Over- and Under-Experimentation in a Patent Race with Private Learning

Kaustav Das* Nicolas Klein†

October 2026

Abstract

We study experimentation in a patent race with private learning and irreversible switching to a productive safe method. Research choices are observable, so continued experimentation can signal favorable information, while abandonment reveals its absence. For every fixed prior, all symmetric sequential equilibria converge to the same on-path behavior as decision intervals vanish. Relative to a planner observing both private signals, the symmetric outcome features premature immediate abandonment at some priors but excessive persistence on the no-news path at sufficiently optimistic priors. Near the private initiation cutoff, the unique symmetric outcome assigns positive probability to stopping both before and after the planner. Faster private learning increases the welfare loss in the immediate-abandonment region, although the width of that prior region is hump-shaped in the signal rate. A dominance argument establishes excessive no-news persistence at sufficiently optimistic priors in every equilibrium, including asymmetric equilibria.

Keywords: strategic experimentation, private learning, patent races, irreversible switching, sequential equilibrium.
JEL classification: C73, D83, O31.

1 Introduction

Firms choosing a research direction must assess both the scientific promise of a project and what their competitors have learned about it. A rival's continued investment may reflect favorable evidence that has not been disclosed. Abandonment may reveal that this evidence never arrived. These inferences are particularly consequential when abandoning a research direction is irreversible: information released by one firm's departure can change the other firm's decision, but cannot bring the departing firm back. Does competition then sustain unpromising research for too long, or terminate promising research too soon?

We study this question in a patent race between two firms. Each initially chooses between an innovative method of uncertain quality and an established method with a known invention rate. The innovative method is either productive for both firms or unproductive for both. A firm using it can receive a private signal establishing that it is productive. Inventions and research choices are public; signals are private. The first invention ends the race and earns its inventor the patent prize.

*University of Leicester; email: daskaustav84@gmail.com
†Université de Montréal and CIREQ; email: kleinnic@yahoo.com

A switch to the established method is irreversible, but the switching firm continues to compete for the invention. This productive fallback distinguishes abandoning a research direction from leaving the race.

A motivating example is the development of RNA interference (RNAi), which offers a way to treat diseases by silencing specific genes. Its commercial development faced substantial uncertainty, particularly about delivering RNA molecules into target cells (Bouchie, 2014). Firms could continue investing in this uncertain avenue or devote resources to more established approaches, such as small-molecule drugs and therapeutic antibodies. In November 2010, Roche announced the end of its internal RNAi research after substantial investment; Pfizer closed its RNAi drug-discovery programme in February 2011 (Zeliadt, 2014). Alnylam continued developing RNAi therapies, and its patisiran became the first FDA-approved RNAi therapeutic in August 2018 (FDA, 2018).

These contrasting choices illustrate how firms can differ in their willingness to persist with an uncertain technology. One interpretation consistent with our model is that persistence may reflect encouraging private experimental results, while abandonment may reveal their absence. The historical decisions need not have been driven by private information alone; the example motivates how observable research choices can convey information to competitors. Closing a specialised programme can also disperse expertise and infrastructure, making re-entry costly. Our irreversible-switching assumption captures this friction over the relevant research horizon.

We derive the planner's allocation and establish a unique symmetric equilibrium outcome as the interval between decisions vanishes. Symmetry means that the firms use the same contingent strategy after exchanging their identities; their stopping randomizations remain independent. Uniqueness holds for every fixed prior and set of research and discount rates. An uninformed firm may randomize over its switch to the established method, while a firm with favorable information continues innovative research. A rival's persistence therefore raises the probability that it has received good news, offsetting the pessimism caused by the absence of an invention.

The direction of the experimentation distortion depends on prior optimism. For an interval of priors, the planner starts innovative research while the symmetric equilibrium approaches immediate use of the established method. At sufficiently optimistic priors, every equilibrium keeps both firms experimenting longer than the planner would without news. Just above the private initiation cutoff, the symmetric equilibrium instead assigns positive probability to both early abandonment and excessive persistence. The forces behind these results differ. At moderate priors, each firm would like to learn from its rival's research choices, but neither captures the full value of the information generated by both laboratories. At very optimistic priors, firms expect research to continue long enough that a persistent rival has probably received favorable private evidence. An uninformed firm then remains optimistic even after a planner observing both laboratories would have abandoned research without news.

Private learning also has a welfare implication. In the region where the symmetric equilibrium abandons immediately, faster signals make coordinated experimentation more valuable and increase the welfare loss. The interval of priors exhibiting this inefficiency is nevertheless widest at an intermediate signal rate. Faster learning lowers both the efficient and private initiation cutoffs: it creates a larger gap when signals are scarce, but eventually makes experimentation worthwhile for both decision makers even at very low priors.

The mechanism combines information aggregation with competition for the prize. The planner can use either firm's signal to keep both firms on the productive method. A firm captures only its

own success probability and also cares about how the rival responds to its choices. Adopting the established method reveals the absence of favorable evidence but leaves the switching firm in the race. Its payoff therefore depends on the rival's subsequent research, while irreversibility prevents it from returning to the innovative method if the rival later receives good news.

1.1 Related literature

Our paper connects strategic experimentation to research races and irreversible stopping. Bolton and Harris (1999) and Keller et al. (2005) establish the interaction between learning spillovers and experimentation incentives when information is public. The latter characterize a unique symmetric Markov equilibrium as well as asymmetric equilibria. Hörner et al. (2022) show why the equilibrium restriction matters: with a diffusion component, efficient outcomes can be sustained beyond Markov strategies, including as limits of discrete-time equilibria. Our use of sequential equilibrium and an explicit vanishing decision interval keeps the information and commitment restrictions visible.

Private learning in stopping games is studied by Murto and Välimäki (2011) and Rosenberg et al. (2013). The latter characterize symmetric and asymmetric behavior when experimentation outcomes are private. Margaria (2020) studies irreversible investment with private signals about project quality and a second-mover advantage. Her unique symmetric equilibrium reduces to a stochastic war of attrition: waiting provides both information and a payoff advantage from investing second. In our race, abandoning the innovative method leaves the firm competing on the established method. Private signals are favorable, and the follower's experimentation option depends on the belief revealed by abandonment. Our focus is how these incentives produce premature abandonment or excessive persistence relative to coordinated research. Thomas (2020) studies privately observed successes when firms compete for a single shared safe option; here both firms can use the established method and continue competing for the terminal prize.

Guo and Roesler (2023) study a collaboration problem with private learning, unobservable effort, and irreversible public exit. Players who privately learn that the project is bad may stop working but delay exit to free-ride on an uninformed collaborator's effort, thereby delaying information transmission. Both their model and ours make continued participation informative about private information. Our firms receive conclusive good news, make observable research choices, and remain competitors after adopting the safe method. We study how these features shape experimentation relative to a planner, including premature abandonment and excessive persistence.

Heidhues et al. (2015) study communication with private experimentation payoffs and reversible choices. Achim and Klein (2022) study how initial private information affects investment and social learning. Dong (2025) shows that initial asymmetric information can increase learning and welfare through a mutual encouragement effect in observable experimentation. Here favorable private information is acquired during research, choices are irreversible, and continued competition after a switch determines the value of revealing that no good news has arrived.

Research-race antecedents include Moscarini and Squintani (2010), who show how private information can generate a survivor's curse and nearly simultaneous departures from a winner-take-all race. Here firms acquire conclusive private information while conducting research and abandon a method without abandoning the prize. Akcigit and Liu (2016) study irreversible choice between safe and risky research lines with unobserved research activities and private dead-end discoveries, showing both wasteful replication and premature abandonment. We instead study publicly observed choices, private favorable signals, and a single race-ending invention, which

make continued risky research informative about undisclosed optimism and allow a comparison of symmetric sequential-equilibrium behavior across all priors. Chatterjee and Evans (2004) study the allocation and duplication of research across projects. Das and Klein (2024) analyze established and innovative methods and show how patent protection affects clustered search. Chatterjee et al. (2026) study the disclosure of intermediate discoveries in a two-stage research race. Disclosing a discovery earns an immediate reward but lets the rival copy it and compete for the final prize. In our model, firms reveal information through observable research choices, and disclosure changes beliefs about a common technology. Our contribution is to establish a unique symmetric limiting outcome in this environment and to identify how prior optimism and the rate of private learning determine departures from efficient research.

Section 2 describes the environment and derives the planner's allocation. Section 3 characterizes symmetric behavior and its economic implications. Section 4 concludes, including the implications for asymmetric equilibria. The appendices contain the proofs and a robust asymmetric counterexample.

2 Research technology and the planner's benchmark

2.1 The environment

Two firms compete to produce an invention. The first inventor receives a prize of one, the other firm receives zero, and the race ends. Both discount at rate $r > 0$. A firm can use an established, *safe* research method, denoted S, or an innovative, *risky* method, denoted R. Safe research produces an invention at rate $\lambda_0 > 0$. The risky method has a common quality $\theta \in \{G, B\}$, with prior probability $p_0 \in (0, 1)$ of G. In the good state it produces inventions at rate $\lambda_2 > \lambda_0$; in the bad state it never produces an invention. Research has no additional flow cost. Its opportunity cost is the invention chance forgone on the alternative method.

While using risky research in the good state, each firm also receives a private, conclusive good-news signal at rate $\lambda_1 > 0$. There are no private signals in the bad state. The invention and signal processes are independent conditional on the state and research choices. A private signal establishes that the common risky method is productive but earns no prize. Final inventions and research choices are public; signals are observed only by their recipient. Thus private learning can influence the rival only through subsequent research choices.

The first switch to safe research is irreversible. A firm that makes this choice remains a competitor, producing inventions at rate λ_0 until the race ends. Returning to an abandoned research line may require rebuilding a specialized team or facilities. We represent a prohibitive cost of return by excluding a switch back to the risky method. It is a commitment restriction in the model: a firm cannot return to risky research even if its rival subsequently reveals favorable evidence. Both firms initially assign probability p_0 to the good state, with neither holding additional private information at the outset.

Firms choose simultaneously at dates $0, \eta, 2\eta, \ldots$ and hold their actions fixed between dates. Signals and inventions continue to arrive within each interval. Consequently a firm that chooses risky research cannot respond to the rival's simultaneous safe choice until the next date. This reaction interval also specifies the treatment of nearly simultaneous switches. We use sequential equilibrium on every positive grid and study its on-path implications as $\eta \downarrow 0$, holding $p_0, \lambda_0, \lambda_1, \lambda_2, r$ fixed. The firms randomize independently; there is no public correlating device.

A firm is *informed* if it has received its private signal and *uninformed* otherwise. An uninformed firm may nevertheless learn the state from the rival's action. If a firm has not already switched to safe research and knows that the state is good, risky research is strictly optimal. Indeed, staying risky guarantees at least $\lambda_2/(r+2\lambda_2)$, whereas switching safely yields at most $\lambda_0/(r+2\lambda_0)$, regardless of the rival's continuation. The former is strictly larger. All strategic stopping decisions can therefore be assigned to uninformed firms.

Our welfare benchmark is the sum of the firms' discounted prizes, equivalently the expected discounted value of the first invention. A centralized-information planner observes both firms' signals and chooses both research methods, subject to the same irreversibility restriction. It learns the state from evidence. Comparing decentralized behavior with this benchmark therefore captures both incentive effects and the loss of information aggregation.

2.2 Efficient experimentation and the value of joint learning

The planner maximizes the expected discounted value of the invention, irrespective of which firm makes it. It observes both firms' private signals and faces the same irreversible research choices as the firms. The planner initially has the same prior as the firms. The irreversibility restriction is nonbinding in its continuous-time optimum: the optimal policy would be unchanged if returning to the risky method were allowed.

Before either firm receives a signal, two risky firms generate two sources of information. Conditional on no invention and no signal, the planner's posterior at time t is

$$p(t) = \frac{p_0 e^{-2(\lambda_2+\lambda_1)t}}{1-p_0+p_0 e^{-2(\lambda_2+\lambda_1)t}}. \tag{1}$$

A signal from either firm establishes that the risky method is good. The planner can then keep both firms on risky research. This ability to use one firm's information for both research decisions is central to the welfare comparison.

Proposition 1 (Efficient experimentation)**.** *Define*

$$p^P = \frac{\lambda_0(r+2\lambda_2)}{\lambda_2(r+2\lambda_2)+2\lambda_1(\lambda_2-\lambda_0)}. \tag{2}$$

In the continuous-time benchmark, the planner assigns both firms to risky research when its posterior exceeds p^P and to safe research when it is weakly below p^P. After a good-news signal, both firms remain risky until the invention. Conditional on no invention and no signal, the switching date is

$$T^P(p_0) = \begin{cases} \dfrac{1}{2(\lambda_2+\lambda_1)} \ln\left(\dfrac{p_0(1-p^P)}{p^P(1-p_0)}\right), & p_0 > p^P, \\ 0, & p_0 \le p^P. \end{cases} \tag{3}$$

There is no interval during which the planner assigns only one firm to risky research. For fixed primitives and prior, every optimal reaction-grid policy has its first no-news safe-switch date converge to $T^P(p_0)$ as $\eta \downarrow 0$. On every sufficiently fine grid, the planner also switches both firms to safe research together.

Risky research has two social benefits: it can produce the invention directly, and it can reveal that both firms should continue with the more productive method. Its cost is the safe invention

opportunity forgone while the state remains uncertain. The planner's return to an additional risky firm is linear in the number already assigned to risky research, conditional on its belief and continuation value. Consequently it uses both risky firms while experimentation is worthwhile and switches both once its value falls to that of safe research. After a favorable signal, the planner uses the more productive method until the invention. If it instead switches both firms safely without a signal, neither firm generates further information about the risky method. There is consequently no later event that makes a return to risky research worthwhile.

It is useful to compare this cutoff with the decision of a firm whose rival has already switched irreversibly to safe research. Denote this solo cutoff by p^S; it is

$$p^S = \frac{\lambda_0(r + \lambda_0 + \lambda_2)}{\lambda_2(r + \lambda_0 + \lambda_2) + \lambda_1(\lambda_2 - \lambda_0)}. \tag{4}$$

The surviving firm experiments exactly when its posterior exceeds p^S. Its continuation problem is also efficient *conditional on the rival's irreversible safe switch.* Indeed, if the survivor's expected discounted prize is X, the safe firm's is $\lambda_0(1 - X)/(r + \lambda_0)$. Their sum is therefore increasing in X. An inefficient allocation can nevertheless result because the first exit has removed one firm from future risky research.

3 The symmetric equilibrium and its economic implications

3.1 A unique limiting prediction

To describe an uninformed firm's strategy before a first safe switch, draw a *planned switching date* at the start of the race. This is the date at which the firm adopts safe research if it receives no private signal and its rival has not switched first. A favorable signal makes it stay risky, while an earlier rival switch leads it to the optimal solo continuation. Drawing the two dates independently reproduces the firms' behavioral randomizations.

Our equilibrium selection result concerns the limit of short reaction intervals. It leaves the prior and all technological parameters fixed. Recall the solo cutoff p^S in (4).

Theorem 1 (The unique symmetric equilibrium in the small-grid limit)**.** *Fix $\lambda_2 > \lambda_0 > 0$, $\lambda_1, r > 0$, and $p_0 \in (0, 1)$. A symmetric sequential equilibrium exists for every reaction interval. As the interval tends to zero, every selection of symmetric equilibria induces the same limiting law of planned switching dates and the same limiting law of realized switching times.*

1. *If $p_0 \leq p^S$, both firms switch immediately in the limit.*
2. *If $p_0 > p^S$, the planned switching law assigns positive probability to stopping after any finite date, but zero probability to never stopping. Its only possible atom occurs at*

$$T^S = \frac{1}{2(\lambda_2 + \lambda_1)} \ln \frac{p_0(1 - p^S)}{p^S(1 - p_0)}.$$

At sufficiently late dates, uninformed firms mix continuously between continuing and switching.

For fixed primitives, different symmetric equilibria therefore cannot predict different limiting distributions of abandonment. When early and late stopping both occur, as shown below, they reflect randomization within this one outcome rather than equilibrium selection.

The cutoff p^S has a strategic interpretation. An initial switch reveals that its sender has not received good news. Below the solo cutoff this information leads an uninformed survivor to switch as well. Above the cutoff, the survivor retains a strictly valuable experimentation option. This follower advantage rules out a common switching atom before T^S: a firm would gain by waiting just beyond such an atom. At T^S the option value vanishes smoothly, so an atom need not create that profitable deviation. The proof treats this boundary separately.

Uniqueness is proved by noting that a firm places switching probability only on optimal dates, with indifference throughout any mixing interval. Every candidate equilibrium has the same late mixing behavior after an appropriate normalization. Working backward, any ordering between two candidate survival distributions must persist: reversing it would require one firm to switch at a date where switching is strictly inferior. A probability identity then excludes a persistent strict ordering. Appendix B.7 proves the result, allowing pooling to resume after mixing.

The *late experimentation phase* is the final interval, starting at a finite date, on which an uninformed firm switches with a continuous, strictly positive hazard while both firms remain risky. The theorem establishes that such a final interval exists, allowing alternating pooling and mixing phases before it. The next result shows that late mixing approaches a steady state with a constant belief and switching hazard.

Corollary 1 (Selection in the late experimentation phase)**.** *For $p_0 > p^S$, the switching hazard of an uninformed active firm tends to $2\lambda_2 + \lambda_1$ along the history on which both firms remain risky. Its posterior probability of the good state converges to*

$$p^\infty = \frac{r\lambda_0(r+\lambda_0+\lambda_2)(r+2\lambda_2)}{r\lambda_0(r+\lambda_0+\lambda_2)(r+2\lambda_2)+(r+2\lambda_0)(\lambda_2-\lambda_0)(r+\lambda_2)(r+2\lambda_2+\lambda_1)}. \tag{5}$$

Conditional on the good state and the same history, the probability that its rival is still uninformed tends to zero. The limiting posterior is increasing in λ_0 and r, and decreasing in λ_1 and λ_2.

Continued experimentation becomes informative because an uninformed rival would eventually switch. At late dates, a rival that still uses risky research is therefore almost surely informed conditional on the good state. This selection exactly offsets the adverse inference from continued silence: the absence of the two inventions and the firm's own signal lowers its likelihood ratio at rate $2\lambda_2 + \lambda_1$, while its rival's continued experimentation raises it at the same limiting rate. Each firm's mixing thus sustains the other's indifference: switching by uninformed rivals makes continued risky research increasingly suggestive of good news. In the limiting steady state, this inference balances the discouragement from continued silence. Beliefs remain bounded away from zero despite the absence of favorable evidence in the firm's own laboratory. A more productive safe method raises the opportunity cost of remaining risky, so late indifference requires a higher belief in the good state. Greater risky productivity lowers that required belief. A higher discount rate reduces the value of waiting for future information and also raises the belief required to continue. Faster private signals make an uninformed rival rarer among firms that persist in the good state. Indifference is restored at a lower private belief and a higher switching hazard. These are properties of the final mixing interval, whose starting date is itself endogenous.

3.2 Premature abandonment and welfare

The planner comparison has to distinguish three questions. Do firms abandon risky research before the planner would on the history without inventions or private signals? Do they use more risky

research on average, accounting for all signal histories? And how much surplus is lost? These questions need not have the same answer. The following results concern the first question and, where explicitly stated, the third. They do not infer an ordering of expected research exposure from a stopping-date comparison.

The strict ordering $p^P < p^S$ reflects the value of coordinated learning. The planner uses either firm's signal to keep both firms on the productive method. Once a firm has irreversibly adopted safe research, its rival's information can no longer change that allocation. This makes continued experimentation socially worthwhile at beliefs at which neither firm wants to undertake it alone against a safe rival.

Proposition 2 (Under-experimentation in the symmetric equilibrium)**.** *If $p^P < p_0 \leq p^S$, the unique symmetric equilibrium in the small-grid limit has both firms adopt safe research immediately. The planner instead starts both firms on risky research and keeps them there until a strictly positive date T^P, unless a signal or invention arrives. The equilibrium surplus is $2\lambda_0/(r + 2\lambda_0)$, strictly below the planner's value.*

At an interior prior $p^P < p_0 < p^S$, an increase in the private signal rate λ_1, holding the prior and all other primitives fixed, strictly increases both the no-news experimentation shortfall T^P and the welfare loss, locally while the prior remains in this interval. The welfare loss also increases strictly with the prior within this interval.

Faster learning does not automatically improve decentralized research. In this region neither firm starts experimenting, so faster private signals generate no equilibrium benefit. They do make joint experimentation more valuable to the planner, who can act on either signal. The resulting welfare loss and the amount of forgone no-news experimentation both rise. This is a statement about a fixed prior within the immediate-safe region; it does not claim monotonicity once the equilibrium changes regimes.

3.3 How private learning changes the experimentation wedge

Corollary 2 (The informational gap between the cutoffs)**.** *The solo cutoff strictly exceeds the planner's cutoff:*

$$p^S - p^P = \frac{\lambda_0\lambda_1(\lambda_2 - \lambda_0)(r + 2\lambda_0)}{\left[\lambda_2(r + \lambda_0 + \lambda_2) + \lambda_1(\lambda_2 - \lambda_0)\right]\left[\lambda_2(r + 2\lambda_2) + 2\lambda_1(\lambda_2 - \lambda_0)\right]} > 0. \tag{6}$$

Both cutoffs decrease with the private-signal rate λ_1 and increase with the discount rate r. Holding the other primitives fixed, their difference is strictly increasing in λ_1 up to

$$\lambda_1^* = \frac{\lambda_2}{\lambda_2 - \lambda_0}\sqrt{\frac{(r + \lambda_0 + \lambda_2)(r + 2\lambda_2)}{2}}, \tag{7}$$

and strictly decreasing thereafter. The difference converges to zero as $\lambda_1 \downarrow 0$ and as $\lambda_1 \to \infty$.

Faster signals make experimentation worthwhile at lower beliefs for both the planner and the solo researcher. The distance between their cutoffs, however, need not increase. When signals are very rare, learning contributes little and both cutoffs approach the myopic threshold λ_0/λ_2. When signals are very frequent, both cutoffs approach zero. The difference is largest at an intermediate signal rate. This comparative static concerns the interval of priors between the cutoffs; it is not a monotonicity claim about an equilibrium's total research effort or welfare loss.

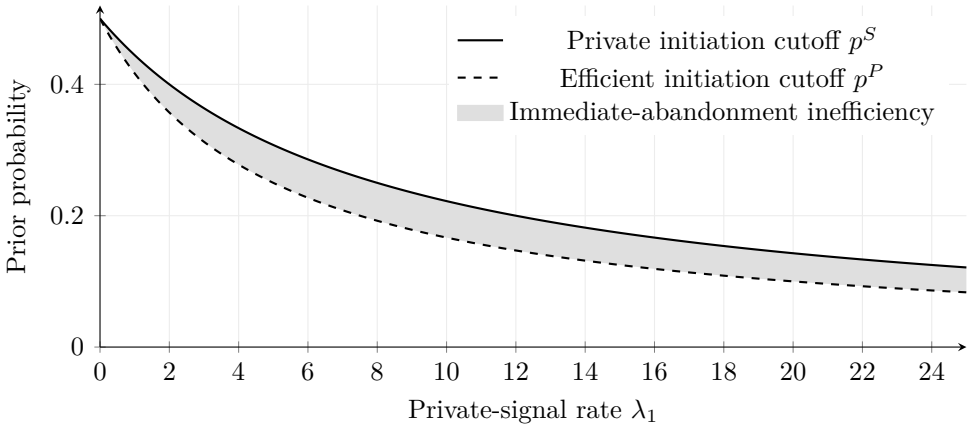


Figure 1: Private learning and the interval of immediate-abandonment under-experimentation. Parameters are $\lambda_0 = 1$, $\lambda_2 = 2$ and $r = 1$. In the shaded interval the planner starts risky research while the unique symmetric limit is immediately safe. Both cutoffs decrease; their vertical distance is largest at $\lambda_1 = \sqrt{40}$. The curves are analytical, not simulated equilibrium approximations.

The distinction between beliefs and calendar time also matters. At a fixed prior $p_0 > p^P$, differentiation of (3) gives

$$\frac{\partial T^P}{\partial \lambda_1} = \frac{1}{\lambda_2 + \lambda_1} \left\{ \frac{1}{r + 2\lambda_2 + 2\lambda_1} - T^P \right\}, \tag{8}$$

$$\frac{\partial T^P}{\partial r} = -\frac{\lambda_1}{(\lambda_2 + \lambda_1)(r + 2\lambda_2)(r + 2\lambda_2 + 2\lambda_1)} < 0. \tag{9}$$

Thus faster signals lower the efficient stopping belief but can shorten the no-news experimentation period: the planner also learns more rapidly from their absence. A higher discount rate reduces the value of favorable evidence that arrives later. The planner therefore requires a higher current probability of a productive risky method and stops sooner without news. For the signal-rate comparison, two effects compete: better prospects of favorable news lower the stopping cutoff, while each unit of silence becomes stronger adverse evidence. The displayed derivative identifies which effect dominates at a given prior.

3.4 Early and late stopping at more optimistic priors

For a comparison of cumulative research on the no-news path, perform the following thought experiment. Draw the firms' planned dates from the equilibrium distribution, then let no private signal or invention arrive. The firms follow their plans and, after the first safe switch, the survivor follows its solo policy. Total risky research is the sum of the two safe-adoption dates, measured in firm-time. Its expectation averages over the stopping draws only. The planner's corresponding quantity is $2T^P$. This separates the stopping randomization from the arrival of news, which can cancel a plan or end the race.

Proposition 3 (A symmetric equilibrium can stop both too early and too late)**.** *For every admissible vector of research rates and discount rate, there is an $\varepsilon > 0$ such that, for every $p_0 \in (p^S, p^S + \varepsilon)$, the unique symmetric limiting equilibrium has both of the following properties on histories without an invention or a private signal up to the planner's stopping date:*

1. *with positive probability both firms adopt safe research strictly before the planner's stopping date;*

2. *with positive probability both firms remain on risky research strictly beyond that date.*

The probabilities are over firms' independent randomization. Each event also has positive probability jointly with the stipulated absence of news. Thus neither allocation dominates the other in risky research at every date and every realization of the randomization.

Nevertheless, expected total risky firm-time in the thought experiment just defined is strictly below the planner's $2T^P$.

The shortfall in average research has an informational free-riding interpretation. For beliefs below the myopic cutoff λ_0/λ_2, each firm would like its rival to bear the opportunity cost of testing the risky method and to reveal what it has learned through its research choice. While either firm values the information only for its own research choice, the planner instead values the information from either laboratory for *both* research decisions. Just above p^S, this difference leaves most equilibrium stopping close to date zero, even though coordinated experimentation remains worthwhile for a strictly positive time. Some firms persist much longer: a firm thinking that a rival that continues may have received good news sustains research after the planner would stop on the no-news path. Near the cutoff, these long delays are too rare to offset early abandonment in average risky firm-time. Expected realized research also depends on favorable signals and on when an invention ends the race.

At sufficiently optimistic priors, there is a stronger result that extends to every equilibrium, including asymmetric ones. A firm cannot abandon risky research while continuing forever gives a higher payoff than even the most favorable possible safe continuation. This dominance argument forces both firms to pool for a sufficiently long initial interval.

Proposition 4 (Optimistic priors force excessive no-news experimentation)**.** *Fix $\lambda_2 > \lambda_0 > 0$ and $r > 0$. For every $\lambda_1 > 0$, there exists a threshold $\bar{p}_0(\lambda_1) \in (p^P, 1)$ such that, for every fixed prior $p_0 > \bar{p}_0(\lambda_1)$, every sequential equilibrium, symmetric or asymmetric, keeps both firms on risky research strictly beyond the planner's stopping date on the history without inventions or private signals. This strict ordering holds relative to the planner optimum on the same grid for all sufficiently small $\eta > 0$, and the delay remains bounded away from zero as $\eta \downarrow 0$. A guaranteed lower bound on the limiting delay increases strictly with the prior.*

An optimistic prior gives private good news time to accumulate. As the prior approaches one, even a planner receiving no news experiments for an increasingly long time. By that date, an uninformed firm is nearly certain that a rival that has remained risky has received a favorable signal. The firm knows that it has received no signal itself, but cannot observe the absence of a signal in the other laboratory. It therefore remains optimistic about the common technology and continues researching. The planner, by contrast, knows that neither laboratory has received good news and switches both firms safely. This difference forces excessive no-news persistence for every positive signal rate, however small, provided that the prior is sufficiently optimistic.

The appendix gives an explicit sufficient threshold and a lower bound on the delay, using payoff bounds that apply to every equilibrium.

4 Conclusion

Private learning and irreversible research choices generate both premature abandonment and excessive persistence in a patent race. The symmetric equilibrium has a unique limiting on-path outcome for every prior. For some priors firms abandon immediately even though coordinated experimentation is worthwhile; in this region faster private learning increases the welfare loss. At sufficiently optimistic priors, firms instead persist beyond the planner's stopping date on the path without inventions or private signals. Near the private initiation cutoff, early and late abandonment both occur with positive probability, while average risky firm-time in the no-news thought experiment is below the planner's.

The two distortions reflect different consequences of private learning. The planner uses either firm's signal to direct both research choices, whereas firms would like to learn from their rival's actions and do not capture the full informational return to coordinated experimentation. This informational externality can cause premature abandonment. With sufficiently optimistic priors, however, research lasts long enough that each uninformed firm becomes nearly certain that its persistent rival has received good news. Both then remain willing to experiment after a planner observing no news in either laboratory would stop. Irreversibility makes the first safe switch consequential: it removes one laboratory from risky research even if later information would have justified its return.

Two conclusions extend beyond symmetry. Proposition 4 forces excessive no-news persistence in every equilibrium at sufficiently optimistic priors. Positive initial pooling alone has no such implication. Appendix C.3 constructs an asymmetric equilibrium with positive pooling in which, on the no-news path, both firms switch safely before the planner. The example persists on sufficiently fine grids and under small perturbations of the primitives. Its continuation is efficient conditional on the first switch, but that switch removes one source of research and learning too soon. Thus initial pooling can coexist with insufficient joint experimentation.

References

Achim, P. and Klein, N. (2022). Strategic investment and learning with private information. *Journal of Economic Theory*, 204:105523.

Akcigit, U. and Liu, Q. (2016). The role of information in innovation and competition. *Journal of the European Economic Association*, 14(4):828–870.

Bolton, P. and Harris, C. (1999). Strategic experimentation. *Econometrica*, 67(2):349–374.

Bouchie, A. (2014). Markets, venture investors and big pharma interest in RNAi soars. *Nature Biotechnology*, 32:203–204. doi:10.1038/nbt0314-203.

Chatterjee, K., Das, K., Dong, M., and Klein, N. (2026). Strategic disclosure in research races. Working paper, SSRN 4803642.

Chatterjee, K. and Evans, R. (2004). Rivals' search for buried treasure: Competition and duplication in R&D. *RAND Journal of Economics*, 35(1):160–183.

Das, K. and Klein, N. (2024). Do stronger patents lead to faster innovation? The effect of clustered search. *International Economic Review*, 65(2):915–954.

Dong, M. (2025). Strategic experimentation with asymmetric information. Working paper, SSRN 5257187.

U.S. Food and Drug Administration (2018). FDA approves first-of-its kind targeted RNA-based therapy to treat a rare disease. Press release, August 10. Available via PR Newswire.

Guo, Y. and Roesler, A.-K. (2023). Private learning and exit decisions in collaboration. Working paper, November 16.

Heidhues, P., Rady, S., and Strack, P. (2015). Strategic experimentation with private payoffs. *Journal of Economic Theory*, 159:531–551.

Hörner, J., Klein, N., and Rady, S. (2022). Overcoming free-riding in bandit games. *Review of Economic Studies*, 89(4):1948–1992.

Keller, G., Rady, S., and Cripps, M. (2005). Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68.

Margaria, C. (2020). Learning and payoff externalities in an investment game. *Games and Economic Behavior*, 119:234–250.

Moscarini, G. and Squintani, F. (2010). Competitive experimentation with private information: The survivor's curse. *Journal of Economic Theory*, 145(2):639–660.

Murto, P. and Välimäki, J. (2011). Learning and information aggregation in an exit game. *Review of Economic Studies*, 78(4):1426–1461.

Rosenberg, D., Salomon, A., and Vieille, N. (2013). On games of strategic experimentation. *Games and Economic Behavior*, 82:31–51.

Thomas, C. (2020). Stopping with congestion and private payoffs. *Journal of Mathematical Economics*, 91:18–42.

Zeliadt, N. (2014). Big pharma shows signs of renewed interest in RNAi drugs. *Nature Medicine*, 20:109. doi:10.1038/nm0214-109.

A The centralized benchmark

A.1 The planner and the solo continuation

Proof of Proposition 1. First consider the policy that assigns both firms to risky research until a planned no-news date t, switches both to safe research at that date if neither signal has arrived,

and keeps both risky forever if a signal arrives first. Its value at prior p is

$$W_t(p) = p\left\{\frac{2\lambda_2}{r+2\lambda_2} + \left(\frac{2\lambda_0}{r+2\lambda_0} - \frac{2\lambda_2}{r+2\lambda_2}\right)e^{-(r+2\lambda_2+2\lambda_1)t}\right\} + (1-p)\frac{2\lambda_0}{r+2\lambda_0}e^{-rt}. \tag{10}$$

The derivative with respect to t has the sign of

$$\frac{p}{1-p}e^{-2(\lambda_2+\lambda_1)t} - \frac{\lambda_0(r+2\lambda_2)}{(\lambda_2-\lambda_0)(r+2\lambda_2+2\lambda_1)}.$$

It follows that the maximizing date is uniquely $T^P(p)$ in (3), including the date zero when $p \le p^P$. Define $W(p) = W_{T^P(p)}(p)$. Then $W(p) = 2\lambda_0/(r+2\lambda_0)$ below the cutoff, is strictly larger above it, and is continuously differentiable across the cutoff with $W'(p^P) = 0$.

To verify optimality over *all* allocations, temporarily permit reversible switching. With $k \in \{0,1,2\}$ risky firms, the aggregate invention rate is $2\lambda_0 + k(p\lambda_2 - \lambda_0)$, the good-news rate is $kp\lambda_1$, and the no-news posterior drift is $-k(\lambda_2+\lambda_1)p(1-p)$. The candidate therefore must satisfy

$$rW(p) = \max_{k\in\{0,1,2\}}\Bigg[2\lambda_0[1-W(p)] + k\Bigg\{(p\lambda_2-\lambda_0)[1-W(p)] + p\lambda_1\left[\frac{2\lambda_2}{r+2\lambda_2} - W(p)\right] - (\lambda_2+\lambda_1)p(1-p)W'(p)\Bigg\}\Bigg]. \tag{11}$$

For $p > p^P$, the $k=2$ equation follows by differentiating (10) along the candidate policy. The coefficient of k is then

$$\frac{rW(p) - 2\lambda_0[1-W(p)]}{2} = \frac{(r+2\lambda_0)W(p) - 2\lambda_0}{2} > 0.$$

Thus $k=2$ attains the maximum. For $p \le p^P$, substitute the constant value and its zero derivative into (11). The coefficient of k is nonpositive exactly when $p \le p^P$, so $k=0$ attains the maximum. At $p=1$, assigning both firms to the more productive risky method attains $W(1) = 2\lambda_2/(r+2\lambda_2)$.

The usual discounted verification can be stated directly. Apply the piecewise-deterministic change-of-variables formula to $e^{-rt}W(p_t)$ before the first invention, including the jump to belief one after a signal. Equation (11) implies that expected discounted invention rewards plus the candidate continuation value form a supermartingale under any admissible allocation. They form a martingale under the stated policy. The bounded continuation term vanishes at an infinite horizon because $r > 0$. This proves optimality in the reversible relaxation. The optimal policy makes at most one switch, from both risky to both safe, and hence is feasible under irreversibility. Strictness of the maximizing allocation on either side of p^P also rules out an interval with only one risky firm. The instantaneous tie at p^P does not produce a different allocation on any interval of positive length: continued risky research would immediately lower the no-news belief below the cutoff.

It remains to connect the benchmark to the reaction-grid model. Before a signal or first safe choice there is a single no-news history. A deterministic planner policy is consequently characterized by its first no-news safe date $t \in \eta\mathbb{N}$ and by whether one or both firms then choose safe research. Signals before t make continued joint risky research optimal. If both firms switch at t, the exact grid value is already (10). If only one switches, the constant-hazard identity in Lemma 2 shows

that maximizing total continuation value is equivalent to maximizing the risky survivor's own value, subject to its compulsory current risky interval.

For a survivor with belief p, the continuous optimum has planned duration

$$\frac{1}{\lambda_2+\lambda_1}\left[\ln\left(\frac{p(1-p^S)}{p^S(1-p)}\right)\right]_+, \tag{12}$$

where $[x]_+ = \max\{x, 0\}$. Its optimal forced grid duration differs from this duration by at most η; the same bound applies to both adjacent choices at a grid indifference. To see this, the grid stopping threshold in posterior odds is

$$\frac{\dfrac{\lambda_0}{r+2\lambda_0}(1-e^{-(r+\lambda_0)\eta})}{\left(\dfrac{\lambda_2}{r+\lambda_0+\lambda_2}-\dfrac{\lambda_0}{r+2\lambda_0}\right)(1-e^{-(r+\lambda_0+\lambda_2+\lambda_1)\eta})}.$$

It lies strictly between $p^S/(1-p^S)$ and $e^{(\lambda_2+\lambda_1)\eta}p^S/(1-p^S)$. Successive grid differences have the sign of current posterior odds minus this threshold. Rounding the crossing date, and imposing at least one interval when necessary, gives the asserted duration bound.

The survivor's statewise values in (14) are uniformly Lipschitz in its planned duration, including the discounted infinite tail. Consequently both possible first-safe-date objectives converge uniformly to their continuous-time counterparts, on the compactified time domain $[0, \infty]$. Randomization only forms convex combinations of these objectives. Continuous-time verification gives the unique first no-news safe date $T^P(p_0)$; any strictly earlier irreversible exit loses value where the maximizing coefficient is strictly positive, and any strictly later exit loses value where it is strictly negative. The value at $t = \infty$ is also strictly below the optimum for an interior prior. Uniform convergence and the strict separation of the continuous maximum outside each neighborhood of $T^P(p_0)$ imply that every optimal grid first-safe date converges to $T^P(p_0)$. Finally, the no-news belief at that date converges to p^P when $p_0 > p^P$, and to $p_0 \le p^P$ otherwise. The strict inequality $p^P < p^S$ and convergence of the solo grid cutoff imply that, at every optimal first-safe date on every sufficiently fine grid, a survivor compelled to remain risky for one interval obtains strictly less than $\lambda_0/(r+2\lambda_0)$. By the constant-hazard identity, assigning one firm to safe and the other to risky research at that date is strictly worse than assigning both to safe research. Thus every optimal sufficiently fine-grid policy switches both together. □

Corollary 3 (The planner on sufficiently fine reaction grids)**.** *For each fixed interior prior and fixed primitives, on every sufficiently fine grid an optimal planner keeps both firms risky until the no-news posterior first becomes weakly smaller than*

$$p^P_\eta = \frac{\dfrac{\lambda_0}{r+2\lambda_0}(1-e^{-r\eta})}{\dfrac{\lambda_0}{r+2\lambda_0}(1-e^{-r\eta})+\left(\dfrac{\lambda_2}{r+2\lambda_2}-\dfrac{\lambda_0}{r+2\lambda_0}\right)(1-e^{-(r+2\lambda_2+2\lambda_1)\eta})}. \tag{13}$$

It then switches both firms to safe research. If the posterior equals p^P_η at a decision date, stopping then and stopping at the next date are both optimal, as are randomizations between these dates. These are the only no-news stopping-date multiplicities. Moreover, $p^P_\eta > p^P$ and $p^P_\eta \to p^P$.

Proof. Proposition 1 has established that every optimal sufficiently fine-grid policy switches both firms together. The difference between the value of this switch at $(n+1)\eta$ and its value at $n\eta$ is obtained directly from (10). Its sign is that of the no-news posterior at $n\eta$ minus p^P_η. That posterior strictly decreases with n while both firms remain risky. Hence the maximizing dates and the possible single adjacent-date tie are exactly as stated. The function $(1-e^{-x})/x$ is strictly decreasing for $x>0$, which gives $p^P_\eta > p^P$ after comparing the two cutoff odds. Taking $\eta \downarrow 0$ gives convergence. □

Proof of Corollary 2. Subtracting the two cutoff expressions gives (6). Each denominator is positive, so the gap is strictly positive. Direct differentiation gives

$$\begin{aligned}
\frac{\partial p^P}{\partial \lambda_1} &= -\frac{2\lambda_0(r+2\lambda_2)(\lambda_2-\lambda_0)}{[\lambda_2(r+2\lambda_2)+2\lambda_1(\lambda_2-\lambda_0)]^2} < 0,\\
\frac{\partial p^S}{\partial \lambda_1} &= -\frac{\lambda_0(r+\lambda_0+\lambda_2)(\lambda_2-\lambda_0)}{[\lambda_2(r+\lambda_0+\lambda_2)+\lambda_1(\lambda_2-\lambda_0)]^2} < 0,\\
\frac{\partial p^P}{\partial r} &= \frac{2\lambda_0\lambda_1(\lambda_2-\lambda_0)}{[\lambda_2(r+2\lambda_2)+2\lambda_1(\lambda_2-\lambda_0)]^2} > 0,\\
\frac{\partial p^S}{\partial r} &= \frac{\lambda_0\lambda_1(\lambda_2-\lambda_0)}{[\lambda_2(r+\lambda_0+\lambda_2)+\lambda_1(\lambda_2-\lambda_0)]^2} > 0.
\end{aligned}$$

The derivative of the gap with respect to λ_1 has the sign of

$$\lambda_2^2(r+\lambda_0+\lambda_2)(r+2\lambda_2) - 2\lambda_1^2(\lambda_2-\lambda_0)^2.$$

This gives the unique maximum in (7) and its monotonicity on either side. The limits follow from (6). Finally, differentiating the logarithm in (3) proves (8)–(9). □

B Proofs for the symmetric equilibrium

B.1 Continuation values and beliefs

In the proofs, an *exit* means irreversible adoption of safe research; the firm remains in the patent race. The proofs use five discounted prize values under fixed research allocations: s for two safe firms, g for two risky firms in the good state, v for a safe firm facing a risky rival in the good state, w for the reverse allocation, and z for a safe firm facing an unproductive rival. Thus

$$s=\frac{\lambda_0}{r+2\lambda_0}, \qquad g=\frac{\lambda_2}{r+2\lambda_2}, \qquad v=\frac{\lambda_0}{r+\lambda_0+\lambda_2}, \qquad w=\frac{\lambda_2}{r+\lambda_0+\lambda_2}, \qquad z=\frac{\lambda_0}{r+\lambda_0},$$
$$a=\lambda_2+\lambda_1, \quad u=e^{-\lambda_1\eta}, \quad A=e^{-(r+\lambda_0+\lambda_2+\lambda_1)\eta}, \quad B=e^{-(r+\lambda_0)\eta},$$
$$O_0=\frac{p_0}{1-p_0}, \qquad Q_n=O_0e^{-2an\eta}.$$

Lemma 1 (Research after conclusive information)**.** *At every history at which a firm has not adopted safe research and knows the good state, risky research is strictly optimal.*

Proof. Risky research forever guarantees at least g, since the rival's invention rate is at most λ_2. Irreversible safe research yields at most s, since the rival's rate is at least λ_0 in the good state. The inequality $g>s$ is strict. □

B.2 Continuation after a first safe switch

Write U for a firm that has received no private signal and I for a firm that has received its conclusive favorable signal. Suppose one firm is already safe. An informed survivor uses R forever and obtains w. If the survivor is U with odds q of G, a policy that uses R for k intervals unless a signal arrives, and then chooses S if still U, has conditional-state values

$$M_k^G = w + (s-w)A^k, \qquad M_k^B = sB^k, \qquad V_k(q) = \frac{qM_k^G + M_k^B}{1+q}. \tag{14}$$

Here $k = 0$ is immediate S, and $k = \infty$ means R forever, with $(M_\infty^G, M_\infty^B) = (w, 0)$. These expressions include possible signals before the planned stopping date.

Lemma 2 (Solo cutoff and a payoff identity)**.** *Define*

$$q_S = \frac{s(1-B)}{(w-s)(1-A)}. \tag{15}$$

A free-to-choose U survivor chooses R if $q > q_S$, chooses S if $q < q_S$, and can mix between them at equality. If it continues at equality and receives no signal, it exits at the following date. If its current action has already been fixed at R, its optimal value pair is any convex combination of the pairs in (14) *maximizing $V_k(q)$ over $k \geq 1$; there are at most two adjacent maximizing integers for finite q.*

In continuous time the cutoff odds are $Q^S = \lim_{\eta \downarrow 0} q_S = p^S/(1-p^S)$, and the optimal uninformed risky duration from odds q is $a^{-1}[\ln(q/Q^S)]_+$.

For any survivor policy, let X be its discounted prize probability and Y that of its safe rival, evaluated under the same distribution of the state. Then

$$Y = z(1-X). \tag{16}$$

This identity also holds conditional on either state.

Proof. After an unsuccessful risky interval without a signal, survivor odds become $qe^{-a\eta}$. Direct subtraction in (14) shows that $V_{k+1}(q) - V_k(q)$ has the sign of $qe^{-ak\eta} - q_S$. Thus the maximizing stopping date is the first date at which odds are weakly below q_S, with the stated tie. Before a signal the private history has no other random observation; behavioral randomization only randomizes this stopping date. After a signal, Lemma 1 applies. This verifies optimality against all policies, not just deterministic stopping dates. In continuous time, the corresponding statewise values are $w + (s-w)e^{-(r+\lambda_0+\lambda_2+\lambda_1)t}$ and $se^{-(r+\lambda_0)t}$. The derivative of their posterior-weighted average has the sign of $qe^{-at} - Q^S$, proving the continuous cutoff and duration as well.

For (16), let $D(t)$ be the probability that neither firm has invented by time t, under the chosen policy. The safe firm's constant hazard gives $Y = \lambda_0 \int_0^\infty e^{-rt} D(t)\, dt$. Accounting for discounting and the two possible inventions gives $X + (r+\lambda_0)\int_0^\infty e^{-rt} D(t)\, dt = 1$. Elimination of the integral proves the identity. The same calculation applies state by state. □

If the first safe action occurs on path at date n, its sender must be U. A surviving U firm's odds immediately after observing it are Q_n, irrespective of earlier mixing. The survivor has already chosen R for the current interval, so it must solve the constrained problem $k \geq 1$, not the problem $k \geq 0$. This reaction lag matters for the equilibrium inequalities below.

B.3 Beliefs and disclosure at a first switch

Before a first safe switch, let $x_i(n)$ be the probability that firm i switches at date $n\eta$ if it has received no signal. Its private history on this branch consists only of the absence of a signal. Conditional on G and the public history of joint risky choices, let $m_i(n)$ be the probability that firm i is still uninformed. Then

$$m_i(n+1) = u\frac{m_i(n)[1-x_i(n)]}{1-m_i(n)x_i(n)}, \tag{17}$$

$$O_i(n) = \frac{Q_n}{m_j(n)}, \quad j \neq i. \tag{18}$$

Here $O_i(n)$ is an uninformed firm's posterior odds. If $m_j(n) = 0$, these odds are infinite. Initially $m_i(0) = 1$; in the exceptional case of certain safe adoption at date zero, an unexpected risky choice gives $m_i(1) = u$, because no private signal was possible at the initial decision. For $n \geq 1$, $m_i(n) \leq u < 1$. The first formula removes the uninformed firms that switch and then accounts for new signals. Conditional on both firms having no signal, their past action likelihoods cancel between states, giving Q_n; dividing by the rival's conditional-G uninformed probability gives the second formula. An on-path safe switch reveals that its sender has received no signal, so an uninformed survivor's odds become Q_n.

The following technical lemma permits this interpretation also at unexpected safe choices in the timing-game representation. It is used only where both initial switching probabilities are below one.

Lemma 3 (Disclosure of missing good news)**.** *Any sequential equilibrium with $x_1(0) < 1$ and $x_2(0) < 1$ has the same on-path behavior as a sequential equilibrium in which every first safe choice, including an off-path choice, reveals that its sender has received no private signal. After that choice, a still uninformed survivor optimizes its solo continuation at odds Q_n, including its compulsory current risky interval.*

Proof. First consider an actually uninformed sender and an uninformed survivor at date n. Their true conditional odds are Q_n. A survivor choosing an optimal policy at these odds maximizes its payoff under that distribution. The identity (16) consequently minimizes the sender's expected payoff from switching safely. An informed survivor always uses risky research. Thus assigning the posterior Q_n to every unexpected first safe choice weakly lowers the uninformed sender's stopping reward, while preserving on-path stopping rewards. Solo cutoff ties do not affect that sender's expected payoff.

Until a first certain uninformed switch, uninformed histories on the joint-risky branch have positive probability in B, so lowering unexpected stopping rewards preserves their incentives. A first certain switch cannot occur at date zero under the lemma's hypothesis. If it is prescribed at a positive date, observing the sender remain risky reveals G to its rival. When just one firm was prescribed to switch, its own uninformed history after disobeying that prescription is off path and can retain uncertainty. Give that sender an optimal remaining stopping policy against a rival that stays risky until a first safe choice. At such a choice, a rival without a private signal updates to Q_n and follows the constrained solo optimum. These are limiting beliefs at zero-probability histories, not additional conclusive public signals. Each stopping reward has weakly fallen, so its optimized deviation payoff cannot rise. If both firms were prescribed to switch, a firm observing its rival remain risky also knows G, and risky research is optimal. These continuations are sequentially rational and do not weaken any earlier incentive.

For consistency, use one common sequence of fully mixed strategies. For $\delta \downarrow 0$, replace each uninformed safe probability by $(1-\delta)x_i(n) + \delta/2$. Recursively calculate the sender's perturbed conditional-G uninformed probability m_δ along the public history. At a limiting zero-probability first safe choice and $n \geq 1$, put

$$a_\delta = m_\delta[(1-\delta)x_i(n) + \delta/2], \qquad s_{I,\delta} = \frac{\delta a_\delta}{1 - m_\delta},$$

where $s_{I,\delta}$ is an informed sender's safe tremble. Cap it at $1/2$ if necessary. Since $m_\delta \leq u < 1$ and $a_\delta \to 0$ at such a history, the cap eventually does not bind, and the conditional probability that a safe sender is uninformed is $1/(1+\delta) \to 1$. At a positive-probability first safe choice, an informed safe tremble tending to zero gives the same disclosure. At date zero only an uninformed sender is feasible. Perturb all solo choices with full support as well. Recursion at every finite date defines these trembles jointly, ensuring consistency across histories. The strict informed choice and the prescribed optimal uninformed continuations are therefore supported by a consistent assessment. One-stage optimality implies full sequential rationality: finite deviations can be eliminated backward, and rewards after L additional intervals are uniformly bounded by $e^{-rL\eta}$. □

B.4 Symmetric existence

Theorem 2 (Existence of a symmetric assessment)**.** *For every $\lambda_2 > \lambda_0 > 0$, $\lambda_1, r, \eta > 0$ and $p_0 \in (0,1)$, the irreversible-switching game admits a symmetric sequential equilibrium.*

Proof. We construct a symmetric assessment directly. Fix a finite number K of intervals and award zero if no invention has occurred by $K\eta$. Fix $\varepsilon \in (0, 1/2)$. At each informed decision prescribe safe-action probability ε, independently of the date of the signal. At each feasible uninformed decision constrain its safe-action probability to $[\varepsilon, 1-\varepsilon]$.

This auxiliary problem has finitely many relevant information sets. Before an invention, a signal received within an interval affects no action until the next date. Its occurrence can therefore be represented by a binary informedness flag at decision dates. Its exact arrival time has no effect on the payoffs or on the prescribed informed tremble. An active uninformed firm has no signal history and remembers the unique public action history and its past risky choices. All feasible uninformed information sets have positive probability under the constrained strategies; their Bayesian beliefs and action payoffs are continuous functions of these strategies.

Use one coordinate for each pair of uninformed information sets that are interchanged by swapping the firms. The resulting symmetric strategy space is a finite product of intervals. At a symmetric profile, consider an individual decision agent's best response, holding every other decision rule, including its counterpart's, fixed. Its expected payoff is affine in its own current safe-action probability. Its maximizing set in the allowed interval is nonempty, compact and convex, and varies upper hemicontinuously with the profile. The counterpart agent has exactly the same maximizing set, by symmetry. The product of these sets therefore defines a correspondence from the symmetric strategy space into itself satisfying the hypotheses of Kakutani's fixed-point theorem. At a fixed point every uninformed decision agent is optimal. This step constructs symmetric behavioral strategies directly; it is not a lottery over asymmetric equilibria.

Let $K \to \infty$ and $\varepsilon \to 0$. Extend each finite profile beyond its terminal date by any symmetric, completely mixed strategy. There are only countably many relevant finite decision histories, with finite state and flag spaces at each. A diagonal subsequence makes all behavioral probabilities and

all Bayesian beliefs converge. The limit assessment is symmetric and consistent, since these are completely mixed strategies in the original infinite game. The binary-flag construction lifts to the original private signal histories by using the same informed action probability at every signal arrival time.

At a fixed uninformed information set, the conditional gain from any one-stage deviation converges to its gain in the infinite-game limit. Indeed, the contribution from rewards after L further intervals is uniformly at most $e^{-rL\eta}$, independently of beliefs, the finite terminal date, and strategies. Payoffs over the preceding finite number of intervals are continuous in the converging assessment. Taking the limit in the constrained decision-agent optimality conditions shows that every action used with positive limiting probability is optimal. An informed firm uses R in the limit, which is strictly optimal by Lemma 1. Thus all one-stage deviations are unprofitable. For a deviation lasting L intervals, backward elimination of its last changed action shows that its gain cannot exceed zero. An arbitrary infinite deviation differs from such a finite deviation by a payoff bounded by $e^{-rL\eta}$, which tends to zero. Thus the one-stage conditions also exclude multistage deviations. The limit is a symmetric sequential equilibrium. □

B.5 The continuous-time game of planned switching dates

A timing kernel gives a firm's payoff as a function of its own planned switching date t and its rival's planned date ℓ. For the derivations in this subsection, let $d_1 = r + 2\lambda_2 + \lambda_1$ and $d_2 = r + 2\lambda_2 + 2\lambda_1$. Put $q = p_0/(1-p_0)$, $m(t) = e^{-\lambda_1 t}$, $C_1(t) = e^{-d_1 t}$, $C_2(t) = e^{-d_2 t}$, and $D(t) = e^{-rt}$, interpreting these exponentials as zero at infinity. Remove the constant qg from the unnormalized expected payoff. Write K^0 for the kernel of the auxiliary timing game with instantaneous uninformed response after the first exit. It is, for $t < \ell$,

$$K^0(t,\ell) = q\big[(v-g)C_1(t)(1-m(t)) + (s-g)C_2(t) \\ + (w-g)(1-m(t))C_1(\ell)\big] + sD(t), \tag{19}$$

and, for $t > \ell$,

$$K^0(t,\ell) = q\big[(w-g)C_1(\ell)(1-m(\ell)) + (s-g)C_2(\ell) \\ + (v-g)(1-m(\ell))C_1(t)\big] + sD(\ell). \tag{20}$$

On the diagonal,

$$K^0(t,t) = q\big[(v+w-2g)C_1(t)(1-m(t)) + (s-g)C_2(t)\big] + sD(t). \tag{21}$$

These formulas follow by conditioning on which planned date comes first and whether a private signal has cancelled either exit. They agree on the diagonal and define a continuous kernel on $[0,\infty]^2$. Actual payoffs are $(qg + K^0)/(1+q)$, an irrelevant positive affine transformation for best responses.

Let $L = r + \lambda_0 + \lambda_2$, $d = r + \lambda_0$ and $J = r + \lambda_0 + \lambda_2 + \lambda_1$. The solo cutoff in odds is $Q^S = \lambda_0 L/[(\lambda_2 - \lambda_0)J]$. Write $Q(t) = qe^{-2at}$, $D(t) = e^{-rt}$, and let $j(Q)$ be the continuous-time optimal solo value with free immediate choice. Define its unnormalized option value by

$$\Delta(Q) = (1+Q)[j(Q) - s].$$

For $Q \le Q^S$, $\Delta(Q) = 0$. For $Q > Q^S$, the optimal uninformed risky duration is $a^{-1}\ln(Q/Q^S)$, so

$$\Delta(Q) = (w-s)Q - s + kQ^{-d/a}, \qquad k = \frac{sa}{J}(Q^S)^{d/a}. \tag{22}$$

This expression follows directly from the continuous-time limits of (14); both its value and its first derivative vanish at $Q = Q^S$.

Put $E(t) = D(t)\Delta(Q(t))$. Using the kernel K^0 defined above, the complete limiting planned-date kernel is, away from the diagonal,

$$K(t,\ell) = \begin{cases} K^0(t,\ell) - zE(t), & t < \ell, \\ K^0(t,\ell) + E(\ell), & t > \ell. \end{cases} \tag{23}$$

On the diagonal $K(t,t) = K^0(t,t)$, because simultaneous uninformed exit leaves both firms safe. The first correction is the loss of the safe leader caused by the follower's additional experimentation; the second is the follower's option gain. Both are unnormalized, and the constant qg and factor $1/(1+q)$ are omitted as before. When $Q(t) > Q^S$, the one-sided values at the diagonal therefore need not coincide. Any passage from grid equilibria to this kernel must address possible concentration of stopping dates in that region.

B.6 The immediate-safe limit

The next result closes the pessimistic-prior part of the limiting uniqueness question without imposing a shape on symmetric behavior on any positive grid. Its conclusion concerns the small-grid limit.

Lemma 4 (The low-prior timing game)**.** *For fixed primitives and prior odds $0 < q \le Q^S$, the continuous-time timing game has the unique symmetric equilibrium law δ_0. Every selection of symmetric grid-equilibrium planned switching laws converges weakly to δ_0 as $\eta \downarrow 0$. Both conclusions include $q = Q^S$.*

Proof. Write $L = r + \lambda_0 + \lambda_2$, $d_0 = r + 2\lambda_2$, $d_1 = d_0 + \lambda_1$, $d_2 = d_0 + 2\lambda_1$, and retain $s = \lambda_0/(r+2\lambda_0)$, $g = \lambda_2/(r+2\lambda_2)$, $v = \lambda_0/L$, $w = \lambda_2/L$. The grid solo cutoff satisfies $q_S > Q^S$. Hence, after any failure-revealing first exit, an uninformed survivor optimally chooses S at the next date: its posterior odds $Q(t) = qe^{-2(\lambda_2+\lambda_1)t}$ are at most $Q^S < q_S$.

If a symmetric equilibrium exits certainly at date zero, its planned exit law is already δ_0. Otherwise, both initial safe probabilities are below one, and Lemma 3 permits the failure-revealing completion without altering on-path behavior. Here $E = 0$, so the limiting kernel $K = K^0$ is continuous on X^2, where $X = [0,\infty]$ is the compactified half-line. The uniform grid approximation proved in Lemma 5 remains valid in this region. Consequently weak convergence of the planned-date laws passes both average payoffs and pure-date deviation payoffs to the limit. Every subsequential limit is therefore a symmetric mixed equilibrium of this timing game. We prove that its only symmetric equilibrium is δ_0.

Put $f(t) = e^{-d_1 t}$, $m(t) = e^{-\lambda_1 t}$, $D(t) = e^{-rt}$, $\alpha = g - v > 0$, $\beta = w - g > 0$, and $\gamma = g - s > 0$. For $t < \ell$ the limiting kernel can be written

$$K(t,\ell) = A(t) + q\beta[1 - m(t)]f(\ell), \qquad A(t) = -q\alpha f(t) + q(s-v)f(t)m(t) + sD(t).$$

In particular,

$$K(0,\ell) = K(\ell,0) = k_0 := s - q\gamma \quad \text{for every } \ell \in X.$$

Thus δ_0 is a symmetric equilibrium.

First observe that $t \mapsto K(t,t)$ is strictly decreasing on $(0,\infty)$. Indeed, define

$$B_0 = 2g - v - w = \frac{r(\lambda_2 - \lambda_0)}{Ld_0} > 0, \qquad C_0 = v + w - g - s = \frac{r(\lambda_2 - \lambda_0)^2}{Ld_0(r + 2\lambda_0)} > 0.$$

Direct differentiation and simplification give

$$\frac{d}{dt}K(t,t) = e^{-rt}\left\{q\left[B_0 d_1 e^{-(d_1 - r)t} + C_0 d_2 e^{-(d_2 - r)t}\right] - rs\right\}, \qquad B_0 d_1 + C_0 d_2 = \frac{rs}{Q^S}.$$

The expression in braces is strictly negative for every $t > 0$ when $q \le Q^S$. It follows that $K(t,t) < k_0$ for every $t \in (0,\infty]$, interpreting the value at infinity by continuity.

Suppose $\mu \neq \delta_0$ is a symmetric timing equilibrium. Any atom at zero can be removed. More precisely, if $x = \mu(\{0\})$ lies in $(0,1)$ and $\nu = (\mu - x\delta_0)/(1-x)$, then against μ the payoff from any date t is $xk_0 + (1-x)\int K(t,\ell)\nu(d\ell)$. Since zero is supported, the equilibrium value is k_0. The best-response inequalities and equality on the support of ν therefore show that ν is itself a symmetric timing equilibrium with value k_0. We may consequently assume $\mu(\{0\}) = 0$.

Let $T = \inf \operatorname{supp} \mu$. If $0 < T < \infty$, continuity of the expected-payoff function and the equilibrium support condition imply that T is a best response. But all rival dates are at least T, and

$$K(T,\ell) \le K(T,T) < k_0, \qquad \ell \ge T,$$

where the first inequality follows from $f(\ell) \le f(T)$. This contradicts the payoff k_0 available at date zero. If $T = \infty$, then $\mu = \delta_\infty$ and the same argument uses $K(\infty,\infty) < k_0$.

It remains to exclude $T = 0$ with no atom at zero. Define $U_\mu(t) = \int K(t,\ell)\mu(d\ell)$. The kernel is continuous, piecewise continuously differentiable, and its first derivatives are uniformly bounded near zero. Since $\mu(\{0\}) = 0$, differentiation under the integral at the right endpoint is justified by dominated convergence and gives

$$\begin{aligned} U'_\mu(0+) &= q\{\alpha d_1 - (s-v)d_2\} - rs + q\beta\lambda_1 \int f(\ell)\mu(d\ell) \\ &= rs\left(\frac{q}{Q^S} - 1\right) - q\beta\lambda_1\left[1 - \int f(\ell)\mu(d\ell)\right] < 0. \end{aligned}$$

Here the second equality uses $\alpha d_1 - (s-v)d_2 + \beta\lambda_1 = rs/Q^S$. The final inequality is strict also when $q = Q^S$, because $f(\ell) < 1$ for all $\ell > 0$ and $\mu(\{0\}) = 0$. Thus $U_\mu(t) < U_\mu(0) = k_0$ throughout some interval $(0,\varepsilon)$, which cannot contain support. This contradicts $\inf \operatorname{supp} \mu = 0$.

Every subsequential limit is therefore δ_0, establishing weak convergence of the entire family of planned-date laws. For actual on-path actions, fix $\varepsilon > 0$ and $0 < \delta < \varepsilon$. When $\eta < \varepsilon - \delta$, the probability that the race is still running at ε and at least one firm remains risky is at most

$$2\mu_\eta((\delta,\infty]) + 2\lambda_1\delta.$$

Outside these exceptional events, both planned dates are at most δ and neither firm's plan is cancelled by a private signal before then. A prior rival exit makes an uninformed survivor safe after at most one additional grid interval. First let $\eta \downarrow 0$ and then $\delta \downarrow 0$ to obtain the asserted convergence of actual on-path actions. Invention before a safe choice terminates the race and therefore does not count as continued experimentation. □

B.7 Uniqueness for all priors

We prove the theorem in three steps: passage from grid equilibria, uniqueness of the late mixing tail up to scale, and a global comparison using complementarity. Lemma 4 already settles $q \le Q^S$, so throughout the proof below $q > Q^S$. Use

$$\alpha = g - v > 0, \quad \beta = w - g > 0, \quad d = r + \lambda_0, \quad J = L + \lambda_1, \quad \kappa = rs/Q^S,$$

where $h = 2\lambda_2 + \lambda_1$; the other quantities were defined in the timing-game derivation.

Lemma 5 (Grid limits and support)**.** *Every weak subsequential limit of symmetric grid-equilibrium planned exit laws is a symmetric equilibrium of the full limiting timing game. For $q > Q^S$, every such limit, and every symmetric equilibrium of that timing game, has unbounded support, no mass at infinity, and no finite atom except possibly at T^S.*

Proof. For sufficiently small grids, $q > q_S$ makes certain simultaneous exit at zero unprofitable: continuing against a safe rival gives strictly more than s. Hence the initial exit probability is below one. Lemma 3 permits the failure-revealing completion without altering on-path behavior. The planned-date laws are therefore symmetric equilibria of the exact grid timing game. If there is a first certain uninformed exit date, it is common to both firms. A unilateral deviation past it faces a rival who is already safe or has received a private signal and stays risky forever. Thus unexpected persistence does not add a missing uninformed-rival continuation to this symmetric timing representation.

The solo cutoffs satisfy $0 < \ln(q_S/Q^S) < a\eta$. An optimal grid solo duration, with the current risky interval compulsory, is within η of the continuous duration $a^{-1}[\ln(Q/Q^S)]_+$. The statewise solo values have derivative bounds $(w - s)J$ and sd. Since $z < 1$, the exact and limiting kernels obey the uniform bound on grid pairs

$$|K_\eta(t, \ell) - K(t, \ell)| \le \{q(w - s)J + sd\}\eta. \tag{24}$$

The diagonal values agree exactly. This is a uniform approximation to a kernel that is itself discontinuous on $t = \ell < T^S$.

Compactify the half-line by adding infinity and take $\mu_\eta \Rightarrow \mu$. Here weak convergence means convergence of integrals against every bounded continuous function on the compactified date space. We first rule out an atom of mass $A > 0$ at $t < T^S$, including $t = 0$. Choose a small interval around t, with mass converging to A as its width tends to zero, and compare the equilibrium payoff conditional on drawing one's own planned date inside that interval with a deviation just beyond its upper endpoint. For two distinct dates in the interval, the sum of the two kernel payoffs tends uniformly to $2K^0(t, t) + (1 - z)E(t)$. At equal dates the correction is zero. Thus the average payoff against the rival's dates inside the interval is at most $K^0(t, t) + (1 - z)E(t)/2 + o(1)$, whereas the delayed deviation gives $K^0(t, t) + E(t) + o(1)$. The conditional own-plan payoff equals the equilibrium payoff, since every supported grid plan is optimal. To control rival dates outside this shrinking interval, first fix a larger neighborhood of t. Outside that neighborhood, the kernel is uniformly continuous as the firm's own date approaches t, so the payoff change tends uniformly to zero. Inside the larger neighborhood but outside the stopping cluster, its absolute value is bounded by a constant times the probability of those dates. That probability tends to zero when the larger

neighborhood shrinks to the atom. Taking the grid limit, shrinking the cluster, and then shrinking the larger neighborhood gives a deviation gain of at least

$$A(1+z)E(t)/2 > 0,$$

a contradiction. The identical argument applies directly to a putative atom of a continuous timing equilibrium.

The discontinuity set of K now has zero $\mu \otimes \mu$ mass. Boundedness and almost-everywhere continuity therefore imply convergence of equilibrium payoffs. For any pure deviation t, choose adjacent grid dates converging to t; the same argument applies to their product with μ_η, since $\mu(\{t\}) = 0$ when $t < T^S$, and K is continuous at the diagonal otherwise. Together with (24), this passes all best-response inequalities to the limit. The expected payoff $U_\mu(t) = \int K(t,\ell)\mu(d\ell)$ is continuous, so it equals its maximum V_μ throughout the support.

If μ has mass $A > 0$ at infinity, choosing a large finite T instead of infinity gains at least $Ase^{-rT} - Cqe^{-d_1 T} > 0$ for a constant C independent of T. Indeed the bad-state safe gain is nonnegative against every rival date and equals se^{-rT} against infinity; all adverse good-state terms are bounded by a constant times $e^{-d_1 T}$. This contradicts optimality of a supported date at infinity.

For $T > T^S$, an atom of mass A gives an upward jump in the one-sided derivative of U_μ equal to

$$AD(T)\{rs - \kappa Q(T)\} > 0.$$

Such a kink cannot maximize a continuous function. Thus finite atoms are possible only at T^S, where both this coefficient and E, E' vanish.

Finally, support cannot be bounded. The law δ_0 is not an equilibrium, since a positive planned date against it gains $\Delta(q) > 0$. Hence $I_\infty := \int(1 - m(\ell))\mu(d\ell) > 0$. If T were the last support point, then for every $t > T$ direct differentiation of the kernel gives

$$U'_\mu(t) = q\alpha d_1 C_1(t) I_\infty > 0.$$

Continuity at T would make a later date strictly better, a contradiction. This proves all the claims. □

Lemma 6 (A common late mixing tail up to scale)**.** *Every symmetric timing equilibrium for $q > Q^S$ has full support on a final interval. Any two such equilibrium laws, divided by their respective positive values of I_∞, agree on a final interval. Their payoff gaps agree and equal zero there.*

Proof. For a finite measure μ write, at atomless dates,

$$S(t) = \mu((t,\infty)), \quad I(t) = \int_{[0,t]} (1 - m(\ell))\mu(d\ell), \quad J_1(t) = \int_{(t,\infty)} C_1(\ell)\mu(d\ell).$$

Put $A = I + S$, $a_1 = \alpha d_1$, $a_2 = (s - v)d_2$ and $k_3 = \beta\lambda_1$. For $t > T^S$, differentiation of the kernel gives

$$U'_\mu = qa_1 C_1 A - (qa_2 C_2 + rsD)S + qk_3 m J_1. \tag{25}$$

Lemma 5 gives unbounded support, so $S > 0$. Set

$$O = qe^{-ht}A/S, \qquad y = J_1/(C_1 S) \in [0,1],$$
$$N(Q,y) = hrs - Q(\lambda_1 a_2 + (r + 2\lambda_2)k_3 y).$$

On an open gap in support, A, S, J_1 are constant. At any zero of U'_μ in such a gap, differentiating (25) and eliminating O yields

$$U''_\mu = -DSN(Q, y).$$

Uniformly over $y \in [0, 1]$, $N(Q, y) > 0$ for all sufficiently small Q. A bounded gap sufficiently far out would have payoff V_μ at both endpoints and at most V_μ inside. It would therefore contain a local minimum or a constant interval, both inconsistent with the displayed strictly negative second derivative at every stationary point in the gap. Unbounded support consequently contains a final interval.

On that interval $U'_\mu = 0$. The Stieltjes derivative of (25) has coefficient $D(rs - \kappa Q) > 0$ on $\mu(dt)$. All remaining terms are absolutely continuous. Hence μ is absolutely continuous there. Writing its hazard as $\lambda = -S'/S$, differentiation of indifference gives

$$a_1 O = rs + a_2 Q - k_3 Q y, \qquad \lambda = \frac{N(Q, y)}{rs - \kappa Q}, \qquad y' = d_1 y - \lambda(1 - y). \tag{26}$$

In particular $\lambda > 0$ sufficiently far out. The derivative in y of the scalar right-hand side in the last equation is

$$d_1 + \lambda + \frac{Q(r + 2\lambda_2)k_3(1 - y)}{rs - \kappa Q},$$

which tends uniformly to $d_1 + h > 0$ as $t \to \infty$. The difference between two bounded solutions therefore solves $z' = \zeta(t)z$, with $\zeta(t)$ bounded below by a positive constant eventually. A nonzero difference would become unbounded. Thus their y, λ and O coincide on their common final interval. Their survival functions are proportional there, and equality of O makes $I + S$ proportional by the same factor. Letting t tend to infinity identifies that factor as the ratio of their I_∞ values. Division by I_∞ therefore gives exactly matching tails. Both payoff gaps are zero on the common full-support interval, including after this positive rescaling. □

Proof of Theorem 1. Lemma 4 settles $q \le Q^S$. For $q > Q^S$, we first prove that the limiting timing game has at most one symmetric equilibrium law. The argument permits arbitrary gaps in mixing and a possible atom at the solo cutoff.

For each candidate define its payoff gap by $W(t) = V_\mu - U_\mu(t) \ge 0$. Optimality gives $W = 0$ on the support of its stopping measure. By Lemma 6, any two candidates have proportional final survival and moment states. Put $M = I_\infty$ and write $Z(t) = J_1(t)$ for the discounted future stopping moment. Divide each candidate measure, moment state, and payoff-gap state by its own M. This normalization preserves all best-response and complementarity conditions by homogeneity. The resulting candidates have identical future states on a common terminal half-line and $I_1(\infty) = I_2(\infty) = 1$. It is crucial here that the tail is proportional to M: proportionality of hazards, without this moment normalization, would not imply identical state vectors.

The timing kernel gives the following forward-time Stieltjes equation:

$$dW = -\operatorname{base}(t)\, dt - \Gamma(t)\, \mu(dt),$$

where, using the coefficients defined below,

$$\operatorname{base}(t) = S(t)A'_F(t) + q\beta\lambda_1 m(t)Z(t) + q\alpha d_1 C_1(t)I(t),$$

$$A_F(t) = -q\alpha C_1(t) + q(s - v)C_1(t)m(t) + sD(t) - zE(t).$$

Also $dS = -d\mu$, $dI = (1-m)d\mu$, and $dZ = -C_1 d\mu$. These equations hold in the Stieltjes sense, including for singular-continuous stopping measures, and classically where densities exist. At the only possible atom, $Q = Q^S$ and $\Gamma = 0$. For the difference argument, write

$$\mathcal{S} = S_1 - S_2, \quad \mathcal{W} = W_1 - W_2, \quad F = \Gamma\mathcal{S} - \mathcal{W}.$$

Equal values of $I(\infty)$ and integration by parts give

$$I_1(t) - I_2(t) = -(1-m(t))\mathcal{S}(t) - \lambda_1 \int_t^\infty m(u)\mathcal{S}(u)\,du,$$

$$Z_1(t) - Z_2(t) = C_1(t)\mathcal{S}(t) - d_1 \int_t^\infty C_1(u)\mathcal{S}(u)\,du.$$

Fix a starting date T beyond which the candidates agree and set $x = T - t$. In the following backward-time formulas, a coefficient such as $m(x)$ means the forward coefficient evaluated at $T - x$; the same convention applies to C_1, Γ, χ and the state differences. The exact difference equation is

$$\begin{aligned} F'(x) &= \chi(x)\mathcal{S}(x) \\ &\quad + q\lambda_1 d_1 \left\{ \beta m(x) \int_0^x C_1(y)\mathcal{S}(y)\,dy + \alpha C_1(x) \int_0^x m(y)\mathcal{S}(y)\,dy \right\}, \end{aligned} \tag{27}$$

with $\chi = -(B_{\text{gap}} + \Gamma') \geq 0$, strictly positive away from the unique date T^S at which $Q = Q^S$. The coefficients are

$$\alpha = g - v > 0, \quad \beta = w - g > 0, \quad m = e^{-\lambda_1 t}, \quad C_1 = e^{-d_1 t},$$

$$E = D\Delta(Q), \qquad \Gamma = (1+z)E, \qquad B_{\text{gap}} = rsD(Q/Q^S - 1) - zE',$$

where $\Delta(Q) = (1+Q)(j(Q)-s)$ and $D = e^{-rt}$. The simplification uses $\alpha d_1 + \beta\lambda_1 - (s-v)d_2 = rs/Q^S$. The prime in Γ' denotes forward time differentiation. On $Q < Q^S$, $\chi = rsD(1 - Q/Q^S) > 0$. On $Q > Q^S$,

$$\chi = (r + 2\lambda_0)\left\{ \frac{a}{r+\lambda_0} q(w-s)e^{-d_2 t} - kq^{-(r+\lambda_0)/a} e^{(r+2\lambda_0)t} \right\} > 0,$$

where $a = \lambda_2 + \lambda_1$ and $k = sa(Q^S)^{(r+\lambda_0)/a}/(r + \lambda_0 + a)$. The ratio of the second term in braces to the first equals $(Q^S/Q)^{(r+\lambda_0+a)/a} < 1$.

Complementarity gives the signed-measure inequality

$$\mathcal{W}\,d\mathcal{S} = -W_1\,dS_2 - W_2\,dS_1 \leq 0 \tag{28}$$

in increasing backward time. Both dS_i are nonnegative measures.

We first justify that any first divergence supplies an initial one-sided survival ordering, without assuming finitely many switches. Restart backward time at the end of their maximal initial backward interval of agreement. All earlier differences and memory integrals then vanish. At jump dates use the corresponding one-sided limits. If it occurs where $Q < Q^S$ (equivalently, $t > T^S$), then $\Gamma = 0$ and $F = -\mathcal{W}$. On a sufficiently short backward interval, (27) has the form $F' = \chi\mathcal{S} + \mathcal{L}\mathcal{S}$ with $\min\chi > 0$ and $\|\mathcal{L}\|_{L^2 \to L^2}$ arbitrarily small. If F vanishes at both endpoints, integration by parts and (28) imply

$$0 \leq \int F\,d\mathcal{S} = -\int \chi\mathcal{S}^2 - \int \mathcal{S}\mathcal{L}\mathcal{S} < 0$$

unless the paths agree throughout the interval. More explicitly, the first inequality follows from $F = -\mathcal{W}$ and (28); the boundary term in integration by parts is zero because F is zero at both endpoints. The last integral has absolute value at most $\|\mathcal{L}\| \int \mathcal{S}^2$. On a short enough interval this is smaller than $\min \chi \int \mathcal{S}^2$, giving strict negativity whenever the survival paths differ. Hence zeros of F cannot accumulate immediately after a nontrivial first departure. Its first nonzero interval has a constant sign. If $F > 0$, then $W_1 < W_2$ and $W_2 > 0$, so $dS_2 = 0$ there and $\mathcal{S}$ increases backward from zero. The reverse sign gives the reverse ordering. An unequal atom at the departure date already gives an immediate one-sided ordering.

At a date with $Q > Q^S$ (equivalently, $t < T^S$), $\Gamma > 0$ and local backward continuation is unique by a direct first-attainment argument. Integrating (27) expresses $F(x)$ as a weighted integral of $\mathcal{S}(y)$ for $0 \le y \le x$, with nonnegative weights. Such an integral is a Volterra operator: only earlier backward-time values enter. Denote the integral of all its weights by $R(x)$. Thus $|\mathcal{S}(y)| \le h$ on $[0, x]$ implies $|F(x)| \le hR(x)$. The total positive mass $R(x)$ mapping $\mathcal{S}$ to $F(x)$ is $O(x)$, whereas Γ is bounded away from zero. Thus $R(x) < \Gamma(x)$ on a sufficiently short backward interval. If $|\mathcal{S}|$ first attains a positive level h there, positivity of the Volterra kernel gives $|F| \le hR < h\Gamma$. At a first positive attainment candidate 1 must have stopping mass arbitrarily close, so $W_1 = 0$ and $F = \Gamma h - (W_1 - W_2) \ge \Gamma h$, a contradiction. At a first negative attainment use candidate 2. Atoms are impossible when $Q > Q^S$, since an infinitesimal delay against an atomic rival exit yields the strictly positive follower gain $D\Delta(Q)$. Thus first attainments are continuous, and no regularity or finite-switching assumption is hidden in this argument.

At the cutoff itself, unequal atoms again supply a one-sided initial ordering: the nonzero backward right limit of the survival difference has a fixed sign on a sufficiently short interval. If the atom sizes agree, there is unique passage through the singularity. To see this directly, put $x = T^S - t$ and

$$G = 2s(\lambda_2 + \lambda_1)(r + 2\lambda_0)e^{-rT^S}, \qquad C = \frac{\lambda_1 d_1}{2(\lambda_2 + \lambda_1)(L + \lambda_1)} < 1.$$

Indeed $aJ - \lambda_1 d_1 = \lambda_2 L + \lambda_0 \lambda_1 > 0$, so even $C < 1/2$. Smooth fit and $\Delta''(Q^S) = sd/[a(Q^S)^2]$ give $\Gamma = Gx^2 + O(x^3)$ and $\chi = Gx + O(x^2)$. The memory kernel at the cutoff has value $q\lambda_1 d_1(\alpha + \beta)m(T^S)C_1(T^S) = GC$. Consequently the total positive mass $R(x)$ of the Volterra operator mapping $\mathcal{S}$ to $F(x)$ satisfies

$$R(x)/\Gamma(x) \longrightarrow (1 + C)/2 < 1.$$

If the difference departs from zero, let x_h be its first attainment of a positive absolute level h. There are no atoms on $Q > Q^S$: delaying such an atom gives the strictly positive follower gain $D\Delta(Q)$ against its atomic mass. Thus the first attainment is continuous, and $|\mathcal{S}| \le h$ beforehand. Consequently $|F(x_h)| \le hR(x_h) < h\Gamma(x_h)$. If $\mathcal{S}(x_h) = h$, candidate 1 must place stopping mass arbitrarily close to that first attainment, hence $W_1(x_h) = 0$. This gives $F(x_h) \ge \Gamma(x_h)h$, a contradiction. The negative case uses $W_2(x_h) = 0$ and yields the opposite contradiction. This rules out even infinitely oscillatory departures at the cutoff.

We next propagate the resulting initial ordering. Suppose, without loss, $\mathcal{S} \ge 0$ on a nontrivial initial backward interval. Equation (27) then gives $F > 0$ as soon as a nonzero survival difference has accumulated, and keeps it strictly positive while the ordering holds. At a first attempted crossing $\mathcal{S} = 0$, we would have $\mathcal{W} = -F < 0$, so $W_2 > W_1 \ge 0$. Candidate 2 therefore has no stopping mass in a neighborhood of the crossing. Candidate 1's survival is nondecreasing backward, making the crossing impossible. An opposing atomic jump is also impossible: its only possible date is T^S,

where $\Gamma = 0$ gives $\mathcal{W} = -F < 0$ regardless of how far $\mathcal{S}$ jumps. Such a jump requires $W_2 = 0$, a contradiction. Thus the survival ordering persists all the way to date zero, across every pooling gap and through the solo-cutoff date.

Finally the equal normalization of $I(\infty)$ gives the exact identity

$$I_1(t) - I_2(t) = -(1 - e^{-\lambda_1 t})\mathcal{S}(t) - \lambda_1 \int_t^\infty e^{-\lambda_1 u}\mathcal{S}(u)\, du.$$

At date zero both left-hand sides vanish. A nontrivial persistent ordering would therefore imply

$$0 = -\lambda_1 \int_0^\infty e^{-\lambda_1 u}\mathcal{S}(u)\, du < 0,$$

a contradiction. Hence the normalized candidates agree on $(0, \infty]$. They cannot differ only at zero: since $q > Q^S$, an atom at zero also admits the strictly profitable infinitesimal-delay deviation just used. Both original candidates are probability measures, so their normalizing constants M agree as well. Their timing laws are identical.

It remains to connect this unique law to the reaction-grid equilibria. For $q \leq Q^S$, use Lemma 4. For $q > Q^S$, symmetric grid equilibria exist for every grid by Theorem 2. Their laws have subsequential limits on the compactified half-line. Lemma 5 makes every such limit a symmetric timing equilibrium, and Lemma 6 and the global comparison just proved make that equilibrium unique. Thus every sequence and every equilibrium selection converge to the same law. The support assertions and eventual continuous mixing follow from Lemmas 5 and 6.

For on-path behavior, couple independent planned dates with their limiting dates and use common Poisson processes. Private signals have continuous arrival laws, so equality between a planned date and a cancelling signal has probability zero. Planned ties before T^S also have probability zero. The only possible finite atom occurs at T^S, where the optimal uninformed solo duration is zero: any ordering of grid exits converging to that date therefore gives the same limiting switch times. At all other first-exit dates, the solo duration is continuous in the date, and the grid duration is within η of it. This proves convergence of the joint on-path switching laws, with a never-realized switch recorded as infinity if the race ends first. It also gives convergence of finite-dimensional action laws at continuity dates. Actions evaluated exactly at a possible atom date require an endpoint convention and are not a pointwise convergence assertion of the theorem. □

The degenerate priors cause no additional multiplicity. At $p_0 = 0$ the low-prior proof applies with $q = 0$ and gives immediate safe behavior (indeed risky research cannot produce either an invention or information). At $p_0 = 1$, Lemma 1 makes every active firm choose risky research forever.

Proof of Corollary 1. Use the notation in the uniqueness proof. On the final mixing interval, (26) gives

$$\lambda(t) = h + O(Q(t)), \qquad a_1 O(t) = rs + O(Q(t)), \qquad Q(t) \longrightarrow 0.$$

These estimates are uniform over $y \in [0, 1]$. The planned-date hazard is also the firm's behavioral switching hazard conditional on remaining uninformed and on both firms remaining risky: private signals and the rival's randomization are independent of the firm's own planned-date randomization. Its posterior odds are exactly

$$O(t) = qe^{-(2\lambda_2 + \lambda_1)t}\frac{I(t) + S(t)}{S(t)}.$$

Indeed, the rival remains risky conditional on G with probability $I(t)+S(t)$, and conditional on B with probability $S(t)$. The firm's own absence of a signal and the two absent inventions contribute the likelihood factor $e^{-(2\lambda_2+\lambda_1)t}$. Consequently

$$O(t) \longrightarrow \frac{rs}{(g-v)(r+2\lambda_2+\lambda_1)} = \frac{r\lambda_0(r+\lambda_0+\lambda_2)(r+2\lambda_2)}{(r+2\lambda_0)(\lambda_2-\lambda_0)(r+\lambda_2)(r+2\lambda_2+\lambda_1)}.$$

Converting odds to probability yields the stated expression. The conditional-G probability that the rival is uninformed is $e^{-\lambda_1 t}S(t)/(I(t)+S(t)) = Q(t)/O(t)$, which tends to zero.

For completeness, put $L = r+\lambda_0+\lambda_2$ and $O_\infty = \lim O(t)$. Logarithmic differentiation gives

$$\begin{aligned}
\partial_{\lambda_0} \ln O_\infty &= \frac{r}{\lambda_0(r+2\lambda_0)} + \frac{1}{L} + \frac{1}{\lambda_2-\lambda_0} > 0,\\
\partial_{\lambda_1} \ln O_\infty &= -\frac{1}{r+2\lambda_2+\lambda_1} < 0,\\
\partial_{\lambda_2} \ln O_\infty &= -\frac{\lambda_0}{L(r+\lambda_2)} - \frac{1}{\lambda_2-\lambda_0} + \frac{2\lambda_1}{(r+2\lambda_2)(r+2\lambda_2+\lambda_1)} < 0,\\
\partial_r \ln O_\infty &= \frac{2\lambda_0}{r(r+2\lambda_0)} - \frac{\lambda_0}{L(r+\lambda_2)} + \frac{\lambda_1}{(r+2\lambda_2)(r+2\lambda_2+\lambda_1)} > 0.
\end{aligned}$$

The third inequality uses $2\lambda_1/[(r+2\lambda_2)(r+2\lambda_2+\lambda_1)] < 2/(r+2\lambda_2) < 1/(\lambda_2-\lambda_0)$. For the fourth, $2L(r+\lambda_2) > r(r+2\lambda_0)$ makes the first two terms have a positive sum. The probability $O_\infty/(1+O_\infty)$ inherits all four signs. □

C Experimentation comparisons

C.1 Proofs of the experimentation comparisons

Proof of Proposition 2. Write $q = p_0/(1-p_0)$, and recall $s = \lambda_0/(r+2\lambda_0)$ and $g = \lambda_2/(r+2\lambda_2)$. The odds corresponding to the two cutoffs are

$$Q^P = \frac{\lambda_0(r+2\lambda_2)}{(\lambda_2-\lambda_0)(r+2\lambda_2+2\lambda_1)}, \qquad Q^S = \frac{\lambda_0(r+\lambda_0+\lambda_2)}{(\lambda_2-\lambda_0)(r+\lambda_0+\lambda_2+\lambda_1)}.$$

They obey

$$\frac{Q^S}{Q^P} = 1 + \frac{\lambda_1(r+2\lambda_0)}{(r+2\lambda_2)(r+\lambda_0+\lambda_2+\lambda_1)} > 1. \tag{29}$$

The symmetric-limit theorem gives immediate safe research for $q \le Q^S$. The planner theorem gives, for $q > Q^P$,

$$T^P = \frac{\ln(q/Q^P)}{2(\lambda_2+\lambda_1)} > 0.$$

The planner's value at a proposed no-news stopping date t is

$$W(t,p_0) = 2p_0 g - 2p_0(g-s)e^{-(r+2\lambda_2+2\lambda_1)t} + 2(1-p_0)se^{-rt}.$$

At $q > Q^P$, its unique maximizing date is $T^P > 0$ and its value strictly exceeds $W(0,p_0) = 2s$. This establishes the first claim.

For the duration comparison, differentiating at fixed p_0 gives

$$\frac{\partial T^P}{\partial \lambda_1} = \frac{2(\lambda_2+\lambda_1)/(r+2\lambda_2+2\lambda_1) - \ln(q/Q^P)}{2(\lambda_2+\lambda_1)^2}.$$

Within $Q^P < q \le Q^S$, the numerator is strictly positive. Indeed,

$$\ln(q/Q^P) \le \ln(Q^S/Q^P) < \frac{\lambda_1(r+2\lambda_0)}{(r+2\lambda_2)(r+\lambda_0+\lambda_2+\lambda_1)} < \frac{\lambda_1}{r+\lambda_0+\lambda_2+\lambda_1} < \frac{2(\lambda_2+\lambda_1)}{r+2\lambda_2+2\lambda_1}.$$

The first strict inequality is $\ln(1+x) < x$ for $x > 0$; the second uses $\lambda_0 < \lambda_2$; and the third follows from $r + 2\lambda_2 + 2\lambda_1 < 2(r + \lambda_0 + \lambda_2 + \lambda_1)$ and $\lambda_1 < \lambda_2 + \lambda_1$.

The envelope theorem, valid because the optimum is strictly interior, now gives the signal-rate derivative of the welfare loss:

$$\frac{\partial}{\partial \lambda_1}\big[W(T^P, p_0) - 2s\big] = 4p_0(g-s)T^P e^{-(r+2\lambda_2+2\lambda_1)T^P} > 0. \tag{30}$$

Its prior derivative is

$$\frac{\partial}{\partial p_0}\big[W(T^P, p_0) - 2s\big] = 2(g-s)\big[1 - e^{-(r+2\lambda_2+2\lambda_1)T^P}\big] + 2s\big[1 - e^{-rT^P}\big] > 0.$$

The derivatives describe local changes that keep the prior inside the stated region; no differentiability of equilibrium welfare across its boundary is asserted.

□

Lemma 7 (Continuity at the immediate-safe boundary)**.** *Fix the research rates and discount rate. As the prior decreases to p^S from above, the unique symmetric limiting law of planned uninformed safe-adoption dates converges weakly to δ_0, and its mean converges to zero.*

Proof. Use the odds notation of the preceding proof. Write K_q for the full limiting timing kernel, without its positive affine payoff normalization. The timing-kernel representation in the equilibrium appendix has the form

$$K_q(t,\ell) = K_q^0(t,\ell) + \begin{cases} -ze^{-rt}\Delta(qe^{-2(\lambda_2+\lambda_1)t}), & t < \ell, \\ 0, & t = \ell, \\ e^{-r\ell}\Delta(qe^{-2(\lambda_2+\lambda_1)\ell}), & t > \ell, \end{cases} \qquad z = \frac{\lambda_0}{r+\lambda_0}.$$

Here K_q^0 is jointly continuous on $[0,\infty]^2$, depends affinely on q, and represents immediate following after an uninformed first exit. The function $\Delta(Q)$ is the optimized solo payoff gain over immediate safe research, multiplied by $1+Q$. It is continuous, nondecreasing, and zero for $Q \le Q^S$. Consequently

$$\sup_{t,\ell} |K_q(t,\ell) - K_{Q^S}(t,\ell)| \le C|q - Q^S| + \Delta(q) \longrightarrow 0 \qquad (q \downarrow Q^S).$$

The compactification includes infinite planned dates. Take any sequence $q_n \downarrow Q^S$ and any weakly convergent subsequence of equilibrium laws, with limit μ. Uniform kernel convergence and continuity of K_{Q^S} pass both the expected equilibrium payoff and every pure deviation payoff to the limit. Thus μ is a symmetric equilibrium of the timing game at Q^S. The low-prior uniqueness proof establishes

that the only such equilibrium is δ_0. Compactness then implies convergence of the entire family to δ_0.

It remains to control the tails uniformly. Fix an upper odds bound $\bar{q} > Q^S$, and write $S_q(t) = \mu_q((t,\infty))$. For $q \in (Q^S, \bar{q}]$, choose T_0 larger than every solo-cutoff date $T^S(q)$. The stopping law has no atoms on $[T_0,\infty)$. At every supported date $t > T_0$, the expected stopping payoff is differentiable and has derivative zero. In the differentiated timing kernel, the past-signal moment plus stopping survival is at most one, and the discounted future-stopping moment is bounded above by $e^{-(r+2\lambda_2+\lambda_1)t}S_q(t)$. The resulting indifference equation therefore implies

$$\begin{aligned}\frac{r\lambda_0}{r+2\lambda_0}S_q(t) \le \bar{q}\left(\frac{\lambda_2}{r+2\lambda_2} - \frac{\lambda_0}{r+\lambda_0+\lambda_2}\right)(r+2\lambda_2+\lambda_1)e^{-(2\lambda_2+\lambda_1)t} \\ + \bar{q}\lambda_1\left(\frac{\lambda_2}{r+\lambda_0+\lambda_2} - \frac{\lambda_2}{r+2\lambda_2}\right)e^{-2(\lambda_2+\lambda_1)t}S_q(t).\end{aligned}$$

Increase T_0, if necessary, until the coefficient of $S_q(t)$ on the right is at most $r\lambda_0/[2(r+2\lambda_0)]$ for every $t \ge T_0$. This gives a constant C, independent of q, such that

$$S_q(t) \le Ce^{-(2\lambda_2+\lambda_1)t} \qquad \text{at every supported date } t > T_0.$$

At a date outside the support, let $u \ge t$ be the next supported date. Such a date exists because support is unbounded; there is no mass between t and u or at u. Hence $S_q(t) = S_q(u)$, and the same bound holds at t. This argument does not require a uniform date at which final mixing begins. Weak convergence to δ_0 gives $S_q(t) \to 0$ at every fixed $t > 0$. The uniform exponential bound and $S_q \le 1$ now give

$$\int_0^\infty S_q(t)\,dt \longrightarrow 0 \qquad (q \downarrow Q^S)$$

by dominated convergence. The integral is the mean planned stopping date, proving the claim. □

Proof of Proposition 3. Continue to use the local odds notation: $T^P(q)$ and $T^S(q)$ denote the stopping dates at prior probability $q/(1+q)$. At $q = Q^S$, the planner's no-news stopping date is strictly positive because $Q^S > Q^P$. Fix

$$\bar{t} = \frac{1}{2}T^P(Q^S) > 0.$$

For $q > Q^S$ sufficiently close to Q^S,

$$2T^S(q) := \frac{\ln(q/Q^S)}{\lambda_2+\lambda_1} < \bar{t} < T^P(q).$$

Lemma 7 also makes $\mu_q([0,\bar{t})) > 0$ for all such priors. Take independent planned dates with law μ_q. On the counterfactual no-news history, their minimum t is the first safe-adoption date. If $t < \bar{t}$, the uninformed survivor adopts safe research at

$$t + \frac{[\ln(qe^{-2(\lambda_2+\lambda_1)t}/Q^S)]_+}{\lambda_2+\lambda_1} = \max\{t, 2T^S(q)-t\} < T^P(q).$$

This proves positive probability of both safe adoptions before the planner's date. The unbounded support of μ_q gives positive probability that both planned dates exceed $T^P(q)$, establishing the opposite comparison with positive probability as well.

For precision about conditioning, these are initially events in the independent stopping randomizations. Each has strictly positive probability jointly with no invention or private signal up to $T^P(q)$. For example, the bad state has positive prior probability, has no private signals or risky inventions, and has strictly positive probability of no safe invention before this finite date under every research allocation. Hence both events retain positive probability conditional on no news. Their conditional probabilities need not equal their probabilities under the independent planned-date laws.

For the expected counterfactual no-news exposure, let τ be the minimum of the two independent planned dates. The two actual adoption dates along this counterfactual history are τ and $\max\{\tau, 2T^S - \tau\}$. Total risky firm-time is therefore $2\max\{\tau, T^S\}$, with expectation

$$2T^S + 2\int_{T^S}^{\infty} S_q(t)^2\, dt.$$

Lemma 7 and $T^S \to 0$ imply that this expression tends to zero as $q \downarrow Q^S$. The planner's counterfactual exposure instead tends to $2T^P(Q^S) > 0$. Shrinking the prior interval in the proposition if necessary gives the strict expected-exposure comparison for every prior in it. This expectation is over stopping randomizations with news suppressed, and does not assert an ordering of expected realized research across all information and invention histories. □

Proof of Proposition 4. Fix $\lambda_2 > \lambda_0 > 0$, $r > 0$, and $\lambda_1 > 0$. An explicit sufficient threshold is

$$\bar{p}_0(\lambda_1) = \left[1 + \frac{r(\lambda_2 - \lambda_0)(r + \lambda_0)}{\lambda_0(r + 2\lambda_2)(r + 2\lambda_0)} \left\{\frac{r(r + \lambda_0)}{(r + 2\lambda_0)(r + 2\lambda_2 + 2\lambda_1)}\right\}^{\frac{2\lambda_2 + \lambda_1}{\lambda_1}}\right]^{-1}. \tag{31}$$

The other primitives are held fixed in this notation. We establish both its location in $(p^P, 1)$ and the claimed equilibrium implication.

At private belief p, risky research forever guarantees at least $p\lambda_2/(r + 2\lambda_2)$. Irreversible safe research yields at most $p\lambda_0/(r + 2\lambda_0) + (1 - p)\lambda_0/(r + \lambda_0)$. Thus no uninformed firm can switch safely while its private odds exceed

$$K = \frac{\lambda_0(r + 2\lambda_2)(r + 2\lambda_0)}{r(\lambda_2 - \lambda_0)(r + \lambda_0)}.$$

Write $q = p_0/(1 - p_0)$. While both firms have pooled, an uninformed firm's odds at date t are $qe^{-(2\lambda_2 + \lambda_1)t}$: the absence of either invention and of its own signal is informative, whereas the absence of a signal to its rival is not observed. Induction on the reaction dates therefore forces both firms to continue at every date below

$$\tau(p_0) = \frac{1}{2\lambda_2 + \lambda_1} \ln\left[\frac{p_0}{1 - p_0} \frac{r(\lambda_2 - \lambda_0)(r + \lambda_0)}{\lambda_0(r + 2\lambda_2)(r + 2\lambda_0)}\right]. \tag{32}$$

No symmetry or restriction on subsequent mixing is used.

For comparison, the planner's cutoff odds are

$$Q^P = \frac{p^P}{1 - p^P} = \frac{\lambda_0(r + 2\lambda_2)}{(\lambda_2 - \lambda_0)(r + 2\lambda_2 + 2\lambda_1)}, \qquad \frac{Q^P}{K} = \frac{r(r + \lambda_0)}{(r + 2\lambda_0)(r + 2\lambda_2 + 2\lambda_1)} < 1.$$

Formula (31) is equivalent to

$$\frac{\bar{p}_0(\lambda_1)}{1 - \bar{p}_0(\lambda_1)} = K\left(\frac{K}{Q^P}\right)^{\frac{2\lambda_2 + \lambda_1}{\lambda_1}} > K > Q^P.$$

The threshold is therefore strictly between p^P and one. For $p_0 > p^P$, the planner's no-news stopping date is $T^P = \ln(q/Q^P)/(2(\lambda_2 + \lambda_1))$, and direct subtraction gives

$$\tau(p_0) - T^P(p_0) = \frac{\lambda_1}{2(\lambda_2 + \lambda_1)(2\lambda_2 + \lambda_1)} \ln\left[\frac{p_0(1 - \bar{p}_0(\lambda_1))}{\bar{p}_0(\lambda_1)(1 - p_0)}\right]. \tag{33}$$

This expression is strictly positive exactly when $p_0 > \bar{p}_0(\lambda_1)$, and its derivative with respect to p_0 is

$$\frac{\lambda_1}{2(\lambda_2 + \lambda_1)(2\lambda_2 + \lambda_1)p_0(1 - p_0)} > 0.$$

This is a comparative static of the guaranteed delay, without requiring monotonicity of every equilibrium's actual first-adoption date.

An optimal same-grid planner makes its first safe adoption by $T^P + \eta$: at the first grid date when the no-news belief is at or below p^P, immediate joint safe research achieves the unrestricted continuous-time optimum, and a positive risky interval is strictly inferior. In equilibrium the first safe adoption cannot precede τ. Hence its delay relative to any same-grid optimum is at least $\tau - T^P - \eta$. For every fixed $p_0 > \bar{p}_0(\lambda_1)$ this is positive for all sufficiently small grids, with a strictly positive limiting lower bound given by (33).

The same calculation formalizes the intuition about private good news. At T^P, pooling implies that an uninformed firm's odds are

$$\frac{p_0}{1 - p_0} e^{-(2\lambda_2 + \lambda_1)T^P} = \left(\frac{p_0}{1 - p_0}\right)^{\frac{\lambda_1}{2(\lambda_2 + \lambda_1)}} (Q^P)^{\frac{2\lambda_2 + \lambda_1}{2(\lambda_2 + \lambda_1)}}.$$

They diverge as $p_0 \uparrow 1$. Conditional on the good state, the probability that its rival has received a signal by then is $1 - e^{-\lambda_1 T^P}$, which also tends to one. Thus the firm's conditional probability that its rival is informed tends to one, even along the actual history on which neither firm has received news. The planner observes that history and stops at belief p^P. □

C.2 Verification of an asymmetric construction

We give a sufficient test for the equilibrium used in the counterexample. Fix a positive grid date $N\eta$. Until that date both firms choose risky research. At $N\eta$, firm 1 switches safely if it has no signal, while firm 2 remains risky. If firm 1 instead remains risky, an uninformed firm 2 infers G and stays risky until a first safe choice. Even if that choice is off path, an uninformed survivor then updates to Q_n and follows its constrained solo policy. A firm with its own conclusive signal stays risky throughout. Choices are deterministic conditional on information histories, even though signals and inventions are random.

Use the continuation values defined in Appendix B and put

$$m_n = e^{-\lambda_1 n\eta}, \qquad O_n = Q_n/m_n, \qquad k_n = \min\{k \geq 1 : Q_n e^{-ak\eta} \leq q_S\},$$

$$f_G(n) = v + m_n(s - v)A^{k_n}, \qquad f_B(n) = z + (s - z)B^{k_n}, \tag{34}$$

$$\Phi(n) = p_0 g + p_0 e^{-(r + 2\lambda_2 + \lambda_1)n\eta}[f_G(n) - g] + (1 - p_0)e^{-rn\eta} f_B(n), \qquad \Phi(\infty) = p_0 g. \tag{35}$$

Here $f_\theta(n)$ is the conditional-state payoff from a first uninformed switch at $n\eta$ against a rival that has remained risky. The quantity $\Phi(n)$ is firm 1's ex ante payoff from this plan, including the possibility that a private signal cancels it. Select the earlier date at a solo cutoff tie.

Lemma 8 (A sufficient test for the asymmetric construction)**.** *The prescribed behavior has a sequential-equilibrium completion if the following three conditions hold. First,*

$$\Phi(N) \geq \Phi(n) \quad \textit{for every } n \in \{0, 1, \ldots, \infty\}. \tag{36}$$

Second, for each $0 \leq n < N$*, put* $h = (N - n)\eta$ *and define*

$$C_F^G(n) = g + e^{-(r+2\lambda_2)h} m_N \{(w - g) + e^{-\lambda_1 h}(s - w)A^{k_N}\},$$
$$C_F^B(n) = e^{-rh} s B^{k_N}.$$

Require

$$O_n[C_F^G(n) - f_G(n)] + C_F^B(n) - f_B(n) \geq 0. \tag{37}$$

Finally require

$$O_N(1 - m_N)(g - v) + Q_N(w - s)(1 - A^{k_N}) - s(1 - B^{k_N}) \geq 0. \tag{38}$$

Proof. Firm 2 stays risky before firm 1's switch. Planning to switch at $n\eta$ gives firm 1 exactly $\Phi(n)$, so (36) excludes every earlier or later switch and risky research forever. Randomizing its date cannot improve on the maximum. At an off-path uninformed history after firm 1 stays risky at $N\eta$, firm 2 knows G. Give firm 1 an optimal remaining stopping date against that rival, allowing an uninformed rival to revise its inference after a first safe choice as just specified. A maximizer exists because $\Phi(n) \to \Phi(\infty)$; choose a deterministic maximizer.

For firm 2, conditioning on its signal status and firm 1's type at $N\eta$ gives the displayed continuation values C_F^G, C_F^B. Inequality (37) compares continuing with an earlier safe switch. At $N\eta$, an informed firm 1 reveals G by remaining risky. Otherwise its safe switch leaves firm 2 facing the compulsory-current-interval solo problem. Subtracting the payoff from switching safely gives (38). Afterward, risky research after a private signal, risky research while the public history implies G, or the optimal solo continuation after a first safe choice excludes all further deviations.

At every unexpected first safe choice, prescribe beliefs revealing that the sender has no signal. These beliefs are consistent by the common tremble construction in Lemma 3; both initial switch probabilities here are zero. The continuations just specified cover all remaining decision histories. One-stage optimality and the uniform bound $e^{-rL\eta}$ on rewards after L additional intervals exclude arbitrary multistage deviations. Thus the strategy has a consistent, sequentially rational completion. □

C.3 An asymmetric counterexample

The following result gives a strict comparison on all sufficiently fine reaction grids. On the no-news path, both safe switches occur strictly before the planner would abandon risky research.

Proposition 5 (Positive pooling with premature abandonment)**.** *Fix*

$$\lambda_0 = 1, \qquad \lambda_1 = 3, \qquad \lambda_2 = 4, \qquad r = 4, \qquad p_0 = \tfrac{1}{4}. \tag{39}$$

For every sufficiently small $\eta > 0$*, there is an asymmetric sequential equilibrium with a positive pooling phase and deterministic choices at every information history. Its first and last safe-switch dates on the path without inventions or private signals converge respectively to*

$$T_* \simeq 0.0178023351, \qquad T_* + \ell_* \simeq 0.0232951038, \tag{40}$$

where $\ell_* > 0$. Every optimal planner's first no-news safe date converges to

$$T^P = \frac{\ln(3/2)}{14} \simeq 0.0289617934. \tag{41}$$

Consequently, on every sufficiently fine grid and on the no-news path, both equilibrium firms switch safely before the same-grid planner first does so. The inequalities are strict, and the existence and strict ordering persist under sufficiently small perturbations of the primitives and prior.

Proof. For these primitives the economic payoff constants are

$$s = \tfrac{1}{6}, \quad g = \tfrac{1}{3}, \quad v = \tfrac{1}{9}, \quad w = \tfrac{4}{9}, \quad z = \tfrac{1}{5}, \quad A = e^{-12\eta}, \quad B = e^{-5\eta}.$$

The limiting solo cutoff is $Q^S = 1/4$. Put

$$\ell_0 = \frac{\ln(4/3)}{7}, \qquad t_M = \ell_0/2, \qquad \rho = (3/4)^{5/7}.$$

At a first-exit date t, the optimal solo duration in continuous time is $\ell(t) = \max\{0, \ell_0 - 2t\}$. The leader's limiting objective therefore simplifies to

$$\Phi_0(t) = \begin{cases} \frac{1}{12} - \frac{1}{18}e^{-15t} + \frac{3}{20}e^{-4t} - \frac{7}{480}\rho e^{6t}, & 0 \le t \le t_M, \\ \frac{1}{12} - \frac{1}{18}e^{-15t} + \frac{1}{72}e^{-18t} + \frac{1}{8}e^{-4t}, & t \ge t_M. \end{cases} \tag{42}$$

On the first interval its derivative has the sign of the strictly decreasing function

$$H_0(t) = \tfrac{5}{6}e^{-11t} - \tfrac{3}{5} - \tfrac{7}{80}\rho e^{10t}.$$

It is positive at 0.017 and negative at $0.019 < t_M$. On the second interval the corresponding function is $\frac{5}{6}e^{-11t} - \frac{1}{4}e^{-14t} - \frac{1}{2}$, also strictly decreasing. The derivatives match at t_M. Hence Φ_0 has a unique global maximizer $T_* \in (0.017, 0.019)$, the root of H_0. In particular $\ell_* := \ell_0 - 2T_* > 0$ and $T_* + \ell_* < \ell_0 - 0.017 < T^P$.

The exact grid objectives converge uniformly to (42) at their grid dates, including the infinite tail. Indeed, $q_S \to 1/4$ and the forced solo duration $k_n\eta$ differs from $\ell(n\eta)$ by at most $\eta + |\ln(q_S/(1/4))|/7$, which tends to zero uniformly in n. The formulas for f_G, f_B, Φ give the claimed uniform convergence. Since the limiting maximum is unique and strictly exceeds the value at infinity, any global grid maximizer $T_\eta = N_\eta\eta$ is finite and converges to T_*. For sufficiently small grids it is positive, and the leader's global test is satisfied by construction.

It remains to check the follower uniformly, so convergence of the leader's choice alone is not being used as an existence argument. For $0 \le t < T < t_M$, the continuous limit of the follower's early comparison is

$$D_F(t;T) = \tfrac{2}{27}e^{-11t} + \tfrac{1}{27}e^{-15T+t} + \tfrac{7}{72}\rho e^{6T+4t} - \tfrac{1}{5} + \tfrac{7}{360}\rho e^{10t}. \tag{43}$$

Extend this expression continuously to $t = T$ from below. For $0 \le t \le T \le 0.019$ its derivative in t is at most

$$-\tfrac{22}{27}e^{-11(0.019)} + \tfrac{1}{27} + \tfrac{7}{12}\rho e^{10(0.019)} < -0.0497.$$

Thus its minimum is at $t = T$. At $T = T_*$ this left limit of the early comparison can be written as

$$\frac{2}{27}e^{-11T_*}(1 - e^{-3T_*}) + \frac{1}{9}e^{-14T_*}(1 - e^{-12\ell_*}) - \frac{1}{5}(1 - e^{-5\ell_*}).$$

Its first term is the strictly positive gain from continuing against an informed rival. The last two terms are the uninformed survivor's optimized solo gain over immediate safe research, multiplied by its Bayesian denominator and the positive factor $1 + z = 6/5$. Their sum is strictly positive because its posterior odds $e^{-14T_*}/3$ exceed the solo cutoff $1/4$ and its optimal remaining duration is $\ell_* > 0$. Numerically $D_F(T_*; T_*) > 0.0032731$.

At $t = T_*$ itself, a simultaneous safe deviation instead leaves both uninformed firms safe. The continuous limit of the separate terminal test (38) is

$$D_F^{\text{term}}(T_*) = \frac{2}{27}e^{-11T_*}(1 - e^{-3T_*}) + \frac{5}{54}e^{-14T_*}(1 - e^{-12\ell_*}) - \frac{1}{6}(1 - e^{-5\ell_*}) > 0. \tag{44}$$

The first term is the same informed-rival gain. The remaining terms are the same positive solo option gain, now with factor one rather than $1 + z$: a terminal safe deviation does not leave a risky follower. Numerically $D_F^{\text{term}}(T_*) > 0.0032555$.

Uniform convergence of the early comparisons on the compact region $n\eta < N_\eta\eta$ preserves their strict positive margin. Separately, the terminal comparison at $N_\eta\eta$ converges to (44), including the compulsory current risky interval, and is also strictly positive for sufficiently small η. Lemma 8 therefore supplies a sequential equilibrium for every sufficiently fine grid. Its last no-news safe date converges to $T_* + \ell_*$.

Proposition 1 gives convergence of every optimal grid planner's first no-news safe date to T^P. The strict gap between $T_* + \ell_*$ and T^P therefore holds against the planner on the same grid for all sufficiently small η.

Finally, the inequalities used above have strictly positive margins: the leader's derivative changes sign strictly inside $(0, t_M)$, its late objective is decreasing, the follower's minimum payoff difference is positive, and the planner's stopping date strictly exceeds the last equilibrium exit. All the payoff formulas and their derivatives depend continuously on the primitives on the compact intervals in question. The same derivative-sign arguments continue to hold under sufficiently small perturbations; for the late tail the strict exponential coefficient inequality remains valid. Uniform grid approximation then provides the same equilibria and date ordering for every sufficiently fine grid under those perturbed primitives. □