

The Last Shall Be First: Innovation as a Head-to-Head Race*

Patrick Arnold[†] Marc Möller[‡] Catherine Roux[§]

Abstract

Uncertainty about the value of a contested innovation induces leaders and laggards to update their expectations in opposite directions. We characterize situations in which firms that have obtained an initial advantage are not the most likely to achieve final success. In spite of amplifying a leader's advantage, greater contest intensity facilitates this effect, challenging the view that laggards require support to remain competitive.

Keywords: Innovation contests; learning; competitive balance; leapfrogging.
JEL classification: C72, D82.

1. Introduction

While Neil Armstrong's iconic words—“*That's one small step for man, one giant leap for mankind*”—are etched into history, it is easily overlooked that only a few years earlier, the USSR had established a significant lead in the Space Race. After the successful launch of Vostok, which made Yuri Gagarin the first human to orbit our planet, President Kennedy acknowledged that “[...] *the Soviet Union gained an important advantage by securing these large boosters*” (Logsdon 2010). However, instead of discouraging U.S. efforts, the Vostok

*This work was supported by the Swiss National Science Foundation [grant number 100018_207874].

[†]University of Basel. E-mail: patrick.arnold@unibas.ch

[‡]University of Bern. E-mail: marc.moeller@unibe.ch, corresponding author

[§]University of Basel. E-mail: catherine.roux@unibas.ch

marked a turning point in Kennedy’s support of the space program. After years of downplaying its importance, Kennedy made NASA the largest civilian agency in history which helped the U.S. secure ultimate victory in the race to the Moon.¹

Examples where an innovator’s success sparks the R&D efforts of its rivals are ubiquitous, ranging from the design of twin aircraft-engines by Boeing and Douglas (Sutton 1998) and commercial mainframe computers by Remington Rand and IBM (Khanna 1995), to the more recent development of touchscreen phones by Apple and Samsung or conversational artificial intelligence by OpenAI and Google. Lerner (1997) provides further historical evidence suggesting that, in technology-intensive industries, catch-up behavior is the norm and firms deliberately increase their R&D expenditures when trailing behind. This observation contrasts with the common concern that a laggard’s incentive to innovate might be muted by a rival’s lead, even if that lead is small (ϵ -preemption) and many steps remain to be completed (discouragement effect). In this paper, we argue that it is an innovator’s ability to observe a rival’s outputs (e.g, technological advancements, patents) but not inputs (e.g., R&D spending, efforts), that induces the intense competition for innovation we often observe.

Our argument is based on the idea that, in a race for an innovation of uncertain value, leading serves as a signal for a large investment from a leader motivated by high expectations, whereas lagging indicates a small investment by a laggard, likely to be the consequence of a more pessimistic view. Leaders and laggards thus update their beliefs in opposite directions, and a laggard can have a stronger incentive to invest, despite a leader’s initial technological advantage. We show that, as a result, innovation can become a *head-to-head race*, in which laggards are equally or even more likely than initial leaders to achieve final success.

Our main contribution is to characterize the environments most conducive to making innovation a head-to-head race. Such environments offer sufficient rewards for intermediate technological advances, for example, in the form of easily filed patents, and they eliminate noise from the determination of the final winner. Perhaps surprisingly, and contrasting with existing results (Wang 2010), noise is not suitable for a head-to-head race, because although noise makes a leader’s technological advantage less decisive, it also obstructs the informational effect that induces a laggard to catch-up. A policy implication is that innovation contests should reward contestants for intermediate advances and offer clear-cut specification of the desired product.²

Our findings challenge a long-held view that laggards require external support to maintain competitive balance in innovative industries (Fudenberg et al. 1983; Harris and Vickers

1. At its peak, NASA employed up to 400,000 people and its spending increased to 4.4% of federal budget (Kantor and Whalley 2025).

2. An example meeting these requirement is the XPRIZE Carbon Removal Competition which offered \$1million to the 15 most advanced contestants after the first year and a thirty page competition guideline specifying how to measure each contestant’s performance with great precision.

1987; Lippman and McCardle 1987). This view is in line with the more general idea of “leveling the playing field” in competitive environments and the associated call for affirmative action (Möller 2012; Chowdhury, Esteve-González, and Mukherjee 2023). We provide a countervailing argument by showing that innovation races may be more likely to be won by laggards than by leaders, even without external support, suggesting the existence of a *self-correcting mechanism*. This mechanism complements leapfrogging (Fudenberg et al. 1983; Aoki 1991), the opportunity for a laggard to overtake a leader through a sufficiently large technological jump, in securing competitive balance. Rather than being strategic, our mechanism is informational. We thereby contribute to a more recent literature highlighting the relevance of learning effects in dynamic innovation contests (Choi 1991; Doraszelski 2003; Halac, Kartik, and Liu 2017; Awaya and Krishna 2021).³ A distinguishing feature of our setting, crucial to the insight that learning improves competitive balance, is that contest outcomes induce leaders and laggards to update their beliefs in opposite directions. Although this feature is present in Möller and Beccuti (2025), a head-to-head race cannot arise. In their best-of-N-battle contest, a laggard might invest more than a leader but is never more likely to win.

2. The model

Two homogeneous, risk-neutral players engage in a two-stage innovation contest that awards a prize $V_t \geq 0$ to the winner of stage $t \in \{1, 2\}$. We let $\nu = \frac{V_1}{V_2}$ denote the value-ratio of winning the initial stage (being the first in space) relative to winning the final stage (being the first on the Moon). The probability that player $i \in \{A, B\}$ wins stage t is given by the generalized Tullock (1980) contest success function:

$$p_{i,t}(e_{i,t}, e_{j,t}) = \frac{e_{i,t}^r}{e_{i,t}^r + e_{j,t}^r}, \quad (1)$$

where $e_{i,t}$ denotes player i 's R&D expenditure, or *effort*, in stage t and $r \in [0, 1]$ is a parameter determining the sensitivity of the contest outcome with respect to players' efforts as opposed to *noise*. We denote the first-stage winner and the first-stage loser as the *leader* (L) and the *follower* (F), respectively.

Assumption 1. *The leader wins stage 2 with probability $p_{L,2}(\frac{1}{\alpha}e_{L,2}, e_{F,2})$, where $\frac{1}{\alpha} > 1$ denotes a technological advantage the leader obtains by winning stage 1.*

3. Learning effects also play a role in a related literature on entry- and exit-dynamics in competitive industries (e.g., Awaya and Krishna 2021; Chen, Ishida, and Mukherjee 2023; Cetemen and Margaria 2024).

A key feature of our model is that players are privately informed about the uncertain value of innovation. Following Möller and Beccuti (2025), we model this uncertainty by postulating the existence of a “good” state $S = g$ and a “bad” state $S = b$, equally likely and such that $V_1, V_2 > 0$ if and only if $S = g$, and we make the following assumption:

Assumption 2. *Each player i privately observes a signal $s_i \in \{B, G\}$. Players’ signals are independent with $\mathbb{P}(s_i = G | S = g) = 1$ and $\mathbb{P}(s_i = B | S = b) = \sigma \in (0, 1)$.*

Note that signals are informative and that a bad signal is even conclusive, since $s_i = B$ can only be received if $S = b$. Together with the fact that innovation has no value in the bad state, this assumption greatly simplifies the analysis because it implies that upon the observation of $s_i = B$, player i will choose $e_{i,1} = e_{i,2} = 0$.⁴

In addition to their private signals, players observe contest outcomes, i.e., who won and who lost each stage, but cannot observe their rival’s efforts. This assumption is important as it implies that players will learn about their rival’s signals only through their observation of contest outcomes. All uncertainty is resolved at the end of stage 2 when values V_1 and V_2 materialize.⁵

Player i ’s payoff is given by the difference between his prize earnings and his aggregate R&D expenditure, $e_{i,1} + e_{i,2}$, i.e., players’ costs of “effort” are linear, and identical across players and periods. Our setting constitutes a dynamic Bayesian game, with players’ *types* given by their signals. We thus use Perfect Bayesian equilibrium as our solution concept. As players choose zero effort conditional on a bad signal, a symmetric equilibrium is completely characterized by a triplet (e_1^*, e_L^*, e_F^*) denoting players’ first-stage effort, leader’s effort, and follower’s effort, all upon receipt of a good signal.

3. Learning

How do players learn about their rival’s signal? Consider a player who received a good signal. Let β denote the player’s prior belief that his rival’s signal is also good:

$$\beta \equiv \mathbb{P}(s_j = G \mid s_i = G) = \frac{\mathbb{P}(s_i = s_j = G)}{\mathbb{P}(s_i = G)} = \frac{1 + (1 - \sigma)^2}{2 - \sigma} \in (0, 1). \quad (2)$$

After observing the stage 1 outcome, the player updates this belief as follows. After losing, the player concludes that his rival exerted positive effort in stage 1 which only happens

4. By the continuity of players’ payoffs, our results remain valid when a bad signal is *almost* conclusive, i.e., when $\mathbb{P}(s_i = G | S = g)$ is sufficiently close to one.

5. In an extension, we allow V_1 and V_2 to be correlated and assume players learn the value of V_1 before stage 2. Results remain qualitatively unchanged unless prizes are perfectly correlated.

after the rival received a good signal. Hence, the follower updates his belief upwards to $\beta_F^* = 1 > \beta$. In contrast, after winning stage 1, the leader updates his belief downwards to

$$\beta_L^* = \frac{\frac{1}{2}\beta}{\frac{1}{2}\beta + 1 - \beta} < \beta, \quad (3)$$

because, in equilibrium, he would have defeated a rival with a good signal only with probability one half. Note that innovation is expected to have positive value only when both players' signals are good and that this value is

$$V_t^G \equiv \mathbb{E}[V \mid s_i = s_j = G] = \frac{V_t}{1 + (1 - \sigma)^2}. \quad (4)$$

It thus follows from $\beta_F^* > \beta > \beta_L^*$ that learning induces players' expectations about the value of innovation to move in opposite directions. In particular, learning induces the follower's expectation about the value of innovation to be higher than the leader's: $\beta_F^* V_t^G > \beta_L^* V_t^G$.

4. Head-to-head races

With the aim to identify learning about the uncertain value of innovation as a self-correcting mechanism, this section derives the conditions under which innovation becomes a head-to-head race. We formalize the notion of a head-to-head race as follows:

Definition 1. *Innovation constitutes a head-to-head race if, conditional on the value of innovation being positive, the follower's probability of winning the race is at least as high as the leader's:*

$$p_F^* \equiv \frac{(\alpha e_F^*)^r}{(e_L^*)^r + (\alpha e_F^*)^r} \geq \frac{(e_L^*)^r}{(e_L^*)^r + (\alpha e_F^*)^r} \equiv p_L^*. \quad (5)$$

The fact that we condition on the innovation having positive value captures the idea that a “winner” of an innovation race would not be observed as such if he had obtained an innovation with no value.

To derive the conditions under which (5) holds, we now characterize the unique candidate for a symmetric pure-strategy equilibrium (e_1^*, e_L^*, e_F^*) .⁶ In stage 2 it has to hold that

$$e_F^* \in \arg \max_{e_F \geq 0} \frac{(\alpha e_F)^r}{(e_L^*)^r + (\alpha e_F)^r} \beta_F^* V_2^G - e_F \quad \text{and} \quad e_L^* \in \arg \max_{e_L \geq 0} \frac{e_L^r}{e_L^r + (\alpha e_F^*)^r} \beta_L^* V_2^G - e_L. \quad (6)$$

6. Equilibrium existence can be shown along the lines of Möller and Beccuti (2025).

The unique effort levels satisfying these conditions are

$$e_L^* = \frac{V_2^G \beta_L^* r (\alpha \beta_L^*)^r}{(\alpha^r + (\beta_L^*)^r)^2} \quad \text{and} \quad e_F^* = \frac{V_2^G r (\alpha \beta_L^*)^r}{(\alpha^r + (\beta_L^*)^r)^2}, \quad (7)$$

and they imply that

$$\frac{p_F^*}{p_L^*} = \frac{\alpha^r}{(\beta_L^*)^r} \quad \text{and} \quad p_F^* \geq p_L^* \Leftrightarrow \beta_L^* \leq \alpha. \quad (8)$$

This means that for a head-to-head race to exist, learning has to be sufficiently strong relative to the leader's technological advantage. However, condition (8) is only necessary but not sufficient for a head-to-head race to exist. Remember from Section 3 that learning requires that players condition their stage 1 efforts on their private information. More specifically, in stage 1, a player should exert positive effort if and only if the player has observed a good signal.

To determine the condition under which $e_1^* > 0$, let U_L^* and U_F^* denote the players' continuation values of entering stage 2 as the leader or the follower, conditional on having received a good signal. We have

$$U_L^* - U_F^* = \frac{(\beta_L^*)^{2r+1} - \alpha^{2r} - (1 - \beta_L^*)(\alpha \beta_L^*)^r(1 - r)}{[\alpha^r + (\beta_L^*)^r]^2} V_2^G, \quad (9)$$

and the equilibrium effort in stage 1 has to satisfy

$$e_1^* \in \arg \max_{e_1 \geq 0} \beta \frac{e_1^r}{e_1^r + (e_1^*)^r} (U_L^* - U_F^* + V_1^G) - e_1 \implies e_1^* = \frac{r}{4} \beta (U_L^* - U_F^* + V_1^G). \quad (10)$$

Note that, off the equilibrium path, the leader's continuation value U_L depends on stage 1 efforts through the updated belief $\beta_L = \frac{e_1 \beta}{e_1 \beta + (1 - \beta)(e_1 + e_1^*)}$, but (10) can be obtained easily by use of the Envelope Theorem.⁷ Also note that if $p_F^* \geq p_L^*$, then (8) implies $U_L^* - U_F^* \leq 0$. Hence, for innovation to be a head-to-head race, the prize for winning stage 1 has to be sufficiently large. Only when $V_1^G > U_F^* - U_L^*$, players exert positive effort in stage 1 conditional on having received a good signal, making the contest's stage 1 outcome informative. We are now ready to formulate our result:

7. For a deviation to $e_1 = 0$, Bayesian updating leads to different beliefs but such a deviation can be shown to be dominated by a deviation to $e_1 = \epsilon$ for $\epsilon > 0$ sufficiently small.

Proposition 1. *There exist a non-empty interval $(\underline{\sigma}, \bar{\sigma}) \subset (0, 1)$ of information structures for which learning induces innovation to be a head-to-head race if and only if*

$$\frac{1}{\alpha} \leq \sqrt{2} \quad \text{and} \quad 1 + \frac{1}{\nu} = \frac{V_1 + V_2}{V_1} < \frac{4}{(1 - \alpha)(2 - r)} + 1, \quad (11)$$

that is, when both the technological advantage of winning the initial stage and the fraction of innovation value reserved to the winner of the final stage are not too large.

Proposition 1 characterizes the conditions under which innovation becomes a head-to-head race. Learning effects have to be sufficiently strong to overcome the leader's technological advantage. This requires the players' information to be neither too precise nor too imprecise, because for $\sigma \rightarrow 0$ or $\sigma \rightarrow 1$, the rival's signal adds no additional information to a player's own signal. However, in an intermediate range of σ , losing stage 1 induces the follower to exert more effort than the leader in stage 2, eventually overcoming the leader's technological advantage and equalizing both players' chances to claim final victory. Proposition 1 shows that this requires players to be able to secure a sufficiently large part of the value of innovation *upfront*. It suggests that environments, where inexpensive and rapidly filed patents for intermediate technological advances are available, are conducive to head-to-head competition.

How does noise influence whether innovation is head-to-head? Our model sheds light on this question, because an increase in the parameter r of the contest success function makes contest outcomes less noisy and more sensitive to players' efforts. Intuitively, one might expect that increasing r tilts the race towards the leader by making his technological advantage more relevant. However, this intuition neglects the fact that an increase in r also raises players' incentives to provide effort in the initial stage of the contest, which, as argued above, constitutes a prerequisite for learning. The following result shows that this indirect informational effect is strong enough to overcome the direct technological effect, making a reduction in noise conducive for head-to-head competition.

Corollary 1. *Reducing the noisiness with which the innovation contest determines its winners is suitable for inducing a head-to-head race, i.e., an increase in r relaxes (11).*

In innovation contests, intermediate and final winners are determined in accordance with evaluation criteria specified by the contest's organizer. Corollary 1 suggests that such criteria should leave little indeterminacy to promote the head-to-head nature of innovation.

5. Concluding remark

Admittedly, the present model is stylized, but it enhances our understanding of a basic learning mechanism through which innovation can become a head-to-head race. Our main result

highlights that, to promote a balanced race, technological advances should be sufficiently credited and attributed with minimal error. While the basic implications for patent policy are straightforward, future research should extend the model to explore how learning varies with the number of competitors and the length of the time horizon.

Appendix

Proof of Proposition 1

Innovation is a head-to-head race if and only if $e_1^* > 0$ and e_F^* and e_L^* are such that $p_F^* \geq p_L^*$. Formally

$$p_F^* \geq p_L^* \Leftrightarrow \beta_L^* \leq \alpha \quad (12)$$

$$e_1^* > 0 \Leftrightarrow \frac{V_1}{V_2} > \frac{-(\beta_L^*)^{2r+1} + \alpha^{2r} + (1 - \beta_L^*)(\alpha\beta_L^*)^r(1 - r)}{(\alpha^r + (\beta_L^*)^r)^2}. \quad (13)$$

Note that β_L^* is a U -shaped function of $\sigma \in (0, 1)$, taking the minimum value $\frac{1}{\sqrt{2}}$. Hence, for $\alpha \geq \frac{1}{\sqrt{2}}$ there exists a non-empty interval such that (12) holds for all σ in that interval, whereas for $\alpha < \frac{1}{\sqrt{2}}$ (12) fails to hold for any σ . Further note that the RHS of (13) is strictly decreasing in β_L^* . Hence, amongst all β_L^* that satisfy (12), (13) becomes easiest to satisfy for $\beta_L^* = \alpha$. Substitution of $\beta_L^* = \alpha$ into (13) thus determines a lower bound

$$\underline{\nu} \equiv \frac{(1 - \alpha)(2 - r)}{4} \quad (14)$$

on the value ratio $\nu = \frac{V_1}{V_2}$ above which β_L^* satisfies both (12) and (13) for *some* $\sigma \in (0, 1)$. ■

References

- Aoki, Reiko. 1991. “R&D competition for product innovation: An endless race.” *The American Economic Review* 81 (2): 252–256.
- Awaya, Yu, and Vijay Krishna. 2021. “Startups and upstarts: disadvantageous information in R&D.” *Journal of Political Economy* 129 (2): 534–569.
- Cetemen, Doruk, and Chiara Margaria. 2024. “Exit dilemma: the role of private learning on firm survival.” *American Economic Journal: Microeconomics* 16 (1): 110–154.

- Chen, Chia-Hui, Junichiro Ishida, and Arijit Mukherjee. 2023. "Pioneer, early follower or late entrant: Entry dynamics with learning and market competition." *European Economic Review* 152:104360.
- Choi, Jay. 1991. "Dynamic R&D competition under" hazard rate" uncertainty." *The RAND Journal of Economics*, 596–610.
- Chowdhury, Subhasish, Patricia Esteve-González, and Anwesha Mukherjee. 2023. "Heterogeneity, leveling the playing field, and affirmative action in contests." *Southern Economic Journal* 89 (3): 924–974.
- Doraszelski, Ulrich. 2003. "An R&D race with knowledge accumulation." *The RAND Journal of Economics*, 20–42.
- Fudenberg, Drew, Richard Gilbert, Joseph Stiglitz, and Jean Tirole. 1983. "Preemption, leapfrogging and competition in patent races." *European Economic Review* 22 (1): 3–31.
- Halac, Marina, Navin Kartik, and Qingmin Liu. 2017. "Contests for Innovation." *Journal of Political Economy* 125 (5): 1523–1569.
- Harris, Christopher, and John Vickers. 1987. "Racing with uncertainty." *The Review of Economic Studies* 54 (1): 1–21.
- Kantor, Shawn, and Alexander Whalley. 2025. "Moonshot: Public R&D and Growth." *The American Economic Review* 115 (9): 2891–2925.
- Khanna, Tarun. 1995. "Racing behavior technological evolution in the high-end computer industry." *Research Policy* 24 (6): 933–958.
- Lerner, Josh. 1997. "An empirical exploration of a technology race." *The RAND Journal of Economics*, 228–247.
- Lippman, Steven A, and Kevin F McCardle. 1987. "Dropout behavior in R&D races with learning." *The RAND Journal of Economics*, 287–295.
- Logsdon, John M. 2010. "John F. Kennedy and the race to the Moon." Springer.
- Möller, Marc. 2012. "Incentives versus competitive balance." *Economics Letters* 117 (2): 505–508.
- Möller, Marc, and Juan Beccuti. 2025. "Fighting for lemons: The balancing effect of private information on incentives in dynamic contests." *The Economic Journal* 135 (669): 1641–1676.

- Sutton, John. 1998. *Technology and market structure: Theory and history*. MIT press.
- Tullock, Gordon. 1980. "Efficient rent seeking." In *J. Buchanan, R. Tollison, and G. Tullock, eds., Toward a theory of the rent seeking society*. Texas A&M Press.
- Wang, Zhewei. 2010. "The optimal accuracy level in asymmetric contests." *The BE Journal of Theoretical Economics* 10 (1).