

Startups and Upstarts: Disadvantageous Information in R&D*

Yu Awaya[†] and Vijay Krishna[‡]

June 12, 2020

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. The channel by which better information becomes a competitive disadvantage appears to be new—it does not stem from a negative value of information or from a second-mover advantage. Rather, it stems from the fact that better information dulls the incentive to learn from one’s rival.

*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, Silvana Krasteva, Erik Madsen, John Morgan, Tsz-Ning Wong and Shmuel Zamir for helpful discussions. We learnt a lot about R&D in the pharmaceutical industry from Joshua Krieger. Finally, we have greatly benefited from the comments of a fine set of referees.

[†]Department of Economics, University of Rochester, E-mail: yuawayaya@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 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. 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 Newbery, 1982). This "preemption effect" is at odds with the replacement effect. There are other forces that may favor incumbents as well—it may be better at R&D or have deeper pockets. 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.¹

In this paper, we identify an entirely new effect that, like the replacement effect, works to the detriment of the established firm. We suppose that the established firm has better information about the feasibility of an R&D project than a startup/entrant. We will show that better information is a competitive *disadvantage* and as a result, the less-informed entrant succeeds more often and has a higher payoff than the better-informed incumbent.

The following simple example illustrates how better information can be a competitive disadvantage.² Two firms compete in an R&D race with a random return

¹Echoing Arrow to some extent, there is empirical work showing that large, established firms tend to engage in R&D which is incremental—aimed at improving the quality of their existing products. The R&D activity of smaller firms tends to be radical—aimed at new products. This finding goes back to the 1980s and the recent work of Akcigit and Kerr (2018) confirms this.

²We are grateful to a referee for generously providing this example.

M that is equally likely to be 0 or 1. There are two periods and R&D costs c per period. In the first period, a firm decides whether or not to enter the race by investing and, if it does, in the second period, decides whether or not to invest again. In each period, firms move simultaneously but first-period decisions are observed prior to the second period. If one firm stays longer than the other, it gets M and the other 0. If both stay until the end, each gets $\frac{1}{2}M$. Otherwise, both get 0. The only asymmetry between the firms is that firm 1 receives a noisy signal about M . Specifically, prior to any decision, it gets either good news, g , or bad news, b . The probability that the signal is g when $M = 1$ is q , where $\frac{2}{3} < q < 1$, and the probability that the signal is b when $M = 0$ is also q . Firm 2 is completely uninformed. Conditional on the signals, the expected returns are $E[M | g] = q$ and the $E[M | b] = 1 - q$. Good news makes firm 1 optimistic while bad news makes it pessimistic. Suppose that the parameters satisfy (i) $\frac{1}{2} - \frac{1}{4}q > 2c$; (ii) $2c > 1 - q > c$.

The *unique* equilibrium of this game is: an optimistic firm 1 invests in the first period and invests in the second period only if it sees that firm 2 also invested. A pessimistic firm 1 stays out of the race. Firm 2 invests in the first period and invests in second period only if it sees that firm 1 also invested.

To see that this is an equilibrium, notice that if firm 2 sees that 1 invested in the first period, it knows that 1 is optimistic and so it is optimal for 2 to invest in the second period as well. This is because $q > \frac{2}{3}$ implies $\frac{1}{2}q > 1 - q > c$. Moreover, it is optimal for firm 2 to invest in the first period. This is because if 1 is pessimistic then 2's payoff from investing is $1 - q - c > 0$ by (ii); and if 1 is optimistic, then 2's payoff is $\frac{1}{2}q - 2c > 0$ which follows from (i) when $q > \frac{2}{3}$. Now, given firm 2's strategy, a pessimistic firm 1 does not want to enter. This is because if 1 enters, then it knows that firm 2 will invest in the second period as well and thus firm 1 will make a loss because $2c > 1 - q$, again by (ii). On the other hand, firm 1 with signal g wants to enter the race again because $\frac{1}{2}q - 2c > 0$.

In equilibrium, when 1 is optimistic both firms get a payoff of $\frac{1}{2}q - 2c > 0$. When 1 is pessimistic, firm 1 does not enter the race and gets 0 whereas firm 2 enters and gets $1 - q - c > 0$. Thus, the uninformed firm's overall expected payoff is *higher* than that of the informed firm.

The key here is that the uninformed firm 2 enters the race because by doing so, it can learn firm 1's information from the latter's actions. To see that learning is crucial, consider a variant of the example in which firms' first-period decisions are *not* observable. Now there is no possibility of learning and in the unique equilibrium, firm 1 enters and stays when it is optimistic and does not enter when pessimistic. Because (i) is the same as $\frac{1}{4}q + \frac{1}{2}(1 - q) > 2c$, firm 2 enters and stays for two periods because it does not know whether 1 is optimistic or pessimistic. Now the informed firm's expected payoff is higher than that of the uninformed firm.

While the example illustrates some of the ideas in this paper, it misses some important aspects of R&D races—that the overall chance of success is greater the longer one stays in the race, that firms may learn from lack of success, that firm 2

may also have some information, etc. Moreover, the result relies on some restrictions on the parameters.

In what follows, we study a standard continuous-time, winner-take-all R&D race. Firms must invest to engage in R&D on a project whose feasibility is unclear. Precisely, 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 infeasible. Firms do not know the state but *both* receive informative private signals about it and so each can learn from the other. The only difference between the firms is that the established firm's signal is more accurate than that of the startup. As the race proceeds, lack of success causes both firms to become increasingly pessimistic about the feasibility of the innovation. Each must then decide when to quit, a decision that is observed by its rival³ and is irrevocable⁴. Such models are rather standard—the basic structure originates in Choi (1991) and has been studied by others—and the only new ingredient we add is comparable asymmetric information.

To isolate this new effect, we assume that the firms are alike in all other respects—the costs and benefits of R&D as well as their R&D abilities are the same. Because the gains from R&D are the same, the replacement and preemption effects are absent.⁵

Our main result is⁶

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 has a lower payoff as well. The startup is favored to win precisely because it is less informed!

We call this 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. The learning is so unbalanced that the startup gains an advantage over the incumbent.

³Pharmaceutical companies must register drug trials with the Food and Drug Administration and report progress or lack thereof publicly (see Krieger, 2019 for details of the reporting process). In other industries, R&D activity must be reported to investors.

⁴We base this on the assumption that once a project is abandoned, it is prohibitively costly to restart it. This seems closer to reality than assuming that firms can shut down R&D and then costlessly restart it.

⁵The equilibrium studied here is robust and introducing small asymmetries in payoffs, R&D costs or abilities would not overturn the results.

⁶Theorem 1 summarizes the findings of Propositions 1 and 2 below. Proposition 3 identifies when the payoff ranking is strict.

The incumbent suffers from a "curse of knowledge"—its superior knowledge hinders learning.

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.⁷ 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 primary mechanism that drives our result is that firms revise their expectations of success from the exit decisions of their rivals. Is this true? In an interesting recent paper, Krieger (2019) uses data on over 10,000 drug-development projects that reached the stage of human trials (what is called phase II) to address this very question. The industry is particularly apt because regulations require that the results of drug trials be disclosed and so firms are well aware of others' R&D activity. Krieger finds that news about a competitor's failure indeed leads to a significant increase in the quitting rate of firms. In one instance, a firm announced that it was ending its development effort of a cancer drug and explicitly gave as its reason that a *competitor* had abandoned a similar project! Moreover, Krieger finds that smaller and inexperienced firms (with little prior drug development) are *less* likely to quit a project than larger and experienced incumbents. These findings echo our theoretical prediction that startups learn more from established firms than the other way around.

A new information "paradox"? Intuition suggests that information should confer a strategic advantage. In our model, it is a disadvantage. The fact that information can have paradoxical consequences in multi-person settings is, of course,

⁷Around 2006, a new kind of technology for television display screens—emission display—seemed very promising. The main competitors to develop this were Sony and a Canon/Toshiba joint venture. Sony abandoned the technology publicly in 2009 and other firms followed soon after.

well-understood and one may rightly wonder whether our main result is just a manifestation of a known phenomenon.

It is known that information may have a negative value. Hirshleifer (1971) showed that publicly available information may make all agents worse off *ex ante*. There are also examples of games in which information privately available to an agent reduces his or her *ex ante* payoff (see Bassan et al., 2003 or Maschler et al., 2013 for examples). One might rightly wonder then whether, in the game we study, the value of private information is negative as well. This is not the case. We show below that in the upstart equilibrium, the value of information is *positive* for both firms—each firm’s expected payoff is increasing in 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.

It is also known that in many situations, there is a second-mover advantage. Does the competitive advantage of the startup stem from the fact that it moves second and is able to learn the incumbent’s information? This is not the case either. In our model, the order of moves (when to exit) is not specified exogenously; rather it is determined in equilibrium. For some signal realizations, the incumbent exits early and the startup learns from the incumbent. In others, the startup exits early and the incumbent learns from the startup.

In our model, the incumbent has a competitive disadvantage because its superior information dulls its incentives to learn relative to the startup. The marginal value to the incumbent of learning the startup’s poor information is smaller than the marginal value to the startup of learning the incumbents better information. The startup is willing to wait-and-see whereas the incumbent is not. But how do we know that this is the reason for the "paradox"? At the end of Section 4 we study a variant of the main model in which firms cannot observe each other’s exit decisions. All other aspects of the model remain unchanged but now there is no possibility of learning from each other. We show that once there is no possibility of learning, the "paradox" disappears—in the *unique* equilibrium, the expected payoff of the better-informed firm is now higher than that of its rival.

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.⁸ As one observer of the startup phenomenon has written:

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

⁸See, 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.

Even though 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.

The incumbent may also be a victim of apparent overconfidence—it is so sure of its own information that it does not find it worthwhile to try to learn what the startup knows. This is the main driving force of our result but again, its basis is not psychological. Rather, it is the result of rational calculation.

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.⁹ 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 (2020) 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

⁹In 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.

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. Our model differs from that of Moscarini and Squintani 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. In equilibrium, the better-informed firm is subject to both curses while the less-informed firm is never subject to the survivor's curse. Finally, we derive circumstances in which there is a unique equilibrium outcome and these too depend on the relative quality of the firms' information.

Strategic experimentation Our model is related to those of strategic experimentation, especially with exponential bandits as in Keller, Rady and Cripps (2005). Unlike our model, 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 we make. Strategic experimentation models typically have multiple Markov equilibria whereas ours has a unique Nash equilibrium. Another difference is that whereas equilibria of strategic experimentation models display under-investment relative to a planner's solution—there is free-riding—in our model firms over-invest.

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.¹⁰ In the equilibrium she studies, this asymmetry induces more experimentation than if the situation were symmetric.

Wars of attrition Our model also shares important features with the war 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

¹⁰Klein and Wagner (2018) study a bandit problem where the quality of information of the players is the same.

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 then study the case of two competing firms and exhibit the upstart equilibrium mentioned above (Proposition 1). We then show that this equilibrium is unique when the asymmetry in the quality of information is large (Proposition 2). These two propositions establish Theorem 1, while Proposition 3 identifies conditions under which there is a unique equilibrium with a strict payoff ranking. We also show that despite the fact that in equilibrium the less-informed firm has lower payoff, it would still prefer higher quality information. In other words, the value of information is positive. Section 5 compares the equilibrium outcome to a planner’s solution and finds that relative to the planner, competition leads to over-investment in R&D. Section 6 concludes. An online appendix to this paper shows that the main results generalize when the firms may get more than two signals and so have finer information.

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.¹¹ 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.¹² Conditional on the state, the

¹¹This 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).

¹²In an online appendix we show that the analysis of the paper extends to the case where the firms’ information is finer—that is, they receive one of multiple signals.

signals of the two firms are independent and

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

We will refer to q_i as the *quality* of i 's signal or information.¹³ 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,

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

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,

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

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

$$p(b_1, b_2) < p(b_1, g_2) < p(g_1, b_2) < p(g_1, g_2) \quad (1)$$

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

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

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

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 that¹⁴

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

¹³The assumption that $\Pr [g_i | G] = \Pr [b_i | B]$ 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.

¹⁴This is just Bayes' rule in terms of odds ratios: given any event \mathcal{E} , we have

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

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,*

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

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

$$T(s_i) \equiv T(p(s_i)) \quad (3)$$

and

$$T(s_1, s_2) \equiv T(p(s_1, s_2)) \quad (4)$$

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

$$T(b_1, b_2) \leq T(b_1, g_2) \leq T(g_1, b_2) \leq T(g_1, g_2)$$

and each of the inequalities is strict unless both sides are 0.

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. Firm 1's payoff when alone is also higher.*

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




Figure 1: Belief Decay

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

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.

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

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

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 initial belief $p_0 > p^*$ is just $2T(p_0)$ since from the definition of T in (2),

$$e^{-2\lambda T(p_0)} \frac{p_0}{1 - p_0} = \frac{p^*}{1 - p^*} \tag{5}$$

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 $2T(p_0)$ is the optimal quitting time of a single firm with initial belief p_0 , using (5), the probability of success given p_0 is

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

Of course, if $p_0 \leq p^*$, then it is optimal to not enter. Thus, the probability of success given p_0 is

$$\Pr[\mathcal{S}_i] = \max\left(0, \frac{p_0 - p^*}{1 - p^*}\right)$$

which is a *convex* function of its belief p_0 .

Since $q_1 > q_2$, we have that firm 1's posterior beliefs at time 0 ($p(b_1), p(g_1)$) constitute a mean-preserving spread of firm 2's posterior beliefs ($p(b_2), p(g_2)$). Now it follows that

$$E[\Pr[\mathcal{S}_1]] \geq E[\Pr[\mathcal{S}_2]]$$

and so the ex ante probability of firm 1 succeeding when alone is at least as large as that of firm 2.

Finally, since this is a single-firm decision problem and firm 1's signal is more informative than that of firm 2 (in the sense of Blackwell), it follows that firm 1's expected payoff is also higher.

This completes the proof of Proposition 0.

4 Upstart equilibrium

In this section, we first establish that with two firms there is always an equilibrium that results in the upstart outcome. Unlike in the case with a single firm, both firms can learn a rival's signal via its actions. Indeed, as we will see, this learning is asymmetric and the less-informed firm learns more from its rival than the better-informed firm. Formally,

Proposition 1 *There exists a perfect Bayesian equilibrium in which the less-informed firm 2's winning probability and payoff are both no less than those of the better-informed firm 1.*

In this game, a *strategy* for firm i is a pair of functions (τ_i, σ_i) where $\tau_i : \{g_i, b_i\} \rightarrow \mathbb{R}_+ \cup \{\infty\}$ and $\sigma_i : \{g_i, b_i\} \times \mathbb{R}_+ \rightarrow \mathbb{R}_+ \cup \{\infty\}$. The first component $\tau_i(s_i)$ 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. The second component determines

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

Since this is a game in continuous time it is worthwhile to explicitly state how strategies translate into outcomes. If the signals are (s_i, s_j) , and $\tau_i(s_i) < \tau_j(s_j)$, then of course firm i exits at $\tau_i(s_i)$ and firm j at $\sigma_j(s_j, \tau_i(s_i))$. If $\tau_i(s_i) = \tau_j(s_j)$, then both firms exit simultaneously.¹⁵

The following strategies result in an equilibrium in which firm 2's payoff is no less than that of firm 1. Firm 1's unilateral exit times are

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

(see (3) and (4)) and firm 1's exit times following an exit by firm 2 at some time t_2 are

$$\begin{aligned} \sigma_1^*(g_1, t_2) &= \begin{cases} T(g_1, b_2) & \text{if } t_2 = T(g_1, b_2) \\ 2T(g_1, b_2) - t_2 & \text{if } t_2 < T(g_1, b_2) \\ 2T(g_1, g_2) - t_2 & \text{if } T(g_1, b_2) < t_2 < T(g_1, g_2) \end{cases} \\ \sigma_1^*(b_1, t_2) &= \begin{cases} \max(t_2, 2T(b_1, b_2) - t_2) & \text{if } t_2 \leq T(b_1) \\ t_2 & \text{if } t_2 > T(b_1) \end{cases} \end{aligned}$$

If firm 2 exits at a time $t_2 \leq T(g_1, b_2)$, then firm 1 believes that $s_2 = b_2$ and if firm 2 exits at $t_2 > T(g_1, b_2)$, believes that $s_2 = g_2$.

Firm 2's unilateral exit times are

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

and its exit times following an exit by firm 1 at some time t_1 are

$$\begin{aligned} \sigma_2^*(g_2, t_1) &= \begin{cases} 2T(b_1, g_2) - t_1 & \text{if } t_1 \leq T(b_1) \\ 2T(g_1, g_2) - t_1 & \text{if } T(b_1) < t_1 < T(g_1, g_2) \end{cases} \\ \sigma_2^*(b_2, t_1) &= \begin{cases} \max(t_1, 2T(b_1, b_2) - t_1) & \text{if } t_1 \leq T(b_1) \\ 2T(g_1, b_2) - t_1 & \text{if } T(b_1) < t_1 < T(g_1, b_2) \end{cases} \end{aligned}$$

If firm 1 exits at a time $t_1 \leq T(b_1)$, then firm 2 believes that $s_1 = b_1$ and if firm 1 exits at $t_1 > T(b_1)$, believes that $s_1 = g_1$.

These strategies result in the "upstart outcome" depicted in Figure 2. When the signals are (b_1, b_2) , firm 1 exits early at $T(b_1)$, as defined in (3), and firm 2 exits immediately afterwards.¹⁶ When the signals are (g_1, g_2) , both firms exit 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 firm 1 follows immediately. Finally, when the signals are (b_1, g_2) , firm 1 exits at $T(b_1)$ and firm 2 exits at $2T(b_1, g_2) - T(b_1)$.

¹⁵We 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.

¹⁶It could be, of course, that $T(b_1) = 0$ and in that case exiting at $T(b_1)$ is the same as not entering.




Figure 2: Upstart Equilibrium

Firm 1 (top) with signal b_1 exits at $T(b_1)$. The exit decision of firm 1 with g_1 depends on firm 2. If firm 2 exits at $T(g_1, b_2)$, then firm 1 with signal g_1 follows immediately, depicted as a U-turn. Otherwise, it stays until $T(g_1, g_2)$. Similarly, the exit decisions of firm 2 (bottom) depend on firm 1. If firm 1 exits at $T(b_1)$, firm 2 with signal b_2 follows immediately and with g_2 , exits at $2T(b_1, g_2) - T(b_1)$. Otherwise, firm 2 with b_2 exits at $T(g_1, b_2)$ and with g_2 , exits at $T(g_1, g_2)$.

Why is this an equilibrium? First, consider firm 1 with signal b_1 . Based on its information alone b_1 would want to quit unilaterally at $T(b_1)$. The only reason it may want to stay longer is if by doing so, it could learn 2's signal. However, no matter what 2's signal is, it cannot be optimal for b_1 to stay after $T(b_1, g_2)$. But firm 2's strategy dictates that with either signal, it will not be the first to quit before $T(b_1, g_2)$. This means that it is unprofitable for b_1 to stay long enough to learn 2's signal. Thus, b_1 quits unilaterally based only on its own signal, that is, at time $T(b_1)$ which is just $\tau_1^*(b_1)$.

The situation is different if firm 1 has signal g_1 . Now it is certainly worthwhile for 1 to wait until $T(g_1, b_2)$ because $T(g_1) > T(g_1, b_2)$. If 2 quits at this time, then 1 learns that its signal is b_2 and then it is optimal for 1 to follow 2 immediately. If 2 does not quit at $T(g_1, b_2)$, then 1 infers that 2's signal is g_2 and then it is optimal for g_1 to stay until $T(g_1, g_2)$, again as dictated by the equilibrium strategy.

Now consider firm 2 with signal b_2 . It knows that firm 1's signal will be revealed at $T(b_1)$ and since $T(b_1) < T(b_2)$, it is optimal for b_2 to wait until $T(b_1)$ to learn 1's signal. If 1 quits at $T(b_1)$, then its signal must be b_1 and then since $T(b_1) > T(b_1, b_2)$, it is optimal for b_2 to quit immediately. If 1 does not quit at $T(b_1)$, then its signal must be g_1 and then it is optimal for b_2 to quit at $T(g_1, b_2)$ since g_1 will surely stay until that time as well. This is exactly what is prescribed by τ_2^* .

Finally, consider firm 2 with signal g_2 . As above, it knows that firm 1's signal will be revealed at $T(b_1)$ and it is optimal to wait at least until then to learn 1's signal. As above, if 1 quits at $T(b_1)$, then its signal must be b_1 and to quit at $2T(b_1, g_2) - T(b_1)$

is then optimal for g_2 . This is because the belief of firm 2 reaches p^* at that time.¹⁷ If 1 does not quit at $T(b_1)$, then its signal must be g_1 and since 1 will then stay until $T(g_1, g_2)$, it is optimal for 2 to stay until then as well. Again, this is what τ_2^* dictates.

It is clear that given their off-equilibrium beliefs, both firms are optimizing as well.¹⁸ Thus, the strategies (σ_i^*, τ_i^*) constitute an equilibrium.

The equilibrium outcome has the feature that except in the case when the signals are (b_1, g_2) , the firms exit at the *same* time and, when they do so, have the same winning probabilities and expected profits. When the signals are (b_1, g_2) , however, firm 1 exits first at $T(b_1)$ and after seeing this, firm 2 exits optimally at time $\max(T(b_1), 2T(b_1, g_2) - T(b_1))$. Again, firm 2's winning probability and payoff are at least as great as those of firm 1.

This establishes Proposition 1.

It is worthwhile to note that when the signals are (b_1, g_2) , the learning is asymmetric. Firm 2 finds it in its interest to wait until firm 1's signal is revealed—which occurs at time $T(b_1)$ —but firm 1 does not find it worthwhile to wait until firm 2's signal will be revealed—which occurs at $T(g_1, b_2) > T(b_1)$. It is precisely the fact that firm 1's signal is of higher quality than firm 2's signal that leads to this asymmetry. Ex post, in this case, firm 1 suffers from regret. Had it known that firm 2's signal was g_2 , it would have liked to stay until $T(b_1, g_2)$ but knowing only its own signal b_1 , it is optimal to exit at $T(b_1) < T(b_1, g_2)$.

Uniqueness We now argue 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 q_2^+ such that for all $q_2 < q_2^+$, there is a unique Nash equilibrium outcome.*

Proposition 2 is proved by showing that the iterated elimination of dominated strategies leaves a single outcome. The argument is somewhat involved—there are many rounds of elimination—and so a formal proof is relegated to Appendix A. Here we indicate the basic idea and the most important step.

Let us consider firm 1's unilateral exit times $\tau_1(b_1)$ and $\tau_1(g_1)$. It is clear that firm 1 with signal g_1 should stay at least until $T(g_1, b_2)$, the optimal exit time when the 2's information is as bad as possible. In other words, for g_1 to exit before $T(g_1, b_2)$ is

¹⁷Formally, by definition $\frac{p(b_1, g_2)}{1-p(b_1, g_2)} e^{-\lambda T(b_1)} e^{-\lambda(2T(b_1, g_2) - T(b_1))} = \frac{p^*}{1-p^*}$. The first exponential term is the probability that firm 1 did not succeed and the second is the probability that firm 2 did not do so either.

¹⁸It can be argued that the particular choice of off-equilibrium beliefs does not affect the equilibrium outcome.

dominated. Similarly, it is clear that firm 1 with signal b_1 should not stay longer than $T(b_1, g_2)$, the optimal exit time when 2's information is as good as possible. Thus, for b_1 to stay longer than $T(b_1, g_2)$ is also dominated. Since $T(b_1, g_2) < T(g_1, b_2)$, we have $\tau_1(b_1) \leq T(b_1, g_2) < T(g_1, b_2) \leq \tau_1(g_1)$. In other words, firm 1's strategy fully reveals its information at the latest by time $T(b_1, g_2)$. Thus firm 2 can learn firm 1's signal for sure by investing until $T(b_1, g_2)$. The key step is that when the quality of firm 2's information is poor relative to that of firm 1, that is, q_2 is small relative to q_1 , $T(b_1, g_2)$ is small and firm 2 can learn 1's information at a relatively small cost. Thus, it is dominated for firm 2 to exit before learning 1's information. The remaining rounds of elimination are rather straightforward and do not require any conditions on the relative quality of the firms' information.

As shown in Appendix A, this process of iterated elimination of dominated strategies results in a *single* outcome. In the iterated process in some rounds we eliminate *weakly* dominated strategies. As a final step, it is shown that there cannot be any other Nash equilibrium outcome—the weakly dominated strategies that were eliminated cannot be part of any Nash equilibrium.

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. In particular, in the *symmetric* situation with $q_1 = q_2$, there are at least two asymmetric equilibria. One is the upstart equilibrium specified above in which firm 2's payoff is higher than that of firm 1. But now there is also an equilibrium which is a "mirror image" of the upstart equilibrium with the roles of firms 1 and 2 interchanged. The "mirror equilibrium" persists as long as $q_1 - q_2$ is small.

Strict ranking of payoffs When is firm 2's payoff strictly higher than that of firm 1? To begin, it is useful to consider two cases where the payoff ranking is *not* strict—the two firms' payoffs are the same. First, if firm 1 is perfectly informed ($q_1 = 1$), then it will enter the race if and only if the state is G . Then firm 2 cannot do better than to follow firm 1. Firm 2 should enter and if it sees that 1 did not, then it should exit immediately. If firm 1 entered, then firm 2 should stay. The second extreme case is when firm 2 has no information ($q_2 = \frac{1}{2}$). Firm 1 should then base its actions only on its own information and again firm 2 cannot do better than to follow firm 1. In both of these extreme cases, in every circumstance, the two firms stay for the same amount of time and so their payoffs are the same.

Thus, for the payoffs to be strictly ranked it is necessary that firm 1 not be perfectly informed and that firm 2 have some useful information. So suppose that $q_1 < 1$ and $q_2 > \frac{1}{2}$. Indeed the payoff ranking is strict when, for fixed q_1 , the quality of firm 2's information, q_2 , is relatively high and the "curse of knowledge" is particularly acute.

To see this formally, recall that in the upstart equilibrium, the two firms' payoffs




Figure 3: Uniqueness and Strict Payoff Ranking

There is a unique equilibrium outcome below the upper curve. The payoff ranking is strict above the lower curve.

are potentially different only when the signals are (b_1, g_2) . In this case, firm 1 exits at $T(b_1)$ and firm 2 stays until $2T(b_1, g_2) - T(b_1)$. Firm 2 stays strictly longer if $T(b_1, g_2) > T(b_1)$. But given $q_1 < 1$ and $q_2 > \frac{1}{2}$, this is the case if and only if $T(b_1, g_2) > 0$ which is the same as $p(b_1, g_2) > p^*$. Note that a necessary condition for this is that the prior probability $\pi > p^*$. Define

$$q_2^- = \min \{q_2 \in [\frac{1}{2}, q_1] : p(b_1, g_2) \geq p^*\} \quad (6)$$

Note that if $T(b_1) > 0$, then $T(b_1, g_2) > 0$, and the payoff ranking is strict. In this case, $q_2^- = \frac{1}{2}$. Fixing all other parameters, $T(b_1) > 0$ when q_1 is relatively small so that firm 1 enters even with b_1 . Moreover, q_2^- is non-decreasing in q_1 .

Thus, in the upstart equilibrium, if q_2 is *large* enough, specifically, $q_2 > q_2^-$, firm 2's winning probability and payoff are both strictly higher than those of firm 1. At the same time Proposition 2 shows that the upstart equilibrium outcome is unique if q_2 is *small* enough, specifically $q_2 < q_2^+$ (defined in (7)). Are the two conditions on q_2 compatible? The answer is yes. Formally,

Proposition 3 *Assuming $\pi > p^*$, for any $q_1 \in (\frac{1}{2}, 1)$, it is the case that $q_2^- < q_2^+$. In other words, there is an open interval of q_2 's for which there is a unique equilibrium with a strict payoff ranking.*

Figure 3 depicts how q_2^- and q_2^+ vary with q_1 .

In the argument for uniqueness outlined above (and formally proved in Appendix A), the only place that q_2 is required to be small relative to q_1 is to show that when 2 has the signal b_2 , it is willing to wait as long as $T(b_1, g_2)$ for 1's information to be revealed. The threshold was named q_2^+ , as defined in (7) in Appendix A. We will show that $q_2^- < q_2^+$.

As above, if $T(b_1) > 0$, then by definition, $q_2^- = \frac{1}{2}$ which is certainly smaller than q_2^+ . On the other hand, if $T(b_1) = 0$, and $q_2 = q_2^-$, then $p(b_1, g_2) = p^*$ and so $T(b_1, g_2) = 0$ as well. Now we will argue that when $q_2 = q_2^-$, the equilibrium is unique. To see this, recall if $T(b_1, g_2) = 0$ then it is dominated for firm b_1 to enter. This means that firm 2 can learn firm 1's signal at zero cost. Now from the definition of q_2^+ , $q_2^- < q_2^+$.

This establishes Proposition 3.

A numerical example It is useful to see the workings of the upstart equilibrium in a numerical example.

Suppose the flow payoff from the patent $m = 2$ and the flow cost of R&D $c = 0.5$. The arrival rate of success $\lambda = 0.1$ and the interest rate $r = 0.05$. Thus, $p^* = (rc/\lambda m) = 0.125$.

The prior probability of state G is $\pi = \frac{1}{2}$. The quality of firm 1's signal $q_1 = 0.85$. With these parameters, there is a unique equilibrium for all $q_2 < q_2^+ \simeq 0.77$, the largest value of q_2 for which firm 2 with signal b_2 is willing to wait until $T(b_1, g_2)$ to see 1's signal (q_2^+ is defined in (7) in Appendix A). There is a strict payoff ranking for all $q_2 > q_2^- = \frac{1}{2}$ (q_2^- is defined in (6)).

When $q_2 = 0.75$, the upstart equilibrium outcomes are as follows. When the signals are (b_1, b_2) , firm 1 exits at $T(b_1) = 1.06$ and firm 2 follows immediately. When the signals are (b_1, g_2) firm 1 exits at $T(b_1)$ and firm 2 stays until $2T(b_1, g_2) - T(b_1) = 12.04$. When the signals are (g_1, b_2) , firm 2 exits at $T(g_1, b_2) = 6.55$ and firm 1 follows immediately. Finally, when the signals are (g_1, g_2) , both exit at $T(g_1, g_2) = 23.89$.

Figure 4 traces the evolution of firm 2's beliefs when its signal is g_2 and it learns firm 1's signal at time $T(b_1)$. If it is g_1 , then g_2 's beliefs jump up and if it is b_1 , they jump down.

The unique equilibrium payoffs of the firms when $(q_1, q_2) = (0.85, 0.75)$ are $\Pi_1^* = 0.275$ and $\Pi_2^* = 0.296$. Firm 2's payoff is 7% higher than that of firm 1. The ex ante probability that firm 1 wins the patent is 0.43 and the ex ante probability that firm 2 wins is 0.50.

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. This suggests perhaps that firm 1 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




Figure 4: Evolution of Firm 2's Beliefs with Signal g_2

Firm 2's initial belief is $p(g_2)$ and decays at the rate 2λ with the lack of success. Firm 2's belief is revised at time $T(b_1)$. If firm 1 exits at that time, firm 2 infers that 1's signal was b_1 and 2's beliefs jump down to the lower curve and then decay at a slower rate λ since now it is the only firm in the race. If firm 1 does not exit at that time, firm 2 infers that 1's signal was g_1 and 2's beliefs jump up to the upper curve where they continue to decay at a rate 2λ since both firms are still in the race.

information for both firms is *positive*.¹⁹

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 .*

The formal proof of this proposition is in Appendix B but the fact that firm 1's payoff is increasing in q_1 can be understood by reasoning as follows. Consider the following artificial situation:

Firm 1 has to decide when to exit after *exogenously* receiving three signals b_1 , (g_1, b_2) and (g_1, g_2) with the appropriate probabilities in each state (in state G , these are $1 - q_1$, $q_1(1 - q_2)$ and q_1q_2). Firm 2's behavior is also *exogenously* specified—it never exits.

It is clear that in this artificial situation, the optimal stopping times for firm 1 are $T(b_1)$, $T(g_1, b_2)$ and $T(g_1, g_2)$, the *same* as in the upstart equilibrium. The payoffs are also the *same*. This reason is that, in the upstart equilibrium, firm 2 never exits

¹⁹Bassan 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.




Figure 5: Value of Information

Equilibrium payoffs of both firms are depicted as functions of q_1 . The kink in the Π_2^* curve occurs when q_1 is high enough so that $T(b_1) = 0$ and the two curves merge once $T(b_1, q_2) = 0$ as well.

before firm 1 and what happens after a firm exits has no effect on its payoffs. But this artificial situation is a single-person decision problem for firm 1 and so information has a positive value. Thus, it has positive value in the upstart equilibrium as well.

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. See Figure 5 which depicts, for fixed q_2 , the upstart equilibrium payoffs Π_1^* and Π_2^* as functions of q_1 . Notice that in the upstart equilibrium $\Pi_1^* < \Pi_2^*$ even when $q_1 = q_2$. Of course, as discussed above, in that case there is a corresponding "mirror equilibrium" as well in which the payoff ranking is reversed.

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 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 $(q_1, q_2) = (\frac{1}{2}, q'')$ so that firm 1 is now less informed than firm 2. In this situation, there is again a unique




Figure 6: Willful Ignorance

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

equilibrium, but this time it is firm 1 which is the upstart.²⁰ This equilibrium is what we have called a "mirror equilibrium" 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 6) that $\Pi_1^{**}(\frac{1}{2}, q'') = \Pi_2^*(q'', \frac{1}{2})$. 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, $\Pi_2^*(q'', \frac{1}{2}) = \Pi_1^*(q'', \frac{1}{2})$. But in the region where the quality of firm 1's information is higher than that of firm 2, Π_1^* is increasing in *both* qualities (Proposition 4 and Corollary 1 in Appendix B). Thus,

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

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 $p''(g_1, b_2) = \Pr[G | g_1, b_2]$ 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 here is: Don't fire the experienced researchers. Willful ignorance does not pay!

²⁰Any attempt to carry out this exercise when there are multiple equilibria is, of course, fraught with peril.

Unobserved exit Why is better information a competitive disadvantage in the R&D race studied in this paper? To explore this question consider a variation of the model of the earlier sections in which a firm's exit decisions are *unobserved* by its rival. In all other respects, the model is the same as outlined in Section 2. If exit is unobserved, then firms can no longer learn each other's private information and now it is no longer the case that better information is a competitive disadvantage. This then serves to isolate the reason why better information is a disadvantage. The better-informed firm has a smaller incentive to learn from the less-informed firm than the other way around and it is this, and only this, that leads to the surprising conclusion that the incumbent firm is at a disadvantage.

With unobserved exit, a strategy for a firm is merely a function $\tau_i : \{b_i, g_i\} \rightarrow \mathbb{R}_+$, that is, a pair of unilateral exit times. We then have

Proposition 5 *With unobserved exit, there is a unique Nash equilibrium $\bar{\tau}$ of the R&D race. In this equilibrium,*

$$\bar{\tau}_1(b_1) \leq \bar{\tau}_2(b_2) \leq \bar{\tau}_2(g_2) \leq \bar{\tau}_1(g_1)$$

and the payoff of the better-informed firm 1 is greater than that of the less-informed firm 2.

A formal proof of the proposition is in the online appendix to the paper.

5 Planner's solution

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.

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 $s \leq T(p_0)$ and then one of the firms exits. We now argue that it is optimal for the planner to ensure that both firms are active until $T(p_0)$. (See Appendix C for a formal proof.)

To see why, let us compare the situation where the planner shuts down one firm's R&D activity at time $s < T(p_0)$ to a situation where the planner continues with both firms for an infinitesimal time longer until $s + dt$. When only one firm is active in the interval $(s, s + dt)$, the planner's payoff is approximately

$$p_s \times \lambda \frac{m}{r} - c$$

where p_s is the planner's belief at time s given that neither firm has succeeded until then. Note that this is positive because $p_s > p^*$. When both firms are active in the interval $(s, s + dt)$, the planner's payoff is approximately

$$p_s \times 2\lambda \frac{m}{r} - 2c$$

In other words, when two firms are active both the arrival rate of success as well as the costs of R&D are doubled. Thus, it is not optimal to shut down one of the firms at any time $s < T(p_0)$. This means that 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 $p_t > p^*$ or alternatively, until time $T(p_0)$.

How does the upstart equilibrium compare to the planner's optimum?

Recall that in the upstart equilibrium, when the signals are (g_1, b_2) , both firms exit at $T(g_1, b_2)$ which is also the planner's optimum exit time. The same is true when the signals are (g_1, g_2) . When the signals are (b_1, b_2) the upstart equilibrium may involve too much R&D when $T(b_1, b_2) < T(b_1)$ because in the upstart equilibrium both firms exit at $T(b_1)$ while the planner would want them to exit at $T(b_1, b_2)$. Finally, when the signals are (b_1, g_2) firm 1 exits at $T(b_1)$ while firm 2 invests until time $2T(b_1, g_2) - T(b_1)$. Conditional on (b_1, g_2) , the probability of success in this event is then $1 - e^{-2\lambda T(b_1, g_2)}$ 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 exits early—at time $T(b_1)$. This causes "learning-from-failure" to slow down relative to the case when two firms invest, which is the planner's solution.

Proposition 6 *The overall probability of R&D success is higher in the upstart equilibrium than in the planner's optimum.*

Thus, the overall probability of R&D success is higher in the upstart equilibrium than in the planner's optimum—there is over-investment in R&D. 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 competition enhances the chances of R&D success. This runs counter to the sentiments expressed by Schumpeter (1942). This is also counter to the findings in the literature on strategic experimentation where typically, there is under-investment (see Keller, Rady and Cripps, 2005).

6 Conclusion

We have argued that, somewhat paradoxically, informational asymmetry favors startups over incumbents. This purely informational effect serves to enhance Arrow's replacement effect. The effect appears to be new—it does not stem from a negative value of information or from a second-mover advantage. Rather it stems from the fact that better information diminishes the incentives to learn from one's rival.

At a theoretical level, we have shown that a natural R&D game, in which the only asymmetry is informational, has the following features. There is a unique equilibrium in which information is a competitive disadvantage even though it has positive value.

The equilibrium is robust—it is almost the unique rationalizable outcome of the game—and so the finding is not a knife-edge result. Introducing small asymmetries in R&D costs, abilities or the returns to invention would not overturn the results.

A Appendix: Uniqueness

In this appendix, we provide the formal proof that when the discrepancy in the two firms' information quality is large, there is a unique Nash equilibrium (Proposition 2). We do this via a process of iterated elimination of weakly/strongly dominated strategies.²¹

In what follows, 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 A.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. ■

With the lemma in hand, we are ready to begin the iterated procedure for removing dominated strategies.

²¹In fact the upstart equilibrium is also the unique outcome resulting from iterated elimination of *conditionally dominated* strategies (Shimoji and Watson, 1998).

A.1 Step 1

Denote by Γ the original game and by $\Gamma(n)$ the game after n rounds of elimination. In what follows, Lemma A.1 will be 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 .*

Proof. Quitting at $\tau_1(g_1) < T(g_1, b_2)$ is weakly dominated by quitting at $\bar{\tau}_1(g_1) = T(g_1, b_2)$. 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 of 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 (a) *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 $T(b_1, g_2)$. ■

Claim 2 (b) *Any strategy of firm 2 such that for $s_2 = b_2$ or g_2 , $\tau_2(s_2) < T(b_1, b_2)$ is **strictly dominated** in $\Gamma(1)$.*

Proof. Clearly, the worst possible case for firm 2 is if firm 1 has the signal b_1 . For firm 2 with either signal exit before $T(b_1, b_2)$ is no better than staying until $T(b_1, b_2)$. But if firm 1's signal is g_1 , then from Claim 1 (a) $\tau_1(g_1) \geq T(g_1, b_2) \geq T(b_1, b_2)$. Since with positive probability firm 1 will stay until $T(b_1, b_2)$, it is strictly dominated for firm 2 to quit before then. ■

IEDS Round 3

Claim 3 *Given all other parameters, there exists a q_2^+ such that for all $q_2 < q_2^+$, 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. Claim 1 (a) and Claim 2 (a) 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 if $\tau_2(b_2) < T(b_1, g_2)$ then (τ_2, σ_2) is strictly dominated by $(\bar{\tau}_2, \bar{\sigma}_2)$ such that $\bar{\tau}_2(b_2) = T(g_1, b_2)$ and $\bar{\sigma}_2(b_2, t_1) = t_1$ for all $t_1 \leq T(b_1, g_2)$. In words, the strategy $(\bar{\tau}_2, \bar{\sigma}_2)$ says that if 1 exits at time $t_1 \leq T(b_1, g_2)$, then b_2 believes that firm 1's signal is b_1 and exits immediately after firm 2. If firm 1 does not exit by $T(b_1, g_2)$, then b_2 believes that firm 1's signal is g_1 . From Claim 2 (b) we already know that to exit before $T(b_1, b_2)$ is strictly dominated for b_2 . We will now establish that if q_2 is small enough, a unilateral exit time $\tau_2(b_2) < T(g_1, b_2)$ is also strictly dominated.

Given $\tau_1(b_1) = t_1$, b_2 's flow profit from the strategy $(\bar{\tau}_2, \bar{\sigma}_2)$ when evaluated at time $T \leq \tau_1(b_1)$ is

$$V_T(t_1) = \lambda m \int_T^{t_1} e^{-r(t-T)} \Pr[\mathcal{S}_0(t) \mid b_2, \mathcal{S}_0(T)] (p_t(b_2) - p^*) dt \\ + \Pr[g_1 \mid b_2, \mathcal{S}_0(T)] \times \lambda m \int_{t_1}^{T(g_1, b_2)} e^{-r(t-T)} \Pr[\mathcal{S}_0(t) \mid g_1, b_2, \mathcal{S}_0(T)] (p_t(g_1, b_2) - p^*) dt$$

where $\mathcal{S}_0(t)$ is the event that neither firm has succeeded until t and 2's beliefs at time $t < t_1$ are defined by

$$\frac{p_t(b_2)}{1 - p_t(b_2)} = e^{-2\lambda t} \frac{p(b_2)}{1 - p(b_2)}$$

If firm 1 does not exit at time t_1 , then firm 2 knows that $s_1 = g_1$ and its beliefs at time $t > t_1$ become

$$\frac{p_t(g_1, b_2)}{1 - p_t(g_1, b_2)} = e^{-2\lambda t} \frac{p(g_1, b_2)}{1 - p(g_1, b_2)}$$

If firm 1 exits at time t_1 , then firm 2 follows and its subsequent payoff is 0.

For any T , firm 2's payoff $V_T(t_1)$ is decreasing in t_1 and so is minimized at $t_1 = T(b_1, g_2)$. This is because b_2 learns 1's signal at $\tau_1(b_1) = t_1$ and the sooner this

information arrives, the better it is for firm 2. Formally, by differentiating $V_T(t_1)$ with respect to t_1

$$V'_T(t_1) \propto \Pr[\mathcal{S}_0(t_1) | b_2, \mathcal{S}_0(T)](p_{t_1}(b_2) - p^*) \\ - \Pr[g_1 | b_2, \mathcal{S}_0(T)] \Pr[\mathcal{S}_0(t_1) | g_1, b_2, \mathcal{S}_0(T)](p_{t_1}(g_1, b_2) - p^*)$$

Now using the fact that at time t_1 , the "prior" before the information arrives $p_{t_1}(b_2)$ is the expectation of the posteriors after the information arrives (Bayes plausibility),

$$V'_T(t_1) \propto \Pr[b_1 | b_2, \mathcal{S}_0(T)] \Pr[\mathcal{S}_0(t) | b_1, b_2, \mathcal{S}_0(T)](p_{t_1}(b_1, b_2) - p^*) < 0$$

since $t_1 > T(b_1, b_2)$.

We now argue that the minimum value $V(T(b_1, g_2))$ is positive once q_2 is small enough. To see this, note that while the first term of $V_T(T(b_1, g_2))$ may be negative, the second is surely positive. As $q_2 \downarrow \frac{1}{2}$, $p(b_1, g_2) \downarrow p(b_1, b_2)$, or equivalently, $T(b_1, g_2) \downarrow T(b_1, b_2)$, and since $T \in [T(b_1, b_2), T(b_1, g_2)]$ the first term approaches zero while the second is strictly positive when $T(g_1, b_2) > 0$. Now define

$$q_2^+ = \max \left\{ q_2 : \min_{T \leq T(b_1, g_2)} V_T(T(b_1, g_2)) \geq 0 \right\} \quad (7)$$

For any $q_2 < q_2^+$, the payoff from $(\bar{\tau}_2, \bar{\sigma}_2)$ 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 firm 2 is indifferent at all $\tau_2(b_2) > T(b_1, g_2)$. But if $s_1 = g_1$, $\bar{\tau}_2(b_2) = T(g_1, b_2)$ is strictly better than $\tau_2(b_2) < T(b_1, g_2)$. Since the latter occurs with positive probability, $(\bar{\tau}_2, \bar{\sigma}_2)$ 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 $q_2 < q_2^+$. Note that q_2^+ depends on the other parameters, in particular on q_1 .

IEDS Round 4

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

Proof. We will first argue that $\tau_1(b_1) > T(b_1)$ is strictly dominated. 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 exits at any time $t_1 \in (T(b_1), T(b_1, g_2)]$, its flow payoff after $T(b_1)$ is negative (recall that $T(b_1)$ is the optimal exit time for firm 1 with only his own signal b_1). If firm 1 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. Thus, staying after $T(b_1)$ is strictly dominated for b_1 .

Clearly there is no reason for b_1 to quit before $T(b_1)$. ■

IEDS Round 5

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

Proof. Given all previous rounds, we know that firm 1 with b_1 will exit at $T(b_1)$ and with signal g_1 will exit no earlier than $T(b_1, g_2) > T(b_1)$. Thus, if firm 2 sees at time $T(b_1)$ that firm 1 exits, it knows that 1's signal was b_1 . If firm 2's signal is g_2 , it is now strictly dominated to quit at a time other than $2T(b_1, g_2) - T(b_1)$, the time when g_2 's beliefs will reach p^* . ■

Claim 5 (b) *Any strategy of firm 2 such that $\sigma_2(b_2, T(b_1)) \neq T(b_1)$ is strictly dominated in $\Gamma(4)$.*

Proof. Given all previous rounds, we know that firm 1 with b_1 will exit at $T(b_1)$ and with g_1 will stay longer. Thus, if firm 2 sees at time $T(b_1)$ that firm 1 exits, it knows that 1's signal was b_1 . Clearly, given that 2's own signal is b_2 , staying any longer 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 with b_1 will exit at $T(b_1)$ and with g_1 will stay longer. Thus, if firm 2 sees at time $T(b_1)$ that firm 1 did not exit, 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 with signal b_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 with b_1 will exit at $T(b_1)$ and with g_1 will stay longer. Thus, if firm 2 sees at time $T(b_1)$ that firm 1 did not exit, 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 stayed beyond $T(b_1)$, 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. ■

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

We first show that if a pure strategy for firm 2 is only weakly dominated in Round 1 (Claim 1 (b)) it cannot be played with positive probability.

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

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

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

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 A.1). On the other hand, if $s_2 = b_2$, it is no better.

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

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

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 $(\tilde{\tau}, \sigma)$ is not a Nash equilibrium. ■

Now we claim that if a pure strategy for firm 1 is only weakly dominated in Round 1 (Claim 1 (a)) it cannot be played with positive probability either.

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

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

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

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 $\Pr[\tilde{\tau}_2(g_2) \geq T(b_1, g_2)] = 1$ (Claim 8) and if $s_2 = b_2$, it is no better. Thus, $\Pr[\tilde{\tau}_1(g_1) \leq T(b_1, g_2)] = 0$.

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

First, note that $\Pr[\tilde{\tau}_1(b_1) > T(b_1, g_2)] = 0$ as well. This is because from Claim 8, $\Pr[\tilde{\tau}_2(g_2) \geq T(b_1, g_2)] = 1$ 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 $\Pr[\tilde{\tau}_1(b_1) > T(b_1, g_2)] = 0$ and $\Pr[\tilde{\tau}_1(g_1) \leq T(b_1, g_2)] = 0$ (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: $\Pr[\tilde{\tau}_2(g_2) \geq \tilde{\tau}_1(g_1) \text{ and } T(b_1, g_2) < \tilde{\tau}_1(g_1) < T(g_1, b_2)] > 0$.

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

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

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

From Claim 9, $\Pr[\tilde{\tau}_1(g_1) \geq T(g_1, b_2)] = 1$ and from Claim 4 $\Pr[\tilde{\tau}_1(b_1) = T(b_1)] = 1$. This means that if firm 1 is active at any time $t > T(b_1)$, 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: $\Pr[T(g_1, b_2) \leq \tilde{\tau}_1(g_1) < \tilde{\tau}_2(g_2) < T(g_1, g_2)] > 0$.

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

The last claim shows that if $(\tilde{\tau}, \sigma)$ is a Nash equilibrium, the probability that a pure strategy in the support of $\tilde{\tau}_2(g_2)$ 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.

B Appendix: Value of information

In this appendix, we formally establish that each firm's expected payoff in the upstart equilibrium is strictly increasing in the quality of its own information (Proposition 4). In other words, for each firm the value of its private information is positive. In fact, it turns out that each firm's payoff is also increasing in the quality of its rival's information.

Firm 1 equilibrium expected payoff can be written as

$$\begin{aligned} \Pi_1^* &= \Pr[g_1, g_2] \times v(p(g_1, g_2)) + \Pr[g_1, b_2] \times v(p(g_1, b_2)) \\ &\quad + \Pr[b_1] \times v(p(b_1)) \end{aligned}$$

where $v(p_0)$ is the flow profit of a firm when both firms hold a common belief p_0 at time 0 and *both* are in the race until this belief reaches p^* . For the usual reasons, v is an increasing and convex function and is strictly increasing and strictly convex when $p_0 > p^*$ (see the online appendix for a proof of these properties). This follows from the fact that, in the upstart equilibrium, when the signals are (g_1, g_2) both firms exit at time $T(g_1, g_2)$ and this is exactly the optimal stopping time given an initial common belief $p_0 = p(g_1, g_2)$. The same is true when the signals are (g_1, b_2) . When 1's signal is b_1 , it exits at $T(b_1)$ and firm 2 does not exit before this time. Again, $T(b_1)$ is the optimal stopping given an initial common belief $p_0 = p(b_1)$.

For a fixed q_1 , define the probability distribution function $F : [0, 1] \rightarrow [0, 1]$

$$F(x) = \begin{cases} 0 & 0 \leq x < p(b_1) \\ \Pr[b_1] & p(b_1) \leq x < p(g_1, b_2) \\ \Pr[b_1] + \Pr[g_1, b_2] & p(g_1, b_2) \leq x < p(g_1, g_2) \\ 1 & p(g_1, g_2) \leq x < 1 \end{cases}$$

and similarly, for $\hat{q}_1 > q_1$ define $\hat{F} : [0, 1] \rightarrow [0, 1]$ analogously. It may be readily confirmed that \hat{F} is a mean-preserving spread of F . Since v is a convex function, it follows that firm 1's equilibrium payoff when his signal quality is \hat{q}_1 , $\hat{\Pi}_1^* > \Pi_1^*$.

Firm 2's expected payoff can be written as

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

where $v(p_0)$ is defined as above and $u(p_0)$ is analogously defined as the flow profit of a firm when with belief p_0 at time 0 and it is the *only* firm present. The first three terms are the same as in Π_1^* . The only additional term occurs when the signals are (b_1, g_2) where firm 2 now stays longer than firm 1. In this case, firm 2's belief at time $T(b_1)$ is $p_{T(b_1)}(b_1, g_2)$ and the resulting payoff is discounted back to time 0.

To establish that Π_2^* is an increasing function of q_2 , 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. The third term does not depend on q_2 .

Lemma B.1 $\Pr[g_1, g_2] \times v(p(g_1, g_2)) + \Pr[g_1, b_2] \times v(p(g_1, b_2))$ is increasing in q_2 .

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

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

is increasing in q_2 .

Now if $\hat{q}_2 > q_2$, then

$$\hat{p}(g_1, b_2) < p(g_1, b_2) < p(g_1, g_2) < \hat{p}(g_1, g_2)$$

where $\widehat{p}(g_1, \cdot)$ denotes the posterior derived from \widehat{q}_2 . Moreover, the mean of $p(g_1, \cdot)$ is $p(g_1)$ and this is the same as the mean of $\widehat{p}(g_1, \cdot)$ (since the expectation of the posteriors is the prior). Thus, the distribution of $\widehat{p}(g_1, \cdot)$ is a mean-preserving spread of the distribution of $p(g_1, \cdot)$.

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

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

Lemma B.2 $\Pr[b_1, g_2] \times u(p_{T(b_1)}(b_1, g_2))$ is increasing in q_2 .

Proof.

$$\begin{aligned} & \frac{\partial (\Pr[b_1, g_2] \times u(p_{T(b_1)}(b_1, g_2)))}{\partial q_2} \\ &= \frac{\partial \Pr[b_1, g_2]}{\partial q_2} u(p_{T(b_1)}(b_1, g_2)) + \Pr[b_1, g_2] \frac{\partial p_{T(b_1)}(b_1, g_2)}{\partial q_2} u'(p_{T(b_1)}(b_1, g_2)) \\ &> \frac{\partial \Pr[b_1, g_2]}{\partial q_2} u(p_{T(b_1)}(b_1, g_2)) + \Pr[b_1, g_2] \frac{\partial p_{T(b_1)}(b_1, g_2)}{\partial q_2} \frac{u(p_{T(b_1)}(b_1, g_2))}{p_{T(b_1)}(b_1, g_2)} \end{aligned}$$

since u is an increasing and convex function that is non-negative and strictly positive for $p > p^*$ and $u(0) = 0$ (see the online Appendix for a detailed proof). Thus, $u'(p) > \frac{1}{p}u(p)$. The sign of the right-hand side of the inequality is the same as the sign of

$$\begin{aligned} & \frac{\partial \Pr[b_1, g_2]}{\partial q_2} p_{T(b_1)}(b_1, g_2) + \Pr[b_1, g_2] \frac{\partial p_{T(b_1)}(b_1, g_2)}{\partial q_2} \\ &= \frac{\partial}{\partial q_2} (\Pr[b_1, g_2] \times p_{T(b_1)}(b_1, g_2)) \\ &= \frac{\partial}{\partial q_2} \left(\Pr[b_1, g_2] \times \frac{e^{-2\lambda T(b_1)} \pi (1 - q_1) q_2}{e^{-2\lambda T(b_1)} \pi (1 - q_1) q_2 + (1 - \pi) q_1 (1 - q_2)} \right) \end{aligned}$$

Dividing the numerator and denominator of the second term by $\Pr[b_1, g_2] = \pi (1 - q_1) q_2 + (1 - \pi) q_1 (1 - q_2)$, we obtain

$$\begin{aligned} \frac{\partial}{\partial q_2} (\Pr[b_1, g_2] \times p_{T(b_1)}(b_1, g_2)) &= \frac{\partial}{\partial q_2} \left(\Pr[b_1, g_2] \times \frac{e^{-2\lambda T(b_1)} p(b_1, g_2)}{e^{-2\lambda T(b_1)} p(b_1, g_2) + 1 - p(b_1, g_2)} \right) \\ &= \frac{\partial}{\partial q_2} \left(\frac{e^{-2\lambda T(b_1)} \Pr[G, b_1, g_2]}{1 - (1 - e^{-2\lambda T(b_1)}) p(b_1, g_2)} \right) \end{aligned}$$

and since both $\Pr[G, b_1, g_2]$ and $p(b_1, g_2)$ are increasing in q_2 , we have that $\Pr[b_1, g_2] \times p_{T(b_1)}(b_1, g_2)$ is increasing in q_2 as well. ■

This completes the proof of Proposition 4.

C Appendix: Planner's solution

In this appendix we formally show that it is optimal for a planner with initial belief p_0 to ensure that both firms are active until $T(p_0)$. This is the key step in establishing Proposition 6.

The per-firm expected flow profit from switching from two firms to one firm at time s is

$$w(s) = \lambda m \int_0^s e^{-rt} (e^{-2\lambda t} p_0 + 1 - p_0) (p_t - p^*) dt \\ + \frac{1}{2} \lambda m \int_s^{2T-s} e^{-rt} (e^{-\lambda(s+t)} p_0 + 1 - p_0) (p_t - p^*) dt$$

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

$$\frac{p_t}{1 - p_t} = \begin{cases} e^{-2\lambda t} \frac{p_0}{1 - p_0} & \text{if } t \leq s \\ e^{-\lambda(s+t)} \frac{p_0}{1 - p_0} & \text{if } t \geq s \end{cases} \quad (8)$$

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 (8), $w(s)$ can be explicitly calculated to be

$$w(s) = \frac{\lambda m p_0 (1 - p^*)}{2\lambda + r} ((2\lambda + r) e^{-2T\lambda} (e^{-rs} - 1) - r (e^{-(2\lambda+r)s} - 1)) \\ + \frac{\lambda m p_0 (1 - p^*)}{2(\lambda + r)} (\lambda e^{-2\lambda T} (e^{-r(2T-s)} - e^{-rs}) - r e^{-rs} (e^{-2\lambda T} - e^{-2\lambda s}))$$

Differentiating with respect to s then yields

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

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

$$w''(s) \\ = \lambda m p_0 (1 - p^*) \times \frac{r^2 e^{rs} (\lambda e^{-2(\lambda+r)T} + (\lambda + r) e^{-2(\lambda T+rs)} - (2\lambda + r) e^{-2(\lambda+r)s})}{2(\lambda + r)} \\ < \lambda m p_0 (1 - p^*) \times \frac{r^2 e^{rs} (\lambda e^{-2(\lambda+r)s} + (\lambda + r) e^{-2(\lambda s+rs)} - (2\lambda + r) e^{-2(\lambda+r)s})}{2(\lambda + r)} \\ = 0$$

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 that 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 $p_t > p^*$ or alternatively, until $T(p_0)$.

References

- [1] Akcigit, U. and W. Kerr (2018): "Growth through Heterogeneous Innovations," *Journal of Political Economy*, 126 (4), 1374–1443.
- [2] 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.
- [3] 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.
- [4] Benoît, J. P., & Dubra, J. (2011): "Apparent Overconfidence," *Econometrica*, 79 (5), 1591–1625.
- [5] 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.
- [6] 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>
- [7] Choi, J. (1991): "Dynamic R&D Competition under 'Hazard Rate' Uncertainty," *Rand Journal of Economics*, 22 (4), 596–610.
- [8] Christiansen, C. (1997): *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*, Harvard Business School Press, Boston, MA.
- [9] Das, K. and N. Klein (2020): "Do Stronger Patents Lead to Faster Innovation? The Effect of Duplicative Search," Department of Economics, University of Exeter, <https://drive.google.com/file/d/184aiXJOXhUQpQSMKiUVsP0QdUVsPzRVv/view>
- [10] Dong, M. (2018): "Strategic Experimentation with Asymmetric Information," Penn State University, <https://sites.google.com/site/miaomiaodong0/1-delimit>
- [11] Gilbert, R. and D. Newbery (1982): "Preemptive Patenting and the Persistence of Monopoly," *American Economic Review*, 76, 238–242.
- [12] Hirshleifer, J. (1971): "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review*, 61, 561–574.

- [13] 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.
- [14] Keller, G., S. Rady and M. Cripps (2005): "Strategic Experimentation with Exponential Bandits," *Econometrica*, 73, 39–68.
- [15] Klein, N. and P. Wagner (2018): "Strategic Investment and Learning with Private Information," University of York, <https://sites.google.com/site/peterachimwagner/research>
- [16] Krieger, J. (2019): "Trials and Terminations: Learning from Competitors' R&D Failure," Harvard Business School.
- [17] Maschler, M., E. Solan and S. Zamir (2013): *Game Theory*, New York: Cambridge University Press.
- [18] Malueg, D. and S. Tsutsui (1997): "Dynamic R&D Competition with Learning," *Rand Journal of Economics*, 28, 751–772.
- [19] Moscarini, G. and F. Squintani (2010): "Competitive Experimentation with Private Information: The Survivor's Curse," *Journal of Economic Theory*, 145, 639–660.
- [20] 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>.
- [21] Schumpeter, J. (1942): *Capitalism, Socialism and Democracy*, New York: Harper.
- [22] Shimoji, M. and J. Watson (1998): "Conditional Dominance, Rationalizability and Game Forms," *Journal of Economic Theory*, 83, 161–195.
- [23] Surowiecki, J. (2014): "Epic Fails of the Startup World," *The New Yorker*, May 19.
- [24] Wong, T-N. (2018); "Free Riding and Duplication in R&D," University of Basel, <https://sites.google.com/site/tzszningwongecon/research>
- [25] Wu, B., and A. Knott (2006): "Entrepreneurial Risk and Market Entry," *Management Science*, 52 (9), 1315–1330.