Understanding Perpetual R&D Races

by Yves Breitmoser, Jonathan H.W. Tan and Daniel J. Zizzo

Electronic Appendices

This file contains five electronic appendices:

- Electronic Appendix A: relegated proofs and the computation of equilibria for experimental games;
- Electronic Appendix B: experimental instructions and experimental screenshots;
- Electronic Appendix C: further details on R&D monopoly entrenchment;
- Electronic Appendix D: payoff transformations and rivalry motive;
- Electronic Appendix E: further details on Markov quantal response equilibria.

Details on the references cited in the electronic appendices are provided at the end of this document.

Appendix A – Relegated Proofs and Computation of Equilibria

A1. Relegated Proofs

Lemma 1 *Assume* $\alpha_H, \alpha_L \in (0,1)$ *and there exists a* $k > 0$ *such that* $\tau_2(k') = 0 \, \forall k' \geq k$. *Then, for all* τ_1, $V_2(\tau_1, \tau_2 \mid k')$ *is increasing in* k' *for all* $k' \geq k$.

Proof: Fix a state $k' \geq k$. Let ρ_t denote the probability of the following event under the strategy profile (τ_1, τ_2): assuming the current state is k', the state in the next $t-1$ rounds will be some $k'' \geq k'$ and the state in round t (counted from now) will be $k'-1$. Define $\rho_\infty := \Sigma_{t>0}\rho_t$. Thus, the valuation $V_2(k')$ under (τ_1, τ_2) satisfies

$$V_2(k') = (1-\delta)R + \delta\left((1-\rho_\infty)R + \sum_{t>0} \rho_t * \left((1-\delta^{t-1})R + \delta^{t-1} V_2(k'-1) \right) \right)$$

Clearly, $\rho_t > 0 \, \forall t$ and $V_2(k'-1) < R$. As a result, we have $V_2(k') > (1-\delta)R + \delta(1-\rho_\infty)R + \delta V_2(k'-1) > V_2(k'-1)$ for all $k' \geq k$. QED

Lemma 2 *Assume* $\alpha_H, \alpha_L \in (0,1)$ *and there exists a* $k < 0$ *such that* $\tau_2(k') = 0 \, \forall k' \le k$ *. Then, for all* τ_1*, the following holds: if* $V_2(\tau_1, \tau_2 \mid k+1) > -R$*, then* $V_2(\tau_1, \tau_2 \mid k')$ *is increasing in* k' *for all* $k' \le k$*, otherwise it is decreasing in* k' *for all* $k' \le k$*.*

Proof: Fix a state $k' \le k$. As in the proof of Lemma 1, let ρ_t denote the probability of the following event under the strategy profile (τ_1, τ_2): assuming the current state is k', the state in the next $t-1$ rounds will be some $k'' \le k'$ and the state in round t will be $k'+1$, and define $\rho_\infty := \Sigma_{t>0} \rho_t$. Thus, the valuation $V_2(k')$ under (τ_1, τ_2) satisfies

$$V_2(k') = -(1-\delta) R + \delta\left(-(1-\rho_\infty) R + \sum_{t>0} \rho_t * \left(-\left(1-\delta^{t-1}\right) R + \delta^{t-1} V_2(k'+1) \right) \right)$$

Assume first $V_2(k+1) > -R$. As a result, we have $V_2(k) < -(1-\delta)R - \delta(1-\rho_\infty)R + \delta V_2(k+1) < V_2(k+1)$, and also $V_2(k) > -R$. Iteratively, this implies that, for all $k' \le k$, $V_2(k'+1) > V_2(k') > -R$ holds, i.e. $V_2(k')$ is increasing. Secondly, consider the case where $V_2(k+1) < -R$. We now have $V_2(k) > -(1-\delta)R - \delta(1-\rho_\infty)R + \delta V_2(k+1) > V_2(k+1)$ and $V_2(k) < -R$ are implied, and iteratively applied, this shows that $V_2(k')$ is decreasing in k'. QED

Proposition 1 *Fix a strategy profile* (τ_1, τ_2) *and states* $\underline{K} < \overline{K}$ *such that 2 exerts high effort only in states* k *satisfying* $\underline{K} \le k \le \overline{K}$*. Upper boundaries* $\overline{V}_2(k)$ *of the valuation function of player 2 satisfy the following equation system.*

$$\begin{aligned}
\overline{V}_2\left(\overline{K}\right) &= m_{\overline{K}}^{+1} \delta R + m_{\overline{K}}^{0} \delta \overline{V}_2\left(\overline{K}\right) + m_{\overline{K}}^{-1} \delta \overline{V}_2\left(\overline{K}-1\right) + r_i\left(\overline{K}, \tau_i\right) \\
\overline{V}_2(k) &= m_k^{+1} \delta \overline{V}_2(k+1) + m_k^{0} \delta \overline{V}_2(k) + m_k^{-1} \delta \overline{V}_2(k-1) + r_i(k, \tau_i) \qquad \forall k : \underline{K} < k < \overline{K} \\
\overline{V}_2\left(\underline{K}\right) &= m_{\underline{K}}^{+1} \delta \overline{V}_2\left(\underline{K}+1\right) + m_{\underline{K}}^{0} \delta \overline{V}_2\left(\underline{K}\right) + m_{\underline{K}}^{-1} \delta \max\left\{-R, \overline{V}_2\left(\underline{K}\right)\right\} + r_i\left(\underline{K}, \tau_i\right)
\end{aligned}$$

Proof: This equation system mainly relies on the equation system defining the valuation function. This system is described through

$$V_2(k) = m_k^{+1} \delta V_2(k+1) + m_k^{0} \delta V_2(k) + m_k^{-1} \delta V_2(k-1) + r_i(k, \tau_i) \qquad \forall k .$$

Now consider the following truncated variant of the equation system.

$$\overline{V}_2\left(\overline{K}\right)=m_{\overline{K}}^{+1}\delta V'+m_{\overline{K}}^{0}\delta\overline{V}_2\left(\overline{K}\right)+m_{\overline{K}}^{-1}\delta\overline{V}_2\left(\overline{K}-1\right)+r_i\left(\overline{K},\tau_i\right)$$
$$\overline{V}_2(k)=m_k^{+1}\delta\overline{V}_2(k+1)+m_k^{0}\delta\overline{V}_2(k)+m_k^{-1}\delta\overline{V}_2(k-1)+r_i(k,\tau_i)\qquad\forall k:\underline{K}<k<\overline{K}$$
$$\overline{V}_2\left(\underline{K}\right)=m_{\underline{K}}^{+1}\delta\overline{V}_2\left(\underline{K}+1\right)+m_{\underline{K}}^{0}\delta\overline{V}_2\left(\underline{K}\right)+m_{\underline{K}}^{-1}\delta V''+r_i\left(\underline{K},\tau_i\right)$$

Let $V^*(k;V',V'')$ denote the solution of this equation system, given the truncation parameters V' and V'' (for all k in the range). For all k, the value of V^* is increasing and continuous in V' and V'', and as a result of that, for all V', we can define V'' such that $V''=\max\left\{-R,V^*\left(\underline{K,V',V''}\right)\right\}$. Let $\tilde{V}''(V')$ denote the respective value of V'' as a function of V', and note that it is increasing in V'. Thus, the solutions of the equation system claimed to characterize upper bounds can be written and reformulated as, using $V_2(k)$ as the true valuations,

$$\begin{aligned}\overline{V_2}(k)&=V^*\left(k,V',\tilde{V}''(V')\right)=V^*\left(k,R,\tilde{V}''(R)\right)\\&>V^*\left(k,V_2\left(\overline{K}+1\right),\tilde{V}''\left(V_2\left(\overline{K}+1\right)\right)\right)\geq V^*\left(k,V_2\left(\overline{K}+1\right),V_2\left(\underline{K}-1\right)\right)\equiv V_2(k)\end{aligned}$$

The first greater-than sign applies, because R is a strict upper bound for $\overline{V}_2\left(\overline{K}+1\right)$, and because V^* is correspondingly monotonous, and the second ones applies, because $\max\left\{-R,\overline{V}_2\left(\underline{K}\right)\right\}$ is an upper bound for $\overline{V}_2\left(\underline{K}-1\right)$, which is implied by Lemma 2 as shown next. On the one hand, Lemma 2 shows that if $\overline{V}_2\left(\underline{K}\right)>-R$, then $\overline{V}_2\left(\underline{K}-1\right)<\overline{V}_2\left(\underline{K}\right)$. Hence, if $\overline{V}_2\left(\underline{K}\right)>-R$, then $\overline{V}_2\left(\underline{K}-1\right)<\max\left\{-R,\overline{V}_2\left(\underline{K}\right)\right\}$ will be satisfied. On the other hand, if $\overline{V}_2\left(\underline{K}\right)<-R$, then $\overline{V}_2\left(\underline{K}-1\right)>\overline{V}_2\left(\underline{K}\right)$, but also $\overline{V}_2\left(\underline{K}-1\right)<-R$ (which is used then). QED

The proof of Proposition 2 (stated in the text) is very similar to the proof of Proposition 1 and therefore skipped.

A2. Computation of equilibria for experimental games

Treatment A: $\alpha_H=0.5$ and $\alpha_L=0.25$. In this treatment, exerting high effort is iteratively dominated in all states. In iteration 1, we can show that this applies to all states except -1 and 0, and in iteration 2, we can show this for the states -1 and 0. In turn, let us also show that "exerting low effort in all states" is a Markov perfect equilibrium. We do so by showing that Eq. (1) is satisfied for all states k. Let us define

$$DV_2(k):=\left(1-\sigma_1^k\right)\left(V_2(k+1)-V_2(k)\right)+\sigma_1^k\left(V_2(k)-V_2(k-1)\right)$$

Thus, we have to show that $DV_2(k) < \frac{4}{9}$ for all k. For most states, this is obvious, since $V_2(0) = 0$ must hold under the hypothesized strategy profile. As a result, $DV_2(k) < \frac{3}{4} * \frac{1}{2} < \frac{4}{9}$ must hold for all $k \neq 0$. To show that the players would neither deviate in state $k = 0$, boundaries for the payoffs in the states $k = -1$ and $k = 1$ are required (under the hypothesized strategy profile). Using the above equation systems for $\underline{K} = -1$ and $\overline{K} = 1$, we obtain $V_2(-1) > -\frac{3}{8}$ and $V_2(1) < \frac{3}{8}$ (conservatively rounded), which implies $DV_2(0) < \frac{4}{9}$.

Treatment B: $\alpha_H = 0.9$ **and** $\alpha_L = 0.25$**.** In iteration 1, we can eliminate high effort in the states $k \leq -5$ and $k \geq 5$, in iteration 2 high effort in state $k = -4$, and in iteration 3 in state $k = -3$. High effort in the remaining states is not dominated. The unique symmetric equilibrium in pure strategies implies to exert high effort in the states $k = -1$ and $k = 2$, and low effort otherwise. To prove this, we have to show that $DV_2(k) > \frac{20}{117}$ in states $k = -1, 2$, and $DV_2(k) < \frac{20}{117}$ otherwise (under the hypothesized strategy profile). When we solve the respective equation systems for $\underline{K} = -4$ and $\overline{K} = 4$, this results immediately. Namely, we obtain $V_2(-2) \approx -0.492$ and $V_2(3) \in (0.363, 0.388)$ (conservatively rounded), which implies $DV_2(k) < \frac{20}{117}$ for $k \leq -3$ and for $k \geq 4$. The remaining bounds are

$$V_2(-1) \approx -0.35 \quad V_2(0) \approx -0.156 \quad V_2(1) \approx -0.054 \quad V_2(2) \in (0.232, 0.253),$$

which is enough information to show that the claimed strategy profile is an MPE. Note that the three approximations are given with an accuracy higher than 10^{-4}.

Treatment C: $\alpha_H = 0.5$ **and** $\alpha_L = 0.1$**.** In iteration 1, we can eliminate high effort in the states $k \leq -4$ and $k \geq 5$, in iteration 2 for the states $k = -3, 4$, in iteration 3 for the states $k = -2, 3$, and finally in state $k = 2$. High effort is rationalizable in the states $k = -1, 0, 1$; the unique symmetric equilibrium in pure strategies is to exert high effort if and only if the state is $k=0$. Thus, we have to show that $DV_2(k)>\frac{5}{18}$ if and only if $k=0$. When we solve the equation systems for $\underline{K}=-2$ and $\overline{K}=2$, we obtain (conservatively rounded)

$$V_2(-1)\in(-0.42,-0.41) \quad V_2(0)\in(-0.26,-0.25) \quad V_2(1)\in(0.22,0.25) \quad V_2(2)<0.43$$

This provides the required information.

Treatment D: $\alpha_H=0.9$ and $\alpha_L=0.1$. In iteration 1, we can eliminate high effort in the states $k\le-5$ and $k\ge5$, in iteration 2 high effort in state $k=-4$. High effort is rationalizable in all other states. The unique symmetric equilibrium implies high effort in the states $k=-2,-1$, and low effort otherwise. To prove this, we have to show that $DV_2(k)>\frac{5}{36}$ for $k=-2,-1$, and $DV_2(k)<\frac{5}{36}$ otherwise. We calculate the boundaries using equation systems based on $\underline{K}=-4$ and $\overline{K}=4$. We obtain, conservatively rounded,

$$V_2(-4)<-0.491 \quad V_2(-3)\in(-0.468,-0.466) \quad V_2(-2)\approx-0.4058 \quad V_2(-1)\approx-0.255 \quad V_2(0)\approx-0.0823$$
$$V_2(1)\approx-0.0114 \quad V_2(2)\approx0.0534 \quad V_2(3)\in(0.339,0.348) \quad V_2(4)\in(0.428,0.453)$$

Here, it appears that the jump in the valuation function from state $k=2$ to $k=3$ justifies high effort either in $k=2$ (to reach the more valuable state $k=3$) or in state $k=3$ (to defend it). This impression is misleading. In state $k=2$, player 1 (who is behind) would exert high effort, which corrupts the chances of 2 to progress to state $k=3$. Formally,

$$DV_2(2)<(1-\alpha_H)*(0.348-0.0534)+\alpha_H*(0.0534+0.0114)<0.09<\frac{5}{36}$$

In state $k=3$, in turn, player 1 gives up, implying that 2 needs not to exert high effort any more. Formally,

$$DV_2(3)<(1-\alpha_L)*(0.453-0.339)+\alpha_L*(0.348-0.0534)<0.133<\frac{5}{36}$$

Similarly, we can show for the other states that the above strategy profile is an equilibrium.

Appendix B – Experimental Instructions and Sample Screenshots

B1. Experimental Instructions

This is the English version of the experimental instructions. The experimental instructions used in the actual experiment were in German, and are available from the authors upon request.

General Instructions

You are about to participate in an experiment on decision-making. The experiment is divided into a number of stages, and each stage is divided into rounds. During the experiment you will earn *experimental points*. At the start of each stage you are assigned an *initial endowment* of 8 experimental points. Each experimental point you earn in the experiment is worth € 1.

At the end of the experiment, a *winning stage* will be randomly chosen by the computer. Your final payment will be equal to what you earn in the winning stage. You will not know which stage is the winning stage until the end of the experiment.

There are 10 participants in the experiment. Each round you choose actions that could affect your earnings and those of a single coparticipant, and similarly he or she takes actions that could affect your earnings and his or hers. At the start of each stage, your coparticipant will be chosen at random among all other participants. Once chosen, the coparticipant remains the same throughout the stage. At the start of the following stage, however, your coparticipant will again be chosen at random. You will not be told who your coparticipant is.

The number of rounds in each stage is determined as follows. There is one chance out of ten (that is, a 10% probability) that the stage you are in terminates at the end of each round. If the stage does not terminate, you simply move on to the next round. Therefore, you will not know how many rounds there are in a stage until it terminates, and this could vary from stage to stage.

Your Decision

Each round you need to decide whether to make a *low investment* or a *high investment*. The cost of making a low investment is 0 points. The cost of making a high investment is 1 point. These costs remain the same throughout the experiment. The costs will come out of the earnings that you have in the stage.

The investment is necessary to make *progress steps*. You can make up to one progress step each round. The progress steps you have made accumulate as the stage proceeds. The computer display shows the number of progress steps that you, and your coparticipant, have made so far in the stage.

When you choose a high investment, there is a higher probability that your investment is successful than if you choose a low investment. The probabilities of success for a low and for a high investment stay constant throughout each stage and are displayed on the computer screen. You are informed at the end of each round whether your investment is successful or not. If your investment is successful for a given round, you move forward by one progress step.

To choose a low investment for the round, click the “Low Investment” button and then, if you are sure of your choice, the “Confirm” button. To choose a high investment for the round, click the “High Investment button” and then, if you are sure of your choice, the “Confirm” button.

After both you and your coparticipant have made your choices, the computer checks whether, so far in the stage, you have accumulated more progress steps than your coparticipant or otherwise:

1. If you have accumulated more progress steps than your coparticipant, you get a *high prize* and your coparticipant gets a *low prize*. The low prize is worth 1 point. The high prize is worth 2 points.
2. If you have accumulated the same number of progress steps as your coparticipant, the computer will decide randomly who gets the high prize and who gets the low prize, and so there is a 50% probability that you get the low prize and a 50% probability that you get the high prize.
3. If you have accumulated less progress steps than your coparticipant, you get the low prize and your coparticipant gets the high prize.

New prizes get assigned every round. Low prizes are always worth 1 point, and High prizes are always worth 2 points. The cost of making a low investment is always 0 points, and the cost of making a high investment is always 1 point.

The number of progress steps and points earned from prizes starts from 0 points, and the initial endowment starts at 8 points, at the beginning of each stage. The probabilities of success may, or may not, change as you move from one stage to the next.

Before starting stage 1, we ask you to answer a brief questionnaire, with the only purpose of checking whether you have understood the instructions. Raise your hand when you have completed the questionnaire.

Many thanks for your participation to the experiment.

Please raise your hand if you have any questions.

B2. Sample Screenshots

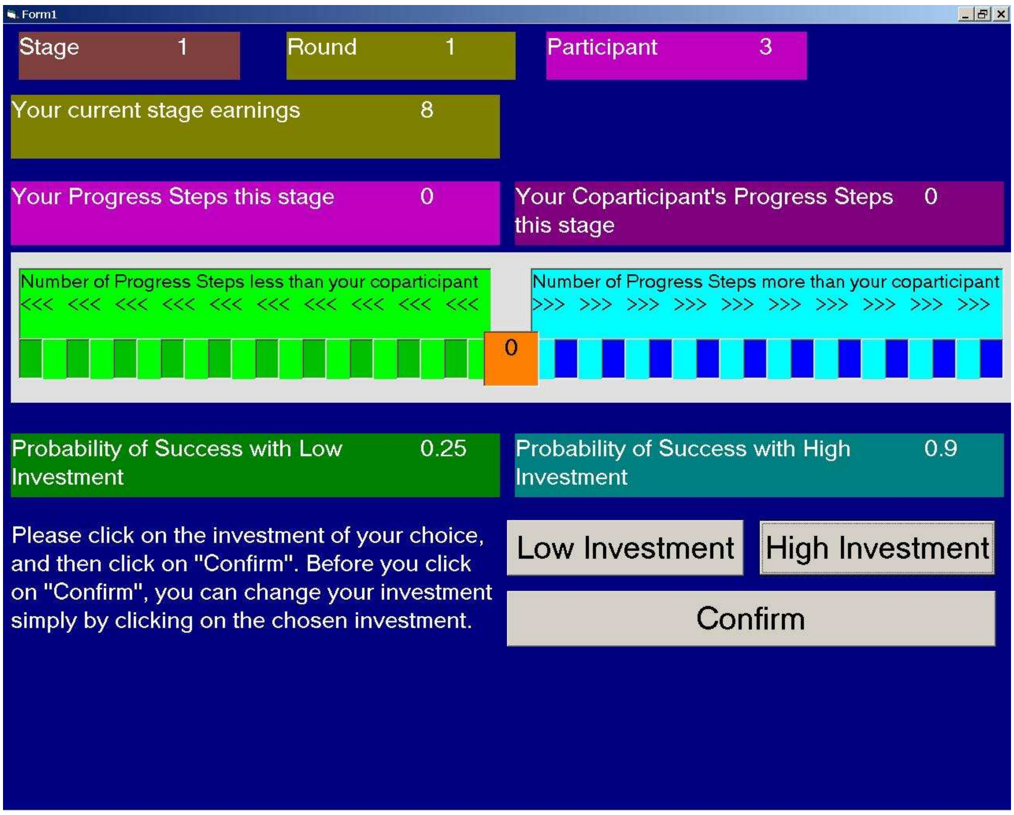

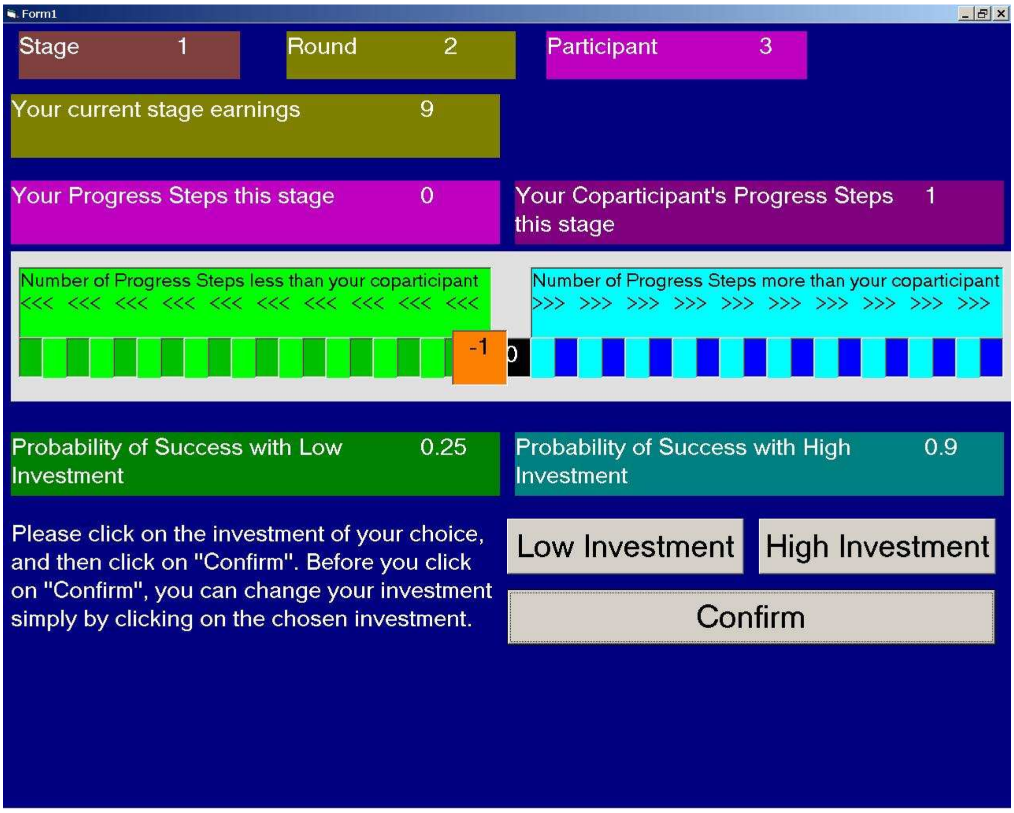

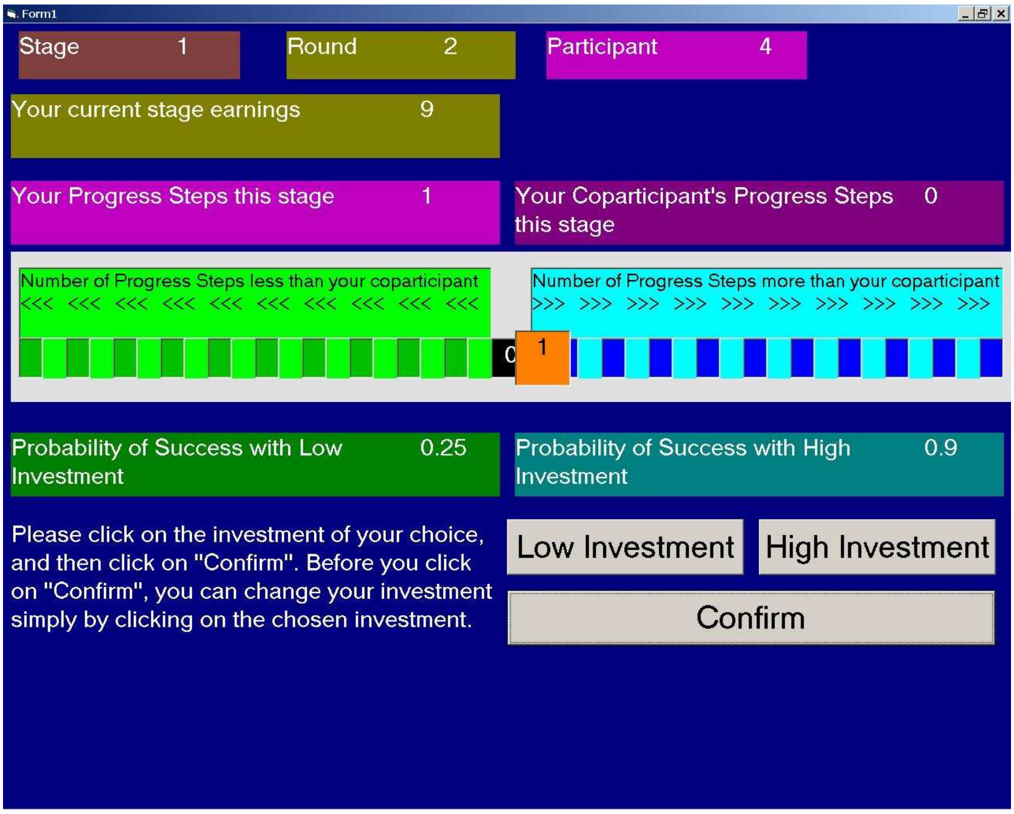

The first screenshot depicts the initial state of a stage, where $k = 0$. The second screenshot shows the state $k = -1$ following a low investment by Participant 3, which resulted in 0 progress steps. The third screenshot shows the complementary state $k = 1$ of the co-participant, Participant 4, following a high investment, which resulted in 1 progress step.

Appendix C – Further Details on R&D Monopoly Entrenchment

Is there a tendency for the market to become an R&D leadership monopoly? Figure 2 in the paper suggests that, for any given treatment and relative position, the average investment by the leader is at least as large as that of the follower. While the regression analysis in Table 1 in the paper suggests instead that the answer is negative for treatment A (with the caveats discussed in a footnote in section 3.2.1 of the paper), it also implies that, in the other treatments, the market does tend to become a R&D leadership monopoly as the gap in relative position becomes large. To shed further light on this while controlling for individual propensities to invest,[1] we ran Spearman correlations between investment and Positive Gap and between investment and Negative Gap for each subject and treatment. It is then possible to compare, for each subject, the two correlation values and see whether the Positive Gap correlation is higher than the Negative Gap correlation. This would imply that, for any given subject, as a follower she reduces high investment at a quicker pace than as a leader as the relative gap in relative position increases. We find that this is not the case for treatment A, whereas it is so for the other treatments (see Table 2 in the paper).

[1] A limitation of this test is that it does not control for session level effects, but the regressions in Table 1 in the paper do control for both individual level and session level effects. This test is just a simpler illustration of the pattern that we observe in Table 1.

Appendix D – Payoff Transformations and Rivalry Motive

D1. Theoretical Framework

An unexplained stylized fact in our experiment is the prevalent over-investment. As noted by Cohen and Levin (1989), other motives beyond strategic incentives to invest in innovation may influence investment decisions, and Brenner (1987) discusses how a rivalry motive can make competition desirable to increase R&D innovation. One may postulate that the perpetual race setting elicits a competitive mindset in the minds of (at least some) agents, making them wish to win the high prize more than they would purely on the basis of the monetary payoffs.

Reinterpret monetary payoffs (profits) as subjective utility Π_i of agent *i* and consider three kinds of payoff transformations, namely additive envy, ratio envy and direct envy:[2]

$\Pi_i = \pi_i + \beta (\pi_i - \pi_j)$ $\beta > 0$ (*additive envy*)

$\Pi_i = \pi_i + \beta (\pi_i / \pi_j)$ $\beta > 0$ (*ratio envy*)

$\Pi_i = \pi_i - \beta \pi_j$ $\beta > 0$ (*direct envy*)

For each of the three models, we considered β values in the 0.1, 0.2, …, 1 interval and computed corresponding R / c values. The additive transformation is allowed the largest increase in R / c value, up to 1.5 for $\beta = 1$; whereas $\beta = 1$ yields an R / c value of 1.25 and 1 in the ratio envy and direct envy transformations, respectively.

We tried to compute the pure symmetric equilibrium strategy for the full grid of 10 β values (between 0.1 and 1) × 3 transformations × 4 treatments. However, pure symmetric equilibrium strategies for all four treatments are defined in only six cases: namely, with β = 0.7, 0.8, 0.9 and 1 for the additive transformations (mapped into R / c values equal to 1.2, 1.3, 1.4 and 1.5 respectively) and with β = 0.9 and 1 for the ratio transformation (mapped into R / c values equal to 1.175 and 1.25 respectively). Moreover, all six of these cases lead to the same equilibrium strategy, which we label Transform 1 in the paper.

We also estimated R / c values equal to 1.5, 2, 3 and 4, corresponding to a (stronger) rivalry motive. The first two values do not yield symmetric pure equilibrium strategies for all four treatments. The last two do, and their strategies are identical for treatments B, C and D. In relation to treatment A, predictions are identical except for o = - 2 and + 3, where, unlike $R / c = 3$, $R / c = 4$ predicts high investment. Given that for this treatment Spearman ρ (high

[2] Zizzo (2008) contains an overview of the literature on envy.

investment predicted, high investment observed) = 0.225 for *R / c* = 3 but only 0.089 for *R / c* = 4, and that the two models have perfectly multicollinear predictions otherwise, we decided to discard *R / c* = 4 and treat *R / c* = 3 as our Transform 2 model in the paper.

By raising the revenue-to-cost ratio *R/c*, we can model this, and as a consequence, predicted investment may approach realistic levels. If so, then by controlling for payoff transformations we can indirectly identify rivalry concerns as a motive of innovation behavior.

D2. Empirical Analysis of Markov Perfect Equilibria Based on a Rivalry Motive

A troubling feature of this exercise noted above is that, for a number of payoff transformations, Markov perfect equilibria (MPE) cannot be computed in the exact way described above. We focus on two well-differentiated payoff transformations for which pure equilibria exist throughout: Transform 1 can be obtained with *R/c* between 1.2 and 1.5, Transform 2 with *R/c* = 3.[3] Transform 1 predicts high investment for a gap between – 1 and 2 in treatment C, and for a gap *o* = 0, 1 in the other treatments. Transform 2 predicts high investment for a gap between – 2 and 3 in treatment C, and for a gap *o* = 0, 1, 2 in the other treatments. Table 4 compares the predictive success for these models with that using MPE with *R/c* = 0.5.

High investment is predicted for more cases in the transformed models and so we may expect better predictive power. Qualitatively, however, Transform 1 and 2 lose out on the across-treatments variety of dynamic paths of MPE: they uniformly predict regions of high investment clusters where relative progress gaps are not too large.[4] Transform 1 average investments values were 0.670, 0.539, 0.761 and 0.531 in treatments A, B, C and D respectively (0.669 overall); the corresponding numbers for Transform 2 were 0.913, 0.752, 0.880 and 0.735 (0.803 overall) and, it will be recalled, 0.686, 0.611, 0.706 and 0.683 (0.669 overall) for the observed data. So the transformed models meet the primary target of hitting average investment values much closer to home, although Transform 2 has a systematic

[3] Unlike Transform 2, Transform 1 can be supported by negative spite parameters in the admissible range for our additive payoff transformation, namely between 0.7 and 1.

[4] We checked the rationalizability of high (and low) investment under Transform 1 and 2 for *o* = -5 … +5, and found that rationalizability places even less constraints here than with the baseline model. Most notably, high investment in treatment A is always rationalizable with both Transform 1 and 2. Additional details are in sub-section D3 below.

tendency of overshooting the target. Table 4 contrasts the empirical fit of MPE with that of Transform 1 and 2. Transform 1 predicts roughly 2/3 of the choices (0.656), and Transform 2 slightly more (0.714). While these results are far from sensational, and may simply be a by-product of fitting the key stylized fact of overinvestment, they do imply that the models predict better than 50% chance success in all sessions ($P < 0.005$).

The strategic problem faced by the subjects can be summarized as follows. In all cases, there exist unique equilibria in pure strategies, and these equilibria are symmetric. Dominance arguments show that high effort is rationalizable only in a connected set of states around $k = 0$. In the most extreme case (treatment 1), high effort is dominated in all states. Thus, equilibrium play can be expected under comparably weak epistemic conditions (MPEs are technically simple, as they can be defined in terms of the Nash equilibrium concept). However, our experimental observations cannot be organized by arguments of dominance or equilibrium play. Even after accounting for the possibility of rivalry concerns, hardly any evidence of a qualitatively accurate relationship between predictions and observations was found.

D3. Rationalizability and Payoff-Transformed Models

We checked the rationalizability of high and low investment under Transform 1 and 2 for o = -5, … +5. We found that very few restrictions are placed in this range (which includes 97.5% of relative positions actually faced by the subjects in the experiment). High investment is not rationalizable only in treatment B, in relation to o = 4, 5, - 4, - 5 for both Transform 1 and 2 and also o = 2, 3 for Transform 2 only. Low investment is always rationalizable under both the baseline model and Transform 1, but, in relation to Transform 2, it is not rationalizable under o = - 1, 0, 1 in treatment A and also under o = - 2, - 3 in treatment 2. On the basis of this analysis, Table D1 classifies average high or low investment according to whether it is rationalizable.

The anomaly of high investment in treatment A, observed with the baseline model, has now been addressed at least to the extent that high investment is always rationalizable. In addition, in treatments A and B, lower investment is observed under Transform 2 when low investment is *not* rationalizable ($P < 0.001$).[5]

TABLE D1: Percentage of high investment choices and rationalizability in the extended model

Transform 1 Model					
Rationalizability	Treatment				All
of high investment	A	B	C	D	Treatments
Yes	-	0.610	0.703	0.677	0.667
No	0.686	0.527	-	-	0.527
Transform 2 Model					
Rationalizability	Treatment				All
of high investment	A	B	C	D	Treatments
Yes	-	0.681	0.703	0.677	0.687
No	0.686	0.289	-	-	0.289
Transform 2 Model					
Rationalizability	Treatment				All
of low investment	A	B	C	D	Treatments
Yes	0.449	0.289	-	-	0.653
No	0.708	0.681	0.703	0.677	0.689

Values are the percentages of high (low) investment choices (made under relative position *o* in the range -5, …, 5) classified according to treatment and to whether high (low) investment is rationalizable in the extended models. Low investment is always rationalizable in the Transform 1 model.

[5] Statistical significance can be obtained not just in univariate tests but also, as for the corresponding *P* values in the main text, by using suitable logistical regressions that control for session level and individual level random effects.

Appendix E – Further Details on Markov Quantal Response Equilibria

Computations of Markov QRE. The computation of Markov quantal response equilibriua (QRE) is possible using the approach of Turocy (2005, 2006), which is implemented in Gambit (McKelvey et al., 2007). Essentially, two groups of entities are required: $\pi_{i,\omega}(s_i,\sigma)$ as defined already, and $\partial\pi_{i,\omega}(s_i,\sigma)/\partial\sigma_j(\omega')(s_j)$. This derivative is obtained by evaluating

$$\text{(E1)}\qquad \begin{aligned}\frac{\partial\left(L^{-1}\times r\right)}{\partial\sigma_j(\omega')(s_j)} &= \frac{\partial L^{-1}}{\partial\sigma_j(\omega')(s_j)}\times r + L^{-1}\times\frac{\partial r}{\partial\sigma_j(\omega')(s_j)}\\ &= L^{-1}\frac{\partial L}{\partial\sigma_j(\omega')(s_j)}L^{-1}\times r + L^{-1}\times\frac{\partial r}{\partial\sigma_j(\omega')(s_j)}\end{aligned}$$

As both $L(i,\omega,s_i;\sigma)$ and $r(i,\omega,s_i;\sigma)$ are linear in $\sigma_j(\omega')(s_j)$, this computation is straightforward. The value of $\partial\pi_{i,\omega}(s_i,\sigma)/\partial\sigma_j(\omega')(s_j)$ is equal to the element x of the vector defined in Eq. (E1).

In contrast to the classes of games studied by Turocy, one should note that $\partial\pi_{i,\omega}(s_i,\sigma)/\partial\sigma_j(\omega')(s_j)$ is generally not zero, i.e. not even if $(i,\omega)=(j,\omega')$. This difference is a consequence of the character of the normal form game considered here – its definition depends on the assumed strategy profile σ. Hence, as σ changes, the normal form game changes, which affects the information required to trace the QRE correspondence. Loosely speaking, the above derivatives contain this information and may therefore not equate with zero. Further guidance and our implementation is available upon request.

In the computations made for the paper, we have to restrict the state space in order to apply the above definition of Markov QREs. The implied inaccuracies in the computations are negligible, however, if the cut-off points are chosen conservatively enough. According to our data set, 95% of the observations are for states $k\in(-5,\ldots,5)$, and 99% of the observations are for states $k\in(-10,\ldots,10)$. We chose the cut-off points $k_1=-15$ and $k_2=15$, and inspections of the predictions suggest that this is sufficient here (essentially, for intermediate values of λ, the predicted probability of high effort *near* the upper limit $k_2=15$ is exaggerated by a few percentage points). Thus, we analyze games with 2 players, 31 states, and 2 actions per state.

In each iteration of the tracing procedure, payoffs have to be calculated for all agents and all actions. In addition, the respective derivatives with respect to all actions of all agents are required. Assume that neither the block triangularity of L is exploited, nor, in our case, that the upper diagonal block is tridiagonal. Then, each of these steps requires the inversion of a matrix with dimension $(|\Omega|+1)^2$. For instance, if we have 2 players, 31 states, and 2 actions per player, this requires 62^2 inversions of 32×32 matrices. Given current technology, this is achieved in 10-20 seconds. Depending on the step length used, 50 or more such steps may be required to obtain the QRE correspondence. While the required amount of time is not negligible, one should take into account that this is actually an analysis of a game with $2^{62} \approx 4.6*10^{18}$ strategy profiles. The mere enumeration of these strategy profiles is not possible using current technology. The above efficiency is achieved by exploiting that (i) the homotopy approach requires the above evaluations only for a few strategy profiles σ (those obtained along the QRE correspondence), and that (ii) in dynamic games, the respective evaluations for mixed σ do not require knowledge of the payoffs associated with pure strategy profiles. The latter, in particular, is a characteristic of dynamic games that does not apply in other classes of games. In turn, by considering dynamic games, we are able to investigate complex environments (and phenomena) at computationally small costs.

Levels of reasoning. We consider a variant on the QRE model allowing for limited depth of reasoning, and which relies on Kübler and Weizsäcker (2004). Kübler and Weizsäcker estimated the depth of reasoning in cascade formation, i.e. in finite horizon games, within a quantal response equilibrium approach. They estimated models with three to five levels of reasoning, where the reasoning had been increasingly “fuzzy” (lower λ) the higher the level. Similar results might be expected here, not only because such limitations may appear intuitive, but also because of reasons that appear technical at first glance: the equilibria of these limited models are much easier and quicker to compute, and in relation to the QRE predictions, numerical inaccuracy is not an issue here (note that the limited depth leads effectively to a finite state space). Nonetheless, we have not found limited depth of reasoning empirically relevant. We estimated models for three to five levels of reasoning, using a variety of optimization algorithms (besides the algorithm of Nelder and Mead, 1965, which was used throughout, this included the BFGS method), but obtained quadratic scores between 4100 and 4200 in all cases. For instance, the estimated set of parameters for a four-level model are as follows (note that the payoffs are scaled differently here, which leads to differently scaled values of λ):

$$(\lambda_1, \lambda_2, \lambda_3, \lambda_4) = (6.028, 66.506, 64.799, 45.429)$$

for a quadratic score of 4173.04. Here, λ_1 refers to the first level of reasoning, λ_2 refers to the second level (the actions played in the next round), λ_3 refers to the third level (the round after next), and λ_4 refers to the fourth level. For the next level, the beginning of the backward induction, $\lambda_5 = 0$ was set (see Kübler and Weizsäcker, 2004, for an in-depth discussion of this approach). The latter, $\lambda_5 = 0$, implies that thus limited models differ structurally from the above QRE models. Combined, these results suggest that the subjects' actions display two kinds of bounded rationality, payoff perturbations as in QRE models and a rivalry motive, but, given these facts, they are able to assess all states equally well. That is, (i) their evaluations of given states are independent of how many rounds would be needed to reach this state, and (ii) the evaluation of state k is the same whether it is reached in the next round (second reasoning level) or in any later round (higher level).

References Cited in Electronic Appendices

Kübler, D. and Weizsäcker, G. (2004). Limited depth of reasoning and failure of cascade formation in the laboratory. Review of Economic Studies, 71(2):425–41.

McKelvey, R. D., McLennan, A. M., and Turocy, T. L. (2007). Gambit: Software tools for game theory. Version 0.2007.01.30 http://econweb.tamu.edu/gambit.

Nelder, J. and Mead, R. (1965). A simplex algorithm for function minimization. Computer Journal, 7:308–13.

Turocy, T. (2005). A dynamic homotopy interpretation of the logistic quantal response equilibrium correspondence. Games and Economic Behavior, 51(2):243–63.

Turocy, T. (2006). Using quantal response to compute Nash and sequential equilibria. Working Paper.

Zizzo, D.J. (2008). The cognitive and behavioral economics of envy. In R.H. Smith (ed.), Envy: Theory and Research, New York: Oxford University Press, pp. 190-210.