

The Causal Effects of Competition on Innovation: Experimental Evidence[†]

Philippe Aghion, Collège de France, France
Stefan Bechtold, ETH Zurich, Switzerland
Lea Cassar, University of Cologne, Germany
Holger Herz, University of Fribourg, Switzerland

published in 2018 in the Journal of Law, Economics and Organization
<https://doi.org/10.1093/jleo/ewy004>

Abstract

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.

1 Introduction

The relationship between competition and innovation has long been of interest to economists, and over the past twenty five years substantial progress has been made to understand this relationship, both theoretically and empirically (e.g., Arrow, 1962; Reinganum, 1983, 1985; Vickers, 1986; Hart, 1980; Schmidt, 1997; Aghion et al., 2001; Vives, 2008; Schmutzler, 2009, 2013; Nickell, 1996; Blundell et al., 1999; Aghion et al., 2005; Aghion and Griffith, 2006). However, the existing empirical studies on competition and innovation suffer from at least two limitations. First, they suffer from endogeneity problems: the degree of competition is influenced by the rate of innovation, and vice versa. Finding exogenous variation in the degree of competition in the

[†]philippe.aghion@college-de-france.fr; sbechtold@ethz.ch; lcassar@uni-koeln.de; holger.herz@unifr.ch. We would like to thank Daniel Chen, Paolo Falco, Ernst Fehr, Bernhard Ganglmair, David Huffman, Klaus Schmidt, Roberto Weber, the editor and two anonymous referees, as well as participants at seminars and conferences at Barcelona GSE, ETH Zurich, and the universities of California at Berkeley, Carlos III of Madrid, Pennsylvania, Oxford and Zurich for helpful feedback.

field in order to identify a causal effect of competition on innovation is difficult. Where such attempts have been made, doubts with respect to the identification of the effect remain.¹ Second, the relationship between competition and innovation is moderated by several factors that are not directly observed in the field, such as firms' rate of time preference or the technological gap between firms within an industry at any point in time. The lack of these data implies that field studies cannot fully disentangle how the relationship between competition and innovation is moderated by these factors.

This is, however, of particular importance since endogenous growth models with step-by-step innovation (Aghion et al., 1997, 2001; Aghion and Howitt, 2009) predict that increased competition should have three types of effects on innovation incentives. First, it 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". Second, competition should have a negative (short-run) "Schumpeterian effect" on laggard firms' innovation incentives in unleveled sectors: increased competition reduces the post-innovation rents of laggard firms and thus their incentive to catch up with the leader. The interplay of these effects implies the "composition effect", that is, that sectors are less likely to be neck-and-neck the larger the degree of competition within a sector. Third, competition should have a positive "anticipated escape-competition effect" on laggard firms' innovation incentives in unleveled sectors: namely, a far-sighted laggard firm anticipates that once it has caught up with the current leader in the sector, it will gain more from innovating further and thereby becoming the new leader in the sector. Existing empirical studies cannot precisely identify and assess these effects as they can not distinguish between firms with different discount rates, and as they do not track the same firms over time due to these firms switching back and forth between being leaders, neck-and-neck, and laggards in their sector. And even if empirical studies could identify these effects, a firm's technological position within a sector would be endogenous, which again complicates identification of causal effects.

To address these 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 laboratory setting provides exogenous control over key variables, such as the degree of competition and the technological gap between firms, and therefore allows to establish causal relationships. At the same time, the laboratory setting is very stylized and many assumptions of the step-by-step innovation model are directly imposed. In this sense, our paper allows to test some key predictions of the theory, under the condition that some of its basic assumptions are met.²

We design two experiments in which pairs of subjects are matched for a number of periods. In each period, one of the two subjects can choose an R&D investment which

¹For example, competition is sometimes measured by market shares or the number of competitors, which is affected by incumbent innovation through entry and exit of firms (for a discussion, see Blundell et al. (1999); Aghion et al. (2005)). For various empirical approaches to identify causal relationships between patenting activities and innovation, see Murray and Stern (2007); Williams (2013); Galasso and Schankerman (2015); Sampat and Williams (2015); Budish et al. (2015).

²The accuracy of these assumptions needs to be assessed in field studies. See, e.g., Blundell et al. (1999); Aghion et al. (2005).

determines the probability of a successful innovation in that period.³ 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’ 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. We implement the infinite horizon by adding a positive stopping probability to the game: After each period, the interaction between two paired subjects ends with a positive probability. We exogenously vary the expected time horizon across sessions, by varying this stopping probability. Pairs either face a short time horizon – a 80% probability of ending the game after each period – or a long time horizon – a 10% probability of ending the game after each period. This setup allows us to test the anticipated escape competition effect: Far-sighted managers react differently to competition than short-sighted managers. We should therefore 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, our focus is on long run predictions of the theory, and thus our objective in designing this experiment is to generate data from sectors interacting for long and comparable periods of time. Hence, all subjects face 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

³We do not aim to mirror the decision of an inventor, but rather the decision of a manager choosing the budget of an R&D department in order to encourage innovation.

sectors, and, overall, competition increases aggregate innovation. All these results are consistent with the predictions of step-by-step innovation models.

The paper relates to the experimental literature on competition and R&D investments. Isaac 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. Suetens and Potters (2007) find 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. The experimental literature has also studied variations of competition in one-period investment games, such as exogenous cooperation agreements (Østbye and Roelofs, 2013), communication (Suetens, 2005), Bertrand versus Cournot competition (Suetens and Potters, 2007) and commitment (Suetens, 2008). Finally, Cantner et al. (2009) investigate the role of competition versus the desire to enhance productivity in a patent race, focussing on various economic and individual psychological criteria that correlate with innovation decisions.

The main distinction between our study and this experimental literature lies in the core characteristic of the step-by-step innovation models compared to previous Schumpeterian models, namely that the innovation incentives do not depend on the post-innovation rents only, but rather on the difference between post-innovation and pre-innovation rents of incumbent firms. In terms of the experimental design, this means that innovation is cumulative over time and pre-innovation rents are manipulated exogenously. The latter was achieved by varying subjects' starting positions between leveled and unleveled. None of the above-mentioned studies share such characteristics. Hence, these studies can only test the (static) Schumpeterian effect, while our design allows to examine the escape competition and Schumpeterian effects in a dynamic investment environment with different time horizons, and to assess the composition effect.⁴ Hence, to our knowledge, we are the first to design a laboratory experiment to test the overall set of predictions of the step-by-step innovation models.

The remaining part of the paper is organized as follows. Section 2 lays out the basic properties of step-by-step innovation models and their main predictions. Section 3 describes the experimental framework for the two experiments. Sections 4 and 5 present the experimental details and results, and Section 6 concludes.

⁴An exception is Isaac and Reynolds (1992) who compare innovation investments over time in a monopoly with a four-seller market. However, the authors do not distinguish between the escape-competition and the Schumpeterian effect, nor do they test the composition effect of competition. Moreover, Zizzo (2002) and Breitmoser et al. (2010) look at innovation investments in a dynamic environment, e.g. in a multi-stage patent race, but they do not investigate the effect of competition. Neither study varies the time horizon.

2 Theoretical predictions

2.1 Basic setup of step-by-step innovation models

Our predictions are based on the family of growth models with step-by-step innovations (Aghion et al., 1997, 2001, 2005; Aghion and Howitt, 1998, 2009). This section summarizes the main features of these models and provides an intuitive account of their main predictions.⁵

In the basic model setup, an industry consists of two firms which produce the same good and compete over selling the good to a customer. Firms can invest into technology, which lowers their cost of production. In particular, if one firm has better technology (lower production costs), the other firm is driven out of the market.⁶ 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.

Firms can improve their technology by investing in R&D, which costs $\psi(z) = z^2/2$. A neck-and-neck firm (firms are neck-and-neck if the current technology level is the same for both firms) then moves one technological step ahead, with probability $z \in [0, 1]$. Let z be the "innovation rate" or "R&D intensity" of the firm. A laggard firm (a firm is a laggard if its current level of technology is smaller than the one of the other firm (the leader)) 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 $z + h \in [0, 1]$. The crucial assumption at the core of all these models is the step-by-step innovation, namely that a "laggard firm" must catch up with the leader before becoming a leader itself.

2.2 The Schumpeterian, escape-competition, and anticipated escape competition effects

The model setup implies that, in an unlevelled sector, the leader gets the whole market and earns a monopoly rent ($\pi_1 > 0$), whereas the laggard makes 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-and-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 unlevelled sector, so that each firm will earn $\pi_1/2$. More generally, if the degree of product market competition is modeled 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:

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

⁵For a detailed derivation of these predictions, please refer to Appendix A. The Appendix uses a simple version of a step-by-step innovation model whose results are broadly generalizable.

⁶For simplicity, the model assumptions are such that the profit of the firms only depends on the technological gap m between the two firms, and not on the absolute technology level (see Aghion and Howitt, 1998, 2009).

where Δ parameterizes product market competition.

This implies that the research intensities of neck-and-neck and laggard firms, respectively, vary with the measure of competition Δ . Hence, an increase in Δ increases the innovation intensity of a neck-and-neck firm. This is the escape competition effect:

Prediction 1 (Escape-competition effect): Innovation by neck-and-neck firms is 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. Step-by-step innovation models also predict that the escape competition effect varies with the rate of time preference:

Prediction 2 (Escape-competition effect by rate of time preference): The escape-competition effect is weaker for firms with a 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, step-by-step innovation models make predictions about the effect of an increase in competition Δ on the innovation intensity of a laggard. This effect is ambiguous in general. In particular, for a sufficiently high discount rate, 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 the discount rate, i.e. for patient laggards, 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 a 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.

2.3 Composition effect

The fraction of unleveled sectors depends itself upon the innovation intensities in both types of sectors. This fraction is increasing in competition as measured by Δ since a higher Δ increases R&D intensity in neck-and-neck sectors (the escape-competition effect) whereas it tends to reduce R&D intensity in unleveled sectors (the Schumpeterian effect). This positive effect of competition on the 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 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.

Altogether, predictions 1 to 5 represent the main predictions of step-by-step innovation models. Different specific versions of these models have been presented in the literature, proving their generalizability to various setups. The aim of our experimental study is to test the general validity of these predictions in a setup that captures the nature of step-by-step innovation. We explain our experimental framework in the next section.

3 Experimental framework

To assess the escape-competition and the Schumpeterian effect, we conducted an infinite time-horizon experiment with both a long and a short time horizon.⁷ To assess the composition effect, we conducted another experiment with a long, but finite time horizon.

The reason why our model predictions are tested in different experiments is that the optimal experimental setup to identify the respective effects differs in fundamental ways. On the one hand, to cleanly identify the escape-competition and Schumpeterian effects, random assignment to leveled and unleveled sectors is of utmost importance. Further, to assess the impact of the rate of time preference on behavior, the experiment must feature infinite time horizons with varying stopping probabilities after each period. On the other hand, the composition effect relates to the long run properties of a sector, so here observing sectors for long and *comparable* periods of time is key, which is very difficult to achieve with random stopping probabilities. Hence, a long and finite time horizon is best suited to assess the predictions.

⁷Exogenous variation in the time horizon can equivalently be interpreted as exogenous variation in firms' rate of time preference, the preferred interpretation in these models.

We will discuss further specifics of our design choices in the respective sections. In Section 4, we will focus on the infinite time-horizon experiment first in which we identify the escape-competition and the Schumpeterian effect. In Section 5, we will then discuss the finite time horizon experiment in which we analyze the composition effect.⁸

4 Experiment 1: The escape competition and Schumpeterian effect

4.1 The basic step-by-step innovation game

At the beginning of the experiment, two subjects i and j were randomly matched with each other, forming a sector. They subsequently interacted in a computerized step-by-step innovation game for an ex-ante unknown number of periods. After each period, termination of the match was determined by the computer, based on a commonly known stopping probability. The timing of a period was as follows: First, one of the two subjects could make a costly R&D investment. Second, based on the investment, the computer randomly determined whether R&D was successful. If she was successful, the investing subject earned a point. Points were accumulated over all periods, and the point balance between subjects reflected the technological distance between the subjects. Third, payoffs were determined conditional on the point balance. Finally, the computer randomly determined whether the game would stop or another period would be played.

More specifically, one of the two subjects could choose an R&D investment n_k , which determined the probability of a successful innovation in this period. $k \in \{-1, 0\}$ denotes the point position of the investor, with $k = -1$ denoting a laggard position, and $k = 0$ denoting neck-and-neck, i.e. an even point balance. If a subject was in a leading position, she could not invest.⁹ Each subject was informed that the common quadratic R&D cost function was given by

$$C(n_k) = 600 \left(\frac{n_k}{100} \right)^2,$$

and that the R&D costs were carried by the subject choosing the R&D investment. In situations in which subjects were neck-and-neck ($k = 0$), subjects could choose an R&D investment from $n_0 \in \{0, 5, 10, \dots, 80\}$.¹⁰ If the investing subject was a laggard, she was automatically granted with a costless additional innovation probability of $h = 30$, which mimics the existence of technological spillovers from leaders to laggards. Thus, the overall innovation probability was given by $n_{-1} + h$, where n_{-1} could be chosen from $n_{-1} \in \{0, 5, 10, \dots, 50\}$, so that the maximal probability of innovating was still given by $n_{-1} + h = 80$.

⁸Instructions for all experiments and treatments are available in the supplementary material.

⁹We will later discuss the choice of the investor in detail.

¹⁰We 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.

Payments after each period to subject i were determined by the following function:

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

where θ_i and θ_j are the accumulated points of subjects i and j . $C(n_k)$ is equal to zero if a subject did not have the possibility to invest into R&D. Subjects were symmetric, and payments to subject j were determined in exactly the same way. This profit function implies that a leader always earned a monopoly rent of 200 points. If subjects were neck-and-neck, their payoffs depended on the degree of competition Δ in the match, which was one of our treatment variables. Finally, if a subject was lagging behind his competitor, her profit was 0. The costs associated with R&D were always subtracted from the rent earned.

The right to invest in R&D in a given period was determined as follows: If a subject was a laggard, she was automatically selected to innovate. Hence, leaders never had the right to invest in R&D.¹¹ If subjects were neck-and-neck, it was randomly determined which of the two subjects was allowed to innovate, with equal probability.¹²

After each period, each subject received information about the preceding period, in particular the chosen investment level, the success of the innovation, period payments, as well as the current technology level and the overall income of both subjects over the entire interaction. They were then also told whether the game would be played for another period or not.

4.2 Treatment variations

The experiment contained three different treatment variations, some of which within subjects, and others across subjects. Our first treatment variation was to randomly assign subjects to a leveled or an unleveled point balance before the first period of the interaction. This random assignment allowed us to causally identify the escape-competition and Schumpeterian effect using first period data. In subsequent periods, the point balance was endogenous to previous play, prohibiting the identification of causal effects independent of potential selection. For this reason, our focus in this experiment is on first period decisions of subjects.

More precisely, it was randomly determined whether the point balance was equal, $\theta_i = \theta_j = 0$, or whether one of the subjects started as a laggard and the other as a leader (with a point difference of 1 between them). It was common knowledge that a match was equally likely to start out leveled or unleveled, and both subjects were equally likely to be in the laggard or leader position in case the point balance was unleveled.

Before subjects were informed about their starting position and who could invest in the first period, we applied the strategy method to maximize the number of observa-

¹¹This design choice mirrors the automatic catch-up assumption in the basic version of step-by-step innovation models (see section 2.1 and Appendix A), which here is equivalent to not allowing for subjects to innovate in situations in which they were leaders.

¹²A side effect of only allowing one subject to innovate is that it 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.

tions: Both subjects made investment decisions for the case in which they were chosen to invest being a laggard and for the case in which they were chosen to invest being neck-and-neck with the other subject. After choices were made, the computer randomly determined the initial point balance and the investing firm, and the investment choice associated with the chosen state was then automatically implemented. The strategy method was only used for first period investment choices. Thereafter, from period 2 onwards, the game proceeded as described before.

The other two treatment variations are summarized in Table 1. The second treatment variation was the degree of competition when firms were neck-and-neck, measured by Δ . In the *no competition treatment*, Δ was equal to 0.5, so that subjects were able to split the monopoly rent between them when they were neck-and-neck. In the *full competition treatment*, Δ was equal to 1, so that subjects faced perfect competition and rents in neck-and-neck states were 0.¹³

Table 1: Treatment variation: Infinite horizon experiment

		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

Δ is the degree of competition when firms are neck-and-neck. p is the stopping probability after each period.

Our third treatment variation was with respect to the time horizon. In some sessions subjects faced a short time horizon – 80% probability of ending the game after each period – while in other sessions the subjects faced a long time horizon – 10% probability of ending the game after each period. In the short time horizon treatment, a game lasted on average 1.2 rounds, whereas in the long time horizon treatment, a game lasted on average 9.2 rounds. This design feature allows us to test whether the escape-competition and Schumpeterian effects vary conditional on the time horizon.¹⁴

¹³We 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.

¹⁴It is important to note that our step-by-step innovation game is as close as possible to the model described in Appendix A. However, our game is in discrete time. While this adjustment does not lead to closed-form solutions as they are presented in the Appendix, we decided to nonetheless impose this design feature. This is due to significant practical difficulties associated with a continuous time implementation, such as implementing and explaining to subjects a random stopping rule and the relation between R&D investments

Some of our treatments varied within subjects and others across subjects. Our first treatment variation, the starting position, varied within session and within subjects. As explained before, each subject made first period decisions for the case that she could invest in a laggard position and for the case that she could invest in a neck-and-neck position. Our second treatment variation, the degree of competition Δ , also varied within subjects. All subjects participated in the full competition ($\Delta = 1$) and in the no competition ($\Delta = 0.5$) treatment, which were played sequentially in varying order across sessions. Variation in the time horizon, our third treatment variation, was done across subjects because we felt that this is the most complicated treatment variable to vary within a session. Moreover, varying all three dimensions within session would have been an overload of variation and would have likely caused substantial confusion among subjects.

Overall, subjects played the game six times, three times in each treatment. Once a match terminated, subjects had to wait until all matches in a session terminated and were then re-matched with another subject with whom they had not been matched with previously and the game started again.¹⁵

Since we repeated each competition treatment three times, we collected three first-period investment decisions with leveled and three first-period decisions with unleveled point balances for each subject and each competition treatment. 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 three minutes before the experiment started.¹⁶

4.3 Experimental procedures

Between 18 and 22 subjects participated in each experimental session. In total, four experimental sessions 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.

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, 2015). The experiment took place in December 2013, and 86 subjects participated. Payment was determined by the sum of the final amounts of points a subject received in all treatments played during a session. In

and innovation success in a continuous time setting.

¹⁵More specifically, we created two matching groups per session, 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.

¹⁶They were informed that the computer's investment decisions would be determined randomly (if the computer had the possibility to invest into R&D in a period, it would randomly choose one of the available investment probabilities, each being equally likely). The computer's strategy was deliberately random so that 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.

addition, each subject received a show-up fee of 10 SFr. On average, an experimental session lasted 1.5 hours. 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. The average payment was 40.8 SFr (\$44.70).¹⁷

4.4 Results

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.

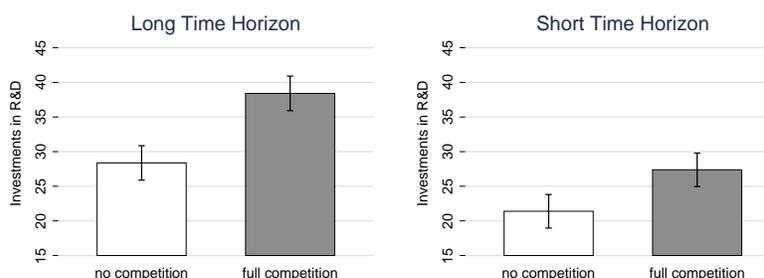


Figure 1: **Average R&D investments in neck-and-neck industries.** Averages are calculated using the average individual first period investments in neck-and-neck states in each treatment. The bars display one standard deviation of the mean.

Figure 1 shows the average first period investments in neck-and-neck states in the infinite horizon experiment by competition and time horizon. It can be seen that average first-period 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-period 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 period R&D investment in the full competition treatment and in the no competition treatment. This

¹⁷Please note that that our experimental setup provided substantial marginal incentives, and potential losses were real. An investment of 80% in one single round cost our subjects SFr 1.90, and it was risky.

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 step-by-step innovation models, 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 period 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 Saten (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): The escape-competition effect is large in magnitude in the long time horizon treatment, but the 4 percentage point difference in the effect is not statistically significant.

Table 2: Neck-and-neck investments

	(1)	(2)	(3)
full competition	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 competition*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 period investments in neck-and-neck states are used as observations. Standard errors are clustered at the individual level (86 clusters). *Full competition* is a dummy variable for the treatment with $\Delta = 1$, in which revenues are zero if firms are neck-and-neck. *Long horizon* is a dummy variable for the treatment with $p = 0.1$, i.e. a 10 percent stopping probability after each period. The baseline is the no competition / short horizon treatment, with $\Delta = 0.5$, i.e., firms share monopoly revenues if they are neck-and-neck and $p = 0.8$. 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 2.¹⁸ It can be seen that the escape-competition effect is very robust across different

¹⁸Table B1 in Appendix B reports results from the same regression with standard errors clustered more conservatively on the matching group level. Significance levels remain similar.

regression specifications. More specifically, average first-period investments in the full competition treatment are between 6 and 8 percentage points higher than average first-period investments in the no competition treatment, and the difference is highly significant. This holds whether we use all first period observations, or whether we restrict our analysis to the third repetition only, after some learning has taken place. Finally, the coefficient of the interaction term of competition and time horizon is large and positive in regression specification (2) and (3). The additional effect of full competition on investment when firms are neck-and-neck is between 3.9 and 5.3 percentage points. However, probably due to limited statistical power, we cannot establish significance of this interaction effect.¹⁹

Next, we turn to the Schumpeterian effect: 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 investments of laggards.

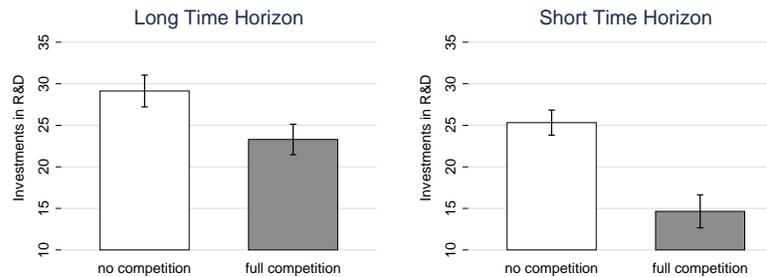


Figure 2: **Average R&D investments of laggards.** Averages are calculated using the average individual first period investments of laggards in unleveled states in each treatment. The bars display one standard deviation of the mean.

Figure 2 shows the average first period 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 period 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

¹⁹In Figure B1 in Appendix B, we also show the R&D investments of neck-and-neck firms in the finite horizon experiment, which is discussed in detail in section 5, conditional on the competition treatment. It can be seen that R&D investments increase in the degree of competition throughout all 50 periods. This data pattern is also consistent with the predictions of the escape competition effect. However, remember that firms are not randomly assigned to neck-and-neck states as the rounds progress, making causal inference impossible.

42.3 percent) lower 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 (Datta and Satten, 2008) and use within subject differences in investments across the competition treatments as observations, as we did in the neck-and-neck case. We find that these differences are highly significant in both the short and long time horizons ($p < 0.01$). Thus, consistent with the predictions of these models, we find evidence of a Schumpeterian effect in unlevelled 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 these models, the Schumpeterian effect should decrease as the time horizon increases. To test this prediction, we compare the differences in first period 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 ($p = 0.02$).

Table 3: Laggard investments

	(1)	(2)	(3)
full competition	-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 competition*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 period investments of laggards are used as observations. Standard errors are clustered at the individual level (86 clusters). *Full competition* is a dummy variable for the treatment with $\Delta = 1$, in which revenues are zero if firms are neck-and-neck. *Long horizon* is a dummy variable for the treatment with $p = 0.1$, i.e. a 10 percent stopping probability after each period. The baseline is the no competition / short horizon treatment, with $\Delta = 0.5$, i.e., firms share monopoly revenues if they are neck-and-neck and $p = 0.8$. 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: *** $p < .01$, ** $p < .05$, * $p < .1$.

Similar results are found by running OLS regressions. These are reported in Table 3.²⁰ Average first period investments in the full competition treatment are between 10

²⁰Table B2 in Appendix B reports results from the same regression with standard errors clustered more conservatively on the matching group level. Significance of the full competition treatment dummy is reduced to the 5% level in this case, and the interaction term between competition and the time horizon only remains

and 13 percentage points lower than average first period 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. The interaction term in regressions (2) and (3) show a reduction of approximately 4.8 to 5.5 percentage points, and these estimates are marginally significant ($p < 0.1$).²¹

Our infinite horizon experiment therefore provides substantial evidence in favor of the escape competition and the Schumpeterian effect. Indeed, investments of laggards are reduced as competition increases, and this effect is significantly more pronounced when the time horizon is short. The escape-competition effect is also very present in our data. However, while the interaction is indeed positive, here we fail to establish a significant interaction between the effect of competition and the time horizon.

5 Experiment 2: The composition effect

As described before, the composition effect is best assessed using a different experimental setup. The composition effect is about long run predictions of these models, and hence requires observation of sectors over long and comparable periods of time. For this reason, experiment 2 features a finite and long time horizon.²² Here, we present the differences in the experimental design as well as our results with respect to the composition effect.

5.1 Experimental design

The experimental design was very similar to the one described in section 4, with a few differences which we explain here. The first difference to the infinite horizon experiment was that the finite horizon experiment involved three different competition treatments: No Competition ($\Delta = 0.5$), Intermediate Competition ($\Delta = 0.75$) and Full Competition ($\Delta = 1$).²³ We again employed a within subjects design, i.e., all subjects participated in all three treatments, and the order of the treatment was varied and balanced across sessions. Each treatment was played exactly once.

significant in column (3).

²¹In Figure B2 in Appendix B, we also show the R&D investments of laggard firms with a lag of 1 in the finite horizon experiment, discussed in more detail in section 5, conditional on the competition treatment. It can be seen that, in early periods, R&D investments are increasing in the degree of 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 restrict the causal analysis of these effects to the infinite horizon experiment.

²²Normann and Wallace (2012) shows that behaviour is highly similar in long finitely-repeated games with a known end and in games with an unknown end that are long in expectation.

²³The intermediate competition treatment was added in order to study potential non-monotonicities in the aggregate relationship between competition and innovation, as in Aghion et al. (2005). To identify the escape-competition and the Schumpeterian effect, however, the theoretical predictions with respect to the degree of competition are monotonic, and hence a third competition treatment was superfluous.

Moreover, subjects started leveled in all treatments, i.e. with $\theta_i = \theta_j = 0$. One subject was then randomly determined to invest in the first period. Thereafter, subjects alternated in their ability to invest, independent of the point balance. This means that also leaders could invest, and we did not impose the automatic catch-up, i.e., subjects could potentially move more than one technological step apart from each other.²⁴

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 300:1.²⁵ Subjects within a session were initially divided into group A and group B. In each treatment, a subject from group A was matched with a subject from group B. Between treatments, subjects were randomly rematched with another subject from the other group whom they had not been matched with previously.

5.2 Experimental procedures

Again, between 18 and 22 subjects participated in each experimental session. In total, six experimental sessions 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, 2015). The experiment took place in February 2012, and 118 subjects participated. Payment was determined by the sum of the final amounts of points a subject received in all treatments played 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 was 45 SFr (\$50.00).²⁶

5.3 Results

The exogenous variation of competition across treatments in the finite horizon experiment allows us to identify the causal effect of competition on industry composition.

According to the step-by-step innovation models, 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.²⁷

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.

²⁴It is noteworthy, however, that in about 80% of our experimental observations, subjects are at most one technological gap apart from each other.

²⁵Because 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.

²⁶Instructions for all treatments are available in the supplementary material.

²⁷Remember that a sector is equal to one duopoly in the experiment, formed by two subjects.

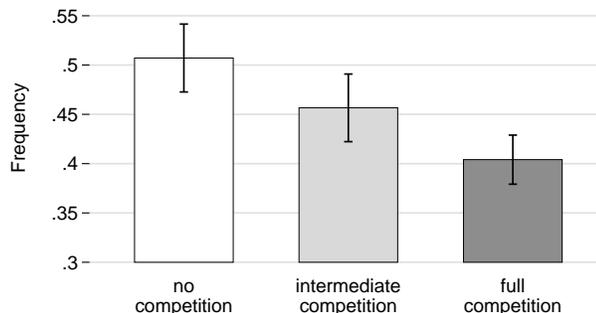


Figure 3: **Average frequency of neck-and-neck states.** The frequency of neck-and-neck states in a sector over 50 periods constitutes one observation. The bars display one standard deviation of the mean.

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, conditional 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.

Table 4: Composition effect

	(1)	(2)
Degree of Competition (Δ)	-0.205*** (0.044)	-0.210** (0.052)
Constant	0.588*** (0.075)	0.578*** (0.087)
Adj. R^2	0.019	0.014
Observations	177	177

The outcome variable in these regressions is the fraction of periods in which a sector was neck-and-neck. Regression (1) uses the frequency of neck-and-neck states in a sector during the 50 periods as the outcome variable. Regression (2) uses the frequency of neck-and-neck states in a sector during periods 11-40. A sector denotes one pair of subjects in one of the three different competition treatments. Each subject participated in all three different competition treatments, each time matched with a different subject. Treatment order fixed effects are included. The degree of competition Δ affects profits when a sector is neck-and-neck. Depending on the treatment $\Delta \in \{0.5, 0.75, 1\}$. Revenues in neck-and-neck sectors are defined by $R_i = (1 - \Delta)200$. Standard errors are clustered on the session level. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

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 4. Column (1) considers the fraction of neck-and-neck states over all 50 periods, and we find that when our competition measure increases by 0.25, the rela-

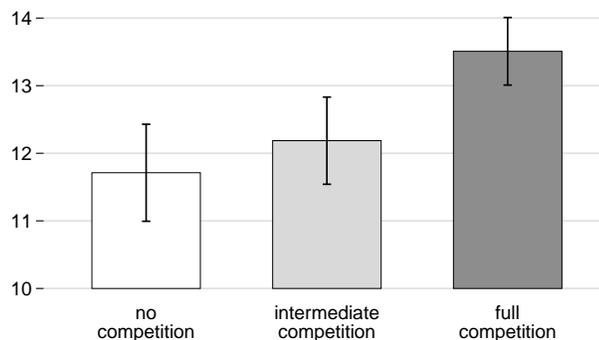


Figure 4: **Average point level reached within a sector.** The maximum level of points reached in a sector over 50 periods constitutes one observation. The bars display one standard deviation of the mean.

tive frequency of sectors being neck-and-neck decreases by 5.1 percentage points. This decrease is highly significant.²⁸ Column (2) accounts for potential end game effects as well as potential effects that stem from early periods, given that sectors started in the same leveled state in period 1 independent of the degree of competition. Therefore, only the fraction of observed neck-and-neck states between periods 11 to 40 within a relationship is considered. The results are very robust to this adjustment and the strongly significant decrease in observed neck-and-neck states is confirmed even when the first and the last 10 periods (which accounts to 40% of all periods) is removed from the analysis.

Finally, we can look at the effect of competition on aggregate R&D investments and profits in our finite horizon experiment. We find that competition increases average R&D investments and, as a consequence, the average level of points that is ultimately reached in a sector in our experiment. Figure 4 shows the average final points level of the leading firm within a sector across competition treatments. The figure shows that the average final points level of the leading firm increases by 0.5 steps, from 11.7 to 12.2, when competition increases from no competition to intermediate competition. The average final points level increases by another 1.4 steps 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 5.

Column (1) in Table 5 provides additional empirical support for this result. As Δ increases by 0.25 points, the final points level of the leading firm increases by 0.9 steps, and this increase is significant. Column (2) in Table 5 shows results from a regression of

²⁸Figure B3 in Appendix B 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 lower in the full competition treatment than in the intermediate or in no competition treatment.

Table 5: Aggregate results

	(1)	(2)	(3)
	Max. Tech. Level	Avg. R&D Investment	Final Profits
Δ	3.597** (1.141)	11.916*** (3.822)	-5040.8*** (389.7)
Constant	10.310*** (0.915)	29.142*** (2.949)	4747*** (352.7)
Adj. R^2	0.07	0.04	0.20
Observations	177	8850	354

The outcome variable in regression (1) is the maximum technological level achieved in a sector after 50 periods. Regression (2) uses all individual R&D investment choices as observations. Regression (3) uses final profits (excluding the endowment) of every subject in each treatment. A sector denotes one pair of subjects in one of the three different competition treatments. Each subject participated in all three different competition treatments, each time matched with a different subject. The degree of competition Δ affects profits when a sector is neck-and-neck. Depending on the treatment $\Delta \in \{0.5, 0.75, 1\}$. Revenues in neck-and-neck sectors are defined by $R_i = (1 - \Delta)200$. Treatment order fixed effects and session fixed effects are included. Standard errors are clustered on the session level. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

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 period 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 points level reached within a sector.

Finally, we can assess the impact of competition on average profits earned by firms in this experiment. In every round, firms could at most earn revenues of 200. Technological progress had, by design, no positive effect on these potential earnings. If competition was fierce, these revenues were reduced to 100 or, depending on the treatment, to 0 if firms were neck-and-neck. Moreover, R&D investments were costly and directly reduced firm profits. Consequently, in our experimental setup, being neck-and-neck has a negative impact on firm profits in the intermediate and full competition treatments, and investing more in R&D has a negative effect on profits as well.

In principle, the effect of competition on final profits is ambiguous. More competition reduces neck-and-neck rents. On the other hand, neck-and-neck states are less likely if competition is higher, and laggards have lower incentives to invest. This saves on R&D costs, while the monopoly rent is independent of the degree of competition. The apparent advantage of higher neck-and-neck rents in the no competition treatment may therefore not necessarily translate into larger average firm profits.

Column (3) in table 5 shows regressions of final profits on the degree of competition. It can be seen that the effect of competition on profits is large and negative. Increasing Δ from 0.5 to 0.75 decreases total profits by approximately 1250 points, and increasing it from 0.5 to 1 decreases total profits by approximately 2500 points.²⁹

²⁹An alternative regression specification with competition treatment dummies confirms that the effect is

Consequently, the negative effects of competition on profits clearly dominate in our setting. Given our previous findings, this is not surprising. While competition indeed reduces the average fraction of neck-and-neck sectors by approximately 10 percentage points, this composition effect is not sufficient to compensate for the reduction in profits when firms are neck-and-neck. Moreover, we have seen that average R&D investments are increasing in competition, implying increasing costs as well. Both these factors jointly imply that profits decrease in the degree of competition.³⁰

6 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 incentives. 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 the prediction of the step-by-step innovation models, we show that the Schumpeterian effect is stronger in the short horizon treatment than in the 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.

Studying innovation in the laboratory inevitably also has its downsides. It grants exogenous control over key variables to study incentive effects in step-by-step innovation models, but the same control implies that many features and assumptions that are made in step-by-step innovation models are directly imposed. Our experiment shows that, given that key assumptions of the step-by-step innovation model are met, the model predictions are met in the empirical data. The laboratory does not provide a test of whether these assumptions are met in reality. We believe that this is an equally important question, but it can only be answered with field data. In this sense, laboratory and field methods are complementary in their strengths to inform theories of competition and innovation.

Moreover, our experiment still treats the actual innovation process as a black box. We are focusing on investments into R&D, assuming that higher investments convert into higher likelihoods to innovate. We therefore do not aim to mirror the decision of an inventor (for such work, see, e.g., Bradler et al. (2016), Erat and Gneezy (2016) and Charness and Grieco (forthcoming)), but rather the decision of a manager choosing the budget of an R&D department in order to encourage innovation.

The experimental methods used in the paper can be used to sort out other open debates in industrial organization as well as law and economics. For example, one could

indeed almost linear in the degree of competition.

³⁰Of course, this analysis ignores potential welfare effects of better technology on consumers, which were not part of our experimental implementation.

use experiments to study the effects of patent protection or R&D subsidies,³¹ or the relative performance of various intellectual property legislations or the impact of various antitrust policies,³² or the effects of various contractual or institutional arrangements, on innovation and entry. The industrial organization as well as the law and 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. They can also be helpful when integrating behavioral insights from decision-making under uncertainty into models of innovation and competition. This and other extensions of our analysis in this paper are left to future research.

³¹In Acemoglu and Acigit (2012), this corresponds to our parameter h .

³²Experimental testing may be particularly useful given the homogeneity of intellectual property regimes across countries resulting from international harmonization efforts.

A The model

In this Appendix, we provide a simplified version of a step-by-step innovation model to provide an intuition on how the predictions on the effects of competition on innovation are generated.

A.1 Basic environment

Time is continuous and there is a continuous measure L of infinitely-lived individuals each of whom supplies 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

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

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 .³³ 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.

A.2 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.

³³To 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.³⁴

For expositional simplicity, it is assumed 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.³⁵

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.

A.3 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-and-neck) firms' profits down to zero. At the other extreme, if the two firms collude so

³⁴The 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).

³⁵See 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.

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 .

A.4 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:³⁶

$$\rho V_0 = \pi_0 + \bar{z}_0(V_{-1} - V_0) + z_0(V_1 - V_0) - wz_0^2/2 \quad (\text{A1})$$

$$\rho V_{-1} = \pi_{-1} + (z_{-1} + h)(V_0 - V_{-1}) - wz_{-1}^2/2 \quad (\text{A2})$$

$$\rho V_1 = \pi_1 + (z_{-1} + h)(V_0 - V_1) \quad (\text{A3})$$

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 (A1) and z_{-1} maximizes the RHS of (A2), we have the first order conditions:

$$wz_0 = V_1 - V_0 \quad (\text{A4})$$

$$wz_{-1} = V_0 - V_{-1}. \quad (\text{A5})$$

In Aghion et al. (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 $w = 1$.

³⁶Note 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.

Then, using (A4) and (A5) to eliminate the V 's from the system of equations (A1)-(A3), we obtain a system of two equations in the two unknowns z_0 and z_{-1} :

$$z_0^2/2 + (\rho + h)z_0 - (\pi_1 - \pi_0) = 0 \quad (\text{A6})$$

$$z_{-1}^2/2 + (\rho + z_0 + h)z_{-1} - (\pi_0 - \pi_{-1}) - z_0^2/2 = 0 \quad (\text{A7})$$

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 (A6) 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 (A7), 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.

A.5 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:

$$\mu_1 = \frac{2z_0}{z_{-1} + h + 2z_0}. \quad (\text{A8})$$

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.

B Additional tables and figures

Table B1: Neck-and-neck investments

	(1)	(2)	(3)
full competition	8.139*** (1.547)	6.148*** (1.335)	6.729*** (1.722)
long horizon	9.005** (2.731)	7.060** (2.221)	9.880*** (2.682)
full competition*long horizon		3.890 (2.530)	5.316 (3.149)
Constant	21.980*** (2.079)	22.952*** (1.984)	21.462*** (2.126)
Adj. R^2	0.100	0.101	0.148
Observations	516	516	172

First period investments in laggard states are used as observations. Subjects within sessions were divided into 2 matching groups, and only interacted with other subjects within their respective matching group. Standard errors are clustered at the matching group level (8 clusters). Regressions (1) and (2) consider all repetitions. Regression (3) only considers the third repetition. *Full competition* is a dummy variable for the treatment with $\Delta = 1$, in which revenues are zero if firms are neck-and-neck. *Long horizon* is a dummy variable for the treatment with $p = 0.1$, i.e. a 10 percent stopping probability after each period. The baseline is the no competition / short horizon treatment, with $\Delta = 0.5$, i.e., firms share monopoly revenues if they are neck-and-neck and $p = 0.8$. Treatment order fixed effects and repetition fixed effects (only in regressions (1) and (2)) are included. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

Table B2: Laggard investments

	(1)	(2)	(3)
full competition	-8.170**	-10.621**	-12.902***
	(2.417)	(3.100)	(1.644)
long horizon	6.232*	3.838	3.067
	(2.679)	(2.336)	(3.007)
full competition*long horizon		4.788	5.516*
		(4.563)	(2.473)
Constant	24.658***	25.855***	26.576***
	(2.486)	(1.877)	(2.391)
Adj. R^2	0.114	0.119	0.136
Observations	516	516	172

First period investments in laggard states are used as observations. Subjects within sessions were divided into 2 matching groups, and only interacted with other subjects within their respective matching group. Standard errors are clustered at the matching group level (8 clusters). Regressions (1) and (2) consider all repetitions. Regression (3) only considers the third repetition. *Full competition* is a dummy variable for the treatment with $\Delta = 1$, in which revenues are zero if firms are neck-and-neck. *Long horizon* is a dummy variable for the treatment with $p = 0.1$, i.e. a 10 percent stopping probability after each period. The baseline is the no competition / short horizon treatment, with $\Delta = 0.5$, i.e., firms share monopoly revenues if they are neck-and-neck and $p = 0.8$. Treatment order fixed effects and repetition fixed effects (only in regressions (1) and (2)) are included. Significance levels: *** $p < .01$, ** $p < .05$, * $p < .1$.

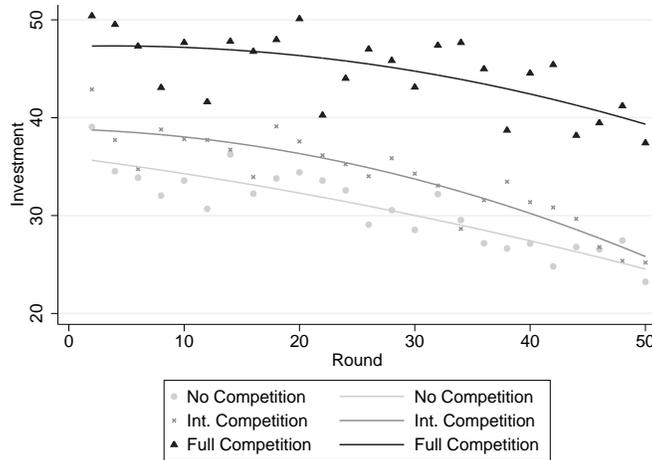


Figure B1: **Neck-and-neck investments over time.** Two-period averages shown by the dots for the respective competition treatment, quadratic time trend shown by the solid lines for each competition treatment. In total, 118 subjects participated in each treatment.

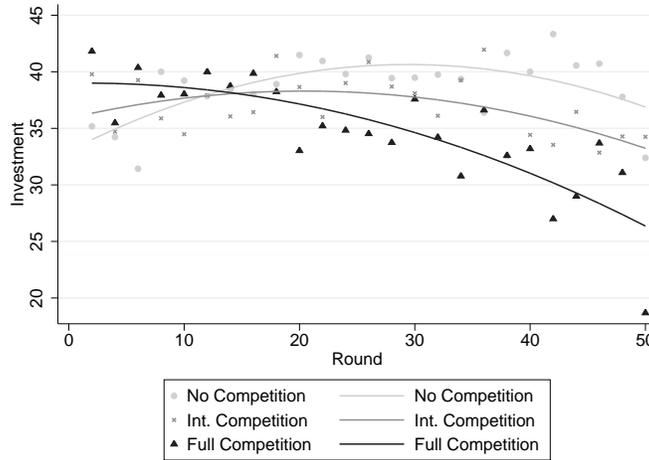


Figure B2: **Laggard investments over time.** Two-period averages shown by the dots for the respective competition treatment, quadratic time trend shown by the solid lines for each competition treatment. In total, 118 subjects participated in each treatment. Only observations with a lag of 1 are included.

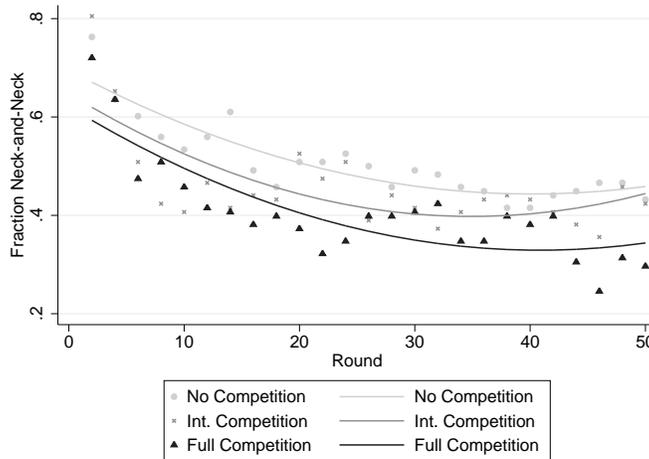


Figure B3: **Composition effect over time.** Two-period averages shown by the dots for the respective competition treatment, quadratic time trend shown by the solid lines for each competition treatment. In total, there were 59 sectors in each treatment.

References

- Acemoglu, Daron and Ufuk Akcigit. 2012. "Intellectual Property Rights: Policy, Competition and Innovation," 10 *Journal of the European Economic Association* 1–42.
- Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt. 2005. "Competition and Innovation: an Inverted-U Relationship," 120 *Quarterly Journal of Economics* 701–728.
- Aghion, Philippe and Rachel Griffith. 2006. *Competition and Growth: Reconciling Theory and Evidence*. Cambridge: MIT Press.
- Aghion, Philippe, Christopher Harris, Peter Howitt, and John Vickers. 2001. "Competition, Imitation and Growth with Step-by-Step Innovation," 68 *Review of Economic Studies* 467–492.
- Aghion, Philippe, Christopher Harris, and John Vickers. 1997. "Competition and Growth with Step-by-Step Innovation: An Example," 41 *European Economic Review* 771–782.
- Aghion, Philippe and Peter Howitt. 1998. *Endogenous Growth Theory*. Cambridge: MIT Press.
- Aghion, Philippe and Peter Howitt. 2009. *The Economics of Growth*. Cambridge: MIT Press.
- Arrow, Kenneth. 1962. "Economic Welfare and the Allocation of Resources for Invention," in Universities-National Bureau Committee for Economic Research and Committee on Economic Growth of the Social Science Research Council, ed., *The Rate and Direction of Inventive Activity: Economic and Social Factors* 609–626. Princeton: Princeton University Press.
- Blundell, Richard, Rachel Griffith, and John Van Reenen. 1999. "Market Share, Market Value and Innovation in a Panel of British Manufacturing Firms," 66 *Review of Economic Studies* 529–554.
- Bradler, Christiane, Susanne Neckermann, and Arne J. Warnke. 2016. "Incentivizing creativity: A large-scale Experiment with Tournaments and Gifts," ZEW Discussion Paper 16-040, Zentrum für Europäische Wirtschaftsforschung.
- Breitmoser, Yves, Jonathan H. Tan, and Daniel J. Zizzo. 2010. "Understanding Perpetual R&D Races," 44 *Economic Theory* 445–467.
- Budish, Eric, Benjamin N. Roin, and Heidi Williams. 2015. "Do Firms Underinvest in Long-term Research? Evidence from Cancer Clinical Trials," 105 *American Economic Review* 2044–2085.
- Cantner, Uwe, Werner Güth, Andreas Nicklisch, and Torsten Weiland. 2009. "Competition in Product Design: An Experiment Exploring Innovation Behavior," 60 *Metroeconomica* 724–752.

- Charness, Gary and Daniela Grieco. Forthcoming. "Creativity and Financial Incentives," *Journal of the European Economic Association*.
- Darai, Donra, Dario Sacco, and Armin Schmutzler. 2010. "Competition and Innovation: an Experimental Investigation," 13 *Experimental Economics* 439–460.
- Datta, Somnath and Glen Satten. 2005. "Rank-Sum Tests for Clustered Data", 100 *Journal of the American Statistical Association* 908–915.
- Datta, Somnath and Glen Satten. 2008. "A Signed-rank Test for Clustered Data," 64 *Biometrics* 501–507.
- Erat, Sanjiv and Uri Gneezy. 2016. "Incentives for Creativity," 19 *Experimental Economics* 269–280.
- Fischbacher, Urs. 2007. "z-Tree: Zurich Toolbox for Ready-made Economic Experiments", 10 *Experimental Economics* 171–178.
- Galasso, Alberto and Mark Schankerman. 2015. "Patents and Cumulative Innovation: Causal Evidence from the Courts," 130 *Quarterly Journal of Economics* 317–369.
- Greiner, Ben. 2015. "Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE," 1 *Journal of the Economic Science Association* 114–125.
- Hart, Oliver D. 1980. "Perfect Competition and Optimal Product Differentiation," 22 *Journal of Economic Theory* 279–312.
- Isaac, R. Mark and Stanley S. Reynolds. 1988. "Appropriability and Market Structure in a Stochastic Invention Model," 103 *Quarterly Journal of Economics* 647–671.
- Isaac, R. Mark and Stanley S. Reynolds. 1992. "Schumpeterian Competition in Experimental Markets," 17 *Journal of Economic Behavior & Organization* 59–100.
- Murray, Fiona and Scott Stern. 2007. "Do Formal Intellectual Property Rights Hinder the Free Flow of Scientific Knowledge? An Empirical Test of the Anti-commons Hypothesis," 63 *Journal of Economic Behavior & Organization* 648–687.
- Nickell, Stephen J. 1996. "Competition and Corporate Performance," 104 *Journal of Political Economy* 724–746.
- Normann, Hans-Theo and Brian Wallace. 2012. "The Impact of the Termination Rule on Cooperation in a Prisoner's Dilemma Experiment," 41 *International Journal of Game Theory* 707–718.
- Østbye, Stein E. and Matthew R. Roelofs. 2013. "The Competition and Innovation Debate: Is R&D Cooperation the Answer?," 22 *Economics of Innovation and New Technology* 153–176.
- Reinganum, Jennifer F. 1983. "Uncertain Innovation and the Persistence of Monopoly," 73 *American Economic Review* 741–748.

- Reinganum, Jennifer F. 1985. "Innovation and Industry Evolution," 100 *Quarterly Journal of Economics* 81–99.
- Sacco, Dario and Armin Schmutzler. 2011. "Is There a U-shaped Relation Between Competition and Investment?," 29 *International Journal of Industrial Organization* 65–73.
- Sampat, Bhaven and Heidi Williams. 2015. "How Do Patents Affect Follow-on Innovation? Evidence from the Human Genome," <http://economics.mit.edu/files/10782>.
- Schmidt, Klaus M. 1997. "Managerial Incentives and Product Market Competition," 64 *Review of Economic Studies* 191–214.
- Schmutzler, Armin. 2009. "Is Competition Good for Innovation? A Simple Approach to an Unresolved Question," 5 *Foundations and Trends in Microeconomics* 355–428.
- Schmutzler, Armin. 2013. "Competition and Investment: A Unified Approach," 31 *International Journal of Industrial Organization* 477–487.
- Suetens, Sigrid. 2005. "Cooperative and Noncooperative R&D in Experimental Duopoly Markets," 23 *International Journal of Industrial Organization* 63–82.
- Suetens, Sigrid. 2008. "Does R&D Cooperation Facilitate Price Collusion? An Experiment," 66 *Journal of Economic Behavior & Organization* 822–836.
- Suetens, Sigrid and Jan Potters. 2007. "Bertrand Colludes More than Cournot," 10 *Experimental Economics* 71–77.
- Vickers, John. 1986. "The Evolution of Market Structure When There is a Sequence of Innovations," 35 *Journal of Industrial Economics* 1–12.
- Vives, Xavier. 2008. "Innovation and Competitive Pressure," 56 *Journal of Industrial Economics* 419–469.
- Williams, Heidi. 2013. "Intellectual Property Rights and Innovation: Evidence from the Human Genome," 121 *Journal of Political Economy* 1–27.
- Zizzo, Daniel J. 2002. "Racing with Uncertainty: a Patent Race Experiment," 20 *International Journal of Industrial Organization* 877–902.