

Baccara, Mariagiovanna; Levy, Gilat; Razin, Ronny

Working Paper
Research Waves

CESifo Working Paper, No. 12248

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Baccara, Mariagiovanna; Levy, Gilat; Razin, Ronny (2025) : Research Waves, CESifo Working Paper, No. 12248, Munich Society for the Promotion of Economic Research - CESifo GmbH, Munich

This Version is available at:

<https://hdl.handle.net/10419/334604>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

CES ifo

12248
2025

November 2025

Working Papers

Research Waves

Mariagiovanna Baccara, Gilat Levy, Ronny Razin

CES ifo

Imprint:

CESifo Working Papers

ISSN 2364-1428 (digital)

Publisher and distributor: Munich Society for the Promotion
of Economic Research - CESifo GmbH

Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740

Email office@cesifo.de
<https://www.cesifo.org>

Editor: Clemens Fuest

An electronic version of the paper may be downloaded free of charge

- from the CESifo website: www.ifo.de/en/cesifo/publications/cesifo-working-papers
- from the SSRN website: www.ssrn.com/index.cfm/en/cesifo/
- from the RePEc website: <https://ideas.repec.org/s/ces/ceswps.html>

Research Waves

Mariagiovanna Baccara* Gilat Levy[†] Ronny Razin^{‡§}

July 11, 2025

Abstract

Competing research waves start and grow as scientists choose their specialization driven by career incentives. We build a strategic experimentation framework where agents irreversibly choose between two risky fields, and information arrives faster as more agents specialize in a field. In the ‘bad news’ case, if no news arrives, all agents join a bandwagon wave into one field. In the ‘good news’ case, both fields are explored in two sequential surges, followed by slow entry into the initially inferior field. We describe how the equilibrium depends on the information-production technology, and assess the impact of first-mover advantages, congestion, and deadlines.

Keywords: Poisson Bandits, Irreversibility, Experimentation Waves, Researchers’ Incentives.

*Olin School of Business, Washington University in St. Louis. Email: mbaccara@wustl.edu

[†]London School of Economics. Email: g.levy1@lse.ac.uk

[‡]London School of Economics. Email: r.razin@lse.ac.uk

[§]We thank Arjada Bardhi, Steve Callendar, Florian Ederer, Johannes Hörner, Meg Meyer, Pietro Ortoleva, Collin Raymond, Julia Salmi, Anna Sanktjohanser, and audiences at various seminars and conferences for valuable feedback.

1 Introduction

Annual spending on research in the U.S. was, according to the latest estimate from the National Science Foundation, 885 billion dollars in 2022 (above 3% of the GDP), out of which 91.5 billion dollars were spent in academic institutions.¹ Research creates vast externalities on the rest of society, while often scholars sort themselves into research areas in a rather decentralized manner. Researchers can be driven by an intrinsic interest to explore a specific field, or by career incentives such as promotion chances, salary increases, as well as reputation and legacy.

In this paper we study how researchers choose which field to specialize in, and the resulting intensity at which competing fields are explored. We assume that scientists’ specialization choices depend on the perception among the scientific community about which field is more promising. This perception depends however on the intensity with which different fields have been explored. As more and more research output is produced and published in one area, the community obtains more information on how valuable that particular methodology is. In contrast, less information may be available about the viability of other fields that have been under-explored. The endogeneity of information uncovered in different fields plays a crucial role in understanding scientists’ equilibrium specialization incentives.

To study the equilibrium forces that determine the size, shape, and length of such endogenous “research waves”, we construct a continuous-time Poisson bandit model in which we allow agents to decide at which point in time to join one of *two risky* research fields, A or B . Considering the choice between two risky fields allows us to analyze how researchers sort themselves into alternative lines of investigation. To keep the analysis tractable, we assume that once an agent joins a field, her decision is *irreversible*. Such an assumption captures settings with large specialization costs.

At the outset, agents have prior beliefs about whether each field is ‘viable’ or not, and, while the two fields are symmetric in any other dimension, we assume that field A is initially perceived to be viable with a higher probability than B . We capture researchers’ incentives by assuming that, once an agent joins a field, she enjoys a flow payoff that depends on the perceived viability of that field. This reflects the idea that scientists’ rewards increase in the reputation of the field they are in. As an agent waits before joining a field, the flow payoff is normalized to zero.

As time passes, research output conveys information about the fields’ viability. We assume

¹See the report from the National Center for Science and Engineering Statistics (NCSES) at <https://nces.nsf.gov/pubs/nsf24317>.

that each field’s information arrival rate at any given time increases in the total mass of agents that joined that field before. Therefore, the strategic choice that agents face is, at any instance, whether to commit to a particular field or wait for others’ research outcomes in order to make a more informed choice.

An important feature of our model is the nature of the information that is revealed as more research output is generated. We study the two information environments that the literature has focused on, the “*bad news*” case and the “*good news*” case.² In the “bad news” case, news indicate that the field is not viable, and so as long as such information does not arrive, *the posterior beliefs about the field’s viability increase over time*. This scenario fits a research environment such as the pharmaceutical one, where, even if its benefits are minor, a new drug is granted FDA approval as long as its side effects are not too severe. As the new drug navigates through the sequence of experimental rounds required to be approved, the beliefs on its market viability increase over time.

In the “good news” case, information arrival conveys that the field is viable, and so in the absence of news, *the posterior beliefs about the field’s viability decrease over time*. This environment fits lines of research characterized by positive breakthroughs, such as scenarios in which scientists aim to turn a theoretical idea into a practical solution or a new product. One example is the development of vaccines (or other drugs) derived from the technology of producing proteins by introducing cells in a mixture of fat and mRNA, which was an idea first raised in the late 1980s.³ While scientists worked on this concept for decades, hopes were fading, and the scientific community started regarding mRNA as too unreliable to be used to develop new drugs or vaccines. Until, of course, the breakthrough during the Covid-19 pandemic proved the technology conclusively viable. Another recent class of examples are A.I. applications such as large language models or neural networks; The recent Physics Nobel laureate, Geoffrey Hinton, has been described as having “*dogged persistence with which he continued to believe in the potential of neural networks as a key to artificial intelligence long after the idea had been discredited by the discipline. Given how academia works, especially in a fast-developing discipline like computer science, that required exceptional determination and self-confidence.*”⁴

We show how the nature of the information environment affects the emergence and growth

²As common in this literature we focus on *conclusive* information arrival: the first arrival of news reveals whether the project is non-viable (in the bad news case) or viable (in the good news case).

³For more detail, see Nature 9/14/2021 (available at <https://www.nature.com/articles/d41586-021-02483-w>).

⁴See The Guardian, 10/12/2024 (available at <https://www.theguardian.com/commentisfree/2024/oct/12/nobel-winner-geoffrey-hinton-is-the-godfather-of-ai-heres-an-offer-he-shouldnt-refuse>).

of research fields. In the bad news case, analyzed in Section 3, in the unique equilibrium, as long as no bad news arrives on either project, all researchers specialize in field A in *one slow research wave*. The intuition is simple. Initially, field A is perceived more likely to be viable. As agents start joining A , more research is produced in this field and, in the absence on bad news on it, the beliefs about its viability increase faster than the beliefs about B 's viability. Thus, the more researchers join A , the more others want to join the same “bandwagon.”

Our main result, in Section 4, uncovers the rich equilibrium dynamics that arise in the good news case. We show that for some parameters, the unique equilibrium involves the exploration of both fields A and B in *two sequential surges and one additional slow wave*. As agents join A first, the posterior beliefs of field A 's viability decrease faster than the posterior beliefs of B 's viability, until a reversal of the posterior beliefs' ranking occurs. Then, some time after such reversal, a surge of agents into field B arises, immediately followed by more agents joining B gradually. Notably, despite the ranking reversal, the total mass of agents specializing in field B remains bounded above by the size of the initial surge into A . This equilibrium upper bound on the size of field B has two significant implications. First, no further posteriors' ranking reversal can occur in equilibrium implying that, as long as no good news arrives on A , field B remains the more promising one. Second, since field B 's size is bounded, the mass of active researchers in both risky fields may remain small relative to the whole researchers' pool. We also illustrate how the information-production technology plays a crucial role in determining the relative magnitude of the fields' sizes.

Next, in Section 5, we discuss how scientists' career incentives influence the structure of the equilibrium research waves.⁵ We allow for a first-mover advantage, according to which scientists benefit from joining a field early on. We find that a first-mover advantage reduces the likelihood of equilibrium exploration of both A and B . Next, we consider congestion externalities, assuming that the benefit of conducting research decreases when the field becomes too crowded. Finally, we consider the effects of a deadline within which scientists have to select a risky field. In both these environments we show how the likelihood of the research community exploring both A and B may increase or decrease.

We conclude our analysis in Section 6 by illustrating, using a simple example of information-production technology, how the cooperative solution (i.e., the utilitarian welfare maximizing policy) may be substantially different from the equilibrium behavior. Specifically, a utilitarian planner experiments with more agents and potentially both fields. As a result, the value of

⁵In the paper we analyze the effects of career incentives in the good news case only, while in the Online Appendix we present the corresponding analysis in the bad news case.

waiting for news is higher for the remaining agents, who, in some cases, wait longer than in the equilibrium solution.

Yet, going back to our initial motivation, the vast externalities that research generates on society as a whole may call for a more general welfare approach. In particular, society may target an optimal mix of diversification (“breadth”) and intensity (“depth”) in research. Our positive analysis highlights how the information-production technology as well as the design of career incentives provide tools to achieve socially optimal research portfolios.

Related Literature. Methodologically, our paper builds on the experimentation with bandits settings, first introduced by Thompson (1933) and Gittins (1979). Our specific setup is closer to the strategic social experimentation models explored in Bolton and Harris (1999), Keller, Rady, and Cripps (2005), Keller and Rady (2010) and Keller and Rady (2015). In all of these papers, once an agent experiments with a risky arm, she provides information externalities to others. Our model departs from this literature as we consider a choice between two risky arms, a natural modelling choice for our application. The presence of two risky arms has a substantial impact on the equilibrium analysis (particularly in the good news case) as well on the comparative statics. We also depart from the set-up of these papers by assuming that agents’ decisions are irreversible, in the sense that once an agent commits to use one risky arm, she cannot switch to the other later on. This assumption simplifies the analysis as, given the equilibrium behavior of others, each agent’s best response becomes an optimal stopping problem. In the context of social experimentation with one risky project, Laiho, Murto, and Salmi (2023a and 2023b) and Frick and Ishii (2024) also assume choices’ irreversibility.

Recent papers that account for the presence of two risky arms are Bardhi, Guo, and Strulovici (2024), who study an equilibrium application to workers’ discrimination, and Lizzeri, Shmaya, and Yariv (2024), in which a unique decision-maker is allowed to separate exploration from exploitation. Forand (2015) also focuses on a single decision-maker rather than on equilibrium analysis.

Thematically, this paper relates to two strands of the literature. The first one is the work on R&D adoption, innovation, and research collaboration. For recent examples, see Halac, Kartik, and Liu (2017), Fonseca (2024), Hopenhayn and Squintani (2021), Bonatti and Horner (2011), Callander, Lambert, and Matouschek (2023), and Bobtcheff, Lévy, and Mariotti (2024).⁶ Cetemen, Urgan, and Yariv (2022) analyze exit waves in collective search environments. The

⁶Chen et al. (2024) and Knoepfle and Salmi (2024) investigate different aspects of the problem of a central authority deciding how much information to release to the public to overcome free riding in social experimentation environments.

second strand of the literature, building on Weber (1918) and Kuhn (1962), explores how scientific knowledge evolves over time. Kitcher (1990) discusses the division of scientific labor across competing themes. Sunstein (2001) highlights how reputational concerns affect researchers’ choices, leading to scientific fads. Brock and Durlauf (1999) highlight how conformity may play a role in scientists’ specialization choices, which is an idea explored more recently, by Akerlof and Michailat (2017), and Siniscalchi and Veronesi (2021).⁷ More recent work on how new scientific ideas arise include Bramoullé and Saint-Paul (2010), Bobtcheff, Bolte, and Mariotti (2017), and Carnehl and Schneider (2024).

Finally, a few papers in the social learning literature examine information aggregation outcomes in settings where agents decide at what time to take an irreversible action when they observe the actions taken by others, but not their private information (see for example Gul and Lundholm, 1995, Rosenberg, Solan, and Vieille, 2007, and Murto and Välimäki, 2011). In our framework agents have no private information, and generate information externalities once they join a field. Therefore, our focus and analysis differ substantially from this line of literature.

2 The Model

Consider a continuous-time, infinite horizon setting, with $t \in [0, \infty)$ and a mass $\bar{m} > 0$ of identical agents present from the onset. There are two, independently distributed, risky fields, A and B , each with flow value $v > 0$ if viable, and zero if not viable. We denote the prior beliefs on the viability of the fields at $t = 0$ as p_0^A and p_0^B , respectively, and without loss of generality, we assume $p_0^A > p_0^B > 0$.

At any time t , each agent can wait, or *irreversibly* join a field $i = A, B$, getting a flow payoff $p_\tau^i v$, with $\tau \geq t$, from then on, where p_τ^i is the posterior belief of field i at τ . This assumption captures the idea that researchers may either intrinsically value conducting research in a viable field, or they may be incentivized through their promotion chances, salary increases, etc. As long as an agent does not join any field, she obtains a flow payoff of zero. This is a normalization that captures the gains from working in a ‘safe’ research field or an older research agenda which turned out to be not viable. The common discount rate is $r > 0$.

For $i = A, B$, let m_t^i be the mass of agents who already joined field i at time t . Naturally, at any point in time $m_t^A + m_t^B \leq \bar{m}$, and the irreversibility assumption implies that m_t^i weakly

⁷Recent work, such as Dossi (2024), Dossi and Morando (2023), Koning et al. (2021), and Einio et al. (2023), documents that scientists’ attributes such as gender and race might also influence their specialization choices. Also, Adda and Ottaviani (2023) highlight drawbacks in awarding grants based on relative performances.

increases in t . Information arrival is endogenous, and determined by mass of agents who already joined a field. Specifically, news on viability of each field $i = A, B$ arrives over time $t \in [0, \infty)$ following a Poisson process with arrival rate $\lambda(m_t^i)$, which is strictly increasing, continuous, and twice differentiable in m_t^i . The function $\lambda(\cdot)$ is a primitive of the environment and represents how the research technology maps agents working in a field into information production. We denote $\lambda(0) \equiv \underline{\lambda} \geq 0$, and $\lambda(\bar{m}) \equiv \bar{\lambda} > \underline{\lambda}$, so that $\underline{\lambda}$ and $\bar{\lambda}$ are the information arrival rates when nobody has joined the field, and once all agents have joined the same field, respectively. When no confusion occurs, we use λ_t^i to denote $\lambda(m_t^i)$, and since there is a one-to-one correspondence between $m_t^i \in [0, \bar{m}]$ and $\lambda(m_t^i) \in [\underline{\lambda}, \bar{\lambda}]$, we interchangeably refer to $\lambda_t^i \geq \bar{\lambda}$ as the *arrival rate* at time t for field i , and also as the total *mass of agents* necessary in a field at t to generate the arrival rate λ_t^i —that is, $\lambda^{-1}(\lambda_t^i)$.⁸

As common in this class of models, we assume that a viable field can only generate positive updates (“good news”), while a non-viable field can only generate negative updates (“bad news”), so that the first information arrival is *conclusive* evidence of the viability, or non-viability, of a field. To streamline the analysis further, we focus on two simplified information environments separately. Specifically, in the “*bad news case*,” a viable field does not generate any news, so information only arrives as negative updates. Hence, in the absence of news, *the posteriors p_t^i increase over time*. In particular, given an arrival rate λ_t^i , p_t^i evolves according to the Bayesian updating process $\dot{p}_t^i = \lambda_t^i p_t^i (1 - p_t^i)$ for $i = A, B$. In the “*good news case*,” a non-viable field generates no news, so information only arrives as positive updates. Therefore, in the absence of news, *the posteriors p_t^i decrease over time*, according to the process $\dot{p}_t^i = -\lambda_t^i p_t^i (1 - p_t^i)$.⁹

In each of these information environments, agents’ strategies consist of an optimal stopping rule—that is, when to join a field—and a choice of a field to join. We focus our analysis on Markov Perfect Equilibria in which agents’ strategies only depend on the state variables $\{p_t^A, p_t^B, m_t^A, m_t^B\}$. Thus agents’ strategies and utilities depend on others’ behavior only through the effects of others’ information externalities and hence the evolution of the posteriors. We specify expected utilities and strategies more formally for each information environment analyzed below.

⁸In our setup all agents are ex-ante identical. Agents may be heterogenous in how inclined or suited they are to join a specific field. A strictly positive ‘baseline’ information arrival rate $\underline{\lambda} > 0$ can capture the presence of some agents (who remain outside the model), whose inclination for a field is so strong that they prefer to specialize in it from the outset.

⁹One could consider a more general setting in which, for any fixed mass m_t^i of agents in field i , both negative news (conditional on non-viability) and positive news (conditional on viability) arrive at rates $\lambda^B(m_t^i)$ and $\lambda^G(m_t^i)$, respectively. We discuss this extension in Section 6.

We conclude this section with two simple lemmas that apply to both the bad news and the good news case. The first one is an immediate consequence of the law of iterated expectations.

Lemma 1 (Value of Joining a Field) *In both the bad news and the good news case, when an agent joins a field $i = A, B$ at time $t \geq 0$, her expected continuation payoff is $p_t^i \frac{v}{r}$.*

Lemma 1 guarantees that, despite the fact that the flow payoffs an agent receives after joining a field depend on the future evolution of that field's posterior, an agent's expected payoff upon joining a field does not depend on the other agents' future behavior (and the information they will generate). This observation readily implies that conditional on joining a field at a given time $t \geq 0$, it is always optimal to join the one with the highest current posterior p_t^i .

Lemma 2 (Ranking of Posterior Beliefs) *In both the bad news and the good news case, if no news have arrived, and no positive mass of agents have joined any field up to $t > 0$ (i.e., $\lambda_t^A = \lambda_t^B = \underline{\lambda}$), then $p_\tau^A > p_\tau^B$ for any $\tau \leq t$.*

Lemma 2, whose proof is straightforward and therefore omitted, ensures that, as long as no positive mass of agents have joined any field yet (and therefore the information arrival rate is still $\underline{\lambda}$ for both fields), A remains the field that is perceived the more viable one. Lemma 1 and Lemma 2 together imply that the first field to be explored in equilibrium is always field A .

Finally, in our model, agents wait in order to make a more informed choice on which field to join. Since news is conclusive, in both the bad news and the good news case, once news arrives, all agents that are still waiting immediately join a field. Naturally, if bad news arrives on a field, it is optimal to join the other, while if good news arrives on a field, it is optimal to join that same field. Hence, as optimal strategies upon news arrival are straightforward, our analysis focuses on describing agents' equilibrium behavior when no news on either field has arrived yet.

3 The Bad News Case

We start with two main observations about the bad news case. First, recall that as long as no agents have joined any field yet, if an agent joins a field she always joins A . Upon entry of some agents into A , we have $\lambda_t^A > \lambda_t^B = \underline{\lambda}$, causing field A 's posterior to increase faster than B 's one. Therefore, joining A becomes increasingly more attractive relative to joining B , creating a

“bandwagon” equilibrium feature: early researchers’ exploration of field A “encourages” others to join the same field later.

Second, since conditional on joining a field, an agent always prefers to join field A , the only reason to wait is to find out whether bad news about A arrives, which would cause a swap of the chosen field from A to B . This incentive drives the equilibrium behavior, and also suggests that there is limited strategic interaction between the fields: agents planning to join A wait for bad news coming through agents that are already in that field.

To characterize the equilibrium wave into A , we fix the strategies followed by others, which generate the arrival processes $\{\lambda_t^A\}$ and $\{\lambda_t^B\}$. Since, in the absence of news, there is no entry into field B , the process $\{\lambda_t^B\}$ remains constant at $\underline{\lambda}$. Given irreversibility, an agent’s strategy is an optimal stopping problem—that is, wait until time $s \geq 0$, and then, in the absence of news, via Lemma 1, join the field with the highest posterior, which is field A . While in the Appendix we fully specify the utility from such a strategy, here we consider how an agent evaluates an infinitesimal increase in the wait, Δs . In particular, at the optimal s , the marginal cost and benefit from such an additional wait have to be equal and satisfy

$$\lambda_s^A(1 - p_s^A)p_s^B \frac{v}{r} = p_s^A v. \quad (1)$$

In particular, the LHS of (1) represents the benefit of the additional wait: with probability $(1 - p_s^A)\lambda_s^A \Delta s$, the agent receives bad news about field A during that time, which induces her to switch her decision from joining field A to joining field B instead, yielding an expected payoff of $p_s^B \frac{v}{r}$. However, while waiting Δs , the agent is foregoing the flow income $p_s^A v \Delta s$. Hence, the RHS of (1) represents the cost of the additional wait.

Condition (1) implies that at no point $s > 0$ we can have an atom of agents joining field A at the same time. While (1) has to hold for any individual agent who chooses optimally to join A at s , once agents join in an atom at s we have an instantaneous jump upward in λ_s^A . This implies that each agent planning to join A at s should deviate and wait for the discrete information benefit generated by the others instead.

This suggests that the equilibrium must involve gradual entry, or “*trickling*,” into A for any $s > 0$. If agents are relatively impatient, so that at $s = 0$ we have $\underline{\lambda}(1 - p_0^A)p_0^B \frac{v}{r} < p_0^A v$, the equilibrium wave starts with an atom of agents $\hat{\lambda}$ joining A . The size of such an atom must ensure that (1) is satisfied from $s = 0$ onward. Letting $\bar{r} \equiv \frac{\lambda p_0^B(1 - p_0^A)}{p_0^A} \geq 0$, Proposition 1 characterizes the unique equilibrium arising in the bad news environment.

Proposition 1 (Bad News - Bandwagon Wave) *Unless bad news about A arrives early enough, no agent ever joins field B . Specifically:*

1. *If $r \leq \bar{r}$, the unique equilibrium is characterized by \underline{t}, \bar{t} , with $\bar{t} > \underline{t} \geq 0$, such that, unless bad news arrives, (i) All agents wait until \underline{t} ; (ii) In the interval $[\underline{t}, \bar{t}]$ agents gradually join field A ; (iii) At \bar{t} , everybody has joined A .*
2. *If $r > \bar{r}$, the unique equilibrium is characterized by $\bar{t} \geq 0$, such that, unless bad news arrives, (i) An atom of agents $\hat{\lambda} \in (0, \bar{\lambda}]$ joins field A at $t = 0$; (ii) If there are agents left, they gradually join field A in the interval $(0, \bar{t}]$; (iii) At \bar{t} , everybody has joined A .*

In the “bandwagon wave” equilibrium described in Proposition 1, unless bad news about A arrives before \bar{t} , field B is not explored at all. Naturally, in order to achieve gradual entry in equilibrium, (1) has to be satisfied throughout $[\underline{t}, \bar{t}]$, making the agents indifferent about joining A at any point along the wave. Therefore, λ_s^A has to increase in a specific manner in order to sustain such indifference.¹⁰

As there is limited strategic interaction between fields A and B in the bad news case, the result above is qualitatively similar to one in a model in which B is a safe project (e.g., Keller and Rady 2015, and Frick and Ishii, 2024).¹¹ Still, the analysis of the bad news scenario is instructive to build intuition and contrast condition (1) with the good news case, in which the strategic interaction between the two risky fields is a crucial driver of the equilibrium characterization.

4 The Good News Case

The analysis of the good news case departs from the bad news case in two main aspects. First, when an agent is planning to join field $i = A, B$, the incentives to wait arise from the potential arrival of good news on the other field, $-i$. The arrival of such news would induce the agent

¹⁰In the Appendix, we fully characterize the equilibrium process $\{\lambda_s^A\}$, as well as the resulting posterior beliefs $\{p_s^A\}$.

¹¹The only implication of the second field being risky is the fact that, when $\underline{\lambda} > 0$, the expected payoff obtained from switching to B upon bad news arriving on A increases over time. In the Online Appendix we show how this affects the length and speed of the bandwagon wave. As field B becomes more attractive over time, the incentive to wait for bad news about A increases. This generates a slower and longer bandwagon wave with respect to a scenario in which the second field’s payoff is fixed at a ‘safe’ level $p_0^B v$.

to switch course and join $-i$ instead. Therefore, the analysis of the good news case features a strategic interaction between the two fields that is absent from the bad news case.

Second, recall that, by Lemmas 1 and 2, the first field to be joined by any agent must be A . After an initial positive mass of agents join A , we have $\lambda_t^A > \lambda_t^B = \underline{\lambda}$. Therefore, as long as no news arrives, the posterior beliefs of field A decrease faster than the posterior beliefs of field B . Hence, the posteriors' ranking must necessarily reverse at some point. This implies that field B , which becomes the more promising one, may be explored even in the absence of news arriving on either field.

However, the potential exploration of B may yield further posteriors' ranking reversals, leading to complex equilibrium patterns. In this section we characterize the equilibrium in the good news case, uncovering rich dynamics that are driven by the presence of two risky fields. We then discuss the central role of the information-production technology $\lambda(\cdot)$ in these dynamics.

4.1 Equilibrium Analysis

Given the behavior of others, which generate the processes $\{\lambda_s^A\}$ and $\{\lambda_s^B\}$, consider an agent's strategy of *waiting until s (unless news arrives), and then join field i* , for $i = A, B$. Let $V_i(s | \{\lambda_t^A\}, \{\lambda_t^B\})$ denote the value of such strategy evaluated at time 0. For $s \geq 0$, we have:

$$V_i(s | \{\lambda_t^A\}, \{\lambda_t^B\}) \equiv \frac{v}{r} \int_0^s (p_t^A \lambda_t^A + p_t^B \lambda_t^B) e^{-\int_0^t (r + p_z^A \lambda_z^A + p_z^B \lambda_z^B) dz} dt + \frac{v}{r} \left(e^{-\int_0^s (r + p_z^A \lambda_z^A + p_z^B \lambda_z^B) dz} \right) p_s^i. \quad (2)$$

The first term of (2) describes the benefits from the arrival of good news on a field before time s . Specifically, good news arrives on field $i = A, B$ at time $t < s$ with probability $p_t^i \lambda_t^i$, upon which the agent joins that field and obtains $\frac{v}{r}$. The second term of (2) captures the scenario in which no news has arrived by time s , and the agents joins field i . Conditional on joining a field at time s , in equilibrium an agent joins the field with the highest current posterior p_s^i .¹²

¹²We use $V_i(s | \lambda^A, \lambda^B)$ to denote the utility of the same strategy conditional on *the information arrival rates in the fields A and B held constant at λ^A and λ^B* , respectively, from $t = 0$ to s .

We build the (unique) equilibrium in several steps. While the discussion below illustrates the main intuitions underlying the equilibrium structure, we provide the full characterization in Proposition 2, and further details in its proof in the Appendix.

Step 1: Initial surge into A. Consider the problem faced by an agent when no one has joined any field yet. Such first entry, which must be into field A , could occur through gradual trickling, as in the bad news case, or as an atom (i.e., a “surge”).

Consider $V_A(s|\underline{\lambda}, \underline{\lambda})$, the expected utility from joining A at time $s \geq 0$ when nobody has joined any field. At the optimal s , we have $V'_A(s|\underline{\lambda}, \underline{\lambda}) = 0$, which is equivalent to

$$\underline{\lambda} p_s^B (1 - p_s^A) \frac{v}{r} = v p_s^A. \quad (3)$$

The LHS of (3) captures the benefits of a small delay in joining A , Δs , associated with the possibility of receiving good news on field B , which occurs with probability $\underline{\lambda} p_s^B \Delta s$. Such news allows an agent to switch away from her default plan of joining field A , yielding a net benefit of $(1 - p_s^A) \frac{v}{r}$. The RHS of (3) represents the flow utility forgone while waiting before joining A , $v p_s^A \Delta s$. In contrast to the bad news case, the benefit for waiting for information is now tied to the arrival rate on B , which is fixed at $\underline{\lambda}$ as nobody has joined B yet.

We now address the possibility of a gradual trickling of agents into field A . For such a gradual entry to occur, we need λ_s^A to be constructed in such a way to satisfy (3) over a time interval. Since the benefit of waiting is to receive news on B , there is no direct effect of λ_s^A on (3). The effect of λ_s^A on (3) is instead indirect, as an increase of λ_s^A decreases p_s^A , and therefore the RHS of (3), making waiting relatively more attractive. The only way in which indifference could be maintained over time is by making p_s^B decrease fast enough relative to p_s^A , requiring $\lambda_s^B > \lambda_s^A$. Specifically, as we show in Lemma A4 in the Appendix, to satisfy (3), we need $\lambda_s^A = \lambda_s^B (1 - p_s^B)$. This however is not possible since $\lambda_s^A > \lambda_s^B = \underline{\lambda}$. Thus, gradual trickling is not feasible, implying that agents must join field A as an atom of some size $\lambda^A > \underline{\lambda}$.

Note also that if at any point in time s an agent is not willing to join A —that is, $V'_A(s|\underline{\lambda}, \underline{\lambda}) > 0$ —then, since p_s^A decreases over time, the incentive to join A becomes even lower later on, implying that the agent would rather wait indefinitely for news. Hence, $V(s|\underline{\lambda}, \underline{\lambda})$ is either maximized at $s = 0$, or at $s = \infty$. This implies that if agents are relatively patient, they all wait for news indefinitely. Otherwise, an atom of them must join field A at $t = 0$. These agents obtain a utility of $V_A(0|\underline{\lambda}, \underline{\lambda}) = p_0^A \frac{v}{r}$.

In the next steps, we focus on this latter case in which an atom of agents join A at $t = 0$, generating some information arrival rate $\lambda^A > \underline{\lambda}$.

Step 2: Potential entry into B. Given the information arrival rates $\lambda^A > \underline{\lambda}$ and $\underline{\lambda}$ for field A and B , respectively, for any fixed posteriors' level, p_s^A decreases faster than p_s^B . Then, there is a time $\widehat{s}(\lambda^A) > 0$ at which the posteriors cross, and their ranking reverses.¹³ This opens the door to the potential exploration of B at some $s \geq \widehat{s}(\lambda^A)$ despite no good news arriving on it.

We now check the further equilibrium conditions that ensure entry into B . Consider $V_B(s|\lambda^A, \underline{\lambda})$, the utility from waiting and joining field B at time s if no news has arrived yet. The potential arrival of good news on A with relatively higher frequency implies that it may be beneficial for agents to adopt such strategy. Fixing $s > 0$, note that $V_B(s|\lambda^A, \underline{\lambda})$ increases in λ^A , since the more agents join A at $s = 0$ the more attractive it is to wait for information. Yet, for agents to be willing to wait and join B at a future time s rather than to join A at $t = 0$, we need $V_B(s|\lambda^A, \underline{\lambda})$ to reach some level $V_B(s|\lambda^A, \underline{\lambda}) \geq p_0^A \frac{v}{r} = V_A(0|\underline{\lambda}, \underline{\lambda})$. In words, the value of the information accrued until that point needs to compensate them for the wait for B rather than joining A immediately.

Similarly, to ensure that the agents who join A at $t = 0$ do not regret their choice, it must also be that $p_0^A \frac{v}{r} \geq \sup_{s>0} V_B(s|\lambda^A, \underline{\lambda})$. Thus, for equilibrium entry into B to occur, we need to find some arrival rate $\hat{\lambda}^A \in (\underline{\lambda}, \bar{\lambda})$ and entry time $t^* \in [\widehat{s}(\hat{\lambda}^A), \infty)$ which jointly satisfy the following two conditions:

$$\begin{aligned} V_B(t^*|\hat{\lambda}^A, \underline{\lambda}) &= p_0^A \frac{v}{r}; \\ t^* &\in \arg \max_{s>0} V_B(s|\hat{\lambda}^A, \underline{\lambda}). \end{aligned} \tag{4}$$

In the Appendix, we show that if $\bar{\lambda}$ is sufficiently high, there is a $\hat{\lambda}^A \in (\underline{\lambda}, \bar{\lambda})$, which allows the peak of $V_B(\cdot|\hat{\lambda}^A, \underline{\lambda})$ to equal $p_0^A \frac{v}{r}$. Otherwise, no entry into B is possible in equilibrium, which results in a unique equilibrium in which all agents join A at $t = 0$. For the next steps, we assume that $\bar{\lambda}$ is sufficiently high, and hence a solution to (4) exists.

Step 3: Initial surge into B. We now establish how entry into B occurs. First, we determine the size of an initial atom λ^B joining B at t^* . For any $\lambda^B > \underline{\lambda}$, let us denote by $\tilde{\lambda}_t^B$ a process equal to $\underline{\lambda}$ for any $\tau < t$, and equal to λ^B at t .

For agents to be willing to join B at t^* it must be the case that $V_B'(t^*|\hat{\lambda}^A, \tilde{\lambda}_{t^*}^B) \leq 0$, or, equivalently, $\hat{\lambda}^A p_{t^*}^A (1 - p_{t^*}^B) \frac{v}{r} - v p_{t^*}^B \leq 0$. Otherwise, an agent that is supposed to join B at t^* would rather delay her entry. This requires $p_{t^*}^B$ not to decrease (under the rate λ^B) too fast

¹³See Lemma A1 in the Appendix for the characterization of the crossing time $\widehat{s}(\lambda^A)$.

with respect to $p_{t^*}^A$ (under the rate $\hat{\lambda}^A$) after t^* . In Lemma A4 in the Appendix we show that $V'_B(t^*|\hat{\lambda}^A, \tilde{\lambda}_{t^*}^B) \leq 0$ is equivalent to $\lambda^B \leq \hat{\lambda}^A(1 - p_{t^*}^A) < \hat{\lambda}^A$. This condition defines an upper bound on λ^B .

If all remaining agents in the market after $\hat{\lambda}^A$ joined A at $t = 0$ are below the upper bound $\hat{\lambda}^A(1 - p_{t^*}^A)$, all of them join field B at t^* . Therefore, the equilibrium consists of two sequential surges, a large one into A at $t = 0$, and a second smaller one into B at t^* .

In the next step we focus on the scenario in which, after an atom $\hat{\lambda}^B = \hat{\lambda}^A(1 - p_{t^*}^A)$ of agents joining B at t^* , there are still agents left on the market.

Step 4: Gradual trickling into B . Suppose that the conditions in (4) are satisfied, so that entry into B arises at t^* , but, after a second surge of size $\hat{\lambda}^B = \hat{\lambda}^A(1 - p_{t^*}^A)$, there are still agents left on the market. Going forward, note that, while there has been entry into B , we still have $\hat{\lambda}^B < \hat{\lambda}^A$. This guarantees the posterior of B to decrease more slowly than that of A . Therefore, B remains the more promising field, implying that the next entry must be into B .

For any additional entry we need that the utility of entrants to be exactly the same as that of all agents that joined any field before them. The continuing process $\{\lambda_s^B\}$ must then satisfy $V_B(s|\hat{\lambda}^A, \{\lambda_s^B\}) = p_0^A \frac{v}{r}$ or alternatively $V'_B(s|\hat{\lambda}^A, \{\lambda_s^B\}) = 0$ for all $s \geq t^*$ at which such entry occurs. This implies that, following t^* , agents gradually trickle into B .

Note that such gradual entry into B is in contrast with Step 1, in which we ruled out trickling into A . To see why trickling into B is feasible, note that the indifference condition underlying trickling into B over some interval $[t^*, t)$ is

$$\hat{\lambda}^A p_s^A (1 - p_s^B) \frac{v}{r} = v p_s^B \quad (5)$$

for any $s \in [t^*, t)$. Similarly to Step 1, λ_s^B does not enter (5) directly, but through its influence on the decline of p_s^B . To maintain (5) for any $s \in [t^*, t)$, we need to counter the effect of $\hat{\lambda}^A$ on p_s^A with a sufficiently slow decrease in p_s^B . As we show in the Appendix, this can be achieved by setting the process $\lambda_s^B = \hat{\lambda}^A(1 - p_s^A) < \hat{\lambda}^A$. In this way, agents trickle into B over an interval of time.

Step 5: No further reversals and behavior at the limit. In the previous step, we illustrated how the equilibrium entry process must satisfy $\lambda_s^B < \hat{\lambda}^A$ for any $s \geq t^*$. The fact that the information arrival rate in field B remains lower than the one in A allows us to complete the equilibrium construction. Specifically, since $p_{t^*}^B > p_{t^*}^A$, the posterior beliefs of A and B induced by $\hat{\lambda}^A$ and $\{\lambda_s^B\}$, respectively, do not cross again. Thus, there are *no more posterior*

reversals, and, in the absence of news, B remains the more promising field. This implies that, in the absence of news, all remaining agents prefer to join B rather than A . Let \bar{t} be the time at which the trickling process $\lambda_s^B = \widehat{\lambda}^A(1 - p_s^A)$ exhausts all remaining agents.

Next, the fact that the entry process into B is bounded implies that there could be agents waiting indefinitely for news on either field. In particular, in such scenario, \bar{t} grows arbitrarily large, and p_s^A and λ_s^B converge to zero and $\widehat{\lambda}^A$, respectively. The continued trickling into field B guarantees that indifference is satisfied at the limit.

Proposition 2 summarizes our equilibrium construction and fully characterizes the unique equilibrium in the good news case.¹⁴

Proposition 2 (Good News - Sequential Entry) *There exists $\tilde{r} \geq \bar{r}$ such that*

1. *If $r < \tilde{r}$, all agents wait indefinitely for news;*
2. *If $r > \tilde{r}$, there exists $\widehat{\lambda}^A > \underline{\lambda}$ such that:*
 - (a) *If $\bar{\lambda} \leq \widehat{\lambda}^A$, all agents join field A at $t = 0$;*
 - (b) *If $\bar{\lambda} > \widehat{\lambda}^A$, then (i) An atom of agents $\widehat{\lambda}^A$ joins A at $t = 0$; (ii) Unless news arrives, an atom of agents $\widehat{\lambda}^B < \widehat{\lambda}^A$ join B at some $t^* > \widehat{s}(\widehat{\lambda}^A)$; (iii) If there are agents left, they join B gradually in the interval $[t^*, \bar{t}]$, with $\bar{t} \in (t^*, \infty]$ and $\lambda_s^B < \widehat{\lambda}^A$ for all $s \in (t^*, \bar{t})$; (iv) If $\bar{t} = \infty$, a remaining atom of agents waits indefinitely for news.*

Note that in the absence of news, the equilibrium mass of agents in field B is bound to remain smaller than the one in field A . This occurs despite the fact that, after the equilibrium ranking reversal at time $\widehat{s}(\widehat{\lambda}^A)$, field B is indefinitely perceived as the most promising one. Therefore, the ranking of the priors determines the fields' relative sizes for the long run.

We continue to explore the equilibrium field sizes in the next section, by illustrating the central role of the information-production technology $\lambda(\cdot)$ in determining the equilibrium dynamic patterns. To do so, we henceforth denote the equilibrium in which both fields A and B are explored, described in Part 2(b) of Proposition 2, as the “ (A, B) equilibrium.”

¹⁴The proof in the Appendix shows that the equilibrium is unique, and the cases accounted for in Proposition 2 exhaust all parameters. For uniqueness, we ignore the non-generic case $r = \tilde{r}$, for which both equilibria described in Parts 1 and 2 of Proposition 2 exist.

4.2 Waves Magnitudes

The previous section highlights the rich dynamics that unfold in equilibrium in the good news case. In this section we highlight how such equilibrium structure depends on the shape of the information-production technology $\lambda(\cdot)$, associated with the specific features of the research environment, and on the prior beliefs.

From Proposition 2, it is apparent that the information-production technology efficacy $\bar{\lambda}$ (or equivalently the population size \bar{m}) ultimately determines the extent to which field B is explored. If $\bar{\lambda}$ is too small, all agents join field A at $t = 0$, and the (A, B) equilibrium does not arise.

The shape of the information-production technology $\lambda(\cdot)$ also plays an important role in determining the type of (A, B) equilibrium that arises. To illustrate such a role, given \bar{m} and $\bar{\lambda}$, the left panel of Figure 1 depicts a convex $\lambda(\cdot)$. For example, in the hard sciences, a convex $\lambda(\cdot)$ captures laboratories needing some critical mass of personnel overlooking the day-to-day operations before being able to conduct meaningful research. Thus, a relatively large mass of agents need to join A to obtain the arrival rate $\hat{\lambda}^A$, which is determined in equilibrium by condition (4). This leaves only a few agents available to join field B , as in case (ii) of Part 2(b) of Proposition 2. Hence, the resulting equilibrium distribution of agents into the two fields is relatively asymmetric. Moreover, all agents always join one of the two fields in a finite time.

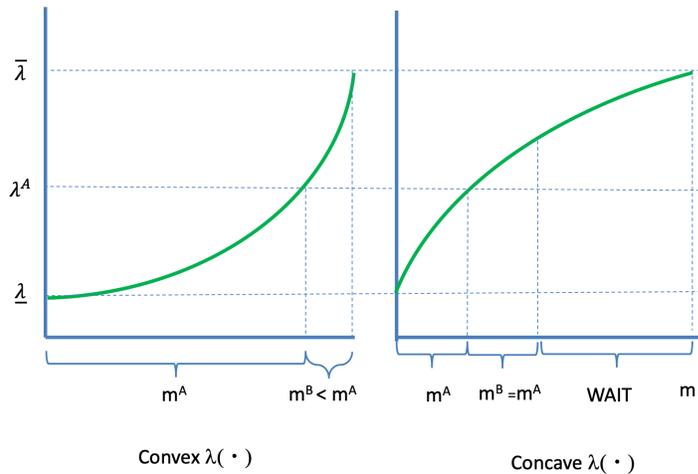


Figure 1: Impact of $\lambda(\cdot)$ on the equilibrium in the good news case. m^A and m^B denote the equilibrium sizes of fields A and B , respectively.

Alternatively, for the same \bar{m} and $\bar{\lambda}$, consider a sufficiently concave $\lambda(\cdot)$, as in the right panel of Figure 1, characterizing a research environment in which significant progress can be achieved by a relatively smaller set of agents. Since in this case fewer agents are required to reach $\hat{\lambda}^A$, in the absence of news, field B could eventually grow to the same size of A , as described in case (iv) of Part 2(b) of Proposition 2. The equilibrium distribution of agents into the two fields is then relatively symmetric. However, such equilibrium involves fewer active agents, and a significant amount of them waiting indefinitely for news.

Next, we fix the technology and change the prior belief on B . This exercise directly affects the value of $\hat{\lambda}^A$, as determined by (4).

Proposition 3 (Good News - Change in Prior Beliefs) *Consider the (A, B) equilibrium.*

As p_0^B increases, $\hat{\lambda}^A$ decreases. As a result, as long as no news arrives, the overall amount of active researchers decreases.

Note that an increase of p_0^B does not affect the payoff from joining A at $t = 0$, $V_A(0|\underline{\lambda}, \underline{\lambda}) = p_0^A \frac{v}{r}$. However, everything else held constant, researchers are more inclined to wait before they choose between fields— i.e. $V_B(s|\hat{\lambda}^A, \underline{\lambda})$ shifts upward for all $s > 0$. This implies that a lower $\hat{\lambda}^A$ is required to guarantee indifference between joining field A at $t = 0$ and waiting to join B later.

Proposition 3 highlights the fact that a decrease of $\hat{\lambda}^A$ constitutes a *double-edged sword* with respect to the exploration of field B : A smaller $\hat{\lambda}^A$ allows for a larger set of agents that can explore B later on. However, since in equilibrium the size of field B is bounded above by the size of field A , a decrease of $\hat{\lambda}^A$ reduces the maximal size that field B can reach. Therefore, an increase of p_0^B causes a reduction in the amount of researchers who join any of the two fields in the absence of news arrival.

5 Career Incentives

As seen in the previous sections, our model allows for a tractable equilibrium analysis in the presence of two risky fields. We now further illustrate the richness of the model by incorporating several institutional features that influence researchers' career incentives, and evaluating their impact on our results. First, we examine two alternative payoff structures. In the first one we allow for a *first-mover advantage* experienced by agents who join a field early, and in the second one we introduce *congestion externalities* experienced by agents in densely populated

fields. Then, we consider *deadlines* within which agents must make a decision on which field to join. Throughout the analysis, we focus on the good news case.¹⁵

5.1 First-Mover Advantage

Researchers who are among the first to engage in a new field of investigation tend to obtain more recognition. For example, these researchers are cited more often, and citations' counts are sometimes used to assess the impact of a researcher's work and to determine promotions and salary increases.

Consider the following simple model of first-mover advantage for joining an unexplored field first. Suppose that if a researcher joins a field $i = A, B$ at time $t \geq 0$, her payoff is, for some $\alpha > 1$,

$$U_t^i = \begin{cases} \alpha p_{t\tau}^i v & \text{if } m_\tau^i = 0 \text{ for } \tau < t; \\ p_{t\tau}^i v & \text{otherwise.} \end{cases}$$

In words, if no positive mass of agents joined i before t , her flow utility from then onward is $\alpha p_s^i v$ for $s \geq t$, where $\alpha > 1$. If a researcher joins a field where a positive mass of researchers is already present, her payoffs are the same as in the standard model described in Section 2 (which corresponds to the $\alpha = 1$ case). For simplicity, we assume that if multiple researchers join field i first at the same time they each receive the same flow utility $\alpha p_\tau^i v$ —that is, there are no congestion externalities (which we study separately in the next section).

At first glance, one could conjecture that after a positive mass of agents joined A , since joining A does not entail the benefit α any longer, while B still does, it becomes more attractive for others to join B . Then, according to this intuition, a first-mover advantage should make the (A, B) equilibrium more likely to occur. As it turns out, this is not the case. In the (A, B) equilibrium, those who join A at time 0 gain α . Those who wait to join B do so in order to receive good news on A . However, entry into A as second-movers is now less valuable, resulting in lower incentives to wait. Therefore, the (A, B) equilibrium is harder to sustain.

Proposition 4 (Good News - First-Mover Advantage) *In the presence of a first-mover advantage $\alpha > 1$, the equilibrium in which all agents join field A immediately is more likely to hold, while the (A, B) equilibrium is less likely to hold.*

¹⁵The analysis of these institutional environments in the bad news case, as well as the description of the impact of *subsidies*, such as scientific grants, on the equilibrium structure both in the bad and the good news case are all available in the Online Appendix.

5.2 Congestion Externalities

Another way in which researchers' payoffs may depend on the presence of others in their field is through congestion externalities. For example, when a field becomes too crowded, it may become harder for a scholar to differentiate themselves, find new ideas, or be cited and acknowledged. Suppose that if a researcher joins a field $i = A, B$ at time $t \geq 0$, her flow payoff is, for some $\gamma \in [0, 1)$, $\hat{m} \in (0, \bar{m})$, and for any $\tau \geq t$,

$$U_\tau^i = \begin{cases} p_\tau^i v & \text{if } m_\tau^i \leq \hat{m}; \\ \gamma p_\tau^i v & \text{otherwise.} \end{cases}$$

In words, the flow payoff the agent receives after joining a field decreases when too many agent join the same field.

The strategic incentives in the presence of congestion externalities are substantially more nuanced than in the benchmark model. First, note that congestion affect an agent's utility from waiting for news. In particular, if the mass of agents who are waiting with her is larger than \hat{m} , she will experience congestion externalities upon good news arrival. This reduces the benefit of waiting. Second, the utility from joining a field is also adversely affected by congestion. Specifically, an agent who is already active in a field will experience a lower utility if too many join the same field in the future.¹⁶

The combined effects of congestion externalities also depend on the information-production technology $\lambda(\cdot)$, as the next results illustrates.

Proposition 5 (Good News - Congestion Externalities) *In the presence of congestion externalities $\gamma \in (p_0^A, 1)$, we have:*

1. *The equilibrium in which all wait indefinitely for news is less likely to hold;*
2. *If $\lambda(\cdot)$ is sufficiently convex, the (A, B) equilibrium is more likely to hold;*
3. *If $\lambda(\cdot)$ is sufficiently concave, and if $\underline{\lambda}$ and p_0^A are sufficiently small and large, respectively, the (A, B) equilibrium is less likely to hold.*

To understand the intuition of Proposition 5, note that when congestion externalities are small in size (i.e., γ is sufficiently close to 1), upon arrival of good news on a field, all remaining agents still join that field.

¹⁶This implies that in the presence of congestion externalities, Lemma 1 does not hold.

Part 1 of Proposition 5 is an immediate consequence of the fact that, if all agents wait indefinitely for news, all agents join the field for which good news arrives first. As congestion externalities always arise, a deviation of joining A earlier becomes more appealing relative to the $\gamma = 1$ case.

Next, consider the (A, B) equilibrium, which is the focus of Parts 2 and 3 of Proposition 5. Let $\lambda(\cdot)$ be sufficiently convex (as in the left panel of Figure 1), so that a large mass of agents, exceeding \hat{m} , is required to attain $\hat{\lambda}^A$. The payoff of joining A at $t = 0$ is strictly lower than in the benchmark model, since agents who join A always experience congestion. However, agents who wait to join B do not always experience congestion. If good news arrives on A while they wait, they experience congestion when they join A . If instead news arrives on B , they avoid congestion. This is because only a small mass of agents remain following the initial entry to A . Hence, joining field B becomes more attractive relative to the benchmark case, implying that the (A, B) equilibrium is more likely to hold.

Assume now a sufficiently concave $\lambda(\cdot)$, such that the mass of agents required to achieve $\hat{\lambda}^A$ is strictly lower than \hat{m} . Moreover, as in the (A, B) equilibrium $\lambda_s^B < \hat{\lambda}^A$ for all s , a sufficiently concave $\lambda(\cdot)$ also implies that throughout the equilibrium process the mass of agents waiting for news is always larger than \hat{m} . Thus, whenever news arrives on some field, all these agents join that field, generating congestion.

Therefore, the agents who join A at $t = 0$ only experience congestion when good news on A arrives before it arrives on B . Conversely, agents who join B experience it when good news arrives on B first. In the Appendix we show that in some environments (e.g., when p_0^A is not too close to p_0^B), news on B is more likely to arrive first. Since in these cases the utility from joining B is lower, the (A, B) equilibrium is harder to sustain.

5.3 Deadlines

Under a tenure-clock system, researchers are on a track leading to a “up or out” promotion decision. Hence, they are incentivized to select a high-stake risky field of investigation rather than work in a safe one. Such choice needs to be made early enough for them to be able to build a record in their chosen field before their “clock” expires. Similarly, Ph.D. students must decide early in which area to specialize, in order to be able to seek employment at the end of the program.

To explore these scenarios, we introduce a deadline $0 < T < \infty$ such that if an agent does not join either A or B by time T , they receive a zero flow payoff from that point on. We assume

that all agents in our model belong to a specific cohort to which such deadline applies.

One obvious implication of a deadline is that no agent can wait for news indefinitely. Therefore, whenever B is explored in equilibrium, all remaining agents at T join B at that time. This implies that *the size of field B is no longer bounded above by the size of field A* . For example, for concave enough $\lambda(\cdot)$, field B is larger than field A from T onwards.

The presence of a deadline has other, more subtle, implications on the equilibrium structure before the deadline's expiration. First, consider the environments for which, as in Part 1 of Proposition 2, absent a deadline, agents wait for news indefinitely rather than join field A at $t = 0$. In the presence of a deadline, agents now compare the utility of joining A at $t = 0$ —that is, $p_0^A \frac{v}{r}$ —to the utility of joining A at T , $V_A(T|\underline{\lambda}, \underline{\lambda})$. If $V_A(T|\underline{\lambda}, \underline{\lambda}) < p_0^A \frac{v}{r}$, some agents join A at $t = 0$, changing the whole equilibrium structure.

Next, consider the environments in which some agents join A at $t = 0$. In Part 2(b) of Proposition 2, some agents plan to join field B from t^* onwards. If $T < t^*$, $V_B(s|\widehat{\lambda}^A, \underline{\lambda})$ needs to reach the value $\frac{p_0^A v}{r}$ in a shorter period of time to enable entry in B . As a result, the (A, B) equilibrium is harder to sustain.

Moreover, if T is extremely short, it could fall before the posteriors of A and B can ever cross. In this case, a new form of equilibrium sometimes arises, constituted by two sequential surges into A . Specifically, in such equilibrium some agents join A at $t = 0$ and all others join A at T .

The following result formalizes the discussion above.

Proposition 6 (Good News - Deadlines) *In the presence of a deadline $0 < T < \infty$, there exist $\tilde{r}(T)$ strictly increasing in T , such that the unique equilibrium is as follows:*

1. *If $r < \tilde{r}(T)$, all agents join field A at T .*
2. *If $r \geq \tilde{r}(T)$, and $T \in (0, t^*)$, then the equilibrium takes one of the following forms:*
 - (a) *All agents join field A , either all at $t = 0$, or some at $t = 0$ and all others at $t = T$;*
 - (b) *An atom $\widehat{\lambda}^A(T) < \bar{\lambda}$ of agents join field A at $t = 0$, and all others join field B at $t = T$;*
3. *If $r \geq \tilde{r}(T)$, and $T \geq t^*$, the equilibrium follows Part 2 of Proposition 2, except that if there are agents still remaining at T , they all join field B at $t = T$.*

6 Discussion

We conclude the paper by first discussing some of the normative implications of our model, and then suggesting some directions for future research.

6.1 Welfare

In this paper we have taken a positive approach and analyzed a model in which researchers, driven by career concerns, sort themselves into fields of investigation. The natural next step is to study the normative properties of our model. The standard normative approach in this line of literature is a utilitarian one, which considers the aggregate utility of all researchers, and accounts for the information externalities they generate for each other. An additional possibility, which is natural for our application, is to think about social welfare, which incorporates the spillovers of research on the society as a whole.

6.1.1 Utilitarian Welfare

The utilitarian approach to welfare amounts to taking the perspective of an academic or professional association, which is tasked with maximizing the aggregate utilities of scientists over their careers. To explore this approach we assume that the social planner is able to allocate researchers to fields over time, while being constrained by all the other features of our environment (e.g., irreversibility and the information-production technology $\lambda(\cdot)$).

To illustrate the key externalities in our model and how a social planner can improve on the equilibrium welfare, we focus on the good news case, and we use a simple example of information-production technology. In particular, let us normalize $\bar{m} = 1$, and, given some $\varepsilon \in (0, 1)$, consider the following step function $\check{\lambda}(\cdot)$:

$$\check{\lambda}(m) = \begin{cases} 0 & \text{if } 0 \leq m < \varepsilon; \\ \bar{\lambda} & \text{if } \varepsilon \leq m \leq 1. \end{cases}$$

With this technology a measure ε of researchers into a field is sufficient—and necessary—to extract information on that field. Since $\underline{\lambda} = 0$ implies that waiting without experimenting is wasteful, in any optimal solution the planner must allocate at least a ε -measured mass of researchers to a field at $t = 0$.

Since $\check{\lambda}(\cdot)$ is not continuous, to compare the planner's solution to equilibrium behavior, we refer to the equilibrium analysis for a class of continuous technologies $\lambda(\cdot)$ converging to $\check{\lambda}(\cdot)$.

From Proposition 2, when $\underline{\lambda} = 0$, either all agents join A at $t = 0$, or, if $\bar{\lambda}$ is high enough, the (A, B) equilibrium arises. Also, recall that it is always the case that the equilibrium utilitarian welfare is $p_0^A \frac{v}{r}$ (which is independent on $\lambda(\cdot)$). The proposition below illustrates how the planner can improve on the equilibrium welfare level. Naturally, when good news arrives on a field, the planner allocates all remaining agents to that field. Hence, we describe the utilitarian optimal policy as long as no news has arrived yet.

Proposition 7 (Utilitarian Optimal Policy) *With the information-production technology*

$\check{\lambda}(\cdot)$, *there exist $\bar{r}^{opt} > 0$ and $\tilde{r}^{opt} > \bar{r}^{opt}$ such that:*

1. *If $r > \tilde{r}^{opt}$, or $r \in [\bar{r}^{opt}, \tilde{r}^{opt}]$ and ε is sufficiently large, a no-experimentation policy, in which all agents are allocated to field A at $t = 0$, is optimal. The utilitarian welfare is $p_0^A \frac{v}{r}$;*
2. *If $r < \tilde{r}^{opt}$ and ε is sufficiently small, a simultaneous double-experimentation policy is optimal. In such a policy, a measure ε of agents are allocated to both A and B at $t = 0$, and all others wait for news indefinitely. The utilitarian welfare is strictly higher than $p_0^A \frac{v}{r}$;*
3. *If $r < \bar{r}^{opt}$ and ε is sufficiently large, a single-experimentation policy is optimal. In such a policy, a measure ε of agents are allocated to A at $t = 0$ and all others to B at some t_B^{opt} . The utilitarian welfare is strictly higher than $p_0^A \frac{v}{r}$.*

Proposition 7 highlights two ways in which the planner's optimal policy differs from equilibrium behavior. First, the planner can experiment across more fields: As described in Part 2, when ε is small enough, the planner experiments simultaneously both in field A and B . The fact that experimentation is cheap allows the planner to maximize the probability of information arrival and hence her ability to allocate the rest of the agents to a viable field as soon as possible. This also implies that, in the absence of news, the remaining agents wait longer than in the equilibrium solution. Second, Part 3 of Proposition 7 illustrates that for a large ε , even if the planner uses a policy reminiscent of the (A, B) equilibrium, its specific features are different. In particular, the planner experiments with more agents at $t = 0$. To see why, note that the planner is not restricted by the equilibrium indifference condition of keeping the expected utility of those who wait to join B equal to $p_0^A \frac{v}{r}$. As seen in Proposition 2, such a constraint bounds the equilibrium information arrival rate in field A to some $\hat{\lambda}^A < \bar{\lambda}$.

6.1.2 Societal Welfare

Scientific research generates externalities to the society at large which go far beyond the boundaries of the researchers' community. This is partly the reason why governments are substantially engaged in research funding, directly through public institutions, or indirectly through grants and other channels.

When we think of the society at large, a critical aspect of welfare is that of research diversification. Researchers may not necessarily internalize the social value of developing a diverse portfolio of research methodologies, each of which could be useful to face a future crisis, such as climate change, or a potential viral pandemic. In other words, society might want to hedge optimally using different research fields.¹⁷

The equilibria explored in this paper have very different properties with respect to research diversification. Thus, our analysis sheds light on what features of the information-production technology and what types of incentives within the scientific community influence research diversification, and could therefore become useful tools to design socially optimal research portfolios.¹⁸

6.2 Extensions and Further Research

Methodologically, our framework provides a tractable setup to study strategic exploration in the presence of two risky fields. In what follows, we discuss several directions in which our study can be extended.

Value-Added Bonuses. Researchers often obtain higher payoffs for contributions that have a big impact on how the rest of the scientific community views the field. To capture this feature in our model, we can assume that there is an additional benefit when an agent joins a field before news arrives, as opposed to after. Both in the bad and in the good news case, the agents' incentives to wait for news relative to joining a field are reduced. This makes an agent *less patient* and so less willing to wait. It is easy to check then that in the bad news case the bandwagon wave is faster than in our set-up, while in the good news case, the (A, B) equilibrium is less likely to hold.

¹⁷As an example, consider research about vaccines: Since it is uncertain which virus will be the source of the next pandemic, it could be optimal from a social perspective to conduct parallel research agendas on vaccines against, say, virus A and virus B . For the optimality of policies that can mitigate multiple social catastrophes, see Martin and Pindyck (2015).

¹⁸Recent work by Hill and Stein (2023 and 2024) empirically quantifies the impact of scientists' incentives on "breadth vs. depth" of research.

Editorial influence. Scientific journals play a critical role in determining what researchers choose to work on. Editorial choices about publishing “breakthroughs” vs. “breakdowns” are one source of such influence (e.g., journals may decide to encourage submissions of negative results). In our model this would translate into endogenizing the nature of the information arrival in the two fields, which is an interesting extension for future research.

Arrival of both bad and good news. As mentioned in Section 2, we could consider a more general setting in which, for any mass m of agents in a field, both conclusive negative news (conditional on non-viability) and conclusive positive news (conditional on viability) can arrive at rates $\lambda^B(m)$ and $\lambda^G(m)$, respectively. Whenever $\lambda^B(m) > \lambda^G(m)$ ($\lambda^B(m) < \lambda^G(m)$) for all $m \in [0, \bar{m}]$ posteriors increase (decrease) in the absence of news, which corresponds to a bad (good) news environment. However, additional subtleties stem from the fact that when an agent waits for information while planning to join, say, field A , she will change her course of action both if she receives bad news on A and good news on B . While this consideration affects the incentives to wait in both the bad and the good news case, the specific results would also depend on the assumptions made on the functions $\lambda^B(\cdot)$ and $\lambda^G(\cdot)$.

Cross-cohorts dynamics. Our model can be interpreted as capturing the decisions of a cohort of researchers over time, and is characterized by a symmetric ‘baseline’ arrival rate for the two fields, $\underline{\lambda} \geq 0$. One could think of an intergenerational model with potentially asymmetric rates $\underline{\lambda}^A \geq \underline{\lambda}$ and $\underline{\lambda}^B \geq \underline{\lambda}$, which depend on the information uncovered by more senior researchers who made their specialization choices in the past, and are still actively engaged in research. While older fields could cease to exist either because bad news arrived on them, or because, in the good news case, their posteriors have converged over time to very low levels, new research fields can potentially emerge with information arrival rates ‘reset’ at $\underline{\lambda} \geq 0$, until some young researchers start their exploration. Such extension could yield intergenerational bandwagon waves in the bad news case, or, in the good news case, to scenarios in which, since $\underline{\lambda}^A > \underline{\lambda}^B$ at the offset of a new generation, the choices made by more senior scientists facilitate entry into newer fields in the future. This can allow us to obtain a natural lifecycle of fields across cohorts.¹⁹

¹⁹Relatedly, one can think of a constant flow of new agents entering the scientific community. As long as this flow is robust enough to sustain the equilibrium arrival rates that we identified throughout, our results still hold.

7 Appendix

In this section we present all the proofs of the results presented in the paper. Additional analysis and results can be found in the Online Appendix.

7.1 Proofs for the Bad News Case

Proof of Proposition 1: We fix the strategies followed by others, which generate the arrival processes $\{\lambda_t^A\}$ and $\{\lambda_t^B\}$. Given irreversibility, the only possible plan of action for an agent is to wait until time $s \geq 0$, and then, in the absence of news, by Lemma 1, to join the field with the highest posterior p_s^i . The utility from such a strategy evaluated at $t = 0$ is

$$V(s) \equiv \frac{v}{r} \int_0^s [(1 - p_t^A)\lambda_t^A p_t^B + (1 - p_t^B)\lambda_t^B p_t^A] e^{-\int_0^t [r + (1 - p_z^A)\lambda_z^A + (1 - p_z^B)\lambda_z^B] dz} dt + \frac{v}{r} \left[e^{-\int_0^s [r + (1 - p_z^A)\lambda_z^A + (1 - p_z^B)\lambda_z^B] dz} \right] \max_{i=A,B} p_s^i. \quad (6)$$

In words, the first term is associated with the arrival of bad news on a field before time s , when the agent immediately joins the other field. Specifically, bad news arrives on field $i = A, B$ at any time $t < s$ with probability $(1 - p_t^i)\lambda_t^i$, at which point the agent joins $-i$ and obtains the expected utility $p_t^{-i} \frac{v}{r}$. The second term captures the scenario in which no news has arrived by period s , and the agents joins the field with the highest posterior, obtaining the expected utility $\frac{v}{r} \max_{i=A,B} p_s^i$.

(1) Let $r \leq \bar{r}$, and conjecture an equilibrium as the one described in the claim. In such equilibrium $p_t^A > p_t^B$ for all $t \geq 0$, and therefore, by Lemma 1, nobody has an incentive to ever join field B . Since we need agents to be indifferent between joining A at any $s \in [\underline{t}, \bar{t}]$, we need (6) to be constant over the interval—that is, $V(s) = \kappa$ for all $s \in [\underline{t}, \bar{t}]$. A necessary condition for this to be true is $V'(s) = 0$. After some algebraic steps, such condition yields the ODE

$$\frac{\dot{p}_s^A}{(p_s^A)^2} = \frac{r}{p_s^B}, \quad (7)$$

with solution

$$p_s^A = \frac{1}{-\int \frac{r}{p_s^B} ds - C} \quad (8)$$

for some constant C . We can use (8) to solve for the equilibrium λ_s^A , obtaining

$$\lambda_s^A = \frac{\dot{p}_s^A}{p_s^A(1-p_s^A)} = \frac{r}{-p_s^B \left(\int \frac{r}{p_s^B} ds + C + 1 \right)}. \quad (9)$$

Using $p_t^B = \frac{p_0^B}{(1-p_0^B)e^{-\lambda t} + p_0^B}$, we can express p_s^A and λ_s^A as a function of C and primitives, as

$$p_s^A(C) = \frac{\underline{\lambda} p_0^B}{-\underline{\lambda} p_0^B r s + r(1-p_0^B)e^{-\lambda s} - \underline{\lambda} p_0^B C},$$

and

$$\lambda_s^A(C) = \frac{-r\underline{\lambda} [(1-p_0^B)e^{-\lambda s} + p_0^B]}{r s \underline{\lambda} p_0^B - r e^{-\lambda s} (1-p_0^B) + \underline{\lambda} p_0^B (C+1)}. \quad (10)$$

Next, note that condition (7) needs to be satisfied at the two interval extremes. At \underline{t} , (7) implies $\frac{\underline{\lambda}(1-p_{\underline{t}}^A)}{p_{\underline{t}}^A} = \frac{r}{p_{\underline{t}}^B}$. From $\frac{p_{\underline{t}}^j}{1-p_{\underline{t}}^j} = \frac{p_0^j}{1-p_0^j} e^{\int_0^{\underline{t}} \lambda dz} = \frac{p_0^j e^{\underline{\lambda} \underline{t}}}{1-p_0^j}$ for $j = A, B$, we obtain

$$\frac{r p_0^A}{\underline{\lambda}(1-p_0^A) - r p_0^A e^{\underline{\lambda} \underline{t}}} = \frac{p_0^B}{1-p_0^B}. \quad (11)$$

which, solving for \underline{t} , yields

$$\underline{t} = \frac{1}{\underline{\lambda}} \ln \left[\frac{(1-p_0^A) \underline{\lambda}}{r p_0^A} - \frac{1-p_0^B}{p_0^B} \right] \geq 0, \quad (12)$$

where the last inequality is guaranteed by $r \leq \bar{r}$. Next, (10) at \bar{t} implies

$$C(\bar{t}) = -\frac{r}{\bar{\lambda}} - r\bar{t} + \frac{r(1-p_0^B)e^{-\bar{\lambda}\bar{t}}}{p_0^B} \left(\frac{1}{\bar{\lambda}} - \frac{1}{\underline{\lambda}} \right) - 1.$$

Finally, to find \bar{t} and C , note that (7) at \bar{t} implies

$$\frac{\bar{\lambda} p_0^B}{\bar{\lambda} p_0^B + r [p_0^B + (1-p_0^B)e^{-\bar{t}\bar{\lambda}}]} = \frac{p_0^A}{p_0^A + (1-p_0^A)e^{-\int_{\underline{t}}^{\bar{t}} \lambda dz}}.$$

Since

$$\frac{\partial \ln(-r \underline{\lambda} p_0^B z + r(1-p_0^B)e^{-\lambda z} - \underline{\lambda} p_0^B (C+1))}{\partial z} = \frac{-r\underline{\lambda} [(1-p_0^B)e^{-\lambda z} + p_0^B]}{-r \underline{\lambda} p_0^B z + r(1-p_0^B)e^{-\lambda z} - \underline{\lambda} p_0^B (C+1)},$$

we have

$$\begin{aligned}
& \frac{p_0^A}{p_0^A + (1 - p_0^A)e^{-\ln\left[\frac{(1-p_0^A)\lambda}{rp_0^A} - \frac{1-p_0^B}{p_0^B}\right] + \int_{\underline{t}}^{\bar{t}} \frac{-r\lambda[(1-p_0^B)e^{-\lambda z} + p_0^B]}{-r\lambda p_0^B z + r(1-p_0^B)e^{-\lambda z} - \lambda p_0^B(C+1)} dz}} \\
&= \frac{p_0^A}{p_0^A + (1 - p_0^A) \frac{(-r\lambda p_0^B \bar{t} + r(1-p_0^B)e^{-\lambda \bar{t}} - \lambda p_0^B(C+1))}{(-r\lambda p_0^B \underline{t} + r(1-p_0^B)e^{-\lambda \underline{t}} - \lambda p_0^B(C+1))} e^{-\lambda \underline{t}}}.
\end{aligned}$$

Therefore, we obtain the condition

$$\begin{aligned}
& \bar{\lambda} p_0^B (1 - p_0^A) \left[-r\lambda p_0^B \bar{t} + r(1 - p_0^B) e^{-\lambda \bar{t}} - \lambda p_0^B (C + 1) \right] e^{-\lambda \underline{t}} \\
&= p_0^A r \left[p_0^B + (1 - p_0^B) e^{-\lambda \bar{t}} \right] \left[-r\lambda p_0^B \underline{t} + r(1 - p_0^B) e^{-\lambda \underline{t}} - \lambda p_0^B (C + 1) \right].
\end{aligned}$$

Since

$$C(\bar{t}) = -\frac{r}{\lambda} - r\bar{t} + \frac{r(1 - p_0^B) e^{-\lambda \bar{t}}}{p_0^B} \left(\frac{1}{\lambda} - \frac{1}{\bar{\lambda}} \right) - 1,$$

the condition above becomes

$$\frac{p_0^B (1 - p_0^A)}{p_0^A} = \frac{-r\lambda p_0^B \underline{t} + r(1 - p_0^B) e^{-\lambda \underline{t}} - r(1 - p_0^B) \left(1 - \frac{\lambda}{\bar{\lambda}}\right) e^{-\lambda \bar{t}} + \lambda p_0^B r \bar{t} + \frac{\lambda}{\bar{\lambda}} r p_0^B}{\lambda e^{-\lambda \underline{t}}}. \quad (13)$$

Recall that (11) implies

$$\frac{\lambda(1 - p_0^A) p_0^B}{p_0^A} = e^{t\lambda} r p_0^B + r(1 - p_0^B).$$

Hence, (13) becomes

$$r p_0^B = -r\lambda p_0^B \underline{t} - \lambda p_0^B \left[-\frac{r}{\lambda} - r\bar{t} + \frac{r(1 - p_0^B) e^{-\lambda \bar{t}}}{p_0^B} \left(\frac{1}{\lambda} - \frac{1}{\bar{\lambda}} \right) \right],$$

which is equivalent to the following condition to identify \bar{t}

$$\left[p_0^B + (1 - p_0^B) e^{-\lambda \bar{t}} \right] \left(1 - \frac{\lambda}{\bar{\lambda}} \right) = \lambda p_0^B (\bar{t} - \underline{t}). \quad (14)$$

Note that the LHS of (14) decreases in \bar{t} , and the RHS increases in \bar{t} . At \underline{t} , it is easy to see that $[p_0^B + (1 - p_0^B) e^{-\lambda \underline{t}}] \left(1 - \frac{\lambda}{\bar{\lambda}}\right) > 0$, while for $t \rightarrow \infty$,

$$\lim_{t \rightarrow \infty} \lambda p_0^B (t - \underline{t}) > \lim_{t \rightarrow \infty} [p_0^B + (1 - p_0^B) e^{-\lambda t}] \left(1 - \frac{\lambda}{\bar{\lambda}}\right).$$

Therefore, (14) identifies a unique solution for \bar{t} . In particular, such solution satisfies $\bar{t} = \underline{t} + \frac{(\bar{\lambda} - \lambda)}{\lambda \lambda p_{\bar{t}}^B}$. From here, we can get expressions for the equilibrium processes. From (13), we have $r p_0^B = -r \lambda p_0^B \underline{t} - \lambda p_0^B [C + 1]$, implying $C + 1 = \frac{-r(\lambda \underline{t} + 1)}{\lambda}$ and $C = \frac{-r(\lambda \underline{t} + 1) - \lambda}{\lambda}$. Therefore, from (10), we obtain, for $s \in [\underline{t}, \bar{t}]$

$$\lambda_s^A = \frac{\lambda [(1 - p_0^B) e^{-\lambda s} + p_0^B]}{e^{-\lambda s} (1 - p_0^B) - p_0^B \lambda (s - \underline{t}) + p_0^B} \quad (15)$$

$$\text{and } p_s^A = \frac{\lambda p_0^B}{r (1 - p_0^B) e^{-\lambda s} - p_0^B r \lambda (s - \underline{t}) + p_0^B r + \lambda p_0^B}. \quad (16)$$

To guarantee that this is an equilibrium, we check the second order conditions. In particular, we must have $V(s) < \kappa$ for any $s < \underline{t}$ and $s > \bar{t}$. To do so, we will show that $\frac{dV(s)}{ds} > 0$ for $s < \underline{t}$, and that $\frac{dV(s)}{ds} < 0$ for $s > \bar{t}$. By continuity, this implies that $V(s)$ attains its maximum when $\frac{dV(s)}{ds} = 0$, which justifies our first-order conditions approach. For $s < \underline{t}$, we have:

$$V(s) = \frac{v}{r} \left[\int_0^s [(1 - p_t^A) \lambda p_t^B + (1 - p_t^B) \lambda p_t^A] e^{-\int_0^t (r + (1 - p_z^A) \lambda + (1 - p_z^B) \lambda) dz} dt + (e^{-\int_0^s (r + (1 - p_z^A) \lambda + (1 - p_z^B) \lambda) dz}) p_s^A \right],$$

yielding

$$\begin{aligned} \frac{dV(s)}{ds} &= \frac{v}{r} [(1 - p_s^A) \lambda p_s^B + (1 - p_s^B) \lambda p_s^A] e^{-\int_0^s (r + (1 - p_z^A) \lambda + (1 - p_z^B) \lambda) dz} \\ &\quad + \frac{v}{r} \dot{p}_s^A e^{-\int_0^s (r + (1 - p_z^A) \lambda + (1 - p_z^B) \lambda) dz} \\ &\quad - \frac{v}{r} p_s^A [r + (1 - p_s^A) \lambda + (1 - p_s^B) \lambda] e^{-\int_0^s (r + (1 - p_z^A) \lambda + (1 - p_z^B) \lambda) dz}. \end{aligned}$$

After some algebra, it is easy to check that $\frac{dV(s)}{ds} > 0$ if and only if $\frac{p_s^B (1 - p_s^A)}{p_s^A} > \frac{r}{\lambda}$. From (11), we have $\lambda p_{\underline{t}}^B (1 - p_{\underline{t}}^A) - p_{\underline{t}}^A r = 0$, which is equivalent to $\frac{p_{\underline{t}}^B (1 - p_{\underline{t}}^A)}{p_{\underline{t}}^A} = \frac{r}{\lambda}$. Therefore, $\frac{dV(s)}{ds} > 0$ if and only if

$$\frac{p_s^B (1 - p_s^A)}{p_s^A} > \frac{p_{\underline{t}}^B (1 - p_{\underline{t}}^A)}{p_{\underline{t}}^A} \Leftrightarrow \frac{p_s^B}{p_{\underline{t}}^B} > \frac{p_s^A (1 - p_{\underline{t}}^A)}{p_{\underline{t}}^A (1 - p_s^A)}.$$

The last condition is equivalent to

$$\frac{p_0^B + (1 - p_0^B)e^{-t\lambda}}{p_0^B + (1 - p_0^B)e^{-s\lambda}} > \frac{(p_0^A + (1 - p_0^A)e^{-t\lambda}) \frac{(1-p_0^A)e^{-t\lambda}}{p_0^A+(1-p_0^A)e^{-t\lambda}}}{(p_0^A + (1 - p_0^A)e^{-s\lambda}) \frac{(1-p_0^A)e^{-s\lambda}}{p_0^A+(1-p_0^A)e^{-s\lambda}}},$$

which, in turn, is equivalent to $e^{-s\lambda} > e^{-t\lambda}$, which holds for $s < \underline{t}$. For $s > \bar{t}$, $\frac{dV(s)}{ds} < 0$ if and only if $(1 - p_s^A)\bar{\lambda}p_s^B - p_s^A r < 0$. Recall that (10) at \bar{t} implies $\frac{p_{\bar{t}}^B(1-p_{\bar{t}}^A)}{p_{\bar{t}}^A} = \frac{r}{\bar{\lambda}}$. Therefore, $\frac{dV(s)}{ds} < 0$ if and only if

$$\frac{p_s^B(1 - p_s^A)}{p_s^A} < \frac{p_{\bar{t}}^B(1 - p_{\bar{t}}^A)}{p_{\bar{t}}^A} \Leftrightarrow \frac{p_s^B}{p_s^A} < \frac{p_{\bar{t}}^B(1 - p_{\bar{t}}^A)}{p_{\bar{t}}^A(1 - p_s^A)}.$$

The last condition is equivalent to

$$\begin{aligned} \frac{p_{\bar{t}}^B}{p_{\bar{t}}^B (p_{\bar{t}}^B + (1 - p_{\bar{t}}^B)e^{-(s-\bar{t})\bar{\lambda}})} &< \frac{p_{\bar{t}}^A}{p_{\bar{t}}^A (p_{\bar{t}}^A + (1 - p_{\bar{t}}^A)e^{-(s-\bar{t})\bar{\lambda}})} \frac{(1 - p_{\bar{t}}^A)}{\frac{(1-p_{\bar{t}}^A)e^{-(s-\bar{t})\bar{\lambda}}}{p_{\bar{t}}^A+(1-p_{\bar{t}}^A)e^{-(s-\bar{t})\bar{\lambda}}}} \Leftrightarrow \\ e^{-(s-\bar{t})\bar{\lambda}} &< p_{\bar{t}}^B + (1 - p_{\bar{t}}^B)e^{-(s-\bar{t})\bar{\lambda}}, \end{aligned}$$

which is equivalent to $e^{-(s-\bar{t})\bar{\lambda}} < 1$, which holds for $s > \bar{t}$. Finally, for equilibrium uniqueness, note that if $r \leq \bar{r}$, and therefore, $V'(0) \geq 0$, since $r > 0$, as long as $\lambda_s^A = \lambda_s^B = \underline{\lambda}$, $V(s)$ must reach a peak at a finite $\underline{t} \geq 0$. If an atom of agents joins A at \underline{t} , (1) implies that $V'(\underline{t}) > 0$, and therefore any agent in the atom would have an incentive to wait rather than to join A . For the same reason, we cannot have an atom of agents joining A at any $s \in (\underline{t}, \bar{t}]$. Hence, the one described in Part 1 of Proposition 1 is the only equilibrium possible.

(2) Suppose that $r > \bar{r}$, implying $\underline{t} < 0$. Let $\hat{\lambda} > \underline{\lambda}$ solve: $\frac{\hat{\lambda}(1-p_0^A)}{rp_0^A} = \frac{1}{p_0^B}$. Since $r > \bar{r}$ implies $V'(0) < 0$, an atom of agents is willing to join field A at $t = 0$ and takes the arrival rate from $\underline{\lambda}$ to $\hat{\lambda}$. From there, the equilibrium p_t^A , λ_t^A , and \bar{t} can be obtained following a construction similar to Part 1. We present such construction in detail, and we fully characterize the equilibrium processes for $r > \bar{r}$ in the Online Appendix. Since $V(s)$ is constant for all $s \in (0, \bar{t}]$, this is an equilibrium. Moreover, since $\lim_{s \rightarrow 0^+} V(s) = p_0^A \frac{v}{r}$, the agents in the initial atom are indifferent between joining immediately (and obtain $p_0^A \frac{v}{r}$), and waiting to join later in the interval. Finally, for uniqueness, note that if a larger atom $\tilde{\lambda} > \hat{\lambda}$ joins A at $t = 0$, (1) would imply $V'(0) > 0$, generating a profitable deviation for any agent in the atom. Also, any atom of agents joining A at any $s \in (0, \bar{t}]$ creates a similar profitable deviation. Therefore, if there are agents left on the market after an atom of size $\hat{\lambda}$ joins A , they must join gradually after the initial atom. ■

7.2 Proofs for the Good News Case

In (2) we define $V_i(s | \{\lambda_t^A\}, \{\lambda_t^B\})$ for $i = A, B$ —the utility of waiting until time $s > 0$ (unless good news arrives), and then joining field i —conditional on the strategies followed by the other agents, which determine the processes $\{\lambda_t^A\}$ and $\{\lambda_t^B\}$. As mentioned in Section 4, to simplify the notation, we denote by $V_i(s | \lambda^A, \lambda^B)$ the utility of the same strategy conditional on *the information arrival rates in the fields A and B held constant at λ^A and λ^B , respectively, from $t = 0$ to s* . Formally, we have:

$$V_j(s | \lambda^A, \lambda^B) \equiv \frac{v}{r} \left[\int_0^s [p_t^A \lambda^A + p_t^B \lambda^B] e^{-\int_0^t (r + p_z^A \lambda^A + p_z^B \lambda^B) dz} dt + p_s^j e^{-\int_0^s (r + p_z^A \lambda^A + p_z^B \lambda^B) dz} \right] \quad (17)$$

We start with a series of Lemmas which are useful in the subsequent proofs.

Lemma A1 *In the good news case, if $\lambda_t^A = \lambda^A > \underline{\lambda}$ and $\lambda_t^B = \underline{\lambda}$ for any $t \geq 0$, there exists $\widehat{s}(\lambda^A) > 0$ such that $p_{\widehat{s}(\lambda^A)}^A = p_{\widehat{s}(\lambda^A)}^B$, and $p_s^A < p_s^B$ for any $s > \widehat{s}(\lambda^A)$. In particular,*

$$\widehat{s}(\lambda^A) = \frac{1}{(\lambda^A - \underline{\lambda})} \ln \left[\frac{p_0^A (1 - p_0^B)}{p_0^B (1 - p_0^A)} \right].$$

Proof of Lemma A1: The claim follows immediately from the fact that $\widehat{s}(\lambda^A)$ satisfies

$$p_s^A = \frac{p_0^A}{p_0^A + (1 - p_0^A) e^{s \lambda^A}} = \frac{p_0^B}{p_0^B + (1 - p_0^B) e^{s \underline{\lambda}}} = p_s^B. \quad \blacksquare$$

Lemma A2 *For $j = A, B$, for any $s \geq 0$, we have:*

1. $\text{sign} [V_j'(s | \lambda^A, \lambda^B)] = \text{sign} [p_s^{-j} \lambda^{-j} (1 - p_s^j) - r p_s^j]$, and
2. $\text{sign} \left[\left(\frac{p_s^{-j} (1 - p_s^j)}{p_s^j} \right)' \right] = \text{sign} [\lambda^j - \lambda^{-j} (1 - p_s^{-j})]$.

Proof of Lemma A2: For Part 1, for $j = A, B$, (17) implies

$$V_j'(s | \lambda^A, \lambda^B) = \frac{v}{r} \left[[p_s^A \lambda^A + p_s^B \lambda^B] e^{-\int_0^s (r + p_z^A \lambda^A + p_z^B \lambda^B) dz} + \dot{p}_s^j e^{-\int_0^s (r + p_z^A \lambda^A + p_z^B \lambda^B) dz} - p_s^j (r + p_s^A \lambda^A + p_s^B \lambda^B) e^{-\int_0^s (r + p_z^A \lambda^A + p_z^B \lambda^B) dz} \right].$$

Therefore, $\text{sign} [V'_j(s|\lambda^A, \lambda^B)] = \text{sign} [p_s^{-j} \lambda^{-j} (1 - p_s^j) - r p_s^j]$. To verify Part 2, note that

$$\begin{aligned} \text{sign} \left[\left(\frac{p_s^{-j} (1 - p_s^j)}{p_s^j} \right)' \right] &= \text{sign} \left[(1 - p_t^j) \frac{\dot{p}_t^{-j}}{p_t^j} - (1 - p_t^j) \frac{p_t^{-j} \dot{p}_t^j}{(p_t^j)^2} - \frac{p_t^{-j}}{p_t^j} \dot{p}_t^j \right] \\ &= \text{sign} [\lambda^j - \lambda^{-j} (1 - p_t^{-j})]. \end{aligned}$$

■

Lemma A3 When λ^A and λ^B are fixed, and $\lambda^A > \lambda^B$, $\lim_{s \rightarrow \infty} V'_B(s|\lambda^A, \lambda^B) < 0$.

Proof of Lemma A3: By Lemma A2 (1), $\text{sign} [V'_B(s|\lambda^A, \lambda^B)] = \text{sign} \left[\lambda_A (1 - p_s^B) \frac{p_s^A}{p_s^B} - r \right]$.

Observe that

$$\frac{p_s^A}{p_s^B} = \frac{\frac{p_0^A e^{-\lambda^A s}}{1 - p_0^A + p_0^A e^{-\lambda^A s}}}{\frac{p_0^B e^{-\lambda^B s}}{1 - p_0^B + p_0^B e^{-\lambda^A s}}} \rightarrow \frac{1 - p_0^B}{1 - p_0^A} \frac{p_0^A}{p_0^B} e^{(\lambda^B - \lambda^A)s} \rightarrow 0$$

implies that $\lambda_A (1 - p_s^B) \frac{p_s^A}{p_s^B} - r$ is negative for large s . ■

Lemma A4 Suppose that, fixing λ^A and λ^B as constant from $t = 0$, there is an interval of time, starting at $\underline{t} > 0$ and ending at $\bar{t} > \underline{t}$, along which there is a continuous trickle into one field $i = A, B$. Then: (i) $V_i(\underline{t}|\lambda^A, \lambda^B)$ is at its max, i.e., $\frac{\lambda^{-i}(1-p_{\underline{t}}^i)}{p_{\underline{t}}^i} = \frac{r}{p_{\underline{t}}^{-i}}$; (ii) $p_t^i > p_t^{-i}$ for any $t \in [\underline{t}, \bar{t}]$; (iii) $\lambda_t^i = \lambda^{-i}(1 - p_s^{-i})$ for any $t \in [\underline{t}, \bar{t}]$; (iv) $\frac{\lambda^{-i}(1-p_t^i)}{p_t^i} < \frac{r}{p_t^{-i}}$ for some $t \in (\bar{t}, \bar{t} + \varepsilon)$.

Proof of Lemma A4: Assume that there is an interval $[\underline{t}, \bar{t}]$ in which agents gradually join a field, and suppose, without loss of generality, that such field is B . (i) For any $t \in [\underline{t}, \bar{t}]$, it must be the case that

$$\frac{\lambda^A (1 - p_t^B)}{p_t^B} = \frac{r}{p_t^A}. \quad (18)$$

Note that along this interval, λ^A is fixed. So at the initial time, \underline{t} , we need (18) to be satisfied, implying that $V_B(t|\lambda^A, \lambda^B)$ is at its max at \underline{t} . (ii) $p_t^B > p_t^A$ for any $t \in [\underline{t}, \bar{t}]$ is an immediate consequence of Lemma 1. To show (iii), since by Part 2 of Lemma A2, we have that $\text{sign} \left[\left(\frac{(1-p_t^B)}{p_t^B} p_t^A \right)' \right] = \text{sign} [\lambda_t^B - \lambda^A (1 - p_t^A)]$. Hence, to sustain (18), we need $\lambda_t^B = \lambda^A (1 - p_t^A)$ which implies that $\lambda_t^B < \lambda^A$ for all $t \in [\underline{t}, \bar{t}]$. Hence, in the interval $[\underline{t}, \bar{t}]$, we must have $p_t^B > p_t^A$ and $\lambda_t^B < \lambda^A$, including at the end of the interval. After time \bar{t} , as p_t^A is decreasing, we have that

$\lambda_{\bar{t}}^B - \lambda^A(1 - p_{\bar{t}}^A) < 0$ for any $t \in (\bar{t}, \bar{t} + \varepsilon)$ (ε chosen so that there are no other atoms of agents joining the fields within $(\bar{t}, \bar{t} + \varepsilon)$). By Part 2 of Lemma A2, this implies that $(\frac{(1-p_{\bar{t}}^B)}{p_{\bar{t}}^B} p_{\bar{t}}^A)' < 0$ so that, by Part 1 of Lemma A2, $\frac{\lambda^A(1-p_{\bar{t}}^B)}{p_{\bar{t}}^B} < \frac{r}{p_{\bar{t}}^A}$ for $t \in (\bar{t}, \bar{t} + \varepsilon)$. \blacksquare

Lemma A5 For any $\lambda^B \geq \underline{\lambda}$, $s > 0$, both $V_A(s|\lambda^A, \lambda^B)$ and $V_B(s|\lambda^A, \lambda^B)$ strictly increase in λ^A .

Proof of Lemma A5: Consider $V_A(s|\lambda^A, \lambda^B)$ first. For any fixed $\lambda^B \geq \underline{\lambda}$, consider the derivative of $V_A(s|\lambda^A, \lambda^B)$ with respect to s . From Part 1 of Lemma A2, we have

$$\frac{\partial V_A(s|\lambda^A, \lambda^B)}{\partial s} = \lambda^B(1 - p_s^A)p_s^B - rp_s^A.$$

Note that for any s , and any $\lambda^A > \hat{\lambda}^A$, if p_s^A and \hat{p}_s^A are the posteriors at s under λ^A and $\hat{\lambda}^A$, respectively, we have $p_s^A < \hat{p}_s^A$, implying $\frac{\partial V_A(s|\lambda^A, \lambda^B)}{\partial s} > \frac{\partial V_A(s|\hat{\lambda}^A, \lambda^B)}{\partial s}$. Since $V_A(0|\lambda^A, \lambda^B) = V_A(0|\hat{\lambda}^A, \lambda^B)$, the claim follows. Next, consider $V_B(s|\lambda^A, \lambda^B)$. Note that $V_B(s|\{\lambda_t^A\}, \{\lambda_t^B\})\frac{r}{v}$ can be rewritten as

$$\begin{aligned} V_B(s|\{\lambda_t^A\}, \{\lambda_t^B\})\frac{r}{v} &= p_0^A p_0^B \left[\int_0^s (\lambda_t^A + \lambda_t^B) e^{-\int_0^t (r + \lambda_z^A + \lambda_z^B) dz} dt + (p_s^B) e^{-\int_0^s (r + \lambda_z^A + \lambda_z^B) dz} \right] \\ &+ p_0^A (1 - p_0^B) \left[\int_0^s \lambda_t^A e^{-\int_0^t (r + \lambda_z^A) dz} dt + p_s^B e^{-\int_0^s (r + \lambda_z^A + \lambda_z^B) dz} \right] \\ &+ p_0^B (1 - p_0^A) \left[\int_0^s \lambda_t^B e^{-\int_0^t (r + \lambda_z^B) dz} dt + p_s^B e^{-\int_0^s (r + \lambda_z^B) dz} \right] \\ &+ (1 - p_0^B)(1 - p_0^A) p_s^B e^{-rs}. \end{aligned}$$

The last two terms do not depend on λ^A . Setting $\lambda_z^B = \lambda^B$ and $\lambda_z^A = \lambda^A$ and simplifying, the first two terms can be written as

$$p_0^A p_0^B \left[(\lambda^A + \lambda^B) \frac{1 - e^{-(r + \lambda^A + \lambda^B)s}}{(r + \lambda^A + \lambda^B)} + e^{-(r + \lambda^A)s} e^{-\lambda^B s} \right] + p_0^A (1 - p_0^B) \lambda^A \int_0^s e^{-(r + \lambda^A)t} dt$$

The last term is proportional to that of a c.d.f. of an exponential distribution with frequency

λ^A , which increases in λ^A . Finally, consider then the derivative of the expression

$$H(\lambda^A) \equiv (\lambda^A + \lambda^B) \frac{1 - e^{-(r+\lambda^A+\lambda)}}{(r + \lambda^A + \lambda^B)} + e^{-(r+\lambda^A)s} e^{-\lambda^B s}$$

with respect to λ^A , which is

$$H'(\lambda^A) = \frac{r}{(r + \lambda^A + \lambda^B)^2} (1 - e^{-s(r+\lambda^A+\lambda^B)} (1 + s(r + \lambda^A + \lambda^B))) > 0,$$

as for any $\kappa > 0$, $1 - e^{-\kappa}(1 + \kappa) > 0$. ■

Proof of Proposition 2:

Proof of Part 1. First, assume that $r < \frac{\lambda p_0^B (1-p_0^A)}{p_0^A}$. Then, Part 1 of Lemma A2 implies $\text{sign}[V'_A(0)] = \text{sign}\left[\frac{\lambda p_0^B (1-p_0^A)}{r} - p_0^A\right] > 0$. Moreover, by Part 2 of Lemma A2,

$$\text{sign}\left[\left(\frac{(1-p_t^A)}{p_t^A} p_t^B\right)'\right] = \text{sign}[\underline{\lambda} - \underline{\lambda}(1-p_t^B)] > 0,$$

implying $V'_A(s|\underline{\lambda}, \underline{\lambda}) > 0$ for any $s \geq 0$. By Lemma 2, since $p_t^A > p_t^B$ for any $t \geq 0$, we have $V_A(t|\underline{\lambda}, \underline{\lambda}) > V_B(t|\underline{\lambda}, \underline{\lambda})$, and therefore all agents wait for news indefinitely. Next, assume that $r \geq \frac{\lambda p_0^B (1-p_0^A)}{p_0^A}$ but that $V_A(s|\underline{\lambda}, \underline{\lambda}) > \frac{p_0^A v}{r}$ for some $s > 0$. Since $V_A(0) = \frac{p_0^A v}{r}$, this must imply that $V'_A(s'|\underline{\lambda}, \underline{\lambda}) > 0$ at some $s' < s$. As noted above, by Part 2 of Lemma A2, this implies $V'_A(s|\underline{\lambda}, \underline{\lambda}) > 0$ for all $s > s'$. At the limit, we have

$$\lim_{s \rightarrow \infty} V_A(s|\underline{\lambda}, \underline{\lambda}) = \frac{v}{r} \underline{\lambda} \left[\frac{2p_0^A p_0^B}{r + 2\underline{\lambda}} + \frac{p_0^A + p_0^B - 2p_0^A p_0^B}{r + \underline{\lambda}} \right].$$

Therefore, $\lim_{s \rightarrow \infty} V_A(s|\underline{\lambda}, \underline{\lambda}) \geq \frac{v}{r} p_0^A$ is equivalent to

$$p_0^A r^2 + r \underline{\lambda} (2p_0^A - p_0^B) - 2\underline{\lambda}^2 p_0^B (1 - p_0^A) \leq 0. \quad (19)$$

Since the LHS of (19) is increasing in r , consider $r = \bar{r} = \frac{\lambda p_0^B (1-p_0^A)}{p_0^A}$, and note that (19) is satisfied for such r . Then, consider

$$p_0^A r^2 + r \underline{\lambda} (2p_0^A - p_0^B) - 2\underline{\lambda}^2 p_0^B (1 - p_0^A) = 0. \quad (20)$$

Since the LHS of (20) is increasing for $r > 0$, such equation cannot have two positive solutions. Therefore, letting \tilde{r} be the (only) positive solution of (20), we have $\tilde{r} \geq \bar{r} = \frac{\lambda p_0^B (1-p_0^A)}{p_0^A}$, and therefore we can conclude that all agents wait indefinitely for news for any $r < \tilde{r}$.

Proof of Part 2. Assume that $r > \tilde{r}$, where \tilde{r} is defined as in the proof of Part 1. From the discussion above, this implies $r \geq \frac{\lambda p_0^B(1-p_0^A)}{p_0^A}$ and $V_A(s|\underline{\lambda}, \underline{\lambda}) < \frac{p_0^A v}{r}$ for all s . Note that, by Lemma 2, $V_B(s|\underline{\lambda}, \underline{\lambda}) < V_A(s|\underline{\lambda}, \underline{\lambda})$ for all $s \geq 0$. Suppose further that $\bar{\lambda}$ is such that

$$\max \left\{ \sup_{s>0} V_A(s|\bar{\lambda}, \underline{\lambda}), \sup_{s>0} V_B(s|\bar{\lambda}, \underline{\lambda}) \right\} < \frac{p_0^A v}{r},$$

implying that all agents join A immediately. Next, assume that $r \geq \frac{\lambda p_0^B(1-p_0^A)}{p_0^A}$, $V_A(s|\underline{\lambda}, \underline{\lambda}) < \frac{p_0^A v}{r}$ for all $s > 0$, and that

$$\sup_{s>0} V_A(s|\bar{\lambda}, \underline{\lambda}) > \frac{p_0^A v}{r} > \sup_{s>0} V_B(s|\bar{\lambda}, \underline{\lambda}).$$

By Lemma A1, if $\lambda^A = \bar{\lambda}$, there exists some \hat{s} for which the posteriors cross, and hence for all $s \geq \hat{s}$, $V_B(s|\bar{\lambda}, \underline{\lambda}) > V_A(s|\bar{\lambda}, \underline{\lambda})$. Note that $V'_A(0|\bar{\lambda}, \underline{\lambda}) < 0$, and also, by definition, $V_A(0|\bar{\lambda}, \underline{\lambda}) = \frac{p_0^A v}{r}$. Thus, it has to be that from some point s' onwards, $V'_A(s'|\bar{\lambda}, \underline{\lambda}) > 0$, and in particular for s'' for which $V_A(s''|\bar{\lambda}, \underline{\lambda}) = \frac{p_0^A v}{r}$. As a result, it has to be that there exists $\bar{s} > \max\{s'', \hat{s}\}$ for which $V_B(\bar{s}|\bar{\lambda}, \underline{\lambda}) > \frac{p_0^A v}{r}$, in contradiction with our assumption. Hence, the only scenario left to check is when $r \geq \frac{\lambda p_0^B(1-p_0^A)}{p_0^A}$, $V_B(s|\underline{\lambda}, \underline{\lambda}) < V_A(s|\underline{\lambda}, \underline{\lambda}) < \frac{p_0^A v}{r}$ for all $s > 0$, and $\sup_{s>0} V_B(s|\bar{\lambda}, \underline{\lambda}) > \frac{p_0^A v}{r}$. In particular, defining $\bar{\lambda}^*$ to be such that $\sup_{s>0} V_B(s|\bar{\lambda}^*, \underline{\lambda}) = \frac{p_0^A v}{r}$, our assumptions imply $\bar{\lambda} > \bar{\lambda}^*$. In the remaining of the proof, we characterize the unique equilibrium that arises in this case.

(1) Next, we argue that *if any agent ever joins a field in equilibrium, an atom of them has to join A at $s = 0$* . By Lemmas 1 and 2, the first agents to join a field must always join A . Lemma A4 implies that a gradual entry into A is not possible, as it would require $\lambda_t^A = \underline{\lambda}(1 - p_s^B) < \underline{\lambda}$. Then, $V_A(s|\underline{\lambda}, \underline{\lambda})$ has to peak where an atom of agents join A . By Lemma A2, the sign of $V'_A(s|\underline{\lambda}, \underline{\lambda})$ is the same as the sign of $\underline{\lambda}(1 - p_t^A) \frac{p_t^B}{p_t^A} - r$, and $\left(\frac{(1-p_t^A)}{p_t^A} p_t^B \right)' = \underline{\lambda} - \underline{\lambda}(1 - p_t^B) > 0$. This implies that $V_A(s|\underline{\lambda}, \underline{\lambda})$ has its peak either at $s = 0$ or at $s = \infty$, proving our claim that if any agent joins a field ever in equilibrium, an atom of them has to join A at $s = 0$.

(2) Note that *we cannot have an equilibrium in which an atom of agents join A at $s = 0$, and the rest wait indefinitely*. Indeed, for s high enough, by Lemmas A1 and A3, we have $V_B(s|\lambda^A, \underline{\lambda}) > V_A(s|\lambda^A, \underline{\lambda})$, and $V'_B(s|\lambda^A, \underline{\lambda}) < 0$. If some individuals join A and some wait indefinitely, it has to be that $\lim_{s \rightarrow \infty} V_B(s|\lambda^A, \underline{\lambda}) = \frac{p_0^A v}{r}$, but if $\lim_{s \rightarrow \infty} V'_B(s|\lambda^A, \underline{\lambda}) < 0$ then we have a contradiction as it must mean that $V_B(s|\lambda^A, \underline{\lambda}) > \frac{p_0^A v}{r}$ for some s , implying that some agents should have joined B instead.

(3) Next, we show that *there cannot be a second atom of agents joining A following the*

first atom at $s = 0$. Again, this is because, by Lemma A2, the sign of $V'_A(s|\lambda^A, \underline{\lambda})$ is the same as $\underline{\lambda}(1 - p_t^A)\frac{p_t^B}{p_t^A} - r$ and $(\frac{(1-p_t^A)}{p_t^A}p_t^B)' = \lambda^A - \underline{\lambda}(1 - p_t^B) > 0$. Hence, there cannot be a second atom of agents joining A as $V_A(s|\lambda^A, \underline{\lambda})$ does not peak for any finite $s > 0$.

(4) We are now ready to proceed with the only remaining possibility—that is, equilibria in which at $s = 0$ an atom of agents joins A , and then some other agents join B (followed potentially by some other entry into either field).

Let $\lambda^B(\lambda^A) \equiv \lambda(\bar{m} - \lambda^{-1}(\lambda^A))$ —that is, $\lambda^B(\lambda^A)$ is the information arrival rate for field B when λ^A is the arrival rate generated by the atom of agents that joined field A , and all remaining agents have joined B . Also, let $V(\infty|\lambda^A, \lambda^B)$ be the utility from waiting forever under the (constant) arrival rates λ^A and λ^B , and note that such function is increasing in both λ^A and λ^B . We now show that there is an equilibrium with: (i) An initial atom of agents joining A , yielding the arrival rate $\hat{\lambda}^A$. (ii) After the crossing where $p_{\hat{s}(\hat{\lambda}^A)}^A = p_{\hat{s}(\hat{\lambda}^A)}^B$ (see Lemma A1), at a point in time $t^* \in (\hat{s}(\hat{\lambda}^A), \infty)$ where $V_B(t^*|\hat{\lambda}^A, \underline{\lambda}) = \sup_s V_B(s|\hat{\lambda}^A, \underline{\lambda}) = \frac{v}{r}p_0^A$, there is an atom of agents $\hat{\lambda}^B$ that join B , where (a) If $\lambda^B(\hat{\lambda}^A) - \hat{\lambda}^A(1 - p_{t^*}^A) \leq 0$ (implying \bar{m} being low enough) then we set $\hat{\lambda}^B = \lambda^B(\hat{\lambda}^A)$. Hence, all remaining agents join B at t^* . (b) If $\lambda^B(\hat{\lambda}^A) - \hat{\lambda}^A(1 - p_{t^*}^A) > 0$, then set the atom of agents joining B at $\hat{\lambda}^B = \hat{\lambda}^A(1 - p_{t^*}^A)$. Some of the remaining agents join B gradually, so that λ_t^B for $t > t^*$ is such that the posteriors satisfy $\hat{\lambda}_A(1 - p_t^B)\frac{p_t^A}{p_t^B} = r$. For this to be possible, by Lemma A4 we need $\lambda_t^B < \hat{\lambda}^A$, which implies that this process may last forever.

To guarantee that the one described above is an equilibrium, recall that we are under the assumption $\sup_s V_B(s|\underline{\lambda}, \underline{\lambda}) < \frac{v}{r}p_0^A < \sup_s V_B(s|\bar{\lambda}, \underline{\lambda})$. Because of Lemma A3, $\sup_s V_B(s|\bar{\lambda}, \underline{\lambda})$ is achieved either at $s = 0$, or at some finite s in which case $\sup_s V_B(s|\bar{\lambda}, \underline{\lambda}) = \max_s V_B(s|\bar{\lambda}, \underline{\lambda})$.

Now, by Lemma A3, when λ^A and λ^B are fixed, and such that $\lambda^A > \lambda^B$, $\sup_s V_B(s|\lambda^A, \lambda^B) = \max_s V_B(s|\lambda^A, \lambda^B)$, i.e., $\arg \sup_s V_B(s|\lambda^A, \lambda^B)$ includes only finite elements. Therefore, by continuity and Lemma A5, there exists $\hat{\lambda}^A < \bar{\lambda}$ such that $\max_s V_B(s|\hat{\lambda}^A, \underline{\lambda}) = \frac{v}{r}p_0^A$. If t^* is defined to be the lowest element in $\arg \max_s V_B(s|\hat{\lambda}^A, \underline{\lambda})$, we have guaranteed that agents are indifferent between joining A at $t = 0$ and joining B at t^* , which is necessary for this equilibrium to hold. Therefore, $\hat{\lambda}^A < \bar{\lambda}$ is the arrival rate generated by the atom of agents joining field A at $t = 0$ —that is, in equilibrium the atom of agents $m_0^A = \lambda^{-1}(\hat{\lambda}^A) < \bar{m}$ joins field A at $t = 0$. Next note that for the parameters we consider, $V'_A(0|\hat{\lambda}^A, \underline{\lambda}) < 0$ and so it is indeed sustainable for an atom of agents to join A at $t = 0$.

By Lemma 1, for any agent to be willing to join B at t^* , it must be the case that $p_{t^*}^A \leq p_{t^*}^B$. We already know that $V_A(s|\hat{\lambda}^A, \underline{\lambda})$ can only have corner maximizers at either $s = 0$ or $s = \infty$.

Since

$$\sup_s V_B(s|\widehat{\lambda}^A, \underline{\lambda}) = \frac{v}{r}p_0^A \geq \lim_{s \rightarrow \infty} V_B(s|\widehat{\lambda}^A, \underline{\lambda}) = \lim_{s \rightarrow \infty} V_A(s|\widehat{\lambda}^A, \underline{\lambda}),$$

implying that as long as there is no new atom of agents in B , $V_A(s|\widehat{\lambda}^A, \underline{\lambda}) \leq \frac{v}{r}p_0^A$ for all s . Since $V_B(s|\widehat{\lambda}^A, \underline{\lambda}) \geq V_A(s|\widehat{\lambda}^A, \underline{\lambda})$ if and only if $p_s^A \leq p_s^B$, it must be the case that at t^* , where

$$V_B(t^*|\widehat{\lambda}^A, \underline{\lambda}) = \frac{v}{r}p_0^A \geq V_A(t^*|\widehat{\lambda}^A, \underline{\lambda}),$$

we have that $p_{t^*}^A \leq p_{t^*}^B$ —that is t^* occurs after the crossing $\widehat{s}(\widehat{\lambda}^A)$. Also, for this to be an equilibrium, we need $V_B(s|\widehat{\lambda}^A, \widehat{\lambda}^B)$ to be weakly decreasing after t^* . A sufficient condition for this is $\lambda_t^B \leq \widehat{\lambda}^A(1 - p_t^A)$ for any $t > t^*$, as p_t^A is decreasing in t . Such condition is satisfied by our construction.

(5) Finally, we show that there is *no equilibrium with further atoms of agents joining any field*. For agents to join B , we need $V_B(s|\widehat{\lambda}^A, \underline{\lambda})$ to have a peak at t^* such that $\max V_B(t^*|\widehat{\lambda}^A, \underline{\lambda}) = \frac{v}{r}p_0^A$, i.e., $\widehat{\lambda}^A(1 - p_{t^*}^B)\frac{p_{t^*}^A}{p_{t^*}^B} = r$. Then, $V_B(t|\widehat{\lambda}^A, \lambda_t^B)$ is going to weakly decrease for any $t > t^*$, ruling out the possibility of additional atoms of agents joining B . In particular, note that by Lemma A2, once $V_B'(t|\widehat{\lambda}^A, \lambda_t^B) \leq 0$, we cannot have another peak of $V_B(t|\widehat{\lambda}^A, \lambda_t^B)$. Moreover, as $\lambda_t^B < \widehat{\lambda}^A$ for all $t \geq 0$, the posteriors cannot cross again. Hence, $V_B(t|\widehat{\lambda}^A, \lambda_t^B) > V_A(t|\widehat{\lambda}^A, \lambda_t^B)$ for all $t > t^*$, ruling out any further entry into A . \blacksquare

Proof of Proposition 3: From Part 2 of Proposition 3, consider the equilibrium indifference condition

$$V_B(t^*(p_0^B, \widehat{\lambda}^A, \underline{\lambda}), p_0^B, \widehat{\lambda}^A, \underline{\lambda}) = \frac{v}{r}p_0^A \quad (21)$$

Total differentiation of (21), together with the envelope theorem and $\frac{\partial V_B(t^*, p_0^B, \widehat{\lambda}^A, \underline{\lambda})}{\partial t^*} = 0$, yield

$$\frac{\partial V_B(t^*, p_0^B, \widehat{\lambda}^A, \underline{\lambda})}{\partial p_0^B} + \frac{\partial V_B(t^*, p_0^B, \widehat{\lambda}^A, \underline{\lambda})}{\partial \widehat{\lambda}^A} \frac{d\widehat{\lambda}^A}{dp_0^B} = 0. \quad (22)$$

We can write

$$\begin{aligned}
V_B(t^*, p_0^B, \widehat{\lambda}^A, \underline{\lambda}) &= p_0^A p_0^B [(\widehat{\lambda}^A + \underline{\lambda}) \int_0^{t^*} e^{-(r+\widehat{\lambda}^A+\underline{\lambda})t} dt + p_{t^*}^B e^{-(r+\widehat{\lambda}^A+\underline{\lambda})t^*}] \\
&+ p_0^A (1 - p_0^B) \widehat{\lambda}^A \int_0^{t^*} e^{-(r+\widehat{\lambda}^A)t} dt + p_{t^*}^B e^{-(r+\widehat{\lambda}^A+\underline{\lambda})t^*} \\
&+ p_0^B (1 - p_0^A) [\underline{\lambda} \int_0^{t^*} e^{-(r+\underline{\lambda})t} dt + p_{t^*}^B e^{-(r+\underline{\lambda})t^*}] \\
&+ (1 - p_0^B)(1 - p_0^A) p_{t^*}^B e^{-rt^*}.
\end{aligned}$$

First, note that

$$\begin{aligned}
\frac{\partial V_B(t^*)}{\partial p_0^B} &= (1 - e^{-(r+\underline{\lambda}+\widehat{\lambda}^A)t^*}) p_0^A \frac{(\widehat{\lambda}^A + \underline{\lambda})}{(r+\underline{\lambda}+\widehat{\lambda}^A)} - (1 - e^{-(r+\widehat{\lambda}^A)t^*}) p_0^A \frac{\widehat{\lambda}^A}{(r+\widehat{\lambda}^A)} \\
&+ (1 - e^{-(r+\underline{\lambda})t^*}) \underline{\lambda} \frac{(1-p_0^A)}{(r+\underline{\lambda})} + p_0^A e^{-(r+\underline{\lambda}+\widehat{\lambda}^A)t^*} + (1 - p_0^A) e^{-(r+\underline{\lambda})t^*} \\
&> 0.
\end{aligned}$$

Second, Lemma A5 guarantees $\frac{\partial V_B(t^*)}{\partial \widehat{\lambda}^A} > 0$. Using (22), we obtain $\frac{d\widehat{\lambda}^A}{dp_0^B} < 0$. ■

Proof of Proposition 4: For any $\alpha > 1$ and arrival rate processes $\{\lambda_t^A\}, \{\lambda_t^B\}$, let us denote the utility of joining field i at time s unless news arrives first in the presence of a first-mover advantage as

$$\begin{aligned}
V_i^\alpha(s | \{\lambda_t^A\}, \{\lambda_t^B\}) &\equiv \frac{v}{r} \int_0^s [p_t^A \lambda_t^A U_t^A + p_t^B \lambda_t^B U_t^B] e^{-\int_0^t [r+p_z^A \lambda_z^A + p_z^B \lambda_z^B] dz} dt \\
&+ \frac{v}{r} \left(e^{-\int_0^s [r+p_z^A \lambda_z^A + p_z^B \lambda_z^B] dz} \right) U_s^i,
\end{aligned} \tag{23}$$

where

$$U_t^i = \begin{cases} \alpha p_t^j & \text{if } m_\tau^i = 0 \text{ for all } \tau < t; \\ p_t^j & \text{otherwise.} \end{cases}$$

Most of this proof follows the arguments of the proof of Proposition 2, with a few differences. First, note that for any $s \geq 0$, we have

$$(V_A^\alpha(s | \underline{\lambda}, \underline{\lambda}))' = \alpha \underline{\lambda} p_s^B (1 - p_s^A) - \alpha r p_s^A = V_A'(s | \underline{\lambda}, \underline{\lambda}). \tag{24}$$

Then, as long as no one has joined any field yet, the analysis is identical to Proposition 2.

Consider now an agent who, after a mass λ^A of other agents have joined A , is considering

joining B . We have:

$$(V_B^\alpha(s|\lambda^A, \underline{\lambda}))' = \underline{\lambda} p_s^A (1 - \alpha p_s^B) - \alpha r p_s^B < \underline{\lambda} p_s^A (1 - p_s^B) - r p_s^B,$$

where the last expression corresponds to the $\alpha = 1$ case. Hence, an agent waiting to join B becomes relatively *more impatient* with respect to the standard model. Moreover, note that the payoff of an agent joining A at $t = 0$ is

$$V_A^\alpha(0|\underline{\lambda}, \underline{\lambda}) = \alpha p_0^A \frac{v}{r} > p_0^A \frac{v}{r} = V_A(0|\underline{\lambda}, \underline{\lambda}).$$

These last two observations together imply that the (A, B) equilibria hold only if $\bar{\lambda} > \hat{\lambda}^A(\alpha)$, where $\hat{\lambda}^A(\alpha)$ is strictly increasing in α . Thus, as α increases, the (A, B) equilibrium is less likely to exist. Note that after entry in B occurs at some $t^*(\alpha)$, the rest of the equilibrium for any $s > t^*(\alpha)$ is the same as the one in Proposition 2. In particular, all the remaining agents joining a field after $t^*(\alpha)$ face incentives similar to the $\alpha = 1$ case, as there is no field left with a first-mover advantage still available (note however that, since $\hat{\lambda}^A(\alpha) > \hat{\lambda}^A$, the equilibrium entry in B is adjusted accordingly). \blacksquare

Proof of Proposition 5: Consider $\gamma \in (p_0^A, 1)$. Since $p_0^A > p_0^B$, and in the good news case posteriors decrease over time, this guarantees $\gamma > p_t^i$ for any $i = A, B$ and $t \geq 0$. In turn, this implies that *upon arrival of good news on any field, all agents still join that field*.

Given $\gamma \in (p_0^A, 1)$, $\hat{m} \in (0, \bar{m})$, and arrival rate processes $\{\lambda_t^A\}, \{\lambda_t^B\}$, let us denote by $V_i^\gamma(s|\{\lambda_t^A\}, \{\lambda_t^B\})$ the utility of joining field $i = A, B$ at time s unless news arrives first in the presence of congestion externalities, where for any $\tau \geq s$ such that $m_\tau^i > \hat{m}$, the flow payoff is reduced by $\gamma < 1$. Formally, we have

$$\begin{aligned} V_i^\gamma(s|\{\lambda_t^A\}, \{\lambda_t^B\}) &\equiv \frac{1}{r} \left[\int_0^s [p_t^A \lambda_t^A Q_t^A + p_t^B \lambda_t^B Q_t^B] e^{-\int_0^t [r + p_z^A \lambda_z^A + p_z^B \lambda_z^B] dz} dt \right. \\ &\quad \left. + \left(e^{-\int_0^s [r + p_z^A \lambda_z^A + p_z^B \lambda_z^B] dz} \right) Q_s^i \right], \end{aligned} \tag{25}$$

where $Q_t^i = \int_t^\infty e^{-r\tau} U_\tau^i d\tau$, and

$$U_\tau^i = \begin{cases} p_\tau^i v & \text{if } m_\tau^i \leq \hat{m}; \\ \gamma p_\tau^i v & \text{otherwise.} \end{cases}$$

The proof of Proposition 2 can be replicated in this environment, except for the differences highlighted below.

Proof of Part 1. For any $\gamma \in (p_0^A, 1)$, $\hat{m} < \bar{m}$, and $\lambda(\cdot)$, we have

$$(V_A^\gamma(s|\underline{\lambda}, \underline{\lambda}))' = \underline{\lambda} p_s^B (\gamma - p_s^A) - r p_s^A.$$

Since the benefits of waiting are lower than in the $\gamma = 1$ case, the equilibrium in which all wait indefinitely for news is less likely to hold.

Proof of Part 2. Consider $\gamma \in (p_0^A, 1)$, and the equilibrium arrival rate $\hat{\lambda}^A$ derived in the (A, B) equilibrium with $\gamma = 1$. With a sufficiently convex $\lambda(\cdot)$, we have $\lambda^{-1}(\hat{\lambda}^A) > \hat{m}$ and $\bar{m} - \lambda^{-1}(\hat{\lambda}^A) < \hat{m}$. Since the agents who join A at $t = 0$ incur congestion externalities, we have $V_A^\gamma(0|\hat{\lambda}^A, \underline{\lambda}) = \gamma p_0^A \frac{v}{r} < V_A(0|\hat{\lambda}^A, \underline{\lambda})$. Conversely, $V_B^\gamma(s|\hat{\lambda}^A, \underline{\lambda}) > \frac{1}{\gamma} V_B(s|\hat{\lambda}^A, \underline{\lambda})$, as once in B , agents will not face congestion. This implies that there is a wider range of $\bar{\lambda}$ for which we can sustain the (A, B) equilibria.

Proof of Part 3. Again, consider $\gamma \in (p_0^A, 1)$, and the equilibrium arrival rate $\hat{\lambda}^A$ derived in the (A, B) equilibrium with $\gamma = 1$. With a sufficiently concave $\lambda(\cdot)$, we have $\lambda^{-1}(\hat{\lambda}^A) < \hat{m}$ and $\bar{m} - 2\lambda^{-1}(\hat{\lambda}^A) > \hat{m}$. This implies that as soon as news arrives on a field, a mass of agents greater than \hat{m} joins that field, causing congestion to arise.

(a) First, consider the payoff of an agent who joins A at $t = 0$, $V_A^\gamma(0|\hat{\lambda}^A, \underline{\lambda})$. Such a payoff exhibits no congestion externalities as long as there are no news on A , and will be reduced by congestion indefinitely if news on A arrives prior to news on B . This implies

$$\begin{aligned} V_A^\gamma(0|\hat{\lambda}^A, \underline{\lambda}) &= V_A(0|\hat{\lambda}^A, \underline{\lambda}) - (1 - \gamma) \frac{v}{r} \int_0^\infty p_t^A \hat{\lambda}^A e^{-\int_0^t (r + \hat{\lambda}^A p_z^A + p_z^B \lambda_z^B) dz} dt \\ &< p_0^A \frac{v}{r} = V_A(0|\hat{\lambda}^A, \underline{\lambda}) \end{aligned} \quad (26)$$

(b) Next, consider $V_B^\gamma(t^*|\hat{\lambda}^A, \underline{\lambda})$, which is the utility of an agent waiting until t^* to join B . Consider first the expected payoff flow up to t^* . This is reduced by congestion since if news arrives on a field the agent joins that field with all other remaining agents. However, once the agent joins B , from t^* onwards, the flow payoff is affected by congestion indefinitely only in the event in which news on B arrives prior to news on A . Formally, we have:

$$\begin{aligned} V_B^\gamma(t^*|\hat{\lambda}^A, \underline{\lambda}) &= V_B(t^*|\hat{\lambda}^A, \underline{\lambda}) - (1 - \gamma) v \int_0^{t^*} \left(p_t^A \hat{\lambda}^A + p_t^B \underline{\lambda} \right) e^{-\int_0^t (r + p_z^A \hat{\lambda}^A + p_z^B \underline{\lambda}) dz} dt \\ &\quad - (1 - \gamma) v \int_{t^*}^\infty p_t^B \lambda_t^B e^{-\int_{t^*}^t (r + p_z^A \hat{\lambda}^A + p_z^B \lambda_z^B) dz} dt \end{aligned} \quad (27)$$

(c) Recall that in a (A, B) equilibrium, we have $V_A(0|\hat{\lambda}^A, \underline{\lambda}) = V_B(t^*|\hat{\lambda}^A, \underline{\lambda})$ by construction. Therefore, from (26) and (27), the (A, B) equilibrium is less likely to hold if $V_B^\gamma(t^*|\hat{\lambda}^A, \underline{\lambda}) <$

$V_A^\gamma(0|\hat{\lambda}^A, \underline{\lambda})$, and therefore, if

$$\begin{aligned}
\int_0^\infty p_t^A \hat{\lambda}^A e^{-\int_0^t (r + \hat{\lambda}^A p_z^A + p_z^B \lambda_z^B) dz} dt &= \int_0^{t^*} p_t^A \hat{\lambda}^A e^{-\int_0^t (r + \hat{\lambda}^A p_z^A + p_z^B \lambda_z^B) dz} dt \\
&+ \int_{t^*}^\infty p_t^A \hat{\lambda}^A e^{-\int_0^t (r + \hat{\lambda}^A p_z^A + p_z^B \lambda_z^B) dz} dt \\
&< \int_0^{t^*} (\underline{\lambda} p_t^B + p_t^A \hat{\lambda}^A) e^{-\int_0^t (r + \hat{\lambda}^A p_z^A + \underline{\lambda} p_z^B) dz} dt \\
&+ \int_{t^*}^\infty p_t^B \hat{\lambda}^B e^{-\int_0^t (r + \hat{\lambda}^A p_z^A + p_t^B \hat{\lambda}^B) dz} dt
\end{aligned} \tag{28}$$

A sufficient condition for (28) to be satisfied is $p_t^B \hat{\lambda}^B > p_t^A \hat{\lambda}^A$ for all $t > t^*$. Since we know that the trickling into B is governed by $\hat{\lambda}_t^B = \hat{\lambda}^A (1 - p_t^A)$, we need to show that, for all $t \geq t^*$,

$$p_t^B (1 - p_t^A) > p_t^A. \tag{29}$$

To do that and finish the proof, we show the following claim:

Claim A: *If $\underline{\lambda} = 0$, and p_0^A is sufficiently large relative to p_0^B , then $p_t^B (1 - p_t^A) > p_t^A$ for all $t \geq t^*$.*

Proof of Claim A: Note that to show $p_t^A < p_t^B (1 - p_t^A)$ for all $t \geq t^*$, it is sufficient to show that $\frac{p_{t^*}^A}{(1 - p_{t^*}^A)} < p_{t^*}^B = p_0^B$ (since $\underline{\lambda} = 0$ and, from t^* onwards, the LHS decreases faster than the RHS). Let us set p_0^B such that $\frac{p_0^B + (p_0^B)^2}{1 + (p_0^B)^2} < p_0^A$ (note that this is feasible since $\frac{p_0^B + (p_0^B)^2}{1 + (p_0^B)^2} > p_0^B$). From (5) we know that $\frac{p_{t^*}^A (1 - p_0^B)}{p_0^B} = \frac{r}{\hat{\lambda}^A}$. Therefore, we have

$$e^{-rt^*} = e^{\ln \left[\frac{\frac{r}{\hat{\lambda}^A} p_0^B (1 - p_0^A)}{p_0^A \left[(1 - p_0^B) - \frac{r}{\hat{\lambda}^A} p_0^B \right]} \right]^{\frac{r}{\hat{\lambda}^A}}} = \left[\frac{\frac{r}{\hat{\lambda}^A} p_0^B (1 - p_0^A)}{p_0^A \left[(1 - p_0^B) - \frac{r}{\hat{\lambda}^A} p_0^B \right]} \right]^{\frac{r}{\hat{\lambda}^A}} \tag{30}$$

Using the equilibrium indifference condition

$$V_B(t^*|\hat{\lambda}^A, 0) = \frac{\hat{\lambda}^A}{r + \hat{\lambda}^A} (p_0^A + e^{-rt^*} p_0^B (1 - p_0^A)) = p_0^A \frac{v}{r},$$

we obtain

$$\frac{r}{\hat{\lambda}^A} = \frac{p_0^B (1 - p_0^A)}{p_0^A e^{rt^*}}. \tag{31}$$

Plugging (30) into (31), we observe that in the (A, B) equilibrium, $\frac{r}{\hat{\lambda}^A}$ solves

$$\frac{r}{\hat{\lambda}^A} = \frac{p_0^B (1 - p_0^A)}{p_0^A} \left[\frac{r}{\hat{\lambda}^A} \frac{p_0^B (1 - p_0^A)}{p_0^A (1 - (1 + \frac{r}{\hat{\lambda}^A}) p_0^B)} \right]^{\frac{r}{\hat{\lambda}^A}} \tag{32}$$

Since we are focusing on solutions to $\frac{r}{\hat{\lambda}^A}$ in $(0, 1)$, note that when $\frac{r}{\hat{\lambda}^A} \rightarrow 0$, the LHS of (32) is smaller than the RHS. Also, it is easy to check that when $\frac{r}{\hat{\lambda}^A} = \frac{(1-p_0^B)}{(1+p_0^B)}$, the RHS of (32) is smaller than the LHS. To see this, substitute $\frac{r}{\hat{\lambda}^A} = \frac{(1-p_0^B)}{(1+p_0^B)}$ in the RHS, and notice that the RHS reduces to $\frac{p_0^B(1-p_0^A)}{p_0^A} \frac{2}{(1+p_0^B)}$. Moreover, note that $\frac{2}{(1+p_0^B)} > 1$, implying $\frac{p_0^B(1-p_0^A)}{p_0^A} \frac{2}{(1+p_0^B)} < \frac{p_0^B(1-p_0^A)}{p_0^A}$. From our assumption on p_0^B , we have $\frac{p_0^B(1-p_0^A)}{p_0^A} < \frac{(1-p_0^B)}{(1+p_0^B)}$. Then, the RHS of (32) is indeed smaller than the LHS—that is,

$$\frac{p_0^B(1-p_0^A)}{p_0^A} \frac{2}{(1+p_0^B)} < \frac{(1-p_0^B)}{(1+p_0^B)}.$$

Thus, by continuity of the RHS, a solution is such that $\frac{r}{\hat{\lambda}^A} < \frac{(1-p_0^B)}{(1+p_0^B)}$. Recall that we want to show that $p_0^B > \frac{p_{t^*}^A}{(1-p_{t^*}^A)}$. From the optimality of t^* , we have $p_{t^*}^A = \frac{r}{\hat{\lambda}^A} \frac{p_0^B}{(1-p_0^B)}$, implying that $p_{t^*}^A$ is increasing in $\frac{r}{\hat{\lambda}^A}$. Since $\frac{r}{\hat{\lambda}^A} < \frac{(1-p_0^B)}{(1+p_0^B)}$, we have

$$\frac{p_{t^*}^A}{1-p_{t^*}^A} = \frac{\frac{r}{\hat{\lambda}^A} \frac{p_0^B}{(1-p_0^B)}}{1 - \frac{r}{\hat{\lambda}^A} \frac{p_0^B}{(1-p_0^B)}} < \frac{\frac{(1-p_0^B)}{(1+p_0^B)} \frac{p_0^B}{(1-p_0^B)}}{1 - \frac{(1-p_0^B)}{(1+p_0^B)} \frac{p_0^B}{(1-p_0^B)}} = p_0^B,$$

which completes the proof of Claim A. By continuity, Claim A still holds for sufficiently small $\underline{\lambda}$, and moreover the condition in (28) is easier to satisfy. \blacksquare

Proof of Proposition 6:

Proof of Part 1. If $r < \bar{r}$, by Part 2 of Lemma A2, $V_A'(s|\underline{\lambda}, \underline{\lambda}) > 0$ for all $s \geq 0$. Hence, it is optimal for all agents to wait until the deadline T . Next, assume that $r \geq \frac{\underline{\lambda} p_0^B (1-p_0^A)}{p_0^A}$ but that $V_A(s|\underline{\lambda}, \underline{\lambda}) > \frac{p_0^A v}{r}$ for some $s \in (0, T]$. Since $V_A(0) = \frac{p_0^A v}{r}$, this must imply that $V_A'(s'|\underline{\lambda}, \underline{\lambda}) > 0$ at some $s' < s$. As noted above, by Part 2 of Lemma A2, this implies $V_A'(s|\underline{\lambda}, \underline{\lambda}) > 0$ for all $s > s'$, so that the highest value of $V_A(s|\underline{\lambda}, \underline{\lambda})$ is achieved at the deadline T . Therefore, define $\tilde{r}(T)$ to be such that $V_A(T|\underline{\lambda}, \underline{\lambda}) = \frac{v}{r} p_0^A$. From the discussion above, we have $\bar{r} \leq \tilde{r}(T)$ for all T . Hence, for any $r < \tilde{r}(T)$, all agents join A at the deadline T . To see that $\tilde{r}(T)$ is increasing in T , consider $T' < T$, and $\tilde{r}(T')$ such that $V_A(T'|\underline{\lambda}, \underline{\lambda}) = \frac{v}{r} p_0^A$. For such a $\tilde{r}(T')$, we have $V_A(T'|\underline{\lambda}, \underline{\lambda}) < V_A(T|\underline{\lambda}, \underline{\lambda})$. Therefore, to obtain $V_A(T'|\underline{\lambda}, \underline{\lambda}) = \frac{v}{r} p_0^A$, since $V_A(T|\underline{\lambda}, \underline{\lambda})$ decreasing in r , we need $\tilde{r}(T') < \tilde{r}(T)$.

Proof of Part 2. Assume that $r > \tilde{r}(T)$ (implying that for all $s \in (0, T]$, $V_A(s|\underline{\lambda}, \underline{\lambda}) < \frac{p_0^A v}{r}$), and $T \in (0, t^*)$. Recall that $\hat{s}(\lambda)$ is the posteriors' crossing time when $\lambda > \underline{\lambda}$ agents are in A , and $\underline{\lambda}$ are in B from $t = 0$. Also, recall that we denote by $\hat{\lambda}^A$ and t^* the equilibrium size of the atom joining A at $t = 0$, and the time of the atom joining B in Proposition 2,

respectively. Also, define $\mu_A^A(T)$ be such that $V_A(T|\mu_A^A(T), \underline{\lambda}) = \frac{v}{r}p_0^A$ and $\mu_B^A(T)$ be such that $V_B(T|\mu_B^A(T), \underline{\lambda}) = \frac{v}{r}p_0^A$. Let us construct the equilibrium following Steps 1 and 2 below, which provide an algorithm to identify the (unique) equilibrium for each set of parameters.

Step 1: Suppose that $V_A(T|\bar{\lambda}, \underline{\lambda}) > \frac{v}{r}p_0^A$, or equivalently, by Lemma A5, $\bar{\lambda} > \mu_A^A(T)$. Then, two cases are possible:

(a) if $\widehat{s}(\mu_A^A(T)) > T$, then in equilibrium $\mu_A^A(T)$ agents join A at $t = 0$, and the remaining agents join A at T . To see that this equilibrium is unique, note that if more agents join A at $t = 0$, agents would prefer to join A at T instead. If fewer agents join A at $t = 0$, then the remaining agents would prefer to join A at $t = 0$ themselves.

(b) If $\widehat{s}(\mu_A^A(T)) < T$, then we must have $V_B(T|\mu_A^A(T), \underline{\lambda}) > V_A(T|\mu_A^A(T), \underline{\lambda}) = \frac{v}{r}p_0^A$, implying that, by Lemma A5, we need $\mu_B^A(T) < \mu_A^A(T)$ to achieve $V_B(T|\mu_B^A(T), \underline{\lambda}) = \frac{v}{r}p_0^A$. Then in the (unique) equilibrium an atom $\tilde{\lambda}^A = \mu_B^A(T)$ joins A at $t = 0$, and the remaining agents join B at T . Such equilibrium is unique for arguments similar to Step 1(a) (note that in the non-generic case $\widehat{s}(\mu_A^A(T)) = T$, both equilibria described in Steps 1(a) and 1(b) exist).

Step 2: Suppose that $V_A(T|\bar{\lambda}, \underline{\lambda}) \leq \frac{v}{r}p_0^A$, or equivalently, $\bar{\lambda} \leq \mu_A^A(T)$. Then, two cases are possible: (a) If $V_B(T|\bar{\lambda}, \underline{\lambda}) < \frac{v}{r}p_0^A$, then all agents join A at $t = 0$; (b) If $V_B(T|\bar{\lambda}, \underline{\lambda}) > \frac{v}{r}p_0^A$, then $\tilde{\lambda}^A = \mu_B^A(T)$ agents join A at $t = 0$, and the remaining join B at T . Again, note that in the non-generic case $V_B(T|\bar{\lambda}, \underline{\lambda}) = \frac{v}{r}p_0^A$, both equilibria in Steps 2(a) and 2(b) exist.

Finally, note that in Steps 1(b) and 2(b) above, in which field B is explored in equilibrium, since in Proposition 2, for $r > \tilde{r}$, $V_B(s|\tilde{\lambda}^A, \underline{\lambda})$ peaks at t^* , we must have $\tilde{\lambda}^A > \widehat{\lambda}^A$.

Proof of Part 3. If $r \geq \tilde{r}(T)$, and $T \geq t^*$, the (unique) equilibrium follows Proposition 2, except that is at T there are still agents left, they all join B at T . ■

Proof of Proposition 7:

Let us denote $\bar{r}^{opt} = \frac{\bar{\lambda}p_0^B(1-p_0^A)}{p_0^A}$, and let \tilde{r}^{opt} be the unique solution of

$$p_0^A r^2 + r\bar{\lambda} (2p_0^A - p_0^B) - 2\bar{\lambda}^2 p_0^B (1 - p_0^A) = 0.$$

Let us start the proof with a series of useful observations.

Fact 1: Since $\underline{\lambda} = 0$, the planner must allocate at least a ε -measured set of agents to a field at $t = 0$.

Fact 2: With arguments similar to the ones used in the proof of Part 1 of Proposition 2 in relation to (20), the definitions of \bar{r}^{opt} and \tilde{r}^{opt} imply that $V_A(\cdot|\bar{\lambda}, \bar{\lambda})$ is maximized at $t = 0$ for $r > \tilde{r}^{opt}$, and at infinity for $r \leq \tilde{r}^{opt}$. Moreover, $V_A(0|\bar{\lambda}, \bar{\lambda}) > 0$ if and only if $r < \bar{r}^{opt}$.

Fact 3: Consider the utility of waiting to join field $i = A, B$ following a time $t \geq 0$ in which

at least a ε -measured mass of agents have already joined both fields—that is, $V_i(s|\bar{\lambda}, \bar{\lambda})$ for any $s \geq t$. Since the proof of Part 1 of Proposition 2 holds for any $\underline{\lambda}$, p_0^A , and p_0^B , the same holds for $V_A(s|\bar{\lambda}, \bar{\lambda})$, and also, following the same argument, for $V_B(s|\bar{\lambda}, \bar{\lambda})$, from t onwards, using the posteriors p_t^B or p_t^A instead of p_0^B or p_0^A , respectively. Therefore, both $V_A(s|\bar{\lambda}, \bar{\lambda})$ and $V_B(s|\bar{\lambda}, \bar{\lambda})$ are maximized either at infinity or at t .

Fact 4: If $r > \bar{r}^{opt}$, then $\max_{s \geq 0} V_B(s|\bar{\lambda}, 0) < \frac{v}{r} p_0^A$, and $\max_{s \geq 0} V_A(s|0, \bar{\lambda}) = \frac{v}{r} p_0^A$. To see this, denoting by $t_B^{opt} \equiv \arg \sup_{s \geq 0} V_B(s|\bar{\lambda}, 0)$, and using simple algebra to obtain $V_B(t_B^{opt}|\bar{\lambda}, 0)$, observe that

$$V_B(t_B^{opt}|\bar{\lambda}, 0) = \frac{\bar{\lambda}}{r + \bar{\lambda}} \left[p_0^A + e^{-rt_B^{opt}} p_0^B (1 - p_0^A) \right] < p_0^A$$

is equivalent to $e^{-rt_B^{opt}} p_0^B (1 - p_0^A) < p_0^A \frac{r}{\bar{\lambda}}$. This condition is satisfied since

$$e^{-rt_B^{opt}} p_0^B (1 - p_0^A) < p_0^B (1 - p_0^A) < p_0^A \frac{r}{\bar{\lambda}},$$

where the first inequality follows from $e^{-rt_B^{opt}} < 1$ and the last one from $r > \bar{r}^{opt}$. Moreover, $\max_{s \geq 0} V_A(s|0, \bar{\lambda}) = \frac{v}{r} p_0^A$ since, as $r > \bar{r}^{opt}$, we have

$$V'_A(s|0, \lambda) = \bar{\lambda} p_s^B (1 - p_0^A) - r p_0^A < \bar{\lambda} p_0^B (1 - p_0^A) - r p_0^A < 0.$$

Fact 5: Given the value of additional information, for any $s \geq 0$, we have

$$V_A(s|\bar{\lambda}, \bar{\lambda}) \geq \max\{V_A(s|0, \bar{\lambda}), V_B(s|\bar{\lambda}, 0)\}.$$

Because of Facts 1-3, the optimal policies have to take one of the following forms:

- (i) *No Experimentation:* All agents are allocated to field A at $t = 0$;
- (ii) *Single Experimentation:* A measure ε of agents are allocated to field $i \in (A, B)$ at $t = 0$ and all others are allocated to $-i$ at some optimal t_{-i}^{opt} ;
- (iii) *Simultaneous Double Experimentation:* A measure ε of agents are allocated to both A and B at $t = 0$ and all others wait indefinitely for news;
- (iv) *Sequential Double Experimentation:* A measure ε of agents are allocated to $i \in (A, B)$ at $t = 0$, a measure ε of agents are allocated to $-i$ at some optimal time t_{-i}^{opt} , and all others wait indefinitely for news.

We are now ready to show the optimality of the policies described in Parts 1-3 of the proposition.

Proof of Part 1. When $r > \tilde{r}^{opt}$, $V_A(s|\bar{\lambda}, \bar{\lambda})$ is maximized at $t = 0$ with the value $p_0^A \frac{v}{r}$. Hence, no experimentation is optimal. Next, consider $r \in [\bar{r}^{opt}, \tilde{r}^{opt}]$. Because of Fact 4 above, $V_A(s|0, \bar{\lambda})$ and $V_B(s|\bar{\lambda}, 0)$ are bounded above by $p_0^A \frac{v}{r}$ for all s . Therefore, a single experimentation policy cannot be optimal. Clearly, when ε is large enough, double experimentation cannot be optimal either, and therefore all agents must be allocated to field A at $t = 0$.

Proof of Part 2. The utilitarian welfare from policy (iii), simultaneous double experimentation, is

$$\varepsilon p_0^A \frac{v}{r} + \varepsilon p_0^B \frac{v}{r} + (1 - 2\varepsilon) \sup_{s \geq 0} V_A(s|\bar{\lambda}, \bar{\lambda}).$$

If $r < \tilde{r}^{opt}$, then $V_A(s|\bar{\lambda}, \bar{\lambda})$ is maximized at $s = \infty$, and for ε small enough, because of Fact 5, this has to be the optimal policy.

Proof of Part 3. When $r < \bar{r}^{opt}$, if ε is large enough, policies (iii) and (iv) cannot be optimal. In particular, experimenting in B is too costly as such agents receive $\varepsilon p_0^B \frac{v}{r}$ rather than $\varepsilon p_0^A \frac{v}{r}$. Thus, the uniquely optimal policy is single experimentation (i.e., policy (ii)) in field A followed by all remaining agents joining B at $t_B^{opt} = \arg \sup_{s \geq 0} V_B(s|\bar{\lambda}, 0) \neq t^*$. ■

8 References

Adda, Jérôme, and Marco Ottaviani, 2023, “Grantmaking, Grading on a Curve, and the Paradox of Relative Evaluation in Nonmarkets,” *Quarterly Journal of Economics*, 139(2), 1255–1319.

Akerlof, George and Pascal Michaillat, 2017, “Beetles: Biased Promotions and Persistence of False Belief,” NBER Working Paper 23523.

Bardhi, Arjada, Yingni Guo, and Bruno Strulovici, 2024, “Early-Career Discrimination: Spiraling or Self-Correcting?,” mimeo.

Bobtcheff, Catherine, Jérôme Bolte, and Thomas Mariotti, 2017, “Researcher’s Dilemma,” *Review of Economic Studies*, 84, 969-1014.

Bobtcheff, Catherine, Raphaël Lévy, and Thomas Mariotti, 2024, “Information Disclosure in Preemption Races: Blessing or (Winner’s) Curse?,” mimeo.

Bolton, Patrick and Chris Harris, 1999, “Strategic Experimentation,” *Econometrica*, 67(2), 349-374.

Bonatti, Alessandro and Johannes Horner, 2017, “Career Concerns with Exponential Learning,” *Theoretical Economics*, 12(1), 425-475.

- Bramoullé, Yann and Gilles Saint-Paul, 2010, “Research Cycles,” *Journal of Economic Theory*, 145, 1890–1920.
- Brock William A. and Steven N. Durlauf, 1999, “A Formal Model of Theory Choice in Science,” *Economic Theory*, 14(1) 113-130.
- Callander, Steven, Nicolas Lambert, and Niko Matouschek, 2024, “Innovation and Competition on a Rugged Technological Landscape,” mimeo.
- Carnehl Christoph and Johannes Schneider, 2024, “A Quest for Knowledge,” mimeo.
- Cetemen, Doruk, Can Urgan, and Leeat Yariv, 2022, “Collective Progress: Dynamics of Exit Waves,” *Journal of Political Economy*, 131(9), 2402-2450.
- Chen Chia-Hui, Hülya Eraslan, Junichiro Ishida, and Takuro Yamashita, 2024, “Optimal Feedback Dynamics Against Free-Riding in Collective Experimentation,” mimeo.
- Dossi Gaia, 2024,, “Race and Science,” Ph.D. Thesis, London School of Economics.
- Dossi, Gaia, and Marta Morando, 2023, “Political Ideology and Innovation,” mimeo.
- Einio, Elias, Josh Feng, and Xavier Jaravel, 2023, “Social Push and the Direction of Innovation.” mimeo.
- Fonseca, Ricardo, 2024, “Staggered Rollouts for Innovation Adoption,” mimeo.
- Forand, Jean Guillaume, 2015, “Keeping Your Option Open,” *Journal of Economic Dynamics and Control*, 53, 47-68.
- Frick, Mira, and Yuhta Ishii, 2024, “Innovation Adoption by Forward-Looking Social Learners,” *Theoretical Economics*, 19, 1505-1541.
- Gittins, John C., 1979, “Bandit Processes and Dynamic Allocation Indices,” *Journal of the Royal Statistical Society*, 41(2), 148-177.
- Gul, Faruk, and Russell Lundholm, 1995, “Endogenous Timing and the Clustering of Agents’ Decisions,” *Journal of Political Economy*, 103(5), 1039-1066.
- Halac, Marina, Navin Kartik, and Qingmin Liu, 2017, “Contests for Experimentation,” *Journal of Political Economy*, 125(5), 1523-1569.
- Hill, Ryan, and Carolyn Stein, 2024, “Race to the Bottom: Competition and Quality in Science,” mimeo.
- Hill, Ryan, and Carolyn Stein, 2023, “Scooped! Estimating Rewards for Priority in Science,” forthcoming *Journal of Political Economy*.
- Hopenhayn, Hugo and Francesco Squintani, 2021, “On the Direction of Innovation,” *Journal of Political Economy*, 129(7), 1991-2022.
- Keller, Godfrey and Sven Rady, 2010, “Strategic Experimentation with Poisson Bandits,” *Theor-*

etical Economics, 5(2), 275-311.

Keller, Godfrey and Sven Rady, 2015, “Breakdowns,” *Theoretical Economics*, 10(1), 175-202.

Keller, Godfrey, Sven Rady, and Martin Cripps, 2005, “Strategic Experimentation with Exponential Bandits,” *Econometrica*, 73(1), 39-68.

Kitcher, Philip, 1990, “The Division of Cognitive Labor”, *The Journal of Philosophy*, 87(1), 5–22.

Knoepfle, Jan, and Julia Salmi, 2024, “Dynamic Evidence Disclosure: Delay the Good to Accelerate the Bad,” mimeo.

Koning, Rembrand, Sampsa Samila, and John-Paul Ferguson, 2021, “Who Do We Invent For? Patents by Women Focus More on Women’s Health, but Few Women Get to Invent,” *Science*, 372(6548), 1345–1348.

Kuhn, Thomas S., 1962, *The Structure of Scientific Revolutions*, University of Chicago Press.

Laiho, Tuomas, Pauli Murto, and Julia Salmi, 2023a, “Gradual Learning from Incremental Actions,” mimeo.

Laiho, Tuomas, Pauli Murto, and Julia Salmi, 2023b, “Slow Social Learning: Innovation Adoption under Network Externalities,” mimeo.

Lizzeri, Alessandro, Eran Shmaya, and Leeat Yariv, 2024, “Disentangling Exploration from Exploitation,” mimeo.

Martin, Ian, and Robert Pindyck, 2015, “Averting Catastrophes: The Strange Economics of Scylla and Charybdis,” *American Economic Review*, 105(10), 2947-2985.

Murto, Pauli and Juuso Välimäki, 2011, “Learning and Information Aggregation in an Exit Game,” *The Review of Economic Studies*, 78(4), 1426-1461.

Popper, Karl Raimund, 1999, *All Life is Problem Solving*, Routledge.

Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille, 2007, “Social Learning in One-Arm Bandit Problems,” *Econometrica*, 75(6), 1591-1611.

Siniscalchi, Marciano and Pietro Veronesi, 2021, “Self-image Bias and Lost Talent,” NBER Working Paper 28308.

Sunstein, Cass R., 2001, “On Academic Fads and Fashions,” *Michigan Law Review*, 99(6), 1251-1264.

Thompson, William, 1933, “On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples,” *Biometrika*, 25(3/4), 285-294.

Weber, Max, 1919, “Science as a Vocation,” published as “Wissenschaft als Beruf,” *Gesammelte Aufsätze zur Wissenschaftslehre* (Tübingen, 1922), 524-55.