Startups and Upstarts*

Yu Awaya† and Vijay Krishna‡

April 20, 2019

Abstract

We study a continuous-time R&D race between an established firm and a startup under asymmetric information. R&D investment brings success stochastically but only if the innovation is feasible. The only asymmetry between the firms is that the established firm has better information about the feasibility of the innovation. We show that there is an equilibrium in which the poorly-informed startup wins *more* often, and has higher expected profits, than the better-informed incumbent. When the informational asymmetry is large, this is the *unique* equilibrium outcome. Even though better information is a competitive disadvantage, the value of information is positive.

*The research reported here was supported by a grant from the National Science Foundation (SES-1626783). We are grateful to Nageeb Ali, Gaurab Aryal, Kaustav Das, John Morgan and Tsz-Ning Wong for helpful discussions.

†Department of Economics, University of Rochester, E-mail: yuawaya@gmail.com.

‡Department of Economics, Penn State University, E-mail: vkrishna@psu.edu.

1 Introduction

Why Tesla and not GM or Toyota? Why Amazon and not Sears or Wal-Mart? Why are startups the source of so many innovations instead of, and at the expense of, established firms? In his landmark history of the hard-disk industry over two decades, Christiansen (1997) found that the market for each new generation of disk drives—typically, smaller in size—was dominated by a different set of firms. Of the 17 firms in the industry in 1976, only IBM's disk-drive division survived until 1995. In the same period, there were 129 entrants but 109 of these failed to make the transition to later generations (Christiansen, 1997, p. 22). Many technological innovations came from startups.

What advantage does a startup have over an established firm? In one of his many classics, Arrow (1962) argued that because of the "monopolist's disincentive created by his preinvention profits" (p. 622) an entrant would have more to gain from an innovation. This is sometimes called the "replacement effect" because by successfully innovating, the monopolist would only be replacing himself while the entrant would be replacing the monopolist.[1] Running counter to Arrow's reasoning are the strong incentives that an incumbent has to protect its monopoly position. This stems from the Econ 101 $m > 2d$ inequality—monopoly profits exceed total profits in a duopoly—which can be cleverly rearranged as $m - d > d$. In this form, it says that the incentive of the incumbent to preserve its monopoly is greater than the incentive of the startup to enter as a duopolist (Gilbert and Newberry, 1982). This "preemption effect" is at odds with the replacement effect. There are other forces that may favor incumbents as well—lower R&D costs or an existing stock of R&D capital.[2] Whether the balance of all these forces favors incumbents or startups is then an empirical question. In a recent paper, Igami (2017) went back to the disk-drive industry and constructed a structural model to try to answer this question. A large fraction of firms failed to make the transition from 5.25- to 3.5-inch drives. Igami found evidence that Arrow's replacement effect played a substantial role, explaining about 60% of the turnover.

In this paper, we study a continuous-time, winner-take-all R&D race between an established firm and a startup in which we identify an entirely new effect that works to the detriment of the established firm—a "curse of information." Incumbent firms have more experience and so have better information about the likelihood of success of the new innovation. The firms are alike in all other respects and so the replacement and preemption effects are absent. Formally, there are two states of nature. In one, the innovation is feasible and R&D brings success stochastically—in the manner of exponential bandits. In the other state, the innovation is not feasible. Firms do not know the state but receive informative private signals about it and the established

[1] Arrow (1962) considered cost-reducing innovations and also the possibility of post-innovation competition.

[2] Schumpeter (1942) argued that monopolies were more conducive to innovation than perfect competition.

firm's signal is more accurate than that of the startup. A good signal makes a firm optimistic and a bad one pessimistic. As the race proceeds, lack of success causes both firms to become increasingly pessimistic about the feasibility of the innovation. R&D is costly and each firm must then decide when to quit, a decision is observed by its rival and is irrevocable.

Our main result is

Theorem 1 *There is an equilibrium of the R&D game in which the less-informed startup wins* more *often, and has a higher payoff, than the better-informed incumbent. Moreover, if the quality of the incumbent's information is* much *better than that of the startup, then this is the* only *equilibrium.*

Our result shows that in an otherwise symmetric situation, the incumbent's informational advantage becomes a competitive disadvantage—it wins the R&D race less often than the startup and, as we will see, has a lower payoff as well. The startup is favored to win precisely because it is less informed!

We call such an equilibrium an "upstart equilibrium." In such an equilibrium, the less-informed startup is, quite naturally, willing to learn from the incumbent. But because of its superior information, the incumbent is unwilling to learn from the startup/upstart. This unbalanced learning is why the startup wins more often and the better information available to the established firm becomes a curse.

Precisely, both the incumbent and the startup play strategies that reveal over time whether or not they are optimistic. But since the incumbent's information is of higher quality than that of the startup, when pessimistic it exits early in the race based solely on its own information. The reason is that while the startup also reveals its signal during the play of the game, this comes too late to make it worthwhile for a pessimistic incumbent to stay and learn. On the other hand, the information does not come too late for the optimistic incumbent for whom it is worthwhile to stay and learn the startup's signal. Thus a pessimistic incumbent exits early while an optimistic one stays. This means that the startup can learn the incumbent's information at low cost. During the play of the game, both the optimistic and the pessimistic startup learn the incumbent's information but only the optimistic incumbent learns the startup's information.

It is then not too hard to argue that if both firms are optimistic or both are pessimistic, they exit at the same time. The same is true when the incumbent is optimistic and the startup pessimistic—this is because they both learn each other's signal. The remaining case is one with a pessimistic incumbent and an optimistic startup. The incumbent exits early and so the startup learns that it is pessimistic. But its own optimism causes the startup to continue with R&D nevertheless. Now the startup has a *greater* chance of winning than does the incumbent.

The upstart equilibrium outcome has some salient features. While it can be supported as a perfect Bayesian equilibrium, it does not rely on any particular choice of off-equilibrium beliefs. More important, when the informational advantage of the incumbent is large, it is the unique Nash equilibrium outcome. The formal argument relies on the iterated elimination of dominated strategies—in our game, this procedure leaves a single outcome. Some idea of the reasoning can be gauged by noting that in these circumstances there cannot be a "mirror equilibrium" in which the roles of the two firms are reversed and the incumbent learns more from the startup than the other way around. Because the startup's information is of very poor quality, it is not worthwhile for the incumbent to invest in learning this. So when the startup's information is very poor, a mirror equilibrium does not exist. In our formal analysis, we rule out not only the mirror equilibrium but all others as well.

In our equilibrium, a firm may suffer from ex post regret—had it known the other's signal, it may have wanted to stay longer in the race or may have wanted to exit earlier.[3] In equilibrium, the established firm never regrets staying too long but may regret exiting too early. The startup, on the other hand, never regrets exiting too early but may regret staying too long.

Intuition suggests that information should confer a strategic advantage. In our model, it is a disadvantage. One might rightly wonder whether this is because in the game we study, the value of information is negative.[4] This is not the case. We show below that in the upstart equilibrium, the value of information is positive for both firms. In other words, neither firm can increase its equilibrium payoff by decreasing the quality of its own information. Theorem 1 above is a comparison of payoffs across firms and does not contradict the fact that each firm has the individual incentive to become better informed. Finally, we also ask whether it might be in the incumbent firm's interest to decrease the quality of its information so drastically that it is completely uninformed. We show that such "willful ignorance" cannot be profitable.

Overconfidence The popular press is full of stories of brash Silicon Valley entrepreneurs who embark on risky projects that established firms deem unworthy. Most of these startups fail but some do succeed and perhaps lead to the kinds of disruption that is observed. Some studies have argued that this over-investment in risky projects stems not from risk-loving preferences but rather from overconfidence.[5] As one observer of the startup phenomenon has written:

> "In the delusions of entrepreneurs are the seeds of technological progress."
> (Surowiecki, 2014)

[3]Moscarini and Squintani (2010) call these the "quitter's curse" and "survivor's curse," respectively.

[4]As is the case, for instance, in the classic "lemons problem." There is no trade when the seller is informed but there would be if she were not.

[5]See, for example, Wu and Knott (2006). Another study found that entrepreneurs are prone to overestimate their own life spans relative to the rest of the population (Reitveld et al. 2013)!

In this view, the Elon Musks of the world drive innovation because of unwarranted self-confidence. They remain optimistic in environments that the GMs of the world are pessimistic about, and perhaps realistically so.

While our model and analysis has no behavioral or psychological elements, it can be seen as providing a rational reinterpretation of such behavior. When the incumbent firm's information is not favorable to the project while the startup's is, the former is pessimistic and the latter optimistic. The startup invests in R&D while the better-informed incumbent does not. In these circumstances, the rational optimism of the startup would be observationally equivalent to overconfidence. In single-person problems, Benoît and Dubra (2011) argued that in many situations a fully rational Bayesian agent may end up with beliefs that, to an outside observer, would seem overconfident. They showed that this "apparent overconfidence" could be generated solely by the structure of information available to the agent. Our model and equilibrium can be interpreted as doing the same, but now in a strategic situation with more than one agent. The postulated information structure and the upstart equilibrium results in behavior that an outside observer may well attribute to overconfidence.

Related literature The basic model of this paper is rather standard. R&D races where the arrival times of success are exponentially distributed and there is uncertainty about the arrival rates were first studied by Choi (1991). Malueg and Tsutsui (1997) extend Choi's model to allow for flexibility in the intensity of R&D. In a variant of Choi's model, Wong (2018) examines the consequences of imperfect patent protection thereby relaxing the winner-take-all structure common to most of the literature.[6] Chatterjee and Evans (2004) introduce another kind of uncertainty—there are two alternative paths to success and it is not known which is the correct one. Firms may switch from one path to another based on their beliefs. Das and Klein (2018) study a similar model and show that there is a unique Markov perfect equilibrium which is efficient when firms are symmetric in R&D ability and not otherwise.

In all of these models, however, there is no asymmetry of information—firms' equilibrium beliefs are identical. In our model, firms receive private signals prior to the race and the resulting asymmetry of beliefs is the key to our results.

The model of Moscarini and Squintani (2010) is, in its basic structure, most closely related to ours. These authors study a very general set-up with arbitrary distributions of arrival times (not necessarily exponential), continuous signals and differing costs and benefits of R&D. They show the possibility that the exit of one firm leads the other to regret staying as long—the firm suffers from a "survivor's curse"—and so it also exits as soon as possible.[7] Our model differs from that of Moscarini and Squintani

[6] In Wong's model the feasibility of the projects is independent across firms and so one firm cannot learn from the other firm's lack of success. In our model, and the others mentioned, the feasibility is perfectly correlated.

[7] Moscarini and Squintani (2010) mention the so-called fifth-generation computers initiative as an example. When the Japanese consortium abandoned this once promising technology, firms in the

in that we have discrete states and signals. At the same time, it specializes their model by assuming exponentially distributed arrival times, identical costs and benefits of R&D and *comparable* information. Moscarini and Squintani also point to a "quitter's curse"—regret at exiting too early. When the firms' information is comparable, as we assume, even the curses are asymmetrically distributed. The better-informed firm never suffers from the survivor's curse but may suffer from the quitter's curse. The opposite is true for the less-informed firm. Finally, we derive circumstances in which there is a unique equilibrium outcome and these too depend on the relative quality of the firms' information.

R&D race models are cousins of strategic experimentation problems, especially those with exponential bandits as in Keller, Rady and Cripps (2005). Unlike the R&D models, the latter are not winner-take-all as one person's success does not preclude the other's. Also, in these models it is possible to switch back and forth between the risky and safe arms, unlike the irrevocable exit assumption in R&D race models. While most of these models were studied under symmetric information, in recent work, Dong (2018) has studied a variant with asymmetric and comparable information—one person has a private signal but the other is completely uninformed.[8] She finds that this asymmetry induces more experimentation than if the situation were symmetric.

R&D race models also share important features with wars of attrition—in particular, the winner-take-all and irrevocable exit assumptions. There is, of course, a vast literature on wars of attrition with and without incomplete information. A related paper in this vein is by Chen and Ishida (2017), who study a model which combines elements from strategic experimentation with wars of attrition. As in strategic experimentation models, one firm's successful innovation does not preclude successful innovation by the other firm. As in the war of attrition, exit by one firm ends the game. Firms are asymmetric in how efficient they are at R&D. There is a mixed strategy equilibrium and Chen and Ishida (2017) exhibit the possibility that the less efficient firm may win more often.

The remainder of the paper is organized as follows. The model of an R&D race is outlined in the next section. Section 3 studies, as a benchmark, the case of a single firm without competition. There is no surprise here—if alone, the better informed firm is more likely to succeed than the less informed firm. In Section 4, we study the case of two competing firms and exhibit the upstart equilibrium mentioned above. Section 5 then shows that this equilibrium is unique when the asymmetry in the quality of information is large. Equilibrium behavior is compared to the joint-profit maximizing solution in the next section. In Section 7, we show that despite the fact that in equilibrium the less informed firm wins more often, the value of information is

US and UK followed.

[8]Klein and Wagner (2018) study a bandit problem where the quality of information of the players is the same.

positive for both firms. Finally, in Section 8 we show that the main result generalizes when the firms may get more than two signals and so have finer information. An appendix considers the special case when there is no asymmetric information and the firms hold common beliefs throughout.

2 Preliminaries

Two firms compete in an R&D race to produce an innovation. Time runs continuously, the horizon is infinite and the interest rate is $r > 0$. The firm that succeeds first will obtain a patent that yields flow monopoly profits of m forever after. Each firm decides on how long it wants to actively participate in the race, if at all, and must incur a flow cost of c while it is active. A firm only chooses whether or not to be active, and not its intensity of R&D. Once a firm quits, it cannot rejoin the race. Also, if a firm quits at time t, say, then this is immediately observed by the other firm.[9] The game ends either if one of the firms succeeds or once both firms quit.

Whether or not the innovation is worth pursuing is uncertain, however, and depends on an unknown state of nature that may be G ("good") or B ("bad") with prior probabilities π and $1-\pi$, respectively. In state B, the innovation is not technologically feasible and all R&D activity is futile. In state G, it is feasible and success arrives at a Poisson rate $\lambda > 0$ per instant, independently for each firm provided, of course, that the firm is still active. This means that the distribution of arrival times of success is *exponential*, that is, the probability that in state G a firm will succeed before time t is $1 - e^{-\lambda t}$.

The two firms are alike in all respects but one—firm 1 (the "incumbent" or established firm) is better informed about the state of nature, G or B, than is firm 2 (the "startup" or entrant firm). Specifically, before the race starts, each firm i receives a noisy private signal $s_i \in \{g_i, b_i\}$ about the state. Conditional on the state, the signals of the two firms are independent and

$$\Pr[g_i \mid G] = \Pr[b_i \mid B] = q_i > \tfrac{1}{2}$$

We will refer to q_i as the *quality* of i's signal or information.[10] Throughout, we will assume that firm 1's signal is of higher quality than that of firm 2 in the sense that $q_1 > q_2$ and so firm 1 is better informed.

Denote by $p(s_i)$ the posterior probability that the state is G conditional on the signal s_i, that is,

$$p(s_i) = \Pr[G \mid s_i]$$

[9] This could happen with a delay $\Delta > 0$ so that if a firm quits at time t, the other firm learns of this only at time $t + \Delta$. We have chosen to set $\Delta = 0$ to simplify the exposition but our analysis is robust to the case when Δ is small (details are available from the authors).

[10] The assumption that $\Pr[g_i \mid G] = \Pr[b_i \mid B]$ is made only for simplicity. It would be enough to assume that firm 1's signals were more informative than firm 2's signals in the sense of Blackwell.

and similarly, denote by $p(s_1, s_2)$ the posterior probability that the state is G conditional on the signals (s_1, s_2), that is,

$$p(s_1, s_2) = \Pr[G \mid s_1, s_2]$$

It is easy to see that since firm 1's signal is more accurate than firm 2's signal, that is, $q_1 > q_2$,

$$p(b_1, b_2) < p(b_1, g_2) < p(g_1, b_2) < p(g_1, g_2) \tag{1}$$

It is useful to define p^* to be such that if a firm believes that the probability that the state is G is p^*, then the flow expected gain is the same as the flow cost. Thus, p^* is defined by

$$\underbrace{p^*\lambda}_{\text{success rate}} \times \underbrace{\frac{m}{r}}_{\text{gain}} = \underbrace{c}_{\text{cost}}$$

and so

$$p^* = \frac{rc}{\lambda m} \tag{2}$$

and we will suppose that $0 < p^* < 1$.

We will assume that firm 1's information is accurate enough so that if it is the only firm, with signal b_1 it would not want to engage in R&D while with signal g_1 it would.

Assumption 1 *The quality of firm* 1*'s is such that*

$$p(b_1) < p^* < p(g_1)$$

Assumption 1 is made solely to allow a sharper statements of our results and to make the consideration of many trivial cases unnecessary. Without it, many of our results would involve only weak inequalities which would become strict if the condition above were to hold.

The following definition will prove useful in the subsequent analysis. Suppose both firms have a common belief at time 0 that the probability of state G is p_0 and with this belief both engage in R&D at time 0. As time elapses and both firms are active but neither firm has been successful, the firms become increasingly pessimistic that the state is G and the posterior probability that the state is G decreases. At time t, the common belief p_t is such that[11]

$$\frac{p_t}{1-p_t} = e^{-2\lambda t}\frac{p_0}{1-p_0}$$

[11] This is just Bayes' rule in terms of odds ratios: given any event $\mathcal{E}$, we have

$$\frac{\Pr[G \mid \mathcal{E}]}{\Pr[B \mid \mathcal{E}]} = \frac{\Pr[\mathcal{E} \mid G]}{\Pr[\mathcal{E} \mid B]} \times \frac{\Pr[G]}{\Pr[B]}$$

since, conditional on the state being G, the probability that neither firm has been successful until time t is $e^{-2\lambda t}$.

Definition 1 *If the initial belief $p_0 > p^*$, $T(p_0)$ is the time when, absent any success by either firm, this belief will decay to p^*, that is,*

$$e^{-2\lambda T(p_0)} \frac{p_0}{1-p_0} = \frac{p^*}{1-p^*} \tag{3}$$

If the initial belief $p_0 \leq p^*$, then $T(p_0) = 0$.

Equivalently, for $p_0 > p^*$,

$$T(p_0) = \frac{1}{2\lambda} \ln\left(\frac{p_0}{1-p_0}\right) - \frac{1}{2\lambda} \ln\left(\frac{p^*}{1-p^*}\right)$$

To save on notation, we will write

$$T(s_i) \equiv T(p(s_i)) \tag{4}$$

and

$$T(s_1, s_2) \equiv T(p(s_1, s_2)) \tag{5}$$

3 Single-firm benchmark

Before studying the situation in which the two firms are competing against one another, it is useful to consider the case where each firm acts in isolation. Comparing the situation in which firm 1 is alone to the situation in which firm 2 is alone, we obtain

Proposition 0 *The probability that firm 1 is successful when alone is greater than the probability that firm 2 is successful when alone.*

To establish the proposition, first note that if firm i gets a signal $s_i \in \{g_i, b_i\}$, then its belief that the state is G is $p(s_i)$ at time 0. If $p(s_i) \leq p^*$ then the firm should not engage in R&D at all since its expected profits from R&D are non-positive. But if $p(s_i) > p^*$ then it is worthwhile to engage in R&D at time 0 and continue to do so as long as its belief $p_t(s_i)$ at time t remains above p^*. In terms of odds ratios, this means that a solitary firm should remain active as long as

$$\frac{p_t(s_i)}{1-p_t(s_i)} = e^{-\lambda t} \frac{p(s_i)}{1-p(s_i)} > \frac{p^*}{1-p^*}$$

reflecting the fact that the probability that a single firm does not succeed until time t is just $e^{-\lambda t}$. The following result is immediate.

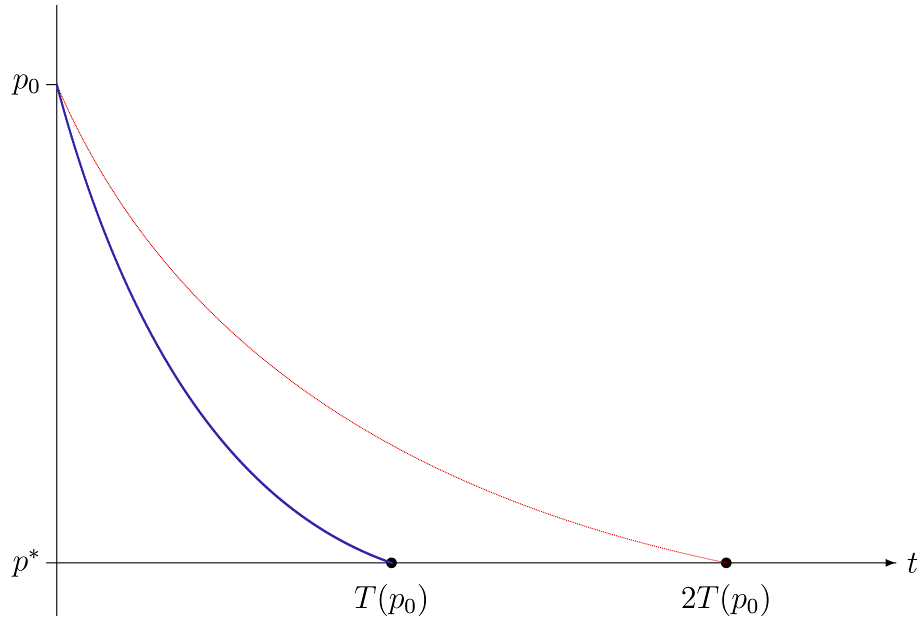

Figure 1: Belief Decay

When two firms are active, beliefs decay twice as fast (lower curve) as with one firm (upper curve).

Lemma 3.1 *A single firm with signal s_i should quit at the earliest time t such that $p_t(s_i) \leq p^*$.*

Proof. If firm i with signal s_i quits at time t_i, its flow profit is

$$r\int_0^{t_i} e^{-rt}\Pr\left[\mathcal{S}_0(t)\right]\left(p_t(s_i)\frac{\lambda m}{r}-c\right)dt = \lambda m\int_0^{t_i} e^{-rt}\Pr\left[\mathcal{S}_0(t)\right]\left(p_t(s_i)-p^*\right)dt$$

where $\Pr[\mathcal{S}_0(t)] = e^{-\lambda t}p(s_i)+1-p(s_i)$ is the probability that there has been no success until time t. Recall that $p^* = rc/\lambda m$. The result obviously follows. ■

The optimal quitting time for a firm with signal s_i is just $2T(s_i)$ since from the definition of T in (3) and (4),

$$e^{-2\lambda T(s_i)}\frac{p(s_i)}{1-p(s_i)} = \frac{p^*}{1-p^*} \tag{6}$$

Since the beliefs of a single firm decay at one-half the rate of decay with two firms—two failures constitute worse news than one failure—it takes twice as long to reach p^*, as depicted in Figure 1. Since $2T(s_i)$ is the single-firm optimal quitting time, using (6), the probability of success given the initial belief $p(s_i)$ is

$$p(s_i)\left(1-e^{-2\lambda T(s_i)}\right) = \frac{p(s_i)-p^*}{1-p^*}$$

Consider firm i when alone. If $p^* < p(b_i) < p(g_i)$, the firm would enter regardless of its signal and its ex ante probability of success is

$$\begin{aligned}\Pr[\mathcal{S}_i] &= \Pr[g_i]\frac{p(g_i)-p^*}{1-p^*} + \Pr[b_i]\frac{p(b_i)-p^*}{1-p^*} \\ &= \frac{\pi - p^*}{1-p^*}\end{aligned}$$

If $p(b_i) \leq p^* < p(g_i)$, the firm would enter only if its signal were g_i and now the ex ante probability of success is

$$\begin{aligned}\Pr[\mathcal{S}_i] &= \Pr[g_i]\frac{p(g_i)-p^*}{1-p^*} + \Pr[b_i]\times 0 \\ &= \frac{\pi q_i - (\pi q_i + (1-\pi)(1-q_i))p^*}{1-p^*} \\ &= \frac{\pi q_i(1-p^*) - (1-\pi)(1-q_i)p^*}{1-p^*}\end{aligned}$$

The proof of Proposition 0 is divided into two cases.

Case 1: $p(b_1) < p^* < p(b_2)$ Now firm 1 would enter only with a good signal whereas firm 2 would enter regardless of its signal. Thus,

$$\begin{aligned}\Pr[\mathcal{S}_1] - \Pr[\mathcal{S}_2] &= \frac{\pi q_1(1-p^*) - (1-\pi)(1-q_1)p^*}{1-p^*} - \frac{\pi - p^*}{1-p^*} \\ &= \frac{(1-\pi)q_1 p^* - \pi(1-q_1)(1-p^*)}{1-p^*} \\ &= (1-\pi)q_1\left(\frac{p^*}{1-p^*} - \frac{\pi(1-q_1)}{(1-\pi)q_1}\right) \\ &= (1-\pi)q_1\left(\frac{p^*}{1-p^*} - \frac{p(b_1)}{1-p(b_1)}\right) \\ &> 0\end{aligned}$$

where the last inequality follows from Assumption 1.

Case 2: $p(b_2) \leq p^* < p(g_2)$ In this case, both firms would enter only if they had good signals and some routine calculations show that the difference in success probabilities

$$\begin{aligned}\Pr[\mathcal{S}_1] - \Pr[\mathcal{S}_2] &= \frac{\pi(1-p^*) + (1-\pi)p^*}{1-p^*}(q_1 - q_2) \\ &> 0\end{aligned}$$

This completes the proof of Proposition 0. ■

4 Upstart equilibrium

In this section, we exhibit an equilibrium of the R&D game in which the established firm enters the race if and only if it receives a favorable signal whereas the startup enters the race regardless of its signal. In this equilibrium, the probability that the startup wins the race is *greater* than or equal to the probability that the established firm wins and is strictly greater whenever $p(b_1, g_2) > p^*$.

In the next section, we will show that when the established firm 1 is much better informed than the startup firm 2, this is the *unique* Nash equilibrium outcome.

Recall from (3) and (5) that if $p(s_1, s_2) = \Pr[G \mid s_1, s_2] > p^*$, the two-firm threshold time $T(s_1, s_2)$ is defined by

$$e^{-2\lambda T(s_1,s_2)} \frac{p(s_1, s_2)}{1 - p(s_1, s_2)} = \frac{p^*}{1 - p^*} \tag{7}$$

and if $p(s_1, s_2) \leq p^*$, then $T(s_1, s_2) = 0$.

The ranking of the posterior probabilities (see (1)) implies

$$0 = T(b_1, b_2) \leq T(b_1, g_2) \leq T(g_1, b_2) < T(g_1, g_2) \tag{8}$$

and the inequalities are strict unless both sides are 0.

Consider following "upstart outcome" depicted in Figure 2. When the signals are (b_1, b_2), firm 1 does not enter and firm 2 exits immediately upon entering and learning that 1 did not enter. When the signals are (g_1, g_2), both firms exit simultaneously at time $T(g_1, g_2)$. When the signals are (g_1, b_2), firm 2 exits at time $T(g_1, b_2)$ and upon learning this, firm 1 follows immediately. Finally, when the signals are (b_1, g_2), firm 1 does not enter and firm 2 stays until $2T(b_1, g_2)$.

In the first three cases the chance that firm 1 will win is the same as the chance that firm 2 will win. But in the last case, firm 1 does not enter and when $T(b_1, g_2) > 0$, firm 2 has a positive probability of winning. Thus, ex ante firm 2 has a greater chance of obtaining the patent than does firm 1—the startup is an upstart. We will first establish

Proposition 1 *There exists a perfect Bayesian equilibrium in which the less-informed firm* 2 *wins more often than the better-informed firm* 1*.*

Strategies A *strategy* for firm i is a pair of functions (τ_i, σ_i) where $\tau_i : \{g_i, b_i\} \to \mathbb{R}_+ \cup \{\infty\}$ and $\sigma_i : \{g_i, b_i\} \times \mathbb{R}_+ \to \mathbb{R}_+ \cup \{\infty\}$.

First, $\tau_i(s_i)$ is the time at which firm i with signal s_i decides to quit unilaterally—that is, if he or she has not received any information that the other firm has quit. If $\tau_i(s_i) = \infty$, this means that the firm decides to never quit unilaterally.

Second, $\sigma_i(s_i, t_j)$ is the time at which firm i with signal s_i quits after learning that the other firm quit at time t_j. Of course, $\sigma_i(s_i, t_j) \geq t_j$.

We have only defined pure strategies here as the equilibrium we construct below does not involve any randomization. When we show that the equilibrium outcome is unique, we will introduce and consider randomized strategies as well.

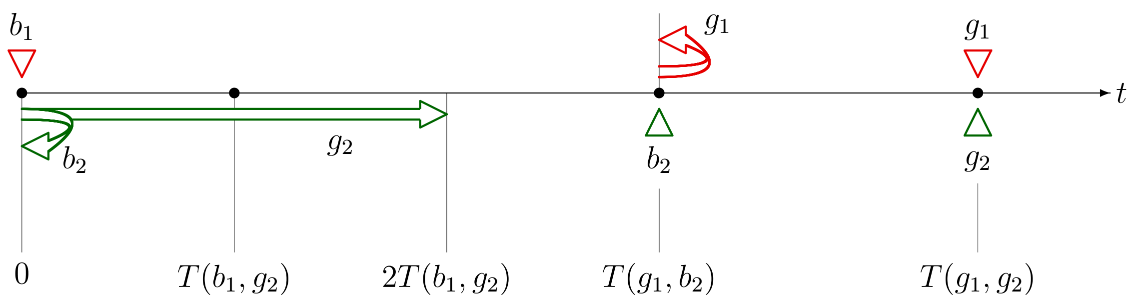

Figure 2: Upstart Equilibrium

Firm 1 (top) enters only with signal g_1 and then if firm 2 exits at $T(g_1, b_2)$, follows immediately, depicted as a U-turn. Otherwise, it stays until $T(g_1, g_2)$. Firm 2 (bottom) enters with either b_2 or g_2. If firm 1 exits at 0, firm 2 with signal b_2 follows immeditiately and with g_2, exits at $2T(b_1, g_2)$. Otherwise, firm 2 with b_2 exits at $T(g_1, b_2)$ and with g_2, exits at $T(g_1, g_2)$.

4.1 Equilibrium strategies

Consider the following strategies:

Firm 1 :

$$\tau_1^*(g_1) = T(g_1, g_2) \text{ and } \tau_1^*(b_1) = 0$$

$$\sigma_1^*(g_1, t_2) = \begin{cases} T(g_1, b_2) & \text{if } t_2 = T(g_1, b_2) \\ 2T(g_1, b_2) - t_2 & \text{if } t_2 < T(g_1, b_2) \\ 2T(g_1, g_2) - t_2 & \text{if } T(g_1, b_2) < t_2 < T(g_1, g_2) \end{cases}$$

with the following beliefs about its rival. If $t_2 \leq T(g_1, b_2)$, then firm 1 believes that $s_2 = b_2$ and otherwise believes $s_2 = g_2$.

Firm 2**:**

$$\tau_2^*(g_2) = T(g_1, g_2) \text{ and } \tau_2^*(b_2) = T(g_1, b_2)$$

$$\sigma_2^*(g_2, t_1) = \begin{cases} 2T(b_1, g_2) & \text{if } t_1 = 0 \\ 2T(g_1, g_2) - t_1 & \text{if } 0 < t_1 < T(g_1, g_2) \end{cases}$$

$$\sigma_2^*(b_2, t_1) = \begin{cases} 0 & \text{if } t_1 = 0 \\ 2T(g_1, b_2) - t_1 & \text{if } 0 < t_1 < T(g_1, b_2) \end{cases}$$

with the following beliefs about its rival. If $t_1 = 0$, then firm 2 believes that $s_1 = b_1$ and otherwise believes $s_1 = g_1$.

4.2 Verification of equilibrium

We now verify that the strategies (τ_i^*, σ_i^*) specified above constitute a perfect Bayesian equilibrium. To do this we will ascertain the optimal quitting time for the two firms in various situations. This quitting time will, as in Lemma 3.1, be determined by the condition that a firm's belief that the state is G is equal to p^*. But when another firm j is present, firm i not only knows its own signal s_i but may learn firm j's signal s_j in the course of play. Thus, it may be the case that even if based on its own signal alone, the belief is below p^*, the possibility of learning s_j in the future is a worthwhile investment. The following analog of Lemma 3.1 is derived under the condition that all such learning has already taken place. Thus we have

Lemma 4.1 *Let p_{it} denote i's belief at time t that the state is G.*
(i) If $p_{it} > p^*$, then i should not quit at t.
(ii) Suppose that at time t firm i believes with probability one that j's signal is s_j. If $p_{it} \leq p^*$, then firm i should quit at t.

Proof. The flow profit of firm i if it quits at time t_i is

$$r \int_0^{t_i} e^{-rt} \Pr[\mathcal{S}_0(t)] \left(p_{it} \frac{\lambda m}{r} - c \right) dt = \lambda m \int_0^{t_i} e^{-rt} \Pr[\mathcal{S}_0(t)] (p_{it} - p^*)\, dt$$

where p_{it} is firm i's belief at time t given all the information it has and $\Pr[\mathcal{S}_0(t)]$ is the probability that there has been no success until time t. This is the payoff because the chance that both firms will succeed at the same instant is zero. Note that firm j's quitting time t_j affects the instantaneous payoff only through its effect on i's belief p_{it}—before t_j the belief p_{it} declines rapidly since there are two unsuccessful firms whereas after j quits at time t_j the belief declines slowly since there is only one unsuccessful firm. ■

Firm 1 Suppose firm 2 follows the strategy (σ_2^*, τ_2^*) specified above.

Firm 1 **with signal** g_1 : We first argue that $\tau_1(g_1) < T(g_1, b_2)$ cannot be a best response. This is because $\tau_2^*(b_2) = T(g_1, b_2) < T(g_1, g_2) = \tau_2^*(g_2)$ and if g_1 exits before $T(g_1, b_2)$, it cannot learn 2's signal and the only information it has until then is g_1. But the posterior probability of G conditional on g_1 alone is $p(g_1) > p(g_1, b_2)$. And if there has been no success until $t < T(g_1, b_2)$, 1's belief $p_{1t} = p_t(g_1) > p_t(g_1, b_2)$ for $t < T(g_1, b_2)$. By Lemma 4.1, it is suboptimal to quit before $T(g_1, b_2)$.

On the other hand, if $\tau_1(g_1) \geq T(g_1, b_2)$, there are two possibilities. Given τ_2^*, either g_1 learns at time $T(g_1, b_2)$ that firm 2 quit and then infers that $s_2 = b_2$ or g_1 learns that firm 2 did not quit and then infers that $s_2 = g_2$. If g_1 learns that 2 quit, then it should also quit as soon as possible, that is, at $T(g_1, b_2)$ (Lemma 4.1 again). Thus, $\sigma_1^*(g_1, T(g_1, b_2))$ is optimal. If g_1 learns that 2 did not quit, then since

$\tau_2^*(g_2) = T(g_1, g_2)$, firm 1 should exit at $T(g_1, g_2)$ as well, that is, $\tau_1^*(g_1)$ is optimal. It is obvious that firm 1's beliefs about s_2 are consistent with firm 2's equilibrium behavior.

By the same reasoning, $\sigma_1^*(g_1, t_2)$ is optimal for all $t_2 \neq T(g_1, b_2)$ given 1's (off-equilibrium) beliefs.

Firm 1 **with signal** b_1 : We will argue that given (τ_2^*, σ_2^*) it is optimal for b_1 to not enter. Suppose $\tau_1(b_1) > 0$. Depending on its signal, firm 2 will quit at either $\tau_2^*(b_2) = T(g_1, b_2)$ or at $\tau_2^*(g_2) = T(g_1, g_2)$. First, $\tau_1(b_1) < T(g_1, b_2)$ is not optimal because 1 will not learn anything about 2's signal and $p(b_1) < p^*$ (Assumption 1). If firm 1 chooses $\tau_1(b_1) \geq T(g_1, b_2)$ and finds that 2 is still active, then it believes that firm 2's signal is g_2 and that the time when firm 2 will quit is $T(g_1, g_2)$. But now by Lemma 4.1 it is best to quit at $T(b_1, g_2) < T(g_1, b_2)$, the time when 1 can learn 2's signal. This means that the value of staying and learning 2's signal at $T(g_1, b_2)$ is negative. Thus, $\tau_1^*(b_1) = 0$ is optimal.

Firm 2 Now suppose firm 1 follows the strategy (σ_1^*, τ_1^*) specified above and consider

Firm 2 **with signal** g_2 : Since $\tau_1^*(g_1) = T(g_1, g_2)$ and $\tau_1^*(b_1) = 0$, if firm 2 enters, it will learn whether 1's signal is b_1 or g_1. If it learns that $s_1 = b_1$, then its optimal response is $\sigma_2^*(g_2, 0) = 2T(b_1, g_2)$. On the other hand, if it learns that $s_1 = g_1$, then by Lemma 4.1 firm 2 should quit at $\tau_2^*(g_2) = T(g_1, g_2)$.

By the same reasoning, $\sigma_2^*(g_2, t_1)$ is optimal for all $t_1 < T(g_1, g_2)$ given 2's (off-equilibrium) beliefs.

Firm 2 **with signal** b_2 : The same reasoning as in the case where firm 2's signal was g_2 shows that again 2's strategy is a best response.

This completes the proof of Proposition 1. ■

The particular choice of off-equilibrium beliefs does not affect the equilibrium outcome—any beliefs will do. Off-equilibrium beliefs could affect a firm's profit only if a deviation to stay longer than expected would cause its rival to drop out earlier. For instance, the equilibrium specifies that firm 1 with signal b_1 should not enter. If it did, then firm 2 would have to assign probability 1 to $s_1 = g_1$, since this is the only belief consistent with the equilibrium path. Thus, by entering b_1 cannot get firm 2 to exit early. In the upstart equilibrium outcome, all such events occur on the equilibrium path.

4.3 Equilibrium payoffs

The expected flow profits of firm 1 in the upstart equilibrium are

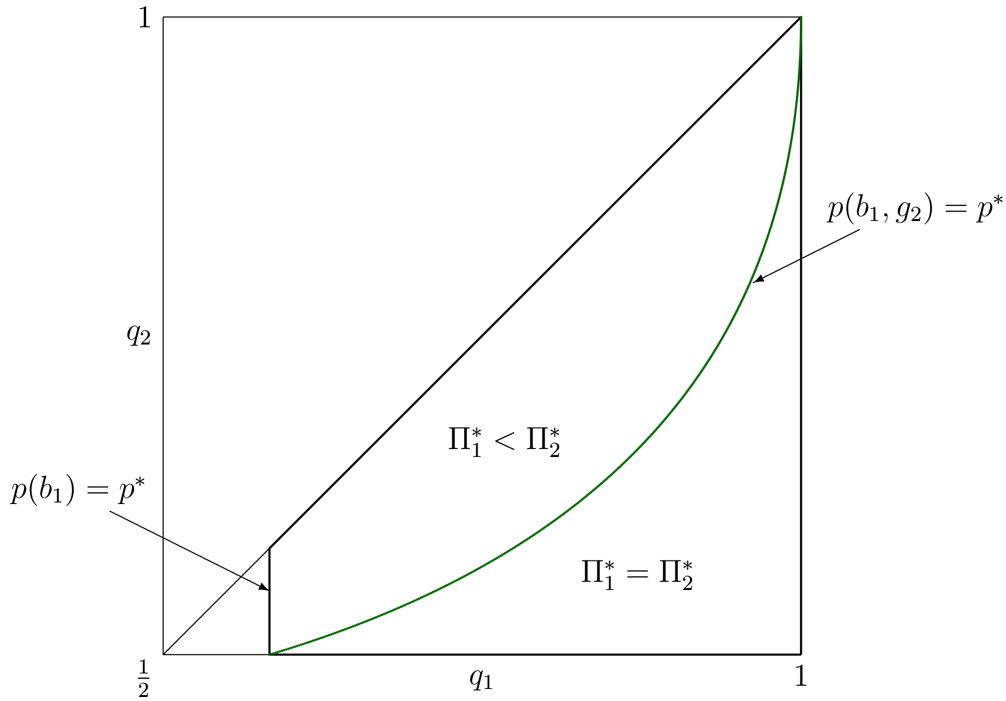

Figure 3: Upstart Equilibrium Payoffs

An upstart equilibrium exists in the quadrilateral region enclosed by dark lines. Above the curve, firm 2's equilibrium payoff (and winning probability) is strictly greater than that of firm 1. Below the curve, they are equal.

$$\Pi_1^* = \Pr[g_1, g_2] \times v(p(g_1, g_2)) + \Pr[g_1, b_2] \times v(p(g_1, b_2)) \quad (9)$$

where $v(p_0)$ is the flow payoff to a firm when both firms have a common belief p_0 at time 0 that the state is G (see Appendix A). The expression for the equilibrium payoff results from the fact that when the signals are (g_1, g_2) or (g_1, b_2), these become commonly known in the course of play of the upstart equilibrium. When the signals are (b_1, g_2) or (b_1, b_2), firm 1 does not enter the race and so its payoff is zero.

The expected profits of firm 2 in the upstart equilibrium are

$$\begin{aligned}\Pi_2^* \quad = \quad & \Pr[g_1, g_2] \times v(p(g_1, g_2)) + \Pr[g_1, b_2] \times v(p(g_1, b_2)) \\ & + \Pr[b_1, g_2] \times u(p(b_1, g_2))\end{aligned} \quad (10)$$

where $u(p_0)$ is the flow payoff to a firm when it is alone with belief p_0 at time 0 (see Appendix A again). The additional term appears because when the signals are (b_1, g_2), firm 1 does not enter, firm 2 thus learns at time 0 that 1's signal is b_1 and stays until $2T(b_1, g_2)$, the optimal quitting time for a single firm with initial belief $p(b_1, g_2)$ (see (6)). When the signals are (b_1, b_2), firm 1 does not enter and as soon as firm 2 learns this, it exits as well. Note that if $p(b_1, g_2) \leq p^*$, $u(p(b_1, g_2)) = 0$.

As long as $p(b_1, g_2) > p^*$,

$$\Pi_2^* > \Pi_1^*$$

These facts are depicted in Figure 3 and we summarize these findings as,

Corollary 1 *In the upstart equilibrium, the less-informed firm 2's payoff is greater than the better-informed firm 1's payoff.*

5 Uniqueness

We now show that when the informational advantage of firm 1 is large, that is, fixing all other parameters, q_2 is small relative to q_1, then the upstart equilibrium outcome is the unique Nash equilibrium outcome.

Proposition 2 *When the established firm's informational advantage is large, there is a unique Nash equilibrium outcome. Precisely, for every q_1 there exists a $\overline{q}_2$ such that for all $q_2 < \overline{q}_2$, there is a unique Nash equilibrium outcome.*

The proof of the Proposition is in two steps. First, we show that iterated elimination of dominated strategies (IEDS) results in a *single* outcome. Here we will use one round of elimination of *weakly* dominated strategies, followed by multiple (actually six more!) rounds of (iteratively) *weakly/strictly* dominated strategies.[12] The resulting outcome will be the same as in (τ^*, σ^*). As a final step, we will show that there cannot be any other Nash equilibrium outcome—the weakly dominated strategies that were eliminated cannot be part of any Nash equilibrium.

5.1 Step 1

Denote by Γ the original game and by $\Gamma(n)$ the game after n rounds of elimination. In what follows, Lemma 4.1 will invoked repeatedly in the following manner: if the two signals are known to be (s_1, s_2), then a firm that exits at $t < T(s_1, s_2)$ would leave some money on the table since that firm's belief time t, $p_{it} > p^*$.

IEDS Round 1

Claim 1 (a) *Any strategy of firm 1 such that $\tau_1(g_1) < T(g_1, b_2)$ is **weakly** dominated in Γ.*

[12] The weakly dominated strategies that we eliminate are so only because of histories that never occur. Below we will show that the upstart outcome is also the result of iterated elimination of *conditionally dominated* strategies.

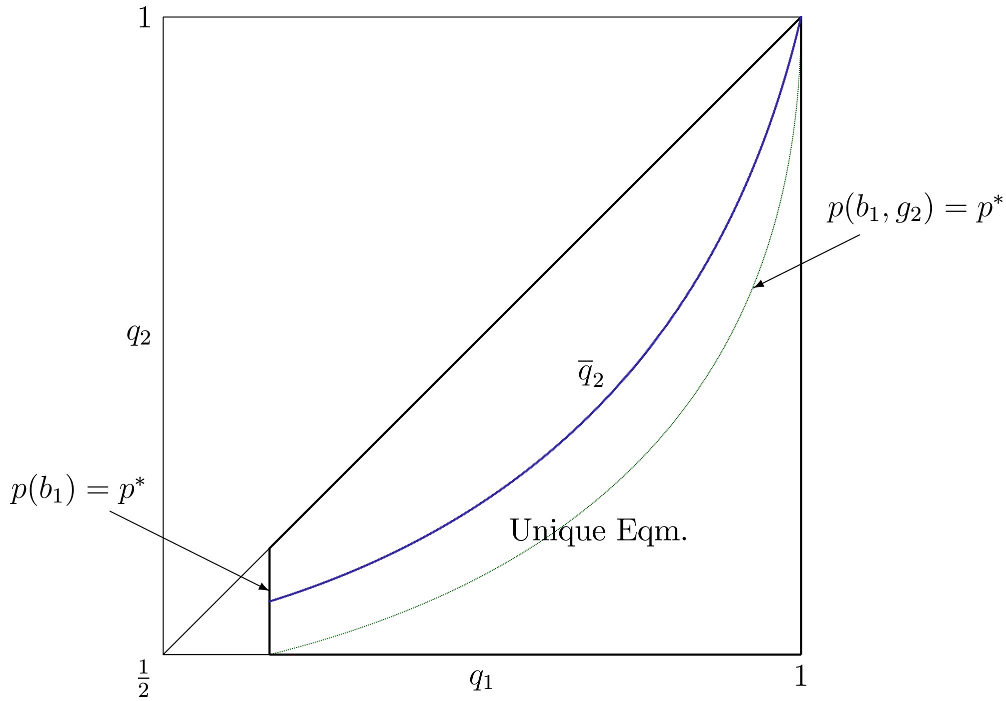

Figure 4: Uniqueness

There is a unique equilibrium outcome below the upper curve. Between the two curves, the startup has strictly higher payoffs in the unique equilibrium outcome.

Proof. Quitting at $\tau_1(g_1) < T(g_1, b_2)$ is weakly dominated by quitting at $\overline{\tau}_1(g_1) = T(g_1, b_2)$. First, if $\max(\tau_2(b_2), \tau_2(g_2)) \geq T(g_1, b_2)$ then quitting at $\tau_1(g_1) < T(g_1, b_2)$ is strictly worse for g_1 than quitting at $T(g_1, b_2)$. If $\max(\tau_2(b_2), \tau_2(g_2)) < T(g_1, b_2)$, then quitting at $\tau_1(g_1) < T(g_1, b_2)$ is strictly worse than quitting at $T(g_1, b_2)$ if $\tau_1(g_1) < \max(\tau_2(b_2), \tau_2(g_2))$ and is equivalent if $\max(\tau_2(b_2), \tau_2(g_2)) < \tau_1(g_1)$. ■

Claim 1 (b) *Any strategy of firm 2 such that $\tau_2(g_2) < T(b_1, g_2)$ is **weakly** dominated in Γ.*

Proof. The proof is the same as in the previous claim with the identities the firms interchanged. ■

It is important to note that in this round the strategies eliminated are not strictly dominated. The reason is that a strategy (τ_1, σ_1) that calls on firm 1 with signal g_1 to quit at a time such that $0 < \tau_1(g_1) < T(g_1, b_2)$ is not strictly worse than quitting at $T(g_1, b_2)$ against a strategy (τ_2, σ_2) such that $\tau_2(b_2) = 0 = \tau_2(g_2)$. Since both types of firm 2 quit at time 0, the choice of $\tau_1(g_1)$ is irrelevant. More generally, such a $\tau_1(g_1)$ is not strictly worse than $T(g_1, b_2)$ against any strategy (τ_2, σ_2) such that $\max(\tau_2(b_2), \tau_2(g_2)) < \tau_1(g_1)$.

IEDS Round 2

Claim 2 *Any strategy of firm* 1 *such that* $\tau_1(b_1) > T(b_1, g_2)$ *is strictly dominated in* $\Gamma(1)$.

Proof. If firm 2's signal is g_2, then from Claim 1 (b) $\tau_2(g_2) \geq T(b_1, g_2)$. In this case, for firm 1 to choose $\tau_1(b_1) > T(b_1, g_2)$ is strictly worse than $\tau_1(b_1) = T(b_1, g_2)$. On the other hand, if firm 2's signal is b_2, then for firm 1 to choose $\tau_1(b_1) > T(b_1, g_2)$ is no better than $\tau_1(b_1) = T(b_1, g_2)$. Thus, the expected payoff from $\tau_1(b_1) > T(b_1, g_2)$ is strictly lower than the expected payoff from quitting at $T(b_1, g_2)$. ■

IEDS Round 3

Claim 3 *Given all other parameters, there exists a* $\overline{q}_2$ *such that for all* $q_2 < \overline{q}_2$, *any strategy of firm 2 such that (i) if* $T(g_1, b_2) > 0$, *then* $\tau_2(b_2) < T(b_1, g_2)$ *is strictly dominated in* $\Gamma(2)$; *and (ii) if* $T(g_1, b_2) = 0$, *then* $\tau_2(b_2) > 0$ *is strictly dominated in* $\Gamma(2)$.

Proof. (i) Claim 1 (a) and Claim 2 imply that $\tau_1(b_1) \leq T(b_1, g_2) < T(g_1, b_2) \leq \tau_1(g_1)$. This means that firm 2 can learn firm 1's signal by staying until $\tau_1(b_1)$.

We will now argue that (τ_2, σ_2) is strictly dominated by $(\overline{\tau}_2, \overline{\sigma}_2)$ such that $\overline{\tau}_2(b_2) = T(g_1, b_2)$ and $\overline{\sigma}_2(b_2, t_2) = t_2$ for all $t_2 \leq T(b_1, g_2)$. Since b_1 will exit no later than $T(b_1, g_2)$, firm 2's flow profit from the strategy $(\overline{\tau}_2, \overline{\sigma}_2)$ when evaluated at any time $T \leq T(b_1, g_2)$ is at least

$$\lambda m \int_T^{T(b_1,g_2)} e^{-r(t-T)} \Pr[\mathcal{S}_0(t)] (p_{2t} - p^*) \, dt$$
$$+ \Pr[g_1 \mid b_2, \mathcal{S}_0(T)] \times \lambda m \int_{T(b_1,g_2)}^{T(g_1,b_2)} e^{-r(t-T)} \Pr[\mathcal{S}_0(t)] (p_{2t} - p^*) \, dt$$

where $\mathcal{S}_0(t)$ is the event that neither firm has succeed until t and firm 2's belief at time t that the state is G is

$$\frac{p_{2t}}{1 - p_{2t}} = \begin{cases} e^{-2\lambda t} \frac{p(b_2)}{1-p(b_2)} & \text{if } t \leq T(b_1, g_2) \\ e^{-2\lambda t} \frac{p(g_1,b_2)}{1-p(g_1,b_2)} & \text{if } t > T(b_1, g_2) \end{cases}$$

Before time $T(b_1, g_2)$, firm 2 cannot learn 1's signal and so its belief p_{2t} results only from its own signal b_2. At time $T(b_1, g_2)$ it learns 1's signal and exits immediately if $s_1 = b_1$. But if firm 1 does not exit at $T(b_1, g_2)$, then firm 2 knows that $s_1 = g_1$ and its belief now results from both its own signal b_2 and firm 1's signal g_1.

Notice that while the first term in the expression for firm 2's payoff above may be negative, the second is surely positive. For q_2 small enough so that $p(b_1, g_2) \leq p^*$, or equivalently, $T(b_1, g_2) = 0$, the first term is zero while the second is strictly positive when $p(g_1, b_2) > p^*$ and so $T(g_1, b_2) > 0$ as well. Thus, there exists a $\overline{q}_2$ such that for all $q_2 < \overline{q}_2$, the payoff from $(\overline{\tau}_2, \overline{\sigma}_2)$ is greater than the payoff from any strategy such that $\tau_2(b_2) < T(b_1, g_2)$.

If $s_1 = b_1$, then 2 is indifferent at all $\tau_2(b_2) > T(b_1, g_2)$. But if $s_1 = g_1$, $\overline{\tau}_2(b_2) = T(g_1, b_2)$ is strictly better than $\tau_2(b_2) < T(b_1, g_2)$. Since the latter occurs with positive probability, $(\overline{\tau}_2, \overline{\sigma}_2)$ is *strictly* better.

(ii) Obvious since in this case $p(g_1, b_2) \leq p^*$. ■

In the rest of the proof, we will assume that $q_2 < \overline{q}_2$. Note that $\overline{q}_2$ depends on the other parameters, in particular on q_1.

IEDS Round 4

Claim 4 *Any strategy of firm 1 such that $\tau_1(b_1) > 0$, is strictly dominated in $\Gamma(3)$.*

Proof. From Claim 1 (b) and Claim 3 we know that firm 2, regardless of its signal, will not be the first to quit before $T(b_1, g_2)$. This means that firm 1 will learn nothing from firm 2 prior to $T(b_1, g_2)$. This implies that if firm 1 with signal b_1 enters and exits before $T(b_1, g_2)$, its payoff is negative (recall that it is not worthwhile for firm 1 to enter just with his own signal b_1). If firm 1 enters, stays until $T(b_1, g_2)$ or longer, the best event is that it learns that firm 2's signal is g_2 at exactly time $T(b_1, g_2)$, the earliest time that he could learn anything about firm 2's signal. But even in this case, it is best to exit immediately after learning firm 2's signal. Thus, even if firm 1 were to learn that firm 2's signal was g_2, it cannot make any use of this information. Then, as before, his payoff from entering is negative. ■

IEDS Round 5

Claim 5 (a) *Any strategy of firm 2 such that $\sigma_2(g_2, 0) \neq 2T(b_1, g_2)$ is strictly dominated in $\Gamma(4)$.*

Proof. Given all previous rounds, we know that firm 1 will enter with g_1 and not with b_1. Thus, if firm 2 sees at time 0 that firm 1 did not enter, it knows that 1's signal was b_1. If firm 2's signal is g_2, it is strictly dominated to quit at a time other than $2T(b_1, g_2)$. ■

Claim 5 (b) *Any strategy of firm 2 such that $\sigma_2(b_2, 0) \neq 0$ is strictly dominated in $\Gamma(4)$.*

Proof. Given all previous rounds, we know that firm 1 will enter with g_1 and not with b_1. Thus, if firm 2 sees at time 0 that firm 1 did not enter, it knows that 1's signal was b_1. Clearly, given that 2's own signal is b_2, staying is strictly dominated. ■

Claim 5 (c) *Any strategy of firm* 2 *such that* $\tau_2(b_2) \neq T(g_1, b_2)$ *is strictly dominated in* $\Gamma(4)$.

Proof. Given all previous rounds, we know that firm 1 will enter with g_1 and not with b_1. Thus, if firm 2 sees that firm 1 entered, it knows that 1's signal is g_1. From Claim 1(a), firm 1 will stay at least until $T(g_1, b_2)$. For firm 2 to quit at a time other than $T(g_1, b_2)$ is strictly dominated. ■

Claim 5 (d) *Any strategy of firm* 2 *such that* $\tau_2(g_2) < T(g_1, g_2)$ *is* ***weakly*** *dominated in* $\Gamma(4)$.

Proof. Given all previous rounds, we know that firm 1 will enter with g_1 and not with b_1. Thus, if firm 2 sees that firm 1 entered, it knows that 1's signal is g_1. If $\tau_1(g_1) \geq T(g_1, g_2)$, then $\tau_2(g_2) < T(g_1, g_2)$ is strictly worse than quitting at $T(g_1, g_2)$. If $\tau_1(g_1) < T(g_1, g_2)$, then all quitting times $\tau_2(g_2)$ such that $\tau_1(g_1) < \tau_2(g_2)$ result in the *same* payoff as quitting at $T(g_1, g_2)$. If $\tau_1(g_1) < T(g_1, g_2)$, then all quitting times $\tau_2(g_2)$ such that $\tau_2(g_2) < \tau_1(g_1)$ results in a payoff strictly worse than from quitting at $T(g_1, g_2)$. ■

Note that for the same reasons as in Round 1, the strategies eliminated in Claim 5 (d) are also only weakly dominated.

IEDS Round 6

Claim 6 (a) *Any strategy of firm* 1 *such that* $\sigma_1(g_1, T(g_1, b_2)) \neq T(g_1, b_2)$ *is strictly dominated in* $\Gamma(5)$

Proof. Given all previous rounds, $\tau_2(b_2) = T(g_1, b_2) < T(g_1, g_2) \leq \tau_2(g_2)$ (Claim 5 (c) and Claim 5 (d)). So if firm 2 quits at $T(g_1, b_2)$, firm 1 knows that 2's signal is b_2. Then it is dominated for firm 1 to continue after $T(g_1, b_2)$. ■

Claim 6 (b) *Any strategy of firm* 1 *such that* $\tau_1(g_1) \neq T(g_1, g_2)$ *is strictly dominated in* $\Gamma(5)$.

Proof. As in the proof of the previous claim, if firm 2 does not quit at $T(g_1, b_2)$, firm 1 knows that 2's signal is g_2. From Claim 5(d), $\tau_2(g_2) \geq T(g_1, g_2)$. Thus, it is dominated for firm 1 to quit at any other time. ■

IEDS Round 7

Claim 7 *Any strategy of firm* 2 *such that* $\tau_2(g_2) > T(g_1, g_2)$ *is strictly dominated in* $\Gamma(6)$.

Proof. If firm 2 with signal g_2 sees that firm 1 entered, it knows that 1's signal is g_1. From Claim 6 (b), thus firm 1 will quit at $T(g_1, g_2)$ and so firm 2 should also quit at that time. ■

5.2 Step 2

The iterated elimination of dominated strategies, weak and strict, carried out above leaves a single *outcome*—the same as that in the upstart equilibrium (τ_i^*, σ_i^*). We now argue that this outcome is the *unique* Nash equilibrium outcome in Γ.

Suppose that $(\widetilde{\tau}, \widetilde{\sigma})$ is a (possibly mixed) Nash equilibrium where $\widetilde{\tau}_i(s_i)$ is a random variable on $[0, \infty)$ and so is $\widetilde{\sigma}_i(s_i, t_j)$. It is clear that there is no point in randomizing once the other player has exited. Thus, we can write $(\widetilde{\tau}, \sigma)$ where σ is pure.

Claim 8 *If* $(\widetilde{\tau}, \sigma)$ *is a Nash equilibrium, then* $\Pr[\widetilde{\tau}_2(g_2) < T(b_1, g_2)] = 0$.

Proof. Suppose to the contrary that $\Pr[\widetilde{\tau}_2(g_2) < T(b_1, g_2)] > 0$. We will sub-divide this event into three cases.

Case 1: $\Pr[\widetilde{\tau}_1(g_1) \leq \widetilde{\tau}_2(g_2) < T(b_1, g_2)] > 0$.

In this case, with positive probability g_1 is the first to quit. But for g_1, quitting at any time $t_1 < T(b_1, g_2)$ is strictly worse than quitting at $T(b_1, g_2)$ in expectation. Note that if $s_2 = g_2$, then quitting at t_1 is strictly worse than quitting at $T(b_1, g_2)$. This is because at any time $t < T(b_1, g_2) < T(g_1, b_2)$, the belief of g_1 is such that $p_{1t} > p^*$ (using Lemma 4.1). On the other hand, if $s_2 = b_2$, it is no better.

Case 2: $\Pr[\widetilde{\tau}_2(g_2) < \widetilde{\tau}_1(g_1) < T(b_1, g_2)] > 0$.

In this case, for g_2, quitting at any time $t_2 < T(b_1, g_2)$ is strictly worse than quitting at $T(b_1, g_2)$ in expectation.

Case 3: $\Pr[\widetilde{\tau}_2(g_2) < T(b_1, g_2) \leq \widetilde{\tau}_1(g_1)] > 0$.

Again, for g_2, quitting at any time $t_2 < T(b_1, g_2)$ is strictly worse than quitting at $T(b_1, g_2)$ in expectation.

Thus, we have argued that $(\widetilde{\tau}, \sigma)$ is not a Nash equilibrium. ■

Claim 9 *If* $(\widetilde{\tau}, \sigma)$ *is a Nash equilibrium, then* $\Pr[\widetilde{\tau}_1(g_1) < T(g_1, b_2)] = 0$.

Proof. Suppose to the contrary that $\Pr[\widetilde{\tau}_1(g_1) < T(g_1, b_2)] > 0$. Again we will sub-divide this event into three cases.

Case 1: $\Pr[\widetilde{\tau}_1(g_1) \leq T(b_1, g_2)] > 0$.

In this case, with positive probability g_1 is the first to quit since by Claim 8, g_2 never quits before $T(b_1, g_2)$. But for g_1 to quit at a time $t_1 < T(g_1, b_2)$ is strictly worse than quitting at $T(g_1, b_2)$ in expectation. This is because if $s_2 = g_2$, this is strictly worse because $\Pr[\widetilde{\tau}_2(g_2) \geq T(b_1, g_2)] = 1$ (Claim 8) and if $s_2 = b_2$, it is no better. Thus, $\Pr[\widetilde{\tau}_1(g_1) \leq T(b_1, g_2)] = 0$.

Case 2: $\Pr[\widetilde{\tau}_2(g_2) < \widetilde{\tau}_1(g_1)$ and $T(b_1, g_2) < \widetilde{\tau}_1(g_1) < T(g_1, b_2)] > 0$.

First, note that $\Pr[\widetilde{\tau}_1(b_1) > T(b_1, g_2)] = 0$ as well. This is because from Claim 8, $\Pr[\widetilde{\tau}_2(g_2) \geq T(b_1, g_2)] = 1$ and so when the signals are (b_1, g_2), for b_1 to stay beyond $T(b_1, g_2)$ is strictly worse than dropping out at $T(b_1, g_2)$. When the signals are (b_1, b_2), either dropping out at some $t_1 > T(b_1, g_2)$ is suboptimal because $t_2 \geq t_1$ or it does not matter because $t_2 < t_1$. Thus to drop out at any $t_1 > T(b_1, g_2)$ is suboptimal for b_1.

Now since $\Pr[\widetilde{\tau}_1(b_1) > T(b_1, g_2)] = 0$ and $\Pr[\widetilde{\tau}_1(g_1) \leq T(b_1, g_2)] = 0$ (Case 1), this means that if firm 1 does not quit by time $T(b_1, g_2)$, then firm 2 knows that $s_1 = g_1$. Then it is suboptimal for firm 2 with signal g_2 to drop out at $t_2 < T(g_1, b_2) < T(g_1, g_2)$. When the signals are (b_1, g_2), $t_2 \geq T(b_1, g_2)$ with probability 1 and $t_1 \leq T(b_1, g_2)$ with probability 0. Thus, firm 1 is the first to drop out and thus for g_2 to quit at any $t_2 \geq T(b_1, g_2)$ is irrelevant. Thus, overall firm 2's strategy is not a best response.

Case 3: $\Pr[\widetilde{\tau}_2(g_2) \geq \widetilde{\tau}_1(g_1) \text{ and } T(b_1, g_2) < \widetilde{\tau}_1(g_1) < T(g_1, b_2)] > 0$.

In this case, for g_1 to quit before $T(g_1, b_2)$ is strictly worse than quitting at $T(g_1, b_2)$ in expectation. This is because if $s_2 = g_2$, it is strictly worse and if $s_2 = b_2$ it is no better. ■

So far we have argued that if $(\widetilde{\tau}, \sigma)$ is a (possibly mixed) Nash equilibrium then almost every pure action τ in its support was not weakly dominated in Round 1 of the IEDS procedure. We complete the proof by showing that the same is true in Round 5.

Claim 10 *If $(\widetilde{\tau}, \sigma)$ is a Nash equilibrium, then* $\Pr[\widetilde{\tau}_2(g_2) < T(g_1, g_2)] = 0$.

Proof. Suppose to the contrary that $\Pr[\widetilde{\tau}_2(g_2) < T(g_1, g_2)] > 0$. Again, we will sub-divide this event into two cases.

Case 1: $\Pr[T(g_1, b_2) \leq \widetilde{\tau}_2(g_2) \leq \widetilde{\tau}_1(g_1) < T(g_1, g_2)] > 0$.

From Claim 9, $\Pr[\widetilde{\tau}_1(g_1) \geq T(g_1, b_2)] = 1$ and from Claim 4 $\Pr[\widetilde{\tau}_1(b_1) = 0] = 1$. This means that if firm 1 is active at any time $t > 0$, then with probability 1, firm 2 believes that $s_1 = g_1$. Thus, it is not optimal for g_2 to quit before $T(g_1, g_2)$.

Case 2: $\Pr[T(g_1, b_2) \leq \widetilde{\tau}_1(g_1) < \widetilde{\tau}_2(g_2) < T(g_1, g_2)] > 0$.

In this case, since Claim 8 implies $\Pr[\widetilde{\tau}_2(g_2) \geq T(b_1, g_2)] = 1$ and Claim 5 (c) implies $\Pr[\widetilde{\tau}_2(b_2) = T(g_1, b_2)] = 1$, at any time $t > T(g_1, b_2)$ firm 1 will believe with probability 1 that $s_2 = g_2$. Thus if $\Pr[\widetilde{\tau}_1(g_1) > T(g_1, b_2)] > 0$, then it is suboptimal for g_1 to quit before $T(g_1, g_2)$. If $\Pr[\widetilde{\tau}_1(g_1) = T(g_1, b_2)] = 0$, then it is better to stay a little longer and learn whether or not $s_2 = g_2$. ■

The last claim shows that if $(\widetilde{\tau}, \sigma)$ is a Nash equilibrium, the probability that a pure strategy in the support of $\widetilde{\tau}_2(g_2)$ is eliminated in Round 5 of the IEDS procedure is zero.

We have thus argued that no Nash equilibrium can have an outcome different from the one in (τ^*, σ^*).

This completes the proof of Proposition 2. ■

Multiplicity with near symmetry We have shown that when firm 1's informational advantage is large, there is a unique equilibrium outcome. When this advantage is small, however, there may be other equilibria as well. To see this, suppose that $q_1 - q_2$ is small. Now the argument for uniqueness no longer holds—in particular, the reasoning in Round 3 of the IEDS procedure fails. Indeed, when $q_1 - q_2$ is small enough, there exists an equilibrium which is "mirror image" of the upstart equilibrium with the roles of firms 1 and 2 interchanged.

In the *mirror equilibrium*, denoted by (τ^{**}, σ^{**}), firm 2 enters the race only if its signal is g_2. Specifically, $\tau_2^{**}(b_2) = 0$ while $\tau_2^{**}(g_2) = T(g_1, g_2)$. Moreover, $\sigma_2^{**}(g_2, T(b_1, g_2)) = T(b_1, g_2)$.

Firm 1 enters the race regardless of its signal and $\tau_1^{**}(b_1) = T(b_1, g_2)$ while $\tau_1^{**}(g_1) = T(g_1, g_2)$. Finally, $\sigma_1^{**}(b_1, 0) = 0$ and $\sigma_1^{**}(g_1, 0) = 2T(g_1, b_2)$.

Here we have not specified off-equilibrium behavior and beliefs but this can be done by mimicking the upstart equilibrium.

When $q_1 - q_2$ is small, $T(b_1, g_2)$ is close to $T(g_1, b_2)$. Moreover, the assumption that $p(b_1) < p^*$ (Assumption 1) implies that $p(b_2) < p^*$ as well. Now the arguments confirming that (τ^*, σ^*) is an equilibrium also confirm that (τ^{**}, σ^{**}) is also an equilibrium.

5.3 Conditional dominance

When q_2 is relatively small, the upstart outcome is not only the unique Nash equilibrium outcome but it is also the unique outcome remaining after iterated elimination of *conditionally dominated* strategies (Shimoji and Watson, 1998).

A strategy for a player is conditionally dominated, if there is an information set for that player that (i) can be reached by the player's own strategy; (ii) is *strictly* dominated by another strategy when measured against only those strategies of other players which can reach the given player's information set. In the iterative procedure carried out above, the strategies that were eliminated in Round 1 and Round 5 were weakly dominated but not strictly dominated. These strategies were, however, conditionally dominated. Thus, the equilibrium outcome we identify is also the only outcome that survives iterated elimination of conditionally dominated strategies.[13]

6 Planner's problem

How does the upstart equilibrium compare to the solution of a "planner" who seeks to maximize the joint expected profits of the two firms? To analyze such a planner's problem, suppose that the belief that the state is G is $p_0 > p^*$ at time 0.

[13]In general games, the iterated elimination of conditionally dominated strategies may leave outcomes that are not Nash equilibria. This is not true in the game considered here, of course.

Since exit is irrevocable and it is never optimal to continue once the belief falls below p^*, the planner's problem reduces to choosing a time S such that both firms are active until $S \leq T \equiv T(p_0)$ and then one of the firms exits. Since both firms engage in R&D until time S, the belief decays at the rate 2λ until S and then at the rate λ after that. Thus per-firm expected flow profit from switching from two firms to one firm at time S is

$$\begin{aligned} w(S) &= \lambda m \int_0^S e^{-rt} \left(e^{-2\lambda t} p_0 + 1 - p_0\right)(p_t - p^*)\, dt \\ &+ \frac{1}{2} \lambda m \int_S^{2T-S} e^{-rt} \left(e^{-\lambda(S+t)} p_0 + 1 - p_0\right)(p_t - p^*)\, dt \end{aligned}$$

where the belief p_t at time t that the state is G is defined by

$$\frac{p_t}{1-p_t} = \begin{cases} e^{-2\lambda t} \frac{p_0}{1-p_0} & \text{if } t \leq S \\ e^{-\lambda(S+t)} \frac{p_0}{1-p_0} & \text{if } t \geq S \end{cases} \tag{11}$$

reflecting the fact that both firms are active until time S and after that only one of the two firms is active. Note that $e^{-2\lambda t} p_0 + 1 - p_0$ is the probability that neither firm is successful until time t. Note also that $p_{2T-S} = p^*$ and that the coefficient $\frac{1}{2}$ in the second term appears because w represents per-firm flow profits and the profit of the firm that exits is 0. After substituting for p_t from (11), $w(S)$ can be explicitly calculated to be

$$\begin{aligned} w(S) &= \frac{\lambda m p_0 (1-p^*)}{2\lambda + r} \left((2\lambda + r) e^{-2T\lambda} \left(e^{-rS} - 1\right) - r\left(e^{-(2\lambda+r)S} - 1\right)\right) \\ &+ \frac{\lambda m p_0 (1-p^*)}{2(\lambda + r)} \left(\lambda e^{-2\lambda T} \left(e^{-r(2T-S)} - e^{-rS}\right) - r e^{-rS} \left(e^{-2\lambda T} - e^{-2\lambda S}\right)\right) \end{aligned}$$

Differentiating with respect to S then yields

$$w'(S) = \lambda m p_0 (1-p^*) \times \frac{r e^{rS} \left(r e^{-2(\lambda+r)S} + \lambda e^{-2(\lambda+r)T} - (\lambda + r) e^{-2(\lambda T + rS)}\right)}{2(\lambda + r)}$$

and note that $w'(T) = 0$. Differentiating again we obtain

$$\begin{aligned} & w''(S) \\ = \;& \lambda m p_0 (1-p^*) \times \frac{r^2 e^{rS} \left(\lambda e^{-2(\lambda+r)T} + (\lambda + r) e^{-2(\lambda T + rS)} - (2\lambda + r) e^{-2(\lambda+r)S}\right)}{2(\lambda + r)} \\ < \;& \lambda m p_0 (1-p^*) \times \frac{r^2 e^{rS} \left(\lambda e^{-2(\lambda+r)S} + (\lambda + r) e^{-2(\lambda S + rS)} - (2\lambda + r) e^{-2(\lambda+r)S}\right)}{2(\lambda + r)} \\ = \;& 0 \end{aligned}$$

whenever $S < T$. Thus, w is a concave function and $w'(T) = 0$. As a result, the joint profits of the firms are maximized when $S = T$, that is, when both firms are active until time T. Thus, we obtain

Proposition 3 *The joint profit-maximizing plan with any initial belief p_0 is for both firms to invest in R&D as long as it is profitable, that is, as long as the updated belief $p_t > p^*$.*

Relative to the planner's optimum, the upstart equilibrium results in strictly lower total profits when $p(b_1, g_2) > p^*$; otherwise, they are the same. It is worth noting, however, that the ex ante probability of R&D success in the upstart equilibrium versus the planner's optimum is always the *same*. To see this, recall that the only difference between the two possibly occurs when the signals are (b_1, g_2). In this case, firm 1 stays out while firm 2 invests until time $2T(b_1, g_2)$. Conditional on (b_1, g_2), the probability of success in equilibrium is then

$$p(b_1, g_2)\left(1 - e^{-2\lambda T(b_1,g_2)}\right)$$

and this is the same as that in the planner's solution. Notice that while the overall probability of success in equilibrium is the same as that for the planner, success arrives later in the former case. This is because in equilibrium, when the signals are (b_1, g_2) only one firm is investing in R&D. This causes "learning-from-failure" to slow down relative to the case when two firms invest, which is the planner's solution. If we interpret the planner's problem as arising from a merger of the two firms to form a monopoly and the equilibrium as arising from competition, then this says that a monopoly would reach R&D success faster than competition, perhaps echoing the sentiments expressed by Schumpeter (1942).

7 Value of information

In the upstart equilibrium, the startup firm 2 not only wins more often than firm 1, it also obtains a higher equilibrium payoff (Corollary 1). This suggests perhaps that firm 1, say, could be better off with less precise information. This is not the case, however. We show next that despite the fact that the equilibrium payoff of the less-informed firm is higher than that of the better-informed firm, the value of information for both firms is *positive*.[14]

Proposition 4 *Suppose $q_1 > q_2$. Then in the upstart equilibrium, firm 1's payoff is increasing in q_1 and firm 2's payoff is increasing in q_2.*

First, consider firm 1. Recall from (9), that

$$\Pi_1^* = \Pr[g_1, g_2] \times v(p(g_1, g_2)) + \Pr[g_1, b_2] \times v(p(g_1, b_2))$$

We will show that each of the terms in the expression above is increasing in q_1. Of course, if $p(g_1, b_2) \leq p^*$, then the second term is zero.

[14] Bassan et. al (2003) exhibit a simple example where in an otherwise symmetric game, the payoff of the uninformed player 2 is higher than that of the informed player 1. In that game, however, the value of information to player 1 is negative.

Lemma 7.1 $\Pr[g_1, b_2] \times v(p(g_1, b_2))$ *is increasing in* q_1.

Proof.

$$\begin{aligned}
&\frac{\partial\left(\Pr[g_1, b_2] \times v(p(g_1, b_2))\right)}{\partial q_1} \\
&= \frac{\partial \Pr[g_1, b_2]}{\partial q_1} v(p(g_1, b_2)) + \Pr[g_1, b_2] \frac{\partial p(g_1, b_2)}{\partial q_1} v'(p(g_1, b_2)) \\
&> \frac{\partial \Pr[g_1, b_2]}{\partial q_1} v(p(g_1, b_2)) + \Pr[g_1, b_2] \frac{\partial p(g_1, b_2)}{\partial q_1} \frac{v(p(g_1, b_2))}{p(g_1, b_2)}
\end{aligned}$$

since $v(p)$ is an increasing and convex function that is non-negative and strictly positive for $p > p^*$ and so $v'(p) > \frac{1}{p} v(p)$ (see Appendix A). Moreover, $\partial p(g_1, b_2) / \partial q_1 > 0$. The sign of the right-hand side of the inequality is the same as the sign of

$$\begin{aligned}
&\frac{\partial \Pr[g_1, b_2]}{\partial q_1} p(g_1, b_2) + \Pr[g_1, b_2] \frac{\partial p(g_1, b_2)}{\partial q_1} \\
&= \frac{\partial}{\partial q_1} \left(\Pr[g_1, b_2] p(g_1, b_2)\right) \\
&= \frac{\partial}{\partial q_1} \Pr[G, g_1, b_2] \\
&= \frac{\partial}{\partial q_1} \left(\pi q_1 (1 - q_2)\right) \\
&> 0
\end{aligned}$$

■

Lemma 7.2 $\Pr[g_1, g_2] v(p(g_1, g_2))$ *is increasing in* q_1.

Proof. The proof is the same as that of the previous lemma with b_2 replaced by g_2.

■

Lemmas 7.1 and 7.2 together imply that firm 1's equilibrium payoff Π_1^* is increasing in q_1.

Next, consider firm 2. From (10),

$$\begin{aligned}
\Pi_2^* \;=\;& \Pr[g_1, g_2] \times v(p(g_1, g_2)) + \Pr[g_1, b_2] \times v(p(g_1, b_2)) \\
&+ \Pr[b_1, g_2] \times u(p(b_1, g_2))
\end{aligned}$$

We will show that the sum of the first two terms is increasing in q_2 and the last term is increasing in q_2 as well.

Lemma 7.3 $\Pr[g_1, g_2] \times v(p(g_1, g_2)) + \Pr[g_1, b_2] \times v(p(g_1, b_2))$ *is increasing in* q_2.

Proof. Since $\Pr[g_1]$ is independent of q_2, it is sufficient to show that

$$\frac{\Pr[g_1, g_2]}{\Pr[g_1]} v(p(g_1, g_2)) + \frac{\Pr[g_1, b_2]}{\Pr[\acute{g}_1]} v(p(g_1, b_2))$$

is increasing in q_2.

Now if $q_2' > q_2$, then

$$p'(g_1, b_2) < p(g_1, b_2) < p(g_1, g_2) < p'(g_1, g_2)$$

where $p'(g_1, \cdot)$ denotes the posterior derived from q_2'. Moreover, the mean $p'(g_1, \cdot)$ is $p(g_1)$ and this is the same as the mean of $p(g_1, \cdot)$ (since the expectation of the posteriors is the prior). Thus, the distribution of $p'(g_1, \cdot)$ is a mean preserving spread of the distribution of $p(g_1, \cdot)$.

Since v is a convex function, the result now follows. ■

Corollary 2 *Suppose $q_1 > q_2$. Then in the upstart equilibrium, firm 1's payoff is increasing in q_2.*

Lemma 7.4 *If $p(b_1, g_2) > p^*$, then $\Pr[b_1, g_2] \times u(p(b_1, g_2))$ is increasing in q_2.*

Proof. In Appendix A it is also established that the single-firm profit function $u(p)$ is also increasing, convex and strictly positive if $p > p^*$ and equal to zero if $p \leq p^*$. Using similar arguments as in the case of firm 1, establishes the result. ■

This completes the proof of Proposition 4. ■

The fact that the value of information is positive for firm 1 does not conflict with the fact that its payoff is lower than that of firm 2. The first is a statement about the derivative of Π_1^* with respect to q_1. The second is a statement comparing the profit levels of the two firms.

7.1 Willful ignorance

Proposition 4 shows that firm 1 cannot increase its equilibrium payoff by decreasing the quality of its information while still remaining better informed than firm 2 (and assuming that the upstart equilibrium is played). Precisely, for all $q_2 < q_1' < q_1$,

$$\Pi_1^*(q_1', q_2) < \Pi_1^*(q_1, q_2)$$

where we have now explicitly indicated the dependence of the equilibrium profits on the qualities of the two firms' signals.

But could firm 1 benefit from a drastic decrease in the quality of its information—say, by replacing all its experienced researchers, who have a good idea of the feasibility of the innovation, with new PhDs, who have none—thus becoming the *less*-informed

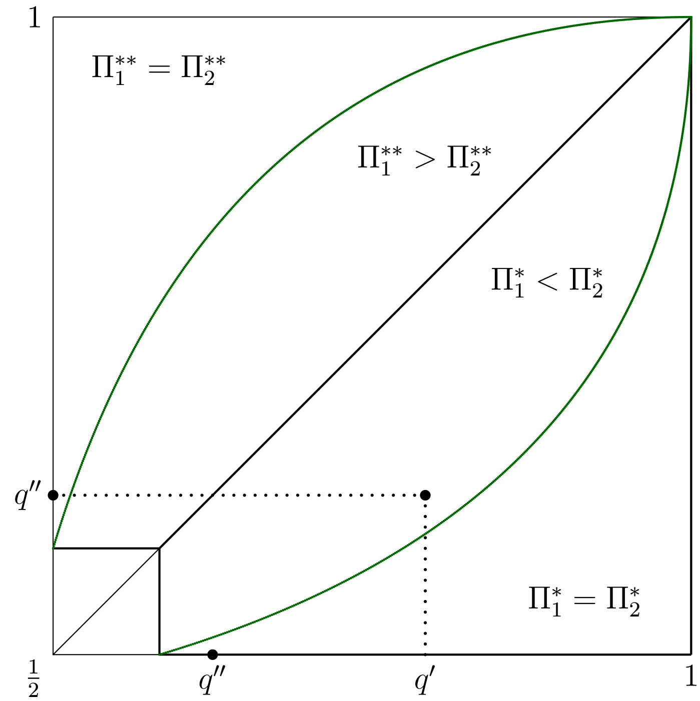

Figure 5: Willful Ignorance

Starting from $(q_1, q_2) = (q', q'')$ firm 1 is worse off by reducing its quality of information to $q_1 = \frac{1}{2}$.

firm? In terms of the model, suppose we start from a situation in which $(q_1, q_2) = (q', q'')$ where $\frac{1}{2} < q'' < q'$ and compare it to a situation in which $(q_1, q_2) = \left(\frac{1}{2}, q''\right)$ so that firm 1 is now less informed than firm 2. In this situation, there is again a unique equilibrium, but this time it is firm 1 which is the upstart.[15] This equilibrium is what we have called a "mirror equilibrium" (see the end of Section 5.2) since the roles of the firms have been reversed. If we denote payoffs in the mirror equilibrium by Π_i^{**}, by symmetry we have (see Figure 5).

$$\Pi_1^{**}\left(\tfrac{1}{2}, q''\right) = \Pi_2^{*}\left(q'', \tfrac{1}{2}\right)$$

But when the quality of firm 2's information is $\frac{1}{2}$, the upstart equilibrium outcome is unique and the expected profits of the two firms are the same, that is,

$$\Pi_2^{*}\left(q'', \tfrac{1}{2}\right) = \Pi_1^{*}\left(q'', \tfrac{1}{2}\right)$$

But in the region where the quality of firm 1's information is higher than that of firm 2, Π_1^* in increasing in *both* qualities (Proposition 4 and Corollary 2). Thus,

$$\Pi_1^{**}\left(\tfrac{1}{2}, q''\right) = \Pi_1^{*}\left(q'', \tfrac{1}{2}\right) < \Pi_1^{*}\left(q', q''\right)$$

[15] Any attempt to carry out this exercise when there are multiple equilibria is, of course, fraught with peril.

since $q' > q'' > \frac{1}{2}$. This means that it is not a good idea for the informationally advantaged but competitively disadvantaged firm 1 to become completely uninformed.

Of course, this argument applies not only to the case of complete ignorance, that is, $q_1 = \frac{1}{2}$. As long as, $q_1 > \frac{1}{2}$, is such that $p''(g_1, b_2) \leq p^*$ the same argument applies (here $p''(g_1, b_2) = \Pr[G \mid g_1, b_2]$ computed using qualities q_1 and $q_2 = q''$). This is because the argument above only relies on the equality, $\Pi_2^*(q_1, q'') = \Pi_1^*(q_1, q'')$.

The message of this subsection is: Don't fire the experienced researchers. Willful ignorance does not pay!

8 Many signals

So far we have assumed that each firm's information is binary—there are only two signals. In this section, we show that the main results are robust to the possibility that the firms' information is finer. Suppose that each of the two firms receives one of a finite number of signals, say, $S_1 = \{x^1, x^2, ..., x^K\}$ and $S_2 = \{y^1, y^2, ..., y^L\}$. As before, given the state, the signals are conditionally independent. We will assume that the signals can be ordered as $x^k < x^{k+1}$ and $y^l < y^{l+1}$ and that the monotone likelihood property is satisfied, that is,

$$\frac{\Pr[x^k \mid G]}{\Pr[x^k \mid B]} \text{ and } \frac{\Pr[y^l \mid G]}{\Pr[y^l \mid B]}$$

are strictly increasing in k and l, respectively. As in previous sections, we denote the posterior probabilities as

$$p(x^k) = \Pr[G \mid x^k] \text{ and } p(y^l) = \Pr[G \mid y^l]$$

and so we have that the posterior probabilities $p(x^k)$ and $p(y^l)$ are strictly increasing in k and l, respectively, as well.

We will use the following terminology to describe firm 2's signals.

Definition 2 *A signal y^l is said to be* optimistic *if*

$$\frac{\Pr[y^l \mid G]}{\Pr[y^l \mid B]} > 1$$

and pessimistic *if* $\Pr[y^l \mid G] / \Pr[y^l \mid B] < 1$.

The monotone likelihood ratio property implies that low signals are pessimistic and high signals optimistic.

In what follows, the following definition will be useful.

Definition 3 *The* quality bound *on the information content of firm 2's signals is*

$$Q_2 = \min\left\{\frac{\Pr[y^1 \mid B]}{\Pr[y^1 \mid G]}, \frac{\Pr[y^L \mid G]}{\Pr[y^L \mid B]}\right\}$$

Note that since

$$\frac{\Pr\left[y^1 \mid G\right]}{\Pr\left[y^1 \mid B\right]} < 1 < \frac{\Pr\left[y^L \mid G\right]}{\Pr\left[y^L \mid B\right]}$$

it is the case $Q_2 > 1$. To see why this is a measure of information quality, observe that if the quality bound Q_2 is close to 1, then for all l, $p\left(y^l\right) = \Pr\left[G \mid y^l\right]$ is close to π, the prior probability—all the posteriors are close to the prior—and so firm 2's signals are rather uninformative. Also, note that if there were only two signals and $\Pr\left[g_2 \mid G\right] = \Pr\left[b_2 \mid B\right] = q_2$, then $Q_2 = q_2/\left(1 - q_2\right)$.

Now observe that since

$$\frac{p\left(x^k, y^l\right)}{1 - p\left(x^k, y^l\right)} = \frac{p(x^k)}{1 - p\left(x^k\right)} \times \frac{\Pr\left[y^l \mid G\right]}{\Pr\left[y^l \mid B\right]}$$

and we have assumed

$$\frac{p\left(x^k\right)}{1 - p\left(x^k\right)} < \frac{p\left(x^{k+1}\right)}{1 - p\left(x^{k+1}\right)}$$

when the quality bound on firm 2's information, Q_2, is close enough to 1, we have

$$p\left(x^k, y^1\right) < ... < p\left(x^k, y^L\right) < p\left(x^{k+1}, y^1\right) < ... < p\left(x^{k+1}, y^L\right) \tag{12}$$

In other words, firm 2's signals are so poor that they cannot reverse the ranking of posteriors based on firm 1's information alone.

Finally, analogous to Assumption 1 in Section 2, we will assume that

$$p\left(x^1\right) < p^* < p\left(x^K\right)$$

that is, firm 1's signals are accurate enough so that, when alone, sometimes it wants to enter and sometimes not.

8.1 Upstart equilibrium

We now demonstrate that, as in Section 4, that there is a perfect Bayesian equilibrium of the R&D race in which firm 2 wins more often than firm 1.

Consider the following strategies. For firm 1,

$$\tau_1^*\left(x^k\right) = \begin{cases} T\left(x^k\right) & \text{if } k < K \\ T\left(x^K, y^L\right) & \text{if } k = K \end{cases}$$

$$\sigma_1^*\left(x^K, T\left(x^K, y^l\right)\right) = T\left(x^K, y^l\right)$$

and

$$\sigma_1^*\left(x^k, t_2\right) = 2T\left(x^k, y^L\right) - t_2 \text{ if } t_2 \neq T\left(x^K, y^l\right)$$

with the off-equilibrium beliefs that if firm 2 exits at a $t_2 \neq T\left(x^K, y^l\right)$, then its signal is y^L.

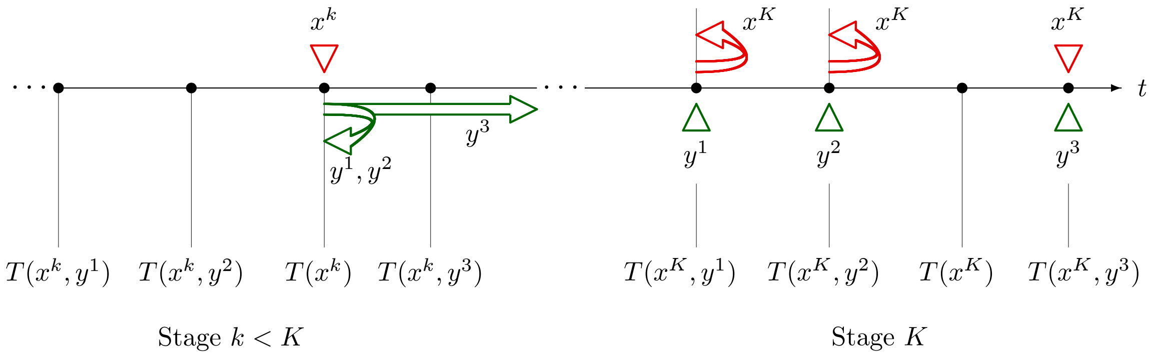

Figure 6: Upstart Equilibrium with Many Signals

Here $L = 3$. There are K stages. In stage $k < K$, firm 1 with signal x^k exits at $T(x^k)$ based only on its own information. Firm 2 follows immediately if it is pessimistic (signal y^1 or y^2) and optimally stays if it is optimistic (signal y^3). The pattern repeats until $k = K - 1$. Firm 2's signal is revealed only in stage K.

For firm 2,

$$\tau_2^*\left(y^l\right) = T\left(x^K, y^l\right)$$
$$\sigma_2^*\left(y^l, T\left(x^k\right)\right) = \max\left(T\left(x^k\right), 2T\left(x^k, y^l\right) - T\left(x^k\right)\right)$$

and

$$\sigma_2^*\left(y^l, t_1\right) = 2T\left(x^K, y^l\right) - t_1$$

with the off-equilibrium beliefs that if firm 1 exits at a $t_1 \neq T\left(x^k, y^L\right)$, then its signal is x^K.

Figure 6 depicts such an equilibrium when $L = 3$. Notice that there are K stages and in stage $k < K$, firm 2 with any signal can learn whether firm 1's signal is x^k or higher. Firm 2's information is revealed only in stage K and so only firm 1 with highest signal, x^K, can learn firm 2's signal. The learning is severely unbalanced. In equilibrium, firm 1 never suffers from the "survivor's curse"—it never regrets staying too long—but may suffer from the "quitter's curse"—it may regret exiting early. Firm 2, on the other hand, never suffers from the quitter's curse but may suffer from the survivor's curse when its signal is pessimistic.

We then have

Proposition 1 (M) *There exists a $Q_2^* > 1$ such that if $1 < Q_2 < Q_2^*$, then the strategies (σ^*, τ^*) constitute a perfect Bayesian equilibrium.*

Proof. First, suppose that $Q_2 > 1$ is small enough so that (12) holds. This implies that the same ranking holds for $T\left(x^k, y^l\right)$ as well. Precisely,

$$T\left(x^k, y^1\right) \leq ... \leq T\left(x^k, y^L\right) \leq T\left(x^{k+1}, y^1\right) \leq ... \leq T\left(x^{k+1}, y^L\right)$$

and the inequalities are strict unless both sides are 0.

Suppose firm 2 follows (τ_2^*, σ_2^*). Consider firm 1 with signal x^k. Then the argument that it is optimal to unilaterally exit at $\tau_1^*\left(x^k\right)$ is exactly the same as in proof of Proposition 1. Clearly, $\sigma_1^*\left(x^k, \cdot\right)$ is a best response given firm 1's beliefs about y^l.

Now suppose firm 1 follows (τ_1^*, σ_1^*). Consider firm 2 with signal y^l. The proposed strategy τ_2^* asks it to unilaterally drop out at $T\left(x^K, y^l\right)$, that is, it believes that firm 1's signal is x^K, the best possible scenario. Certainly, it cannot be a best response for y^l to choose $\tau_2\left(y^l\right) > T\left(x^K, y^l\right)$. Now consider a $\tau_2\left(y^l\right)$ such that $T\left(x^{K-1}\right) < \tau_2\left(y^l\right) < T\left(x^K, y^l\right)$. In this case, at time $\tau_2\left(y^l\right)$ firm 2 will learn if firm 1's signal is x^K or not. If it is x^K, then firm 1 will not quit before firm 2 quits and it cannot be a best response for 2 to then quit at $\tau_2\left(y^l\right)$. Thus, unilaterally quitting at $\tau_2^*\left(y^l\right) = T\left(x^K, y^l\right)$ is a best response to firm 1's strategy.

If y^l is an optimistic signal (as defined above), then $T\left(x^k, y^l\right) > T\left(x^k\right)$ and so $\sigma_2^*\left(y^l, T\left(x^k\right)\right) = 2T\left(x^k, y^l\right) - T\left(x^k\right)$, which is the optimal quitting time once firm 1 with signal x^k quits at $T\left(x^k\right)$ and firm 2 learns that 1's signal is x^k. But if y^l is a pessimistic signal, then $T\left(x^k, y^l\right) < T\left(x^k\right)$ and so $\sigma_2^*\left(y^l, T\left(x^k\right)\right) > T\left(x^k, y^l\right)$. In this case, once firm 1 with signal x^k quits at $T\left(x^k\right)$ and firm 2 learns that 1's signal is x^k, it quits immediately but suffers from some ex post regret for having stayed too long—in the period between $T\left(x^k, y^l\right)$ and $T\left(x^k\right)$ it loses money. But if Q_2 is close enough to 1, then the gap $T\left(x^k\right) - T\left(x^k, y^l\right)$ is small and the loss is small relative to the gain from learning. Thus, (τ_2^*, σ_2^*) is a best response to (τ_1^*, σ_1^*). ■

Proposition 2 generalizes to the case of many signals as well.

Proposition 2 (M) *There exists a $\overline{Q}_2 \in (1, Q_2^*)$ such that if $1 < Q_2 < \overline{Q}_2$, then the outcome in (σ^*, τ^*) is the unique Nash equilibrium outcome.*

Proof. The proof is very similar to the proof in the case of two signals. The iterated elimination of dominated strategies used to establish Proposition 2 can be mimicked. Here we indicate only the basic steps.

First, as in Round 1, any $\tau_1\left(x^k\right) > T\left(x^k, y^L\right)$ and $\tau_2\left(y^l\right) > T\left(x^K, y^l\right)$ are weakly dominated. In Round 2, any $\tau_1\left(x^k\right) < T\left(x^k, y^1\right)$ is strictly dominated. This means that by staying until $T\left(x^k, y^L\right)$ firm 2 can learn whether firm 1's signal is x^k or whether it is higher. In Round 3, if Q_2 is small enough, $\tau_2\left(y^l\right) < T\left(x^K, y^l\right)$ is strictly dominated. This is because the information of whether 1's signal is x^k or higher comes at the latest by $T\left(x^k, y^L\right)$ and when Q_2 is small, waiting for this information is relatively inexpensive. The remainder of the proof follows that of Proposition 2. ■

A Appendix: Common beliefs

Although our main concern in this paper is with asymmetric information, in this Appendix we study a situation in which the firms share a common belief about G

at time 0. This could, for instance, occur if the signals were publicly known. More important, the firms can learn each other's signal in the course of play in the upstart equilibrium (for instance, if the signals are g_1 and g_2, this occurs at time $T(g_1, b_2)$). At that point, they have common beliefs about the state.

Suppose that the common belief at time 0 is $p_0 > p^*$ and that firm 2 remains active indefinitely. We have seen that the optimal strategy for firm 1 is to remain active until time $T(p_0)$ as defined in (3):

$$e^{-2\lambda T(p_0)} \frac{p_0}{1-p_0} = \frac{p^*}{1-p^*}$$

The expected flow profits of firm 1 are then

$$\begin{aligned} v(p_0) &= r \int_0^{T(p_0)} e^{-rt} \left(e^{-2\lambda t} p_0 + 1 - p_0\right) \left(p_t \lambda \frac{m}{r} - c\right) dt \\ &= \lambda m \int_0^{T(p_0)} e^{-rt} \left(e^{-2\lambda t} p_0 + 1 - p_0\right) (p_t - p^*) \, dt \\ &= \lambda m \int_0^{T(p_0)} e^{-rt} \left(e^{-2\lambda t} p_0 (1-p^*) - (1-p_0) p^*\right) dt \end{aligned}$$

which results in

$$v(p_0) = \frac{p_0(1-p^*)}{2\lambda + r} \left((2\lambda + r) e^{-2T(p_0)\lambda} \left(e^{-rT(p_0)} - 1\right) - r\left(e^{-(2\lambda+r)T(p_0)} - 1\right)\right)$$

Using the definition of $T(p_0)$ from above, after some calculation we obtain that

$$v(p_0) = -c + \frac{m+2c}{\mu+2} p_0 + \frac{2c}{\mu+2} (1-p_0) \left(\frac{1-p_0}{p_0}\right)^{\frac{\mu}{2}} \left(\frac{p^*}{1-p^*}\right)^{\frac{\mu}{2}} \tag{13}$$

where $\mu = r/\lambda$. It is easy to see that v is an increasing, convex function and $v(p^*) = 0$. For $p_0 \leq p^*$, $v(p_0) = 0$ since it is optimal for a firm to stay out.

Similarly, if firm 1 were alone and had an initial belief $p_0 > p^*$, the optimal strategy would be to quit at $2T(p_0)$. The single-firm maximized value function in terms of flows is then

$$u(p_0) = -c + \frac{m+c}{\mu+1} p_0 + \frac{c}{\mu+1} (1-p_0) \left(\frac{1-p_0}{p_0}\right)^{\mu} \left(\frac{p^*}{1-p^*}\right)^{\mu} \tag{14}$$

Again, u is an increasing, convex function satisfying $u(p^*) = 0$. For $p_0 \leq p^*$, $u(p_0) = 0$. It can be verified that for all $p_0 > p^*$, $u(p_0) > v(p_0)$, that is, competition decreases profits.

When there is no asymmetric information and the beliefs are common, we have

Proposition 5 *With common beliefs, there is a unique Nash equilibrium outcome.*

Proof. With common beliefs, a strategy for firm i is a pair of functions (τ_i, σ_i) as in the main text (now there are no private signals, however). Suppose the common initial belief is p_0. It is easy to see that if firm j chooses $\tau_j = T(p_0)$, then it is a best response for firm $i \neq j$ to choose $\tau_i = T(p_0)$ as well.

To show uniqueness, first note that any (τ_i, σ_i) such that $\tau_i \neq T(p_0)$ is weakly dominated by $(T(p_0), \sigma_i)$. Thus, the Nash equilibrium outcome above is the only outcome that survives one round of elimination of weakly dominated strategies. The argument that there is no Nash equilibrium in weakly dominated strategies is the same as Step 2 in the proof of Proposition 2. ■

References

[1] Arrow, K. (1962): "Economic Welfare and the Allocation of Resources for Invention," in *The Rate and Direction of Inventive Activity*, R. Nelson (ed.), Princeton University Press, Princeton, NJ.

[2] Bassan, B., O. Gossner, M. Scarsini, and S. Zamir (2003): "Positive Value of Information in Games," *International Journal of Game Theory*, 32 (1), 17–31.

[3] Benoît, J. P., & Dubra, J. (2011): "Apparent Overconfidence," *Econometrica*, 79 (5), 1591–1625.

[4] Chatterjee, K. and R. Evans (2004): "Rivals' Search for Buried Treasure: Competition and Duplication in R&D," *Rand Journal of Economics*, 35 (1), 160–183.

[5] Chen, C-H. and J. Ishida (2017): "A War of Attrition with Experimenting Players," Osaka University, https://ideas.repec.org/p/dpr/wpaper/1014.html

[6] Choi, J. (1991): "Dynamic R&D Competition under 'Hazard Rate' Uncertainty," *Rand Journal of Economics,* 22 (4), 596–610.

[7] Christiansen, C. (1997): *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*, Harvard Business School Press, Boston, MA.

[8] Das, K. and N. Klein (2018): "Competition, Search and Duplication in a Patent Race," Department of Economics, University of Exeter, http://people.exeter.ac.uk/kd320/duplicative_search_1.pdf

[9] Dong, M. (2018): "Strategic Experimentation with Asymmetric Information," Penn State University, https://sites.google.com/site/miaomiaodong0/1-delimit

[10] Gilbert, R. and D. Newberry (1982): "Preemptive Patenting and the Persistence of Monopoly," *American Economic Review*, 76, 238–242.

[11] Igami, M. (2017); "Estimating the Innovator's Dilemma: Structural Analysis of Creative Destruction in the Hard Disk Industry, 1981–1998," *Journal of Political Economy*, 125, 798–847.

[12] Keller, G., S. Rady and M. Cripps (2005): "Strategic Experimentation with Exponential Bandits," *Econometrica*, 73, 39–68.

[13] Klein, N. and P. Wagner (2018): "Strategic Investment and Learning with Private Information," University of York, https://sites.google.com/site/peterachimwagner/research

[14] Malueg, D. and S. Tsutsui (1997): "Dynamic R&D Competition with Learning," *Rand Journal of Economics*, 28, 751–772.

[15] Moscarini, G. and F. Squintani (2010): "Competitive Experimentation with Private Information: The Survivor's Curse," *Journal of Economic Theory*, 145, 639–660.

[16] Rietveld, N., P. Groenen, P. Koellinger, M. Loos, and R. Thurik (2013): "Living Forever: Entrepreneurial Overconfidence at Older Ages," Erasmus Research Institute of Management, http://hdl.handle.net/1765/40673.

[17] Schumpeter, J. (1942): *Capitalism, Socialism and Democracy*, New York: Harper.

[18] Shimoji, M. and J. Watson (1998): "Conditional Dominance, Rationalizability and Game Forms," *Journal of Economic Theory*, 83, 161–195.

[19] Surowiecki, J. (2014): "Epic Fails of the Startup World," *The New Yorker*, May 19.

[20] Wong, T-N. (2018); "Free Riding and Duplication in R&D," University of Basel, https://sites.google.com/site/tszningwongecon/research

[21] Wu, B., and A. Knott (2006): "Entrepreneurial Risk and Market Entry," *Management Science,* 52 (9), 1315–1330.