Research Waves

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

December 15, 2024

Abstract

Competing research lines start and grow as information is gradually uncovered by scientists who choose their fields driven by career incentives. We build a strategic experimentation framework in which agents irreversibly specialize in one of two risky fields, and information updates arrive more frequently as more agents specialize in a field. We describe the equilibrium forces that determine the size, shape, and length of such "research waves." In the 'bad news' case, all researchers specialize in one field, generating a unique bandwagon wave. As the difference in priors increases, such wave starts earlier and grows more slowly. In the 'good news' case, both fields can be explored in equilibrium in two sequential surges. The probability of both fields being investigated increases in the researchers' mass and in the efficacy of technology. Finally, we assess the impact of citations' benefits, tenure clocks, and grants on the structure of the equilibrium waves.

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

Driven by career incentives, researchers sort themselves into different fields of investigation. Along their career paths, motives such as scientific grants, tenure and promotion decisions, salary increases, as well as reputation and legacy, can all potentially influence how scientists choose their research topics.[1] When they are evaluated for promotion, the opinion of more established colleagues, or the number of citations that their work has generated constitute a significant part of this assessment. As a result, researchers have an incentive to specialize in fields that are perceived to be relevant by the profession at large.

The fact that the perception of others affects scientists' investigation choices complicates our understanding of how science evolves. This is because such opinions depend on the intensity with which alternative scientific approaches have been explored before. As more and more research output is produced and published in one field, the community obtains more information on how valuable that particular methodology is. In contrast, less information is available about the viability of other fields that have been explored by relatively fewer researchers. As a result, the endogeneity of information uncovered in different fields and the relative speeds at which such information grows play a crucial role in understanding young scientists' incentives to shift their attention towards areas that are new and relatively under-explored.

In this paper, we study the equilibrium forces that determine the size, shape, and length of such endogenous "research waves." We construct a continuous-time Poisson bandit model (similar to the classic setup of Keller, Rady, and Cripps, 2005) in which: (i) agents decide at which point in time to join one of *two risky* research fields, A or B, and (ii) once the choice of joining a field is made by an agent, it is *irreversible.* The first assumption allows us to analyze how researchers sort themselves into alternative lines of investigation. The second assumption relates to the cost of specializing in a particular scientific field. By taking such a cost to infinity, the analysis becomes significantly more tractable.

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. As time passes, research output conveys further information about the fields' viability. As more researchers generate more output, we assume

[1] The organization of science has been often discussed by scholars over the decades. In 1918, in his "Science as Vocation" lecture, Max Weber discusses the choices of research fields made by young scholars. Kuhn (1962) refers to the phase in which scholars produce research within a mainstream paradigm as "normal science," until overwhelming evidence arises to challenge this perception. More recently, Sunstein (2001) highlights how reputational concerns affect researchers' choices, leading to scientific fads.

that each field's information arrival rate at any given time increases in the total mass of agents that joined that field before.

We capture the 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 in the future. This feature of the model captures the idea that scientists' future rewards increase in the reputation of the field they are in. Initially, we abstract away from any other consideration that may affect researchers' incentives. For example, scientists might benefit from working in more populated fields to align themselves with the interests of established scholars (see Akerlof and Michaillat, 2017, and Siniscalchi and Veronesi, 2021), or alternatively, they could suffer from working in a crowded field where their contributions are less likely to be critical. Also, scientists could be rewarded for being 'first movers' into a new field as they might benefit from the citations that their field generates in the future. Later in the paper, we explore how some of these considerations affect our results.

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, information arrives in the form of negative updates about the viability of a field. In this case, in the absence of news, *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.[3]

In the "good news" case, the information arrival conveys positive updates about the viability of a field. Therefore, 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.[4] While hundreds of scientists worked on this

[2]As common in this literature, to make the analysis simpler, we focus on *conclusive* information arrival: in the bad news case, since a viable field cannot generate bad news, the first arrival determines the non-viability of the field, and similarly for the good news case.

[3]The 'bad news' case is also reminiscent of Popper (1999)'s falsification hypothesis, stating that a scientific theory can only be proven false but never truly confirmed.

[4]For more detail, see a recent news feature on Nature 9/14/2021 (available at https://www.nature.com/articles/d41586-021-02483-w).

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.[5] 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."*[6]

We show how the nature of the information environment affects the emergence and growth of research fields. In the bad news case, there is a unique equilibrium in which, as long as no bad news arrives on field A early enough, all researchers specialize in that field, implying that field B is not explored. The intuition is simple. As long as no agent joins either field, field A is perceived more likely to be viable. As a result, the first time agents join a field, they surely select field A. As more agents join A, more information is generated in that field and, unless bad news arrives on A, the beliefs about its viability will increase faster than the beliefs about B's viability. Thus, the more researchers join A, the more others want to join the "bandwagon." We characterize the dynamic structure of the equilibrium information arrival rate in field A as *one slow research wave.* If agents are patient enough, the research wave into A starts slowly, and accelerates over time. In fact, we show that as the population size grows arbitrarily large, the length of the equilibrium wave is bounded above. Finally, we show that as the difference between the fields' prior beliefs increases, the wave starts earlier, grows more slowly, and lasts longer.

In the good news case, there is also a unique equilibrium which, for some parameters, *involves the exploration of both A and B in two sequential surges.* As long as no agent has joined any field, field A is still perceived more likely to be viable than B. Therefore, field A must still be the one attracting the first agents joining a field. Specifically, we show that the equilibrium involves an immediate surge–i.e., an atom of researchers– into field A. Once field A starts being populated, the posterior beliefs about its viability decrease faster than the

[5]Indeed, research on new vaccines or drugs may be best characterized by an initial "good news" phase, when the project goes from a theoretical idea into actual development, and a subsequent "bad news" phase, when the new drug is tested for the presence of significant side effects on human subjects.

[6]See https://www.theguardian.com/commentisfree/2024/oct/12/nobel-winner-geoffrey-hinton-is-the-godfather-of-ai-heres-an-offer-he-shouldnt-refuse

posterior beliefs about B, until a reversal of the posterior beliefs' ranking occurs. Then, some time after such reversal occurs, a surge into field B arises in equilibrium, 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 most promising. Second, field B's equilibrium upper bound entails that some agents may wait for news indefinitely. Also, our results suggest that the likelihood of the exploration of field B increases in the agents' total pool size and in the efficacy of the information-production technology.

We discuss how scientists' career incentives can influence the structure of the equilibrium research waves. First, we consider an extension of the model in which researchers enjoy a first-mover advantage from joining a field early–for example, because future citations may improve their reputation and facilitate promotions. We also consider the tenure-clock system, which is prevalent in U.S. academic institutions, in which researchers have a limited horizon to produce research output that demonstrates their competence. This setup can be thought as a deadline within which a cohort of scientists is forced to choose a risky field, rather than wait for more information working in a 'safer' investigation area. Finally, we consider direct incentives, such as grants or other subsidies, provided by the wider community (i.e., research institutions, or the government) to promote research in specific areas.

In general, in the bad news case, we find that the bandwagon structure of the equilibrium and the implied lack of research diversification are very robust, and persist across models. Specifically, citation benefits and deadlines cause the bandwagon wave into A to accelerate, while scientific grants can cause the bandwagon wave to switch to B, but never yield research diversification.

In the good news case, first, we find that citations' benefits always make the equilibrium diversification of agents' specializations into both fields that can arise in the standard model less likely to occur. Second, we find that the introduction of a deadline has more nuanced implications on the equilibrium structure. Specifically, since a deadline forces agents to take action rather than to delay, its presence may allow equilibrium diversification to arise in environments in which all agents would have waited for news indefinitely otherwise (i.e., when they are relatively patient). When agents are less patient, we show that introducing a deadline is always detrimental to achieve research diversification in equilibrium. Finally, we find that well-designed scientific grants also have subtle effects on the equilibrium structure, and they can both expand the environments in which diversification arise, and influence the fields' relative equilibrium sizes.

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. Moreover, in our setting, 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. The presence of two risky arms, which is natural for our application, has a substantial impact on the equilibrium analysis (particularly in the good news case) as well on the comparative statics and normative implications of our results. The assumption of irreversibility simplifies the analysis in the sense that, given the equilibrium behavior of others, each agent's best response becomes an optimal stopping problem.

Other 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), where a unique decision-maker is allowed to separate exploration from exploitation. Forand (2015) also analyzes two risky projects in a good-news environment, and also focuses on a single decision-maker rather than on equilibrium analysis. In the context of social experimentation with one risky project, Laiho, Murto, and Salmi (2023a and 2023b) also assume choices' irreversibility.

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).[7] Cetemen, Urgun, and Yariv (2022) analyze exit waves in collective search environments. Second, as discussed above, starting from Kuhn (1962), a large literature on the philosophy of science has addressed how knowledge is created. Kitcher (1990) discusses the individual vs. optimal division of scientific labor across competing themes. Brock and Durlauf (1999) high-

[7] Chen 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.

light how conformity may play a role in scientists' specialization choices, which is an idea explored more recently, as mentioned above, by Akerlof and Michaillat (2017), and Siniscalchi and Veronesi (2021).[8] 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). These papers explore questions different from ours.

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). Clearly, despite a similar trade-off in taking action vs. waiting for more information, the information structure and the focus of this work is different from ours.

2 The Model

Consider a continuous-time, infinite horizon setting, with $t \in [0, \infty)$ and a mass $m > 0$ of identical agents present from the onset. There are two independent, risky fields, A and B, each with flow value $v > 0$ if viable, and zero if not viable. We denote the prior beliefs of 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 the current beliefs about her field's viability affect a researcher's promotion chances, salary, and so on. 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 $m_t^A + m_t^B \leq m$, and the irreversibility assumption implies that m_t^i weakly increases in t. The 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

[8] Recent 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.

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 $\lambda(m) \equiv \overline{\lambda} > \underline{\lambda}$, so that $\underline{\lambda}$ and $\overline{\lambda}$ 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 $m_t^i \in [0, m]$ and $\lambda(m_t^i) \in [\underline{\lambda}, \overline{\lambda}]$, we interchangeably refer to $\lambda_t^i \geq \overline{\lambda}$ 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, $\lambda^{-1}(\lambda_t^i)$.

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, it is easy to see that in our setting 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.

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 $p_t^i \frac{v}{r}$.*

[9] One could consider a more general setting in which, for any fixed mass m_t^i of agents in field i, 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, with $\lambda^B(m_t^i) > \lambda^G(m_t^i)$ for the bad news case, and vice-versa for the good news case. Since as long as no news arrives, posteriors still increase in the bad news case and decrease in the good news case, the analysis would qualitatively follow our setup, while the specific results would depend on the assumptions made on the functions $\lambda^B(\cdot)$ and $\lambda^G(\cdot)$.

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). In particular, Lemma 1 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 $p_t^A > p_t^B$.*

Lemma 2 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), the posterior beliefs processes $\{p_t^A\}$ and $\{p_t^B\}$ cannot cross, implying that their ranking remains the same as the initial one. Lemma 1 and Lemma 2 together imply that, since $p_0^A > p_0^B$, the first field to be explored in equilibrium is always field A.

3 The Bad News Case

3.1 Equilibrium Analysis in the Bad News Case

In the bad news case, once bad news arrives on a field $i = A, B$ at time t, all remaining agents join the other, and, by Lemma 1, obtain an expected continuation payoff of $p_t^{-i}\frac{v}{r}$. Given that the optimal actions upon arrival of bad news on either field are straightforward, we only need to characterize the agents' equilibrium behavior when no news on either field has arrived yet.

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

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

To see how this utility is constructed, note that 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 $\frac{v}{r}\max_{i=A,B} p_s^i$.

Proposition 1 characterizes the unique equilibrium arising in the bad news environment. In particular, let $\overline{r} \equiv \frac{\lambda p_0^B(1-p_0^A)}{p_0^A}$. We have

Proposition 1 (Bad News - Bandwagon Wave) *In the bad news case, unless bad news about A arrives early enough, no agent ever joins field B. Specifically, we have:*

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

In the "*bandwagon wave*" equilibrium described in Proposition 1, unless bad news about A arrives before $\overline{t}$, field B is not explored at all. To see why, note that Lemmas 1 and 2 imply that as long as no agents have joined a field yet, if an agent joins a field, she always joins A. Hence, the arrival rates are such that $\lambda_s^A > \lambda_s^B = \underline{\lambda}$, implying that field A's posterior increases faster than field B's one. Then, joining A becomes increasingly more attractive relative to B, creating the "bandwagon" equilibrium feature–that is, early researchers' exploration of field A "encourages" others to join the same field later.

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 arrive*, which would cause a swap of the chosen field from A to B. This incentive is reflected in the equilibrium conditions characterizing the agents' optimal behavior, as we show below. This observation also implies that in equilibrium there is force limiting the presence of large atoms of agents joining field A at the same time. To see why, observe that an atom of agents joining A causes a discrete increase in the information arrival rate in field A, and, if such increase is large enough, it may induce an agent in the atom to deviate and wait for the information benefit generated by the others instead. This suggests that *the equilibrium must involve some gradual entry into* A.

Specifically, agents must be indifferent between joining A over the interval $[\underline{t}, \bar{t}]$, implying $V(s)$ being equal to a constant–that is, $V(s) = \kappa$ for any $s \in [\underline{t}, \bar{t}]$. Furthermore, we must have $\lambda_t = \underline{\lambda}$ and $V(t) \leq \kappa$ for any $t \leq \underline{t}$, and $\lambda_t = \overline{\lambda}$ and $V(t) \leq \kappa$ for any $t \geq \bar{t}$. To guarantee a constant $V(s)$ in the interval $[\underline{t}, \bar{t}]$, we set $V'(s) = 0$, yielding the ODE for any $s \in [\underline{t}, \bar{t}]$

$$\frac{\dot{p}_s^A}{(p_s^A)^2} = \frac{r}{p_s^B}. \tag{2}$$

To see how the ODE reflects the agents' incentives described above, using $\dot{p}_s^A = \lambda_s^A p_s^A(1-p_s^A)$, (2) can be written as

$$\lambda_s^A(1 - p_s^A)\frac{p_s^B v}{r} = p_s^A v. \tag{3}$$

Consider an agent evaluating an infinitesimal increase in the wait s, Δs. The LHS of (3) represents the benefit of such 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 $p_s^B \frac{v}{r}$ rather than zero. However, while waiting Δs, the agent is foregoing the income $p_s^A v\Delta s$. Hence, the RHS of (3) represents the cost of such an additional wait. Using (2) at the lower cutoff, we obtain the time at which the bandwagon wave starts:

$$\underline{t} = \frac{1}{\underline{\lambda}} \ln\left[\frac{\left(1 - p_0^A\right)\underline{\lambda}}{rp_0^A} - \frac{1 - p_0^B}{p_0^B}\right] \geq 0. \tag{4}$$

From (4), it is apparent that $\underline{t} \geq 0$ if and only if $r \leq \bar{r}$. If this is the case, we obtain the equilibrium processes λ_s^A and p_s^A for $s \in [\underline{t}, \bar{t}]$, and we use λ_s^A to obtain the ending time of the wave by finding the $\bar{t}$ at which all agents have joined field A–that is, $\lambda_{\bar{t}} = \overline{\lambda}$.[10] If $r > \bar{r}$, agents are more impatient and they start joining field A at $t = 0$. Therefore, by setting $\widehat{\lambda} = \min\left\{\frac{rp_0^A}{(1-p_0^A)p_0^B}, \overline{\lambda}\right\}$, we obtain an equilibrium in which an atom $\widehat{\lambda}$ of agents join A at $t = 0$, and, if there are agents left, they join A gradually according to a process similar to Part 1. Note that this scenario always occurs when $\underline{\lambda} = 0$, and therefore in equilibrium $p_t^B = p_0^B$ for all t.

We use the equilibrium characterization of λ_s^A to describe the shape of the equilibrium research wave in the following corollary (whose proof is presented in the Supplemental Ap-

[10] See (15), (16), and (14) in the Appendix for the equilibrium characterizations of λ_s^A, p_s^A, and $\bar{t}$.

pendix).[11]

Corollary 1 (Bad News-Wave Properties) *We have:*

1. *The equilibrium λ_s^A is convex over the interval $\left(\underline{t}, \overline{t}\right]$;*[12]
2. *There is a finite $\widetilde{t} > \underline{t}$ such that $\lim_{\overline{\lambda}\to\infty} \overline{t} = \widetilde{t}$.*

To see the intuition of the convexity in Part 1 of Corollary 1, consider a fixed delay in joining A early vs. later in the wave, and for simplicity, consider $\underline{\lambda} \to 0$, so that $p_t^B \simeq p_0^B$ for all t. From (3), since λ_s^A is increasing in s, the change in $\frac{1-p_s^A}{p_s^A}$, and therefore the change in the cost of delaying joining A relative to B is higher later on. Therefore, to maintain indifference, researchers need to be compensated with increased information–that is, a larger change in λ_s^A–later in the wave, resulting in a convex equilibrium pattern of λ_s^A. Note that when p_t^B increases over time (i.e., $\underline{\lambda} > 0$), the option of waiting becomes increasingly more desirable, and the needed change in λ_s^A relatively smaller. Corollary 1 guarantees the resulting pattern of λ_s^A still to be convex.

To understand Part 2 of Corollary 1, note that neither (4), nor the equilibrium processes λ_s^A, or p_s^A (see (15) and (16) in the Appendix) depend on $\overline{\lambda}$. Hence, as $\overline{\lambda}$ grows, the starting time $\underline{t}$ and the structure of the equilibrium wave remain exactly the same, but its length increases–that is, $\overline{t}$ increases. Nevertheless, Part 2 of Corollary 1 guarantees that as $\overline{\lambda}$ grows large, the length of the wave is bounded above by $\widetilde{t} - \underline{t}$, yielding a arbitrarily fast wave. This is because as p_s^A grows closer to 1, the relative impact of the same change in λ_s^A is lower, and there is a strong incentive to join A. Therefore, to maintain indifference between joining A and waiting, agents need to be compensated with a large amount of additional information, generated by arbitrarily large increases in λ_s^A.

To conclude, it is instructive to compare the results of this section to the previous bandit literature that has analyzed the case of one risky arm with reversible investments in the bad news case. Since in our equilibrium agents have an incentive to wait before joining A only because they can potentially receive bad news about field A (and in that case, switch entry to field B), there is no substantial strategic interaction between the two fields. Therefore, the equilibrium behavior in our setup is *qualitatively* similar to the literature with one risky arm.

[11] Note that convexity pattern and the wave properties described in Corollary 1 and Proposition 2 are expressed in terms of the equilibrium process λ_s^A. The equilibrium mass of researchers joining A over time can be derived as $\left\{\lambda^{-1}(\lambda_s^A)\right\}$, whose pattern depends on the specific assumptions made on the function $\lambda(\cdot)$.

[12] From Part 2 of Proposition 1, if $r \geq \overline{r}$, we have $\underline{t} = 0$.

The only implication of the second field being risky is the fact that the expected payoff obtained from switching to B increases over time. As mentioned above, this affects the length and speed of the wave, as a more attractive prospect for B generates a greater incentive to wait for bad news about A, and therefore an increasingly slower (and longer) bandwagon wave than in a scenario in which the second field's payoff is fixed at a 'safe' level $p_0^B v$. The nature of the strategic interaction between the two risky fields is substantially different, and affects the result in a qualitative way, in the good news case, as we explore in Section 4.

3.2 Comparative Statics in the Bad News Case

The next result illustrates how the bandwagon research wave characterized in Proposition 1 changes in response to changes in the prior beliefs about the fields' viability. Here we focus on the case of agents that are relatively patient–that is, $r \leq \bar{r}$, while in the Supplemental Appendix we present the $r > \bar{r}$ case.[13]

Proposition 2 (Bad News - Comparative Statics for low r) *When $r \leq \bar{r}$, as p_0^A increases or p_0^B decreases locally, we have:*

1. *$\underline{t}$ decreases (wave starts earlier);*
2. *$p_{\underline{t}}^A$ decreases (wave starts at a lower posterior cutoff);*
3. *For any $k > 0$, $\lambda_{\underline{t}+k}^A$ and $p_{\underline{t}+k}^A$ decrease (slower wave);*
4. *$\bar{t} - \underline{t}$ increases (longer wave).*

Proposition 2 states that if agents are relatively patient, as the fields' prior beliefs are increasingly further apart, the bandwagon wave described in Part 1 of Proposition 1 starts earlier and at a lower posterior's cutoff, grows more slowly, and lasts longer. To understand the intuition behind this result, consider $\widetilde{p}_0^A > p_0^A$, while still maintaining $r \leq \bar{r}$ (hence, we are considering a local change around p_0^A). Also, recall that, from Part 1 of Proposition 1 and (3), at any point along the wave, the equilibrium processes have to satisfy

$$\frac{\lambda_s^A(1 - p_s^A)}{p_s^A} = \frac{r}{p_s^B}. \tag{5}$$

First, consider (5) at $s = \underline{t}$, when the arrival rate is $\lambda_s^A = \underline{\lambda}$. As seen previously, at $t = 0$, the LHS of (5) is higher than the RHS, and, as the posteriors increase, they are both moving

[13] The proof of of Proposition 2 is also presented in the Supplemental Appendix.

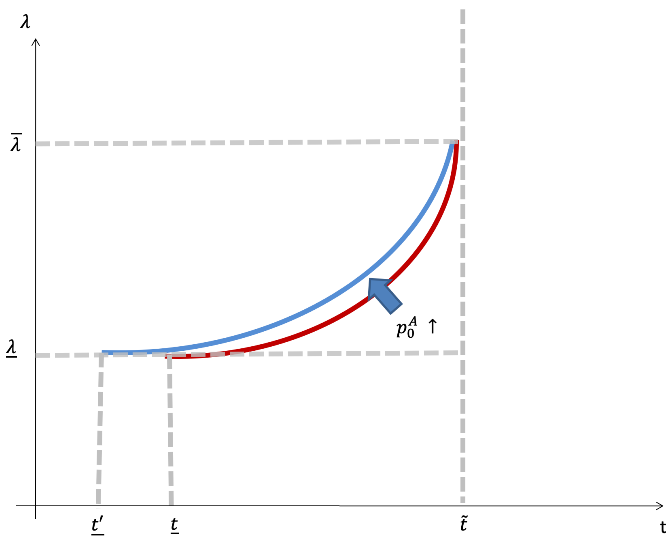

Figure 1: Effect of an Increase in p_0^A on the Equilibrium $\left\{\lambda_s^A\right\}$ in the Bad News Case

downward. A higher prior $\widetilde{p}_0^A$ decreases the LHS, resulting in a shorter time to reach the point $\widetilde{\underline{t}}$ at which this condition is satisfied. Therefore, since $\frac{r}{p_s^B}$ is decreasing at a exogenous rate, the intersection corresponds to a higher $\frac{\lambda(1-p_s^A)}{p_s^A}$, yielding a lower cutoff posterior $p_{\widetilde{\underline{t}}}^A < p_{\underline{t}}^A$.

Consider now a fixed amount of time elapsed in the wave, $k > 0$, and consider an agent considering a small delay in joining A. Since $\widetilde{\underline{t}} < \underline{t}$, and the process $\left\{p_s^B\right\}$ is fixed, $p_{\widetilde{\underline{t}}+k}^B < p_{\underline{t}+k}^B$, resulting in a higher RHS in (5)–that is, a higher incentive to join A immediately rather than waiting, because waiting for bad news on A and joining B became relatively less appealing. Hence, in order to maintain indifference, the relative appeal of joining A has to decrease as well, resulting in $p_{\widetilde{\underline{t}}+k}^A < p_{\underline{t}+k}^A$. Moreover, since the rate of increase of the RHS is lower than the rate of increase of $\frac{1-p_{\underline{t}+k}^A}{p_{\underline{t}+k}^A}$, to maintain equality, $\lambda_{\underline{t}+k}^A$ has to go down. Similar arguments apply to a decrease in p_0^B.[14] Figure 1 illustrates the effect of an increase of p_0^A on the equilibrium process λ_s^A.

As mentioned above, in the Supplemental Appendix, we fully characterize the equilibrium wave in the $r > \overline{r}$ case and we present the comparative statics analysis associated with this case.

[14] Note also that since $\overline{r} \equiv \frac{\lambda p_0^B(1-p_0^A)}{p_0^A}$, a substantial increase in p_0^A or decrease in p_0^B may induce the equilibrium to switch from the scenario described in Part 1 to the one in Part 2 of Proposition 1–i.e, the earlier start of the wave caused by a significant increase of p_0^A or decrease of p_0^B may cause the wave to start with an atom of agents joining A at $t = 0$.

4 The Good News Case

4.1 Equilibrium Analysis in the Good News Case

The analysis of the good news case departs from the bad news case in two main aspects. First, recall that, by Lemmas 1 and 2, the first field to be joined by any agent must be A. In the good news case, as agents join field A, its posterior starts decreasing at a higher rate than field B's posterior, so that their ranking can potentially reverse. This suggests that in the good news case field B may be explored even in the absence of news arriving on either field.

Second, when an agent is planning to join a field $i = A, B$, the incentives to wait stem from the potential arrival on good news on the other field, $-i$. In fact, the arrival of such news would induce the agent 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, as the information arrival process of field $-i$ influences the incentives to join field i.

We start by noting that, once good news arrives on a field, all the remaining agents join that field and obtain the continuation value $\frac{v}{r}$. Therefore, as before, we only need to characterize the equilibrium strategies of the agents while they wait for any news to arrive.

Given the strategies followed by others, which generate the processes $\{\lambda_s^A\}$ and $\{\lambda_s^B\}$, we consider an agent's strategy of waiting until s (unless news arrives), and then, by Lemma 1, join the field with the highest current posterior p_s^i. The value of such a strategy evaluated at $t = 0$ is

$$\begin{aligned} V(s) \equiv \; & \frac{v}{r}\int_0^s \left(p_t^A \lambda_t^A + p_t^B \lambda_t^B\right) e^{-\int_0^t (r + p_z^A \lambda_z^A + p_z^B \lambda_z^B) dz} dt \\ & + \frac{v}{r}\left(e^{-\int_0^s (r + p_z^A \lambda_z^A + p_z^B \lambda_z^B) dz}\right) \max_{i=A,B} p_s^i. \end{aligned} \tag{6}$$

Similarly to the bad news case, the first term of (6) captures the benefits arising if good news about one of the fields arrive before s. Specifically, good news arrives on field $i = A, B$ at time $t < s$ with probability $p_t^i \lambda_t^i$. Then, the agent joins that field and obtains $\frac{v}{r}$. The second term captures the scenario in which no news has arrived by time s, and the agent joins the field with the highest current posterior.

We are now ready to construct the equilibrium in the good news case. Suppose that no agent joins any field initially, so that both posteriors start decreasing under the information arrival rate $\underline{\lambda}$. Since $p_0^A > p_0^B$, by Lemmas 1 and 2, if an agent decides to join a field in this initial period, she always selects field A. In particular, suppose that, similarly to the bad news case, as long as no news arrives, agents gradually join field A along some interval $[\underline{t}, \bar{t}]$. Following techniques similar to the bad news case, it is easy to show that the indifference condition for any $s \in [\underline{t}, \bar{t}]$ would amount to

$$\underline{\lambda} p_s^B (1 - p_s^A) \frac{v}{r} = v p_s^A. \tag{7}$$

The LHS of (7) captures the benefits of a small delay, Δs, in joining A. As discussed above, the benefit of waiting in this phase is associated with the possibility of receiving *good news on field B*, which occurs with probability $\underline{\lambda} p_s^B \Delta s$. In this case, 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 represents the flow utility forgone while waiting before joining A, $v p_s^A \Delta s$. As $\underline{\lambda}$ is fixed and the evolution of p_s^B depends on it, for agents to gradually join A, we need the process λ_s^A to be constructed so that the induced p_s^A satisfies (7). As it turns out, (7) implies that p_s^A would need to decrease more slowly than p_s^B, which is not possible since $\lambda_s^A > \underline{\lambda}$. *Thus, if agents join field A at all, they must join it all together as an atom of some size* $\lambda^A > \underline{\lambda}$.

Suppose that an atom of agents of size $\lambda^A > \underline{\lambda}$ join A at $t = 0$. For any fixed posteriors' level, field A's posterior decreases faster than the posterior of B, causing the two posteriors to cross at some finite $\widehat{s}(\lambda^A) > 0$, and, from then on, their ranking to reverse.[15] Via Lemma 1, 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. Proposition 3 characterizes the (unique) equilibrium in the good news case.[16]

Proposition 3 (Good News - Equilibrium) *There exists $\widetilde{r} > \bar{r}$ such that*

1. *If $r < \widetilde{r}$, all agents wait indefinitely for news;*
2. *If $r > \widetilde{r}$, 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) An atom of agents $\widehat{\lambda}^B < \widehat{\lambda}^A$ joins 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.*

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

[16] For uniqueness, we ignore the non-generic case $r = \widetilde{r}$, for which both equilibria described in Parts 1 and 2 of Proposition 3 exist.

To describe the construction of the equilibrium in Proposition 3, for any $\lambda^A, \lambda^B \geq \underline{\lambda}$, let us define as $V_j(s|\lambda^A, \lambda^B)$ the utility of waiting until time $s > 0$ (unless good news arrives), and then joining field $j = A, B$–that is, $V(s)$ defined in (6)– *when the information arrival rates in the fields A and B are held constant* at λ^A and λ^B, respectively, from $t = 0$ to s.[17] Then, we let $t^*(\lambda^A)$ be the lowest t that satisfies

$$V_B(t|\lambda^A, \underline{\lambda}) = \sup_{s>0} V_B(s|\lambda^A, \underline{\lambda}).$$

In words, fixing the information arrival rates λ^A and $\underline{\lambda}$ in fields A and B, respectively, $t^*(\lambda^A)$ is the (earliest) time at which the value of joining field B peaks.

To understand Part 1 of Proposition 3, assume that no other agent joins either field unless good news arrives. Lemma 2 implies that if an agent joins a field, the field of choice must be A. Note also that if at any point in time s an agent is not willing to join A–that is, $V_A'(s|\underline{\lambda}, \underline{\lambda}) > 0$– 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, if $V'(0) > 0$, or, recalling that $\overline{r} \equiv \frac{\underline{\lambda} p_0^B(1-p_0^A)}{p_0^A}$, $r \leq \overline{r}$, everybody has an incentive to wait indefinitely for news. Moreover, if $V'(0) = \underline{\lambda} p_0^B (1 - p_0^A)\frac{v}{r} - v p_0^A < 0$, or $r > \overline{r}$, then $V(s|\underline{\lambda}, \underline{\lambda})$ is maximized either at $s = 0$, or at $s = \infty$. Since the rate $\widetilde{r}$ makes an agent indifferent between these two choices, for any $r < \widetilde{r}$ in equilibrium, all agents wait indefinitely for news.

In Part 2 of Proposition 3, since $r > \widetilde{r} > \overline{r}$, we have $V_A'(0|\underline{\lambda}, \underline{\lambda}) < 0$–that is, agents are willing to join A as an atom at $s = 0$ rather than wait any small amount of time. As noted above, for any atom of agents $\widehat{\lambda}^A$ that join A at $s = 0$, there is a time in the future $\widehat{s}(\widehat{\lambda}^A) > 0$ at which the posteriors cross, and their ranking reverses.

Next, note that *for any agent to join B at any time $t > 0$, such agent must be indifferent between doing that and join A at $t = 0$*, together with the agents in the atom $\widehat{\lambda}^A$. Hence, the size of the atom $\widehat{\lambda}^A$ must be chosen to guarantee such indifference. Fixing $s > 0$, note that $V_B(s|\lambda^A, \underline{\lambda})$ increases in λ^A. This is intuitive, since the more agents join A at $s = 0$, the more information becomes available for agents waiting to join B in the future, yielding, keeping everything else constant, an increase in the value of the strategy of planning to join B at any $s > 0$. Then, we define $\widehat{\lambda}^A$ to be such that $V_B(t^*(\widehat{\lambda}^A)|\widehat{\lambda}^A, \underline{\lambda}) = \frac{p_0^A v}{r}$. In words, we set $\widehat{\lambda}^A$ so that the peak of $V_B(s|\widehat{\lambda}^A, \underline{\lambda})$ yields exactly the same utility as joining field A at $t = 0$. If $\bar{\lambda} < \widehat{\lambda}^A$, all agents join A at $s = 0$, and the game ends. However, if $\bar{\lambda} > \widehat{\lambda}^A$, we let an atom of size

[17] For a formal definition of $V_j(s|\lambda^A, \lambda^B)$, see (17) in the Appendix.

$\widehat{\lambda}^A$ of agents join A at $t = 0$. Hence, by construction, agents are indifferent between joining A at $s = 0$ and joining B at some time $t^*(\widehat{\lambda}^A)$ in the future, when $V_B(s|\widehat{\lambda}^A, \underline{\lambda})$ reaches its peak. Note that this process simultaneously determines $\widehat{\lambda}^A$ and the time of joining B, $t^* \equiv t^*(\widehat{\lambda}^A)$. To guarantee agents to be willing to join B at t^* rather than A, we also show that t^* occurs after the posteriors' crossing $\widehat{s}(\widehat{\lambda}^A)$.

Next, we quantify the atom of agents $\widehat{\lambda}^B$ joining B at t^*, and we address the possibility that more posteriors' ranking reversals could potentially occur in equilibrium. For agents to be willing to join B at t^* it must be the case that $V_B'(t^*|\widehat{\lambda}^A, \widehat{\lambda}^B) \leq 0$, or, equivalently, $\widehat{\lambda}^A p_{t^*}^A(1 - p_{t^*}^B)\frac{v}{r} - vp_{t^*}^B \leq 0$. This requires $p_{t^*}^B$ not to decrease (under the rate $\widehat{\lambda}^B$) too fast with respect to $p_{t^*}^A$ (under the rate $\widehat{\lambda}^A$) after t^*. We show (see Lemma A4 in the Appendix) that this condition amounts to $\widehat{\lambda}^B \leq \widehat{\lambda}^A(1 - p_{t^*}^A) < \widehat{\lambda}^A$. It could be the case that the agents left in the market after $\widehat{\lambda}^A$ joined A are below such an upper bound for $\widehat{\lambda}^B$. If that is the case, all remaining agents join B at t^* and the game ends. However, if after an atom $\widehat{\lambda}^B = \widehat{\lambda}^A(1 - p_{t^*}^A)$ of agents join B at t^*, there are still agents left on the market, they must gradually join field B after t^*. To guarantee indifference in joining B over an interval of time $[t^*, \overline{t}]$, we need

$$V'(s) = \widehat{\lambda}^A p_s^A(1 - p_s^B)\frac{v}{r} - vp_s^B = 0$$

to be satisfied for any $s \in [t^*, \overline{t}]$. As seen before (again, see Lemma A4 in the Appendix), this requires the trickling into B to be governed by the condition $\lambda_s^B = \widehat{\lambda}^A(1 - p_s^A) < \widehat{\lambda}^A$ for all $s \in [t^*, \overline{t})$. As $\overline{t}$ grows arbitrarily large, if the pool of agents is not exhausted before, p_s^A converges to zero, and therefore λ_s^B converges to $\widehat{\lambda}^A$. Since the gradual entry into B is bounded above by $\widehat{\lambda}^A$, if the total mass of agents in the market is large (i.e., $m > 2\lambda^{-1}(\widehat{\lambda}^A)$), the remaining agents that have joined neither A nor B wait indefinitely for news on either field. The trickling into field B extended to infinity guarantees indifference to be satisfied at the limit.[18,19] Moreover, as a last step in this construction, note that $\lambda_s^B < \widehat{\lambda}^A$ for any $s \geq 0$ guarantees the posteriors beliefs induced by $\widehat{\lambda}^A$ and $\{\lambda_s^B\}$ do not cross again, ruling out further atoms of agents joining any field after the ones we described. In turn, *this implies that no more posteriors' ranking reversals are possible in equilibrium.*

Finally, note that a necessary condition for field B to be explored in equilibrium is $r \geq \overline{r}$. Recall from Proposition 1 that, in the bad news case, this set of parameters is associated with

[18] The proof in the Appendix shows that the equilibrium is unique, and the cases accounted for in Proposition 3 exhaust all parameters.

[19] As r grows arbitrarily large, it is easy to see that $\widehat{\lambda}^A$ becomes arbitrarily large as well. Then, Proposition 3 implies that all agents join field A at $t = 0$.

an atom of agents joining A at $t = 0$, followed by the remaining agents to trickle gradually into A. In the good news case, it is precisely the fact that some agents are willing to join A immediately, that causes that field's ultimate "downfall," resulting in other agents switching away from A and exploring B instead. Put simply, since the agents that join A increase the rate of arrival of good news in that field, absent the arrival of such news, at some point field B is bound to become more attractive than A. Expecting this, some agents do not join A immediately and wait instead, planning to join B at some point in the future if no news arrives before. The value of the information accrued until that point compensates them for the loss of the flow payoff they could have earned in the meantime. However, in the good news case, if no agent is willing to join field A initially, nobody ever will. Then field B never catches up, and all agents wait indefinitely for good news on either field.

4.2 Comparative Statics in the Good News Case

4.2.1 Population Size and Information-Production Technology

From the equilibrium construction in Proposition 3, it is apparent that the population size m (or equivalently, the information-production technology $\lambda(\cdot)$) ultimately determines the extent to which field B is explored. If m is small (i.e., $m \leq \lambda^{-1}(\widehat{\lambda}^A)$), all agents join field A at $t = 0$, and field B remains unexplored. For intermediate levels of m (i.e. $\lambda^{-1}(\widehat{\lambda}^A) < m \leq 2\lambda^{-1}(\widehat{\lambda}^A)$) all agents that have not joined A initially join B either at t^*, or gradually after t^*. If m is large (i.e., $m > 2\lambda^{-1}(\widehat{\lambda}^A)$), field B asymptotically converges to the same size of A, and the remaining agents wait indefinitely for news.

Also, the shape of the information-production technology $\lambda(\cdot)$ affect the relative sizes of the agents' groups that join field A field B, or wait indefinitely for news in equilibrium, denoted by m^A, m^B, and $m - m^A - m^B$, respectively.[20] Specifically, given m and $\overline{\lambda}$, consider a convex $\lambda(\cdot)$, as in the left panel of Figure 2. 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 research. In such a scenario, many agents need to join A before reaching $\widehat{\lambda}^A$, leaving scarce room for field B to grow. Alternatively, consider a concave $\lambda(\cdot)$ (as in the right panel of Figure 2), in a field in which there are decreasing return of scale to information production. Then, a relatively smaller mass of agents is required to reach $\widehat{\lambda}^A$, implying that field B could eventually grow to the same size of A. However, there may be a

[20]Formally, $m^j = \lim_{t\to\infty} m_t^j$ for $j = A, B$.

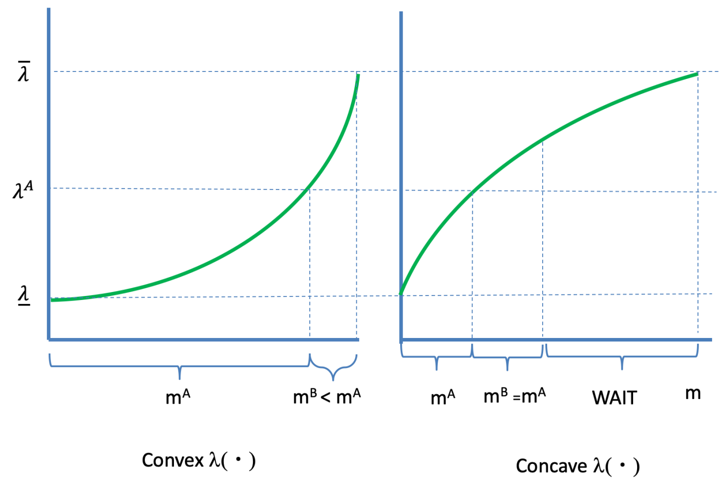

Figure 2: Impact of Information-Production Technology on Equilibrium with Good News

significant amount of agents waiting indefinitely for news.

4.2.2 Prior Beliefs on Fields' Viability

The next result (whose proof is in the Supplemental Appendix) studies the impact of a change in priors on the structure of the equilibrium characterized in Proposition 3.

Proposition 4 (Good News - Change in Prior Beliefs) *We have*

1. *If $r < \widetilde{r}$, a local change in p_0^A or p_0^B has no effect on the equilibrium structure;*
2. *If $r \geq \widetilde{r}$, as p_0^B increases locally, $\widehat{\lambda}^A$ decreases.*

Part 1 of Proposition 4 follows immediately from the fact that, from Part 1 of Proposition 3, if $r < \widetilde{r}$ all agents wait indefinitely for news. Any change in the priors p_0^A and p_0^B which preserves $r < \widetilde{r}$ does not affect the equilibrium structure. Consider now $r \geq \widetilde{r}$, for which the equilibrium follows the structure in Part 2 of Proposition 3. Note that an increase of p_0^B does not affect $V_A(0)$, but, everything else held constant, shifts $V_B(s|\widehat{\lambda}^A, \underline{\lambda})$ up for all $s > 0$. This implies that, in order to guarantee the agents being indifferent between joining field A at $t = 0$ and waiting to join B later, the peak of $V_B(s|\widehat{\lambda}^A, \underline{\lambda})$ has to decrease. This is achieved by decreasing the information the agents receive waiting for news–that is, by decreasing $\widehat{\lambda}^A$. Note that a decrease of $\widehat{\lambda}^A$ constitutes a *double-edged sword* with respect to the exploration of field

B: On the one hand, it makes it more likely to occur, as a smaller $\widehat{\lambda}^A$ is less likely to exhaust all agents on the market. On the other hand, since in equilibrium, in the absence of news, the size of field B is bounded above by the size of field A (i.e. $\lambda_s^B < \widehat{\lambda}^A$ for all $s \geq 0$), a decrease of $\widehat{\lambda}^A$ reduces the maximal size that field B can reach without good news arriving on it.

5 Research Incentives

In this section we consider several policies that are common in the scientific community and we evaluate their impact on our results.

5.1 First-Mover Advantage

Researchers who enter early into a new field of investigation tend to obtain more recognition. This can occur, for example, by using citations' counts as a measure to assess the impact of a researcher's work and to determine their promotions or salary increases. Consider the following simple model of citations, or more in general, first-mover advantage for joining an unexplored field first. Suppose that if a researcher joins a field $i = A, B$ at time $t \geq 0$, and no positive mass of agents joined i before, her flow utility from then onward is $p_\tau^i \gamma v$ for $\tau \geq t$, where $\gamma > 1$. 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 $\gamma = 1$ case). For simplicity, we assume that if multiple researchers join field i first at the same time they each receive the flow utility $p_\tau^i \gamma v$–that is, there are no congestion externalities.[21] We focus the analysis of this section on a large enough γ.

Citations with Bad News. At first glance, one could conjecture that after a positive mass of agents joined A, since joining A does not entail the benefit γ any longer, while B still does, it becomes more attractive for others to join B. Then, according to this intuition, a large enough first-mover advantage–that is, a large γ– may break the bandwagon wave illustrated in Proposition 1, and introduce research diversification in the bad news case. As we show in the next result (whose proof is in the Supplemental Appendix), it turns out that this is not the case.

Proposition 5 (Bad News - Citations) *In the presence of a first-mover advantage, if γ is large enough, all equilibria are as follows: for any $s \in [0, \underline{t}]$ there is an equilibrium in*

[21] At the end of this section, we discuss how the congestion externalities affect the analysis.

which all agents join A at s. Therefore, unless bad news about A arrives early enough, no agent ever joins field B.

Proposition 5 illustrates the multiplicity of equilibria generated by the coordination element present in this environment. Nonetheless, the result guarantees that all equilibria are characterized by one surge into A, and it rules out both slow waves into A and, unless bad news on A arrives early enough, any exploration of B.[22] To see why, note that agents must still join field A first. Then, consider an agent who, once a positive mass of agents have joined A, plans to join B at some point in the future. For such an agent, the cost of increasing the wait by Δs is the foregone flow payoff $v\gamma p_s^B \Delta s$. However, with probability $(1-p_s^B)\underline{\lambda}\Delta s$, the potential arrival of bad news about B induces the agent to switch her plan and join A instead, yielding, since joining A does not entail the benefit γ any longer, a payoff of $p_s^A \frac{\nu}{r}$ rather than zero, and therefore an expected marginal benefit of the additional wait of $p_s^A \frac{\nu}{r}(1-p_s^B)\underline{\lambda}\Delta s$. Hence, for a large enough γ, the marginal cost of waiting is larger then the marginal benefit, implying that, if a positive mass of agents joins A, any entry into B must occur at the same time as the entry into A. However, at that time, joining field A along with the others dominates joining B, making entry into B not feasible in equilibrium. Hence, in equilibrium, all agents must join A together. As seen in Section 3, since the utility obtained from joining A when nobody else has joined peaks at $\underline{t}$, there cannot be an equilibrium in which all agents join A at any time after $\underline{t}$ (as any agent would have a profitable deviation in joining A slightly before the others). However, since the utility from joining A increases over $[0, \underline{t}]$, for any $s \in [0, \underline{t}]$ there is an equilibrium in which all agents join A together at s–constituting a coordination failure for any $s < \underline{t}$.

Citations with Good News. Considerations similar to the ones illustrated in the bad news case prevent entry into B in the good news case also. Therefore, the potential exploration of both fields illustrated in Proposition 3 ceases to exist when γ is large enough. Similarly to the bad news case, at any s such that $V_A(s|\underline{\lambda}, \underline{\lambda})$ is increasing, we can sustain an equilibrium in which all agents join A at s. In addition, if r is low enough, there is an equilibrium in which all agents wait for news indefinitely.[23]

Proposition 6 (Good News - Citations) *In the presence of a first-mover advantage, if γ*

[22] The timeline for the arrival of bad news on A to allow exploration of B is *tighter* than in Proposition 1. In the $\gamma = 1$ case, as long as bad news on A arrives before the end of the wave, $\bar{t}$, some exploration of field B occurs in equilibrium. Here, bad news arriving before $\underline{t} < \bar{t}$ is a necessary condition for any exploration of B to occur.

[23] The proof of Proposition 6 is in the Supplemental Appendix.

is large enough, agents either all join A at the same time, or all wait for news indefinitely.

To summarize, Propositions 5 and 6 imply that in the presence of a large first-mover advantage (e.g., citations' benefits), the bandwagon wave into field A persists and accelerates in the bad news case, and the exploration of both fields is prevented in the good news case.

Finally, note that we have deliberately modeled the first-mover advantage in a very stark manner, as (i) only the very first movers receive an advantage (i.e., there is no advantage in joining a field immediately after a positive mass of first movers), and (ii) the size of the first-mover advantage does not depend on how many agents obtain it (i.e., there are no congestion externalities). Note that the dynamic insights uncovered in Propositions 5 and 6 are robust to the introduction of some congestion externalities: As discussed, the incentives to wait for news before joining B are diminished because A is already populated, and congestion externalities present in field A would tend to lower such incentives even further.[24]In Section 6.1, we briefly discuss a more nuanced way to introduce citations in our model, as a direction for further research.

5.2 Deadlines

Under a tenure-clock system, researchers are on a track leading to a "up or out" promotion decision. Hence, they face a urgency to select their specialization early enough to be able to build a record in their chosen field before their "clock" expires. To explore the implications of such a policy, we now 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.

In the bad news case, the effect of introducing the deadline T is mechanical: in the absence of news arriving in any field, any agents who did not join A before T, joins A at T. The equilibrium characterization preceding the deadline is completely unchanged with respect to Proposition 1. Since this result is straightforward, we present it in the Supplemental Appendix.

We now describe how in the good news case, introducing a deadline has a substantial impact on the equilibrium structure before its expiration. Recall that $\widehat{\lambda}^A$ and t^* are the size of the atom of agents joining field A at $t = 0$, and the time at which an atom of agents join field B, respectively, in the equilibrium described in Part 2 of Proposition 3. Then, we have:

Proposition 7 (Good News - Deadlines) *In the presence of a deadline $0 < T < \infty$, there exist $\widetilde{r}(T) > \overline{r}$ strictly increasing in T, such that the unique equilibrium is as follows:*

[24]However, the presence of congestion externalities may facilitate an equilibrium where all agents join the two fields at the same time, with the relative sizes of the atoms joining each field equalizing their respective utilities.

1. *If* $r < \widetilde{r}(T)$*, all agents join field* A *at* T*.*
2. *If* $r \geq \widetilde{r}(T)$*, 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* $\widetilde{\lambda}^A < \overline{\lambda}$ *of agents join field* A *at* $t = 0$*, and all others join field* B *at* $t = T$*. For* $r \geq \widetilde{r}$*, we have* $\widetilde{\lambda}^A > \widehat{\lambda}^A$*;*
3. *If* $r \geq \widetilde{r}(T)$*, and* $T \geq t^*$*, the equilibrium follows Part 2 of Proposition 3, except that if there are agents still remaining at* T*, they all join field* B *at* $t = T$*.*

First, it could be the case that, when everybody else waits, agents are patient enough to prefer waiting for the deadline to expire than joining A before it. This occurs when $r < \widetilde{r}(T)$, with $\widetilde{r}(T)$ making an agent indifferent between $\frac{v}{r}p_0^A$ and $V_A(T|\underline{\lambda}, \underline{\lambda})$. If that is the case, in the only equilibrium all agents join A at T (Part 1 of Proposition 7).

Let us focus now on the case in which agents are willing to join A at $t = 0$ rather than wait (Parts 2 and 3 of Proposition 7). In the presence of a deadline, this can happen for two reasons: either (i) because $V_A(s|\underline{\lambda}, \underline{\lambda})$ decreases initially, and then reaches its initial level back again at some $\widetilde{s} > 0$, but the deadline T expires before such $\widetilde{s}$ (i.e. $r \in (\widetilde{r}(T), \widetilde{r})$), or (ii) because $V_A(s|\underline{\lambda}, \underline{\lambda})$ stays below its initial value, $\frac{v}{r}p_0^A$, for any $s > 0$ (i.e., $r \geq \widetilde{r}$, as in Part 2 of Proposition 3).

In case (i), Part 1 of Proposition 3 guarantees that, in the absence of a deadline, in equilibrium all agents wait for news indefinitely. Hence, *the presence of a deadline induces agents to be willing to join* A *at* $t = 0$*, and makes equilibrium diversification possible in environments in which agents would have waited for news indefinitely otherwise.*

Let us now focus on case (ii) above, which captures the same environments of Part 2 of Proposition 3.[25] In general, to sustain equilibrium diversification, some time has to elapse between the two surges into fields A and B, respectively, so that the posterior of field B can "catch up" on the posterior of A, and field B can be perceived as the most promising one. Surely, if the deadline T expires after t^*, the sizes of the surges $\widehat{\lambda}^A$ and $\widehat{\lambda}^B$ are the same as in Proposition 3, and the gradual entry into B after t^* occurs following the same process, until either the deadline expires, or no agent remains (Part 3 of Proposition 7).

[25] While in this discussion we distinguish cases (i) and (ii) to compare the equilibrium outcomes to Proposition 3, the proof of Proposition 7 is constructive, and includes a unified algorithm to identify the (unique) equilibrium, and the conditions for diversification to arise, for both cases (i) and (ii).

Moving backwards, consider the range $T \in \left(\widehat{s}(\widehat{\lambda}^A), t^*\right)$, where recall that $\widehat{s}(\widehat{\lambda}^A)$ is the time of the posteriors' crossing in Part 2 of Proposition 3. From Proposition 3, we know that in this range, $V_A(T|\widehat{\lambda}^A, \underline{\lambda}) < V_B(T|\widehat{\lambda}^A, \underline{\lambda}) < V_B(t^*|\widehat{\lambda}^A, \underline{\lambda})$. Hence, to restore the indifference between joining field A at $t = 0$ (and obtaining $\frac{p_0^A v}{r}$), and joining field B at T, the size of the atom of agents joining A at $t = 0$ has to be increased past $\widehat{\lambda}^A$, to some $\widetilde{\lambda}^A > \widehat{\lambda}^A$ such that $V_B(T|\widetilde{\lambda}^A, \underline{\lambda}) = \frac{p_0^A v}{r}$. This implies an increased chance of exhausting all agents on the market and leaving no agents available to explore B. Therefore, under scenario (ii) above, *the presence of the deadline reduces the likelihood of diversification to arise.* At the same time, since in the presence of a deadline no agent can wait for news indefinitely, if diversification arises, all remaining agents join B at T, implying that *the size of field B is no longer bounded above by the size of field A.*

For relatively shorter deadlines (i.e., expiring before $\widehat{s}(\widehat{\lambda}^A)$), for diversification to occur in equilibrium, we still need to guarantee the posteriors to cross before T. When this is not feasible (because the deadline is too short, or the priors are too far from each other), in equilibrium we might have either all agents joining A at $t = 0$, or *two sequential surges of agents joining A, the first at $t = 0$, and the second at the deadline expiration T.* The latter equilibrium structure is more likely to occur when the total pool of agents is large enough.

To summarize, Proposition 7 shows that the presence of a deadline has nuanced effects on the equilibrium structure in the good news case. On the one hand, in environments in which agents are relatively impatient, and diversification can arise in the absence of a deadline, introducing a deadline always makes diversification less likely to occur. We show that, in such environments, a shorter deadline is more harmful to equilibrium diversification than a longer one. On the other hand, when agents are relatively more patient, from Proposition 3 we know that, in the absence of a deadline, they all wait for news indefinitely. Clearly, the presence of a deadline prevents this equilibrium outcome and, for well-chosen deadlines, it may also facilitate equilibrium diversification.

5.3 Subsidies

Research-funding institutions such as the NSF or the NIH often influence future directions of investigation by specifying field-based priorities according to which subsidies, such as scientific grants, are assigned. We now consider the possibility that researchers are awarded grants in one of their potential fields of specializations.

A natural way to model scientific grants is for them to modify the researchers' flow payoffs, by improving scientists' reputations, or by directly increasing their salaries.[26] Formally, we assume that when a grant or subsidy is awarded to one of the two fields $j = A, B$, a researcher who joins j at time $s \geq 0$, receives an expected flow payoff $\theta p_t^j v$ at any $t \geq s$, for some $\theta > 1$. Clearly, the $\theta = 1$ case corresponds to the standard case described in Section 2. In general, the presence of a subsidy increases the cost of waiting before joining the *subsidized* field. However, if an agent is waiting to join the *non-subsidized* field, the benefit from waiting increases (since receiving news may induce the agent to switch course and join the subsidized field instead). Also, if a subsidy for field j is such that $\theta p_0^j > 1$, all agents always prefer joining j at $t = 0$ and receiving $\theta p_0^j \frac{v}{r}$, rather than waiting and receiving at most $\frac{v}{r}$ in the future. Hence, in what follows we focus on $\theta p_0^j \leq 1$.

Grants with Bad News. It is easy to see that if the subsidy is awarded to field A, the equilibrium structure of Proposition 1 remains the same, as the incentives to join field B are reduced even further. Note however that the cost of waiting to join field A is now higher while the benefits of waiting remain the same. This implies that the bandwagon wave into field A starts earlier and is faster than in Proposition 1, resulting in a shorter wave (to maintain indifference, the incentives to wait need to be increased by increasing the information received while waiting, i.e., $\{\lambda_s^A\}$).

Consider now a subsidy awarded to B. If the subsidy is such that $\theta p_0^B > p_0^A$, then the initial ranking of the fields swaps, generating an equilibrium bandwagon wave into B. As it turns out, the equilibrium may still entail a bandwagon wave into B even if $\theta p_0^B < p_0^A$ but $p_0^A - \theta p_0^B$ is relatively small. In fact, as both posteriors p_s^A and p_s^B increase, the difference between them decreases over time, so that θ could compensate for the difference between them at some point in the future. If such a reversal occurs early enough, a bandwagon wave into B arises in equilibrium. Finally, if the subsidy is very small, then the bandwagon wave into A persists, but, since the benefit of waiting for bad news about A is higher, it becomes slower than in the $\theta = 1$ case. Hence, *the presence of the subsidy may alter which field individuals join, but the bandwagon feature of the equilibrium with bad news persists.* These results are summarized in the following result (whose proof is in the Supplemental Appendix).

Proposition 8 (Bad News - Grants) *In the bad news case with grants ($\theta > 1$), the unique equilibrium is as follows:*

1. *If the grant is awarded to A, unless news arrives, all agents gradually join field A.*

[26] Alternatively, we could model research funding by assuming that it affects the shape of the function $\lambda(\cdot)$ (e.g., scientific grants may allow labs to automate some tasks). We illustrate these effects in Section 4.2.2.

Such bandwagon wave starts earlier and is faster than in the $\theta = 1$ case.

2. *If the grant is awarded to B, there exists a $K > 0$ such that a bandwagon wave into B arises in equilibrium if and only if $p_0^A - \theta p_0^B \leq K$. Otherwise a bandwagon wave into A arises, and it is slower than in the $\theta = 1$ case.*

Grants with Good News. In the good news case, scientific grants do not substantially change the structure of the equilibrium described in Proposition 3, but affect both the environments in which research diversification arises, and the relative sizes of the two fields. As mentioned above, a grant awarded to field $i = A, B$ makes agents relatively less patient when they plan to join i (both because the opportunity cost of joining i increases, and the relative benefit if good news on $-i$ arrives decreases), and more willing to wait when they are planning to join $-i$ (because the relative benefit if good news about i arrives increases). Hence, if a grant is awarded to field A, since field A must still be the first one to be explored, the equilibrium outcome in which all agents wait indefinitely for news is less likely to arise as θ grows, and the scenario in which an atom of agents join A at $t = 0$ is more likely to occur. Note that, since the payoff from joining A at $t = 0$ increases in θ, to maintain indifference between joining A at $t = 0$ and joining B later, such atom must increase in θ–that is, as discussed in Section 4.2.1, it is more likely that all agents join A at $t = 0$, but when this is not the case, more exploration of field B can occur in equilibrium.

Next, consider a grant awarded to field B. If such grant is such that $\theta p_0^B > p_0^A$, the same implications just described for a grant awarded to field A hold, with the fields' roles reversed. However, in the case of a smaller grant (i.e., $\theta p_0^B < p_0^A$), field A remains the first one to be explored, but *the atom of agents joining A at $t = 0$ decreases in θ, making research diversification increasingly more likely to arise.* These considerations are summarized in the following result, whose proof is in the Supplemental Appendix.

Proposition 9 (Good News - Grants) *In the good news case with grants, the unique equilibrium is as follows:*

1. *If a grant is awarded to field A, there is $\widetilde{r}^A(\theta)$ decreasing in θ such that the equilibrium follows Proposition 3, with $\lambda^A(\theta)$ increasing in θ.*
2. *If a grant is awarded to field B, we have:*
   (a) *If $\theta p_0^B > p_0^A$, there is $\widetilde{r}^B(\theta)$ decreasing in θ such that the equilibrium follows Proposition 3, with the fields' roles reversed and $\lambda^B(\theta)$ increasing in θ;*

(b) *If* $\theta p_0^B < p_0^A$*, there is* $\widehat{r}^B(\theta)$ *increasing in* θ *such that the equilibrium follows Proposition 3, with* $\lambda^A(\theta)$ *decreasing in* θ*.*

6 Conclusion

6.1 Extensions and Further Research

We study a strategic experimentation setting in which researchers irreversibly choose their specialization between two risky fields over time. The information arrival in each field depends on the mass of researchers who have specialized in that field in the past. We characterize the unique equilibrium in both the bad news and the good news case, highlighting in which environments research diversification occurs.

Several insights follow from our analysis. First, in the bad news case, the unique equilibrium is characterized by a "bandwagon wave" in the dominant field at the offset, implying that no diversification arises unless bad news on that field is uncovered early enough to make the exploration of the second field possible. Such bandwagon wave starts slow, and become faster over time. Moreover, as the priors of the two fields grow further apart, the bandwagon wave starts earlier and becomes increasingly slower.

Second, in the good news case, research diversification can arise in equilibrium when either the agents' pool, or the information-production technology efficacy, are large enough. The fields are investigated in two sequential surges, the first into the dominant field, and the second into the other field. The second surge is more gradual than the first one, and it is bounded above by the size of the first.

Methodologically, our framework provides a tractable setup to study strategic exploration in research. There are several directions in which our study can be extended. Other than the policy tools considered in Section 5, other mechanisms can be implemented to influence scientists' choices. For example, scientific journals play a critical role in determining what researchers work on. Editorial choices focusing on "breakthroughs" vs. "breakdowns" are one margin of such influence. In our model this would translate into endogenizing the nature of the information arrival in the two fields, and also possibly mixing good news and bad news. Also, as wider fields splinter into more specialized smaller ones, one can expect the viability of such subfields to become positively correlated, rather than independent as we assumed. Our framework can be easily altered to account for all these variations and study their impact.[27]

[27] For example, with positive correlation, our qualitative results would not change in the bad news case, while

While the model we study in Section 5.1 is deliberately simplified, as an alternative way to incorporate *citations' benefit* into our model, consider a more nuanced setting in which, once an agent joins a field $i = A, B$ at time $t \geq 0$, his flow payoff is $p^i_\tau \left(1 + f(m^i_\tau - m^i_t)\right) v$ for all $\tau \geq t$, with $f(\cdot)$ strictly increasing, continuous, twice differentiable, and such that $f(0) = 0$. The function $f(\cdot)$ captures the benefit generated, via citations, by the flow of *subsequent* agents joining the same field. In the bad news case, a bandwagon wave into field A still emerges in equilibrium, where agents have a stronger incentive to join field A earlier with respect to our standard setting (corresponding to the $f = 0$ case), as they lose citations if they wait, resulting in a faster equilibrium wave. Moreover there can be an additional coordination equilibrium in which, unless news arrive, nobody ever joins field A, and a bandwagon wave into field B arises.[28] The analysis of this alternative citations' model in the good news case is more subtle and a promising topic for future research.

In our setup all agents are ex-ante identical. Naturally, any *agents' heterogeneity* that makes some better suited or more inclined to join a field rather than others would affect the equilibrium indifference conditions both in the bad news and the good news case, and result in such agents to be more likely to explore their preferred field. The relative sizes of these exogenous subgroups, and their preferences' intensities are likely to affect the structure of the resulting equilibria.[29] Also, one could consider a setting in which agents are compensated for their '*added value*,' in the sense that they enjoy an additional benefit from being already in a field when news arrives, compared to joining after. While this variation is beyond the scope of this paper, it is an interesting direction for exploration.

Finally, note that our model follows 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 a *intergenerational model* with potentially asymmetric rates $\underline{\lambda}^A \geq \underline{\lambda}$ and $\underline{\lambda}^B \geq \underline{\lambda}$, which capture the information uncovered by more senior researchers who made their specialization choices in the past, and are still actively engaged in research. As news arrives about one or both fields, new research ideas 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

in the good news case, in the context of the equilibrium studied in Proposition 3, it would delay the crossing of the posteriors and lower $V_B(\cdot)$, making field B less likely to be explored.

[28] For such equilibrium to exists, the posteriors must cross before the end of the wave, otherwise late agents, who expect fewer citations, have no incentives to join B rather than A. This is possible for $p^A_0 - p^B_0$ small enough, or large enough $\overline{\lambda}$.

[29] Note that in our setup 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 offset.

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 spontaneously yield diversification in the future.

6.2 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 in this research agenda is to study the normative properties of our model. The standard normative approach in the strategic experimentation literature is a utilitarian one, which considers the aggregate utility of all researchers, and accounts for the information externalities they impose on each other. In our application, this approach amounts to taking the perspective of an academic or professional association, which is tasked with maximizing the aggregate utilities of the scientists over their careers. In general, information externalities cause the equilibrium speed of the specializations' choices to be too slow with respect to the utilitarian optimal solution.

However, in our view, our application requires a wider approach to welfare analysis. In particular, scientific research generates externalities to the society at large which go far beyond the boundaries of the researchers' community. This is the reason why governments are substantially engaged in research funding, directly through public institutions, or indirectly through grants and other channels.

Of course, the *speed, or intensity,* at which research is generated, is still relevant from a broader social perspective. Yet, when we think of the society at large, an additional critical aspect of welfare is that of *research diversity, or its breadth.* 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 social crisis, such as environmental issues, or a potential viral pandemic. In other words, society might want to hedge optimally through different research fields.[30]

The equilibria in this paper have very different properties with respect to research diversification. Specifically, our results suggest that ensuring diversification is particularly problematic to achieve in the bad news case, where, as long as no bad news arrive, only one research field is

[30] As 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).

ever explored. Such lack of diversification persists in the presence of citation benefits, deadline effects, and scientific grants awarded to minor fields. In the good news scenario, potentially both fields can be explored in equilibrium, but societal inefficiencies may still arise in terms of the scope and the timing of the fields' exploration. Such potential diversification may be harder to achieve after the introduction of citation benefits and tenure clocks, while well-designed grants can promote diversification and affect the equilibrium relative sizes of the fields, potentially bringing them closer to the socially optimal research portfolio.[31]

7 Appendix

In this section we present the proofs of the main results of the paper. All remaining proofs can be found in the Supplemental Appendix (at https://mariagiovannabaccara.com/Appendix_RW.pdf).

7.1 Proofs for the Bad News Case

Proof of Proposition 1: (1) Let $r \leq \overline{r}$, and conjecture an equilibrium as the one described in the proposition. 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 $s \in \left[\underline{t}, \overline{t}\right]$, we need (1) to be constant over the interval–that is, $V(s) = \kappa$ for all $s \in \left[\underline{t}, \overline{t}\right]$. A necessary condition for this to be true is $V'(s) = 0$. After some algebraic steps, such condition yields the ODE

$$\frac{\dot{p}_s^A}{(p_s^A)^2} = \frac{r}{p_s^B}, \tag{8}$$

with solution

$$p_s^A = \frac{1}{-\int \frac{r}{p_s^B} ds - C} \tag{9}$$

for some constant C. We can use (9) to solve for the equilibrium λ_s^A, obtaining

$$\lambda_s^A = \frac{\dot{p}_s^A}{p_s^A(1 - p_s^A)} = \frac{r}{-p_s^B \left(\int \frac{r}{p_s^B} ds + C + 1\right)}. \tag{10}$$

[31] Recent work by Hill and Stein (2023 and 2024) empirically quantifies the impact of scientists' incentives on breadth and depth of research.

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

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

and

$$\lambda_s^A(C) = \frac{-r\underline{\lambda}\left[\left(1-p_0^B\right)e^{-\underline{\lambda}s} + p_0^B\right]}{rs\underline{\lambda}p_0^B - re^{-\underline{\lambda}s}(1-p_0^B) + \underline{\lambda}p_0^B(C+1)}. \tag{11}$$

Next, note that condition (8) needs to be satisfied at the two interval extremes. At $\underline{t}$, (8) implies $\frac{\underline{\lambda}(1-p^A_{\underline{t}})}{p^A_{\underline{t}}} = \frac{r}{p^B_{\underline{t}}}$. From $\frac{p^j_{\underline{t}}}{1-p^j_{\underline{t}}} = \frac{p^j_0}{1-p^j_0}e^{\int_0^{\underline{t}}\underline{\lambda}dz} = \frac{p_0^A e^{\underline{t}\underline{\lambda}}}{1-p_0^A}$ for $j = A, B$, we obtain

$$\frac{rp_0^A}{\underline{\lambda}(1-p_0^A) - rp_0^A e^{\underline{t}\underline{\lambda}}} = \frac{p_0^B}{1-p_0^B}, \tag{12}$$

which, solving for $\underline{t}$, yields (4), where the last inequality is guaranteed by $r \leq \bar{r}$. Next, (11) at $\bar{t}$ implies

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

Finally, to find $\bar{t}$ and C, note that (8) at $\bar{t}$ implies

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

Since

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

we have

$$
\frac{p_0^A}{p_0^A + (1-p_0^A) e^{-\ln\left[\frac{\left(1-p_0^A\right)\underline{\lambda}}{r p_0^A} - \frac{1-p_0^B}{p_0^B}\right] + \int\limits_{\underline{t}}^{\bar{t}} \frac{-r\underline{\lambda}\left[\left(1-p_0^B\right)e^{-\underline{\lambda}z} + p_0^B\right]}{-r\underline{\lambda}p_0^B z + r(1-p_0^B)e^{-\underline{\lambda}z} - \underline{\lambda}p_0^B(C+1)} dz}}
$$

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

Therefore, we obtain the condition

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

Since

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

the condition above becomes

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

Recall that (12) implies

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

Hence, (13) becomes

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

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

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

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 $\left[p_0^B + \left(1-p_0^B\right)e^{-\underline{\lambda} t}\right]\left(1-\frac{\underline{\lambda}}{\overline{\lambda}}\right) > 0$, while for $t \to \infty$,

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

Therefore, (14) identifies a unique solution for $\overline{t}$. In particular, such solution satisfies $\overline{t} = \underline{t} + \frac{(\overline{\lambda}-\underline{\lambda})}{\underline{\lambda}\overline{\lambda}p^B_{\overline{t}}}$. From here, we can get expressions for the equilibrium processes. From (13), we have $rp_0^B = -r\underline{\lambda}p_0^B\underline{t} - \underline{\lambda}p_0^B\left[C+1\right]$, implying $C+1 = \frac{-r(\underline{\lambda}\underline{t}+1)}{\underline{\lambda}}$ and $C = \frac{-r(\underline{\lambda}\underline{t}+1)-\underline{\lambda}}{\underline{\lambda}}$. Therefore, from (11), we obtain, for $s \in \left[\underline{t}, \overline{t}\right]$

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

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

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 $s > \overline{t}$. To do so, we will show that $\frac{dV(s)}{ds} > 0$ for $s < \underline{t}$, and that $\frac{dV(s)}{ds} < 0$ for $s > \overline{t}$. By continuity, this implies that $V(s)$ attains it maximum when $\frac{dV(s)}{ds} = 0$, which justifies our first-order conditions approach. For $s < \underline{t}$, we have:

$$\begin{aligned} V(s) = \; & \tfrac{v}{r}\Big[\int_0^s \left[(1-p_t^A)\underline{\lambda}p_t^B + (1-p_t^B)\underline{\lambda}p_t^A\right]e^{-\int_0^t (r+(1-p_z^A)\underline{\lambda}+(1-p_z^B)\underline{\lambda})dz}dt + \\ & (e^{-\int_0^s (r+(1-p_z^A)\underline{\lambda}+(1-p_z^B)\underline{\lambda})dz})p_s^A\Big], \end{aligned}$$

yielding

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

After some algebra, it is easy to check that $\frac{dV(s)}{ds} > 0$ if and only if $\frac{p_s^B(1-p_s^A)}{p_s^A} > \frac{r}{\underline{\lambda}}$. From (12), we have $\underline{\lambda}p^B_{\underline{t}}(1-p^A_{\underline{t}}) - p^A_{\underline{t}}r = 0$, which is equivalent to $\frac{p^B_{\underline{t}}(1-p^A_{\underline{t}})}{p^A_{\underline{t}}} = \frac{r}{\underline{\lambda}}$.Therefore, $\frac{dV(s)}{ds} > 0$ if and only if

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

The last condition is equivalent to

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

which, in turn, is equivalent to $e^{-s\underline{\lambda}} > e^{-t\underline{\lambda}}$, which holds for $s < \underline{t}$. For $s > \bar{t}$, $\frac{dV(s)}{ds} < 0$ if and only if $(1-p_s^A)\overline{\lambda}p_s^B - p_s^A r < 0$. Recall that (11) at $\bar{t}$ implies $\frac{p_{\bar{t}}^B(1-p_{\bar{t}}^A)}{p_{\bar{t}}^A} = \frac{r}{\overline{\lambda}}$. Therefore, $\frac{dV(s)}{ds} < 0$ if and only if

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

The last condition is equivalent to

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

which is equivalent to $e^{-(s-\bar{t})\overline{\lambda}} < 1$, 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 $\underline{t}$, (3) 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 $r > \bar{r}$, so that, from (4), we obtain $\underline{t} < 0$. Let $\widehat{\lambda} > \underline{\lambda}$ solve: $\frac{\widehat{\lambda}(1-p_0^A)}{rp_0^A} = \frac{1}{p_0^B}$. 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 $\widehat{\lambda}$. 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 Supplemental Appendix. Since $V(s)$ is constant for all $s \in (0, \bar{t}]$, this is an equilibrium. Moreover, since $\lim_{s\to 0^+} V(s) = p_0^A \frac{v}{r}$, 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 $\widetilde{\lambda} > \widehat{\lambda}$ joins A at $t = 0$, (3) would imply $V'(0) > 0$, generating a profitable deviation for any agent in the atom. Also, any atom of agents joining A at any $s \in (0, \bar{t}]$ would create a similar profitable deviation. Therefore, the only equilibrium possible if there are agents left on the market after an atom of size $\widehat{\lambda}$ joins A, is for them to join gradually after the initial atom. ■

7.2 Proofs for the Good News Case

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* $p^A_{\widehat{s}(\lambda^A)} = p^B_{\widehat{s}(\lambda^A)}$, *and* $p_s^A < p_s^B$ *for any* $s > \widehat{s}(\lambda^A)$. *In particular,*

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

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

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

■

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

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

Proof of Lemma A2: For Part 1, for $j = A, B$,

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

implies

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

Therefore, $sign\left[V_j'(s|\lambda^A, \lambda^B)\right] = sign\left[p_s^{-j}\lambda^{-j}(1-p_s^j) - rp_s^j\right]$. To verify Part 2, note that

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

■

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

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

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

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 $\overline{t}>\underline{t}$, 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 $t\in\left[\underline{t},\overline{t}\right]$; (iii) $\lambda_t^i = \lambda^{-i}(1-p_s^{-i})$ for any $t\in\left[\underline{t},\overline{t}\right]$; (iv) $\frac{\lambda^{-i}(1-p_t^i)}{p_t^i} < \frac{r}{p_t^{-i}}$ for some $t\in(\overline{t},\overline{t}+\varepsilon)$.*

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

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

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 $t\in\left[\underline{t},\overline{t}\right]$ is an immediate consequence of Lemma 1. To show (iii), since by Part 2 of Lemma A2, we have that $sign\left[\left(\frac{(1-p_t^B)}{p_t^B}p_t^A\right)'\right] = sign\left[\lambda_t^B - \lambda^A(1-p_t^A)\right]$. 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 $t\in\left[\underline{t},\overline{t}\right]$. Hence, in the interval $\left[\underline{t},\overline{t}\right]$, we must have $p_t^B > p_t^A$ and $\lambda_t^B < \lambda^A$, including at the end of the interval. After time $\overline{t}$, as p_t^A is decreasing, we have that $\lambda_{\overline{t}}^B - \lambda^A(1-p_t^A) < 0$ for any $t\in(\overline{t},\overline{t}+\varepsilon)$ (ε chosen so that there are no other atoms of agents joining the fields within $(\overline{t},\overline{t}+\varepsilon)$). By Part 2 of Lemma A2, this implies that $(\frac{(1-p_t^B)}{p_t^B}p_t^A)' < 0$ so that, by Part 1 of Lemma A2, $\frac{\lambda^A(1-p_t^B)}{p_t^B} < \frac{r}{p_t^A}$ for $t\in(\overline{t},\overline{t}+\varepsilon)$. ■

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 $\widehat{p}_s^A$ are the posteriors at s under λ^A and $\hat{\lambda}^A$, respectively, we have $p_s^A < \widehat{p}_s^A$, 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 claims follows. Next, consider $V_B(s|\lambda^A, \lambda^B)$. Note that $V_B(s)\frac{r}{v}$ can be rewritten as

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

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

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

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

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

with respect to λ^A, which is

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

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

Proof of Proposition 3:

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

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

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 $r \geq \frac{\underline{\lambda} p_0^B(1-p_0^A)}{p_0^A}$ but that $V_A(s|\underline{\lambda},\underline{\lambda}) > \frac{p_0^A v}{r}$ for some $s > 0$. 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

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

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

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

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

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

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

Proof of Part 2. Assume that $r > \widetilde{r}$, where $\widetilde{r}$ is defined as in the proof of Part 1. From the discussion above, this implies $r \geq \frac{\underline{\lambda} p_0^B(1-p_0^A)}{p_0^A}$ 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 $r \geq \frac{\underline{\lambda} p_0^B(1-p_0^A)}{p_0^A}$, $V_A(s|\underline{\lambda},\underline{\lambda}) < \frac{p_0^A v}{r}$ 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 $r \geq \frac{\underline{\lambda} p_0^B (1-p_0^A)}{p_0^A}$, $V_B(s|\underline{\lambda}, \underline{\lambda}) < V_A(s|\underline{\lambda}, \underline{\lambda}) < \frac{p_0^A v}{r}$ 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 $\lim_{s\to\infty} V_B(s|\lambda^A, \underline{\lambda}) = \frac{p_0^A v}{r}$, but if $\lim_{s\to\infty} V_B'(s|\lambda^A, \underline{\lambda}) < 0$ 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 $\lambda^B(\lambda^A) \equiv \lambda(m - \lambda^{-1}(\lambda^A))$–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 $\widehat{\lambda}^A$. *(ii)* After the crossing where $p^A_{\widehat{s}(\widehat{\lambda}^A)} = p^B_{\widehat{s}(\widehat{\lambda}^A)}$ (see Lemma A1), at a point in time $t^* \in (\widehat{s}(\widehat{\lambda}^A), \infty)$ where $V_B(t^*|\widehat{\lambda}^A, \underline{\lambda}) = \sup_s V_B(s|\widehat{\lambda}^A, \underline{\lambda}) = \frac{v}{r}p^A_0$, there is an atom of agents $\hat{\lambda}^B$ that join B, where (a) If $\lambda^B(\widehat{\lambda}^A) - \widehat{\lambda}^A(1 - p^A_{t^*}) \leq 0$ (implying m being low enough) then we set $\hat{\lambda}^B = \lambda^B(\widehat{\lambda}^A)$. Hence, all remaining agents join B at t^*. (b) If $\lambda^B(\widehat{\lambda}^A) - \widehat{\lambda}^A(1 - p^A_{t^*}) > 0$, then set the atom of agents joining B at $\hat{\lambda}^B = \widehat{\lambda}^A(1 - p^A_{t^*})$. Some of the remaining agents join B gradually, so that λ^B_t for $t > t^*$ is such that the posteriors satisfy $\widehat{\lambda}_A(1 - p^B_t)\frac{p^A_t}{p^B_t} = r$. For this to be possible, by Lemma A4 we need $\lambda^B_t < \widehat{\lambda}^A$, 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 $\sup_s V_B(s|\underline{\lambda}, \underline{\lambda}) < \frac{v}{r}p^A_0 < \sup_s V_B(s|\bar{\lambda}, \underline{\lambda})$. 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 $\widehat{\lambda}^A < \bar{\lambda}$ such that $\max_s V_B(s|\widehat{\lambda}^A, \underline{\lambda}) = \frac{v}{r}p^A_0$. If t^* is defined to be the lowest element in $\arg\max_s V_B(s|\widehat{\lambda}^A, \underline{\lambda})$, 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, $\widehat{\lambda}^A < \bar{\lambda}$ is the arrival rate generated by the atom of agents joining field A at $t = 0$–that is, in equilibrium the atom of agents $m^A_0 = \lambda^{-1}(\widehat{\lambda}^A) < m$ joins field A at $t = 0$. Next note that for the parameters we consider, $V'_A(0|\widehat{\lambda}^A, \underline{\lambda}) < 0$ 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 $p^A_{t^*} \leq p^B_{t^*}$. We already know that $V_A(s|\widehat{\lambda}^A, \underline{\lambda})$ can only have corner maximizers at either $s = 0$ or $s = \infty$. Since

$$\sup_s V_B(s|\widehat{\lambda}^A, \underline{\lambda}) = \frac{v}{r}p^A_0 \geq \lim_{s\to\infty} V_B(s|\widehat{\lambda}^A, \underline{\lambda}) = \lim_{s\to\infty} V_A(s|\widehat{\lambda}^A, \underline{\lambda}),$$

implying that as long as there is no new atom of agents in B, $V_A(s|\widehat{\lambda}^A, \underline{\lambda}) \leq \frac{v}{r}p^A_0$ 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 $p^A_s \leq p^B_s$, it must be the case that at t^*, where

$$V_B(t^*|\widehat{\lambda}^A, \underline{\lambda}) = \frac{v}{r}p^A_0 \geq V_A(t^*|\widehat{\lambda}^A, \underline{\lambda}),$$

we have that $p^A_{t^*} \leq p^B_{t^*}$–that is t^* *occurs after the crossing* $\widehat{s}(\widehat{\lambda}^A)$. Also, for this to be an equilibrium, we need $V_B(s|\widehat{\lambda}^A, \hat{\lambda}^B)$ to be weakly decreasing after t^*. A sufficient condition for this is $\lambda^B_t \leq \widehat{\lambda}^A(1-p^A_t)$ for any $t > t^*$, as p^A_t 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 $\max V_B(t^*|\widehat{\lambda}^A, \underline{\lambda}) = \frac{v}{r}p^A_0$, i.e., $\widehat{\lambda}^A(1-p^B_{t^*})\frac{p^A_{t^*}}{p^B_{t^*}} = r$. Then, $V_B(t|\widehat{\lambda}^A, \lambda^B_t)$ 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 $V'_B(t|\widehat{\lambda}^A, \lambda^B_t) \leq 0$, we cannot have another peak of $V_B(t|\widehat{\lambda}^A, \lambda^B_t)$. Moreover, as $\lambda^B_t < \widehat{\lambda}^A$ for all $t \geq 0$, the posteriors cannot cross again. Hence, $V_B(t|\widehat{\lambda}^A, \lambda^B_t) > V_A(t|\widehat{\lambda}^A, \lambda^B_t)$ for all $t > t^*$, ruling out any further entry into A. ■

Proof of Proposition 7: **(1)** If $r < \overline{r}$, 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 $r \geq \frac{\lambda p^B_0(1-p^A_0)}{p^A_0}$ but that $V_A(s|\underline{\lambda}, \underline{\lambda}) > \frac{p^A_0 v}{r}$ for some $s \in (0, T]$. Since $V_A(0) = \frac{p^A_0 v}{r}$, 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 $\widetilde{r}(T)$ to be such that $V_A(T|\underline{\lambda}, \underline{\lambda}) = \frac{v}{r}p^A_0$. From the discussion above, we have $\overline{r} \leq \widetilde{r}(T)$ for all T. Hence, *for any* $r < \widetilde{r}(T)$, *all agents join* A *at the deadline* T. To see that $\widetilde{r}(T)$ is increasing in T, consider $T' < T$, and $\widetilde{r}(T)$ such that $V_A(T|\underline{\lambda}, \underline{\lambda}) = \frac{v}{r}p^A_0$. For such a $\widetilde{r}(T)$, we have $V_A(T'|\underline{\lambda}, \underline{\lambda}) < V_A(T|\underline{\lambda}, \underline{\lambda})$. Therefore, to obtain $V_A(T'|\underline{\lambda}, \underline{\lambda}) = \frac{v}{r}p^A_0$, since $V_A(T|\underline{\lambda}, \underline{\lambda})$ decreasing in r, we need $\widetilde{r}(T') < \widetilde{r}(T)$.

(2) Assume that $r > \widetilde{r}(T)$ (implying that for all $s \in (0, T]$, $V_A(s|\underline{\lambda}, \underline{\lambda}) < \frac{p^A_0 v}{r}$), and $T \in (0, t^*)$. Recall that $\widehat{s}(\lambda)$ 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 $\widehat{\lambda}^A$ and t^* the equilibrium size of the atom joining A at $t = 0$, and the time of the atom joining B in Proposition 3, respectively. Also, define $\mu^A_A(T)$ be such that $V_A(T|\mu^A_A(T), \underline{\lambda}) = \frac{v}{r}p^A_0$ and $\mu^A_B(T)$ be such that $V_B(T|\mu^A_B(T), \underline{\lambda}) = \frac{v}{r}p^A_0$. 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 $V_A(T|\overline{\lambda}, \underline{\lambda}) > \frac{v}{r}p^A_0$, *or equivalently, by Lemma A5,* $\overline{\lambda} > \mu^A_A(T)$. 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 $\widetilde{\lambda}^A = \mu_B^A(T)$ 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 $V_A(T|\overline{\lambda}, \underline{\lambda}) \leq \frac{v}{r}p_0^A$, *or equivalently,* $\overline{\lambda} \leq \mu_A^A(T)$. Then, two cases are possible: (a) If $V_B(T|\overline{\lambda}, \underline{\lambda}) < \frac{v}{r}p_0^A$, then all agents join A at $t=0$; (b) If $V_B(T|\overline{\lambda}, \underline{\lambda}) > \frac{v}{r}p_0^A$, then $\widetilde{\lambda}^A = \mu_B^A(T)$ agents join A at $t=0$, and the remaining join B at T. Again, note that in the non-generic case $V_B(T|\overline{\lambda}, \underline{\lambda}) = \frac{v}{r}p_0^A$, 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 3, for $r > \widetilde{r}$, $V_B(s|\widehat{\lambda}^A, \underline{\lambda})$ peaks at t^*, we must have $\widetilde{\lambda}^A > \widehat{\lambda}^A$.

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

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," Theoretical 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,” Gesammlte Aufsaetze zur Wissenschaftslehre (Tubingen, 1922), 524-55.