.

Amazon.com, Inc. Amazon sets new Prime delivery speed record in 2025 with over 13 billion items arriving the same or next day around the world. About Amazon, February 2026e. URL <https://aboutamazon.com/news/retail/amazon-prime-same-day-next-day-delivery-2025>. February 3, 2026. Accessed 2026-10-20.

Anthropic. Claude 3.7 Sonnet and Claude Code, 2025. URL <https://www.anthropic.com/news/claude-3-7-sonnet>. Accessed 2026-10-20; primary company release.

Anthropic. Anthropic Economic Index: New Building Blocks for Understanding AI Use, 2026a. URL <https://www.anthropic.com/research/economic-index-primitives>. Source date: 2026-01-15. Accessed October 20, 2026.

Anthropic. Detecting and Countering Misuse of AI: September 2026, 2026b. URL <https://www.anthropic.com/threat-intelligence-report-september-2026>. Source date: 2026-09-10. Accessed October 20, 2026.

Anysphere, Inc. Series b, january 16, 2025. Cursor Blog, 2025a. URL <https://cursor.com/blog/series-b>. Accessed 2026-10-20.

Anysphere, Inc. Series d, november 13, 2025. Cursor Blog, 2025b. URL <https://cursor.com/blog/series-d>. Accessed 2026-10-20.

Anza. 02-06-24 Solana Mainnet Beta Outage Report, 2024. URL <https://solana.com/news/02-06-24-solana-mainnet-beta-outage-report>. Accessed 2026-10-20; primary developer incident report.

- Apple Inc. Declared age range. Apple Developer Documentation; developer news item of December 23, 2025, 2025. URL <https://developer.apple.com/documentation/declaredagerange>. Accessed 2026-10-19.
- AppLovin Corporation. AppLovin Corporation Form S-1/A (Registration Statement, Amendment), March 2021. URL <https://www.sec.gov/Archives/edgar/data/1751008/000119312521089656/d73804ds1a.htm>. March 2021; Adjust acquisition terms. Accessed 2026-10-20.
- Stuart Armstrong, Nick Bostrom, and Carl Shulman. Racing to the precipice: A model of artificial intelligence development. *AI & Society*, 31(2):201–206, 2016.
- Kenneth J. Arrow. The economic implications of learning by doing. *Review of Economic Studies*, 29(3):155–173, 1962. doi: 10.2307/2295952.
- W. Brian Arthur. Competing technologies, increasing returns, and lock-in by historical events. *Economic Journal*, 99(394):116–131, 1989.
- Associated Press. FINA bans record-setting bodysuits. Associated Press, via Sun Journal, July 25, 2009, 2009. URL <https://www.sunjournal.com/2009/07/25/fina-bans-record-setting-bodysuits/>. Accessed 2026-10-19.
- Associated Press. Chegg’s stock plunges on fears of competition from chatgpt, may 2, 2023. AP via News4Jax, 2023. URL <https://www.news4jax.com/business/2023/05/02/cheggs-stock-plunges-on-fears-of-competition-from-chatgpt/>. Accessed 2026-10-20.
- Autorité de la concurrence. Targeted advertising: the Autorité de la concurrence imposes a fine of €150,000,000 on Apple, March 2025. URL <https://www.autoritedelaconcurrence.fr/en/article/targeted-advertising-autorite-de-la-concurrence-imposes-fine-eu150000000-apple>. March 31, 2025. Accessed 2026-10-20.
- Avalara, Inc. Prospectus (form 424b4), june 14, 2018. U.S. Securities and Exchange Commission, EDGAR, 2018. URL <https://www.sec.gov/Archives/edgar/data/0001348036/000119312518193520/d317509d424b4.htm>. Revenue 2015–2017 and the \$24.00 offer price. Accessed 2026-10-19.
- Avalara, Inc. Form 10-k for the fiscal year ended december 31, 2019. U.S. Securities and Exchange Commission, EDGAR, 2020. URL https://www.sec.gov/Archives/edgar/data/1348036/000156459020007791/avlr-10k_20191231.htm. Also cites the 2018 10-K, https://www.sec.gov/Archives/edgar/data/1348036/000156459019005454/avlr-10k_20181231.htm. Accessed 2026-10-19.
- Avalara, Inc. Form 10-k for the fiscal year ended december 31, 2021. U.S. Securities and Exchange Commission, EDGAR, 2022. URL https://www.sec.gov/Archives/edgar/data/1348036/000156459022006696/avlr-10k_20211231.htm. Accessed 2026-10-19.
- Avery Dennison Corporation. First Quarter 2024 Financial Review and Analysis (Exhibit 99.2 to Form 8-K), April 2024. URL <https://www.sec.gov/Archives/edgar/data/8818/000119312524108400/d827475dex992.htm>. April 24, 2024. Accessed 2026-10-20.
- Zineb Ait Bahajji and Gary Illyes. HTTPS as a ranking signal. Google Online Security Blog, August 7, 2014, 2014. URL https://security.googleblog.com/2014/08/https-as-ranking-signal_6.html. Accessed 2026-10-19.

- Ballpark Digest. Hawk-eye innovations, mlb announce tracking-platform deal. Ballpark Digest, 2020. URL <https://ballparkdigest.com/2020/08/20/hawk-eye-innovations-mlb-annouce-tracking-platform-deal/>. Published 2020-08-20. Accessed 2026-10-20.
- Bank for International Settlements. Project Hertha, 2025. URL <https://www.bis.org/project/hertha>. Accessed October 20, 2026.
- Bank for International Settlements. Project Agora, 2026. URL <https://www.bis.org/project/agora>. Source date: 2026-05-27. Accessed October 20, 2026.
- William P. Barnett and Morten T. Hansen. The red queen in organizational evolution. *Strategic Management Journal*, 17(S1):139–157, 1996. URL <https://www.gsb.stanford.edu/faculty-research/publications/red-queen-organizational-evolution>.
- Frank M. Bass. A new product growth for model consumer durables. *Management Science*, 15(5): 215–227, 1969.
- Joel Becker, Nate Rush, Tom Cunningham, David Rein, and Khalid Mahamud. We Are Changing Our Developer Productivity Experiment Design, 2026. URL <https://metr.org/blog/2026-02-24-uplift-update/>. METR, February 24. Accessed 2026-10-20; primary experimental research update.
- Belgian Federal Government. Pour qui la facturation electronique deviendra-t-elle obligatoire?, 2024. URL <https://efacture.belgium.be/fr/article/pour-qui-la-facturation-electronique-deviendra-t-elle-obligatoire>. 2024-10-09; accessed version 2026-10-20. Accessed 2026-10-20.
- Luisa Beltran. Fintech Plaid raises \$575 million at a \$6.1 billion valuation, less than half its 2021 peak. *Fortune*, April 3, 2025, 2025. URL <https://www.fortune.com/2025/04/03/fintech-plaid-funding-round-fidelity-blackrock-franklin-templeton-nea-perret>. Accessed 2026-10-19.
- Bessemer Venture Partners. The official state of the cloud: Ai, august 13, 2025. bvp.com, 2025. URL <https://www.bvp.com/the-official-state-of-the-cloud-ai>. Accessed 2026-10-20.
- Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, 1992. doi: 10.1086/261849.
- Biometric Update. Socure fields 2.7 billion identity verification requests in 2024, 2025. URL <https://www.biometricupdate.com/202502/socure-fields-2-7-billion-identity-verification-requests-in-2024>. 4 Feb 2025. Accessed 2026-10-19.
- Biometric Update. Discord launches age-adaptive design, says 90% of users won’t need active checks. Biometric Update, September 23, 2026, 2026a. URL <https://www.biometricupdate.com/202609/discord-launches-age-adaptive-design-says-90-of-users-wont-need-active-checks>. Accessed 2026-10-19.
- Biometric Update. Deepfakes are overwhelming businesses unprepared to deal with AI threat: Pindrop. Biometric Update, September 28, 2026, reporting Pindrop’s Deepfake Readiness Index, 2026b. URL <https://www.biometricupdate.com/202609/deepfakes-are-overwhelming-businesses-unprepared-to-deal-with-ai-threat-pindrop>. Accessed 2026-10-19.

- Biometric Update. Socure hits \$364m ARR in Q2 on strong cross-sector growth. Biometric Update, August 12, 2026, reporting Socure’s Q2 2026 results release, 2026c. URL <https://www.biometricupdate.com/202608/socure-hits-364m-arr-in-q2-on-strong-cross-sector-growth>. Accessed 2026-10-19.
- Biometric Update. Socure acquires Fravity, raises \$156m to bring agentic AI to fraud investigations. Biometric Update, August 27, 2026, 2026d. URL <https://www.biometricupdate.com/202608/socure-acquires-fravity-raises-156m-to-bring-agentic-ai-to-fraud-investigations>. Accessed 2026-10-19.
- Bitcoin Developer Documentation. Mining, 2026. URL <https://developer.bitcoin.org/devguide/mining.html>. Accessed 2026-10-20; primary technical documentation.
- Bloomberg Law. JPMorgan tells fintechs to pay up for customer data access. Bloomberg Law, 2025a. URL <https://news.bloomberglaw.com/banking-law/jpmorgan-tells-fintechs-they-have-to-pay-up-for-customer-data>. Accessed 2026-10-19.
- Bloomberg Law. Vanta notches \$4.15 billion valuation with new funding round. Bloomberg Government, July 23, 2025, 2025b. URL <https://news.bgov.com/private-equity/vanta-notches-4-15-billion-valuation-with-new-funding-round>. Accessed 2026-10-19.
- BloombergNEF. Lithium-Ion Battery Pack Prices See Largest Drop Since 2017, Falling to \$115 per Kilowatt-Hour, December 2024. URL <https://about.bnef.com/insights/commodities/lithium-ion-battery-pack-prices-see-largest-drop-since-2017-falling-to-115-per-kilowatt-hour-bloombergnef/>. December 10, 2024. Accessed 2026-10-20.
- BloombergNEF. Lithium-Ion Battery Pack Prices Fall to \$108 per Kilowatt-Hour Despite Rising Metal Prices, December 2025. URL <https://about.bnef.com/insights/commodities/lithium-ion-battery-pack-prices-fall-to-108-per-kilowatt-hour-despite-rising-metal-prices/>. December 9, 2025. Accessed 2026-10-20.
- BlueGamma. Bloomberg terminal cost 2026: About \$32,000 per user a year. BlueGamma blog, 2026. URL <https://www.bluegamma.io/post/bloomberg-terminal-pricing>. Published 2026-10-07. Accessed 2026-10-20.
- Board of Governors of the Federal Reserve System. Economic Well-Being of U.S. Households in 2025, 2026. URL <https://www.federalreserve.gov/publications/files/2025-report-economic-well-being-us-households-202605.pdf>. Accessed 2026-10-20; primary household survey.
- Rishi Bommasani et al. On the Opportunities and Risks of Foundation Models, 2021. URL <https://arxiv.org/abs/2108.07258>. Accessed 2026-10-20; primary research report.
- Masha Borak. Yoti records 62% revenue growth in 2025. Biometric Update, January 7, 2026, 2026. URL <https://www.biometricupdate.com/202601/yoti-records-62-revenue-growth-in-2025>. Accessed 2026-10-19.
- William A. Brock and Steven N. Durlauf. Discrete choice with social interactions. *Review of Economic Studies*, 68(2):235–260, 2001.
- Erik Brynjolfsson, Danielle Li, and Lindsey Raymond. Generative AI at work. *Quarterly Journal of Economics*, 140(2):889–942, 2025. doi: 10.1093/qje/qjae044. URL <https://academic.oup.com/qje/article/140/2/889/7990658>.

- Eric Budish. The economic limits of Bitcoin and the blockchain. Working Paper 24717, National Bureau of Economic Research, 2018. URL <https://www.nber.org/papers/w24717>.
- Eric Budish, Robin S. Lee, and John J. Shim. A theory of stock exchange competition and innovation: Will the market fix the market? NBER Working Paper 25855, May 2019, revised November 2022, 2019. URL https://www.nber.org/system/files/working_papers/w25855/w25855.pdf. Accessed 2026-10-19.
- Jeremy I. Bulow, John D. Geanakoplos, and Paul D. Klemperer. Multimarket oligopoly: Strategic substitutes and complements. *Journal of Political Economy*, 93(3):488–511, 1985.
- Chris Burt. Incode seeks fundraise at \$3b valuation to seize emerging biometrics opportunities. Biometric Update, November 20, 2025, citing Bloomberg, 2025. URL <https://www.biometricupdate.com/202511/incode-seeks-fundraise-at-3b-valuation-to-seize-emerging-biometrics-opportunities>. Accessed 2026-10-19.
- Burton-Taylor International Consulting. Global spend on financial market data & news 2016: market sizing press release. Burton-Taylor International Consulting (hosted by TP ICAP), 2017. URL <https://tpicap.com/tpicap/sites/g/files/escbpb106/files/2020-05/2017%20Market%20Sizing%20press%20release%20-%20vFINAL.pdf>. Published 2017-03-28. Accessed 2026-10-20.
- Burton-Taylor International Consulting. Financial market data/analysis: Global share & segment sizing 2026. TP ICAP / Burton-Taylor, 2026. URL <https://tpicap.com/burtontaylor/reports/03/2026/financial-market-dataanalysis-global-share-segment-sizing-2026>. Published 2026-03-26. Accessed 2026-10-20.
- Vitalik Buterin, Eric Conner, Rick Dudley, Matthew Slipper, Ian Norden, and Abdelhamid Bakhta. EIP-1559: Fee Market Change for ETH 1.0 Chain, 2019. URL <https://eips.ethereum.org/EIPS/eip-1559>. Accessed 2026-10-20; primary protocol specification.
- Vitalik Buterin et al. EIP-4844: Shard Blob Transactions, 2022. URL <https://eips.ethereum.org/EIPS/eip-4844>. Accessed 2026-10-20; primary protocol specification.
- CA/Browser Forum. SC081v3: reducing certificate validity and data reuse periods, 2025. URL <https://cabforum.org/2025/04/11/ballot-sc081v3-introduce-schedule-of-reducing-validity-and-data-reuse-periods/>. Accessed October 20, 2026.
- California Department of Insurance. Safer from Wildfires, 2026. URL <https://www.insurance.ca.gov/01-consumers/200-wrr/Safer-from-Wildfires.cfm>. Accessed October 20, 2026.
- Hans Carlsson and Eric van Damme. Global games and equilibrium selection. *Econometrica*, 61(5):989–1018, 1993.
- Mark Cartwright. The anglo-german arms race: Riding the waves to wwi. World History Encyclopedia, 2025. URL <https://www.worldhistory.org/article/2812/the-anglo-german-arm-s-race/>. Published 2025-10-01. Accessed 2026-10-20.
- C.H. Robinson. C.H. Robinson scales fleet of AI agents past 30, 2025. URL <https://www.chrobinson.com/fr-fr/about-us/newsroom/news/2025/ch-robinson-scales-fleet-of-ai-agents-past-30/>. Accessed October 20, 2026.
- C.H. Robinson. C.H. Robinson reports 2026 second quarter results, 2026. URL <https://investor.chrobinson.com/Financials/Quarterly-Results/default.aspx>. Accessed October 20, 2026.

Chain Store Age. Deloitte: Holiday shoppers hooked on these two shipping preferences, november 1, 2018. Chain Store Age (reporting Deloitte 33rd annual Holiday Survey), 2018. URL <https://chainstoreage.com/operations/deloitte-holiday-shoppers-hooked-on-these-two-shipping-preferences>. Accessed 2026-10-20.

Athena Chapekis and Anna Lieb. Google users are less likely to click on links when an AI summary appears in the results, 2025. URL <https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/>. Pew Research Center, 2025-07-22. Accessed 2026-10-19.

Chegg, Inc. Chegg reports q1 2023 earnings (exhibit 99.1 to form 8-k), press release, may 1, 2023. SEC EDGAR, 2023. URL <https://www.sec.gov/Archives/edgar/data/1364954/000136495423000067/a9901-financialresultsq120.htm>. Accessed 2026-10-20.

Circle / SEC. Circle Form 10-Q, quarter ended June 30, 2026, 2026. URL <https://www.sec.gov/Archives/edgar/data/1876042/000187604226000248/crc1-20260630.htm>. Accessed October 20, 2026.

Circle Internet Group. Circle Reports Second Quarter 2026 Results, 2026. URL <https://www.circle.com/pressroom/circle-reports-second-quarter-2026-results>. Source date: 2026-08-05. Accessed October 20, 2026.

Cloudflare. 2025 q4 DDoS threat report: A record-setting 31.4 tbps attack caps a year of massive DDoS assaults. Cloudflare blog, February 5, 2026, 2026. URL <https://blog.cloudflare.com/ddos-threat-report-2025-q4/>. Accessed 2026-10-19.

Cloudflare, Inc. Cloudflare announces fourth quarter and fiscal year 2025 financial results. Press release, February 10, 2026; 2019 revenue from the February 13, 2020 release, 2026. URL <https://cloudflare.net/news/news-details/2026/Cloudflare-Announces-Fourth-Quarter-and-Fiscal-Year-2025-Financial-Results/default.aspx>. Accessed 2026-10-19.

CnEVPost. CAAM forecasts China NEV sales to grow 15.2% to 19 million units in 2026, January 2026. URL <https://cnevpost.com/2026/01/14/caam-forecasts-china-nev-sales-19-million-2026/>. January 14, 2026; CAAM data for 2025. Accessed 2026-10-20.

Coalition for Content Provenance and Authenticity. The C2PA launches content credentials 2.3 and celebrates 5 years of impact across the digital ecosystem, 2026. URL <https://c2pa.org/the-c2pa-launches-content-credentials-2-3-and-celebrates-5-years-of-impact-across-the-digital-ecosystem/>. February announcement, including the conformance program. Accessed October 16, 2026.

Coalition for Content Provenance and Authenticity. C2PA and content credentials explainer, n.d. URL <https://spec.c2pa.org/specifications/specifications/2.2/explainer/Explainer.html>. Version 2.2. Accessed October 16, 2026.

Cognizant Technology Solutions Corp. Form 10-k for the fiscal year ended december 31, 2000. U.S. Securities and Exchange Commission, EDGAR, filed March 15, 2001, 2001. URL <https://www.sec.gov/Archives/edgar/data/1058290/000090310001000038/0001.txt>. Accessed 2026-10-19.

Coinbase. CDP SQL API Now Accepts x402 Payments: Onchain Data for Autonomous Agents, 2026. URL <https://www.coinbase.com/developer-platform/discover/launches/sql-api-x402>. Source date: 2026-02-12. Accessed October 20, 2026.

Core Scientific, Inc. Annual Report on Form 10-K for the Year Ended December 31, 2025, 2026. URL <https://www.sec.gov/Archives/edgar/data/1839341/000162828026013305/core-20251231.htm>. Accessed 2026-10-20; primary annual filing.

Coskun et al. Turning AI data centers into grid-interactive assets: field demonstration in Phoenix, 2025. URL <https://arxiv.org/abs/2507.00909>. Accessed October 20, 2026.

Criteo S.A. Criteo Reports Fourth Quarter 2025 Results, February 2026. URL <https://criteo.investorroom.com/2026-02-11-CRITEO-REPORTS-FOURTH-QUARTER-2025-RESULTS>. February 11, 2026. Accessed 2026-10-20.

CrowdStrike Holdings, Inc. Form 10-k for the fiscal year ended january 31, 2021. U.S. Securities and Exchange Commission, EDGAR, 2021. URL <https://www.sec.gov/Archives/edgar/data/1535527/000153552721000007/crwd-20210131.htm>. Accessed 2026-10-20.

CrowdStrike Holdings, Inc. Crowdstrike reports fourth quarter and fiscal year 2025 financial results (form 8-k, exhibit 99.1). U.S. Securities and Exchange Commission, EDGAR, 2025. URL <https://www.sec.gov/Archives/edgar/data/1535527/000153552725000005/crwd-20250304xex991.htm>. Dated 2025-03-04. Accessed 2026-10-20.

CrowdStrike Holdings, Inc. Crowdstrike reports second quarter fiscal year 2027 financial results. CrowdStrike Investor Relations, 2026. URL <https://ir.crowdstrike.com/news-releases/news-release-details/crowdstrike-reports-second-quarter-fiscal-year-2027-financial>. Dated 2026-08-26. Accessed 2026-10-20.

Crunchbase News. Bessemer state of the cloud 2023: ‘every single category of saas software’ will be changed by ai, april 12, 2023. Crunchbase News, 2023. URL <https://news.crunchbase.com/cloud/chatgpt-ai-public-private-cloud-bessemer/>. Accessed 2026-10-20.

CyberArk Software Ltd. CyberArk to acquire Venafi (exhibit 99.1 to form 6-k), may 20, 2024. U.S. Securities and Exchange Commission, EDGAR; completion announced October 1, 2024, 2024. URL https://www.sec.gov/Archives/edgar/data/1598110/000117891324001780/exhibit_99-1.htm. Accessed 2026-10-19.

CYBRA Corporation. What You Need to Know About Walmart’s Latest RFID Tagging Mandate, 2026. URL <https://cybra.com/what-you-need-to-know-about-walmarts-latest-rfid-mandate/>. First published April 27, 2022; updated March 12, 2026. Accessed 2026-10-20.

Dai et al. Autonomous mobile robots for exploratory synthetic chemistry, 2024. URL <https://www.nature.com/articles/s41586-024-08173-7>. Accessed October 20, 2026.

Decagon. Decagon’s \$250 million commitment to the AI concierge future, 2026. URL <https://decagon.ai/blog/series-d-announcement>. Published 2026-01-27. Accessed 2026-10-19.

DeepSeek-AI. DeepSeek-V3 technical report. arXiv:2412.19437v1 [cs.CL], December 27, 2024, 2024. URL <https://arxiv.org/abs/2412.19437v1>. Accessed 2026-10-20.

Deeptrace. The state of deepfakes: Landscape, threats, and impact, september 2019. Deeptrace (PDF hosted by The Register), 2019. URL https://regmedia.co.uk/2019/10/08/deepfake_report.pdf. Accessed 2026-10-20.

Digital Container Shipping Association. Soft barriers to eBL adoption: Addressing the hidden frictions that slow digital progress, n.d. URL <https://dcsa.org/newsroom/understanding-soft-barriers-to-digital-adoption>. Reports a 2025 electronic-issuance estimate. Accessed October 16, 2026.

- Paul J. DiMaggio and Walter W. Powell. The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2):147–160, 1983.
- Anne D’Innocenzio. Target steps up next-day parcel delivery as discounter tries to narrow gap with rivals. Associated Press, via Barchart, September 2025. URL <https://www.barchart.com/story/news/34831135/target-steps-up-next-day-parcel-delivery-as-discounter-tries-to-narrow-gap-with-rivals>. September 16, 2025. Accessed 2026-10-20.
- Avinash K. Dixit and Robert S. Pindyck. *Investment under Uncertainty*. Princeton University Press, Princeton, 1994.
- Drata. Drata celebrates four years of transforming governance, risk and compliance, crossing \$100 million in annual recurring revenue. Press release, February 19, 2025, 2025. URL <https://drata.com/about/news/drata-celebrates-four-years-of-transforming-governance-risk-and-compliance-crossing-100-million-in-annual-recurring-revenue>. Accessed 2026-10-19.
- Scott Dresser. Amazon launches a new AI foundation model to power its robotic fleet and deploys its 1 millionth robot, 2025. URL <https://www.aboutamazon.com/news/operations/amazon-million-robots-ai-foundation-model>. Amazon, June 30. Company deployment report; accessed October 16, 2026.
- edgar.tools. Impinj Inc (PI) – Financials and Filings (10-K revenue 2021–2025; DEF 14A 2026), 2026. URL <https://app.edgar.tools/companies/PI>. Aggregator of SEC filings; FY2025 10-K filed February 9, 2026. Accessed 2026-10-20.
- Jr. Edward W. Kelley. The Federal Reserve’s efforts to address the year 2000 computer problem. Testimony before the Senate Committee on Commerce, Science, and Transportation, April 28, 1998, 1998. URL <https://www.federalreserve.gov/boarddocs/testimony/1998/19980428.htm>. Accessed 2026-10-19.
- Electric Capital. Developer Report: 2024 Analysis, 2024. URL <https://www.developerreport.com/developer-report>. Accessed 2026-10-20; primary research summary.
- Electric Capital. Developer Report: About and Methodology, 2026. URL <https://www.developerreport.com/about>. Accessed 2026-10-20; primary measurement methodology.
- Electric Reliability Council of Texas. Updates to ADER pilot participation limits, 2026. URL https://www.ercot.com/services/comm/mkt_notices/M-A030226-01. Market notice M-A030226-01, March 2. Accessed October 16, 2026.
- eMarketer. Faq on programmatic advertising: Keeping up with automated ad buying. eMarketer, 2026. URL <https://www.emarketer.com/content/faq-on-programmatic-advertising--keeping-up-with-automated-ad-buying>. Published 2026-01-26. Accessed 2026-10-20.
- Entrust. 2026 identity fraud report. Entrust, released November 19, 2025, 2025. URL <https://www.entrust.com/resources/reports/identity-fraud-report>. Accessed 2026-10-19.
- Ethereum.org Contributors. Smart Contract Composability, 2026a. URL <https://ethereum.org/developers/docs/smart-contracts/composability/>. Accessed 2026-10-20; primary project documentation.
- Ethereum.org Contributors. The Merge, 2026b. URL <https://ethereum.org/roadmap/merge/>. Accessed 2026-10-20; primary project documentation.

- European Commission. VAT in the Digital Age (ViDA), 2025. URL https://taxation-customs.ec.europa.eu/taxation/vat/vat-digital-age-vida_en. Accessed 2026-10-19.
- European Commission. eInvoicing in Italy. Digital Building Blocks country factsheet, updated September 30, 2026, 2026. URL <https://ec.europa.eu/digital-building-blocks/sites/display/DIGITAL/eInvoicing+in+Italy>. Accessed 2026-10-19.
- European Parliament and Council of the European Union. Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 (Digital Omnibus on AI), 2026. URL https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=OJ:L_202601744. OJ L, 2026/1744, 24.7.2026; in force 27 July 2026. Accessed 2026-10-19.
- Eurostat. Hourly labour costs — Statistics Explained, 2026a. URL https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Hourly_labour_costs. Data extracted 31 March 2026. Accessed 2026-10-19.
- Eurostat. Labour cost levels by NACE Rev. 2 activity (lc_lci_lev), 2026b. URL https://ec.europa.eu/eurostat/api/dissemination/statistics/1.0/data/lc_lci_lev?format=JSON&lang=EN&geo=FR&nace_r2=B-S_X_0&lcstruct=D1_D4_MD5&time=2025. Dataset updated 2026-04-23. Accessed 2026-10-19.
- Stephen J. Farr, Carol Horton Tremblay, and Victor J. Tremblay. The welfare effect of advertising restrictions in the U.S. cigarette industry. *Review of Industrial Organization*, 18(2):147–160, 2001. doi: 10.1023/A:1007801608741. URL <https://link.springer.com/article/10.1023/A:1007801608741>.
- Joseph Farrell and Paul Klemperer. Coordination and lock-in: Competition with switching costs and network effects. In Mark Armstrong and Robert Porter, editors, *Handbook of Industrial Organization*, volume 3, pages 1967–2072. Elsevier, Amsterdam, 2007. doi: 10.1016/S1573-448X(06)03031-7.
- Joseph Farrell and Garth Saloner. Standardization, compatibility, and innovation. *RAND Journal of Economics*, 16(1):70–83, 1985.
- Joseph Farrell and Garth Saloner. Installed base and compatibility: Innovation, product preannouncements, and predation. *American Economic Review*, 76(5):940–955, 1986.
- Federal Reserve. About the FedNow Service, 2026. URL https://www.federalreserve.gov/paymentsystems/fednow_about.htm. Accessed October 20, 2026.
- Federal Reserve Financial Services. FedNow service volume and value statistics, 2026. URL <https://www.frbservices.org/resources/financial-services/fednow-service-resources/fednow-service-volume-value-statistics/>. Updated July 6. Quarterly series through 2026 Q2; accessed October 16, 2026.
- Federal Trade Commission Staff. Partnerships Between Cloud Service Providers and AI Developers, 2025. URL https://www.ftc.gov/system/files/ftc_gov/pdf/p246201_aipartnerships6breport_redacted_0.pdf. Accessed 2026-10-20; primary regulatory research report.
- FIDO Alliance. Passkey index 2026, 2026. URL <https://fidoalliance.org/passkey-index-2026/>. October 15. Index of participating services, not all internet sign-ins; accessed October 16, 2026.

- Fierce Wireless. Report: Nokia’s smartphone market share dropping, march 11, 2009. Fierce Wireless (Fierce Network), 2009. URL <https://www.fierce-network.com/wireless/report-nokia-s-smartphone-market-share-dropping>. Accessed 2026-10-20; relays Gartner data.
- FINA. PR59: FINA bureau meeting, 2009. URL <https://www.worldaquatics.com/news/1914351/pr59-fina-bureau-meeting>. July 31. Swimsuit rules effective January 1, 2010.
- Flashbots. MEV-Share introduction, 2026. URL <https://docs.flashbots.net/flashbots-mev-share/introduction>. Accessed October 20, 2026.
- Flexport. Why we bought, built, and sold the Convoy Platform, 2025. URL <https://www.flexport.com/blog/why-we-bought-built-and-sold-the-convoy-platform/>. Accessed October 20, 2026.
- Formic. 3D Machine finds new competitive edge with automation partner Formic, 2023. URL <https://formic.co/resources/customer-stories/3d-machine-finds-new-competitive-edge-with-automation-partner-formic>. Accessed October 20, 2026.
- Fortune. Microsoft cancels leases for ai data centers, analyst says, february 24, 2025. Fortune, 2025. URL <https://fortune.com/2025/02/24/microsoft-cancels-leases-for-ai-data-centers-analyst-says>. Accessed 2026-10-20.
- Robert H. Frank. *The Darwin Economy: Liberty, Competition, and the Common Good*. Princeton University Press, Princeton, 2011.
- Robert H. Frank and Philip J. Cook. *The Winner-Take-All Society*. Free Press, New York, 1995.
- Drew Fudenberg and Jean Tirole. Preemption and rent equalization in the adoption of new technology. *Review of Economic Studies*, 52(3):383–401, 1985.
- Andreas Fuster, Matthew Plosser, Philipp Schnabl, and James Vickery. The role of technology in mortgage lending. NBER Working Paper 24500, 2018. URL <https://www.nber.org/papers/w24500>. Accessed October 20, 2026.
- Andrea Galeotti, Sanjeev Goyal, Matthew O. Jackson, Fernando Vega-Redondo, and Leeat Yariv. Network games. *Review of Economic Studies*, 77(1):218–244, 2010.
- Gartner, Inc. Gartner predicts search engine volume will drop 25% by 2026, due to ai chatbots and other virtual agents, february 19, 2024. Gartner Newsroom, 2024. URL <https://www.gartner.com/en/newsroom/press-releases/2024-02-19-gartner-predicts-search-engine-volume-will-drop-25-percent-by-2026-due-to-ai-chatbots-and-other-virtual-agents>. Accessed 2026-10-20.
- Gartner, Inc. Forecast analysis: Information security, worldwide, 2026. Gartner research document (abstract), 2026. URL <https://www.gartner.com/en/documents/7408930>. Published 2026-02-05. Accessed 2026-10-20.
- GeekWire. Amazon stock sinks 10% after q4 profit miss as jassy signals \$200b in capital spending, february 5, 2026. GeekWire, 2026. URL <https://www.geekwire.com/2026/amazon-stock-sinks-10-after-q4-profit-miss-as-jassy-signals-200b-in-capital-spending/>. Accessed 2026-10-20.
- German Federal Ministry of Finance. Fragen und Antworten zur Einfuehrung der obligatorischen E-Rechnung zum 1. Januar 2025, 2026. URL <https://www.bundesfinanzministerium.de/Content/DE/FAQ/e-rechnung.html>. 2026-03-23. Accessed 2026-10-20.

- Sharon Goldberg and Amos Paul. Post-quantum encryption for Cloudflare IPsec is generally available, 2026. URL <https://blog.cloudflare.com/post-quantum-ipsec/>. Cloudflare, April 30. Provider-specific traffic and product report; accessed October 16, 2026.
- Google. Stable channel update for desktop (Chrome 68.0.3440.75). Chrome Releases blog, July 24, 2018, 2018. URL <https://chromereleases.googleblog.com/2018/07/stable-channel-update-for-desktop.html>. Accessed 2026-10-19.
- Google. How we are making data centers more flexible to benefit power grids, 2025. URL <https://blog.google/innovation-and-ai/infrastructure-and-cloud/global-network/how-we-are-making-data-centers-more-flexible-to-benefit-power-grids/>. Accessed October 20, 2026.
- Google. Agent Payments Protocol donation to FIDO Alliance, 2026. URL <https://blog.google/products-and-platforms/platforms/google-pay/agent-payments-protocol-fido-alliance/>. Accessed October 20, 2026.
- Google Workspace. The best of Google AI, now included in Workspace business and enterprise plans. Google Workspace blog, January 15, 2025, 2025. URL <https://workspace.google.com/blog/product-announcements/empowering-businesses-with-AI>. Accessed 2026-10-19.
- Grace’s Guide to British Industrial History. Armstrong whitworth. Grace’s Guide, 2026. URL https://www.gracesguide.co.uk/Armstrong_Whitworth. Version fetched 2026-10-20; quotes The Engineer and The Times. Accessed 2026-10-20.
- Mark Granovetter. Threshold models of collective behavior. *American Journal of Sociology*, 83(6): 1420–1443, 1978.
- Zvi Griliches. Hybrid corn: An exploration in the economics of technological change. *Econometrica*, 25(4):501–522, 1957.
- Omer Guy. The electronic bill of lading and the architecture of trust, 2026. URL <https://dcsa.org/newsroom/bel-interoperability-with-omer-guy>. DCSA. Reports the June 2026 interoperability milestone; accessed October 16, 2026.
- Stephen Hardy. Zayo to buy Spread Networks for \$127 million in cash. Lightwave, November 27, 2017, 2017. URL <https://www.lightwaveonline.com/business/mergers-acquisitions/article/16673519/zayo-to-buy-spread-networks-for-127-million-in-cash>. Accessed 2026-10-19.
- Christopher Harris and John Vickers. Perfect equilibrium in a model of a race. *Review of Economic Studies*, 52(2):193–209, 1985. doi: 10.2307/2297616.
- James Hickey. Walmart, Avery Dennison Expand RFID Use in Grocery Section. RFID Journal, October 2025. URL <https://www.rfidjournal.com/?p=224197>. October 30, 2025. Accessed 2026-10-20.
- Fred Hirsch. *Social Limits to Growth*. Harvard University Press, Cambridge, MA, 1976.
- Jordan Hoffmann et al. Training Compute-Optimal Large Language Models, 2022. URL <https://arxiv.org/abs/2203.15556>. Accessed 2026-10-20; primary research paper.
- Hyperliquid. HyperCore Overview, 2026a. URL <https://hyperliquid.gitbook.io/hyperliquid-docs/hypercore/overview>. Accessed 2026-10-20; primary protocol documentation.

Hyperliquid. HyperEVM, 2026b. URL <https://hyperliquid.gitbook.io/hyperliquid-docs/hyperevm>. Accessed 2026-10-20; primary protocol documentation.

Hyperliquid. Trading Fees, 2026c. URL <https://hyperliquid.gitbook.io/hyperliquid-docs/trading/fees>. Accessed 2026-10-20; primary protocol documentation.

Hyperliquid. HyperEVM Developer Documentation, 2026d. URL <https://hyperliquid.gitbook.io/hyperliquid-docs/for-developers/hyperevm>. Accessed 2026-10-20; primary protocol documentation.

ImageNet. ILSVRC2012 results, october 13, 2012. image-net.org, 2012. URL <https://www.image-net.org/challenges/LSVRC/2012/results.html>. Accessed 2026-10-20.

Impinj, Inc. Impinj Reports Fourth Quarter and Full Year 2025 Financial Results, February 2026. URL <https://investor.impinj.com/news/press-release/2026/Impinj-Reports-Fourth-Quarter-and-Full-Year-2025-Financial-Results/default.aspx>. February 5, 2026. Accessed 2026-10-20.

Indiana Office of Utility Consumer Counselor. Indiana Michigan Power rates and pending cases, 2026. URL <https://www.in.gov/oucc/featured-topics/indiana-michigan-power-i-and-m-rates/>. Accessed October 20, 2026.

Infinera. Infinera and Spread Networks Enable Sub-13.33 ms Route Between New York and Chicago. Company press release, June 22, 2010, 2010. URL <https://www.globenewswire.com/news-release/2010/06/22/1325985/0/en/infinera-and-spread-networks-enable-sub-13-33-ms-route-between-new-york-and-chicago.html>. Company-reported round-trip latency. Accessed 2026-10-20.

Inflection AI. Inflection-2: The next step up, november 22, 2023. Inflection AI, 2023. URL <https://inflection.ai/inflection-2>. Accessed 2026-10-20.

inside.lighting. Insights: Latest LED Adoption Report (DOE 2020 U.S. Lighting Market Characterization), May 2024. URL <https://inside.lighting/news/24-05/insights-latest-led-adoption-report>. May 14, 2024. Accessed 2026-10-20.

Intercom. Fin AI Agent outcomes, n.d. URL <https://www.intercom.com/help/en/articles/8205718-fin-ai-agent-outcomes>. Provider’s chat and email outcome definitions and pricing. Accessed October 16, 2026.

Intercontinental Exchange, Inc. Intercontinental Exchange reports strong full year 2025 results (exhibit 99.1 to form 8-k), february 5, 2026. U.S. Securities and Exchange Commission, EDGAR, 2026. URL https://www.sec.gov/Archives/edgar/data/1571949/000110465926010646/tm264989d1_ex99-1.htm. Accessed 2026-10-19.

International Energy Agency. Key Questions on Energy and AI: Executive Summary, 2026. URL <https://www.iea.org/reports/key-questions-on-energy-and-ai/executive-summary>. Source date: 2026. Accessed October 20, 2026.

International Federation of Robotics. Five Million Robots Now Operate in Factories Globally, 2026a. URL <https://ifr.org/ifr-press-releases/news/five-million-robots-now-operate-in-factories-globally>. Source date: 2026-09-24. Accessed October 20, 2026.

- International Federation of Robotics. Five million robots now operate in factories globally, 2026b. URL <https://ifr.org/ifr-press-releases/news/service-robots-on-the-rise-worldwide>. September 24. World Robotics 2026 release, reporting 2025 installations and stock; accessed October 16, 2026.
- International Renewable Energy Agency (IRENA). Renewable Power Generation Costs in 2024. Technical report, IRENA, Abu Dhabi, July 2025. URL <https://www.irena.org/Publications/2025/Jun/Renewable-Power-Generation-Costs-in-2024>. Summary PDF: https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2025/Jul/IRENA_TEC_RPGC_in_2024_Summary_2025.pdf. Accessed 2026-10-20.
- Internet Security Research Group. 2025 annual report. ISRG, 2025. URL <https://www.abetterinternet.org/documents/2025-ISRG-Annual-Report.pdf>. Accessed 2026-10-19.
- Investing.com. Earnings call transcript: Walmart q4 2026 beats eps forecast, stock rises, february 19, 2026. Investing.com, 2026. URL <https://www.investing.com/news/transcripts/earnings-call-transcript-walmart-q4-2026-beats-eps-forecast-stock-rises-93CH-4513893>. Accessed 2026-10-20.
- John Deere. Advancing weed control with John Deere’s See & Spray technology, 2025. URL <https://edge.machinefinderpro.deere.com/ww/en-US/articles/advancing-weed-control-with-john-deeres-see--spray-technology-13794>. MachineFinder, November 28. Manufacturer-reported 2025 season results; accessed October 16, 2026.
- J.P. Morgan. Tokenized Collateral Network, 2026. URL <https://cws-main-ndc.jpmorgan.com/kinexys/tokenized-collateral-network>. Accessed October 20, 2026.
- JPMorganChase. JPMorganChase and Plaid announce an extension to their data access agreement for sharing of consumer permissioned data. Press release, September 16, 2025, 2025. URL <https://www.jpmorganchase.com/newsroom/press-releases/2025/jpmc-plaid-renewed-data-access-agreement>. Accessed 2026-10-19.
- Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979. doi: 10.2307/1914185.
- Noriaki Kano, Nobuhiko Seraku, Fumio Takahashi, and Shin-ichi Tsuji. Attractive quality and must-be quality. *Journal of the Japanese Society for Quality Control*, 14(2):147–156, 1984. doi: 10.20684/quality.14.2_147. URL https://www.jstage.jst.go.jp/article/quality/14/2/14_KJ00002952366/_article.
- Katten Muchin Rosenman LLP. Katten wins dismissal of antitrust suit on behalf of Speedo. Firm announcement, May 14, 2010 (TYR Sport, Inc. v. Warnaco Swimwear, Inc., C.D. Cal.), 2010. URL <https://katten.com/Katten-Wins-Dismissal-of-Antitrust-Suit-on-Behalf-of-Speedo-05-14-2010>. Accessed 2026-10-19.
- Michael L. Katz and Carl Shapiro. Network externalities, competition, and compatibility. *American Economic Review*, 75(3):424–440, 1985.
- Michael L. Katz and Carl Shapiro. Technology adoption in the presence of network externalities. *Journal of Political Economy*, 94(4):822–841, 1986. doi: 10.1086/261409.
- Keane, Inc. Form 10-k405 for the fiscal year ended december 31, 2000. U.S. Securities and Exchange Commission, EDGAR, filed March 30, 2001, 2001. URL <https://www.sec.gov/Archives/edgar/data/54883/000092701601500101/d10k405.txt>. Accessed 2026-10-19.

- Keyfactor. Keyfactor research reveals digital certificate outages a weekly reality for 1 in 10 enterprises. Press release, September 30, 2025, 2025. URL <https://www.keyfactor.com/press-releases/keyfactor-research-reveals-digital-certificate-outages-a-weekly-reality-for-1-in-10-enterprises/>. Accessed 2026-10-19.
- Paul Klemperer. Markets with consumer switching costs. *Quarterly Journal of Economics*, 102(2): 375–394, 1987. doi: 10.2307/1885068.
- Bruno Koch. Business Case E-Invoicing / E-Billing 2017, 2017. URL <https://www.billentis.com/assets/reports/e-invoicing-businesscase.pdf>. Billentis, 10 February 2017. Accessed 2026-10-19.
- Lily Mae Lazarus. Exclusive: Vanta crosses \$300 million in ARR as shadow AI fuels compliance demand. *Fortune*, Term Sheet, April 2026. URL <https://www.fortune.com/2026/04/29/exclusive-vanta-arr-300-million-sequoia-shadow-ai-claude-cursor/>. Published April 29, 2026. Accessed 2026-10-19.
- Edward P. Lazear and Sherwin Rosen. Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, 89(5):841–864, 1981. doi: 10.1086/261010.
- Major League Baseball. Statcast. MLB.com Glossary, 2026. URL <https://mlb.com/glossary/statcast>. Version fetched 2026-10-20. Accessed 2026-10-20.
- Charles F. Manski. Identification of endogenous social effects: The reflection problem. *Review of Economic Studies*, 60(3):531–542, 1993.
- Maplebear Inc. Instacart Q4 2025 Shareholder Letter (Form 8-K exhibit), February 2026. URL <https://www.sec.gov/Archives/edgar/data/1579091/000157909126000006/cartfourthquarter2025sha.htm>. Earnings call February 12, 2026. Accessed 2026-10-20.
- MARA Holdings, Inc. Fourth Quarter and Fiscal Year 2025 Shareholder Letter, 2026. URL <https://www.sec.gov/Archives/edgar/data/1507605/000150760526000004/q425shareholderletter.htm>. Accessed 2026-10-20; primary shareholder disclosure.
- Marketing Dive. Amazon annual ad revenue passes \$68B, boosted by full-funnel strategy, February 2026. URL <https://www.marketingdive.com/news/amazon-annual-ad-revenue-passes-68b-boosted-by-full-funnel-strategy/811569/>. February 6, 2026. Accessed 2026-10-20.
- Marsh. Cyber Insurance Market Overview: Fourth Quarter 2021. Marsh McLennan, December 2021, 2021. URL <https://www.marsh.com/en/services/cyber-risk/insights/cyber-insurance-market-overview-q4-2021.html>. Reports US cyber pricing up 96% year over year in the third quarter of 2021 and insurers “stipulating that organizations adopt security controls.” Accessed 2026-10-19.
- Mastercard. Using AI and collaboration to prevent consumer payment scams, 2023. URL https://www.mastercard.com/content/dam/mccom/shared/news-and-trends/insights/2025/using-ai-and-collaboration-to-prevent-consumer-payment-scams/pdf/scam_risk_white_paper_october-2023.pdf. Accessed October 20, 2026.
- Craig Matsumoto. Carrier pwns NYC/Chicago latency. *Light Reading*, June 22, 2010, 2010. URL <https://www.lightreading.com/optical-networking/carrier-pwns-nyc-chicago-latency>. Accessed 2026-10-19.

- Harry McCracken. How the bloomberg terminal made history: And stays ever relevant. Fast Company, October 2015. URL <https://www.fastcompany.com/3051883/the-bloomberg-terminal>. Published 2015-10-06. Accessed 2026-10-20.
- McKay Brothers. McKay Brothers announces new milestones in low latency networks. Press release, February 5, 2013, 2013. URL <https://www.prnewswire.com/news-releases/mckay-brothers-announces-new-milestones-in-low-latency-networks-189880081.html>. Accessed 2026-10-19.
- McKinsey and Artemis Analytics. Stablecoins in payments: what the raw transaction numbers miss, 2026. URL <https://www.mckinsey.com/industries/financial-services/our-insights/stablecoins-in-payments-what-the-raw-transaction-numbers-miss>. Accessed October 20, 2026.
- Meta. Llama 3.1 Community License Agreement, 2024. URL https://github.com/meta-llama/llama-models/blob/main/models/llama3_1/LICENSE. Accessed 2026-10-20; primary license text.
- Meta Platforms, Inc. Meta Platforms Q4 2021 Earnings Call Transcript. platformaeronaut.com transcript archive, February 2022. URL <https://transcripts.platformaeronaut.com/transcripts/META-4Q21-transcript>. Call of February 2, 2022. Accessed 2026-10-20.
- Meta Platforms, Inc. Meta reports first quarter 2024 results, april 24, 2024. Meta Investor Relations, 2024. URL <https://investor.atmeta.com/investor-news/press-release-details/2024/Meta-Reports-First-Quarter-2024-Results/default.aspx>. Accessed 2026-10-20.
- Meta Platforms, Inc. Meta reports fourth quarter and full year 2025 results, january 28, 2026. Meta Investor Relations, 2026. URL <https://investor.atmeta.com/investor-news/press-release-details/2026/Meta-Reports-Fourth-Quarter-and-Full-Year-2025-Results/default.aspx>. Accessed 2026-10-20.
- Stephen Mewborn, Jackelyn Livesey, Ruth Zöhrer, and Naomi Alford. No More Easy Money on the Side: Retail Media Enters the Performance Era. Bain & Company, December 2023. URL <https://www.bain.com/insights/no-more-easy-money-on-the-side-retail-media-enters-the-performance-era/>. December 13, 2023. Accessed 2026-10-20 via the /fr/ mirror.
- Microsoft. Microsoft and openai extend partnership, january 23, 2023. Official Microsoft Blog, 2023. URL <https://blogs.microsoft.com/blog/2023/01/23/microsoftandopenaiextendpartnership/>. Accessed 2026-10-20.
- Microsoft Corporation. Fiscal Year 2026 Fourth Quarter Earnings Conference Call, 2026a. URL <https://www.microsoft.com/en-us/investor/events/fy-2026/earnings-fy-2026-q4>. July 29; official transcript; accessed 2026-10-20.
- Microsoft Corporation. Microsoft Cloud and AI Strength Fuels Fourth Quarter Results, 2026b. URL <https://www.microsoft.com/en-us/investor/earnings/fy-2026-q4/press-release-webcast>. July 29; year ended June 30, 2026; accessed 2026-10-20.
- Microsoft Learn. Governing agent identities, 2026. URL <https://learn.microsoft.com/en-us/entra/id-governance/agent-id-governance-overview>. Accessed October 20, 2026.
- Paul Milgrom and John Roberts. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58(6):1255–1277, 1990. doi: 10.2307/2938316.

Ministère de l'Économie. La facturation électronique démarre aujourd'hui: priorité à l'accompagnement des entreprises, 2026. URL <https://presse.economie.gouv.fr/la-facturation-demarre-aujourd'hui-priorite-a-laccompagnement-des-entreprises/>. September 1, press release 980. Includes platform enrollment and the 2026 sanctions grace period; accessed October 16, 2026.

Ministère de l'Économie. Tout savoir sur la facturation électronique pour les entreprises, n.d. URL <https://www.economie.gouv.fr/tout-savoir-sur-la-facturation-electronique-pour-les-entreprises>. Official rollout schedule. Accessed October 16, 2026.

Model Context Protocol Contributors. Model Context Protocol Specification, 2026. URL <https://modelcontextprotocol.io/specification/2026-07-28>. Source date: 2026-07-28 specification. Accessed October 20, 2026.

Modern Retail. With every company entering the ad business, here's what to expect from retail media networks in 2025, December 2024. URL <https://www.modernretail.co/marketing/with-every-company-entering-the-ad-business-heres-what-to-expect-from-retail-media-networks-in-2025/>. December 30, 2024. Accessed 2026-10-20.

Joel Monegro. Fat Protocols, 2016. URL <https://www.usv.com/writing/fat-protocols>. Accessed 2026-10-20; primary thesis essay.

Stephen Morris and Hyun Song Shin. Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, 88(3):587–597, 1998.

Munich Re. Insure AI, 2026. URL <https://www.munichre.com/en/solutions/for-industry-clients/insure-ai.html>. Accessed October 20, 2026.

Satoshi Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System, 2008. URL <https://bitcoin.org/bitcoin.pdf>. Accessed 2026-10-20; primary protocol paper.

National Association of Insurance Commissioners. Insurance topics: Artificial intelligence, 2026. URL <https://content.naic.org/insurance-topics/artificial-intelligence>. Updated April 3. Summarizes surveys with differing observation periods and definitions; accessed October 16, 2026.

National Cyber Security Centre. Timelines for migration to post-quantum cryptography, 2025. URL <https://www.ncsc.gov.uk/guidance/pqc-migration-timelines>. UK guidance. Accessed October 16, 2026.

NeuGroup. Bloomberg terminals: How much more you'll pay next year. NeuGroup Insights, 2024. URL <https://connect.neugroup.com/public/blogs/bloomberg-terminals-how-much-more-youll-pay-next-year>. Published 2024-10-09. Accessed 2026-10-20.

NIST. Comments on software and agentic AI identity concept paper, 2026. URL <https://www.nist.gov/news-events/news/2026/09/comments-software-and-agentic-ai-identity-concept-paper>. Accessed October 20, 2026.

Norris et al., Duke University. Rethinking load growth, 2025. URL <https://nicholasinstitute.duke.edu/sites/default/files/publications/rethinking-load-growth.pdf>. Accessed October 20, 2026.

NVIDIA. NVIDIA Announces Financial Results for Second Quarter Fiscal 2027, 2026. URL <https://investor.nvidia.com/news/press-release-details/2026/NVIDIA-Announces-Financial-Results-for-Second-Quarter-Fiscal-2027/default.aspx>. Source date: 2026-08-26. Accessed October 20, 2026.

NVIDIA Corporation. Nvidia blackwell platform arrives to power a new era of computing, march 18, 2024. NVIDIA Newsroom, 2024a. URL <https://nvidianews.nvidia.com/news/nvidia-blackwell-platform-arrives-to-power-a-new-era-of-computing>. Accessed 2026-10-20.

NVIDIA Corporation. Nvidia ceo jensen huang keynote at computex 2024, june 2, 2024. NVIDIA Blog, 2024b. URL <https://blogs.nvidia.com/blog/computex-2024-jensen-huang/>. Accessed 2026-10-20.

NVIDIA Corporation. Nvidia announces financial results for first quarter fiscal 2025 (form 8-k exhibit 99.1), may 22, 2024. SEC EDGAR, 2024c. URL <https://www.sec.gov/Archives/edgar/data/1045810/000104581024000113/q1fy25pr.htm>. Accessed 2026-10-20.

NVIDIA Corporation. Nvidia announces financial results for fourth quarter and fiscal 2025 (form 8-k exhibit 99.1), february 26, 2025. SEC EDGAR, 2025. URL <https://www.sec.gov/Archives/edgar/data/1045810/000104581025000021/q4fy25pr.htm>. Accessed 2026-10-20.

NVIDIA Corporation. NVIDIA announces financial results for fourth quarter and fiscal 2026 (form 8-k exhibit), february 25, 2026. U.S. Securities and Exchange Commission, EDGAR, 2026. URL <https://www.sec.gov/Archives/edgar/data/1045810/000104581026000019/q4fy26pr.htm>. Accessed 2026-10-19.

Devon O'Brien, Ryan Sleevi, and Andrew Whalley. Chrome's plan to distrust Symantec certificates. Google Security Blog, September 11, 2017, 2017. URL <https://security.googleblog.com/2017/09/chromes-plan-to-distrust-symantec.html>. Accessed 2026-10-19.

Ofcom. Report on the use of age assurance. Report laid before Parliament, July 16, 2026, 2026. URL <https://www.ofcom.org.uk/siteassets/resources/documents/online-safety/information-for-industry/age-assurance-report-2026/report-on-the-use-of-age-assurance.pdf>. Accessed 2026-10-19.

OpenAI. Introducing chatgpt, november 30, 2022. OpenAI, 2022. URL <https://openai.com/index/chatgpt/>. Accessed 2026-10-20.

OpenAI. Gpt-4o mini: advancing cost-efficient intelligence, july 18, 2024. OpenAI, 2024a. URL <https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/>. Accessed 2026-10-20.

OpenAI. Introducing swe-bench verified, august 13, 2024. OpenAI, 2024b. URL <https://openai.com/index/introducing-swe-bench-verified/>. Accessed 2026-10-20.

OpenAI. OpenAI and NVIDIA announce strategic partnership to deploy 10 gigawatts of NVIDIA systems. OpenAI, September 22, 2025, 2025. URL <https://openai.com/index/openai-nvidia-systems-partnership/>. Accessed 2026-10-20.

OpenAI. Enterprise Privacy at OpenAI, 2026. URL <https://openai.com/enterprise-privacy/>. Accessed 2026-10-20; primary company policy.

Osservatori Politecnico di Milano. Fatturazione Elettronica: cos'è e quali sono obblighi per PA e privati, 2026. URL <https://www.osservatori.net/blog/digital-b2b/fatturazione-elettronica-obblighi-pa-privati/>. Updated 3 April 2026; 2019 SdI data. Accessed 2026-10-19.

- Pactum. Customer stories and results, n.d. URL <https://pactum.com/clients>. Company-reported enterprise customers and use cases. Accessed October 20, 2026.
- Palo Alto Networks, Inc. Palo Alto Networks to acquire CyberArk (exhibit 99.1 to form 8-k), July 30, 2025. U.S. Securities and Exchange Commission, EDGAR; merger completed February 11, 2026 (CyberArk Form 6-K), 2025. URL <https://www.sec.gov/Archives/edgar/data/1327567/000119312525168821/d67702dex991.htm>. Accessed 2026-10-19.
- Stavan Parikh and Rao Surapaneni. Powering AI commerce with the new agent payments protocol (AP2), 2025. URL <https://cloud.google.com/blog/products/ai-machine-learning/announcing-agents-to-payments-ap2-protocol>. Google Cloud Blog, September 16.
- Path to Purchase Institute. Trends 2024, Part 1: Retailer Media Network Ratings & Insights, February 2024. URL <https://p2pi.com/trends-2024-part-1-retailer-media-network-ratings-insights>. February 5, 2024; survey fielded Sept 18–Oct 10, 2023, n=76. Accessed 2026-10-20.
- Pay.UK. Fraud detection pilot detects over £112m worth of fraud, 2024. URL <https://newseven.tsinsights.wearepay.uk/media-centre/press-releases/payuk-s-fraud-detection-pilot-exceeds-expectations-detecting-over-112m-worth-of-fraud/>. Accessed October 20, 2026.
- Persona. Series D. Persona blog, April 30, 2025, 2025. URL <https://www.withpersona.com/blog/series-d>. Accessed 2026-10-19.
- Polish Ministry of Finance. Zakres obowiazkowego KSeF, 2026. URL <https://ksef.podatki.gov.pl/informacje-ogolne-ksef-20/zakres-obowiazkowego-ksef/>. accessed version 2026-10-20. Accessed 2026-10-20.
- PPC Land. Italy fines Apple €98.6 million for making developers ask users twice, December 2025. URL <https://ppc.land/italy-fines-apple-eu98-6-million-for-making-developers-ask-users-twice/>. December 23, 2025; AGCM decision of December 22, 2025. Accessed 2026-10-20.
- PPC Land. Stack overflow drops to 1,442 questions in July, down 99% from 2014 peak, August 8, 2026. PPC Land, 2026. URL <https://ppc.land/stack-overflow-drops-to-1-442-questions-in-july-down-99-from-2014-peak/>. Accessed 2026-10-20.
- Ryan Prete and Tripp Baltz. High court tax ruling was a jackpot for software companies. Bloomberg Tax, March 15, 2019, 2019. URL <https://news.bloombergtax.com/daily-tax-report-state/high-court-tax-ruling-was-a-jackpot-for-software-companies>. Accessed 2026-10-19.
- Private Equity Wire. Francisco Partners, KKR take minority stakes in TeamSystem at up to €10bn valuation. Private Equity Wire, relaying a Reuters report, September 2026. URL <https://www.privateequitywire.co.uk/francisco-partners-kkr-take-minority-stakes-in-teamssystem-at-up-to-e10bn-valuation/>. Published September 14, 2026. Accessed 2026-10-19.
- Progressive. Snapshot: Usage-based auto insurance, n.d. URL <https://www.progressive.com/auto/discounts/snapshot/>. Current product description and availability conditions. Accessed October 16, 2026.
- PwC. Experience is everything: Here's how to get it right (consumer intelligence series; future of customer experience survey 2017/18), 2018. PwC, 2018. URL <https://www.pwc.com/us/en/ad>

- visory-services/publications/consumer-intelligence-series/pwc-consumer-intelligence-series-customer-experience.pdf. Accessed 2026-10-20.
- Radial, Inc. Radial consumer delivery expectations survey (conducted with dynata, august 2025), press release, september 8, 2025. Business Wire, 2025. URL <https://www.businesswire.com/news/home/20250908033040/en>. Accessed 2026-10-20.
- Michel Rauchs. Geographic Shift, 2021. URL <https://www.jbs.cam.ac.uk/2021/geographic-shift/>. October 13; accessed 2026-10-20; primary research account.
- Reality Defender. Press. realitydefender.com, 2026. URL <https://www.realitydefender.com/press>. Accessed 2026-10-20; headlines dated July-October 2026.
- RealPage, Inc. Thoma Bravo Completes Acquisition of RealPage (Exhibit 99.1 to Form 8-K), April 2021. URL <https://www.sec.gov/Archives/edgar/data/1286225/000119312521125818/d177605dex991.htm>. April 22, 2021. Accessed 2026-10-20.
- Jennifer F. Reinganum. On the diffusion of new technology: A game theoretic approach. *Review of Economic Studies*, 48(3):395–405, 1981.
- Don Reisinger. Uber Tests New Instant Pay Service for Drivers. *Fortune*, March 2016. URL <https://fortune.com/2016/03/17/uber-instant-pay>. March 17, 2016. Accessed 2026-10-20.
- Riskified Ltd. Prospectus (form 424b4), july 28, 2021. U.S. Securities and Exchange Commission, EDGAR, filed July 30, 2021, 2021. URL <https://www.sec.gov/Archives/edgar/data/1851112/000162828021015042/riskified424b4.htm>. Accessed 2026-10-19.
- Riskified Ltd. Form 20-f for the fiscal year ended december 31, 2025. U.S. Securities and Exchange Commission, EDGAR, filed March 6, 2026, 2026. URL <https://www.sec.gov/Archives/edgar/data/1851112/000185111226000007/rskd-20251231.htm>. Accessed 2026-10-19.
- Mark Roberti. Case Study: Wal-Mart’s Race for RFID. *eWeek*, September 2003. URL <https://www.eweek.com/enterprise-apps/case-study-wal-marts-race-for-rfid/>. September 15, 2003. Accessed 2026-10-20.
- Jean-Charles Rochet and Jean Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029, 2003. doi: 10.1162/154247603322493212.
- Jean-Paul Rodrigue. The Ideal-X, first containership, 1956. The Geography of Transport Systems, last modified April 24, 2026, 2026. URL <https://transportgeography.org/contents/chapter1/the-setting-of-global-transportation-systems/idealx-first-containeriship-1956/>. Sale to R.J. Reynolds in 1969 confirmed by FreightWaves, May 28, 2001. Accessed 2026-10-19.
- Bryan Rognier. Copilot is now included in Microsoft 365 personal and family. Microsoft 365 blog, January 16, 2025, 2025. URL <https://www.microsoft.com/en-us/microsoft-365/blog/2025/01/16/copilot-is-now-included-in-microsoft-365-personal-and-family/>. Accessed 2026-10-19.
- Lisa S. Rotenstein, A. Jay Holmgren, Robert Thombley, et al. Changes in clinician time expenditure and visit quantity with adoption of artificial intelligence–powered scribes: A multisite study. *JAMA*, 335(16):1408–1417, 2026. doi: 10.1001/jama.2026.2253. URL <https://jamanetwork.com/journals/jama/fullarticle/2847319>.

- R.R. Donnelley & Sons. A Guide to Walmart's Latest RFID Mandate, 2024. URL https://rrd.com/uploads/thoughts/pdfs/a-guide-to-walmarts-latest-rfid-mandate_20240404.pdf. Requirements as of February 2024; file dated April 4, 2024. Accessed 2026-10-20.
- John Rust. Optimal replacement of GMC bus engines: An empirical model of harold zurcher. *Econometrica*, 55(5):999–1033, 1987.
- Salesforce. New research: AI service agents are scaling and delivering CSAT, 2026. URL <https://www.salesforce.com/news/stories/ai-service-agents-improve-customer-satisfaction/>. May 20. Company survey conducted March 9–April 4, 2026; accessed October 16, 2026.
- Salesforce, Inc. New salesforce research: Customer experience is as important as products (state of the connected customer, sixth edition), august 28, 2023. Salesforce Newsroom, 2023. URL <https://www.salesforce.com/news/stories/customer-engagement-research-2023/>. Accessed 2026-10-20.
- Emily Schechter. A secure web is here to stay, 2018. URL <https://blog.chromium.org/2018/02/a-secure-web-is-here-to-stay.html>. Chromium Blog, February 8.
- Thomas C. Schelling. Hockey helmets, concealed weapons, and daylight saving: A study of binary choices with externalities. *Journal of Conflict Resolution*, 17(3):381–428, 1973.
- Thomas C. Schelling. *Micromotives and Macrobehavior*. W. W. Norton, New York, 1978.
- Semrush. Semrush, an Adobe Company, Named in Gartner Market Guide for Answer Engine Visibility Tools, 2026. URL <https://www.semrush.com/news/463886-semrush-an-adobe-company-named-in-gartner-market-guide-for-answer-engine-visibility-tools>. 2026-07-07. Accessed 2026-10-19.
- SEMrush Holdings, Inc. Form s-1 registration statement. U.S. Securities and Exchange Commission, EDGAR, 2021. URL <https://www.sec.gov/Archives/edgar/data/1831840/000162828021003658/semrushforms-1.htm>. Filed March 2021. Accessed 2026-10-20.
- Semrush Holdings, Inc. Semrush announces fourth quarter and full year 2025 financial results (form 8-k, exhibit 99.1). U.S. Securities and Exchange Commission, EDGAR, 2026a. URL <https://www.sec.gov/Archives/edgar/data/1831840/000162828026013251/semrush8-kexhibit991q42025.htm>. Dated 2026-03-02. Accessed 2026-10-20.
- Semrush Holdings, Inc. Semrush Holdings Form 10-K for fiscal year 2025, 2026b. URL <https://www.sec.gov/Archives/edgar/data/1831840/000162828026013259/semr-20251231.htm>. SEC EDGAR, filed 2026-03-02. Accessed 2026-10-19.
- Carl Shapiro and Hal R. Varian. *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business School Press, Boston, MA, 1999.
- Cuy Sheffield. Making Sense of Stablecoins, 2025. URL <https://corporate.visa.com/en/sites/visa-perspectives/trends-insights/making-sense-of-stablecoins.html>. Accessed 2026-10-20; primary measurement methodology.
- ShipBob. Shipbob raises \$200 million in series e funding led by bain capital ventures, valuing the ecommerce logistics platform well over \$1 billion, press release, june 29, 2021. Multichannel Merchant (press release copy), 2021. URL <https://multichannelmerchant.com/press-releases/shipbob-raises-200-million-in-series-e-funding-led-by-bain-capital-ventures-valuing-the-ecommerce-logistics-platform-well-over-1-billion>. Accessed 2026-10-20.

Shopify Inc. Holy ship! shopify to acquire deliverr for \$2.1b, building the future of global logistics for independent brands, may 5, 2022. Shopify News, 2022. URL <https://www.shopify.com/news/holy-ship-shopify-to-acquire-deliverr-for-2-1b-building-the-future-of-global-logistics-for-independent-brands>. Accessed 2026-10-20.

Shopify Inc. Shopify announces first-quarter 2023 financial results; agrees to sell shopify logistics to flexport, press release, may 4, 2023. Shopify News, 2023a. URL <https://www.shopify.com/news/shopify-announces-first-quarter-2023-financial-results-agrees-to-sell-shopify-logistics-to-flexport>. Accessed 2026-10-20.

Shopify Inc. Shopify announces second-quarter 2023 financial results, press release, august 2, 2023. Shopify News, 2023b. URL <https://www.shopify.com/news/shopify-announces-second-quarter-2023-financial-results>. Accessed 2026-10-20.

Sierra. Year two in review, 2026. URL <https://sierra.ai/blog/year-two-in-review>. Published 2026-02-06. Accessed 2026-10-19.

SiliconANGLE. Report: Design flaw discovery set to delay launch of nvidia’s new blackwell gpus, august 4, 2024. SiliconANGLE, 2024. URL <https://siliconangle.com/2024/08/04/report-design-flaw-discovery-set-delay-launch-nvidias-new-blackwell-gpus/>. Accessed 2026-10-20.

Shivalika Singh, Yiyang Nan, Alex Wang, Daniel D’Souza, et al. The leaderboard illusion, arxiv:2504.20879, submitted april 29, 2025. arXiv, 2025. URL <https://arxiv.org/abs/2504.20879>. Accessed 2026-10-20.

Solana Documentation. Fees, 2026. URL <https://solana.com/docs/core/fees>. Accessed 2026-10-20; primary protocol documentation.

Sports Illustrated. Suits me: Going for a swim in the record-breaking LZR Racer. Sports Illustrated, August 19, 2008, 2008. URL <https://www.si.com/olympics/2008/08/19/lzr-racersuit>. Accessed 2026-10-19.

Prakhar Srivastava. Tax software company Avalara confidentially files for US IPO, eyes return to public markets. Reuters, via Yahoo Finance, July 21, 2025, 2025. URL <https://finance.yahoo.com/news/tax-software-company-avalara-confidentially-203225384.html>. Accessed 2026-10-19.

Christopher Steiner. Wall street’s speed war. Forbes, September 9, 2010, 2010. URL <https://www.forbes.com/forbes/2010/0927/outfront-netscape-jim-barksdale-daniel-spivey-wall-street-speed-war.html>. The \$300 million construction cost is Forbes’s estimate; the company never disclosed it. Accessed 2026-10-19.

Stripe. Remote Helps Contractors Save Over 1 Percent of Payout Value by Offering Stablecoin Disbursements, 2026. URL <https://stripe.com/customers/remote>. Accessed 2026-10-20; primary company-published customer account.

Succinct. Prover Network FAQ, 2026. URL <https://explorer.succinct.xyz/faq>. Accessed October 20, 2026.

Sumsub. Sumsub research: Global deepfake incidents surge tenfold from 2022 to 2023, november 28, 2023. Sumsub Newsroom, 2023. URL <https://sumsub.com/newsroom/sumsub-research-global-deepfake-incidents-surge-tenfold-from-2022-to-2023/>. Accessed 2026-10-20.

Supply Chain Digest. RFID News: Looking Back at the Wal-Mart RFID Time Line, February 2009. URL https://www.scdigest.com/assets/On_Target/09-02-23-1.php. February 23, 2009. Accessed 2026-10-20.

John Sutton. *Sunk Costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration*. MIT Press, Cambridge, MA, 1991.

Claire Swedberg. Walmart Recommits to RFID. RFID Journal, January 2022. URL <https://www.rfidjournal.com/news/walmart-re-commits-to-rfid-with-supplier-mandates/193702/>. January 28, 2022. Accessed 2026-10-20.

Symantec Corporation. Form 8-k, item 1.01: Purchase agreement with DigiCert, august 2, 2017. U.S. Securities and Exchange Commission, EDGAR, filed August 4, 2017, 2017. URL <https://www.sec.gov/Archives/edgar/data/849399/000119312517248749/d356822d8k.htm>. Accessed 2026-10-19.

Symbotic / SEC. Symbotic Form 10-K FY2024, 2024. URL <https://www.sec.gov/Archives/edgar/data/1837240/000183724024000232/sym-20240928.htm>. Accessed October 20, 2026.

Symbotic / SEC Exhibit 99.1. Symbotic fourth quarter and fiscal year 2025 results, 2025. URL https://www.sec.gov/Archives/edgar/data/1837240/000183724025000277/q4258-k_ex991.htm. Accessed October 20, 2026.

Symbotic Inc. Symbotic to acquire Walmart’s Advanced Systems and Robotics business and develop Accelerated Pickup and Delivery centers (Form 8-K, Exhibit 99.1), 2025. URL <https://www.sec.gov/Archives/edgar/data/1837240/000119312525007370/d916752dex991.htm>. Press release dated 16 January 2025, filed with the SEC. Accessed 2026-10-19.

Szymanski et al. Author correction: An autonomous laboratory for accelerated synthesis, 2026. URL <https://www.nature.com/articles/s41586-025-09992-y>. Accessed October 20, 2026.

TaiyangNews. IRENA: Global Solar PV Additions Exceeded 510 GW In 2025, April 2026. URL <https://taiyangnews.info/business/irena-global-solar-pv-additions-exceeded-510-gw-in-2025>. April 2, 2026; reports IRENA Renewable Capacity Statistics 2026. Accessed 2026-10-20.

Target Corporation. Target to acquire same-day delivery platform shipt, inc. to bolster fulfillment capabilities, press release, december 13, 2017. Target Corporate, 2017. URL <https://corporate.target.com/press/releases/2017/12/target-to-acquire-same-day-delivery-platform-shipt>. Accessed 2026-10-20.

Tancred Taylor. 52.9 Billion UHF RFID Endpoint ICs Shipped in 2024. ABI Research Insight IN-7733, March 2025. URL <https://abiresearch.com/market-research/insight/7785716-529-billion-uhf-rfid-endpoint-ics-shipped-?hsLang=en>. March 10, 2025. Accessed 2026-10-20.

TeamSystem Holdco S.p.A. Consolidated financial statements for the year ended december 31, 2025. Published April 16, 2026, 2026. URL https://www.teamsystem.com/media/files/1689_2026%2004%2016%20-%20Documento%20bilancio%20Consolidato%202025_ENG_fascicolo%20completo.pdf. Accessed 2026-10-19.

TechCrunch. Microsoft says 8.5m windows devices were affected by crowdstrike outage. TechCrunch, 2024. URL <https://techcrunch.com/2024/07/20/microsoft-says-8.5m-windows-devices-were-affected-by-crowdstrike-outage>. Published 2024-07-20. Accessed 2026-10-20.

- TechCrunch. Amazon fulfillment competitor stord raises \$250m at \$3b valuation, may 26, 2026. TechCrunch, 2026. URL <https://techcrunch.com/2026/05/26/amazon-fulfillment-competitor-stord-raises-250m-at-3b-valuation/>. Accessed 2026-10-20.
- David J. Teece. Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research Policy*, 15(6):285–305, 1986. doi: 10.1016/0048-7333(86)90027-2.
- The Clearing House. RTP network marks may day with record-breaking volume and value, 2026. URL <https://www.theclearinghouse.org/payment-systems/Articles/2026/05/RTP-Network-Marks-May-Day-with-Record-Breaking-Volume-and-Value>. May 14. Reports payments on May 1, 2026; accessed October 16, 2026.
- The Motley Fool. Amazon.com inc (amzn) q1 2019 earnings call transcript, april 25, 2019. fool.com, 2019a. URL <https://www.fool.com/earnings/call-transcripts/2019/04/25/amazoncom-inc-amzn-q1-2019-earnings-call-transcrip.aspx>. Accessed 2026-10-20.
- The Motley Fool. Avalara Inc (AVLR) Q4 2018 earnings conference call transcript. Transcript of the February 12, 2019 call, 2019b. URL <https://www.fool.com/earnings/call-transcripts/2019/02/12/avalara-inc-avlr-q4-2018-earnings-conference-call.aspx>. Accessed 2026-10-19.
- The Motley Fool. Why chegg and pearson just popped but duolingo dropped, may 3, 2023. The Globe and Mail (syndicated), 2023. URL <https://www.theglobeandmail.com/investing/markets/stocks/PS0/pressreleases/16550775>. Accessed 2026-10-20.
- The Register. Australian social media ban enforcement report. The Register, March 31, 2026, reporting the eSafety Commissioner’s first compliance report, 2026. URL https://www.theregister.com/2026/03/31/australian_social_media_ban_enforcement_report/. Account removals in the first days: Region Canberra, January 16, 2026, <https://region.com.au/millions-of-under-16s-social-media-accounts-closed-since-ban-came-into-effect/935204/>. Accessed 2026-10-19.
- The Trade Desk, Inc. The trade desk reports second quarter 2026 financial results. The Trade Desk Investor Relations, 2026a. URL <https://investors.thetradedesk.com/news-and-events/news/news-details/2026/The-Trade-Desk-Reports-Second-Quarter-2026-Financial-Results/default.aspx>. Dated 2026-08-06. Accessed 2026-10-20.
- The Trade Desk, Inc. The trade desk reports fourth quarter and fiscal year 2025 financial results. Business Wire (via StockTitan); company IR page also listed, 2026b. URL <https://investors.thetradedesk.com/news-and-events/news/news-details/2026/The-Trade-Desk-Reports-Fourth-Quarter-and-Fiscal-Year-2025-Financial-Results/default.aspx>. Dated 2026-02-25. Text verified at <https://www.stocktitan.net/news/TTD/the-trade-desk-reports-fourth-quarter-and-fiscal-year-2025-financial-dzm3yj8ukxpf.html>. Accessed 2026-10-20.
- The Trade Desk, Inc. Form 10-k for the fiscal year ended december 31, 2025. U.S. Securities and Exchange Commission, EDGAR, 2026c. URL <https://www.sec.gov/Archives/edgar/data/1671933/000167193326000014/ttd-20251231.htm>. Accessed 2026-10-20.
- TheStreet, via AOL. Walmart Just Announced Major Change to Delivery Service, April 2025. URL <https://www.aol.com/walmart-just-announced-major-change-192513412.html>. April 29, 2025; quotes Doug McMillon’s April 2025 shareholder letter. Accessed 2026-10-20.

- TIKR. Coreweave (nasdaq: Crwv) stock sinks almost 8% as weak guidance spooks wall street, november 11, 2025. TIKR Blog, 2025. URL <https://www.tikr.com/blog/coreweave-nasdaq-crwv-stock-sinks-almost-8-as-weak-guidance-spooks-wall-street>. Accessed 2026-10-20.
- Gordon Tullock. Efficient rent seeking. In James M. Buchanan, Robert D. Tollison, and Gordon Tullock, editors, *Toward a Theory of the Rent-Seeking Society*, pages 97–112. Texas A&M University Press, College Station, 1980.
- Patrick Turley, Michelle N. Meyer, Nancy Wang, David Cesarini, Evelyn Hammonds, Alicia R. Martin, Benjamin M. Neale, Heidi L. Rehm, Louise Wilkins-Haug, Daniel J. Benjamin, et al. Problems with using polygenic scores to select embryos. *New England Journal of Medicine*, 385(1):78–86, 2021. doi: 10.1056/NEJMSr2105065. URL <https://www.nejm.org/doi/full/10.1056/NEJMSr2105065>.
- Monica Uprety. What’s new at Pactum: AI agents take on more procurement work, 2026. URL <https://pactum.com/blog/whats-new-pactum-september-2026>. Pactum, September 10. Product release report; accessed October 16, 2026.
- U.S. Census Bureau. Large firms with at least 20 employees biggest AI users. America Counts, May 26, 2026, 2026. URL <https://census.gov/library/stories/2026/05/ai-use-businesses.html>. Accessed 2026-10-19.
- U.S. Department of Energy, Solid-State Lighting Program. 2020 LED Adoption Report. Technical report, U.S. DOE, August 2020. URL <https://www.energy.gov/eere/ssl/articles/2020-led-adoption-report>. Data year 2018. Accessed 2026-10-20.
- U.S. Department of Justice. Visa and Plaid abandon merger after antitrust division’s suit to block. Press release 21-38, January 12, 2021, 2021. URL <https://www.justice.gov/opa/pr/visa-and-plaid-abandon-merger-after-antitrust-division-s-suit-block>. Accessed 2026-10-19.
- U.S. Department of Justice. Justice Department Sues RealPage for Algorithmic Pricing Scheme that Harms Millions of American Renters, August 2024. URL <https://www.justice.gov/usao-mdnc/pr/justice-department-sues-realpage-algorithmic-pricing-scheme-harms-millions-american>. August 23, 2024. Accessed 2026-10-20.
- U.S. Department of Justice. Justice Department Requires RealPage to End the Sharing of Competitively Sensitive Information and Alignment of Pricing Among Competitors, 2025. URL <https://www.justice.gov/opa/pr/justice-department-requires-realpage-end-sharing-competitively-sensitive-information-and>. 2025-11-24. Accessed 2026-10-20.
- U.S. Department of Justice, Antitrust Division. Complaint, United States et al. v. RealPage, Inc., No. 1:24-cv-00710 (M.D.N.C.), August 2024. URL <https://www.justice.gov/usdoj-media/atr/media/1365471/dl>. Filed August 23, 2024. Accessed 2026-10-20.
- U.S. Energy Information Administration. Battery storage capacity averaged 70% growth over the last three years, 2026. URL <https://www.eia.gov/todayinenergy/detail.php?id=67925>. August 7. Operating nameplate power capacity through June 2026; accessed October 16, 2026.
- U.S. Securities and Exchange Commission. SEC approves IEX proposal to launch national exchange, issues interpretation on automated securities prices. Press Release 2016-123, June 17, 2016, 2016. URL <https://www.sec.gov/news/pressrelease/2016-123.html>. Accessed 2026-10-19.

U.S. Securities and Exchange Commission. SEC Votes to End Defense of Climate Disclosure Rules (Press Release 2025-58), 2025. URL <https://www.sec.gov/newsroom/press-releases/2025-58>. March 27, 2025. Accessed 2026-10-19.

U.S. Securities and Exchange Commission. SEC Proposes Rescission of Climate-Related Disclosure Rules, 2026. URL <https://www.sec.gov/newsroom/press-releases/2026-49-sec-propose-s-rescission-climate-related-disclosure-rules>. 2026-05-29. Accessed 2026-10-20.

Vanta. Celebrating new milestones. Vanta blog, January 31, 2024, 2024. URL <https://www.vanta.com/resources/celebrating-new-milestones>. Accessed 2026-10-19.

Vendavo. Pricing and CPQ for industrial and MRO distribution, 2026. URL <https://www.vendavo.com/industries/distribution/industrial-and-mro/>. Accessed October 20, 2026.

VeriFone Systems, Inc. Fourth quarter and fiscal year 2015 results (exhibit 99.1 to form 8-k), december 14, 2015. U.S. Securities and Exchange Commission, EDGAR, 2015. URL <https://www.sec.gov/Archives/edgar/data/1312073/000131207315000052/ex991q42015pressrelease.htm>. Accessed 2026-10-19.

VeriFone Systems, Inc. Fourth quarter and fiscal year 2016 results (exhibit 99.1 to form 8-k), december 12, 2016. U.S. Securities and Exchange Commission, EDGAR, 2016. URL <https://www.sec.gov/Archives/edgar/data/1312073/000131207316000109/ex991q42016pressrelease.htm>. Accessed 2026-10-19.

VeriFone Systems, Inc. Fourth quarter and fiscal year 2017 results (exhibit 99.1 to form 8-k), december 12, 2017. U.S. Securities and Exchange Commission, EDGAR, 2017a. URL <https://www.sec.gov/Archives/edgar/data/1312073/000131207317000037/ex991q42017pressrelease.htm>. Accessed 2026-10-19.

VeriFone Systems, Inc. Form 10-k for the fiscal year ended October 31, 2017. U.S. Securities and Exchange Commission, EDGAR, 2017b. URL <https://www.sec.gov/Archives/edgar/data/0001312073/000131207317000047/pay10k10312017.htm>. Accessed 2026-10-19.

VeriFone Systems, Inc. Verifone to be acquired by Francisco Partners for \$3.4 billion (exhibit 99.1 to form 8-k), april 9, 2018. U.S. Securities and Exchange Commission, EDGAR, 2018. URL <https://www.sec.gov/Archives/edgar/data/1312073/000119312518111709/d546581dex991.htm>. Accessed 2026-10-19.

Vertex, Inc. Vertex Completes Acquisition of ecosio, 2024. URL <https://ir.vertexinc.com/news-releases/news-release-details/vertex-completes-acquisition-ecosio>. Press release, 4 September 2024. Accessed 2026-10-19.

Bob Violino. Wal-Mart Draws Line in the Sand. RFID Journal, June 2003. URL <https://www.rfidjournal.com/?p=84246>. June 12, 2003. Accessed 2026-10-20.

Visa. EMV Liability Shift: Why It Pays to Adopt New Technology, 2016. URL <https://usa.visa.com/content/dam/VCOM/global/partner-with-us/documents/visa-liability-shift.pdf>. 2016 guidance; effective date 2015-10-01. Accessed 2026-10-20.

Visa. Financing the next hundred stablecoin-linked card programs, 2026. URL <https://www.visa.com/en-us/thought-leadership/innovation/financing-stablecoin-linked-card-programs>. Updated September 8. Provider-reported settlement run rate; accessed October 16, 2026.

- Visa Inc. Visa to acquire Plaid. Press release, January 13, 2020, 2020. URL <https://investor.visa.com/news/news-details/2020/Visa-To-Acquire-Plaid/default.aspx>. Accessed 2026-10-19.
- Vista Equity Partners. Avalara to be acquired by vista equity partners for \$8.4 billion. Press release, August 8, 2022; closing announced October 19, 2022 (Avalara Form 8-K), 2022. URL <https://www.vistaequitypartners.com/news/avalara-to-be-acquired-by-vista-equity-partners-for-8-4-billion/>. Accessed 2026-10-19.
- Mike Wackett. Two decades of liner shipping consolidation has destroyed shareholder value, says report. gCaptain / The Loadstar, February 12, 2018, reporting McKinsey, “Container shipping: More mergers, better mergers”, 2018. URL <https://gcaptain.com/two-decades-of-liner-shipping-consolidation-has-destroyed-shareholder-value-says-report/>. Accessed 2026-10-19.
- Walmart Inc. Q3 fy25 earnings call transcript, november 19, 2024. Walmart Corporate, 2024. URL <https://corporate.walmart.com/content/dam/corporate/documents/newsroom/2024/11/19/walmart-releases-q3-fy25-earnings/q3-fy25-earnings-call-transcript.pdf>. Accessed 2026-10-20.
- Walmart Inc. Walmart reports revenue growth of 5.6%, up 4.9% in constant currency (Q4 FY26 earnings release, Form 8-K exhibit), February 2026. URL <https://stock.walmart.com/sec-filings/all-sec-filings/content/0000104169-26-000032/earningsreleasefy26q4.htm>. February 19, 2026. Accessed 2026-10-20.
- WatersTechnology. Bloomberg upsets market data managers with terminal price hike. WatersTechnology, 2000. URL <https://www.waterstechnology.com/node/1622926>. Published 2000-05-15. Accessed 2026-10-20.
- Wikipedia contributors. Hms agincourt (1913). Wikipedia, The Free Encyclopedia, 2026a. URL [https://en.wikipedia.org/wiki/HMS_Agincourt_\(1913\)](https://en.wikipedia.org/wiki/HMS_Agincourt_(1913)). Version fetched 2026-10-20. Accessed 2026-10-20.
- Wikipedia contributors. Prolexic technologies. Wikipedia, 2026b. URL https://en.wikipedia.org/wiki/Prolexic_Technologies. Accessed 2026-10-20; tertiary, citing SecurityWeek (December 2, 2013) and Akamai (February 18, 2014).
- Wikipedia contributors. Vickers limited. Wikipedia, The Free Encyclopedia, 2026c. URL https://en.wikipedia.org/wiki/Vickers_Limited. Version fetched 2026-10-20. Accessed 2026-10-20.
- Wiley Rein LLP. Key developments with state app store accountability acts as Texas act takes effect. Client alert, June 8, 2026; Supreme Court denial of emergency applications July 6, 2026 per CCIA v. Paxton litigation page, 2026. URL <https://www.wiley.law/alert-Key-Developments-With-State-App-Store-Accountability-Acts-as-Texas-Act-Takes-Effect>. CCIA litigation page: <https://ccianet.org/litigation/ccia-v-paxton/>. Accessed 2026-10-19.
- Wise. Wise FY26 results, 2026. URL <https://owners.wise.com/news-releases/news-release-details/wise-fy26-results>. Accessed October 20, 2026.
- T. P. Wright. Factors affecting the cost of airplanes. *Journal of the Aeronautical Sciences*, 3(4): 122–128, 1936. doi: 10.2514/8.155.

x402. Introducing x402 V2: Evolving the Standard for Internet-native Payments, 2025. URL <https://x402.org/x402-v2-launch/>. Source date: 2025-12-11; page updated 2026-06-24. Accessed October 20, 2026.

Ken Yeung. Uber paid \$1.3 billion to drivers using Instant Pay in the past 12 months. VentureBeat, April 2017. URL <https://venturebeat.com/2017/04/05/uber-paid-1-3-billion-to-drivers-using-instant-pay-in-the-past-12-months>. April 5, 2017. Accessed 2026-10-20.

Zayo Group Holdings. Zayo to Acquire Spread Networks. Company press release, November 27, 2017, 2017. URL <https://www.zayo.com/newsroom/zayo-to-acquire-spread-networks/>. Acquisition agreement, not closing; the announcement expected closing in 2018. Accessed 2026-10-20.

Zayo Group Holdings, Inc. Form 10-q for the quarter ended December 31, 2017. U.S. Securities and Exchange Commission, EDGAR, 2018. URL <https://www.sec.gov/Archives/edgar/data/0001608249/000155837018000478/zayo-20171231x10q.htm>. Records the acquisition of Spread Networks. Accessed 2026-10-19.

Mark Zuckerberg. Open Source AI Is the Path Forward, 2024. URL <https://about.fb.com/news/2024/07/open-source-ai-is-the-path-forward/>. Accessed 2026-10-20; primary strategy statement.

A The model in full

This appendix derives the two payoff curves of Section 1, adds the outside option, and works through the four patterns of Figure 1 with numbers.

A.1 Whose decision is modeled

Take a fixed group of rivals and let x be the share that has adopted: $x = 0$ means none and $x = 1$ means all. Each actor is small enough that its own choice does not materially change that share. In a firm-level application, x should reflect the competitors or partners that actually affect the firm's business. Appendix E.10 allows their influence to differ.

Let $U_A(x)$ be the payoff from adopting and $U_N(x)$ the payoff from holding out. Both are measured relative to the world without this technology. Zero means no change from that baseline, not zero revenue or profit. A negative number means the actor is worse off than before.

Use a common period for all amounts. The cost c includes annualized acquisition cost and the cost of operating the technology. For now, consider an actor with the average cost. An individual actor's calculation replaces c with its own c_i . The model later allows those costs to differ across the population.

The illustrative figures use a common payoff unit per period. If one unit represents a pre-technology annual stake of \$1 million, a benefit of 0.1 means \$100,000 per year and a cost of 0.3 means \$300,000 per year. These are amounts relative to that chosen unit, not adoption percentages. The adoption share x and strength ratio r have no monetary units. Multiplying all payoff amounts by the same positive factor changes the vertical scale, not which choice pays more. The numerical examples are chosen to expose mechanisms; they are not estimates of a particular technology.

A.2 Four sources of value or pressure

The model distinguishes four sources of value or pressure.

Direct benefit, γ .

The technology improves your work even if no rival adopts. Lower handling costs from better equipment belong here.

Network benefit, βx .

The technology becomes more useful when others have compatible equipment or practices. More container-ready ports make a container system more useful.

Competitive edge, α and r .

The technology helps you capture a share of something rivals would otherwise receive. The prize might be business, a ranking or a strategic position. An adopter's gain has a matching loss elsewhere in the rival group.

Exclusion, $\varepsilon\Phi(x)$.

A holdout loses access or faces a penalty that no rival receives as a gain. A platform warning or a procurement requirement can work this way. If the lost business goes to another rival, that transferred portion belongs in the contest instead.

These channels can operate together. A productive tool can also redistribute business; a useful standard can also be enforced by a gatekeeper. Keeping the channels separate prevents the same gain or loss from being counted twice.

A.3 The two curves and their gap

First suppose holdouts remain exposed to competition and exclusion. Let $p \in [0, 1]$ be the common belief that the technology works. The denominator

$$D(x) = 1 + (r - 1)x$$

tracks how much competitive strength is already in the field. Here r is an adopter's strength relative to a holdout's, and α is the positional edge available to a lone adopter. The two payoffs are

$$U_A^{\text{exposed}}(x) = p \left[\gamma + \beta x + \frac{\alpha(1-x)}{D(x)} \right] - c, \quad (3)$$

$$U_N^{\text{exposed}}(x) = -p \frac{\alpha x}{D(x)} - \varepsilon \Phi(x). \quad (4)$$

Read the adopter's curve from left to right: direct benefit, compatibility benefit, an edge over the remaining holdouts, then cost. Read the holdout's curve as two losses: ground taken by adopters and an exclusion penalty. The pressure curve Φ rises from zero to one; its location θ is an adoption share and k controls its steepness. Calendar enforcement requires an explicit time path rather than interpreting θ as a date.

Beliefs scale the technology's direct, network and competitive effects. They do not scale the cost or the gatekeeper's penalty. A firm can face a real compliance requirement while remaining doubtful about the underlying technology.

Subtracting the holdout's payoff from the adopter's payoff gives the *gain from adopting*:

$$\underbrace{\Delta(x)}_{\text{adopt minus hold out}} = \underbrace{p\gamma}_{\text{direct}} + \underbrace{p\beta x}_{\text{network}} + \underbrace{\frac{p\alpha}{D(x)}}_{\text{contest}} + \underbrace{\varepsilon\Phi(x)}_{\text{exclusion}} - c.$$

A positive gap means adopting pays better than holding out. It does not mean the adopter is better off than before the technology existed: both curves can lie below zero.

The practical decision also includes the value of waiting, W . A holdout acts when $\Delta(x) - W > 0$, subject to adjustment delays. Waiting may allow another trial to finish or avoid committing to a generation that will soon be replaced. It affects the decision date; it is not an additional resource cost in the welfare calculation.

A.4 A credible alternative to staying exposed

A holdout may be able to ally, specialize or leave the contested segment. Let $Q \leq 0$ be the payoff from that alternative. A value of $Q = -0.1$ means accepting a loss of 0.1 per period relative to the original baseline in exchange for shelter. A better alternative has a Q closer to zero.

The loss from remaining exposed is

$$L(x) = p \frac{\alpha x}{D(x)} + \varepsilon \Phi(x). \quad (5)$$

The holdout shelters once that loss reaches $-Q$. Call the first such adoption share x_Q ; set $x_Q = \infty$ if the loss never gets that large. The notation $Q = -\infty$ means no outside option.

This version assumes a common outside option and that holdouts shelter together. Once they leave, their stakes are no longer available for adopters to capture. The full payoff curves are therefore

$$U_A(x) = p \left[\gamma + \beta x + \frac{\alpha(1-x)}{D(x)} \mathbf{1}\{x < x_Q\} \right] - c, \quad (6)$$

$$U_N(x) = \max\{-L(x), Q\}. \quad (7)$$

The indicator switches the competitive edge off after sheltering. The full adoption gain becomes

$$\Delta(x) = \begin{cases} p[\gamma + \beta x + \alpha/D(x)] + \varepsilon\Phi(x) - c, & x < x_Q, \\ p[\gamma + \beta x] - Q - c, & x \geq x_Q. \end{cases} \quad (8)$$

At the sheltering boundary the gap drops because the adopter loses its edge over the departed rivals. A worked example in Appendix B shows how this can stop further adoption.

The common- Q assumption makes that response abrupt. It selects sheltering only when remaining exposed has become individually worse, given that others have not sheltered yet. Different outside options would spread the exits over time. A rival implementation that preserves compatibility is also a different alternative from leaving the segment; Appendix H.4 discusses that case.

A.5 Why a competitive advantage wears away

The contest divides a fixed prize worth V per rival in proportion to strength. Adopters have strength r , holdouts strength one, so the average strength is $D(x)$. Their shares relative to the original equal split are $r/D(x)$ and $1/D(x)$. Subtract the original prize V and the two competitive payoffs become

$$V \left(\frac{r}{D(x)} - 1 \right) = \frac{\alpha(1-x)}{D(x)}, \quad V \left(\frac{1}{D(x)} - 1 \right) = -\frac{\alpha x}{D(x)}, \quad \alpha = V(r-1).$$

One side's gain is the other side's loss. Across the population, the terms cancel:

$$x \frac{\alpha(1-x)}{D(x)} + (1-x) \left(-\frac{\alpha x}{D(x)} \right) = 0.$$

This is what a zero-sum contest means here. Adopting still costs resources; those costs remain after the transfer cancels.

With $V > 0$ and $r > 1$, a lone adopter has many holdouts to take ground from. As rivals catch up, the opportunity shrinks. The contest's contribution to the adoption gap, $\alpha/D(x)$, falls with x . The paper also uses a linear limit, written $r = 1$ with α fixed. In that limit the gap stays constant even though the adopter's payoff falls. This specification is a limit with a fixed positional edge; it does not describe equal-strength players competing for a finite prize.

A.6 Irreversibility and cost dispersion

The static comparison asks what each actor would choose at a given exposure. The diffusion model adds irreversible commitment: an actor that has adopted cannot undo the decision. The forcing test still asks whether dropping the technology and saving its cost would pay. If acquisition costs are already sunk, dropping saves less; that makes continued use easier to justify.

Cost differences matter too. The model uses a log-logistic distribution with mean c and dispersion σ . At $\sigma = 0$ everyone has the same cost; a larger σ creates both cheap early adopters and costly holdouts. Its exact formula and the full parameter guide are in Appendix L. The distinction between cheap early adopters and actors stranded by an eroding edge will matter for the diffusion results.

B Three tests, worked with numbers

Once the technology is known to work, set $p = 1$ and ask the three questions of Section 1. With no outside option, the answers for a mean-cost actor are

$$\begin{aligned}\Delta(0) &= \gamma + \alpha - c && \text{Would adopting first pay?} \\ \Delta(1) &= \gamma + \beta + \alpha/r + \varepsilon - c && \text{Would dropping out hurt?} \\ y_R &= W_R(1) - W_R(0) = \gamma + \beta - c && \text{Does full adoption benefit rivals?}\end{aligned}\tag{9}$$

The third line adds everyone's payoff, so the competitive transfer disappears. Exclusion also disappears at full adoption because no holdouts remain to bear it. The outsider-adjusted comparison is $y_S = y_R + \eta$, where η is the benefit per adopter to customers or other outsiders.

A technology is *forced* when $\Delta(1) > 0$: a mean-cost actor would lose by dropping it once everyone else uses it. It is *competitively compelled* when $\Delta(0) \leq 0$: that actor would not strictly gain by adopting alone. It is *wasteful for rivals* when $y_R < 0$. These are separate tests. Other cost types can get different answers, and a worthwhile outside option can change the first two.

A yes means a strictly positive private gain. The labels summarize endpoints, not every point between them: a curved gap may cross zero more than once. The dominant label is the endpoint label used by the implementation, not a guarantee that adoption pays at every intermediate share. At $y_R = 0$ the welfare comparison is neutral; strict waste requires $y_R < 0$.

B.1 An arms race: everyone follows, everyone loses

Consider the linear-limit example in Figure 1, with direct benefit $\gamma = 0.1$, competitive edge $\alpha = 0.9$ and cost $c = 0.3$. There is no network benefit, exclusion or outside option. These are illustrative payoff units per period, not estimated historical returns. The curves are

$$U_A(x) = 0.7 - 0.9x, \quad U_N(x) = -0.9x, \quad \Delta(x) = 0.7.$$

A first adopter gains 0.7. When half the field has adopted, an adopter earns 0.25 and a holdout loses 0.45. At full adoption, every adopter loses 0.2 relative to the original world; a lone actor dropping out would lose 0.9. At every stage, adopting beats holding out by 0.7.

Individually, keeping the technology is the right choice. Collectively, abandoning it would save each rival 0.2. The technology is forced and wasteful, even though nobody needed to wait for rivals before adopting. That is the arms-race pattern.

B.2 Critical mass: adoption becomes more useful as others join

Now remove the contest and use a direct benefit of $\gamma = 0.2$, a network benefit of $\beta = 0.8$ and cost $c = 0.6$. Holdouts lose nothing:

$$U_A(x) = -0.4 + 0.8x, \quad U_N(x) = 0, \quad \Delta(x) = -0.4 + 0.8x.$$

A lone adopter loses 0.4. At $x = 0.5$ the two choices pay equally. Above that share, adoption pays and each new adopter makes joining more useful. At full adoption everyone gains 0.4.

This is a coordination problem. A sponsor needs enough compatible use to make adoption self-sustaining. With identical actors and no shocks, a seed below half does not trigger the rest. Cost differences smooth this picture, but the distinction from an arms race remains: compatibility creates value that survives when everyone joins.

B.3 Partial adoption: the edge stops covering its cost

Return to a pure contest, now with a prize $V = 0.3$, relative strength $r = 4$, cost $c = 0.4$ and no direct benefit. The lone adopter's edge is $\alpha = V(r - 1) = 0.9$. Its adoption gap is

$$\Delta(x) = \frac{0.9}{1 + 3x} - 0.4.$$

Adopting first pays: $\Delta(0) = 0.5$. Keeping the technology at full adoption does not: $\Delta(1) = -0.175$. The gap reaches zero at

$$x^\circ = \frac{0.9/0.4 - 1}{3} = \frac{5}{12} \approx 42\%.$$

Below that point the edge still covers the cost; above it, too many rivals have caught up. At the static equilibrium, adopters and holdouts both earn $-1/6$, approximately -0.167 . The advantage has been competed away, leaving the field worse off.

For the same pure contest with a direct benefit γ , the general interior solution is

$$x^\circ = \frac{\alpha/(c - \gamma) - 1}{r - 1}, \quad \frac{\alpha}{r} < c - \gamma < \alpha, \quad r > 1.$$

This is *selective adoption*. It can force a response from the whole field without making universal copying attractive. It is a static reference point; irreversible decisions can carry actual adoption beyond it, as Appendix E.8 shows.

B.4 Pressure on holdouts need not create a cascade

In the pure contest, holdouts lose more as adoption spreads, yet the gap between adopting and holding out gets smaller. A late adopter avoids a loss but no longer gets the large first-mover advantage. The derivative makes the distinction explicit:

$$\Delta'(x) = p\beta - \frac{p\alpha(r - 1)}{D(x)^2} + \varepsilon\Phi'(x) \quad (x < x_Q).$$

The network and exclusion terms push the incentive upward; a positive-prize contest pushes it downward. Economists call the latter *strategic substitutes*: others' adoption reduces your incremental gain from adopting. With only that contest, an upward tipping threshold cannot arise. Network benefits or exclusion are needed to make adoption self-reinforcing. Figure 5 shows the resulting patterns.

B.5 An alternative can stop the race before the field catches up

Give the selective-adoption example an outside option $Q = -0.1$. Holdouts now accept a loss of 0.1 to shelter rather than remain exposed. They reach that point when

$$\frac{0.9x}{1 + 3x} = 0.1 \implies x_Q = \frac{1}{6} \approx 17\%.$$

Once they shelter, adopters have no remaining holdout stake to capture. An adopter earns -0.4 , a sheltered holdout earns -0.1 , and the gap drops to -0.3 . New adoption no longer pays.

Under irreversible adoption, this is a stopping boundary: a discrete round of commitments can carry adoption past it, leaving existing adopters with continuing costs after their competitive advantage has disappeared. It is not an indifference equilibrium at which those adopters would willingly remain if they could switch. The lesson for an investment decision is to evaluate rivals' alternatives before projecting an early competitive return.




Figure 5: How cost and competitive strength change the endpoint pattern. The curves mark whether adopting first pays (blue), whether remaining adopted at full use pays (orange), and whether full adoption benefits rivals (dashed). Held fixed: $V = 0.5$, $\gamma = 0.1$, $\beta = 0.3$, $p = 1$, $\varepsilon = \sigma = W = 0$, $Q = -\infty$. Both c and r vary along the axes, and $\alpha = 0.5(r - 1)$ changes with r : this holds the prize fixed, not the first adopter's edge. A decisive edge opens a region of selective adoption; tipping requires the network benefit to outweigh the erosion of the competitive gap.

At full adoption, a credible outside option caps the private cost of dropping out:

$$\Delta(1) = \gamma + \beta - c + \min\{\alpha/r + \varepsilon, -Q\}.$$

A costly exit reduces pressure but leaves a cost to avoiding the technology. Only a costless alternative, $Q = 0$, removes that extra term entirely. No such option can disarm forcing if direct and network gains already strictly exceed the cost.

B.6 A gatekeeper can make adoption necessary without adding value

Exclusion raises the payoff gap by lowering the holdout curve. While holdouts remain exposed, increasing ε raises the full-adoption forcing margin but leaves y_R unchanged. The penalty is paid during the transition, when some holdouts remain; it vanishes when everyone complies.

The reach of the gatekeeper matters. A supplier may escape a retailer's mandate by leaving that supply chain, at the cost of its business there. If the penalty follows the actor into every feasible alternative, leaving does not provide the modeled shelter. Assess that reach, the transition burden and any benefit to users separately. A faster transition reduces the time penalties are paid, but the model does not include every installation cost of accelerating it.

B.7 What each parameter changes

Figure 1 shows four possible technologies, with a different set of assumptions in each panel. It does not isolate the effect of any one assumption. Figure 6 does: every panel starts from the same selective-forcing example, then changes just the parameter named in its title. The plotted quantity is now the gap $\Delta = U_A - U_N$. Above zero, adopting pays more; below zero, holding out pays more.

Table 21 connects each change to the original two curves. The last column asks a separate question: does that change make *full adoption* more valuable for the rival group? All comparisons hold the other parameters fixed, assume known effectiveness ($p = 1$), and leave holdouts exposed ($Q = -\infty$).

Table 21: How changing one parameter moves the payoffs. Here $d > 0$ is an increase in payoff units per period, $D = 1 + (r - 1)x$, and $y_R = \gamma + \beta - c$. The r row assumes a positive competitive edge.

Change	What moves in the two-curve picture?	Full-adoption payoff y_R
Direct benefit $\gamma + d$	Adopter curve rises by d at every share. The gap widens equally everywhere.	Rises by d
Network benefit $\beta + d$	Adopter curve rises by dx : unchanged at zero adoption, larger gains as use spreads.	Rises by d
Competitive edge $\alpha + d$, fixed r	Adopter curve rises by $d(1 - x)/D$; holdout curve falls by dx/D . The gap widens by d/D .	Unchanged
Faster edge erosion: higher r , fixed α	Once $x > 0$, the gap narrows. The first adopter's gain is unchanged; later adopters recover less of it.	Unchanged
Exclusion penalty $\varepsilon + d$	Holdout curve falls by $d\Phi(x)$. Adoption becomes harder to avoid without improving the adopter's payoff.	Unchanged
Adoption cost $c + d$	Adopter curve falls by d at every share. The gap narrows equally everywhere.	Falls by d

The numerical comparison shows why adoption and value need separate answers. Raising either the direct benefit or the network benefit from zero to 0.5 makes full adoption privately sustainable and raises its collective payoff from -0.4 to $+0.1$. Doubling the competitive edge from 0.9 to 1.8, or adding an exclusion penalty of 0.3, also makes full adoption privately sustainable. In those cases its collective payoff stays at -0.4 : the stronger incentive comes from transfers or pressure on holdouts.

Each panel reports both the stable static adoption share and y_R . These are different comparisons. For the unchanged selective example, the stable share is 42% and welfare there is approximately -0.167 per rival. The displayed $y_R = -0.4$ is what rivals would receive *if everyone adopted*. A negative y_R does not say that the static outcome is full adoption, and a static share does not predict how far irreversible investment will go.

Changing r requires saying what stays fixed. In Figure 6, raising r from 4 to 8 keeps the lone adopter's edge at $\alpha = 0.9$. That edge wears away faster, so stable adoption falls from 42% to 18%. The implied prize $V = \alpha/(r - 1)$ falls from 0.3 to about 0.129. If instead the prize stays at $V = 0.3$, the same increase in r raises α from 0.9 to 2.1 and stable adoption rises to 61%. At fixed V , a stronger technology captures more of the same prize; at fixed α , the experiment changes how quickly the same initial advantage disappears. Figure 5 holds V fixed; the r panels of Figures 6 and 11 hold α fixed.

The remaining parameters answer different questions. The exclusion location θ and steepness k change where and how abruptly pressure arrives, with $\Phi(0) = 0$ and $\Phi(1) = 1$ unchanged.




Figure 6: One parameter changes; the others stay fixed. The baseline is selective forcing with $\gamma = \beta = \varepsilon = 0$, $\alpha = 0.9$, $r = 4$ and $c = 0.4$ ($V = 0.3$ initially). All panels use $p = 1$, $\sigma = W = 0$ and $Q = -\infty$; exclusion, when added, uses $\theta = 0.5$, $k = 8$. Gray dashed curves repeat the baseline; blue solid curves apply the change in the panel title. Dots mark stable static adoption shares. The numbers below each panel separate adoption from the payoff to the rival group at full adoption. In the r panel, α is held fixed, so V changes. These are illustrative scenarios, not fitted estimates.

They do nothing when $\varepsilon = 0$. A better outside option Q can cap holdouts' loss and remove the adopter's contest return when sheltering begins. Cost dispersion σ changes how many actors lie above or below the mean-cost calculation; it does not move that actor's payoff curves or the total cost of full adoption. The outsider benefit η changes social value, $y_S = y_R + \eta$, without changing either private curve.

Belief p scales direct, network and competitive effects, while leaving adoption cost and exclusion unchanged. The later diffusion parameters describe the starting share (x_0), the evidence needed to establish effectiveness (p_0, x_v), the speed of action (λ_0, λ), temporary shocks (τ) and the value of waiting (W). They help determine the path taken through the incentives. Figures 11 and 12 show those path comparisons with their complete baseline settings; Appendix L provides the full reference.

C Formal statements

This appendix collects the assumptions and results behind the walkthrough. The main text explains their use; the next appendix gives the proofs. Static indifference equilibria, sheltering boundaries and rest points under irreversible adoption are distinguished explicitly.

Unless stated otherwise, the economic applications use nonnegative contest and exclusion pressure, $r \geq 1$, and fixed beliefs. A stable static reference means that the gain crosses zero downward (or $g(x) - x$ does so with dispersed costs). This is stability under continuous adjustment toward the best response, not a guarantee that the discrete map $x_{t+1} = g(x_t)$ converges or that irreversible adoption returns after an overshoot.

Lemma 1 (Zero sum). *For all $x < x_Q$, $x \cdot \alpha(1 - x)/D(x) + (1 - x) \cdot (-\alpha x/D(x)) = 0$, and for $x \geq x_Q$ both terms are zero. The contest never enters the aggregate payoff of the rivals in it.*

Lemma 2 (Shape of the contest). *The contest's contribution to the switching gain is $\alpha/D(x)$, with slope $-\alpha(r - 1)/D(x)^2 = -V(r - 1)^2/D(x)^2$.*

- (a) *For a positive prize ($V > 0$) it is strictly decreasing in x whenever $r \neq 1$, and constant when $r = 1$.*
- (b) *As $r \rightarrow 1$ with α fixed, the payoffs converge uniformly to the linear forms $\alpha(1 - x)$ and $-\alpha x$.*
- (c) *As $r \rightarrow \infty$ with V fixed, the holdout's positional payoff converges to $-V$ for every $x > 0$, a step at zero, while the adopter's converges to $V(1 - x)/x$.*

Definition 1. Evaluate Δ at $p = 1$ for the mean-cost rival, and let $y_R = W_R(1) - W_R(0)$ and $y_S = W_S(1) - W_S(0)$.

- (i) The technology is *forced* if $\Delta(1) > 0$: at full adoption, holding it beats the best non-adopting response in realized flow payoff, so a mean-cost adopter would not drop it. Each other cost type has its own test, obtained by replacing c with c_i . Waiting value is excluded because this is a flow-payoff classification, not a solution of the timing decision.
- (ii) It is *competitively compelled* if $\Delta(0) \leq 0$: a mean-cost rival would not strictly gain by adopting alone.
- (iii) It is *wasteful for rivals* if $y_R < 0$ and *if $y_S < 0$.*

Proposition 1 (Contests are strategic substitutes). *The slope of the payoff gap is*

$$\Delta'(x) = p\beta - \frac{p\alpha(r-1)}{D(x)^2} + \varepsilon\Phi'(x) \quad \text{for } x < x_Q, \quad \Delta'(x) = p\beta \quad \text{for } x > x_Q,$$

and at x_Q the gap jumps down by the adopter's edge $p\alpha(1 - x_Q)/D(x_Q)$. If $\beta = \varepsilon = 0$, $\alpha \geq 0$ and $r \geq 1$, then Δ is non-increasing, jump included. The tipping pattern is then impossible, the homogeneous game has no unstable interior equilibrium, and with $\sigma > 0$ a continuous best-response map has a unique stable fixed point in the sense defined above. A downward sheltering jump can skip the fixed point; it then gives a stopping boundary for one-way adoption rather than an indifference equilibrium. Within these restrictions, tipping requires a network gain ($\beta > 0$) or exclusion ($\varepsilon > 0$).

Proposition 2 (Selective forcing). *Let $\beta = \varepsilon = 0$, $\sigma = 0$, $p = 1$, $V > 0$, $r > 1$, let holdouts have no outside option, and write $\kappa \equiv c - \gamma$. If $\alpha/r < \kappa < \alpha$, then $\Delta(0) > 0 > \Delta(1)$, and the unique equilibrium is*

$$x^\circ = \frac{\alpha/\kappa - 1}{r - 1} \in (0, 1),$$

which is stable. At x° holdouts receive $U_N(x^\circ) = -\kappa x^\circ$, and so do adopters net of their cost: rivalry dissipates the edge.

Proposition 3 (Shakeout). *Under the conditions of Proposition 2, give holdouts an outside option with $-\kappa x^\circ < Q < 0$. They shelter at*

$$x_Q = \frac{-Q}{\alpha + (r - 1)Q} \in (0, x^\circ),$$

which is the boundary at which the switching gain changes from positive to negative. At x_Q every holdout has taken the outside option, no adopter holds an edge, and $\Delta(x_Q^+) = -Q - \kappa < 0$. Under irreversible, noiseless adoption, new adoption stops at or beyond this boundary. This is not an interior Nash equilibrium of the reversible payoff game: committed adopters would prefer to shelter if they could switch.

Proposition 4 (Exclusion forces without changing full-adoption welfare). *Raising ε raises $\Delta(1)$ one for one as long as the full-adoption pressure stays below $-Q$, but leaves y_R and y_S unchanged. Its welfare cost, $(1 - x)\varepsilon\Phi(x)$, is borne by exposed holdouts during the transition and vanishes at $x \in \{0, 1\}$.*

Corollary 1 (What an outside option can disarm). *For nonnegative contest and exclusion pressure, $\Delta(1) \geq p(\gamma + \beta) - c$. When the outside option binds, $\Delta(1) = p(\gamma + \beta) - Q - c$; the lower bound is attained by a costless outside option $Q = 0$. An outside option can disarm forcing that comes from the contest or from exclusion, never forcing that comes from standalone or network gains.*

Proposition 5 (Private-social wedge). *Let $p = 1$, $\sigma > 0$, and let the cheapest share x adopt. For $x \neq x_Q$,*

$$W'_S(x) - \Delta(x; c(x)) = \beta x + \eta - \left[\frac{\alpha}{D(x)} + (1 - x)\varepsilon\Phi'(x) \right] \mathbf{1}\{x < x_Q\},$$

where $\Delta(x; c(x))$ is the marginal adopter's private gain. At an interior allocation where the marginal adopter is indifferent, the wedge itself determines whether adoption is socially excessive ($W'_S < 0$). If $\beta x + \eta \geq 0$, this occurs if and only if holdouts are exposed and

$$\frac{\alpha}{D(x)} + (1 - x)\varepsilon\Phi'(x) > \beta x + \eta.$$

Corollary 2 (Rivalry as a Pigouvian subsidy). *At an interior private equilibrium in the linear model ($r = 1$, $\varepsilon = \eta = 0$) with exposed holdouts, adoption is excessive if and only if $\alpha > \beta x$. The marginal private-social wedge is zero when $\alpha = \beta x$; that equality alone does not establish a global welfare optimum.*

Corollary 3 (Pure exclusion over-forces). *With only exclusion pressure ($\alpha = \beta = \eta = 0$, $\Phi' > 0$ and holdouts exposed), every interior private equilibrium has excessive adoption. This is a marginal statement. Whether adoption raises welfare relative to the status quo also depends on standalone gains, real costs and the burden placed on laggards.*

Proposition 6 (A sufficient condition against a complete ban). *Let $p = 1$, $c > 0$, $0 < \sigma < 1$, and let holdouts have no outside option. Then*

$$W'_S(0^+) = \gamma + \eta - \varepsilon \Phi'(0).$$

If this is positive, a complete ban is not welfare-maximizing. If also $W_S(1) = \gamma + \beta + \eta - c < 0$, a welfare-maximizing allocation exists strictly between zero and full adoption.

Proposition 7 (An interior maximum of contest waste). *Consider the linear pure contest: $r = 1$, $\alpha > 0$, $\gamma = \beta = \varepsilon = \eta = W = 0$, no outside option and known effectiveness. Let costs be $c_i = cR_i$, where R_i has the mean-one log-logistic distribution with $0 < \sigma < 1$. At the static equilibrium,*

$$x(c) = \Pr(R_i < \alpha/c), \quad \mathcal{L}(c) = -W_R(x(c)) = c \int_0^{\alpha/c} u f_R(u) du.$$

The loss is positive for every finite $c > 0$ and tends to zero as either $c \downarrow 0$ or $c \uparrow \infty$. It therefore attains a maximum at an interior cost level.

Proposition 8 (Aggregation). *In the continuum limit:*

- (a) *Persistent heterogeneity only ($\tau = 0$). If the initial adopters are the cheapest x_0 and z_t is non-decreasing in t , then*

$$x_{t+1} = x_t + \lambda_t \max\{0, F_c(z_t) - x_t\},$$

where F_c is the cost distribution function.

- (b) *Transitory heterogeneity only ($\sigma = 0$). Then $x_{t+1} = x_t + (1 - x_t) \lambda_t \Lambda((z_t - c)/\tau)$, which with $\lambda_t = 1$ is the discrete-time logit hazard of panel adoption models.*

Proposition 9 (Dormancy). *Let $\sigma > 0$, $\tau = 0$ and $\lambda_0 \in (0, 1)$. Let the cheapest $x_0 < x_v$ be seeded, let there be no exogenous demonstration, and let the waiting value W be constant before the demonstration. Suppose $z(x, p_0)$ is non-decreasing in x , as it is whenever $r = 1$, $\beta, \varepsilon \geq 0$ and no outside option binds below x_L . Let x_L be the smallest $x \geq x_0$ with $g_{p_0}(x) \leq x$. Before the demonstration, adoption rises monotonically toward x_L and stays strictly below it, unless $x_L = x_0$, in which case no one adopts. An endogenous demonstration therefore occurs if and only if $x_v < x_L$, and its date grows without bound as $x_v \uparrow x_L$.*

Observation 1 (Overshoot under an eroding edge). Take $V = 0.5$, $r = 6$, $c = 0.6$ and $\sigma = 0.4$, for which the static stable equilibrium is $x^\circ \approx 0.65$. The exact continuum path ends at 0.85 with $\lambda = 0.6$ and at 0.77 with $\lambda = 0.1$. The closed recursion ends at the static equilibrium (Figure 10). Realized rival welfare is -0.41 on the fast path against -0.22 at the static equilibrium. The overshoot is a property of myopic, irreversible adoption with cost-independent frictions. Whether it survives forward-looking timing in the sense of Fudenberg and Tirole (1985) remains open.

Proposition 10 (Counters). *Let $p = 1$, $\alpha > 0$, $r \geq 1$ and $Q = -\infty$, with $x, y \geq 0$ and $x + y \leq 1$. For $r > 1$ this is a positive-prize contest; $r = 1$ denotes the fixed- α linear limit. Set $\tilde{x} = 0$ at the no-adopter corner $x = 0, y = \varphi = 1$, as in the implementation. Then:*

- (a) *Nesting.* With $y = 0$ or $\varphi = 0$, the payoffs (13) are (6)–(7).
- (b) *Zero sum.* The contest terms sum to zero over adopters, counterers and plain holdouts. Rivals' welfare is therefore $W_R(x, y) = \gamma x + \beta x^2 - (1 - x)\varepsilon\Phi(x)$ minus the costs paid by adopters and counterers.
- (c) *Counters complement adoption and each other.* $z_K(0, y) = 0$. For $x > 0$ and $\varphi > 0$, z_K is increasing in x and in y .
- (d) *Counters substitute for adoption.* For $x > 0$ and $\varphi > 0$, $z_A - z_K$ is decreasing in y . z_A itself is decreasing in y if $r > 1$ and constant if $r = 1$.
- (e) *A full counter restores the split.* With $\varphi = 1$, $x > 0$ and $x + y = 1$, adopters and counterers both receive the no-technology share of the prize.

Corollary 4 (What a counter can disarm). *Under Proposition 10's assumptions, with $c_K \geq 0$, $\Delta_K(1) \geq \Delta(1) - \alpha/r = \gamma + \beta + \varepsilon - c$. A technology whose standalone, network and exclusion gains alone exceed its cost cannot be disarmed by this counter. The counter removes positional pressure; it does not supply compatibility or satisfy the gatekeeper's requirement.*

Proposition 11 (Nesting). (i) *Threshold model.* With $\alpha = \beta = 0$, $r = 1$, $\gamma = B$, $c = C + Q$, $W = W$, $\varepsilon = S$ and no outside option, the switching gain is $\tilde{\Delta}(x) \equiv \Delta(x) - W = S\Phi(x) - G$ with $G = C + W + Q - B$. The threshold model's Q is a constant opportunity cost of adopting, which the unified model folds into c ; it is not the outside option of (7), which caps a holdout's loss instead of shifting every gain. Its four statuses (independent, forced, latent, not forceable) are $\tilde{\Delta}(0) > 0$; $\tilde{\Delta}(0) \leq 0 < \tilde{\Delta}(x)$; $\tilde{\Delta}(x) \leq 0 < \max \tilde{\Delta}$; and $\max \tilde{\Delta} \leq 0$.

(ii) *Linear model.* With $r = 1$, $\varepsilon = W = \eta = 0$ and no outside option, payoffs, welfare, the two tests $F = \Delta(1)$ and $y = y_R$, and the critical mass $x^* = (c - \gamma - \alpha)/\beta$ all coincide with the linear model. Its diffusion recursion is Proposition 8(a), which is exact because z_t is non-decreasing when $\beta \geq 0$ and $\gamma + \alpha \geq 0$.

Proposition 12 (Observational equivalence). *Adoption paths, at any level of aggregation, depend on the payoff parameters through the gross gain $(x, p) \mapsto z(x, p)$ of Appendix E, together with F_c , the waiting-value process, and the diffusion parameters $(p_0, x_v, \lambda_0, \lambda, \tau)$. Holding the initial state, cost distribution, waiting-value process and diffusion rules fixed, parameter vectors that induce the same gain map produce the same adoption paths but in general different welfare. For example, with $r = 1$, no outside option and $W = 0$, moving δ from γ to α leaves every path unchanged and lowers $W_R(x)$ by δx . The result need not hold if that change also alters the edge-dependent waiting value.*

D Proofs

Lemma 1. For $x < x_Q$, $x\alpha(1 - x)/D - (1 - x)\alpha x/D = 0$. For $x \geq x_Q$ the adopter's edge is switched off in (6) and the holdout's contest loss is replaced by Q in (7), so both contest terms are zero.

Lemma 2. The switching gain's contest term is $\alpha(1 - x)/D + \alpha x/D = \alpha/D$. Differentiating gives $-\alpha(r - 1)/D^2$, and $\alpha(r - 1) = V(r - 1)^2$, which gives (a). For (b), $|\alpha(1 - x)/D - \alpha(1 - x)| = |\alpha|(1 - x)x|r - 1|/D \leq |\alpha||r - 1|/\min_x D$, and similarly for the holdout term; both vanish as $r \rightarrow 1$. For (c), $-V(r - 1)x/(1 + (r - 1)x) \rightarrow -V$ for $x > 0$, and $V(r - 1)(1 - x)/(1 + (r - 1)x) \rightarrow V(1 - x)/x$.

Proposition 1. Differentiate each branch of (8). The jump at x_Q follows from $L(x_Q) = -Q$, which gives $\Delta(x_Q^+) - \Delta(x_Q^-) = -p\alpha(1 - x_Q)/D(x_Q)$. With $\beta = \varepsilon = 0$, $\alpha \geq 0$ and $r \geq 1$, $\Delta' = -p\alpha(r - 1)/D^2 \leq 0$ below x_Q and zero above it, and the jump is downward, so $\Delta(1) \leq \Delta(0)$ and the tipping pattern ($\Delta(0) \leq 0 < \Delta(1)$) is impossible. An unstable interior equilibrium of the homogeneous game needs Δ to cross zero from below, which a non-increasing Δ cannot do. With $\sigma > 0$, a continuous non-increasing best-response map $x \mapsto F_c(z(x))$ crosses the diagonal exactly once, from above. At a downward discontinuity the sign can change without equality; this defines a sheltering boundary rather than a fixed point.

Proposition 2. Here $\Delta(x) = \alpha/D(x) - \kappa$, which is strictly decreasing since $r > 1$. $\Delta(0) = \alpha - \kappa > 0$ and $\Delta(1) = \alpha/r - \kappa < 0$, so there is a unique root, solving $D(x^\circ) = \alpha/\kappa$. It lies in $(0, 1)$ because $1 < \alpha/\kappa < r$. It is stable because $\Delta > 0$ below it and $\Delta < 0$ above it, and neither corner is an equilibrium. Finally, $U_N(x^\circ) = -\alpha x^\circ/D(x^\circ) = -\kappa x^\circ$, and $U_A(x^\circ) = U_N(x^\circ)$ because $\Delta(x^\circ) = 0$.

Proposition 3. The exposed loss $\alpha x/D(x)$ is increasing in x and equals κx° at x° (Proposition 2), so $-\kappa x^\circ < Q < 0$ puts x_Q , the root of $\alpha x/D(x) = -Q$, in $(0, x^\circ)$. Solving gives $x_Q = -Q/(\alpha + (r - 1)Q)$, whose denominator is positive because $-Q < \alpha/(r - 1)$, the supremum of the exposed loss. For $x < x_Q < x^\circ$, $\Delta(x) = \alpha/D(x) - \kappa > \alpha/D(x^\circ) - \kappa = 0$. For $x \geq x_Q$ holdouts have sheltered, so $\Delta(x) = \gamma - c - Q = -Q - \kappa$, which is negative because $-Q < \kappa x^\circ < \kappa$. Hence the sign changes at x_Q . Under irreversible, noiseless adoption, new commitments cease once the path reaches or passes it. Since $\Delta(x_Q) < 0$, it is not an interior indifference equilibrium of the reversible game. At x_Q the adopter's edge is switched off by (6).

Proposition 4. By (9), $\Phi(1) = 1$ gives $\partial\Delta(1)/\partial\varepsilon = 1$ while $p\alpha/r + \varepsilon < -Q$. $\Phi(0) = 0$ and $1 - x = 0$ at $x = 1$ remove ε from $W_R(0)$ and $W_R(1)$.

Proposition 5. $C'(x) = c(x)$. For $x < x_Q$, differentiating (15) gives $W'_S(x) = \gamma + 2\beta x + \varepsilon\Phi - (1 - x)\varepsilon\Phi' - c(x) + \eta$, and subtracting $\Delta(x; c(x)) = \gamma + \beta x + \alpha/D + \varepsilon\Phi - c(x)$ gives the bracketed wedge. For $x > x_Q$, $W'_S(x) = \gamma + 2\beta x - Q - c(x) + \eta$ and $\Delta(x; c(x)) = \gamma + \beta x - Q - c(x)$, whose difference is $\beta x + \eta$. At indifference $\Delta(x; c(x)) = 0$, so $W'_S(x)$ equals the wedge.

Corollary 1. At $x = 1$ a lone deviator faces $D = r$ and the pressure $p\alpha/r + \varepsilon$; its best non-adopting payoff is $\max\{-(p\alpha/r + \varepsilon), Q\}$. Thus $\Delta(1) = p(\gamma + \beta) - c + \min\{p\alpha/r + \varepsilon, -Q\}$. With nonnegative pressure and $Q \leq 0$ the minimum is nonnegative. If the outside option binds it equals $-Q$, which reaches zero at $Q = 0$.

Proposition 6. With no outside option, (15) is continuous on $[0, 1]$. The log-logistic quantile satisfies $c(0) = 0$, and $\Phi(0) = 0$. Differentiating from the right gives $W'_S(0^+) = \gamma + \eta - \varepsilon\Phi'(0)$. A positive derivative gives $W_S(x) > W_S(0) = 0$ for sufficiently small positive x . Continuity on the compact interval gives a maximizer. When also $W_S(1) < 0$, neither endpoint can maximize welfare, so every maximizer is interior.

Proposition 7. The private gain is $\alpha - cR_i$, so the stated cutoff defines the unique static adopted share. Contest payoffs cancel, leaving the integral of adopted resource costs. The density is positive on $(0, \infty)$, so the loss is positive and continuous for $c > 0$. Since $E[R_i] = 1$, $\mathcal{L}(c) \leq c \rightarrow 0$ as $c \downarrow 0$. Also $\mathcal{L}(c) \leq \alpha \Pr(R_i < \alpha/c) \rightarrow 0$ as $c \uparrow \infty$, since there is no atom at zero. A positive value at any

finite c and these two limits restrict the global maximum to a compact subinterval of $(0, \infty)$, where continuity guarantees attainment.

Proposition 8. (a) Let $E_t = \{q : c(q) < z_t\}$, the cheapest $F_c(z_t)$. Because z_t is non-decreasing, E_t is increasing, and every adopter is either a seed (the cheapest x_0) or entered from some $E_s \subseteq E_t$. If $F_c(z_t) \geq x_0$, all adopters lie in E_t , so the eligible non-adopters have mass $F_c(z_t) - x_t \geq 0$, and a share λ_t of them adopts. If $F_c(z_t) < x_0$, no non-seed has yet been eligible, so $x_t = x_0$ and no one adopts. (b) With a common cost, every non-adopter faces the same hazard.

Proposition 9. Write $g \equiv g_{p_0}$, which is continuous and non-decreasing because F_c is continuous and increasing for $\sigma > 0$. Before the demonstration $p_t = p_0$, so if $x_0, \dots, x_t$ are non-decreasing then so are $z_0, \dots, z_t$. Proposition 8(a) then gives $x_{t+1} = x_t + \lambda_0 \max\{0, g(x_t) - x_t\} \geq x_t$, and by induction the recursion holds throughout. By continuity $g(x_L) \leq x_L$. If $x_L = x_0$, then $g(x_0) \leq x_0$ and $x_t = x_0$ for all t . Otherwise $g(x) > x$ on $[x_0, x_L)$. For x_t in that interval, $x_{t+1} = (1 - \lambda_0)x_t + \lambda_0 g(x_t)$ is greater than x_t , and less than x_L because $g(x_t) \leq g(x_L) \leq x_L$ and $\lambda_0 < 1$. The path is increasing and bounded by x_L , so it converges. Its limit $\bar{x} \geq x_0$ satisfies $g(\bar{x}) \leq \bar{x}$, so $\bar{x} \geq x_L$ by the definition of x_L , and hence $\bar{x} = x_L$. If $x_v < x_L$ the path crosses x_v in finite time. If $x_v \geq x_L$ it never does, because $x_t < x_L$ for all t (or $x_t = x_0 < x_v$). Finally, for any horizon $\bar{T}$, $\max_{t \leq \bar{T}} x_t < x_L$, so every x_v between that maximum and x_L has a demonstration date later than $\bar{T}$.

Proposition 10. (a) $y = 0$ or $\varphi = 0$ gives $\tilde{x} = x$. (b) The contest terms sum to $(\alpha/D)[x(1 - \tilde{x}) - (1 - \varphi)y\tilde{x} - (1 - x - y)\tilde{x}] = (\alpha/D)[x - \tilde{x}(1 - \varphi y)] = 0$. Welfare is then the mean of (13) without the contest terms. (c) $z_K = p\varphi h(\tilde{x})$ with $h(u) = \alpha u/(1 + (r - 1)u)$, $h(0) = 0$ and $h'(u) = \alpha/D(u)^2 > 0$. Further, $\partial\tilde{x}/\partial x = 1/(1 - \varphi y) > 0$ and $\partial\tilde{x}/\partial y = \varphi x/(1 - \varphi y)^2 > 0$. (d) $z_A - z_K = p[\gamma + \beta x + \alpha(1 - \varphi\tilde{x})/D(\tilde{x})] + \varepsilon\Phi(x)$, whose derivative in $\tilde{x}$ is $-p\alpha(\varphi + r - 1)/D(\tilde{x})^2 < 0$. The derivative of $\alpha/D(\tilde{x})$ in $\tilde{x}$ is $-\alpha(r - 1)/D^2$, which is negative for $r > 1$ and zero for $r = 1$. (e) $\tilde{x} = x/(1 - y) = 1$, so $\alpha(1 - \tilde{x})/D = 0$, and $1 - \varphi = 0$.

Corollary 4 and (14). A single deviator leaves $y = 0$, so $\tilde{x} = 1$ and $D = r$. Then $U_A - U_N = \Delta(1)$ and $U_K - U_N = \varphi\alpha/r - c_K$, so $U_A - \max\{U_N, U_K\} = \Delta(1) - \max\{0, \varphi\alpha/r - c_K\}$. Since $c_K \geq 0$ and $\varphi \leq 1$, the maximum is at most α/r , and $\Delta(1) - \alpha/r = \gamma + \beta + \varepsilon - c$.

Proposition 11. Substitute the restrictions into (6)–(15). For (i), Φ is increasing, so $\max \tilde{\Delta} = \tilde{\Delta}(1) = S - G$, and the statuses reduce to the threshold model's ($G < 0$; $0 \leq G < S$ and $S\Phi(x) > G$; $0 \leq G < S$ and $S\Phi(x) \leq G$; $G \geq S$). For (ii), $D \equiv 1$, and the stated conditions make $z_t = p_t(\gamma + \alpha + \beta x_t)$ non-decreasing.

Proposition 12. With the waiting-value process held fixed, the hazard (11) depends on payoffs only through $z - W_t$. Identical initial states and diffusion rules then produce identical paths by induction, including identical demonstration dates. Welfare (15) depends on γ separately. The edge-dependent specification of W_t must be held fixed or disabled for this argument to apply.

E When adoption starts and where it stops

The endpoint tests and static crossings describe incentives at a given adoption share. They do not determine the share ultimately reached, the time required or whether adoption stops at a static equilibrium. Three ingredients settle those questions: when rivals learn that the technology works, how fast they act on a positive gain, and the fact that adoption cannot be undone. This

section adds them in turn. The main tool is the best-response map of Appendix E.5, which turns the payoff curves of Figure 1 into a picture of motion. For an investor, three results here matter most: a working technology can stay dormant until something demonstrates it (Appendix E.6), a gatekeeper has to apply pressure before adoption is widespread or it never bites (Appendix E.9), and a field can over-invest in an eroding edge and leave late buyers stranded (Appendix E.8).

E.1 What starts adoption

Before a demonstration event, rivals hold belief p_0 . The demonstration happens either when usage reaches a visibility threshold ($x_t \geq x_v$) or at an exogenous date. After it, $p = 1$. Cost differences or a sponsor’s initial adoption can bring usage to the visibility threshold, triggering a common revision in beliefs. A war can supply such evidence. A deadline or mandate may instead change exclusion or expected demand; representing it as a belief jump is a reduced-form approximation. Appendices E.2 and E.3 distinguish these interpretations. A global-games refinement of the learning process (Carlsson and van Damme, 1993; Morris and Shin, 1998) remains outside the baseline.

E.2 A pilot’s return mixes edge and network

A good pilot answers whether the first user gains. It does not by itself answer what happens after competitors catch up. At $p = 1$, with a common cost and no outside option, the first adopter earns $U_A(0) = \gamma + \alpha - c$, whereas each rival at full adoption earns $U_A(1) = y_R = \gamma + \beta - c$. Thus

$$U_A(0) - y_R = \alpha - \beta. \quad (10)$$

In the arms-race walkthrough, the first adopter earns 0.7 while full adoption leaves each rival at -0.2 : the pilot-to-saturation gap is the 0.9 competitive edge. In the coordination walkthrough, the pilot earns -0.4 while full adoption earns 0.4: the missing network is worth 0.8.

The pilot includes the whole first mover’s positional edge and none of the eventual network gain. It overstates the full-adoption payoff when $\alpha > \beta$ and understates it when $\beta > \alpha$. Relative to the standalone payoff $\gamma - c$, its positional component is exactly α . These are different comparisons. Nor is the pilot necessarily representative: if its cost is $c_L < c$, its advantage over the mean rival’s full-adoption payoff is $\alpha - \beta + c - c_L$.

This is an interpretation of the two curves, not a claim that rational rivals must be deceived. A rival that knows the decomposition can anticipate the edge’s erosion. The empirical problem arises when the decomposition is unknown and a pilot’s private return is treated as evidence of persistent productivity. A useful demonstration should therefore measure output and costs, outcomes for exposed holdouts, and the adopter’s return as competitors catch up. A large private return with offsetting holdout losses supports the contest channel; a gain that persists when rivals catch up supports the standalone or network channels. A successful demonstration and a wasteful arms race can coexist.

E.3 Mistaken beliefs and early preparation

The baseline demonstration rule is deliberately one-sided: it reveals a technology that works. It cannot describe a disappointing trial or a mistaken belief about the competitive stakes. For that extension, distinguish a perceived switching gain $\hat{d}_{it}$ from the switching gain d_{it}^{true} evaluated using the realized payoff primitives. An adoption episode is *belief-driven false forcing* when $\hat{d}_{it} > 0$ but $d_{it}^{\text{true}} \leq 0$ at the same exposure and costs. This differs from a wasteful bandwagon: a wasteful technology can confer a real private advantage and be adopted under correct beliefs. An unfavorable realization also does not by itself show that an ex ante decision was irrational.

A more general learning model would let beliefs rise after a good trial and fall after a bad one. It would also account for the pilot's cost and its advantage over unprepared rivals. A downward revision can stop new adoption while leaving earlier commitments stranded. Exclusion may persist after a disappointing trial because the gatekeeper's penalty is not scaled by confidence in the technology.

Firms may also act before rivals deploy. A future deadline can bring orders forward when installation takes time or capacity must be reserved. If installation is instantaneous and freely available, a future penalty alone need not justify today's operating cost. Record order, installation and productive-use dates separately before treating early purchases as evidence against competitive pressure.

Appendix R retains a Bayesian update and a forward-looking timing formulation. The baseline simulations instead use current incentives and a reduced-form waiting value. Whether anticipation removes their overshoot is an open question.

E.4 How a favorable decision becomes an actual adoption

A positive gain does not mean every firm installs the technology immediately. Budgets, build times and implementation capacity delay action. The model represents those delays with λ_t , which controls the speed of adjustment. Temporary shocks can also move a decision around its threshold; their scale is τ . Persistent differences in cost remain attached to the same actor.

The *adoption hazard* is simply the probability that an actor adopts this period, conditional on not having adopted yet. Let $z_t \equiv z(x_t, p_t) - W_t$ be the gross switching gain, where $z(x, p)$ is the payoff gap (8) before the cost at belief p : $p[\gamma + \beta x + \alpha/D(x)] + \varepsilon\Phi(x)$ while holdouts are exposed and $p[\gamma + \beta x] - Q$ once they have sheltered. A rival with cost c_i that has not yet adopted does so in period t with probability

$$h_t(c_i) = \lambda_t \Lambda\left(\frac{z_t - c_i}{\tau}\right), \quad \lambda_t = \begin{cases} \lambda_0 & \text{before the demonstration,} \\ \lambda & \text{after,} \end{cases} \quad (11)$$

where Λ is the logistic function and $\tau \geq 0$ scales a transitory shock. As $\tau \rightarrow 0$ the logistic becomes the indicator $\mathbf{1}\{c_i < z_t\}$. Adoption is irreversible. The split between λ_0 and λ is the model's version of the innovation-imitation distinction in Bass (1969).

The same decision rule can generate an aggregate adoption curve or individual adoption dates. With only persistent cost differences and a non-decreasing gain, the simple update counts eligible actors who have not yet adopted. With identical costs and temporary shocks, every remaining actor faces the same adoption probability. The exact conditions and both formulas are in Proposition 8.

The two limits leave different fingerprints even when payoffs are identical. Figure 7 holds the switching gain fixed at $z = 0.5$, below the mean cost of 0.6, so that no rival's incentive ever changes. With persistent heterogeneity only, the rivals for whom adopting pays adopt first and the rest never do. Adoption plateaus at $F_c(z) = 56\%$, and the hazard among rivals yet to adopt falls from 0.17 to zero. With transitory heterogeneity only, every non-adopter faces the same odds in every period. The hazard is flat at 0.08 and adoption creeps toward saturation. A falling aggregate hazard is therefore not evidence of a weakening incentive, because sorting on persistent costs produces one by itself. Aggregate series can tell the two apart only through the plateau. Panels tell them apart directly, because they follow the same rival over time.




Figure 7: Persistent versus transitory heterogeneity with identical, constant payoffs ($\gamma = 0.3, \alpha = 0.2, r = 1, c = 0.6, \lambda = 0.3$, so $z = 0.5$ throughout). Left: persistent cost differences produce a plateau at the share for whom adopting pays, and transitory shocks produce slow saturation. Right: the hazard falls under persistent heterogeneity only because the rivals who remain are the costly ones. Shared settings: $\beta = \varepsilon = 0, p = 1, x_0 = 0, Q = -\infty$ and $W = 0$; each legend entry gives its σ and τ .

E.5 The best-response map

Let $F_c(z)$ be the fraction of actors whose cost is below z . The *best-response map* counts the share for whom adoption would pay at the current exposure:

$$g_p(x) \equiv F_c(z(x, p) - W), \quad (12)$$

Here $z(x, p)$ is the gain before paying the actor's cost or subtracting the value of waiting. With fixed beliefs, no temporary shocks, the cheapest actors seeded first and a non-decreasing switching gain, the next share is

$$x_{t+1} = x_t + \lambda \max\{0, g_p(x_t) - x_t\}.$$

For example, if 30% have adopted, the gain currently covers the costs of 50%, and $\lambda = 0.25$, the next period adds a quarter of the twenty-point gap: adoption reaches 35%. If the eligible share is already below the adopted share, this irreversible update adds nobody. Proposition 8 gives the exact conditions under which that counting rule works.

Where the map lies above the diagonal, adoption moves a fraction λ of the way up to it. Where the map lies below the diagonal, adoption does not move. Figure 8 draws the map for a network good with dispersed costs, which is the coordination archetype of Figure 1 with heterogeneous costs. Four features of the picture hold in general.

- *Crossings identify the static reference points.* With $W = 0$, a fixed point is a static payoff equilibrium. A positive constant waiting value instead gives a switching threshold. A downward crossing gives the usual stable static reference; an upward crossing gives a tipping threshold. Here the stable equilibria are at 0.6% and 93%, and the tipping point is at 39%.
- *Irreversibility freezes failed seeds.* In the static game the low equilibrium attracts every state below the tipping point. In the dynamic model nothing pulls adoption down. A seed of 38% stays at 38%: the cheapest 38% have adopted, the next-cheapest rival still does not gain, and no adopter leaves. Every seed at which the map lies on or below the diagonal is a rest point. Two seeds two points apart, on either side of the tipping point, end at 38% and 93%.




Figure 8: The best-response map $g(x) = F_c(z(x))$ for a network good with dispersed costs ($\gamma = 0.2$, $\beta = 0.8$, $c = 0.6$, $\sigma = 0.2$, $\lambda = 1$). (a) Adoption grows where the map lies above the 45° line (blue) and stops where it lies below (gray). With $\lambda = 1$ the path is the staircase $x_{t+1} = g(x_t)$. Filled circles are stable equilibria and the open circle is the tipping point. (b) Paths from different seeds. Seeds below the tipping point freeze, because adoption is irreversible. Seeds above it trace an S-curve to the high equilibrium. Other settings: $\alpha = \varepsilon = 0$, $r = p = 1$, $Q = -\infty$, $W = \tau = 0$; seed shares are labeled in the figure.

- *The S-curve comes from the shape of the map.* Just above the tipping point the map barely clears the diagonal, so each period adds few adopters. The gap widens as adoption grows and closes again near the high equilibrium. The seed of 40% takes five periods to reach 50% and four more to reach 90%. The S-shape of diffusion curves needs no separate imitation term here: it comes from the network gain, which makes the map rise, and from the dispersion of costs, which makes it smooth.
- *The map has the slope of the switching gain.* F_c is increasing, so g_p rises where Δ rises and falls where Δ falls. Network and exclusion channels bend the map upward, the contest bends it down, and an outside option, where it binds, drops it by the adopters' edge. A continuous falling map crosses the diagonal once (Proposition 1), so a pure contest has no tipping point. Appendix E.8 shows that a falling map also no longer describes the state completely.

E.6 Why a working technology can remain dormant

Beliefs scale the technology channels, so before the demonstration rivals act on a lower map, g_{p_0} , than after it, g_1 (Figure 9a). The path therefore has two phases. A seed x_0 first creeps at the slow pre-demonstration speed λ_0 toward the first rest point of the pre-demonstration map, the *ceiling* x_L . If usage crosses the visibility threshold x_v on the way, beliefs jump, the map shifts up to g_1 , and adoption takes off at speed λ . If it does not, adoption approaches x_L and stays there. The technology works, but it never reaches the usage needed to show that it works.

Figure 9(b) shows both outcomes with identical payoffs. With $x_v = 15\%$, below the ceiling $x_L = 17.4\%$, the demonstration comes at $t = 19$ and adoption rises to 81%. With $x_v = 18\%$ it never comes, and adoption stalls at 17% until an external demonstration supplies the missing evidence. A mandate could instead change the payoffs; that is a different intervention. Panel (c) traces the demonstration date as x_v rises. The date grows slowly at first, then without bound as

dynamics, but it never enters the payoff gap or the classification. A demonstration then does two things at once: it raises the expected gain and removes the reason to wait. Figure 9(d) shows both effects with an exogenous demonstration date. With $\omega = 0.8$, no one beyond the seed adopts before the demonstration (2%, against 13% without a waiting value). The jump at the demonstration is also larger (21 points against 19). Waiting suppresses adoption before the demonstration and produces a larger jump afterward. The optimal-stopping formulation in Appendix E.3 makes this structural problem explicit. Deriving the reduced-form W_t from equilibrium in a timing game remains open (Reinganum, 1981; Fudenberg and Tirole, 1985; Dixit and Pindyck, 1994).

E.8 Why the field can invest too much before the edge disappears

The simple update works only while the gross switching gain does not fall. Under strategic substitutes (Proposition 1), the map falls and z_t falls as x_t rises. Some rivals then adopt while their cost is below z_t and stay adopted after it is not. The closed recursion subtracts all past adopters from the currently eligible pool. When some past adopters are no longer eligible, that subtraction is too large and it understates the remaining pool. The share x_t then no longer describes the state. It matters which rivals have adopted, not just how many.

For the example in Figure 10, the static equilibrium is about 65%. The exact adoption path reaches 85% with fast adjustment ($\lambda = 0.6$) and 77% with slower adjustment ($\lambda = 0.1$). Rival welfare on the fast path is -0.41 , compared with -0.22 at the static allocation. The example uses $V = 0.5$, $r = 6$, $c = 0.6$ and cost dispersion $\sigma = 0.4$.

Figure 10(b) shows the mechanism. In the first period the edge is at its largest, $z_0 = \alpha = 2.5$, and adopting pays for 99% of rivals. Sixty percent of them adopt, at every cost level. Their adoption shrinks the edge, so in the second period adopting pays only for the cheapest 70%, and 60% of those not yet in adopt. By the end, the edge pays only the cheapest 53%, fewer than at the static equilibrium, yet 85% have adopted. The 32% of rivals to the right of the final threshold are stranded. They would not adopt today, but they adopted while the edge was large and cannot undo it. Each of them pays $c_i > \alpha/D(x_T)$ every period for an edge that no longer covers it. With avoidable flow costs they would drop the technology if they could, so the overshoot and its welfare cost are the price of commitment, and with reversible adoption the path would be free to unwind toward the static equilibrium. With sunk acquisition costs they would stay, because an adopter's contest payoff $\alpha(1 - x)/D(x)$ exceeds a holdout's $-\alpha x/D(x)$ at every x , and the waste is the acquisition cost rather than a perpetual flow. Which reading applies is a fact about the technology, not the model; Section 1 fixes the flow reading throughout. Slower adjustment overshoots less: fewer rivals commit in each period, so the edge has shrunk further by the time most of them decide. The closed recursion loses track of those stranded early adopters and therefore misses the overshoot in this example.

E.9 Comparative dynamics, one parameter at a time

Figures 11 and 12 vary one parameter at a time around a common baseline technology in which all four channels are active and holdouts have no outside option: $\gamma = 0.05$, $\beta = 0.4$, $\alpha = 0.3$, $r = 2$, $\varepsilon = 0.2$ ($\theta = 0.5$, $k = 8$), $c = 0.6$ and $\sigma = 0.3$. A seed of 2% adopts first and $\lambda = 0.3$. The baseline's best-response map lies above the diagonal all the way to its rest point at 72%, but comes within five points of it at 38% adoption. Early on, the erosion of the contest's edge cancels most of the network gain, so the map rises slowly while the diagonal rises at slope one. Exclusion pressure, centered at 50%, steepens the map only later. Adoption crawls through this bottleneck and reaches 50% only in period 22. The bottleneck is why so many panels in Figure 11 show the same pattern. A modest change in one direction closes the gap, and adoption stalls below it. The same change in the other direction widens the gap, and adoption sweeps through. The response is





Figure 10: Selective forcing with an eroding edge ($V = 0.5$, $r = 6$, $c = 0.6$, $\sigma = 0.4$). (a) The exact paths overshoot the static stable equilibrium, and faster adjustment overshoots more. The closed recursion of Proposition 8(a) misses this. (b) Who has adopted by period 30 on the fast path, by cost and by the period in which they adopted. Adopting pays only to the left of the dotted line, given the final edge. The rivals to its right adopted while the edge was large and cannot undo it. Here $\alpha = V(r - 1) = 2.5$, $\gamma = \beta = \varepsilon = 0$, $p = 1$, $Q = -\infty$, $W = \tau = 0$, $x_0 = 0$; adjustment speeds are $\lambda = 0.6$ and 0.1 .

lopsided: lowering γ from 0.05 to 0 cuts final adoption from 70% to 21%, while raising it to 0.15 lifts it only to 84%, although half the field then adopts three times as fast (by period 7 rather than 22).


The panels sort the parameters by *when* in the path they act.

- *The edge acts at the start.* The contest's $\alpha/D(x)$ is largest when adopters are few, so α decides whether adoption gets going. With $\alpha = 0.1$ no one beyond the seed adopts. With $\alpha = 0.5$ half the field has adopted by period 5. The decisiveness r , holding α fixed, decides whether the edge lasts. With $r = 1$ the edge never erodes, and adoption climbs to 86%. With $r = 6$ the same initial edge erodes so quickly that adoption stalls at 15%, slightly above its static rest point of 13%, which is the overshoot of Appendix E.8 on a small scale. In this example, rapid erosion prevents the initial competitive edge from sustaining widespread adoption.
- *Network and exclusion channels act later.* In the β and ε panels, the low-value paths track the baseline for the first few periods and only then fall away, stalling at 23% and 30%. These channels do not start adoption. They decide whether it keeps going once a share of rivals has adopted. With $\beta = 0.6$ or $\varepsilon = 0.4$, adoption sweeps through the bottleneck to 88% or 89%.
- *The timing of exclusion matters as much as its strength.* Pressure centered early ($\theta = 0.2$) lifts the map before the bottleneck, and half the field adopts by period 8. Pressure centered late ($\theta = 0.8$) arrives only after a point that adoption never reaches, and adoption stalls at 32%. In this specification, exclusion must affect incentives before adoption reaches the point at which it would otherwise stall.
- *Level shifters act everywhere.* The cost c shifts the whole map, like γ in the opposite direction. Raising c from 0.6 to 0.7 stalls adoption at 19%.
- *An outside option acts at its shelter point.* With $Q = -0.35$ nothing changes, because holdouts' loss never reaches 0.35 and no one shelters. With $Q = -0.15$ holdouts shelter once



Baseline: $\gamma = 0.05$, $\varepsilon = 0.2$ ($\theta = 0.5$, $k = 8$), $c = 0.6$
 $\varepsilon = 0.2$, $\theta = 0.5$, $k = 8$, $\sigma = 0.3$, $Q = -\infty$, $W = 0$; $p = 1$, $x_0 = 2\%$, $\lambda = 0.3$, $\tau = 0$.
 Solid line = baseline; each panel changes only its named parameter.

Figure 11: Comparative dynamics of the payoff parameters. Each panel varies one parameter around the baseline ($\gamma = 0.05$, $\beta = 0.4$, $\alpha = 0.3$, $r = 2$, $\varepsilon = 0.2$, $\theta = 0.5$, $k = 8$, $c = 0.6$, $\sigma = 0.3$, no outside option; the technology is known to work, $x_0 = 2\%$, $\lambda = 0.3$). Darker lines are higher values; the solid, thicker line is the baseline, repeated in every panel. Waiting and transitory shocks are off ($W = \tau = 0$). In the r panel, $\alpha = 0.3$ stays fixed, so the implied prize changes.



Baseline payoffs: $\gamma = 0.05$, $\varepsilon = 0.2$ ($\theta = 0.5$, $k = 8$), $c = 0.6$
 $\varepsilon = 0.2$, $\theta = 0.5$, $k = 8$, $\sigma = 0.3$, $Q = -\infty$, $W = 0$. $p_0 = 0.8$, $x_v = 8\%$, $x_0 = 2\%$, $\lambda_0 = 0.1$, $\lambda = 0.3$, $\tau = 0$.
Solid line = baseline; each panel changes only its named parameter.

Figure 12: Comparative dynamics of the diffusion parameters. Payoffs stay fixed at $\gamma = 0.05$, $\beta = 0.4$, $\alpha = 0.3$, $r = 2$, $\varepsilon = 0.2$, $\theta = 0.5$, $k = 8$, $c = 0.6$, $\sigma = 0.3$, $Q = -\infty$, $W = 0$. The technology is demonstrated once usage reaches x_v , and the baseline has $p_0 = 0.8$, $x_v = 8\%$, $x_0 = 2\%$, $\lambda_0 = 0.1$, $\lambda = 0.3$ and $\tau = 0$. Open circles mark demonstrations. Darker lines are higher values; the solid, thicker line is the baseline. Each panel changes only the parameter in its title.

41% have adopted, in period 16. The adopters' edge vanishes, the map drops below the diagonal, and adoption stops at 42% instead of 70%. Rivals are worse off, at -0.136 against -0.116 : adopters keep paying for an edge over no one, and holdouts pay 0.15 a period to shelter from it. This is the shakeout of Proposition 3 with every channel on.

- *Cost dispersion can encourage early adoption while limiting its eventual extent.* At low adoption the mean-cost rival does not gain, so someone has to be cheaper than average for adoption to start. With $\sigma = 0.1$ almost no one is, and adoption never leaves the seed. That field has a high equilibrium at 94%, but it lies beyond a tipping point at 59%. With $\sigma = 0.5$ adoption is fastest, with half the field in by period 6, but it ends at 78%, below the 94% a homogeneous field could reach, because dispersion also leaves more rivals whose costs are too high ever to adopt.

Figure 12 adds a demonstration phase to the same technology and varies the diffusion parameters. In the baseline, prior beliefs sustain adoption up to a ceiling of 14.8%, the visibility threshold is 8%, and the demonstration comes in period 9. The parameters split into three groups.

- *Whether: p_0 and x_v .* They set the ceiling and the threshold of Proposition 9. Lowering the prior to 0.65 drops the ceiling to 6.9%, below x_v , so the technology never becomes visible and

adoption stays at 7%. Raising the threshold to 16%, above the ceiling, does the same. In the other direction, $p_0 = 0.95$ or $x_v = 4\%$ brings the demonstration forward to period 4 or 3.

- *When: x_0 and λ_0 .* These move the demonstration date but not the destination. With no seed the demonstration comes in period 10, and a seed of 10%, above x_v , makes it immediate. Slow pre-demonstration adjustment ($\lambda_0 = 0.04$) delays it to period 21, and fast adjustment ($\lambda_0 = 0.3$) brings it to period 3. Every one of these paths heads for the same 72% rest point.
- *How fast, and sometimes how far: λ and τ .* The post-demonstration speed λ sets the steepness of take-off: half the field adopts by period 19 with $\lambda = 0.6$, and fewer than half by period 50 with $\lambda = 0.15$. Because the baseline map rises, λ does not change where adoption ends. Under an eroding edge it would (Appendix E.8). Transitory shocks do change where adoption ends. A rival whose persistent cost exceeds the gain still adopts in a period when its shock is favorable, and it cannot undo that. With $\tau = 0.15$, adoption passes the 72% rest point and reaches 93%. The same shocks bring the demonstration forward, to period 3.

The payoff parameters determine the available incentives; the diffusion parameters determine whether and when actors act on them. The exceptions are λ under an eroding edge and τ , which also change where adoption ends.

E.10 Weighting the rivals and partners that matter

A global adoption percentage can miss the relevant pressure. One large customer, dominant marketplace or close competitor may matter more than dozens of distant users. Assign weight w_{ij} to actor j 's influence on actor i , and let $a_j = 1$ if j has adopted. Then the exposure is

$$x_i = \frac{\sum_j w_{ij} a_j}{\sum_j w_{ij}}.$$

This is the network-weighted counterpart of the share x (Galeotti et al., 2010). Contest and exclusion exposure use a competition network, and network gains use a compatibility network, which need not coincide. A gatekeeper is a node with very large weight, which is why exclusion often looks like a step. On a star network, seeding the hub can tip adoption at parameters where seeding one leaf does not. Who adopts matters more than how many. On general networks the contest is zero-sum only pairwise, so welfare is computed from realized payoffs.

F Counters: a third choice

So far a rival that does not adopt simply holds out, and every other response is folded into the constant Q . Many forcing technologies, however, meet a counter: torpedoes and mines against battleships, ad-blockers against advertising, speed bumps against latency arbitrage, detection against synthetic media. A counter differs from a second positional technology in two ways. It is worthless unless someone has adopted the technology it counters. And it works by taking away the adopter's edge rather than by adding an edge of its own. This section adds the counter as a third choice.

F.1 Protection against the loss

A counter pays for the loss it avoids. If exposure would cost 0.10 per period, a defense that shields 80% of that loss saves 0.08. At a cost of 0.05, it leaves its user 0.03 better off than remaining unprotected. With no exposure it saves nothing and still costs 0.05. These illustrative amounts show why a counter's value depends on deployment of the technology it counters. Each rival adopts

the forcing technology (A), adopts a counter (K), or does neither (N). Let x and y be the shares of adopters and counterers. A counter with *efficacy* $\varphi \in [0, 1]$ shields a share φ of its holder's stake from the contest. Adopters can raid only the unshielded pool, which holds the adopters, the plain holdouts, and the unshielded $1 - \varphi$ of each counterer. Within that pool the prize is divided in proportion to strength, as in Section 1. The contest terms are therefore the base model's, evaluated at the adopters' share of the pool,

$$\tilde{x} \equiv \frac{x}{1 - \varphi y},$$

with $\tilde{x} = 0$ when $x = 0$ (including the fully shielded corner), and the three payoffs are

$$\begin{aligned} U_A &= p \left[\gamma + \beta x + \frac{\alpha(1 - \tilde{x})}{D(\tilde{x})} \right] - c, & U_N &= -p \frac{\alpha \tilde{x}}{D(\tilde{x})} - \varepsilon \Phi(x), \\ U_K &= -p(1 - \varphi) \frac{\alpha \tilde{x}}{D(\tilde{x})} - \varepsilon \Phi(x) - c_K. \end{aligned} \tag{13}$$

Counterers still lack the technology, so exclusion applies to them as to plain holdouts. The outside option of Section 1 is switched off in this section. The full counter $\varphi = 1$ is its within-segment analogue: it shelters the whole stake at cost c_K but leaves its holder in the segment and subject to the gatekeeper, which is why a counter cannot disarm exclusion (Corollary 4) while an outside option can (Corollary 1). Counter costs are log-logistic with mean c_K and dispersion σ_K , independent of adoption costs. Each rival takes the best of the three choices, and both commitments are exclusive and irreversible. The gross gains from adopting and from countering, over doing neither and before a rival's own cost, are

$$z_A(x, y) = p \left[\gamma + \beta x + \frac{\alpha}{D(\tilde{x})} \right] + \varepsilon \Phi(x), \quad z_K(x, y) = p \varphi \frac{\alpha \tilde{x}}{D(\tilde{x})}.$$

The counter is worth the share φ of the loss that a plain holdout would suffer.

The simpler alternative, giving counterers a strength k in the same pool, does not produce a counter. Such a “counter” would pay even when no one has adopted. And adoption would lower its value rather than raise it, because in a proportional-share contest every strength-raising investment substitutes for every other (Proposition 1). That is a second positional technology. Shielding distinguishes a counter from another investment in competitive strength, and the full counter $\varphi = 1$ plays the role of a counter as strong as the edge.

A counter changes the incentives of the rivals who have not protected themselves. Each additional shield concentrates the contest on the remaining exposed stakes, increasing the reason for other holdouts to counter. The defense can become hard to avoid in its own right.

Countering also makes adoption less attractive relative to buying protection. When the edge erodes ($r > 1$), it reduces the adoption gain relative to doing nothing as well. In the linear limit it lowers the adopter's and plain holdout's payoffs by the same amount. The contest still cancels across all three groups, so counter costs remain in the welfare total.

F.2 Whether protection makes adoption avoidable

At full adoption, a rival that drops the technology need not hold out unprotected; it can counter instead. With $p = 1$ and the mean-cost rival, its gain from adopting over the best non-adopting response is

$$\Delta_K(1) = \Delta(1) - \max \left\{ 0, \frac{\varphi \alpha}{r} - c_K \right\}. \tag{14}$$

The technology is forced against the counter if $\Delta_K(1) > 0$. The counter *disarms* it if $\Delta(1) > 0 \geq \Delta_K(1)$, that is, if $c_K \leq \varphi \alpha / r - \Delta(1)$. A counter does not change the first-mover test at $x = 0$: with no deployed technology to defend against, a non-adopter gains nothing from countering.




Figure 13: Selective forcing with a counter ($V = 0.3$, $r = 4$, $c = 0.4$, $\sigma = 0.3$; counter $\varphi = 0.8$, $c_K = 0.05$, $\sigma_K = 0.3$; $x_0 = 2\%$, $\lambda = 0.3$). (a) The state space. Left of the blue curve, adopting pays for more rivals than have adopted. Right of the orange curve, countering pays for more rivals than have countered. The path (one dot per period) moves right first and turns up only once adoption makes the counter pay. It ends far from the static equilibrium, and short of the open circle, where adoption ends without a counter. (b) Shares over time. (c) Contest payoffs of the mean rival of each kind, with and without the counter. Here $\alpha = 0.9$, $\gamma = \beta = \varepsilon = 0$, $p = 1$, $Q = -\infty$, $W = \tau = 0$; no rivals initially hold a counter.

The limit is practical as well as mathematical. A counter can end an arms race over a positional edge. It cannot end a network good or a gatekeeper's mandate, because no rival's counter restores the compatibility or admissibility that holding out forfeits.

F.3 The head start of the original technology

Because a counter is worthless until the technology spreads, the two diffuse in sequence. Figure 13 shows this for the selective-forcing technology of Figure 1 ($V = 0.3$, $r = 4$, $c = 0.4$), with dispersed costs and a counter with $\varphi = 0.8$ and $c_K = 0.05$. In the first period 30% of rivals adopt and no one counters. Counters spread from the second period, once 43% have adopted, and by period 15 almost no plain holdouts remain. The path ends with 52% adopters and 48% counterers, against 68% adopters without the counter. Panel (c) shows what happens to the contest. The adopter's edge erodes further than without the counter, to 0.04 rather than 0.09. Plain holdouts lose *more* than without the counter (0.215 against 0.201), although fewer rivals adopt, because adopters' raids concentrate on them. That displacement is why every holdout ends up countering.

The static equilibrium, which ignores the order in which rivals commit, would hold adoption to 21%. Irreversibility again separates the path from the static equilibrium. Many of the rivals who would rather counter at the end adopted at the start, when the counter was worthless. The gap is larger in the decisive arms race of Figure 14(a) ($V = 0.9$, $r = 2$, $\gamma = 0.1$, $c = 0.3$). There a counter with $\varphi = 0.9$ and $c_K = 0.05$ disarms the technology ($\Delta_K(1) = -0.105$) and would hold adoption to 45% in the static equilibrium. Along the path, adoption still reaches 83%, against 97% without the counter. Slowing adjustment to $\lambda = 0.02$ lowers that only to 81%, because the counter's value grows only as fast as adoption itself. Under these irreversible commitments, a counter can stop further spread but cannot reverse earlier adoption. Whether forward-looking rivals who anticipate the counter would wait for it is the same open question as for overshoot (Observation 1).

F.4 Whether protection pays for itself

Protecting a user from a transfer loss redistributes the contest; it does not create value for the rival group by itself. The counter enters that group's welfare through its cost. It raises rivals' welfare only




Figure 14: Change in rivals' welfare from making a counter available, by efficacy φ and mean cost c_K ($\sigma_K = 0.3$; welfare at period 80 from a 2% seed with $\lambda = 0.3$). (a) Decisive arms race ($V = 0.9$, $r = 2$, $\gamma = 0.1$, $c = 0.3$, $\sigma = 0.3$). Below the dashed line the counter disarms the technology, $\Delta_K(1) \leq 0$. (b) Selective forcing ($V = 0.3$, $r = 4$, $c = 0.4$, $\sigma = 0.3$). The ring marks the counter of Figure 13. Solid curves separate gains (blue) from losses (red). In both panels, $\beta = \varepsilon = W = \tau = 0$, $p = 1$, $Q = -\infty$, and $\gamma = 0$ in (b); counters start at zero and adjust at the same speed as adopters.

by deterring adoption whose cost exceeds its standalone and network value. Figure 14 maps the welfare change over the counter's efficacy and cost. In both technologies, cheap, effective counters raise welfare, by up to 0.06 in the arms race and 0.08 under selective forcing. Under selective forcing, however, a wide band of moderately priced counters lowers welfare, by up to 0.017. Holdouts buy them to cut losses that are transfers to adopters, and they deter little adoption. Such a counter is a second arms race layered on the first. In the arms race the corresponding losses are negligible, because few rivals are left holding out to buy the counter. The question for a regulator weighing a counter is therefore not whether it protects those who buy it, but whether it deters enough adoption to pay for itself.

The extension keeps three restrictions worth relaxing. Commitments are exclusive and irreversible, so counterers cannot later adopt, and adopters cannot abandon the technology once counters spread. The obsolescence that ends many forcing technologies needs both. The model also excludes an outside option alongside the counter, so rivals cannot choose to leave the segment. Finally, it restricts counter adoption to members of the rival group. Counters adopted by outsiders, such as users' ad-blockers, or aimed at threats from outside the rival set, such as quantum attacks for post-quantum cryptography, act instead through a standalone gain that changes over time, which the model does not yet allow.

G Who gains when another rival adopts

For a firm's decision, the relevant comparison is its own two payoffs. For an industry body, an investor assessing the whole market, or a policymaker, add the effects on everyone inside the chosen boundary. That sum is the welfare calculation, and it is the buyer-side accounting that the company arguments in the main text deliberately leave aside.

Suppose the technology works and adoption is assigned to the cheapest share x . Let $C(x) = \int_0^x c(q) dq$ be their total cost. Rival welfare is

$$W_R(x) = \gamma x + \beta x^2 + (1-x)h(x) - C(x), \quad h(x) = \begin{cases} -\varepsilon\Phi(x), & x < x_Q, \\ Q, & x \geq x_Q. \end{cases} \quad (15)$$

Read the terms as direct gains, network gains, holdouts' exclusion or sheltering costs, and adoption costs. Here h records the holdout component left after contest transfers cancel; it is not the exposed holdout's entire payoff. On a historical path where adopters are not the cheapest actors, use the costs actually incurred.

The contest terms cancel from aggregate welfare. Winning business from a rival may justify your investment, but that transfer adds nothing to rivals' combined payoff. Exclusion is different: the lost value is not captured elsewhere in the group. Sheltering also carries a real loss relative to the original baseline.

Add outsiders explicitly:

$$W_S(x) = W_R(x) + \eta x.$$

Customers can gain from lower prices or better service even when competing suppliers' returns fall. A collective wish to stop competing is therefore not enough to establish a social case for stopping adoption.

The marginal question is whether *one more* adoption helps. Before sheltering, the difference between its social gain and its private gain is

$$\underbrace{W'_S(x) - \Delta(x; c(x))}_{\text{social gain minus private gain}} = \underbrace{\beta x + \eta}_{\text{benefits to others}} - \underbrace{\left[\frac{\alpha}{D(x)} + (1-x)\varepsilon\Phi'(x) \right]}_{\text{losses imposed on others}}.$$

The actor misses the network benefit it gives existing adopters and the benefit it gives outsiders. It also ignores the competitive loss it imposes and any extra exclusion pressure on remaining holdouts. At an interior allocation where the next adopter is indifferent, more adoption reduces welfare when the second bracket exceeds the first. This is the private–social wedge proved in Appendix D.

The result allows both too much and too little adoption. A positional incentive can offset an otherwise insufficient network incentive; with only exclusion and no network or outsider benefit, the marginal private equilibrium over-adopts. Once holdouts have sheltered, the competitive and exclusion terms disappear from the marginal comparison. These are statements about an additional adoption, not a substitute for comparing the total value of different adoption levels.

H What changes as a market matures

A decision can change because the technology improves, holdouts leave, customers raise their expectations or suppliers lose pricing power. Those changes can look like ordinary diffusion in an adoption chart. The following extensions separate them. They are scenarios to investigate, not additional channels estimated in the historical fits.

H.1 A technology can move from doubtful to necessary to routine

Let maturity m change benefits and costs. In the threshold restriction with $W = 0$, write $G(m) = c(m) - \gamma(m)$ and $S(m) = \varepsilon(m)$. A possible sequence is

$$G(m) > S(m) \quad \longrightarrow \quad 0 < G(m) < S(m) \quad \longrightarrow \quad G(m) < 0.$$

Initially even maximum exposure does not justify adoption. At intermediate maturity adoption needs competitive pressure. Eventually it pays without that pressure. This is a window of *competitive compulsion*, not necessarily a window of forcing: the final technology can still satisfy $\Delta(1) > 0$. Nor is the sequence inevitable. Costs, outside options and the competitive stake can move in different directions, and a technology may never enter the middle regime.

The related business lifecycle is differentiator, necessity, then infrastructure. In the contest, the differentiating return disappears as rivals catch up even with fixed primitives. Falling costs or new standards can then make the capability routine, while a new generation creates a fresh positional edge. This connects successive forcing episodes to the Red Queen process of competition and organizational learning (Barnett and Hansen, 1996). The distinction is observable: estimate whether a technology's three signs change as its cost and performance change, rather than interpreting every change in its adoption rate as movement along one fixed curve. A forecast of rapid improvement can also raise the value of waiting for the next generation, so improving prospects need not accelerate purchases of the current one.

H.2 Holdouts can adopt or disappear

Suppose twenty firms adopt and eighty hold out. If holdouts then leave faster, adoption among the survivors rises even if no additional firm buys the technology. An industry can become almost fully adopted without most initial holdouts ever adopting. Let A_t and N_t count active adopters and active holdouts, with $\bar{x}_t = A_t/(A_t + N_t)$. In continuous time, let h_t be the holdout adoption hazard and d_A, d_N the exit hazards. With no entry or reversal,

$$\dot{A}_t = h_t N_t - d_A A_t, \quad \dot{N}_t = -h_t N_t - d_N N_t,$$

so

$$\dot{\bar{x}}_t = h_t(1 - \bar{x}_t) + \bar{x}_t(1 - \bar{x}_t)(d_N - d_A). \quad (16)$$

The first term is adoption within the population. The second is selection: if holdouts exit faster, the active-adopter share rises even with $h_t = 0$. With entry rates b_A, b_N , add $[(1 - \bar{x}_t)b_A - \bar{x}_t b_N]/(A_t + N_t)$. Market-share measures require a further decomposition for growth among surviving firms.

The outside option Q describes a response by members of the original field; it does not change the denominator of the baseline's x . Equation (16) describes a different statistic. Estimation should retain the original cohort, report exits and alliances separately, and also report adoption among survivors. Entry, failure and changes in output weights must not be fitted as if they were individual adoption decisions. Strategic anticipation and selection can coexist, but only the former requires that rivals foresee the threat.

H.3 Expectations can outlast the deployments that created them

A capability can become necessary because customers learn to expect it. Fast delivery, easy returns and online service illustrate the mechanism: an attribute can move from an attractive extra to a minimum requirement, in the terminology of Kano et al. (1984). A simple extension introduces an expectation state $b_t \in [0, 1]$,

$$b_{t+1} = b_t + \rho \max\{0, x_t - b_t\}, \quad 0 < \rho \leq 1,$$

and makes a holdout's demand loss depend on b_t . Expectations follow adoption upward and persist if adoption subsequently recedes. The same current share can then exert different pressure after different histories; x_t alone no longer describes the state.

Replacing $\varepsilon\Phi(x_t)$ by $\varepsilon\Phi(b_t)$ approximates a penalty that no rival captures. If lost demand goes to adopting rivals, however, that portion is a transfer and must be entered on their payoff curve as

well. If buyers enjoy better service, their gain belongs in social welfare. Customer pressure should therefore be decomposed rather than automatically counted as exclusion waste. Repeated customer surveys, purchasing behavior and holdout revenues can distinguish an expectation ratchet from a platform’s explicit access rule.

H.4 Why a necessary capability need not pay its company

A customer may need a compatible capability while retaining several competing suppliers. That can make adoption widespread without giving any one company durable pricing power. Forcing describes the user’s alternatives; value capture describes who supplies the response and at what price. Let a company charge f per period, on top of the user’s real implementation cost c_i^{res} . The private cost is $c_i = c_i^{\text{res}} + f$, while company profit is $fx - K_v(x)$ for supply cost K_v . When the company lies inside the social accounting boundary, fees cancel between users and company; implementation and supply costs do not. The forcing inequalities place no lower bound on the company’s markup. A compulsory open standard supplied competitively can be indispensable and earn little rent.

Interoperability can lower implementation costs and increase the compatibility gain while making alternative suppliers more useful. Adoption of the capability may rise while dependence on one implementation falls. In the base model a better outside option caps the holdout’s loss. In a company-choice extension, a rival implementation may preserve network benefits too, which the segment-exit option Q does not represent. The appropriate company test is

$$U(A) - \max\{U(B), U(\text{multihome}), U(\text{exit})\}.$$

It can fall even while the gain from adopting *some* compatible technology rises. Adoption share, switching costs and supplier margins should therefore be measured separately. Neither widespread compliance nor a large forcing margin is an estimate of company profit.

H.5 Different stakes and repeated upgrades change the decision

Cost dispersion is only one source of asymmetry. An incumbent may have more customers and capital at risk, yet also face a larger conversion cost; an entrant may choose an unexposed niche. These effects pull in different directions, so incumbency alone does not order forcing. Actor-level stakes, exposures and outside options are needed. In a heterogeneous contest, multiplying each rival’s payoff by an unrelated stake need not preserve zero sum. A consistent extension must allocate a common prize or specify bilateral transfers whose gains and losses cancel.

Binary adoption also leaves out repeated investment after everyone has joined. For comparison, take $n \geq 2$ symmetric rivals choosing expenditure $e_i \geq 0$ for a fixed total prize $\mathcal{V} > 0$, with payoff

$$\pi_i = \mathcal{V} \frac{e_i}{\sum_j e_j} - e_i.$$

The symmetric interior equilibrium is $e^* = \mathcal{V}(n-1)/n^2$, so total expenditure is $\mathcal{V}(n-1)/n$. The first-order condition is $\mathcal{V}(n-1)e^*/(ne^*)^2 = 1$, and the payoff is strictly concave in own expenditure when rivals spend positively. Everyone participates, yet rivalry continues to consume resources. This standard contest benchmark (Tullock, 1980) shows why universal adoption does not establish that an arms race has ended. The binary model’s fixed r captures one discrete upgrade, not the equilibrium intensity of subsequent upgrades. Expenditure, quality and release frequency are therefore necessary complements to adoption dates.

I Evidence that separates pressure from value

An adoption curve shows that enough actors found switching worthwhile. It does not say whether they gained productivity, won business from holdouts or escaped a penalty. In the linear specification, the private gain depends on $\gamma + \alpha$. Moving value from the direct benefit γ to the contest edge α leaves the adoption incentive unchanged, while lowering aggregate value.

For example, a direct benefit of 0.3 and an edge of 0.2 can produce the same adoption path as a direct benefit of 0.1 and an edge of 0.4, with no outside option, $W = 0$ and the other settings held fixed. The second case creates 0.2 less value per adopter. A better fit to adoption alone cannot choose between them. Proposition 12 states the exact equivalence conditions.

Figure 15 shows three such pairs. In panel (a), moving 0.2 from the standalone gain γ to the edge α leaves adoption unchanged in every period, before the demonstration and after it. Welfare, however, is lower by $0.2 x_t$, so one technology leaves rivals roughly where they started and the other costs them 0.11. In panel (b), a standalone gain lower by 0.1 and a waiting value of 0.1 produce the same path, because both lower the switching gain by the same amount. Adopters realize γ , but no one realizes W , so welfare differs by $0.1 x_t$; and only γ enters the payoff gap, so the technology with the waiting value is forced for more rivals yet adopted no faster. In panel (c), a network gain and an exclusion penalty with $k = 0$ produce the same switching gain once the technology is demonstrated, and they end at the same 92%. Beliefs scale only the network gain, though, so before the demonstration exclusion drives adoption to 81% while the network good reaches 38%. Welfare differs in sign: the network good is worth 0.52 to rivals, while exclusion leaves them slightly worse off.

Diffusion histories therefore cannot classify a technology as wasteful or productive. The model indicates what can:

- *Adopter outcomes* as adoption spreads trace $U_A(x) + c$. Their slope separates network gains (rising) from contest erosion (falling).
- *Holdout outcomes* trace $U_N(x)$. Their shape identifies the contest's (α, r) through the zero-sum restriction that ties it to the adopter curve, and separates the exclusion penalty $\varepsilon\Phi$. A floor at which holdout outcomes stop falling, together with the exits, alliances and specializations that accompany it, identifies the outside option Q .
- *Belief variation* also helps, because beliefs scale the technology channels but not exclusion. Pre-demonstration adoption separates $\varepsilon\Phi(x)$ from $p\beta x$ even when the two coincide after the demonstration.
- *Bans and joint disarmament*, such as the 1971 cigarette television advertising ban or the 2010 swimsuit ban, inform y_R when their effects on costs and outcomes can be separated from changes in pricing, entry and other rules; Farr et al. (2001) show how an advertising ban raised cigarette makers' profits while leaving the social verdict to the smoking externality. Outsider outcomes inform η .
- *Shocks to rivals' adoption* that are excluded from the focal rival's payoff, such as rival-specific infrastructure access, regulatory eligibility or rollout timing, are needed to separate forcing from common shocks. This is the reflection problem (Manski, 1993; Brock and Durlauf, 2001).

Appendix J implements Bayesian estimation case by case for aggregate series, using the continuum dynamics implied by the hazard and an observation-error likelihood. Hierarchical pooling and joint estimation with actor-period panels remain extensions. They should be preceded by recovery experiments on panels simulated by `simulate_agents`, rather than inferred to work from aggregate fit alone.




Figure 15: Parameter pairs that adoption data cannot tell apart. Top: share adopted. Bottom: rivals' mean payoff when the technology works. (a) γ and α trade places ($\beta = 0.2$, $c = 0.6$, $\sigma = 0.3$, $p_0 = 0.5$, $x_v = 8\%$). (b) A standalone gain against a waiting value ($\alpha = 0.2$, $\beta = 0.2$, $c = 0.6$, $\sigma = 0.3$, the technology known to work). (c) A network gain against exclusion with a linear pressure curve ($\gamma = 0.3$, $c = 0.4$, $\sigma = 0.5$, $p_0 = 0.4$, exogenous demonstration at $t = 10$). In (a) and (b) the paths coincide exactly. In (c) the adoption incentives coincide only once beliefs are settled, and the paths converge after the demonstration; welfare remains different. In all panels $r = 1$, $Q = -\infty$, $\tau = 0$, and any unlisted payoff channel or waiting value is zero. Panels (a,b) use $x_0 = 2\%$, $\lambda = 0.3$, with $\lambda_0 = 0.1$ in (a); (c) uses $x_0 = 5\%$ and $\lambda_0 = \lambda = 0.4$.

I.1 What to measure before assigning a forcing label

The unit of analysis is a capability, an actor set and a time period. “AI” or “crypto” is too broad. For each case, record who competes with whom, what counts as productive adoption, which alternatives remain feasible, and whether the denominator is an original cohort, surviving actors, output or transactions. Define the network before looking at the adoption outcome. A regulator’s deadline is an event that may change exclusion, expected benefits or beliefs; it is not automatically evidence that the technology’s productivity has been demonstrated.

Payoffs should first be measured in the domain’s natural flow units, such as annual profit, operating cost per unit of output, or a stated utility scale. A logit hazard identifies utility relative to the shock scale τ : multiplying all payoff amounts, costs, waiting values and τ by the same positive constant leaves the hazard unchanged. Fix a scale or anchor it with outcome data. The normalized stake in Appendix J permits within-case interpretation; it does not make one unit of naval capability equal one unit of customer-service output.

Report several derived quantities rather than one forcing score:

- *Current switching margin:* $M_i(x, p) = \Delta(x; c_i, p) - W_i$, together with the gain from adopting alone. A positive margin can be ordinary standalone adoption.
- *Endpoint lock-in:* $F_i = \Delta(1; c_i, 1)$ against the best specified non-adopting response, or $\Delta_K(1)$ when counters are available. This is distinct from whether the path can reach full adoption.

- *Thresholds and their accessibility*: report every relevant crossing of $M_i = 0$ and, for a population cascade, the unstable crossings of $g_p(x) = x$. A selective contest can have a downward crossing; an outside option can produce a jump. Neither has the threshold model’s single upward tipping exposure.
- *Sensitivity and timing*: the slope of the switching gain at a continuous crossing, the size of a jump when applicable, and time from a defined demonstration to a specified adoption share. Report non-attainment and the observation horizon rather than dropping cases that never reach that share.
- *Welfare and uncertainty*: $W_R(x_T)$, $W_S(x_T)$ and their full-adoption counterparts, with posterior probabilities such as $\Pr(M_i > 0 \mid \text{data})$, $\Pr(F_i > 0 \mid \text{data})$ and $\Pr(y_S < 0 \mid \text{data})$.

Margins and slopes require a common economic normalization for cross-case comparison. Exposure thresholds are dimensionless but still depend on the actor set and network. Calendar times also reflect construction times and adjustment frictions. None of these statistics by itself identifies the strength of competitive causation.

The order of evidence. Start with a record for each actor in each period: orders and use, real costs, output or financial outcomes, relevant rivals’ adoption, counters, alliances and exits. This is the actor-period panel. Estimate an adoption hazard with actor and time controls, then seek a shock that shifts rivals’ adoption without directly shifting the focal actor’s payoff. Lagged exposure alone does not solve common shocks or anticipation. Fit the structural channels only after specifying that causal design. Separate observation error from behavioral shocks; use cross-sectional cost data to distinguish persistent dispersion from slow adjustment.

Validation should test recovery on simulated panels, predictions for later periods or held-out actors, alternative network weights and actor sets, and sensitivity to cost distributions and prior bounds. A hierarchical model may pool diffusion parameters among comparable cases, but should not manufacture precision by pooling unrelated utility scales. Fit linear, threshold and flexible pressure specifications on the same observations. A low in-sample error from a flexible curve is evidence of descriptive fit, not evidence that competitive exposure caused adoption. The case studies below specify outcomes that could reject their proposed channel, so a prospective application can fail rather than merely be relabeled after the fact.

J What the historical adoption curves tell us

The central finding from 14 historical cases is that fitting adoption is much easier than explaining where its value comes from. The fuller model often traces almost the same path as the simpler one, and the split between productive value and competitive transfers remains sensitive to prior assumptions.

The data were rebuilt from traceable sources. The earlier series had 108 observations, including 30 estimates, and used only the linear restriction; Appendix Q retains those reproduction benchmarks. The findings below use the updated series and full-model fits.

J.1 Traceable observations

Each case has a directory with a source registry, the observations, and the snapshots they were read from. Every source is recorded with its URL, its access date, its licence, and the SHA-256 hash of the snapshot. Wikipedia pages are fetched at a pinned revision, so the snapshot never changes.

Every value is a row that cites one source, and there are two kinds:

- A value typed from a page carries a verbatim quote, and a verification step checks that the quote appears in the snapshot’s text.
- A value parsed by code is re-derived from the snapshot by the case’s collector. The machine-readable sources are the Firefox telemetry files behind the Let’s Encrypt statistics, the Census Bureau’s Business Trends and Outlook Survey spreadsheets, the IEA electric-vehicle data, the USDA *Agricultural Statistics* table, Wikipedia infoboxes, and one chart digitized from its image.

Count-based cases (dreadnoughts, frequent-flyer programs, nuclear weapons, armed drones) are recorded actor by actor, including a disarmament event in the nuclear series. Their series is the share of a fixed field currently holding the technology, not the share among survivors.

The rebuilt data use 80 source records. They hold 295 rows: 183 parsed by code, 111 transcribed with verified quotes, and one excluded note row (Table 33). The fitted series have 350 points in all, with 75 dated adoption or disarmament events in the four count cases sampled on an annual or monthly grid. Only three rows are estimates, down from 30, and 58 rows are kept but excluded because they use a different definition. Openly licensed snapshots are committed with the code. Copyrighted pages, and the one oversized scan, are kept locally and can be re-fetched; their hashes and supporting quotations are recorded.

Several inherited series changed materially:

- *Hybrid corn* is now the complete official series, 28 annual values from the 1962 *Agricultural Statistics*. It replaces seven estimates, two of them off by five points.
- *Frequent-flyer programs*: all eight carriers are now dated from sources, and seven of the eight had a program within four months of American’s, not within twenty.
- *Containers*, *swimsuits* and *armed drones* each now use a single definition in place of a composite.
- *Nuclear weapons* now follows one source, Bleek’s coding of acquisition, for both the dates and the field.

J.2 How to read the estimates

The fits describe the recorded adoption paths under a specified model. They do not estimate causal losses imposed on rivals. Some parameters were restricted using prior judgments, and the uncertainty intervals are conditional on those restrictions and the observation-error assumptions.

The model sets the gross gain from staying adopted at full adoption to one, so a mean-cost actor is forced when $c < 1$. It allows a nonlinear contest, temporary shocks and, for HTTPS and RFID, exclusion. It does not fit outside options, counters, changing fundamentals or reversals. Repeated observations of the same adoption event are not independent experiments.

The numerical settings, prior ranges and likelihood are in Appendix K. The reported ΔAIC values are AIC-style fit diagnostics evaluated at posterior modes, not formal model-selection tests. Read them as a check on whether the extra flexibility materially changes the fitted curve.

J.3 Historical customer value across the lifecycle

Section 1 defines the customer budget $B(x, p)$, the ceiling on what a company can charge. The figures below apply it to the historical parameter posteriors. They say how the ceiling moved as each market filled; they say nothing about prices, shares or margins, which the fits do not observe.

Figure 16 follows that customer value from low adoption to a mature market. It holds the current parameter posterior fixed and assumes effectiveness has been demonstrated, $p = 1$. The

orange dot is 5% adoption and the blue square is 95%; the unadopted share falls from 95% to 5% between them. When the curve falls, early customers have an advantage over unprepared rivals that later customers cannot obtain. When it rises, network gains or exclusion make the product more valuable as adoption spreads. The balance, at a fixed belief, is

$$\frac{\partial B}{\partial x} = p \left[\beta - \frac{\alpha(r-1)}{(1+(r-1)x)^2} \right] + \varepsilon \Phi'(x).$$

For example, the container fit’s median customer benefit rises from about 0.49 at 5% adoption to 0.97 at 95%; the dreadnought fit falls from about 1.94 to 1.02. The curves and bands meet at one because the normalization fixes that endpoint. Before proof, use $p < 1$ in (2); the chart’s proven-effectiveness assumption is not an estimate of an early investor’s confidence. The fitted belief model only permits a jump to successful demonstration, so it also misses the risk of a disappointing trial.

J.4 Market maturation versus learning about the market

Two things can change between an early investment and a later review. More customers adopt, moving the opportunity along a lifecycle curve. New evidence can also change the estimated curve itself. The model holds its structural parameters fixed within a case; a revised posterior is a change in our estimate, not evidence that the underlying parameter physically changed.

Figure 17 isolates the second effect. It compares a fit using only the early adoption observations with a fit using the latest recorded history. Both estimates evaluate $B(0.1, 1)$: the same 10% adoption and proven effectiveness. Any movement therefore comes from the parameter estimate, rather than from moving further through the market. The numerical companion also records each parameter’s posterior and customer value at the actual adoption share and fitted belief in each window.

The early window ends at the last recorded point before adoption first exceeds 20%. This defines low penetration, not time since invention or a particular financing round. Nine cases have distinct early and later windows. AI use among all US firms and the RFID mandate remain below that threshold, so their early and latest windows coincide. Gen AI in large organizations, HTTPS and LZR suits have no observations below the threshold; the chart leaves their early estimate absent. Some early windows have only one observation and mostly reflect the priors.

These are retrospective refits with today’s priors held fixed. Observation dates are not verified publication dates, and the cases were selected with hindsight. They show how restricting the adoption history changes this model’s estimates; they do not reconstruct the information available to an investor on that date. A genuine investment backtest would also need archived data vintages, contemporaneous priors, failed technologies, and supplier prices, costs and valuations. Appendix K records the windowing and sampling method.


J.5 The full-adoption endpoint

Figure 2 in Section 1.2 presents the historical endpoint estimates. Payoffs use the within-case normalization $a + \beta + \gamma + \varepsilon = 1$, so a mean-cost actor wants to stay adopted at full adoption when $\Delta(1) = 1 - c > 0$. The rival group’s payoff relative to no adoption is $y_R = \gamma + \beta - c$. Nuclear weapons and RFID lie left of zero, where full adoption would not sustain itself for a mean-cost actor.

The plotted dots pair the marginal posterior medians, and the bars show separate 5th–95th percentile intervals, not joint 90% regions. Distinguishing dominant from tipping adoption also requires $\Delta(0)$. For a supplier investor, the endpoint describes customers’ mature-market incentives and the rival group’s payoff; the lifecycle and value-capture questions above remain necessary.

Customer value from early adoption to a mature market

Gross customer benefit, $B(x, 1)$; log scale; full-adoption value = 1



Current full-history posterior; effectiveness assumed proven ($p = 1$).


Benefits are before all adoption costs. These are lifecycle scenarios, not historical investor beliefs.

Figure 16: Estimated gross adoption benefit at different adoption shares. Each curve is the gross benefit of adopting rather than holding out, before all adoption costs, under the full-history parameter posterior and $p = 1$. Dots mark 5% adoption and squares 95%; bands are pointwise 90% posterior intervals. The vertical axis is logarithmic. Each case is normalized to $B(1, 1) = 1$, so levels compare value with that case's mature-market benchmark, not dollar market sizes. These are conditional lifecycle scenarios, including stages not observed in a case's data; they are not historical investor beliefs, price forecasts or estimated supplier returns.

How later evidence revises the early customer-value estimate

Same 10% adoption and proven effectiveness in both estimates; 90% intervals

—●— Early observation window (up to 20% adoption) —■— Latest observation window



Retrospective window refits with today's priors held fixed; not a point-in-time investment backtest.

Dates are observation dates; n counts fitted points, not independent experiments. Full-adoption benefit = 1.

Figure 17: How adding later observations revises the estimate of early customer value. Circles use the early observation window; squares use the latest window. Both evaluate the same gross benefit at 10% adoption and $p = 1$. Lines span separate 90% posterior intervals, not an interval for the difference between estimates. Missing early coverage is labeled; a single square also marks cases whose latest observation is still in the early window. The same present-day priors and observation model are used in both fits. These retrospective comparisons isolate information-window changes under those assumptions, not investable historical returns.

The estimates come from the stored derived summaries in `data/cases/*/fit.json`; exact plotted values are exported to `paper/tables/calibration_scatter.csv`.

J.6 Three findings to use with care

Figures 18 and 19 and Table 22 support three practical conclusions.

A more elaborate model can fit almost the same curve. Adding contest curvature and exclusion changes the fitted paths little. HTTPS shows the largest improvement: its error falls from 0.015 to 0.011, but not enough to offset the extra-parameter penalty in the AIC-style diagnostic. This says the adoption series offer little reason to prefer the added flexibility. It does not show that those mechanisms are absent.

The data say little about whether the value is productive or positional. In every case, the uncertainty about the positional share changes little from the prior. Most series also say little about how decisive the contest is. A good adoption fit therefore leaves the welfare classification dependent on assumptions supplied before fitting. The detailed comparisons and their dependence caveats are in Appendix K.

Even a high probability of forcing needs its assumptions beside it. Twelve cases give at least 95% probability that a mean-cost actor would keep the technology at full adoption. Several prior cost ranges already implied that result, so these are not twelve independent confirmations. Nuclear weapons and the RFID mandate are the exceptions; within the fitted model, their partial adoption is explained by costs. Outside options or changing enforcement could offer other explanations. HTTPS also splits between dominant and tipping patterns, at 58% and 42%, rather than receiving a certain label.

Use these fits to describe paths and locate unanswered questions. Before relying on a welfare classification, vary the prior ranges, cost distribution and treatment of repeated observations, and add outcomes for adopters and holdouts. Reproducing the stored calculations verifies the arithmetic; it does not establish causal validity or predictive performance on new cases.




Figure 18: Full-model calibration on the rebuilt series. Bands are conditional 90% posterior intervals of the adoption path, the solid line is the posterior median, and the dashed line is the linear model at its posterior mode. Each panel's ΔAIC is the AIC-style diagnostic described in the text, comparing the full model with the linear model plus the transitory shock. Crosses mark alternative definitions excluded from the fit. These curves summarize posterior draws rather than one parameter set: fitted settings and priors are specified in Appendix K, and each case's complete fitted summaries are stored in `data/cases/<case>/fit.json`.

Table 22: Full-model calibration. Panel A: n is the number of fitted series points; RMSE is evaluated at the posterior modes. ΔAIC compares the linear model plus the transitory shock with the linear model (τ), and the full model with the linear model plus the shock (r, ε); negative favours the larger model. Panel B: posterior medians and 90% intervals, with the most probable region of Table 1 and its probability. A dash denotes exclusion fixed at zero. Prior-bound diagnostics remain in the stored fit summaries; source counts are in Table 33.

A. Descriptive fit					
Case	n	RMSE		ΔAIC	
		Linear	Full	τ	r, ε
Clap skates	6	0.007	0.006	+1.8	+2.1
Container	15	0.032	0.031	+1.9	+1.9
Hybrid corn	28	0.017	0.017	+1.8	+2.0
Dreadnought	18	0.049	0.049	+2.0	+2.0
Armed drones	64	0.021	0.021	+2.0	+2.0
Frequent flyer	32	0.059	0.059	+1.8	+2.0
Gen AI (large orgs)	5	0.031	0.031	+2.0	+2.0
AI (all US firms)	54	0.004	0.004	+2.0	+2.0
HTTPS	14	0.015	0.011	+2.8	+6.1
LZR suit	3	0.041	0.041	+2.0	+2.0
MLB analytics	4	0.017	0.017	+2.0	+2.0
China NEVs	11	0.009	0.009	+2.3	+2.0
Nuclear	91	0.031	0.031	+2.0	+1.8
RFID mandate	5	0.000	0.000	+2.0	+8.0

B. Posterior summaries					
Case	r	c	ε	Modal region (P)	
Clap skates	1.30 [1.02, 2.40]	0.25 [0.15, 0.30]	–	arms race (0.97)	
Container	1.50 [1.04, 3.51]	0.66 [0.56, 0.70]	–	coordination (0.93)	
Hybrid corn	1.49 [1.03, 3.34]	0.35 [0.22, 0.47]	–	productive (1.00)	
Dreadnought	1.99 [1.09, 7.33]	0.30 [0.16, 0.48]	–	arms race (0.97)	
Armed drones	2.25 [1.08, 8.52]	0.18 [0.11, 0.36]	–	productive (0.87)	
Frequent flyer	1.86 [1.10, 6.76]	0.20 [0.09, 0.29]	–	productive (0.51)	
Gen AI (large orgs)	1.60 [1.05, 5.08]	0.43 [0.21, 0.50]	–	productive (0.97)	
AI (all US firms)	1.60 [1.04, 4.75]	0.80 [0.56, 0.98]	–	arms race (0.55)	
HTTPS	1.50 [1.02, 3.70]	0.79 [0.66, 0.89]	0.28 [0.03, 0.38]	arms race (0.57)	
LZR suit	1.88 [1.04, 6.82]	0.03 [0.01, 0.05]	–	productive (0.50)	
MLB analytics	1.88 [1.06, 7.82]	0.23 [0.11, 0.34]	–	arms race (0.78)	
China NEVs	1.39 [1.02, 2.80]	0.49 [0.33, 0.59]	–	productive (0.71)	
Nuclear	1.13 [1.01, 1.55]	1.19 [1.03, 1.29]	–	inert (0.53)	
RFID mandate	1.57 [1.04, 3.05]	1.64 [1.13, 1.96]	0.21 [0.02, 0.38]	inert (1.00)	

K Calibration assumptions and estimation settings

The fit normalizes the gross gain from staying adopted at full adoption to one: $a + \beta + \gamma + \varepsilon = 1$ with $a = \alpha/r$. With no outside option, the mean-cost rival’s forcing test is then $c < 1$ for every r . The linear payoff restriction also sets exclusion to zero.

The free parameters and their priors are:

- the positional share a , uniform on the inherited range, with β and γ splitting the remainder in the inherited ratio;
- the contest’s decisiveness r , half-normal on $\log r$ with scale one, truncated at $\log r = 3$;
- the cost level c and dispersion σ ;
- the prior belief p_0 , the visibility threshold x_v , the two adjustment speeds λ_0 and λ , and the seed x_0 ;
- the transitory shock τ , half-normal with scale 0.05.

The exclusion share ε and its shape (θ, k) are free only for the two cases whose history includes a gatekeeper penalty, HTTPS and the RFID mandate; elsewhere ε is fixed at zero. The waiting value W is fixed at zero rather than separately inferred from benefits using adoption shares alone (Proposition 12). The outside option is switched off, because no series records a sheltered share. Counters, learning from unfavorable evidence, endogenous exit and changing technological fundamentals are also outside these fits.

Uniform ranges for the diffusion parameters are the inherited ones, except where the rebuilt data moved them. Each change is recorded, with its reason, in the case’s metadata file. The main ones are a new seed where the series now starts at zero, a rescaled visibility threshold where the field changed, and nuclear effectiveness treated as known from the start in 1945 rather than inferred from a visibility threshold.

Each observation is normal around a numerical approximation to the continuum path, using 400 cost bins and linear interpolation at fractional periods. Its standard deviation is

$$s(x) = \sqrt{0.01^2 + (0.05x)^2}$$

for stated and derived values, with x evaluated at the observed share, and 0.05 for estimates. For actor counts the standard deviation is half an actor; for the two series that stay below 10% it is set to match their precision. These are specified error scales, not survey-based sampling errors.

The posterior is sampled with the ensemble sampler of the `emcee` package, run for 10,000 steps with at least 40 walkers started at the posterior mode. Two restrictions are fitted at their own modes for comparison: the linear model ($r = 1$, $\varepsilon = \tau = 0$), and the linear model with the transitory shock.

The displayed ΔAIC values are AIC-style diagnostics, computed as $2\Delta k - 2\Delta \log L$ at the posterior modes. Because these are not unrestricted maximum-likelihood estimates, the values should not be interpreted as formal AIC model selection. The likelihood also treats sampled series points as conditionally independent. In count-based cases, many points repeat information from the same adoption events; the 350 points are not 350 independent experiments. Serial dependence, changing survey definitions and the nuclear reversal can make uncertainty larger than these intervals suggest. The irreversible model cannot reproduce a disarmament decline exactly.

K.1 Lifecycle curves and retrospective observation windows

The investor figures use separate joint posterior draws stored under `data/lifecycle/`; the original `data/cases/*/fit.json` summaries remain unchanged. The lifecycle curves evaluate $B(x, 1)$ over

$x = 0, 0.01, \dots, 1$ for each draw from a new full-history sample. Quantiles are taken after that transformation. Using joint draws preserves dependence among parameters when calculating the uncertainty bands.

The early observation window includes every fitted row through the last timestamp before the first observed share above 20%. A later decline does not reopen the early window. If the first observation is already above 20%, no early fit is produced; if the threshold is never crossed, the early and latest windows are identical. The likelihood uses only the retained rows and their original error scales, and its simulation horizon is shortened to cover that window. Automated checks verify that changing later observations does not change the early-window likelihood. This prevents later observations from entering the early likelihood but does not remove hindsight from the case priors, data revisions, definitions or case selection.

Each distinct window uses the same case priors, payoff normalization, likelihood and 400 cost bins as the main calibration. An early-window posterior mode is fitted using only that window; the latest-window sampler starts from the stored full-history mode. Each ensemble runs for 10,000 steps with at least 40 walkers, discards the first third and thins by five. Both the initial walker distribution and the sampler’s random state have explicit seeds. A reproducible selection of 4,000 joint draws is saved for the transformations and figures. Acceptance rates, autocorrelation estimates, approximate effective sample sizes, input hashes and code hashes accompany each fit. These finite-run diagnostics do not establish convergence across modes; the comparisons remain exploratory, conditional on the same prior and dependence limitations as the main calibration.

For the observation-window comparison, both posteriors are evaluated at $x = 0.1, p = 1$, which separates an estimate revision from a change in lifecycle position. The accompanying state table instead evaluates the gross benefit, net adoption incentive and adopter payoff at the last *observed* adoption share in each window, using that draw’s fitted belief at the corresponding model period. This is a model-based reconstruction, not an observed investor belief. The state table does not turn the binary demonstration rule into a general learning model.

The scripts `scripts/fit_lifecycle.py` and `scripts/calibration_assets.py` regenerate the window fits, figures and numerical companions. The curve table records the pointwise bands; the window table records dates, sample counts, observed adoption and state estimates; the parameter table records the joint draws’ marginal summaries. These are `paper/tables/lifecycle_curves.csv`, `lifecycle_windows.csv` and `lifecycle_parameters.csv`, respectively. None contains observed supplier prices, margins or investment returns.

K.2 Detailed posterior comparisons

Figure 18 and Table 22 report the fits. Three findings stand out.

The additional channels make little difference to descriptive fit. The full model’s paths are nearly indistinguishable from the linear model’s (Figure 18). Against the linear model with the transitory shock, the contest and exclusion channels lower the AIC-style diagnostic in no case. Their ΔAIC is about +2 almost everywhere. HTTPS and the RFID mandate, the two cases with a free exclusion channel, are higher still (+6.1 and +8.0). HTTPS gains the most: the additional channels cut its RMSE from 0.015 to 0.011, with an exclusion share of 0.28 [0.03, 0.38], but not enough to offset the four-parameter penalty in this diagnostic. Adding the transitory shock alone also improves fit little. These comparisons do not establish that the omitted mechanisms are absent.

Most series weakly constrain the contest’s decisiveness. Figure 19 compares the posterior and prior. For seven cases the ratio of standard deviations of $\log r$ is between 0.89 and 1.06: dreadnoughts, drones, frequent-flyer programs, swimsuits, MLB analytics and both gen-AI series. Six cases have ratios of 0.46 to 0.70: clap skates, the RFID mandate, Chinese electric vehicles, hybrid corn, containers and HTTPS. All six pull r toward one. Nuclear weapons, with 91 sampled points




Figure 19: Posterior of the contest's decisiveness r by case against its prior (gray band, 90%). The right column is the ratio of posterior to prior standard deviation of $\log r$; near one, the posterior has contracted little. The intervals are conditional on the prior and observation model.

over 81 years, has the narrowest interval, $[1.01, 1.55]$, and a ratio of 0.24; the dependence caveat is particularly relevant there. For the positional share a , the ratios remain between 0.96 and 1.02 in every case. Thus the split that determines waste remains strongly dependent on its prior. This is an empirical weak-identification result, not a proof that nonlinear contest parameters can never be identified from a sufficiently informative path.

Endpoint forcing is more concentrated than the channel decomposition, conditional on the priors. Twelve cases assign at least 95% posterior probability to $\Delta(1) = 1 - c > 0$ for the mean-cost rival. This is not twelve independent confirmations of forcing: several prior cost ranges already lie wholly below one. The closest call is AI use among all US firms, whose 5th percentile of $\Delta(1)$ is $+0.02$. Nuclear weapons has median $c = 1.19$, and the RFID mandate has $c = 1.64$. Under the fitted restrictions, cost heterogeneity accounts for their partial adoption; alliance sheltering or changing enforcement could provide other explanations. The waste sign $y = \gamma + \beta - c$ inherits the weakly constrained split a , so the region labels in Table 22 are conditional on prior judgments. The adoption pattern is sensitive too: with weakly constrained r and a prior restricted to $r \geq 1$, $\Delta(0) = \gamma + ar - c$ can be larger than in the linear restriction. HTTPS splits between dominant and tipping patterns (58% and 42%).

The calibration illustrates the value of the outcome-data strategy in Appendix I. Adoption histories constrain some combinations of relative costs and diffusion parameters while leaving welfare classifications sensitive to assumptions. Parameters whose 90% interval approaches a prior bound are recorded with the calibration; they include beliefs, adjustment rates, visibility thresholds and the container cost level. Sensitivity to these bounds, the prior on r , the cost distribution and time-series dependence remains necessary before treating the reported classifications as robust empirical findings. The stored sampler diagnostics and replay checks establish reproducibility of the reported calculations, not convergence under every alternative specification or out-of-sample validity.

L Parameter reference

Section B.7 gives the visual walkthrough: how each main payoff parameter moves the two curves, what it does to adoption, and whether it adds collective value. This appendix supplies the distributional details and a compact reference. For numerical illustrations, the figure panels state changed parameters and their captions state the settings held fixed. The plotted values and complete payoff settings for Figures 1 and 6 are also exported to `paper/tables/payoff_examples.csv`.

The main text uses c for the mean-cost rival. Individual costs are represented by the quantile function

$$c(q) = c \kappa(\sigma) \left(\frac{q}{1-q} \right)^\sigma, \quad \kappa(\sigma) = \frac{\sin(\sigma\pi)}{\sigma\pi},$$

with the continuous limit $\kappa(0) = 1$. This log-logistic distribution has mean c and dispersion $0 \leq \sigma < 1$. The cost of assigning adoption to the cheapest share x is $C(x) = \int_0^x c(q) dq = c I_x(1 + \sigma, 1 - \sigma)$, where I_x is the regularized incomplete beta function. On an irreversible path, use the costs of the actual adopters instead.

The exclusion curve is a logistic centered at θ with steepness k , normalized to $\Phi(0) = 0$ and $\Phi(1) = 1$. Its $k = 0$ limit is $\Phi(x) = x$. The table is a reference for changing one parameter while holding the others fixed; real technologies can change several features at once.

Table 23: Parameter guide. Panel A gives effects on the three summary numbers in (9) for the mean-cost rival, holding other parameters fixed. Except for the Q and p rows, use $p = 1$ and no outside option. The r and p signs assume nonnegative technology benefits; negative positional effects can reverse them. Panel B describes the adoption path, not the payoffs conditional on known effectiveness (Appendix E).

	Meaning	Example	$\Delta(0)$	$\Delta(1)$	Shape of $\Delta(x)$	y_R
<i>A. Payoff parameters</i>						
γ	standalone gain to an adopter	containers' lower handling cost	+	+	level	+
β	network gain from other adopters	more container-ready ports	0	+	rises	+
α	positional edge of a lone adopter	one dreadnought against older fleets	+	+	falls if $r > 1$	0
r	erosion of the edge, $\alpha > 0$ fixed	the same initial advantage wears away sooner	0	−	lower for $x > 0$	0
ε	penalty on holdouts that no rival captures	browsers flagging non-HTTPS sites	0	+	rises near θ	0^a
θ, k	where and how abruptly exclusion bites	a warning introduced at low rather than high market adoption	0	0	moves and sharpens the rise	0^a
c	mean adoption cost	building a new fleet; retooling a port	−	−	level	−
σ	dispersion of adoption costs	incumbents with installed fleets against entrants	0	0	spreads rivals around it	0
Q	value of the outside option (≤ 0): ally, specialize, exit the segment	slower traders leaving the instruments that latency arbitrage contests	0^c	$-^c$	caps the rise; the edge vanishes at x_Q	0
η	gain to outsiders per adopter	shippers' freight rates; users' security	0	0	none	0^b
p	belief that the technology works	an untested battleship design	+	+	scales γ, β, α	0

Continued on the next page

Table 23: Parameter guide (continued)

	Meaning	Example	$\Delta(0)$	$\Delta(1)$	Shape of $\Delta(x)$	y_R
<i>B. Diffusion parameters</i>			Main effect on the adoption path			
p_0	belief before the demonstration	confidence in an untested design	sets the ceiling x_L ; dormancy if $x_L \leq x_v$			
x_v	usage that demonstrates the technology	first use in combat; a public benchmark	dormancy if above x_L ; a later date near it			
x_0	seed, held by the cheapest rivals	a sponsor's first orders	an earlier demonstration; immediate if $x_0 \geq x_v$			
λ_0	adjustment speed before the demonstration	trials and pilot purchases	when the demonstration comes, not whether			
λ	adjustment speed after it	build times; budget cycles	speed of take-off; also the final level if the edge erodes			
τ	scale of transitory shocks	budget windows; changes of leadership	adoption creeps past the rest point			
W_0, ω, κ_W	waiting value: a constant, its rise with belief uncertainty, its erosion by the edge	waiting for a rival's trial to fail	delays adoption and silences it before the demonstration; never enters payoffs or the classification			
<i>C. Counter parameters (Appendix F)</i>			Main effect			
φ	share of a counterer's stake shielded from adopters' edge	mines against battleships; ad-blocking	lowers the gain from adopting; disarms if $c_K \leq \varphi\alpha/r - \Delta(1)$			
c_K, σ_K	mean and dispersion of counter costs	building a mine flotilla	whether holdouts counter, and whether the counter pays for itself			

^aExclusion costs holdouts $(1-x)\varepsilon\Phi(x)$ during the transition, though not at either end. ^bRaises y_S one for one.

^cFor $Q < 0$, $\Delta(0)$ is unchanged; at $Q = 0$ sheltering is immediate and removes the initial contest return. Raising Q lowers $\Delta(1)$ only when the outside option binds. At fixed V rather than fixed α , increasing r raises both endpoint gains because $\alpha = V(r-1)$; see Section B.7.

M Deployment conditions across fourteen markets

Table 24 compares deployment conditions across fourteen markets. It asks why adoption pays, where deployment stands and what constrains it. Those questions establish an adoption opportunity; evidence about holdout margins, retention or access is needed to establish economic forcing. Productive tools such as robots and software assistance can develop that pressure, while a network can attract adopters without yet damaging holdouts. Section 6 and Appendix N assess investment priorities across sixteen opportunities. The sources behind each readiness row are cited in the third column.

Scaling means established deployment in the named segment. *Selective* means expansion depends on local economics or workflow. *Phased* marks an active regulatory or technical transition. *Partners needed* means several parties still need to align.

Table 24: Fourteen markets compared on the buyer’s question, with sources as of October 2026. Deployment status describes maturity within the stated population; it says nothing about company profitability.

Market / decision	Why adoption can pay	Where the market is today	Scale verdict	Main constraint
Support assistance	Cost and service: less work per resolved case; rivals can raise service expectations.	2026 company survey: 66% report AI agents in their organization (Salesforce, 2026); a field study found 15% more issues resolved per hour (Brynjolfsson et al., 2025).	Scaling in repeatable queues.	Prove lower cost per durable resolution after escalation and repeat contacts.
Software assistance	Productivity and speed: useful output survives imitation; release advantages may fade.	Commercial tools; METR’s February 2026 update cannot reliably quantify current gains (Becker et al., 2026).	Selective by workflow.	Review and rework can absorb gains; measure delivery through deployment.
Clinical note drafting	Time savings: less clerical work remains useful even when all clinics adopt.	2026 study of 8,581 clinicians: 13.4 minutes less record work per eight scheduled patient hours; observational (Rotenstein et al., 2026).	Selective with clinician review.	Corrections, integration and a buyer able to retain the benefit of time saved.
Warehouse robots	Cost competition: throughput and unit-cost gains survive imitation.	Amazon reported its millionth robot in June 2025 (Dresser, 2025); five million industrial robots in service globally (International Federation of Robotics, 2026b).	Scaling in structured facilities.	Utilization, site conversion, integration and maintenance costs.
Precision spraying	Input savings: less herbicide at comparable weed control; rivals can lower costs.	Deere reports over five million acres treated in 2025 (John Deere, 2025).	Scaling in suitable crops.	Acreage, weed control, equipment access and local payback.
Instant bank payments	Network reach: faster funds plus access to relevant counterparties.	RTP processed 2.27 million payments on May 1, 2026 (The Clearing House, 2026); FedNow settled about 5.0 million transfers in 2026 Q2 (Federal Reserve Financial Services, 2026).	Scaling on established rails.	Sending access, fraud controls, liquidity and 24-hour operations.
Stablecoin settlement	Route savings: faster settlement pays if full costs beat available bank routes.	Visa: over \$20 billion annualized settlement rate, September 2026 (Visa, 2026).	Selective by corridor.	Redemption, conversion costs and local liquidity; compare full routes.

Continued on next page

Table 24, continued

Market / decision	Why adoption can pay	Where the market is today	Scale verdict	Main constraint
E-invoicing (France)	Coordination and mandate: shared formats save work; phased rules add pressure.	Receiving phase began September 2026; issuance phases continue in 2027 (Ministère de l'Économie, n.d., 2026).	Phased rollout.	Partner activation, data quality and exceptions; no sanctions in 2026.
Electronic bills of lading	Coordination: faster transfer and finance pay when the whole shipment chain accepts them.	DCSA: about 11% electronic issuance in 2025; interoperability advanced in 2026 (Digital Container Shipping Association, n.d.; Guy, 2026).	Partners needed across the chain.	Banks, shippers and carriers must align legal acceptance and operating procedures.
Passkeys	Defense: lower login friction and fraud can benefit every adopter.	FIDO participants: 36% of successful sign-ins in its October 2026 index (FIDO Alliance, 2026).	Scaling in online services.	Enrollment, recovery and weak fallbacks can limit realized protection.
Post-quantum migration	Future protection and compatibility: rivals do not create the quantum threat.	Cloudflare: over two-thirds of human TLS traffic protected, April 2026 (Goldberg and Paul, 2026); NCSC targets 2028/2031/2035 (National Cyber Security Centre, 2025).	Phased beyond supported connections.	Inventory and legacy systems; protected web traffic does not cover every asset.
Insurance telematics	Selection: informed insurers can attract safer risks from less informed rivals.	Commercial programs such as Snapshot; availability varies by state (Progressive, n.d.; National Association of Insurance Commissioners, 2026).	Selective within an established market.	Consent, data access, permitted pricing and policy renewals.
Procurement negotiation	Bargaining and labor: admin savings persist; price gains may shrink as suppliers respond.	Pactum reports over 50 enterprise customers (Pactum, n.d.; Uprety, 2026).	Selective in bounded categories.	Supplier response and retention; test gains after both sides automate.
Grid storage / flexible demand	Energy shifting: useful flexibility persists; competing assets can erode revenue.	US utility batteries: nearly 52 GW in June 2026 (U.S. Energy Information Administration, 2026); aggregation access varies locally (Electric Reliability Council of Texas, 2026).	Scaling storage; selective aggregation.	Contracts, market access and interconnection; stress-test falling price spreads.

Several additional candidates require separate consideration. Agent commerce and content provenance are candidates whose forcing mechanism still lacks a link: a transaction lost without an accepted credential, or an enforced platform rule with a measurable holdout effect. C2PA's

2026 specification work supplies infrastructure rather than that evidence (Coalition for Content Provenance and Authenticity, 2026, n.d.). Further evidence of lost transactions or enforced access requirements is needed to establish a forcing mechanism in agent commerce or content provenance. Polygenic embryo screening has a positional story and no validated effect to buy (Turley et al., 2021). Proof-of-work mining is an established contest benchmark; the main comparison separates miners' returns from security benefits and equipment suppliers' receipts (Section 4).

N Investment directions: parameters, evidence and capital

This appendix sets out the evidence, assumptions and capital requirements for sixteen investment opportunities. Parameters are constructed annual scenarios, not fitted estimates. The four sectors are crypto, AI, fintech and greenfield (other markets). The opportunities include companies, networks, physical assets and financing arrangements; the appropriate investment instrument depends on the rights and returns available.

N.1 The plotted coordinates

For each representative actor, $c = q + \kappa$ and

$$F = \gamma + \beta + \alpha/r + \varepsilon - c, \quad y_R = \gamma + \beta - c, \quad y_R^{ext} = y_R - h(1).$$

The current decision is

$$M = p \left[\gamma + \beta x + \frac{\alpha}{1 + (r - 1)x} \right] + \varepsilon x - c - W.$$

Here $\Phi(x) = x$ is a screening assumption. p, x, W are subjective benefit haircuts, exposure and waiting values; they are not estimated success probabilities, measured adoption or actual calendar deadlines. A positive endpoint does not imply a positive current purchase decision.

Figure 20 plots the scenarios on the base model's endpoint grid. Figure 21 applies the common producer revenue loss. All scenarios for a direction use its original plan central $B_1 = \gamma + \beta + \alpha/r + \varepsilon$ as a fixed normalizer. This preserves comparability across revisions within that direction; it does not compare dollar market size across directions. Joined downside/central/upside points are joint scenarios, not an adoption path or a confidence interval.

The central cases for C1, A1, F2 and G1 make conservative assumptions where evidence is missing. C1 includes operational savings and fees but no network premium; A1 includes administrative savings but no bargaining advantage; F2 assumes net fraud prevention without an additional network benefit; and G1 counts capacity gains only when they generate revenue or reduce costs. The positive upside for A1 still illustrates the hypothesis that an experiment could establish; it is not supported merely by its presence on the figure.




Figure 20: Sixteen directions on the original endpoint grid. Upper right: adoption pays and the rival group gains; lower right: adoption pays and the group loses; lower left: full adoption is not sustained. Upper left is infeasible with these nonnegative channels. Points are constructed scenarios; the source tables provide exact values and direction names.




Figure 21: The same working cases after assumed common price pass-through. Hollow dots are before h ; filled dots subtract $h(1)$ vertically. Adoption incentives do not change. A2, C1 and G1 can have positive adoption incentives while producer profits decline. This is not a finding of social waste or a measured price equilibrium.

Table 25: Working-case annual parameters and endpoints, in \$000 per representative actor. All cases are unvalidated scenarios. Public evidence constrains the assumptions for C1, A1, F2 and G1; M includes the stated discounts to expected benefits and the value of waiting.

ID	γ	β	α	r	ε	c	$h(1)$	F	$y_R - h(1)$	M
C1	400	0	0	1	0	240	320	160	-160	0
C2	100	0	600	3	0	200	0	100	-100	90
C3	800	500	0	1	0	700	200	600	400	-238
C4	40	30	200	4	0	150	0	-30	-80	-92
A1	100	0	0	2	0	200	0	-100	-100	-220
A2	2,400	0	0	1	0	1,000	1,920	1,400	-520	380
A3	800	0	0	1	0	600	0	200	200	-350
A4	50	300	0	1	0	200	100	150	50	-210
F1	400	100	0	1	0	250	240	250	10	30
F2	750	0	0	1	0	600	0	150	150	-300
F3	125	125	0	1	0	175	75	75	0	-172
F4	500	0	600	2	0	500	400	300	-400	0
G1	200	0	0	1	0	160	160	40	-120	-30
G2	1,000	0	0	1	0	600	400	400	0	-100
G3	150	0	0	1	300	250	0	200	-100	-135
G4	1,500	0	3,000	3	0	1,500	600	1,000	-600	-300

N.2 Directions across the four sectors

The following table distinguishes the source of adoption pressure from the business that could earn revenue from it. “Option” means an identifiable missing observation could change the decision. “Stop” applies to the generic proposition specified, not every use of the underlying technology.

Table 26: Mechanisms, value capture and decisions after public research.

Direction	Possible forcing mechanism	Position that can retain value	Decision
C1: corridor payouts	Lower completed-payout cost or better availability erodes rival margin	Local redemption/liquidity plus embedded distribution	Conditional advance
C2: orderflow rights	Solvers compete over a fixed trading contribution pool	Application or wallet with durable orderflow and fee rights	Option with distribution
C3: collateral mobility	Accepted claims free funding and allow tighter quotes	Bank/custodian or neutral connection with acceptance and financing	Requires a partner
C4: proof-carrying services	Buyers favor verifiable computation and exclude weaker claims	Specialized prover or buyer-recognized verification workflow	Stop generic thesis; paid-workload option
A1: seller negotiation	Rival sellers gain contribution or terms	Vertical commercial authority and audited transaction outcomes	Experiment only
A2: operated services	Cost/service competition reduces holdout viability	Service operator’s contracts and reliable completion	Advance narrow diligence
A3: agent authority	Exposed firms suffer attacks or lose required access	Cross-platform enforcement beyond incumbent coverage	Benchmark / option

Table 26: Investment directions (continued).

Direction	Possible forcing mechanism	Position that can retain value	Decision
A4: outcome acceptance	Accepted evidence and guarantees enable trade	Trusted acceptance workflow; separately, capitalized risk carrier	Small option
F1: continuous treasury	Better liquidity use tightens prices and service	Embedded bank/PSP controls and customer access	Requires access to customers
F2: shared fraud signals	Changing threats increase exposed accounts' expected loss	Useful receiving-account labels with operational rights	Requires a partner
F3: agent-payment guarantees	Merchants lose valuable transactions without accepted risk coverage	Guarantor or dispute operator with capital and claims control	Reject a generic protocol fee; investigate a paid guarantee
F4: continuous SME credit	Faster selection takes attractive borrowers from slower lenders	Lender/data partner with outcomes, funding and customer relationship	Requires data access and capital
G1: industrial cells	Lower cost or lead time takes work from comparable plants	Repeatable integrator, field service and equipment finance	Advance one application
G2: flexible power	Usable power determines competitive service capacity	Accepted flexibility and access/payment rights	Advance subject to enforceable contracts
G3: insurability retrofits	Insurer/lender conditions change property economics	Recognized retrofit service and financing channel	Requires enforceable contracts
G4: autonomous formulation	Qualified products undercut cost or improve performance	Qualified IP, validation and industrial scale-up	Requires a qualified product

Crypto. The payoff units are, respectively, a PSP with \$200m annual eligible flow (C1), a solver/market maker (C2), an institution with \$200m eligible collateral (C3), and a supplier of verifiable computation (C4). C2 assumes a \$300k contested contribution pool and $r = 3$, giving $\alpha = 600$; Flashbots' refund design demonstrates why gross orderflow is not retained protocol receipts (Flashbots, 2026). C3 assumes release of 10% of collateral at a 4% funding spread ($\gamma = 800$), plus a separately unverified \$500k acceptance/netting benefit. J.P. Morgan and Agorá establish functionality and experimentation, not those amounts (J.P. Morgan, 2026; Bank for International Settlements, 2026). C4 assumes 400 avoided verification hours at \$100 plus speculative acceptance/contest gains. Succinct proves execution of specified programs; it does not establish true inputs or a token-holder distribution (Succinct, 2026).

AI. A1's actor is a \$50m industrial supplier; its \$100k administration saving alone does not cover \$200k cost. A2's actor is the mature \$20m service firm, distinct from the new-business funnel. A3 assumes \$2m attributable incident loss and 40% incremental prevention, a particularly uncertain expected-loss input. NIST's identity project is not a mandate, and Microsoft offers a licensed incumbent alternative (NIST, 2026; Microsoft Learn, 2026). A4 assumes \$50k direct saving and

\$300k reusable acceptance value against \$200k cost; setting $\beta = 0$ makes $F = -150$. AI insurance supplies a market analogue, not proof of a new acceptance-network toll (Munich Re, 2026).

Fintech. F1 starts with \$50m prefunding: releasing 20% at 4% gives $\gamma = 400$, with another \$100k network benefit left unverified. FedNow’s native functionality is an alternative (Federal Reserve, 2026). F2’s \$5m residual loss must be measured after existing controls and actual loss allocation. F3 uses a hypothetical \$50m agent-payment flow and \$125k each of direct and network benefit; AP2 does not verify that paid volume (Google, 2026). A \$100k gross guarantee fee less \$60k claims, \$10k risk capital/transfer and \$10k transaction operations leaves only \$20k before other delivery costs. F4 assumes a \$100m credit book with 50bp operating/risk benefit and a separate contest edge. Historical mortgage evidence does not validate current SME selection or through-cycle losses (Fuster et al., 2018).

Greenfield. G1’s unit is a production cell, G2’s a 20MW site. G3 uses a \$100m commercial property portfolio: \$100k assumed premium savings plus \$50k uninsured-loss reduction, and a separate \$300k enforced access penalty. California’s mitigation guidance does not establish that commercial contract; without ε , $F = -100$ (California Department of Insurance, 2026). G4 assumes \$100m of affected industrial product sales, 1.5% direct cost improvement and a \$1.5m contest stake with $r = 3$. Laboratory results do not establish qualification, yield or royalties. The A-Lab correction and autonomous chemistry evidence motivate investment in validation and qualified-product ownership, while challenging automatic discovery-value claims (Szymanski et al., 2026; Dai et al., 2024).

N.3 Commercial scale and financing

The assumed customer pools are scenarios, not verified market sizes. Do not sum overlapping budgets: C1/F1 share payment buyers, A3/A4/F3 can sit in one transaction, and F2/F4 can use related data. A token waterfall must separately deduct rebates, execution, security, incentives and dilution before identifying any enforceable holder claim.

The narrower C1 cross-check starts from the dated \$226bn global B2B run rate. Assumed filters of 20% origin/segment relevance, 50% eligible payout use and 40% accessibility leave \$9.04bn flow; 10% entrant share at 6bp yields \$0.54m. Only the starting global estimate is source-derived. The original eventual pool implied \$240bn of adopted eligible flow and therefore cannot be called a verified current subset of that estimate.

Table 27: Constructed month-24 outcomes and capital envelopes, \$m unless stated. Annualized service/equipment receipts are not all software ARR. Debt and funding terms are assumed, not offers.

	Active units	Annualized receipts	Peak equity	With 25% buffer	Unfinanced need
A2	65.4	3.27	3.25	4.06	3.25
G1	28.9	2.89	3.93	4.91	7.10
C1	4.6	0.55	3.76	4.69	4.69
G2	4.0	1.00	5.59	6.99	5.59
F2	7.1	2.12	6.04	7.55	6.04
A1	42.7	4.27	2.06	2.58	2.06

Each monthly path starts with no installed base and includes sales delay, installation capacity, churn, delivery costs, acquisition/onboarding, fixed expense and receivables. Equipment purchases, depreciation, amortizing debt and operating liquidity are accounted for separately. Customer counts

are fractional expected cohorts under assumed conversion, not scheduled installations. No acquisition purchase price, terminal value or data-center construction is included.

The seller-negotiation scenario projects attractive supplier revenue, but its assumed customer savings do not cover adoption costs. All five other paths still consume cash at month 24. A 50% fee cut with unchanged direct costs makes A2 and G1 economic gross profit negative; slow acquisition can reduce asset purchases and funding while also shrinking the business. Neither a fast sales curve nor low burn substitutes for customer economics.

N.4 Falsification and reproducibility

The script removes unmonetized benefit, sets network and exclusion channels to zero, changes waiting and implementation cost, and distinguishes reducing r at fixed α from reducing it at fixed V , where $\alpha = V(r - 1)$. Shelter uses the base model’s capped endpoint loss; a compatible supplier is compared directly. No price ceiling is interpreted as a realized price forecast.

The productivity deadline is fragile. Under the industrial-cell assumptions, $h(1) = 160$ is less than initial profit 180, so the hypothesized holdout break-even lies beyond full adoption and there is no finite deadline. A2’s 15-month latest-start illustration remains conditional on its assumed diffusion rate; it is not a calendar forecast.

Public failures and alternatives are included where found, but no verified profitable non-adopter panel was obtained for the six initial fields. A company’s silence about a technology is not evidence of holding out. The assessment is based on public sources. Interviews, customer invoices, randomized comparisons and private underwriting records are still needed.

Running `scripts/investment_execution.py` with the paper dependency group regenerates 48 endpoint cases, 218 parameter stresses, 432 monthly cohort rows, figures and tables. It checks endpoints against the repository model and reconciles financing, depreciation, fee compression and customer stock/flows. The accompanying executed notebook and public-source ledger are in `data/investment_directions/execution/`.

O AI and crypto: evidence and accounting boundaries

The case study in Section 4 uses primary research, protocol specifications, company disclosures and documented product terms. The research ledger is `data/investor/research/R20_ai_crypto_platforms.md`; source-level claims, dates and limitations are in `data/investor/evidence.json`. The source cutoff is October 20, 2026. Historical observations retain their original periods; documentation without a publication date describes the version accessed that day. The cases were selected for their mechanisms and cannot estimate the success rate of investments across either industry.

O.1 What each receipt measures

A payment must be assigned to its recipient before it can support a claim about capture. Bitcoin block rewards combine issuance and user fees. Issuance transfers purchasing power to miners and is not customer sales. On Ethereum, the base-fee burn reduces supply, whereas priority fees and issuance enter validator economics differently. Solana’s fee documentation likewise distinguishes the burned and validator portions. Hyperliquid’s trading fees, deployer receipts and assistance-fund mechanism need to be separated from HyperEVM gas, whose documentation specifies burning both base and priority fees (Hyperliquid, 2026d). These mechanisms do not make protocol dashboards equivalent to audited income statements. Token appreciation, market capitalization and assets

deposited in contracts are different stocks; trading fees are flows. Neither the stocks nor gross fees establish a holder’s contractual cash-flow entitlement or a supplier’s net profit.

For AI, recognized revenue, annualized revenue, contracted backlog, capital expenditure and financing commitments answer different questions. A run rate is not a completed year’s sales. Capital expenditure buys assets used over multiple periods; depreciation and utilization affect operating returns. Hyperscaler capital budgets also support other workloads. Adding a lab’s purchases to its cloud provider’s investment and the hardware company’s sales counts successive links in the same spending chain. The main section therefore uses NVIDIA’s reported Data Center revenue as evidence of upstream receipts, without allocating all of it to a contest or treating a corporate gross margin as a segment margin.

The same boundary applies to research costs. DeepSeek’s training estimate prices specified GPU-hours; it excludes earlier research and experiments. Comparable lab economics would require personnel, data, failed runs, inference, support, financing and subsequent model generations. The public sources reviewed here do not provide a consistent, audited model-level cost panel for the five providers. We therefore make no cross-provider profitability ranking and no estimate of the share of AI spending that is socially wasteful.

O.2 Cloud investment and the cash that funds it

The large cloud providers can fund expansion from existing businesses as well as external capital. Table 28 reconciles selected full-year cash flows. Each row covers the whole company and its stated fiscal period; none isolates AI investment or the return on a particular model. The final column is operating cash less the specified property-and-equipment outflow, before acquisitions, financing flows and other investing uses.

Table 28: Selected realized cloud-provider cash flows, in US\$ billions. Different fiscal periods and capital-spending definitions are retained explicitly.

Company and year ended	Operating cash	Cash PP&E	Difference
Alphabet, December 31, 2025	164.7	91.4	73.3
Amazon, December 31, 2025	139.5	128.3	11.2
Microsoft, June 30, 2026	182.9	115.9	67.0

Sources: company releases ([Alphabet Inc., 2026](#); [Amazon.com, Inc., 2026a](#); [Microsoft Corporation, 2026b](#)). Amazon nets property sales and incentives against purchases; Alphabet reports purchases and Microsoft reports cash additions. Differences are calculated from unrounded amounts, then rounded. These outflows exclude noncash asset additions.

Leases make the distinction consequential. On its July 2026 call, Microsoft reported quarterly capital expenditure of \$41 billion, finance leases of \$5.6 billion and cash paid for property and equipment of \$35.8 billion. It also explained that revised asset lives would shift some future data-center leases from finance to operating classification: the latter do not enter its headline capex measure ([Microsoft Corporation, 2026a](#)). A lower reported capex path can therefore reflect accounting treatment while future payment obligations remain. Contractual commitments, depreciation, utilization and operating cash must be checked together. These disclosures also show why a claim about loss-making independent laboratories cannot be extended to every participant financing the race.

O.3 Announcements, financing and realized operations

The September 22, 2025 OpenAI–NVIDIA announcement described a letter of intent for at least 10 GW of systems and an intended investment of up to \$100 billion, deployed progressively ([OpenAI,](#)

2025). Those were announcement terms, not a record of delivered capacity or recognized sales; the initial announcement cannot establish its subsequent completion. The FTC’s January 2025 study of cloud/model partnerships identified arrangements requiring developers to spend substantial parts of their partners’ investment on those partners’ cloud services ([Federal Trade Commission Staff, 2025](#)). Such links can finance useful capacity while making the buyer and supplier financially interdependent. They require tracing rather than summing headline commitments.

Core Scientific provides a realized bridge between the two sectors. Its 2025 annual filing separates colocation, self-mining and hosted-mining revenue, and separates passed-through power fees within colocation ([Core Scientific, Inc., 2026](#)). The observed revenue shift supports a change in the use of sites and power. It does not imply that announced multiyear contract values were received in 2025, that all converted sites were operational, or that an ordinary mining site meets an AI customer’s cooling, networking and reliability requirements. Customer funding of a conversion also changes the allocation of construction and credit risk.

O.4 Events that change different parts of the mechanism

China’s 2021 restrictions affected where mining could operate. Cambridge’s October account reports a sharp June fall in global hashrate, partial recovery in July and August, and changes in geographic shares ([Rauchs, 2021](#)). Its geographic estimates came from participating mining pools; they do not fully observe concealed activity. Ethereum’s 2022 Merge changed which method could produce valid blocks on Ethereum itself. The estimated reduction in its energy use belongs to that consensus transition, not to Bitcoin or the whole crypto industry ([Ethereum.org Contributors, 2026b](#)). The two events changed different mechanisms: China’s restrictions relocated mining, while Ethereum replaced proof-of-work.

The analogous AI branches include changes in efficiency, access to weights, chip access and customer demand. The conditional table is not an equilibrium solution for these branches. Capital renewal, entry and exit, heterogeneous capacity, strategic subsidies and changes in the total prize are outside the base adoption model. Its illustrative payoff calculation demonstrates why a positive incentive to adopt cannot identify the sign of the field’s aggregate gain. It cannot estimate consumer surplus, environmental costs, security benefits or the spillovers of model research.

O.5 Measurements still needed to distinguish the hypotheses

A platform-retention study needs fixed cohorts, periods and definitions. For each chain, measure fee-paying users and developers retained after incentives, the applications they actually use, and whether liquidity is duplicated or migrated. Separate trading volume from net deposits and useful settlement; remove known wash trading, automated repetition and consensus votes before comparing activity. For models, measure retained paid workloads and successful outcomes at a fixed quality threshold after credits end or a competing model launches. Record routing across providers, evaluation costs and failures, rather than counting API accounts as exclusive customers.

The public evidence establishes mechanisms and some realized outcomes, but does not supply those matched cohorts across all nine cases. In particular, Hyperliquid’s fee-routing rules establish a possible path from activity to token supply changes; they do not establish sustained external-application retention. Electric Capital’s repository data observe open-source contributions, with overlapping chain attribution, rather than exclusive customer relationships ([Electric Capital, 2026](#)). The absence of a comparable panel is a limit on the conclusion, not evidence that the network effect is absent.

For a payout corridor, measure delivered local purchasing power and availability time after every conversion and fee. Stripe’s Remote account supplies a reported benefit and a concrete workflow;

it supplies neither randomized evidence nor a published cohort of repeat users and complete redemption costs (Stripe, 2026). For AI tasks, measure completion quality and review time as well as speed. The support study and METR’s evolving experiment answer narrower questions than economy-wide productivity. These measurements would test the investment hypotheses.

O.6 Forecast definitions, assumptions and resolution

The October 20, 2026 register contains eleven capability fields and separate entries for October 20, 2031 and October 20, 2036. It is stored in `data/ai_crypto_forecasts/forecasts.csv`; definitions, source records and assumptions are adjacent. The research ledger is `data/investor/research/R21_ai_crypto_forecasts.md`. All twenty-two entries are qualitative conditional forecasts. The observed quantities, external projections, customer claims and posted offers in `observations.csv` are explicitly distinguished. None is treated as a fitted actor-level panel.

The commercial-field event compares annual operating profit with a no-technology counterfactual in constant purchasing power, controlling for common shocks. It follows the original 2026 cohort including exits, with entrants reported separately. Nonprofitable firms, traders’ execution costs and service-access measures need separate denominators and cannot be pooled into this profit-ratio event. The central convention is at least 10% loss for at least 50% of the field over eight consecutive quarterly intervals, after the best feasible non-adopting response and with an economically feasible adoption response. We vary loss thresholds over 5, 10 and 20%, and reach over 25, 50 and 75%. Actual resolution requires matched actor outcomes and implementation costs. Missing observations leave an event unresolved.

Pressure reached by a horizon means that some qualifying eight-quarter window has completed. Pressure active at a horizon requires the final eight quarters to qualify. A 2026 observation does not establish the preceding persistence history, and a declining loss can still represent an active event. The register records current pressure separately from the expected mechanism of further change. It does not claim measured customer counts, aggregate dollars at risk, onset dates or probability calibration.

O.6.1 What the reproducible calculations establish

The script `scripts/ai_crypto_forecasts.py` supplies mechanism stress tests, not parameter estimates behind the qualitative tiers. There are eleven constructed branches across six mechanisms, each observed at the origin and forty quarterly endpoints. Linear interpolation joins the input knots in Table 29. Payoffs remain annual flows; the script does not compound them as quarterly returns. Exposure is imposed conditionally rather than generated by diffusion. Replacement states therefore do not purport to solve ten years of irreversible adoption.

Each artificial field has cost multipliers 0.5, 1, 1.5 with weights 0.25, 0.5, 0.25. These are assumptions, not estimates of firm heterogeneity. Conditional effectiveness is $p = 1$; it is not an event probability. No weights are assigned to branches, which overlap and cannot be added into forecast probabilities. The calendar, thresholds and complete input paths are frozen in JSON; the output includes both the level of loss and its change from 2026.

For ordinary model branches, the script uses the existing payoff functions, with $\Phi(x) = x$, $\alpha = V(r - 1)$ and an additional common price loss $h = \rho\gamma x$. That loss reduces both adopter and holdout profits and cancels from their switching difference. The genuine-exit branch uses the model’s outside option and removes departed contest stakes; a compatible competing provider is not that exit. To avoid an unsupported combination, the script rejects simultaneous common price pressure and a finite outside option. The defense branch instead sets

$$U_N = -L_{\text{attack}}, \quad U_A = -(1 - \varphi)L_{\text{attack}} - c, \quad M = \varphi L_{\text{attack}} - c - W.$$

Here L_{attack} is incremental expected loss after existing controls. This is separate-population accounting, not an application of the strategic-counter theorem.

An artificial cohort qualifies only when holdout loss exceeds the chosen threshold, switching margin is positive and post-adoption operating profit remains positive. The additional viability condition prevents a beneficial but still insolvent response from being counted as feasible. All remaining alternatives are conditional assumptions of the branch; these calculations do not establish that real firms lack other responses.

Table 29: Assumed input knots at 2026 / 2031 / 2036, with linear quarterly interpolation. All payoff flows use baseline-profit units. Unlisted $V, \beta, \varepsilon, \rho, W$ are zero and unlisted $r = 1$. These are mechanism tests, not estimates.

Stress test	Assumed parameter path
Workflow price competition: Savings transmitted	$x=0.15/0.6/0.85$; $\gamma=0.5/0.9/1.2$; $c=0.3/0.35/0.4$; $\rho=0.2/0.6/0.8$; $\text{wait}=0.03/0.03/0.03$
Workflow price competition: Weak transmission	$x=0.15/0.6/0.85$; $\gamma=0.5/0.9/1.2$; $c=0.3/0.6/0.9$; $\rho=0.1/0.1/0.05$; $\text{wait}=0.03/0.03/0.03$
Capacity contest: Continuing contest	$x=0.5/0.85/0.95$; $\gamma=0.15/0.3/0.45$; $V=0.6/0.3/0.2$; $r=2/3/3$; $c=0.25/0.3/0.35$; $\text{wait}=0.02/0.02/0.02$
Capacity contest: Feasible exit	$x=0.5/0.85/0.95$; $\gamma=0.15/0.3/0.45$; $V=0.6/0.3/0.2$; $r=2/3/3$; $c=0.25/0.3/0.35$; $\text{wait}=0.02/0.02/0.02$; $Q=-0.2/-0.08/-0.05$
Payout corridor: Corridor advantage	$x=0.1/0.5/0.75$; $\gamma=0.2/0.4/0.45$; $\beta=0.05/0.1/0.15$; $c=0.12/0.18/0.16$; $\rho=0.25/0.7/0.8$; $\text{wait}=0.01/0.01/0.01$
Payout corridor: Bank substitution	$x=0.1/0.5/0.75$; $\gamma=0.2/0.12/0.08$; $\beta=0.05/0.05/0.04$; $c=0.12/0.18/0.16$; $\rho=0.25/0.4/0.3$; $\text{wait}=0.01/0.01/0.01$
Compatible access: Access dependence	$x=0.1/0.6/0.85$; $\gamma=0.03/0.1/0.1$; $\beta=0.4/0.6/0.8$; $\varepsilon=0/0.3/0.5$; $c=0.1/0.15/0.15$; $\text{wait}=0.01/0.01/0.01$
Compatible access: Optional connection	$x=0.1/0.6/0.85$; $\gamma=0.03/0.1/0.1$; $\beta=0.4/0.6/0.8$; $c=0.1/0.15/0.15$; $\text{wait}=0.01/0.01/0.01$
External attack defense: Adapting threat	$\text{threat}=0.05/0.3/0.6$; $\text{efficacy}=0.7/0.7/0.7$; $c=0.08/0.12/0.15$; $\text{wait}=0.01/0.01/0.01$
External attack defense: Bundled protection	$\text{threat}=0.05/0.08/0.06$; $\text{efficacy}=0.7/0.7/0.7$; $c=0.08/0.12/0.15$; $\text{wait}=0.01/0.01/0.01$
Benefits without forcing: Useful but optional	$x=0.1/0.6/0.9$; $\gamma=0.3/0.3/0.3$; $\beta=0.4/0.4/0.4$; $c=0.1/0.1/0.1$

Table 30: Constructed stress tests, not empirical forecasts. Loss and its change from 2026 are percentage points of baseline annual operating profit. Event: 10% loss, 50% reach, eight completed quarters. By / active distinguishes attainment from persistence.

Assumed mechanism / branch	2031 loss / change	2036 loss / change	2031 by / active	2036 by / active
Workflow price competition: Savings transmitted	32.4 / +30.9	81.6 / +80.1	Yes / Yes	Yes / Yes

Continued on next page

Continued from previous page

Assumed mechanism / branch	2031 loss / change	2036 loss / change	2031 by / active	2036 by / active
Workflow price competition: Weak transmission	5.4 / +4.7	5.1 / +4.3	No / No	No / No
Capacity contest: Continuing contest	18.9 / -1.1	13.1 / -6.9	Yes / Yes	Yes / Yes
Capacity contest: Feasible exit	8.0 / -12.0	5.0 / -15.0	Yes / No	Yes / No
Payout corridor: Corridor advantage	14.0 / +13.5	27.0 / +26.5	No / No	Yes / Yes
Payout corridor: Bank substitution	2.4 / +1.9	1.8 / +1.3	No / No	No / No
Compatible access: Access dependence	18.0 / +18.0	42.5 / +42.5	No / No	Yes / Yes
Compatible access: Optional connection	0.0 / +0.0	0.0 / +0.0	No / No	No / No
External attack defense: Adapting threat	30.0 / +25.0	60.0 / +55.0	Yes / Yes	Yes / Yes
External attack defense: Bundled protection	8.0 / +3.0	6.0 / +1.0	No / No	No / No
Benefits without forcing: Useful but optional	0.0 / +0.0	0.0 / +0.0	No / No	No / No

The stress tests distinguish the size of a loss from its duration. In the constructed payout branch, endpoint loss reaches 14% in 2031 but has not remained material long enough to pass the two-year test. The access branch likewise has a high endpoint before completing the required window. In the exit branch, pressure is reached before 2031 but is no longer active at either horizon. Direct and network benefits alone never create holdout loss. The full 198 threshold-horizon results in `scenario_sensitivity.csv` expose these distinctions; they are not 198 independent forecasts.

Capture requires another comparison. With $c = q + \kappa$, the holdout-based ceiling is $q \leq B - \kappa - W$, but a compatible supplier can impose a tighter bound. For example, if avoiding the holdout produces gross value 0.50 and other implementation cost is 0.10, the ceiling before waiting is 0.40. If an equally effective alternative costs 0.12 in total and the focal product still requires 0.10 of other costs, the focal supplier can charge at most 0.02 at equal quality. These assumed numbers illustrate why capability forcing need not imply protocol rent. Actual gross profit then deducts service costs, and capital-intensive suppliers must fund replacement and earn returns on their invested capital.

O.6.2 Evidence boundaries and the next resolution

The source register retains publication and observation periods. The support-study entry uses the published-version sample and effect, without mixing earlier working-paper numbers. Census's broadened question is a series break. The IEA's 2030 number is an external projection, IFR counts include conventional industrial robots, incident reports are selected cases, and an advertised API price is not sales. Source-specific limitations are retained alongside the observations.

The forecasts lack interviews, a complete panel covering adopters, holdouts and exits, calibrated loss distributions and prospective validation. Consequently no parameter recovery, held-out forecast

accuracy, probability calibration or industry-wide capture estimate is reported. The numerical tests verify economic accounting and event-window logic, not prediction accuracy. A future quantitative edition must first identify the relevant exposure and obtain the outcomes needed to distinguish direct value, contest transfer and exclusion.

The next review should establish cohorts and measurement contracts in 2027; test pass-through, retention and incremental security losses in 2028; and apply the candidate-specific 2029 reversal tests before upgrading broad claims. Dated revisions should retain the original forecasts and record failed candidates and exits, allowing later assessment against the planned milestones.

P Where the framework comes from

The model is a binary choice with externalities in the sense of Schelling (1973, 1978), with heterogeneous thresholds as in Granovetter (1978). The distinction between a rival’s adoption raising one’s *incentive* to adopt and raising one’s *payoff* is the strategic-complements distinction of Bulow et al. (1985), and the threshold structure of games with strategic complements is that of Milgrom and Roberts (1990). The two earlier reduced forms each measure only one side of it. Network and compatibility forcing follow Katz and Shapiro (1985, 1986), Farrell and Saloner (1985, 1986) and Arthur (1989); two-sided pricing of the crossing follows Rochet and Tirole (2003); switching costs and lock-in follow Klemperer (1987), Shapiro and Varian (1999) and Farrell and Klemperer (2007). The positional channel is a proportional-share contest in the tradition of Tullock (1980); its welfare consequences are the positional arms races of Hirsch (1976) and Frank (2011), the rank-order tournaments of Lazear and Rosen (1981), the winner-take-all markets of Frank and Cook (1995) and the races of Harris and Vickers (1985). Adoption timing under rivalry follows Reinganum (1981) and Fudenberg and Tirole (1985), and the waiting value draws on Dixit and Pindyck (1994). Diffusion curves follow Griliches (1957) and Bass (1969); the cost curves of Section 2.3 follow Wright (1936) and Arrow (1962). Teece (1986) explains when innovators can retain returns, while Sutton (1991) links market concentration to endogenous sunk costs. Reference dependence in customers’ expectations follows Kahneman and Tversky (1979), and information cascades, the wave that carries no payoff information, follow Bikhchandani et al. (1992). On the empirical side, identification of peer effects in adoption faces the reflection problem of Manski (1993), and the hazard here is a dynamic relative of the discrete-choice social-interactions model of Brock and Durlauf (2001). Network exposure follows Galeotti et al. (2010). Demonstration thresholds under noisy beliefs connect to global games (Carlsson and van Damme, 1993; Morris and Shin, 1998). Institutional and fad-driven adoption (DiMaggio and Powell, 1983; Abrahamson and Rosenkopf, 1993) motivates the treatment of belief-driven “false forcing”. For artificial intelligence, Armstrong et al. (2016) model a racing game of the arms-race type.

Three further connections clarify the extensions. Competitive learning can repeatedly erode a rival’s advantage, as in the Red Queen account of Barnett and Hansen (1996). Customer requirements can differ between attractive features and minimum expectations (Kano et al., 1984). And evidence that collective restrictions raise producer profits need not establish a social gain, as the cigarette-advertising analysis of Farr et al. (2001) shows. These connections help separate the transmission of competitive pressure from its consequences for customers, entry and successive generations of technology.

Q Earlier models and reproduction benchmarks

The exact parameter restrictions are stated in Proposition 11. The threshold model retains exclusion pressure and the linear model holds the contest gap constant.




Figure 22: Four of the 14 historical diffusion series with the unified model’s fitted paths ($r = 1$, $\varepsilon = 0$, so the paths coincide with the linear model’s fits). Payoff parameters appear in Table 31; the complete case-specific settings, including beliefs, seeds and adjustment speeds, are stored in `data/claude_cases.json` and `data/claude_fits.json`.

Table 31: Inherited calibration of 14 historical cases under the linear restriction ($r = 1$, $\varepsilon = W = 0$, no outside option). $W(x_T)$ is realized rival welfare at the end of the series relative to $x = 0$. [†]Sensitive to the rounding of stored parameters (see text).

Case	α	β	γ	c	σ	RMSE	T_{50}	$\Delta(0)$	$\Delta(1)$	y	$W(x_T)$	Region
Hybrid corn	0.35	0.05	0.60	0.20	0.15	0.019	10 yr	+0.75	+0.80	+0.45	+0.43	productive
HTTPS	0.45	0.25	0.30	0.80	0.15	0.014	4 yr	-0.05	+0.20	-0.25	-0.17	wasteful bandwagon
Dreadnought	0.90	0.00	0.10	0.15	0.10	0.043	4 yr	+0.85	+0.85	-0.05	-0.05	arms race
Container	0.10	0.60	0.30	0.49	0.58	0.035	16 yr	-0.09	+0.51	+0.41	+0.50	coordination
MLB analytics	0.85	0.00	0.15	0.10	0.10	0.015	6 yr	+0.90	+0.90	+0.05	+0.05	productive
Frequent flyer	0.80	0.05	0.15	0.08	0.10	0.035	7 mo	+0.87	+0.92	+0.12	+0.11	productive
LZR suit	1.00	0.00	0.00	0.05	0.34	0.021	3 mo	+0.95	+0.95	-0.05	-0.05	arms race
Clap skates	0.90	0.00	0.10	0.24	0.38	0.005	12 yr	+0.76	+0.76	-0.14	-0.12	arms race
Nuclear [†]	1.00	0.00	0.00	1.00	0.19	0.022	52 yr	-0.002	-0.002	-1.00	-0.39	inert
Armed drones [†]	0.65	0.00	0.35	0.21	0.25	0.057	25 yr	+0.79	+0.79	+0.14	+0.12	productive
China NEVs	0.45	0.20	0.35	0.60	0.15	0.012	10 yr	+0.20	+0.40	-0.05	-0.02	arms race
Gen AI (large orgs)	0.40	0.10	0.50	0.50	0.54	0.025	4 qtr	+0.40	+0.50	+0.10	+0.25	productive
AI (all US firms)	0.25	0.15	0.60	0.99	0.24	0.003	—	-0.14	+0.01	-0.24	+0.02	wasteful bandwagon
RFID mandate	0.20	0.50	0.30	2.00	0.26	0.002	—	-1.50	-1.00	-1.20	-0.001	inert

The threshold calibration reproduces in full; thirteen of the fourteen linear-model fits reproduce to the precision of the stored parameters. All 15 threshold-model technologies match in G , in current pressure $S\Phi(x)$, and in status. Thirteen of the 14 fitted linear-model series reproduce their demonstration date, T_{50} , T_{90} , final share and RMSE within the five-decimal rounding of the stored fits (Figure 22, Table 31).

The fourteenth, nuclear weapons, does not, and the reason is substantive. Its fit sets $x_0 = x_v$ and $F_c(p_0z) \approx x_0$ to five decimals. A 5×10^{-6} change in x_v decides whether a demonstration ever happens, so the fitted demonstration date is not identified. In the terms of Proposition 9, the fit puts the visibility threshold at the pre-demonstration ceiling x_L , which is where the demonstration date in Figure 9(c) diverges.

The inherited calibration should be read with three caveats.

- The static split of (α, β, γ) was set by judgment and ban evidence, not estimated. By Proposition 12 the diffusion fits cannot check it.
- Region labels use the mean-cost rival. Cases with a plateau below full adoption have c pushed up by the fit, and HTTPS’s “wasteful bandwagon” label comes from its assumed positional share $\alpha = 0.45$, not from evidence. That label is a candidate for re-estimation, not a finding.

- These are reproduction benchmarks, not the paper’s updated estimates. Appendix J rebuilds the historical series and reports fits allowing $r \neq 1$, transitory shocks and, where relevant, exclusion.

Why the earlier dreadnought interpretations differed. The two reduced forms disagree on whether a lone dreadnought was worth building. The threshold calibration has $\Delta(0) = -G = -0.33$, while the linear calibration has $\Delta(0) = \gamma + \alpha - c = +0.85$. In the unified model the disagreement is over $\Delta(0) = p[\gamma + \alpha] - c$: whether the first mover’s positional edge α is counted, and whether incumbents’ installed fleets, which are part of their c , outweighed it. That is an empirical question about the adopter curve at low x , Britain’s relative position in 1906–1909, rather than a definitional one.

R Learning and forward-looking timing

Appendix E.3 describes the decision implications. These equations specify possible extensions; they are not the adoption rule used for the historical fits.

One way to allow both favorable and unfavorable evidence is a binary quality state $\zeta \in \{0, 1\}$ and a signal Y_t with state-dependent density ℓ_ζ . Beliefs update by

$$p_{t+1} = \frac{p_t \ell_1(Y_t)}{p_t \ell_1(Y_t) + (1 - p_t) \ell_0(Y_t)}.$$

Signal likelihoods must account for the pilot’s positional advantage and its selected cost. Otherwise the inference confuses success against unprepared rivals with value after they respond. The jump to $p = 1$ is the limiting case of conclusive favorable evidence, not a general learning law. Downward revisions can halt new adoption while leaving earlier commitments stranded. Institutional pressures can also persist after a technological disappointment, because exclusion is not multiplied by p in (8); this connects mistaken technological expectations to institutional bandwagons (Abrahamson and Rosenkopf, 1993; DiMaggio and Powell, 1983).

Anticipation is a separate issue. Let s contain exposure, beliefs, costs and any announced deadline, and let $\delta \in (0, 1)$ discount future flows. With irreversible adoption, an optimal-stopping formulation is

$$J_A(s) = U_A(s) + \delta \mathbb{E}[J_A(s') \mid s], \quad J(s) = \max\{J_A(s), U_N(s) + \delta \mathbb{E}[J(s') \mid s]\}.$$

Here J is the value while still free to wait and J_A is the value after commitment. The transition law must incorporate rivals’ strategies for this to be a timing game (Fudenberg and Tirole, 1985; Rust, 1987). The historical fits use the myopic hazard defined in Appendix E. A future cutoff can bring investment forward when installation takes time or capacity must be reserved. With instantaneous, freely available installation, a future penalty alone need not justify paying today’s operating cost early. Empirically, order dates, installation dates and productive-use dates should therefore be recorded separately. Adoption before measured rival use may reflect anticipation rather than an absent competitive effect.

S Reproducibility

The model is `src/forcing/model.py`. Each result is checked in `tests/`:

Result	Test
Section 1	test_payoffs_are_relative_to_the_status_quo test_outside_option_caps_the_loss_and_withdraws_the_stake
Lemma 1	test_contest_is_zero_sum test_welfare_is_mean_payoff_and_ignores_the_contest
Lemma 2	test_contest_slope_has_the_sign_of_minus_alpha_r_minus_one test_small_edge_limit_is_the_linear_contest test_decisive_contest_makes_the_holdout_loss_a_step
Proposition 1	test_contests_alone_never_tip
Definition 1	test_classification_ignores_the_waiting_value
Proposition 2	test_selective_forcing_closed_form
Proposition 3	test_shakeout_closed_form test_outside_option_stops_adoption_at_the_shakeout
Proposition 4	test_exclusion_raises_forcing_but_not_full_adoption_welfare
Corollary 1	test_outside_option_caps_the_loss_and_withdraws_the_stake
Proposition 5	test_private_social_wedge
Propositions 6, 7	test_ban_criterion_includes_exclusion test_interior_allocation_beats_both_endpoints test_contest_waste_vanishes_at_cost_extremes test_tax_receipts_are_transfers
Appendices E.2, H	test_pilot_return_separates_edge_and_network test_selection_can_raise_adoption_without_conversion test_continuous_contest_equilibrium
Proposition 8	test_linear_dynamics_match test_homogeneous_costs_give_the_logit_hazard test_persistent_heterogeneity_makes_the_hazard_fall
Appendix E.5	test_failed_seeds_freeze_rather_than_unravel
Proposition 9	test_demonstration_needs_visibility_below_the_ceiling
Appendix E.7	test_waiting_value_sharpens_the_demonstration_cliff
Observation 1	test_aggregate_recursion_is_exact_only_for_monotone_thresholds
Proposition 10	test_counter_nests_the_base_model test_counter_dynamics_nest_the_base_dynamics test_counter_contest_is_zero_sum test_counter_welfare_is_the_mean_payoff test_counter_complements_adoption_and_other_counters test_counters_leaves_the_linear_contests_gain_from_adopting_unchanged test_full_counter_restores_the_split
Corollary 4	test_disarmament_test_with_a_counter
Appendix F	test_counter_shares_match_sampled_rivals test_counter_dynamics_match_a_finite_population test_counters_arrive_late_and_stop_rather_than_reverse_adoption test_a_counter_raises_welfare_only_by_deterring_adoption
Proposition 11	test_threshold_model_is_a_special_case test_threshold_calibration_table_reproduces test_linear_model_is_a_special_case test_linear_fits_reproduce test_nuclear_fit_sits_on_a_knife_edge
Proposition 12	test_adoption_data_cannot_split_alpha_from_gamma test_adoption_data_cannot_split_standalone_gain_from_waiting test_belief_separates_exclusion_from_network_gains test_utility_rescaling_preserves_adoption
Appendix J	test_archived_map_likelihoods_reproduce

Figures and Table 31 are generated by `scripts/paper_assets.py`. The four worked cases in Section 7 are regenerated by `scripts/investor_assets.py`. The contest, detection and invoice cases use the planning assumptions in `data/investor/scenarios.json`; their purchase margins are checked against `forcing.gain`, and the contest field path comes from `forcing.simulate`. The fulfillment case uses `data/investor/productivity.json`, the explicit extension in `src/forcing/productivity.py`, and `scripts/productivity_assets.py`. Its deployment trajectory and price response are assumptions, not estimates from the base model. That script can also be run separately to regenerate Table 15, its CSV, the timing sensitivity CSV and the supplier summary JSON. Tests in `tests/test_productivity.py` check the separation of switching gain from holdout loss, the viability ceiling, the customer transfer, and the disappearance of the timing threshold under low pass-through.

Every source cited in the investor sections is registered in `data/investor/evidence.json`. Calibration figures and tables are generated by `scripts/calibration_assets.py` from the stored fit summaries. These record input and code hashes. The input-integrity checks and a replay of all 42 stored posterior-mode likelihoods verify the assembled calibration against the current code; they do not replace convergence or sensitivity checks on the posterior sampling.

T Data sources

Every series in Appendix J is rebuilt by

```
uv run python -m forcing.collect all
```

which, for each case:

1. re-downloads every snapshot;
2. re-derives parsed rows from the snapshots;
3. verifies hashes, quotes, values and schema;
4. builds the fitted series.

`scripts/fit_cases.py` then writes each calibration with the hashes of its inputs and of the code. A test fails if any committed calibration was fitted from inputs that have since changed.

Table 33 summarizes the rows and how they were obtained:

- *Crawled* sources were fetched byte for byte by code.
- *Captured* sources were saved from a browser session, because the publisher blocks scripts.
- *Parsed* rows were derived by code.
- *Quoted* rows were transcribed with a verbatim quote checked against the snapshot.

Flags count rows that are stated, derived, estimated, and excluded. The list that follows gives every source used, by case; the source registries add access dates, hashes and Wayback Machine captures.

Source	Method	Licence
Clap skates Jos J. de Koning (Vrije Universiteit Amsterdam), in Sportscience News (1997). Background to the slap skate: Fifteen years of slap skate history, biomechanical backgrounds, first results and recent developments. http://www.sportsci.org/news/news9703/slapxtra.htm	crawler	copyrighted

Source	Method	Licence
International Skating Union (2026). Speed Skating and the clap skate revolution. https://isu-skating.com/speed-skating/news/speed-skating-and-the-clap-skate-revolution/	crawler	copyrighted
Olympedia (OlyMADMen) (2026). Speed Skating at the 1998 Winter Olympics (overview). https://www.olympedia.org/editions/46/sports/SSK	crawler	copyrighted
Stephen Seiler (Sportscience News) (1997). The new Dutch "Slapskates": Will They Revolutionize Speed Skating Technique?. http://www.sportsci.org/news/news9703/slapskat.htm	crawler	copyrighted
Containerization, North Atlantic & UK		
Daniel M. Bernhofen, Zouheir El-Sahli, Richard Kneller (2012). Estimating the effects of the container revolution on world trade (working paper, February 3, 2012; Figure 2: Changes in container utilization in the UK and Japan). http://economicsarchive.fiu.edu/events/2013/seminar-daniel-bernhofen/bek_container_feb-3-2013.pdf	crawler	copyrighted
Benjamin Bridgman (Bureau of Economic Analysis) (2014). Why Containerization Did Not Reduce Ocean Trade Shipping Costs. https://www.aeaweb.org/conference/2015/retrieve.php?pdfid=3462&tk=k5HEFitY	crawler	copyrighted
Frank Broeze (2002). The Globalisation of the Oceans: Containerisation from the 1950s to the Present, ch. '1970s: conquering the world' (Cambridge Core chapter page with extract). https://www.cambridge.org/core/books/globalisation-of-the-oceans/1970s-conquering-the-world/F87A212081653D8714A7393E7D008209	crawler	copyrighted
Hybrid corn, US		
U.S. Department of Agriculture (Statistical Reporting Service) (1963). Agricultural Statistics 1962 (Table 46: Corn hybrids: Acreage of corn planted with hybrids, United States, 1933-60). https://esmis.nal.usda.gov/sites/default/release-files/j3860694x/vd66w276c/8c97kt28h/Agstat-04-23-1962.pdf	crawler	public-domain
Dreadnought battleships		
Wikipedia contributors (2026). Wikipedia: French battleship Courbet (revision 1374865778). https://en.wikipedia.org/w/index.php?title=French_battleship_Courbet&oldid=1374865778&action=raw	crawler	CC-BY-SA
Wikipedia contributors (2026). Wikipedia: HMS Dreadnought (1906) (revision 1374707991). https://en.wikipedia.org/w/index.php?title=HMS_Dreadnought_(1906)&oldid=1374707991&action=raw	crawler	CC-BY-SA
Wikipedia contributors (2025). Wikipedia: Italian battleship Dante Alighieri (revision 1323091977). https://en.wikipedia.org/w/index.php?title=Italian_battleship_Dante_Alighieri&oldid=1323091977&action=raw	crawler	CC-BY-SA
Wikipedia contributors (2026). Wikipedia: Japanese battleship Settsu (revision 1375023838). https://en.wikipedia.org/w/index.php?title=Japanese_battleship_Settsu&oldid=1375023838&action=raw	crawler	CC-BY-SA
Wikipedia contributors (2026). Wikipedia: Russian battleship Gangut (revision 1379253632). https://en.wikipedia.org/w/index.php?title=Russian_battleship_Gangut&oldid=1379253632&action=raw	crawler	CC-BY-SA
Wikipedia contributors (2026). Wikipedia: SMS Nassau (revision 1378184497). https://en.wikipedia.org/w/index.php?title=SMS_Nassau&oldid=1378184497&action=raw	crawler	CC-BY-SA
Wikipedia contributors (2026). Wikipedia: SMS Viribus Unitis (revision 1373687988). https://en.wikipedia.org/w/index.php?title=SMS_Viribus_Unitis&oldid=1373687988&action=raw	crawler	CC-BY-SA
Wikipedia contributors (2026). Wikipedia: USS Michigan (BB-27) (revision 1377049957). https://en.wikipedia.org/w/index.php?title=USS_Michigan_(BB-27)&oldid=1377049957&action=raw	crawler	CC-BY-SA
Armed drones, states		
Dan Gettinger (Center for the Study of the Drone at Bard College) (2019). The Drone Databook. https://dronecenter.bard.edu/files/2019/10/CSD-Drone-Databook-Web.pdf	crawler	copyrighted
Drone Wars UK (2022). Armed drones proliferation update - May 2022. https://dronewars.net/2022/05/09/armed-drones-proliferation-update-may-2022/	crawler	copyrighted
Drone Wars UK (2023). Armed drone proliferation update - March '23. https://dronewars.net/2023/03/22/armed-drone-proliferation-update-march-23/	crawler	copyrighted
Drone Wars UK (2025). Armed Drone Proliferation: Continued Exports Leading to Civilian Casualties. https://dronewars.net/2025/03/03/armed-drone-proliferation-continued-exports-leading-to-civilian-casualties/	crawler	copyrighted
Drone Wars UK (2025). Who has Armed Drones? (States operating armed MALE drones - updated March 2025). https://dronewars.net/who-has-armed-drones/	crawler	copyrighted
Drone Wars UK (2024). Who has Armed Drones? (Feb 2024 update; Wayback Machine capture 2024-04-10). https://web.archive.org/web/20240410090925id_/https://dronewars.net/who-has-armed-drones/	crawler	copyrighted
New America (2021). Countries that Have Drones (Datawrapper dataset, World of Drones). https://datawrapper.dwcdn.net/GwojV/9/dataset.csv	crawler	CC-BY

Source	Method	Licence
Wikipedia contributors (2026). Wikipedia: General Atomics MQ-1 Predator (revision 1376630226). https://en.wikipedia.org/w/index.php?title=General_Atomics_MQ-1_Predator&oldid=1376630226&action=raw	crawler	CC-BY-SA
Frequent-flyer programs Georgia newspapers, via the Rawson Company clipping service (1981). Clippings: 'EAL Plans Discounts' (Gwinnett Daily News) and 'Eastern To Reward Frequent Travelers' (Atlanta Daily World), clipping-service stamps 6 Sep 1981. https://cdm16905.contentdm.oclc.org/digital/bl/dmwebservices/index.php?q=dmGetItemInfo/eastern/51840/xml	crawler	copyrighted
Randy Petersen (InsideFlyer) (2006). The Big 2-5: Celebrating 25 Years of Frequent Flyer Programs. https://www.insideflyer.com/2006/05/12/the-big-2-5-celebrating-25-years-of-frequent-flyer-programs/	crawler	copyrighted
Randy Petersen (InsideFlyer) (2006). 60 Seconds with Mike Boynton, Delta Air Lines (retired). https://www.insideflyer.com/2006/04/25/60-seconds-with-mike-boynton-delta-air-lines-retired/	crawler	copyrighted
The New York Times (1982). United Schedules Its Boeing 767's (Company News), with correction of 30 Jul 1982. https://web.archive.org/web/20150524120028id_/https://www.nytimes.com/1982/07/29/business/united-schedules-its-boeing-767-s.html	crawler	copyrighted
Carole Shifrin (The Washington Post) (1981). Rewarding Frequent Flyers. https://www.washingtonpost.com/archive/lifestyle/travel/1981/10/25/rewarding-frequent-flyers/28d68e93-122c-4547-9bdb-79549b862223/	browser	copyrighted
Wikipedia contributors (2026). Wikipedia: AAdvantage (revision 1376666546). https://en.wikipedia.org/w/index.php?title=AAdvantage&oldid=1376666546&action=raw	crawler	CC-BY-SA
Generative AI, large organisations McKinsey & Company (2023). The state of AI in 2023: Generative AI's breakout year. https://web.archive.org/web/20260925170608id_/https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2023-generative-ais-breakout-year	crawler	copyrighted
McKinsey & Company (2024). The state of AI in early 2024: Gen AI adoption spikes and starts to generate value. https://web.archive.org/web/20260614200348id_/https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-2024	crawler	copyrighted
McKinsey & Company (2025). The state of AI: How organizations are rewiring to capture value. https://web.archive.org/web/20250416052125id_/https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai	crawler	copyrighted
McKinsey & Company (2025). The state of AI in 2025: Agents, innovation, and transformation (PDF). https://web.archive.org/web/20260408014729id_/https://www.mckinsey.com/~media/mckinsey/business%20functions/quantumblack/our%20insights/the%20state%20of%20ai/the-state-of-ai-in-2025-vf.pdf	crawler	copyrighted
McKinsey & Company (2026). The state of AI in 2026: On the road to ROI (PDF). https://web.archive.org/web/20260915005927id_/https://www.mckinsey.com/~media/mckinsey/business%20functions/quantumblack/our%20insights/the%20state%20of%20ai/the-state-of-ai-in-2026-on-the-road-to-roi-aug-2026.pdf	crawler	copyrighted
AI use, all US businesses U.S. Census Bureau (2025). Business Trends and Outlook Survey: AI Core Questions (Original), 09/11/2023 to 10/05/2025. https://www.census.gov/hfp/btos/downloads/AI%20Core%20Questions.xlsx	crawler	public-domain
U.S. Census Bureau (2026). Business Trends and Outlook Survey: National (current questionnaire). https://www.census.gov/hfp/btos/downloads/National.xlsx	crawler	public-domain
HTTPS page loads Google (2026). Google Transparency Report: Percentage of pages loaded over HTTPS in Chrome by platform (JSON). https://transparencyreport.google.com/transparencyreport/api/v3/httpsreport/chromeusage/byos?series=page-load&groupby=os	crawler	unknown
Internet Security Research Group (Let's Encrypt); data from Mozilla Firefox telemetry (2026). current-https-adoption.csv (Firefox telemetry, daily, by country group). https://d1dfn7jg27m4cf.cloudfront.net/current-https-adoption.csv	crawler	unknown
Internet Security Research Group (Let's Encrypt); data from Mozilla Firefox telemetry (2017). historical-https-adoption.csv (Firefox telemetry HTTP_PAGELOAD_IS_SSL, global, daily). https://d1dfn7jg27m4cf.cloudfront.net/historical-https-adoption.csv	crawler	unknown
Polyurethane swimsuits Lisa Dillman (Los Angeles Times) (2008). As swim records fall, high-tech suit faces scrutiny. https://www.latimes.com/archives/la-xpm-2008-mar-27-sp-swim27-story.html	crawler	copyrighted
Jason Tillotson (Swimming World) (2018). Eight Years Later, the Super Suit Era Still Plagues the Record Books. https://www.swimmingworldmagazine.com/news/eight-years-later-the-super-suit-era-still-plagues-the-record-books/	crawler	copyrighted

Source	Method	Licence
Wikipedia contributors (2026). Wikipedia: LZR Racer (revision 1379269770). https://en.wikipedia.org/w/index.php?title=LZR_Racer&oldid=1379269770&action=raw	crawler	CC-BY-SA
MLB analytics departments ESPN; MLB section by Ben Baumer (2015). The Great Analytics Rankings (MLB: In a Moneyball world, a number of teams remain slow to buy into sabermetrics). https://web.archive.org/web/20160806162303id_/http://www.espn.com:80/espn/feature/story/_/id/12331388/the-great-analytics-rankings	crawler	copyrighted
Ben Lindbergh and Rob Arthur (2016). Statheads Are The Best Free Agent Bargains In Baseball. https://web.archive.org/web/20160426203822id_/http://fivethirtyeight.com/features/statheads-are-the-best-free-agent-bargains-in-baseball/	crawler	copyrighted
New-energy vehicles, China Lei Kang, CnEVPost (2024). China's NEV retail reaches record 945,000 in Dec, 7.74 million for full year 2023. https://cnevpost.com/2024/01/09/china-nev-retail-reaches-945000-in-dec/	crawler	copyrighted
CnEVPost (2025). China NEV retail at record 1.38 million in Dec, preliminary CPCA data show. https://cnevpost.com/2025/01/08/china-nev-retail-dec-2024-preliminary-cpca-data/	crawler	copyrighted
CnEVPost (2026). China NEV retail sales rise 7% to 1.387 million units in Dec, preliminary CPCA data shows. https://cnevpost.com/2026/01/07/china-nev-retail-1-1-387-million-dec-2025-preliminary-cpca/	crawler	copyrighted
Our World in Data, from International Energy Agency, Global EV Outlook 2026 (2026). Share of new cars sold that are electric (grapher CSV). https://ourworldindata.org/grapher/electric-car-sales-share.csv?v=1&csvType=full&useColumnShortNames=true	crawler	CC-BY
Nuclear weapons Philipp C. Bleek (2017). When Did (and Didn't) States Proliferate? Chronicling the Spread of Nuclear Weapons. https://www.belfercenter.org/sites/default/files/pantheon_files/files/publication/When%20Did%20%28and%20Didn%27t%29%20States%20Proliferate%3F1.pdf	crawler	copyrighted
Wikipedia contributors (2026). Wikipedia: List of states with nuclear weapons (revision 1377913291). https://en.wikipedia.org/w/index.php?title=List_of_states_with_nuclear_weapons&oldid=1377913291&action=raw	crawler	CC-BY-SA
Walmart RFID mandate Associated Press (via TELECOM Digest) (2005). Wal-Mart: Radio Tags Keep Shelves Stocked. https://bio.nic.funet.fi/pub/mirrors/lcs.mit.edu/telecom-archives/TELECOM_Digest_Online2005-2/2288.html	crawler	copyrighted
Computerworld (2006). Wal-Mart details its RFID journey. https://www.computerworld.com/article/1720141/wal-mart-details-its-rfid-journey.html	crawler	copyrighted
Perry Trunick (Material Handling & Logistics) (2004). Wal-Mart rolls out RFID pilots. https://www.mhlnews.com/technology-automation/article/22033989/wal-mart-rolls-out-rfid-pilots	crawler	copyrighted
Supply Chain Digest (2009). RFID News: Looking Back at the Wal-Mart RFID Timeline. https://www.scdigest.com/assets/On_Target/09-02-23-1.php	crawler	copyrighted
Michael Garry (Supermarket News) (2005). Wal-Mart: 57 Suppliers Meet January RFID Mandate. https://www.supermarketnews.com/supplier-news/wal-mart-57-suppliers-meet-january-rfid-mandate	crawler	copyrighted
Michael Garry (Supermarket News) (2008). Plugging Along. https://www.supermarketnews.com/grocery-technology/plugging-along	crawler	copyrighted

Table 33: Provenance of the rebuilt data.

Case	Sources				Fitted table	Rows		Flags s/d/e/b
	All	Crawled	Captured	Manual		Parsed	Quoted	
Clap skates	5	5	0	0	7 values	0	7	1/4/1/1
Container	3	3	0	0	18 values	15	3	0/15/0/3
Hybrid corn	1	1	0	0	28 values	28	0	28/0/0/0
Dreadnought	10	10	0	0	8 events	8	0	0/8/0/0
Armed drones	11	11	0	0	48 events	13	52	0/48/0/17
Frequent flyer	11	10	1	0	8 events	1	7	0/8/0/0
Gen AI (large orgs)	5	5	0	0	10 values	0	10	4/0/1/5
AI (all US firms)	2	2	0	0	77 values	77	0	54/0/0/23
HTTPS	5	5	0	0	24 values	24	0	0/14/0/10
LZR suit	4	4	0	0	5 values	0	5	2/1/0/2
MLB analytics	7	7	0	0	5 values	0	5	1/2/1/1
China NEVs	5	5	0	0	15 values	11	4	11/0/0/4
Nuclear	3	3	0	0	11 events	6	5	0/11/0/0
RFID mandate	8	8	0	0	14 values	0	13	4/1/0/9