

NBER WORKING PAPER SERIES

THE CAUSAL EFFECTS OF COMPETITION ON INNOVATION:
EXPERIMENTAL EVIDENCE

Philippe Aghion
Stefan Bechtold
Lea Cassar
Holger Herz

Working Paper 19987
<http://www.nber.org/papers/w19987>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2014

We would like to thank Daniel Chen, Ernst Fehr, Bernhard Ganglmair, David Huffman, Klaus Schmidt, Roberto Weber, as well as the participants at the IMEBE 2013 conference, the ETH Zurich/Max Planck Workshop on the Law & Economics of Intellectual Property and Competition Law 2013 and the CSAE seminar at the University of Oxford for their useful comments and feedback. Financial support from ETH Zurich is gratefully acknowledged. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by Philippe Aghion, Stefan Bechtold, Lea Cassar, and Holger Herz. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Causal Effects of Competition on Innovation: Experimental Evidence
Philippe Aghion, Stefan Bechtold, Lea Cassar, and Holger Herz
NBER Working Paper No. 19987
March 2014
JEL No. C91,L10,O31

ABSTRACT

In this paper, we design two laboratory experiments to analyze the causal effects of competition on step-by-step innovation. Innovations result from costly R&D investments and move technology up one step. Competition is inversely measured by the ex post rents for firms that operate at the same technological level, i.e. for neck-and-neck firms. First, we find that increased competition leads to a significant increase in R&D investments by neck-and-neck firms. Second, increased competition decreases R&D investments by firms that are lagging behind, in particular if the time horizon is short. Third, we find that increased competition affects industry composition by reducing the fraction of sectors where firms are neck-and-neck. All these results are consistent with the predictions of step-by-step innovation models.

Philippe Aghion
Department of Economics
Harvard University
1805 Cambridge St
Cambridge, MA 02138
and NBER
paghion@fas.harvard.edu

Stefan Bechtold
ETH Zurich
Rämistrasse 101
8092 Zürich, Switzerland
sbechtold@ethz.ch

Lea Cassar
University of Zurich
Rämistrasse 71
8006 Zürich, Switzerland
lea.cassar@econ.uzh.ch

Holger Herz
University of Zurich
holger.herz@econ.uzh.ch

An online appendix is available at:
<http://www.nber.org/data-appendix/w19987>

1 Introduction

The relationship between competition and innovation has long been of interest to economists and motivated numerous studies, both theoretical and empirical, over the past three decades (e.g., Hart (1980); Schmidt (1997); Aghion et al. (2001); Vives (2008); Schmutzler (2009, 2013); and Nickell (1996); Blundell et al. (1999); Aghion et al. (2005); Aghion and Griffith (2006)). However, the existing empirical studies on this subject face the issue that the relationship between competition and innovation is endogenous (Jaffe (2000); Hall and Harhoff (2012)).¹ Moreover, clean and direct measures of innovation and competition are usually not available in field data, which can lead to the additional problem of measurement error.

To address these two issues head on, in this paper we employ the methods of experimental economics to analyze the effects of competition on step-by-step innovation. The predictions we submit to our experiments are drawn from the step-by-step innovation models of Aghion et al. (1997, 2001) and Aghion and Howitt (2009). These models predict that product market competition should foster innovation in neck-and-neck sectors where firms operate at the same technological level: in such sectors, increased product market competition reduces pre-innovation rents, thereby increasing the incremental profits from innovating and becoming a leader. This is known as the “escape-competition effect”. On the other hand, these models predict a negative “Schumpeterian effect” on laggard firms in unlevelled sectors: increased competition reduces the post-innovation rents of laggard firms and thus their incentive to catch up with the leader. However, this effect is (partly) counteracted by an “anticipated escape-competition effect” once the laggard has caught up with the current leader in the sector. The escape-competition and Schumpeterian effects, together with the fact that the equilibrium fraction of neck-and-neck sectors depends positively on the laggards’ innovation incentives in unlevelled sectors and negatively on neck-and-neck firms’ innovation incentives in leveled sectors, imply that the equilibrium fraction of sectors where firms are neck and neck should decrease with competition: this is the “composition effect” of competition.

To test these predictions, we design two laboratory experiments. In both

¹See Blundell et al. (1999) and Aghion et al. (2005) for a discussion of the endogeneity issue and for attempts at addressing it. For various empirical approaches to identify causal relationships between patenting activities and innovation, see Murray and Stern (2007); Williams (2013); Galasso and Schankerman (2013).

experiments, pairs of subjects are matched for a number of periods, forming a sector. In each period, one of the two subjects can choose an R&D investment which determines the probability of a successful innovation in that period. Innovation is costly and has an associated cost generated via a quadratic cost function. If innovation is successful, the technological level of the innovative subject increases by one step. At the end of each period, rents are distributed to each subject according to her relative technological location in her sector. If the two subjects are in an unleveled sector, then the subject ahead (the “leader”) receives a positive monopoly rent, whereas the other subject in the same sector (the “laggard”) makes zero profit. If subjects in the same sector are neck and neck, their rents are equal and depend on the degree of competition. In the no competition treatment, these firms are able to split the monopoly rent between them, whereas under the full competition treatment the neck-and-neck firms’ profits are zero. In the intermediate competition treatment, neck-and-neck subjects are able to split half the monopoly rent between them.

We conduct an “infinite horizon” experiment to bring out the escape-competition and the Schumpeterian effects most clearly and a “finite horizon” experiment to assess the composition effect. In the infinite horizon experiment, we exogenously vary the subjects’ (or “firms”) starting positions. That is, some pairs of subjects start as unleveled sectors while other pairs start as neck-and-neck sectors. This design feature, together with the treatment variation in the degree of competition, allows us to causally assess the escape-competition effect and the Schumpeterian effect. After every period, the interaction between two paired subjects ends with a positive stopping probability: in other words, the time horizon is infinite but we exogenously vary the expected time horizon across sessions. Pairs either faced a short time horizon – a 80% probability of ending the game after each round – or a long time horizon – a 10% probability of ending the game after each round. This setup allows us to test an additional prediction from the theory: We should expect a more negative impact of competition on laggards’ innovation intensity in the short horizon treatment than in the long horizon treatment, since the longer the time horizon, the more the anticipated escape-competition effect may counteract the Schumpeterian effect.

In the finite horizon experiment, all subjects faced the same finite horizon of 50 periods. Each pair starts as a neck-and-neck sector, and the ability to innovate alternates between the two subjects across periods. Because of the

exogenous variation of competition across treatments, this design allows us to cleanly identify the causal effect of competition on industry composition and also on aggregate innovation outcomes.

The results can be summarized as follows. First, an increase in competition leads to a significant increase in R&D investments by neck-and-neck firms. Second, an increase in competition decreases R&D investments by laggard firms. Moreover, this Schumpeterian effect is significantly stronger the shorter the time horizon. Third, increased competition affects industry composition by reducing the fraction of neck-and-neck sectors, and, overall, competition increases aggregate innovation. All these results are consistent with the predictions of step-by-step innovation models.

This paper relates to several strands of literature. First, it contributes to an extensive theoretical industrial organization literature on competition, rent dissipation, and research and development (see Tirole (1988)), and to the patent race literature (Harris and Vickers (1985a,b, 1987)). Second, it relates to the endogenous growth literature and more specifically to its Schumpeterian growth branch (see Aghion et al. (2013)).

Third, the paper contributes to the existing empirical literature on the relationship between competition and innovation (see Nickell (1996), Blundell et al. (1999), Aghion et al. (2005); Aghion and Griffith (2006); Acemoglu and Akgigit (2012)). Our laboratory setting and the investment game that we study are stylized, yet competition and the time horizon both vary exogenously and moreover R&D investments are directly observed. This in turn allows us to complement the existing field evidence by shedding further light on the causal rather than correlational relationship between competition and innovation.

Fourth, the paper contributes to an emerging experimental literature on competition and innovation.² Our experimental study differs from this lit-

²Isaac and Reynolds (1988) analyze the effects of competition and appropriability in simultaneous investment, one-period patent races. They show that per-capita investments are decreasing with the number of contestants, whereas the aggregate level of investment increases. Darai et al. (2010) find similar results in a two-stage game in which R&D investment choices are followed by product market competition. Moreover, in a two-stage game with cost-reducing investments followed by a differentiated Cournot duopoly, Sacco and Schmutzler (2011) find a U-shape relationship between competition and innovation, the former being defined as the degree of product differentiation. Finally, Suetens and Potters (2007) finds that tacit collusion is higher in Bertrand competition than in Cournot competition. However, the study does not look at the effect of competition on innovation.

erature in particular by focusing on an environment in which innovation is cumulative over time.³ And to our knowledge, we are the first to design a laboratory experiment to examine the escape competition and Schumpeterian effects in a dynamic investment environment with different time horizons, and to assess the composition effect.

The remaining part of the paper is organized as follows. Section 2 lays out the basic step-by-step innovation model and derives its main predictions. Section 3 describes the setup for the two experiments. Section 4 presents the results, and Section 5 concludes.

2 Theoretical predictions

To motivate and guide our experiment, and to derive our main predictions, we first present a simple version of the model with step-by-step innovations. While we chose to write the simplest possible model for pedagogical purposes, the predictions derived in this Section are robust to generalizations of the environment (see, for example, Aghion et al. (2001)).

In this model, a “laggard firm”, i.e. a firm that is currently behind the technological leader in the same sector, must catch up with the leader before becoming a leader itself. This step-by-step assumption implies that, in a positive fraction of sectors, firms will be neck and neck, i.e. at the same technology level. By making life more difficult for neck-and-neck firms, a higher degree of product market competition will encourage them to innovate in order to acquire a significant lead over their rivals. However, higher competition may have a discouraging effect on innovation by laggard firms in unlevelled sectors.

2.1 Basic environment

We consider a simple Schumpeterian growth model in continuous time. There is a continuous measure L of infinitely-lived individuals each of whom supplies

³Isaac and Reynolds (1992) also study innovation investments over time. However, the authors do not distinguish between the escape-competition and the Schumpeterian effect, nor do they test the composition effect of competition. Similarly, Zizzo (2002) and Breitmoser et al. (2010) study innovation investments in a multi-stage patent race. These papers, however, do not look at the effect of varying competition and therefore cannot address the Schumpeterian, escape-competition and composition effects.

one unit of labor per unit of time. Each individual has intertemporal utility

$$u_t = \int_0^1 \ln C_t e^{-\rho t} dt$$

where

$$\ln C_t = \int_0^1 \ln y_{jt} dj,$$

and where each y_j is the sum of two goods produced by (infinitely-lived) duopolists in sector j :

$$y_j = y_{Aj} + y_{Bj}.$$

The logarithmic structure of the utility function implies that in equilibrium individuals spend the same amount on each basket y_j .⁴ We choose this expenditure as the numeraire, so that a consumer chooses each y_{Aj} and y_{Bj} to maximize $y_{Aj} + y_{Bj}$ subject to the budget constraint: $p_{Aj}y_{Aj} + p_{Bj}y_{Bj} = 1$; that is, she will devote the entire unit of expenditure to the least expensive of the two goods.

2.1.1 Technology and innovation

Each firm takes the wage rate as given and produces using labor as the only input according to the following linear production function,

$$y_{it} = \gamma^{k_{it}} l_{it}, \quad i \in \{A, B\}$$

where l_{jt} is the labor employed, k_{it} denote the technology level of duopoly firm i at date t , and $\gamma > 1$ is a parameter that measures the size of a leading-edge innovation. Equivalently, it takes γ^{-k_i} units of labor for firm i to produce one unit of output. Thus the unit costs of production is simply $c_i = w\gamma^{-k_i}$ which is independent of the quantity produced.

⁴To see this, note that a final good producer will choose the y_j 's to maximize $u = \int \ln y_j dj$ subject to the budget constraint $\int p_j y_j dj = E$, where E denotes current expenditures. The first-order condition for this is:

$$\partial u / \partial y_j = 1/y_j = \lambda p_j \text{ for all } j$$

where λ is a Lagrange multiplier. Together with the budget constraint this first-order condition implies

$$p_j y_j = 1/\lambda = E \text{ for all } j.$$

An industry j is thus fully characterized by a pair of integers (k_j, m_j) where k_j is the leader's technology and m_j is the technological gap between the leader and the follower.⁵

For expositional simplicity, we assume that knowledge spillovers between the two firms in any intermediate industry are such that neither firm can get more than one technological level ahead of the other, that is:

$$m \leq 1.$$

In other words, if a firm already one step ahead innovates, the lagging firm will automatically learn to copy the leader's previous technology and thereby remain only one step behind. Thus, at any point in time, there will be two kinds of intermediate sectors in the economy: (i) *level* or *neck-and-neck* sectors where both firms are at technological par with one another, and (ii) *unleveled* sectors, where one firm (the *leader*) lies one step ahead of its competitor (the *laggard* or *follower*) in the same industry.⁶

To complete the description of the model, we just need to specify the innovation technology. Here we simply assume that by spending the R&D cost $\psi(z) = z^2/2$ in units of labor, a leader firm moves one technological step ahead, with probability z . We call z the "innovation rate" or "R&D intensity" of the firm. We assume that a laggard firm can move one step ahead with probability h , even if it spends nothing on R&D, by copying the leader's technology. In other words, it is easier to reinvent the wheel than to invent the wheel. Thus $z^2/2$ is the R&D cost (in units of labor) of a laggard firm moving ahead with probability $z+h$. Let z_0 denote the R&D intensity of each firm in a neck-and-neck industry; and let z_{-1} denote the R&D intensity of a follower firm in an unleveled industry; if z_1 denotes the R&D intensity of the leader in an unleveled industry, note that $z_1 = 0$, since our assumption of automatic catch-up means that a leader cannot gain any further advantage by innovating.

⁵The above logarithmic final good technology together with the linear production cost structure for intermediate goods implies that the equilibrium profit flows of the leader and the follower in an industry depends only on the technological gap m between the two firms (see Aghion and Howitt (1998, 2009)).

⁶See Aghion et al (2001) for an analysis of the more general case where there is no limit to the technological gap between leaders and laggards in unleveled sectors.

2.2 Equilibrium profits and competition in level and unleveled sectors

One can show that the equilibrium profits are as follows (see Aghion and Howitt (2009)). First, in an unleveled sector, the leader's profit is equal to

$$\pi_1 = 1 - \frac{1}{\gamma},$$

whereas the laggard in the unleveled sector will be priced out of the market and hence will earn a zero profit:

$$\pi_{-1} = 0.$$

Consider now a level (or neck-and-neck) sector. If the two firms engaged in open price competition with no collusion, then Bertrand competition will bring (neck-an-neck) firms' profits down to zero. At the other extreme, if the two firms collude so effectively as to maximize their joint profits and shared the proceeds, then they would together act like the leader in an unleveled sector, so that each firm will earn $\pi_1/2$.

Now, if we model the degree of product market competition inversely by the degree to which the two firms in a neck-and-neck industry are able to collude, the normalized profit of a neck-and-neck firm will be of the form:

$$\pi_0 = (1 - \Delta) \pi_1, \quad 1/2 \leq \Delta \leq 1,$$

where Δ parameterizes product market competition.

We next analyze how the equilibrium research intensities z_0 and z_{-1} of neck-and-neck and laggard firms, respectively, vary with our measure of competition Δ .

2.3 The Schumpeterian and escape-competition effects

Let V_m (resp. V_{-m}) denote the normalized steady-state value of being currently a leader (resp. a laggard) in an industry with technological gap m , and let w denote the steady-state wage rate. We have the following Bellman

equations:⁷

$$\rho V_0 = \pi_0 + \bar{z}_0(V_{-1} - V_0) + z_0(V_1 - V_0) - wz_0^2/2 \quad (1)$$

$$\rho V_{-1} = \pi_{-1} + (z_{-1} + h)(V_0 - V_{-1}) - wz_{-1}^2/2 \quad (2)$$

$$\rho V_1 = \pi_1 + (z_{-1} + h)(V_0 - V_1) \quad (3)$$

where \bar{z}_0 denotes the R&D intensity of the other firm in a neck-and-neck industry (we focus on a symmetric equilibrium where $\bar{z}_0 = z_0$).

In words, the growth-adjusted annuity value ρV_0 of currently being neck-and-neck is equal to the corresponding profit flow π_0 plus the expected capital gain $z_0(V_1 - V_0)$ of acquiring a lead over the rival plus the expected capital loss $\bar{z}_0(V_{-1} - V_0)$ if the rival innovates and thereby becomes the leader, minus the R&D cost $wz_0^2/2$. Similarly, the annuity value ρV_1 of being a technological leader in an unleveled industry is equal to the current profit flow π_1 plus the expected capital loss $z_{-1}(V_0 - V_1)$ if the leader is being caught up by the laggard (recall that a leader does not invest in R&D); finally, the annuity value ρV_{-1} of currently being a laggard in an unleveled industry, is equal to the corresponding profit flow π_{-1} plus the expected capital gain $(z_{-1} + h)(V_0 - V_{-1})$ of catching up with the leader, minus the R&D cost $wz_{-1}^2/2$.

Using the fact that z_0 maximizes the RHS of (1) and z_{-1} maximizes the RHS of (2), we have the first order conditions:

$$wz_0 = V_1 - V_0 \quad (4)$$

$$wz_{-1} = V_0 - V_{-1}. \quad (5)$$

In Aghion, Harris and Vickers (1997) the model is closed by a labor market clearing equation which determines ω as a function of the aggregate demand for R&D plus the aggregate demand for manufacturing labor. Here, for simplicity we shall ignore that equation and take the wage rate w as given, normalizing it at $w = 1$.

Then, using (4) and (5) to eliminate the V 's from the system of equations (1)-(3), we obtain a system of two equations in the two unknowns z_0 and z_{-1} :

$$z_0^2/2 + (\rho + h)z_0 - (\pi_1 - \pi_0) = 0 \quad (6)$$

$$z_{-1}^2/2 + (\rho + z_0 + h)z_{-1} - (\pi_0 - \pi_{-1}) - z_0^2/2 = 0 \quad (7)$$

⁷Note that the left-hand-side of the Bellman equations should first be written as $rV_0 - \dot{V}_0$. Then, using the fact that on a balanced growth path $\dot{V}_0 = gV_0$, and using the Euler equation $r - g = \rho$, yields the Bellman equations in the text.

These equations can be solved recursively for unique positive values of z_0 and z_{-1} , and we are mainly interested by how these are affected by an increase in product market competition Δ . It is straightforward to see from equation (6) and the fact that

$$\pi_1 - \pi_0 = \Delta\pi_1$$

that an increase in Δ will increase the innovation intensity $z_0(\Delta)$ of a neck-and-neck firm. This is the escape competition effect:

Prediction 1 (*Escape-competition effect*): *Innovation by neck and neck firms is always stimulated by higher competition.*

Because competition negatively affects pre-innovation rents, competition induces innovation in neck-and-neck sectors since firms are particularly attracted by the monopoly rent.

One can express $z_0(\Delta)$ as

$$z_0(\Delta) = -(\rho + h) + \sqrt{(\rho + h)^2 + 2\Delta\pi_1},$$

taking the derivative

$$\frac{\partial z_0}{\partial \Delta} = \frac{\pi_1}{\sqrt{(\rho + h)^2 + 2\Delta\pi_1}}.$$

In particular, $\frac{\partial z_0}{\partial \Delta}$ can be shown to decrease when the rate of time preference ρ increases. This generates the next prediction:

Prediction 2 (*Escape-competition effect by rate of time preference*): *The escape-competition effect is weaker for firms with high rate of time preference.*

In other words, patient neck-and-neck firms put more weight on the future post-innovation rents after having become a leader, and therefore, react more positively to an increase in competition than impatient neck-and-neck firms.

Then, plugging $z_0(\Delta)$ into (7), we can look at the effect of an increase in competition Δ on the innovation intensity z_{-1} of a laggard. This effect is ambiguous in general. In particular, for sufficiently high ρ , the effect is negative as then z_{-1} varies like

$$\pi_0 - \pi_{-1} = (1 - \Delta)\pi_1.$$

In this case, the laggard is very impatient and thus looks at its short-term net profit flow if it catches up with the leader, which in turn decreases when competition increases. This is the *Schumpeterian effect*:

Prediction 3 (*Schumpeterian effect*): *Innovation by laggard firms in unleveled sectors is discouraged by higher competition.*

Since competition negatively affects the post-innovation rents of laggards, competition reduces innovation of laggards. However, for low values of ρ , this effect is counteracted by an *anticipated escape-competition effect*:

Prediction 4 (*Anticipated escape-competition effect*): *The effect of competition on laggards' innovation is less negative for firms with low rate of time preference.*

In other words, patient laggards take into account their potential future reincarnation as neck-and-neck firms, and therefore react less negatively to an increase in competition than impatient laggards. The lower the rate of time preference, the stronger the (positive) anticipated escape-competition effect and therefore the more it mitigates the (negative) Schumpeterian effect of competition on laggards' innovation incentives. For laggards with a very low rate of time preference, the anticipated escape-competition effect could even completely dominate the Schumpeterian effect, leading to a positive relationship between competition and innovation for laggards.

Thus the effect of competition on innovation depends on the situation of a sector. In unleveled sectors, the Schumpeterian effect is at work even if it does not always dominate. But in level (or neck-and-neck) sectors, the escape-competition effect is the only effect at work; that is, more competition always induces neck-and-neck firms to innovate more in order to escape from tougher competition.

2.4 Composition effect

In steady state, the fraction of sectors μ_1 that are unleveled is constant, as is the fraction $\mu_0 = 1 - \mu_1$ of sectors that are leveled. The fraction of unleveled sectors that become leveled each period will be $z_{-1} + h$, so the sectors moving from unleveled to leveled represent the fraction $(z_{-1} + h)\mu_1$ of all sectors. Likewise, the fraction of all sectors moving in the opposite

direction is $2z_0\mu_0$, since each of the two neck-and-neck firms innovates with probability z_0 . In steady state, the fraction of firms moving in one direction must equal the fraction moving in the other direction:

$$(z_{-1} + h)\mu_1 = 2z_0(1 - \mu_1),$$

which can be solved for the steady state fraction of unleveled sectors:

$$\mu_1 = \frac{2z_0}{z_{-1} + h + 2z_0}. \quad (8)$$

This fraction is increasing in competition as measured by Δ since a higher Δ increases R&D intensity $2z_0$ in neck-and-neck sectors (the escape-competition effect) whereas it tends to reduce R&D intensity $z_{-1} + h$ in unleveled sectors (the Schumpeterian effect). This positive effect of competition on the steady-state equilibrium fraction of neck-and-neck sectors we refer to as the composition effect of competition:

Prediction 5 (*Composition effect*): *The higher the degree of competition, the smaller the equilibrium fraction of neck-and-neck sectors in the economy.*

More competition increases innovation incentives for neck-and-neck firms whereas it reduces innovation incentives of laggard firms in unleveled sectors. Consequently this reduces the flow of sectors from unleveled to leveled whereas it increases the flow of sectors from leveled to unleveled.

We now proceed to confront these predictions with experimental evidence.

3 The experiments

We conduct two separate experiments, the “infinite horizon” and the “finite horizon” experiments, that are based on the same step-by-step innovation game. The purpose of the infinite horizon experiment is to provide causal evidence for the escape-competition and Schumpeterian effects, both when the time horizon is long and short. Exogenous variations in the time horizon can equivalently be interpreted as exogenous variations in firms’ rate of time preference. The purpose of the finite horizon experiment is to observe duopolies over a long period of time in order to provide causal evidence for the composition effect. We first present the basic features of our laboratory step-by-step innovation game, and will then describe the specific features of the two experiments in more detail.

3.1 The step-by-step innovation game

At the center of our experimental design is a computerized step-by-step innovation game with the following features: Two subjects i and j are randomly matched with each other, forming a sector. Before period 1, both subjects are exogenously assigned to technology levels τ_i and τ_j . In each period, one of the two subjects can choose an R&D investment $n \in \{0, 5, 10, \dots, 80\}$, which determines the probability of a successful innovation in this period.⁸ Each subject is informed that the common quadratic R&D cost function is given by:

$$C(n) = 600 \left(\frac{n}{100} \right)^2. \quad (9)$$

At the beginning of period 1, it is randomly determined which subject can invest in innovation in that period. If innovation is successful, the innovator moves up one step from θ to $\theta + 1$. Thus, in our experiments only one subject could innovate at a time. This reduces the computational complexity of the game by fixing subjects' beliefs about their opponents' current action and therefore about the expected returns in that period from innovating. However, precluding simultaneous moves in our experiment does not qualitatively affect the theoretical predictions of the model.⁹

To account for technology spillovers which may occur from leaders to laggards, laggards are automatically granted with an additional innovation probability of $h = 30$. Thus, the overall innovation probability is given by $n + h$, where n can be chosen from $n \in \{0, 5, 10, \dots, 50\}$, so that the maximal probability of innovating is still given by $n + h = 80$.

At the end of each period, payments are made to both subjects. Payments depend on the technology levels θ_i and θ_j of subjects i and j . Payments after

⁸We have deliberately excluded the possibility to choose an investment in R&D of 100, i.e. to innovate with certainty. We believe that certain innovation would be an unrealistic feature of the environment which we are studying.

⁹Indeed, while our experiments were conducted in discrete time, the model described in Section 2 is in continuous time. A continuous time implementation of the experiments was technically not feasible. Since the probability of simultaneous innovation is negligible when time is continuous, alternating moves are in fact a closer experimental implementation of the continuous time model. Also note that Aghion and Howitt (2009) specify a step-by-step innovation model in which time is discrete and only one firm innovates in each period. The basic theoretical predictions remain the same.

every period to subject i are determined by the following function:

$$\pi_i = \begin{cases} 200 - C(n) & \text{if } \theta_i > \theta_j \\ (1 - \Delta)200 - C(n) & \text{if } \theta_i = \theta_j \\ 0 - C(n) & \text{if } \theta_i < \theta_j. \end{cases}$$

Subjects are symmetric, and payments to subject j are determined in exactly the same way. This profit function implies that the leader is always able to earn a monopoly rent of 200 points. If subjects are neck and neck, their payoffs depend on the degree of competition in the sector, which is measured by Δ . If $\Delta = 0.5$, subjects are able to split the monopoly rent between them, whereas if $\Delta = 1$, subjects face perfect competition and rents in neck-and-neck states are 0.¹⁰ Finally, if a subject is lagging behind his competitor in terms of technology, her profit is 0. The costs associated with R&D are always subtracted from the rent earned.

After each period, each subject receives information about the behavior of the other subject in the preceding period. In particular, she sees which investment level the other subject has chosen, what the costs of innovation were, whether the other subject benefited from a laggard innovation bonus of $h = 30$ in the preceding round, and whether she will benefit from a laggard innovation bonus in the current period. In addition, each subject sees the payment of both subjects in the preceding period as well as the current technology level and the overall income of both subjects over the entire treatment. Based on this information, subjects can decide how much to invest in the current period.

3.2 The infinite horizon experiment

The infinite horizon experiment consisted of three different treatment variations, some of which within subjects, and others across subjects. First, all subjects participated in a full competition ($\Delta = 1$) and in a no competition

¹⁰We chose to capture exogenous variations in the degree of competition by directly modifying the payoffs to subjects in neck-and-neck states. In our view, this is the most direct and cleanest exogenous manipulation of competition in a step-by-step innovation model. An alternative approach in the experimental literature would be to vary the number of subjects competing in the same market. For the purposes of our experiments, however, we are less concerned by the particular channel whereby the competitive process affects competitive outcomes, than by how the outcome of this competitive process affects innovation incentives.

treatment ($\Delta = 0.5$), which were played sequentially in varying order across sessions.

Second, each subject made decisions conditional on being a laggard or being neck-and-neck in the first round. Therefore, the starting position was randomly varied within subject. In each treatment, it was randomly determined whether both subjects were assigned to technology levels of $\tau_i = 0$, or whether one of the subjects was given a head start and assigned to a technology level of $\tau_i = 1$. If the sector was leveled at the start, it was randomly determined which of the two firms could invest in the first round. If the sector was unleveled, the subject that was randomly put into the laggard position could invest in the first period. We elicited the first round investment decisions of all subjects using the strategy method. That is, we asked the subjects how much they want to invest if they start in a neck-and-neck sector, and how much if they start as laggards. After choices were made, the computer would randomly select the initial state of the sector and which of the two firms could invest. The investment choice associated with the chosen state was then automatically implemented. The strategy method was only used for first round investment choices. The exogenous variation in competition and starting position allows us to causally test the escape-competition and Schumpeterian effects.

Third, the time horizon was varied across sessions. At the end of each period, the game would end with a fixed and known stopping probability p . This procedure enables us to exogenously impose an infinite horizon and exogenously vary the stopping probability p , which can be equivalently interpreted as exogenous variation in the rate of time preference ρ . In some sessions subjects faced a short time horizon – 80% probability of ending the game after each round – while in other sessions the subjects faced a long time horizon – 10% probability of ending the game after each round. In the short time horizon treatment, a game lasted on average 2.2 rounds, whereas in the long time horizon treatment, a game lasted on average 10 rounds. This design feature allows us to test whether the escape-competition and Schumpeterian effects vary conditional on the time horizon.

Table 1 summarizes the four treatment variations. Remember that in each treatment, subjects made first period decisions conditional on being in the laggard state and being in the neck-and-neck state. Once the second period was entered into, the game progressed as follows: If a subject is lagging

		Time horizon	
		long ($p = 0.1$)	short ($p = 0.8$)
Competition	no comp. ($\Delta = 0.5$)	long horizon / no competition	short horizon / no competition
	full comp. ($\Delta = 1$)	long horizon / full competition	short horizon full competition

Table 1: Treatment variation: Infinite horizon experiment

behind, that subject is given the right to invest.¹¹ If the sector is neck-and-neck, it is randomly determined at the beginning of the round which subject is able to invest. To make sure that subjects understood the experiment well, in particular the random stopping rule, they practiced the game against a computer opponent for a period of 3 minutes before the experiment started.¹²

The infinite horizon experiment is specifically designed to address the escape competition and the Schumpeterian effects. In order to obtain clean causal evidence, exogenous assignment of subjects to neck-and-neck and laggard states in different competition and time horizon treatments is absolutely crucial. Because subjects are only truly randomly assigned to the neck-and-neck or laggard state in the first period of the interaction, we restrict our attention in the analysis to first period investments. To allow for some learning, we repeated each competition treatment three times. So overall, each subject made six first-round investment decisions as a neck-and-neck firm and six first-round investment decisions as a laggard firm, three for each competition treatment.

Once a game ends, subjects are re-matched with another subject with whom they have not been matched with previously. More specifically, we

¹¹This design feature is equivalent of the automatic catch-up assumption in the theory, that is, firms cannot be more than one technological level apart.

¹²They were informed that the computer's investment decisions would be determined randomly, i.e., nothing could be learned from the computer's strategy. If a game ended within the 3 minutes, they were informed about the final outcomes of that game and a new game would start. This procedure allowed them to get familiar with the computer interface and in particular with the random stopping rule, so that they could form expectations about the length of the game.

designed matching groups and divided subjects within a matching group into a group A and a group B. Each group A subject would only be matched with group B subjects from the same matching group, but no subject would be matched twice with the same subject. 200 points were exchanged to SFr. 1 at the end of the experiment, and subjects were endowed with 5000 points at the beginning of the session, to ensure them against potential losses.

3.3 The finite horizon experiment

The finite horizon experiment involved three different competition treatments: No Competition ($\Delta = 0.5$), Intermediate Competition ($\Delta = 0.75$) and Full Competition ($\Delta = 1$). We again employed a within subjects design, i.e., all subjects participated in all three treatments. This treatment variation is summarized in the table below:

Treatment	Competition
No Competition	$\Delta = 0.5$
Intermediate Competition	$\Delta = 0.75$
Full Competition	$\Delta = 1$

Table 2: Treatment variation: Finite horizon experiment

Our main interest in the finite horizon experiment is in analyzing predictions that stem from the dynamics within sectors over an extended number of periods. An infinite horizon would have yielded a large number of practical problems in terms of conducting the experiment as well as in analyzing the data. The finite horizon experiment consisted of 50 periods in each treatment. To simulate an infinite game with expected time horizon of 50 periods, the probability of ending the game after each period would have needed to be only 2%. This would have implied an expected standard deviation in treatment length of 50 periods. Consequently, the finite horizon experiment is better suited to address the causal effect of competition on industry competition and aggregate innovation.

At the beginning of each treatment, all subjects were assigned to technology levels $\tau_i = 0$, i.e. all sectors started neck-and-neck. One subject was then randomly determined to invest in the first period. Thereafter, subjects alternated in their ability to invest. There was no automatic catch-up, i.e., subjects could theoretically move more than one technological step apart

from each other. This is consistent with the most general version of the model presented in Aghion et al. (2001) and generates the same predictions as simpler versions of the model.¹³

To be ensured against potential losses, each subject was endowed with 3000 points at the beginning of the first period. The exchange rate from points to SFr was 300:1.¹⁴ Subjects within a session are initially divided into group A and group B. In each treatment, a subject from group A is matched with a subject from group B. Between treatments, subjects are randomly rematched with another subject from the other group whom they have not been matched with previously.

3.4 Experimental procedures

Between 18 and 22 subjects participated in each experimental session. In total, four experimental sessions of the infinite horizon experiment were conducted. To control for treatment order effects, each potential sequence of the two competition treatments was used in one session both for the long and for the short time horizon. Moreover, six experimental sessions of the finite horizon experiment were conducted. Again, to control for treatment order effects, sessions were designed such that each potential sequence of the three treatments was used in one session.

The experiments were programmed and conducted with the software z-Tree (Fischbacher (2007)). All experimental sessions were conducted at the experimental laboratory of the Swiss Federal Institute of Technology (ETH) in Zurich. Our subject pool consisted primarily of students at the University of Zurich and ETH Zurich and were recruited using the ORSEE software (Greiner (2004)). The finite horizon experiment took place in February 2012, and the infinite horizon experiment took place in December 2013. 118 subjects participated in the finite horizon experiment, and 86 subjects participated in the infinite horizon experiment. Payment was determined by the sum of the final amounts of points a subject received in all treatments played

¹³See Aghion et al. (2001) for a generalization of the model in Section 2 which allows for more than one gap and derives similar results using calibrations. It is noteworthy, however, that in about 80% of our experimental observations subjects are at most one technological gap apart from each other.

¹⁴Because of the expected number of periods each subject played in the finite and infinite experiments, the exchange rate was modified in order to provide subjects with appropriate earnings for participation in the experiments.

during a session. In addition, each subject received a show-up fee of 10 SFr. On average, an experimental session lasted 1.5 hours. The average payment in the finite horizon experiment was 45 SFr (\$50.00), and 40.8 SFr (\$44.70) in the infinite horizon experiment.

4 Results

In this section, we present our empirical results and compare them to the above predictions. We will first discuss the effects of competition on R&D investments in leveled and unleveled sectors by reporting the findings from the infinite horizon experiment. Thereafter, we will present the effects of competition on industry composition as well as the overall effect of competition on innovation outcomes resulting from the finite horizon experiment.

4.1 Evidence from the infinite horizon experiment

4.1.1 The escape-competition effect

Increased competition should have a positive effect on R&D investments if firms are neck and neck. Empirically, we find in our experiment:

Result 1 (*Escape-competition effect*): *An increase in competition leads to a significant increase in R&D investments by neck-and-neck firms.*

Figure 1 shows the average first round investments in neck-and-neck states in the infinite horizon experiment by competition and time horizon. It can be seen that average first-round R&D investment in the full competition treatment is approximately 10 percentage points (or 35.2 percent) higher than in the no-competition treatment when the time horizon is long, and approximately 6 percentage points (or 28 percent) higher when the time horizon is short. To test the significance of each of these two differences, we use a one-sided clustered version of the signed-rank test proposed by Datta and Satten (2008), which controls for potential dependencies between observations. Recall that the game was repeated three times (repetition) per treatment. Thus, we elicited three first-round R&D investment decisions per subject (conditional on being neck-and neck) per competition treatment, in either of the two time horizons. For each individual and repetition, we then calculated the difference between first round R&D investment in the full

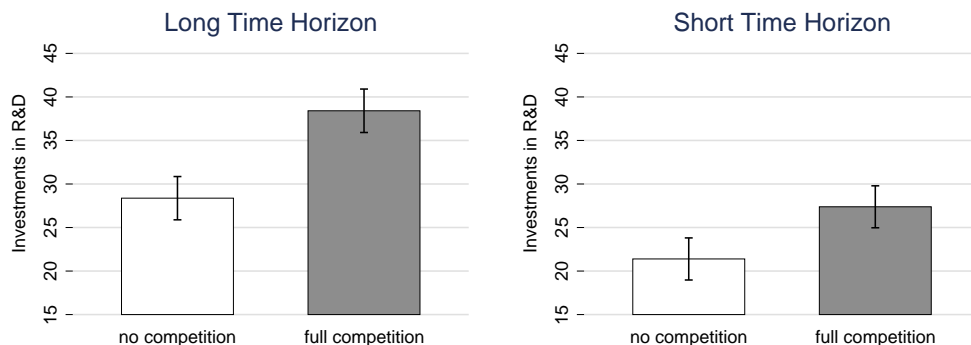


Figure 1: **Average R&D investments in neck-and-neck industries.** Averages are calculated using the average individual first round investments in neck-and-neck states in each treatment. The bars display one standard deviation of the mean.

competition treatment and in the no competition treatment. This generates three observations per subject. Clustering at the individual level, we find that these differences are highly significant in both the short and long time horizons ($p < 0.01$). Thus, consistent with the theory, we find an escape-competition effect for both time horizons. This result confirms the causal nature of the positive effect of competition on the R&D investments of neck-and-neck firms.

Furthermore, we can test if the size of the escape-competition effect varies with the time horizon. According to the theory, the escape-competition effect should be stronger in the long time horizon than in the short time horizon. To test this prediction, we compare the differences in first round R&D investments across competition treatments, as defined above, between subjects that participated in the short-horizon treatment and subjects that participated in the long-horizon treatment, using the one-sided clustered version of the rank-sum test proposed by Datta and Satten (2005). While we do observe a larger effect of competition on investment in the long time horizon than in the short time horizon (10 versus 6 percentage points), clustering at the individual level, we find that this difference is not statistically significant ($p = 0.26$).

Result 2 (*Escape-competition effect and time horizon*): *We find no significant difference in the escape-competition effect between the short time horizon and the long time horizon.*

Table 3: Neck-and-neck investments

	(1)	(2)	(3)
full comp.	8.139*** (2.091)	6.148** (2.340)	6.729** (2.917)
long horizon	9.005*** (2.766)	7.060** (3.443)	9.880** (3.774)
full comp.*long horizon		3.890 (4.117)	5.316 (4.888)
Constant	21.243*** (2.500)	22.216*** (2.481)	21.462*** (3.027)
Adj. R^2	0.097	0.098	0.148
Observations	516	516	172

First round investments in neck-and-neck states are used as observations. Standard errors are clustered at the individual level. Regressions (1) and (2) consider all three repetitions. Regression (3) only considers the third repetition. Treatment order fixed effects and repetition fixed effects (only in regressions (1) and (2)) are included. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

Similar results are found by running OLS regressions, which are reported in Table 3. It can be seen that the escape-competition effect is very robust across different regression specifications. More specifically, average first-round investments in the full competition treatment are between 6 and 8 percentage points higher than average first-round investments in the no competition treatment, and the difference is highly significant. This holds whether we use all first round observations, or whether we restrict our analysis to the third repetition only, after some learning has taken place. Finally, notice that the coefficient of the interaction term of competition and time horizon is large and positive in all regression specifications. However, we cannot establish significance.¹⁵

¹⁵In Figure A.1 in Appendix A, we also show the R&D investments of neck-and-neck firms in the finite horizon experiment, conditional on the competition treatment. It can be seen that R&D investments in the full competition treatment are larger than in the intermediate competition treatment, and R&D investments in the intermediate competition

4.1.2 The Schumpeterian and anticipated escape-competition effects

Increased competition is expected to have a negative effect on R&D investments if firms are lagging behind. Empirically, we find in our experiments:

Result 3 (*Schumpeterian effect*): *An increase in competition decreases R&D in*




Figure 2: **Average R&D investments of laggards in unleveled industries.** Averages are calculated using the average individual first round investments of laggards in unleveled states in each treatment. The bars display one standard deviation of the mean.

Figure 2 shows the average first round investments of laggard firms in unleveled states in the infinite horizon experiment, divided by competition and time horizon. It can be seen that average first round R&D investments in the full competition treatment is approximately 6 percentage points (or 20 percent) lower than in the no competition treatment when the time horizon is long, and approximately 11 percentage points (or 42.3 percent) lower

treatment are larger than in the no competition treatment throughout all 50 periods. This data pattern is also consistent with the predictions of the escape competition effect. However, since firms are not randomly assigned to neck-and-neck states as the rounds progress in the finite horizon experiment, we restrict the causal analysis to the infinite horizon experiment in which the variation in the time horizon, the variation in the degree of competition as well as the assignment to the neck-and-neck state are all fully exogenously varied.

when the time horizon is short. To test the significance of each of these two differences, we again use the one-sided clustered version of the signed-rank test proposed by Datta and Satten (2008). As for the neck-and-neck case, we collected three first round R&D investment decisions per subject (conditional on being a laggard) per competition treatment, in either of the two time horizons. For each individual and repetition, we then calculate the difference between first round R&D investment in the full competition treatment and in the no competition treatment. This generates three observations per subject. Clustering at the individual level, we find that these differences are highly significant in both the short and long time horizons ($p < 0.01$). Thus, consistent with the theory, we find evidence of a Schumpeterian effect in unleveled sectors.

Furthermore, by comparing the effect of competition on laggards' investment across time horizons, we find:

Result 4 (*Anticipated escape-competition effect*): *The Schumpeterian effect is stronger the shorter a firm's time horizon.*

According to the theory, the Schumpeterian effect should decrease as the time horizon increases. To test this prediction, we compare the differences in first round R&D investments across competition treatments, as defined above, between subjects who participated in the short horizon treatment and subjects who participated in the long horizon treatment, again using the one-sided clustered version of the rank-sum test (Datta and Satten (2005)). Clustering at the individual level, we find that this difference is statistically significant ($p = 0.02$).

Similar results are found by running OLS regressions. These are reported in Table 4. Average first round investments in the full competition treatment are between 10 and 13 percentage points lower than average first round investments in the no competition treatment, when the time horizon is short. Consistent with the theory, when the time horizon is long, the Schumpeterian effect is less pronounced by approximately 4.8 to 5.5 percentage points, as reflected in the statistical significance of the interaction term in regressions (2) and (3). This significance is, however, only marginal ($p < 0.1$).¹⁶

¹⁶In Figure A.2 in Appendix A, we also show the R&D investments of laggard firms with a lag of 1 in the finite horizon experiment, conditional on the competition treatment. It can be seen that, in early periods, R&D investments are increasing in the degree of

Table 4: Laggard investments

	(1)	(2)	(3)
full comp	-8.170*** (1.474)	-10.621*** (1.686)	-12.902*** (1.909)
long horizon	6.232*** (2.122)	3.838 (2.440)	3.067 (2.779)
full comp*long horizon		4.788* (2.864)	5.516* (3.256)
Constant	25.414*** (1.761)	26.611*** (1.753)	26.576*** (2.037)
Adj. R^2	0.113	0.117	0.136
Observations	516	516	172

First round investments in laggard states are used as observations. Standard errors are clustered at the individual level. Regressions (1) and (2) consider all repetitions. Regression (3) only considers the third repetition. Treatment order fixed effects and repetition fixed effects (only in regressions (1) and (2)) are included. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

4.2 Evidence from the finite horizon experiment

4.2.1 The composition effect

The exogenous variation of competition across treatments in the finite horizon experiment allows us to also identify the causal effect of competition on industry composition as well as on aggregate innovation outcomes.

According to the theory, we should observe a larger fraction of sectors being neck-and-neck the smaller the degree of competition. This is indeed what we find and is summarized below.¹⁷

competition, i.e., they are largest in the full competition treatment. As time progresses and the time horizon gets shorter, the pattern reverses. In late periods, R&D investments are largest in the no competition treatment. This data pattern is consistent with a strong anticipated escape-competition effect that dominates the Schumpeterian effect in early periods, i.e. when firms face a long time horizon. On the other hand, as time progresses and the time horizon becomes shorter, the anticipated escape-competition effect loses strength and the Schumpeterian effect starts to dominate. However, since firms are not randomly assigned to laggard states as the rounds progress in the finite horizon experiment, we again restrict the causal analysis of these effects to the infinite horizon experiment.

¹⁷Remember that a sector is equal to one duopoly in the experiment, formed by two subjects.




Figure 3: **Average frequency of neck-and-neck states.** The frequency of neck-and-neck states in a sector over 50 periods constitutes one observation. The bars display one standard deviation of the mean.

Result 5 (*Composition effect*): *As competition increases, sectors become less likely to be neck and neck, and subjects are more likely to be technologically apart from each other.*

Table 5: Composition effect

	(1)
Δ	-0.205 *** (0.044)
Constant	0.588 *** (0.075)
Adj. R^2	0.019
Observations	177

A sector in a competition treatment accounts for one observation in these regressions. Regression (1) uses the frequency of neck-and-neck states in a sector during the 50 periods as the outcome variable. Treatment order fixed effects are included. Standard errors are clustered on the session level. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

Evidence for the composition effect can be seen in Figure 3, which shows the average fraction of periods in which sectors were neck and neck, condi-




Figure 4: **Average technology level reached within a sector.** The maximum level of technology reached in a sector over 50 periods constitutes one observation. The bars display one standard deviation of the mean.

tional on the degree of competition. As the Figure demonstrates, the frequency of observing leveled sectors decreases by approximately 5 percentage points as the degree of competition in the industry increases by 0.25.

To assess the statistical significance of these effects, we compare the frequency of neck-and-neck states in a sector across the different competition treatments using regression analysis. The dependent variable is the fraction of observed neck-and-neck states within a sector. Individual and treatment order fixed effects are included. Results are shown in Table 5. We find that when our competition measure increases by 0.25, the relative frequency of sectors being neck and neck decreases by 5.1 percentage points, and this decrease is highly significant.¹⁸

4.2.2 The effect of competition on aggregate outcomes

Next, we can look at the effect of competition on aggregate R&D investments in our finite horizon experiment.

¹⁸Figure A.3 in Appendix A also reports the time path of the frequency in which duopolies are neck-and-neck, across the three competition treatments. It can be seen that throughout all time periods, the frequency of duopolies that are neck-and-neck is higher in the full competition treatment than in the intermediate or in no competition treatment.

Result 6 (*Aggregate innovation*): *Competition increases average R&D investments and, as a result, the average level of technology that is ultimately reached in a sector in our experiment.*

Figure 4 shows the average final technology level of the leading firm within a sector across competition treatments. The figure shows that the average final technology level of the leading firm increases by 0.5 points, from 11.7 to 12.2, when competition increases from no competition to intermediate competition. The average final technology level increases by another 1.4 points to 13.5 when competition increases to full competition. Again, we can evaluate the statistical significance of these effects using regression analysis. The regression includes treatment order fixed effects, and standard errors are clustered on the session level. Results are shown in Table 6.

Table 6: Overall technological progress

	(1)	(2)
	Max. Tech. Level	Avg. R&D Investment
Δ	3.597** (1.124)	11.926*** (3.821)
Constant	9.812*** (1.243)	28.300*** (3.156)
Adj. R^2	0.007	0.012
Observations	177	8850

A sector in a competition treatment accounts for one observation in these regressions. Regression (1) uses the maximum final technology level in a sector as the outcome variable. Regression (2) uses all individual R&D investment choices as observations. Treatment order fixed effects and individual fixed effects are included. Standard errors are clustered on the session level. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

Column (1) in Table 6 provides strong empirical support for Result 6. As Δ increases by 0.25 points, the final technology level of the leading firm increases by 0.9 points, and this increase is significant. Column (2) in Table 6 shows results from a regression of R&D investments on the degree of competition. This regression uses all observations and includes individual and treatment order fixed effects. Therefore, the coefficient on Δ can be interpreted as the average effect of competition on R&D investments over the course of our 50 round experiment. It can be seen that R&D investments on average increase by 3 percentage points (roughly 10 percent) as Δ increases

by 0.25 points. Hence, our setup provides evidence of a positive impact of competition on R&D investments as well as on the maximal technology level reached within a sector.

5 Conclusion

In this paper, we provided a first attempt at analyzing the effect of competition on step-by-step innovation in the laboratory. Using the lab instead of field data has several advantages. First, it addresses the endogeneity issue head on: our results do capture causal effects of competition on innovation. Second, the lab experiment allows us to disentangle the effects of competition on innovation in leveled and unleveled sectors. In particular, we find strong evidence of an escape-competition effect in neck-and-neck sectors. Third, our design allows us to study how the effect of competition on innovation varies with the time horizon. Consistent with theory, we show that the Schumpeterian effect is stronger in the short horizon treatment than in long horizon treatment, suggesting that in the latter case, an anticipated escape-competition effect is also at work. Fourth, we are able to identify the causal effect of competition on industry composition and on aggregate outcomes. We find that, as competition increases, sectors become less likely to be neck and neck, and the average technology level of the leading firm increases.

The methodology used in the paper can be used to sort out other open debates in industrial organization and law & economics. For example, one could use experiments to study the effects of patent protection or R&D subsidies,¹⁹ or the relative performance of various intellectual property legislations or the impact of various antitrust policies,²⁰ or the effects of various contractual or institutional arrangements, on innovation and entry. The industrial organization and law & economics literatures often point to counteracting effects without always spelling out the circumstances under which one particular effect should be expected to dominate. We believe that lab experiments can fill this gap by providing more precise predictions as to when such or such effect should indeed dominate. This and other extensions of our analysis in this paper are left to future research.

¹⁹In Acemoglu and Akcigit (2012), this corresponds to our parameter h .

²⁰Experimental testing may be particularly useful given the homogeneity of intellectual property regimes across countries resulting from international harmonization efforts.

References

- Acemoglu, D. and U. Akcigit (2012). Intellectual Property Rights: Policy, Competition and Innovation. *Journal of the European Economic Association* 10, 1–42.
- Aghion, P., U. Akcigit, and P. Howitt (2013). What Do We Learn From Schumpeterian Growth Theory? Working Paper 18824, National Bureau of Economic Research.
- Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005). Competition and Innovation: an Inverted-U Relationship. *Quarterly Journal of Economics* 120, 701–728.
- Aghion, P. and R. Griffith (2006). *Competition and Growth: Reconciling Theory and Evidence*. MIT Press.
- Aghion, P., C. Harris, P. Howitt, and J. Vickers (2001). Competition, Imitation and Growth with Step-by-Step Innovation. *Review of Economic Studies* 68, 467–492.
- Aghion, P., C. Harris, and J. Vickers (1997). Competition and Growth with Step-by-Step Innovation: An Example. *European Economic Review* 41, 771–782.
- Aghion, P. and P. Howitt (2009). *The Economics of Growth*. MIT Press.
- Blundell, R., R. Griffith, and J. V. Reenen (1999). Market Share, Market Value and Innovation in a Panel of British Manufacturing Firms. *Review of Economic Studies* 66, 529–554.
- Breitmoser, Y., J. H. Tan, and D. J. Zizzo (2010). Understanding Perpetual R&D Races. *Economic Theory* 44 (3), 445–467.
- Darai, D., D. Sacco, and A. Schmutzler (2010). Competition and Innovation: an Experimental Investigation. *Experimental Economics* 13, 439–460.
- Datta, S. and G. Satten (2005). Rank-Sum Tests for Clustered Data. *Journal of the American Statistical Association* 471 (1), 908–915.
- Datta, S. and G. Satten (2008). A Signed-rank Test for Clustered Data. *Biometrics* 64 (2), 501–507.

- Fischbacher, U. (2007). z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics* 10, 171–178.
- Galasso, A. and M. A. Schankerman (2013). Patents and Cumulative Innovation: Causal Evidence from the Courts. <http://ssrn.com/abstract=2247011>.
- Greiner, B. (2004). The Online Recruitment System ORSEE: A Guide for the Organization of Experiments in Economics. Discussion Papers on Strategic Interaction 2003–10, Max Planck Institute of Economics, Strategic Interaction Group.
- Hall, B. H. and D. Harhoff (2012). Recent Research on the Economics of Patents. *Annual Review of Economics* 4, 541–565.
- Harris, C. and J. Vickers (1985a). Patent Races and the Persistence of Monopoly. *Journal of Industrial Economics* 33, 461–481.
- Harris, C. and J. Vickers (1985b). Perfect Equilibrium in a Model of a Race. *Review of Economic Studies* 52, 193–209.
- Harris, C. and J. Vickers (1987). Racing with Uncertainty. *Review of Economic Studies* 54, 1–21.
- Hart, O. D. (1980). Perfect Competition and Optimal Product Differentiation. *Journal of Economic Theory* 22, 279–312.
- Isaac, R. M. and S. S. Reynolds (1988). Appropriability and Market Structure in a Stochastic Invention Model. *Quarterly Journal of Economics* 103, 647–671.
- Isaac, R. M. and S. S. Reynolds (1992). Schumpeterian Competition in Experimental Markets. *Journal of Economic Behavior & Organization* 17, 59–100.
- Jaffe, A. B. (2000). The U.S. Patent System in Transition: Policy Innovation and the Innovation Process. *Research Policy* 29, 531–557.
- Murray, F. and S. Stern (2007). Do Formal Intellectual Property Rights Hinder the Free Flow of Scientific Knowledge? An Empirical Test of the Anti-commons Hypothesis. *Journal of Economic Behavior & Organization* 63, 648–687.

- Nickell, S. J. (1996). Competition and Corporate Performance. *Journal of Political Economy* 104, 724–746.
- Sacco, D. and A. Schmutzler (2011). Is There a U-shaped Relation Between Competition and Investment? *International Journal of Industrial Organization* 29, 65–73.
- Schmidt, K. M. (1997). Managerial Incentives and Product Market Competition. *Review of Economic Studies* 64, 191–214.
- Schmutzler, A. (2009). Is Competition Good for Innovation? A Simple Approach to an Unresolved Question. *Foundations and Trends in Microeconomics* 5(6), 355–428.
- Schmutzler, A. (2013). Competition and Investment: A Unified Approach. *International Journal of Industrial Organization* 31(5), 477–487.
- Suetens, S. and J. Potters (2007). Bertrand Colludes More than Cournot. *Experimental Economics* 10(1), 71–77.
- Tirole, J. (1988). *The Theory of Industrial Organization*. MIT Press.
- Vives, X. (2008). Innovation and Competitive Pressure. *Journal of Industrial Economics* 56, 419–469.
- Williams, H. (2013). Intellectual Property Rights and Innovation: Evidence from the Human Genome. *Journal of Political Economy* 121, 1–27.
- Zizzo, D. J. (2002). Racing with Uncertainty: a Patent Race Experiment. *International Journal of Industrial Organization* 20(6), 877 – 902.

Appendix A: R&D investments over time in the finite horizon experiment




Figure A.1: **Neck-and-neck investments over time.** 2-round averages shown by the dots for the respective competition treatment, quadratic time trend shown by the solid lines for each competition treatment. In total, 118 subjects participated in each treatment.

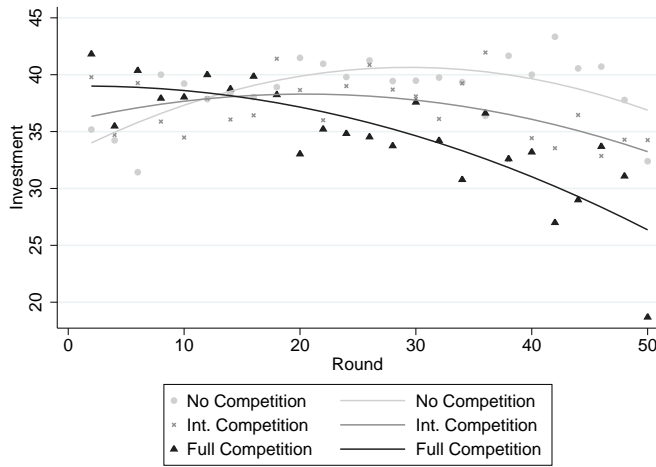


Figure A.2: **Laggard investments over time.** 2-round averages shown by the dots for the respective competition treatment, quadratic time trend shown by the solid lines for each competition treatment. In total, 118 subjects participated in each treatment. Only observations with a lag of 1 are included.




Figure A.3: **Composition effect over time.** 2-round averages shown by the dots for the respective competition treatment, quadratic time trend shown by the solid lines for each competition treatment. In total, there were 59 sectors in each treatment.