The effects of training on own and co-worker productivity: Evidence from a field experiment

Citation for published version (APA):

de Grip, A., & Sauermann, J. (2012). The effects of training on own and co-worker productivity: Evidence from a field experiment. *The Economic Journal*, 122(560), 376-399. <https://doi.org/10.1111/j.1468-0297.2012.02500.x>

Document status and date:

Published: 01/05/2012

DOI:

[10.1111/j.1468-0297.2012.02500.x](https://doi.org/10.1111/j.1468-0297.2012.02500.x)

Document Version:

Publisher's PDF, also known as Version of record

Document license:

Taverne

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

THE EFFECTS OF TRAINING ON OWN AND CO-WORKER PRODUCTIVITY: EVIDENCE FROM A FIELD EXPERIMENT*

Andries De Grip and Jan Sauermann

This article identifies the effects of work-related training on worker productivity by exploiting a field experiment that randomly assigns workers to treatment and control groups combined with data on worker performance before and after training. We find that participation in the training programme leads to a 10% increase in performance. Moreover, we provide experimental evidence for externalities from training: An increase of 10 percentage points in the share of treated peers improves a worker's performance by 0.51%. Furthermore, we find that the performance increase is not due to lower quality provided by the worker.

Estimating the returns of work-related training on productivity is a challenge for two reasons. First, the non-random selection of workers into training makes it difficult to identify the causal effect of training participation on individual productivity from other unobservable factors that drive participation in training as well as productivity. If this selection is not appropriately accounted for, the causal impact of training programmes on productivity can be overestimated (Bassanini *et al.*, 2007). Second, it is difficult to find appropriate proxy variables for productivity. While wages are often used to estimate returns to training participation, the returns of employer-provided training on productivity should be higher than those on wages (Dearden *et al.*, 2006); however, direct measures of productivity at the individual level are scarce.

Our contribution to the literature is twofold: we are the first to provide experimental evidence on the causal effects of investments in training on productivity using individual performance information and the externalities of training on co-worker productivity. A randomised field experiment enables us to measure the causal effects of investments in training on worker productivity. The field experiment was carried out in the call centre of a multi-national telephone company in the Netherlands. We have detailed information on the contents, length and purpose of the training programme, as well as unique panel data with administrative information on individual performance, to estimate the returns to training. Agents were randomly assigned to treatment and control groups. This exogenous variation in training participation is used to estimate the returns that are causally related to training and not to unobservable factors affecting both training participation and productivity. We find that agents perform 10% better after participating in the training programme than non-treated agents. We

* Corresponding author: Jan Sauermann, Research Centre for Education and the Labour Market (ROA), School of Business and Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, the Netherlands. Email: j.sauermann@maastrichtuniversity.nl.

The authors thank Iwan Barankay, Eric Bonsang, Lex Borghans, Thomas Dohmen, Dan Hamermesh, Ben Kriechel, Olivier Marie, Eric Maurin, Raymond Montizaan and three anonymous referees for helpful comments. This article was presented at Maastricht University, Education in Adulthood and the Labour Market at IAB, IWAAE 2010, EEA 2010, RWI, Advances with Field Experiments (Wharton School), RES 2011, SOLE 2011, IZA Summer School 2011 and ZEW. This article has been screened to ensure no confidential information is revealed.

show that this effect is causal and not the result of employee selection into and out of training. Furthermore, we find that the performance increase is not due to a lower quality of work.

Workers belonging to the treatment group were trained over successive weeks. Moreover, due to the restricted capacity of the training centre, the teams had to be split into two training groups. This time-varying treatment of the agents in a team allows for the further random assignment of agents into a first and a second training group. This enables us to identify possible externalities from training on untreated peers within a team. We find that a 10-percentage point increase in the share of treated agents leads to a performance increase of 0.51% for their untreated teammates.

From a policy point of view, an unbiased estimation of the impact of training on productivity is important for assessing the role of further training in the development of human capital. Previous studies have found rather mixed results, mostly depending on the method applied, as well as the measure of productivity used. We therefore categorise the research into three main strands. A large number of studies focus on the wage effects of work-related training and are extensively surveyed in Bartel (2000), Dearden *et al.* (2006) and Bassanini *et al.* (2007). After correcting for selectivity by applying fixed-effects estimation or instrumental variables regressions, these studies find returns of training ranging from 0% (Goux and Maurin, 2000) to 10% (Parent, 2003). The disadvantage of using wages to measure the effects of training, however, is that wage increases after-training only reflect the returns to the worker. The effect on productivity is thus likely to be underestimated under rent sharing.

The second strand of research uses large-scale surveys across firms, establishments or industries. As a measure of firm productivity, most of these studies use the value added or sales of firms (Black and Lynch, 2001; Dearden *et al.*, 2006; Zwick, 2006; Konings and Vanormelingen, 2010) or direct measures of productivity within one sector of industry (Holzer *et al.*, 1993). Both Dearden *et al.* (2006) and Konings and Vanormelingen (2010) show that the effects of training on firm productivity are about twice those on worker wages. When controlling for various sources of worker heterogeneity, the latter authors find that the productivity premium for a trained employee is, on average, 17%. Although large-scale surveys can have the advantage of providing a representative sample for one or more sectors of industry, they inherently suffer from unobservable heterogeneity in the type as well as duration of training programmes and firm production processes (Ichniowski and Shaw, 2009). Moreover, it is difficult, if not impossible, to find direct measures of productivity that are comparable across industries.

The third strand of research on the effects of training on productivity focuses on just a single firm or establishment, or comparable establishments within a sector of industry (Bartel, 1995; Krueger and Rouse, 1998; Liu and Batt, 2007). Using direct performance measures and wages, they find mostly positive effects of training.

Overall, previous studies have estimated a range of treatment effects of training participation on productivity. At the same time, there is consensus that appropriate correction for selection into training matters for estimation. Although the use of field experiments has sharply increased over the last decade, only a few studies in the human capital literature exploit (quasi) exogenous variations in training participation to estimate the effects of training. Leuven and Oosterbeek (2008) use information on the

reasons for non-participation to create treatment and control groups artificially and find no significant effects from participation in training. The authors conclude that the usual methods to correct for selection into training do still overestimate the true effect of training. Similarly, Schwerdt *et al.* (2011) find no significant effects in their study analysing the effect of randomly assigned government-funded training vouchers.

This study exploits random assignment to training by means of a field experiment to overcome selectivity in training participation. By randomly assigning agents to treatment and control groups, we can estimate the average treatment effect (ATE) of training participation on individual productivity. We use panel data on individual worker productivity to estimate the effects of the training programme. To measure productivity, we use the key performance indicator used by the call centre to evaluate its call agents, that is, the average time needed to handle inbound customer calls. This performance measure has also been used in other studies in the call centre sector (Liu and Batt, 2007; Breuer *et al.*, 2010). We measure each worker's performance each week before and after the training, which allows us to analyse short-run performance dynamics in the weeks after training. Using individual productivity as an outcome allows us to capture the total effect of training on productivity, while wage information captures only the share an agent receives.¹

Despite a growing literature on peer effects in the workplace, the externalities of training have hardly been discussed in the literature on the impact of training on productivity.² The literature on human capital externalities originates from growth theory and has been applied to human capital spillover effects from workers' level of education (Moretti, 2004). When the returns to training are estimated at the individual (worker) level without taking knowledge spillovers into account, returns will be underestimated. This also holds for a second type of externalities, peer effects due to social pressure (Falk and Ichino, 2006). The identification of externalities, however, is empirically difficult. Previous studies used either quasi-random variations in group composition (Mas and Moretti, 2009; De Grip *et al.*, 2011) or the random assignment of subjects to groups (Guryan *et al.*, 2009). While many field experiments have exploited changes in management behaviour induced by firm management (Bandiera *et al.*, 2007; Shaw and Lazear, 2008), we exploit the structure of the field experiment to estimate externalities. To achieve identification, teams were randomly split into two training groups. This random assignment of agents from the same team allows us to estimate the causal estimates of within-team externalities.

The training we analyse was a one-week course to train call agents in conversation techniques designed to decrease the average handling time of calls while maintaining call quality. Call centres have become a major sector of employment due to strong growth rates since the 1980s, facilitated by the increasing availability of information and communication technology infrastructure. For the US, Batt *et al.* (2005) estimate that

¹ A second advantage of using individual productivity instead of annual productivity data is that studies using annual data are more likely to underestimate the causal effects of training. If the yearly data do not contain information on the training period and use average productivity over the previous year, the returns to training are likely to be underestimated, even if selectivity is properly controlled for.

² To the best of our knowledge, the only exception is the study of Dearden *et al.* (2006) who aim to estimate externalities by comparing individual-level wage regressions with industry-level wage regressions that should capture externalities.

call centres employed about 4 million employees, roughly 3% of the total workforce. Although the current trend is to outsource call centres, most call centres are in-house (Batt *et al.*, 2009). Work-related training is an important element of the call centre industry. In general, call agents receive hardly any initial vocational training, whereas the heavy use of information technology in in-house centres requires high investments in work-related training. From this perspective, we expect the returns on training in call centres will be high (Sieben *et al.*, 2009).

The remainder of this article is structured as follows: Section 1 provides an overview of the firm analysed in this article and its workers and describes the experiment in detail. Section 2 presents our regression model and the estimation results. Additional evidence and tests on the robustness of our results are discussed in Section 3. Finally, Section 4 summarises the article and presents concluding remarks.

1. Context of the Experiment

1.1. Organisation of the Department and Worker Tasks

The field experiment analysed in this article was implemented in an in-house call centre of a multi-national mobile network operator in the Netherlands. The call centre acts as a service centre for current and prospective customers. It has five departments segmented by customer group. To ensure a homogeneous production process, we focus on the largest department, that for private customers with fixed cell phone contracts.

The main task of call agents in this department is to answer customer phone calls. Customers contact customer service when they have problems, complaints or questions. All agents take part in a training course when entering the department that enables them to handle basic types of calls. Throughout their careers, agents receive further training. These training programmes mainly focus on information in promotional campaigns, communication and information technology skills, as well as on handling more complex calls.

In the first period of our sample, 179 individuals were working in the department. Column (1) of Table 1 shows the descriptive statistics of all agents working in the department at the beginning of the period of observation. Most agents are part-timers. The average number of hours worked per week is 18.0, and only 22.4% of agents work 25 hours (h) or more each week. A total of 29.1% of all agents are men, the average age of the agents is 32.6 years and the agents have an average of 2.7 years of experience working for the firm.

The agents are organised into 10 teams. In general, all teams provide all services, that is, there is no team specialisation in handling certain types of calls or customers. All teams work on the same floor of the building. Work places are organised into work islands, with up to eight agents of a team sitting next to each other. Each team is led by a team leader who reports to the head of the department. The team leaders' main task is supervising and monitoring the agents of their team. They evaluate their agents regularly based on performance scorecards containing detailed information on key performance indicators. Agent pay is based on a single collective agreement. Agents are paid a base wage and receive an annual wage raise of zero to 8%, depending on an

Table 1
Selection into the Experiment

Agents	Selection into the experiment ($N = 179$)				Initial assignment ($N = 86$)				Estimation sample ($N = 74$)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)		
	All	Assigned	Non-assigned	Diff (2)–(3)	TG	CG	Diff (5)–(6)	TG	CG	Diff (8)–(9)		
Gender	0.2905	0.2791	0.3011	–0.0220	0.3513	0.2245	0.1269	0.3824	0.2750	0.1074		
(share of male agents)	(0.4553)	(0.4512)	(0.4612)		(0.4840)	(0.4216)		(0.4933)	(0.4522)			
Age	32.595	35.575	29.431	6.143***	34.245	36.578	–2.332	34.070	36.146	–2.076		
	(11.325)	(11.062)	(10.792)		(10.017)	(11.791)		(10.095)	(11.640)			
Tenure	2.646	4.312	1.106	3.206***	4.233	4.372	–0.139	4.071	3.773	0.298		
(in years)	(3.610)	(4.019)	(2.310)		(3.796)	(4.218)		(3.810)	(3.984)			
Performance (y_{it})	0.3307	0.3716	0.2928	0.0789***	0.3715	0.3717	–0.0002	0.3596	0.3582	0.0013		
	(0.1108)	(0.1017)	(0.1058)		(0.0884)	(0.1116)		(0.0779)	(0.0893)			
Call quality (y_{it}^q)	0.4541	0.4676	0.4416	–0.0260	0.4759	0.4609	0.0150	0.4802	0.4571	0.0230		
	(0.1244)	(0.1447)	(0.1010)		(0.1812)	(0.1087)		(0.1844)	(0.1089)			
Composite productivity (y_{it}^p)	0.2529	0.2873	0.2209	0.0664***	0.2846	0.2895	0.0049	0.2759	0.2780	–0.0022		
	(0.0864)	(0.0818)	(0.0783)		(0.0687)	(0.0916)		(0.0583)	(0.0705)			
Average working hours	18.000	16.593	19.301	–2.708**	15.270	17.592	–2.322	15.706	17.350	–1.644		
	(8.455)	(8.058)	(8.646)		(8.875)	(7.317)		(9.104)	(7.277)			
Share peak hours	0.5541	0.5304	0.5773	–0.0468*	0.5445	0.5197	0.0248	0.5553	0.5286	0.0266		
	(0.1766)	(0.1946)	(0.1547)		(0.1972)	(0.1939)		(0.1926)	(0.1970)			
Number of agents	179	86	93		37	49		34	40			

Notes. Significance levels: difference significant at * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$; two-sided t-test. Standard deviations are shown in parentheses. Descriptive statistics are calculated before agents were assigned to the training (week 50/2008). TG, treatment group; CG, control group.

annual performance rating by the team leader. There are no other individual or team-related bonuses.

1.2. *Training Purpose, Contents and Organisation*

The training programme analysed in this study was intended for all agents of the department with some experience on the job. Its aim was to increase the efficiency of agents answering customer calls. Agents who were recently recruited were not selected for this training (yet). The management had decided to organise the training to decrease the average time needed for handling calls, because the call centre had been performing below the targets set.

The training was organised as a week-long programme. Due to capacity constraints, only one group, with a maximum of 10 agents, could be trained at a time. The training took place in an in-house training centre located on a different floor in the same building. The programme consisted of 10 half-day training sessions that were held on five consecutive days, from Monday to Friday. In half of these sessions, agents were either formally trained by a coach or had group discussions assisted by a coach and their team leader. In these group discussions, agents discussed the skills they lacked, how their skills could be improved and how agents could help each other on the work floor. In formal sessions, agents were trained in conversation techniques designed to decrease the average time needed for handling customer calls without lowering quality. This included, for instance, the way in which call agents gathered information from customers to resolve calls quickly.

The remaining sessions consisted of learning by doing, by either handling regular customer calls or listening to the calls of other call agents. During these sessions, incoming customer calls were routed to the training centre and agents handled these calls under the supervision of the training coach and their team leader.

1.3. *The Field Experiment*

1.3.1. *Identification strategy*

In the economic literature on estimating the effects of training, the most evident problem is the potential correlation of unobservable factors with both training participation and the outcome variable. This study uses exogenous variations in training participation to identify the causal effect of training on individual productivity.

As shown in Figure 1, the field experiment consists of three periods. We observe 32 weeks, from week 45/2008 to week 24/2009. At the beginning of the first period (pre-experiment), which lasted 17 weeks, agents were assigned to treatment and control groups. In this period, neither the treatment group agents nor the control group agents were trained. During the second, experiment period, the treatment group was trained consecutively over five weeks. After the experiment period (post-experiment, lasting 10 weeks), agents from the treatment group as well as their untreated peers from the control group worked as usual. Because the agents of the control groups were trained after the post-experiment period (from week 25/2009 onwards), we use data from week 45/2008 to week 24/2009 only. Agents who are part of the control group thus never participate in training throughout our observation period.

Fig. 1. *Overview of the Field Experiment*

Notes. This Figure shows the design of the field experiment. The observations start in week 45/2008 and end in week 24/2009. The assignment to treatment and control groups took place in week 50/2008. Training of the treatment group was held in the experiment period.

The training had been designed for agents with some experience in the job. Of the 179 individuals working in the department at the beginning of our observation period (week 50/2008), 86 were selected for the training.³ Management focused the training programme on the more tenured agents for two reasons: first, agents should be experienced to benefit from this type of training and, second, management was aiming to lower the risk of losing its training investments through turnover.

Column (4) of Table 1 shows the differences between assigned and non-assigned agents. Assigned agents are an average of six years older, have longer tenure by more than three years and perform better. While this selection is clearly non-random, it does not violate the assumption that assignment to the treatment group is exogenous. This is because assignment to the treatment and control groups is exogenous conditional on being assigned to the training programme. An advantage of the focus on more experienced agents is that the calls they handle are relatively homogeneous: Agents with longer tenure deal with all types of calls, whereas starting agents do not work on all types of calls yet.

Conditional on being selected for the training programme, 37 of the 86 agents were randomly selected for participation in the treatment group. The remaining 49 agents, who were assigned to the control group, were trained after the post-experiment period, that is, after the end of the observation period. The differences in observable characteristics between the assigned treatment and control groups are relatively small, with none significantly different from zero (column 7 of Table 1). For the estimation, however, agents were excluded if the management reassigned them from the control group to the treatment group, or *vice versa*, because they could not participate in the planned training week due to vacation or illness. Three agents from the treatment group and nine agents from the control group are thus excluded from the estimation sample. Column (10) of Table 1 shows that there is no statistical difference in observable variables between the treatment and control groups of our estimation sample.⁴

³ Agents not selected initially were assigned to training weeks after the training of the treatment group. Because these agents could not be randomly assigned to treatment and control groups, we do not consider them to be part of the control group. The effects we find therefore only hold for more experienced agents. As Sieben *et al.* (2009) show, training courses for more experienced agents are highly relevant in the call centre sector.

⁴ In addition, all tests were conducted by using the non-parametric Wilcoxon-rank test, which is suited for comparing discrete measures between small groups (Wilcoxon, 1945). The results do not differ from those of the two-sided t-test.

The resulting estimation sample is an unbalanced panel of 74 agents over 32 weeks. In total, we have 1,859 agent-week observations on performance. On average, agents work 28 weeks out of the 32 weeks. The time-variant treatment in our experiment implies that the treatment and control groups differ during each specific week, that is, agents belong to the control group until they are treated. For this reason, teams were not only randomly assigned to the treatment and control groups but also randomly assigned to the various weeks during the treatment period.

1.3.2. *Identification of externalities*

Although agents carry out individual tasks, social interactions can arise, influencing the treatment effect.⁵ Agents belonging to the same team sit next to each other. Furthermore, they can communicate with their co-workers during team meetings and breaks. In a survey held among the agents, 80% stated that they exchange work-related information either during official team meetings or at their workplace. Therefore, the likelihood of externalities is higher within teams than across teams. We do not preclude externalities between teams; however, given the physical distance on the work floor, these externalities should be less than externalities within teams.

Our field experiment allows us to identify these externalities of the training on untreated peers. For this purpose, we exploit the fact that 8 of the 10 teams had to be split into separate training groups to match the size restriction of training groups in the firm's training centre, and the agents in these teams were randomly assigned to the two training groups.

Divided teams can thus be distinguished in the first and second training groups. In between the training weeks of the two training groups in which a team had been split, one or two groups of a different team were trained. During these weeks, the team consisted of exogenously chosen agents who were treated and randomly chosen agents who were not yet treated. This further randomisation enables us to distinguish between three groups of agents: those who were trained first (first training group); their teammates who were trained later (second training group); and the agents in the teams in the original control group who will not be trained in the period we analyse. We exploit the fact that the agents of one team were randomly assigned to training groups to identify within-team externalities. Table 2 shows the means of observable characteristics for treatment group agents in a team's first training group ($N = 13$) *versus* the means for agents who are in a team's second training group ($N = 21$).⁶

All agents were informed about the training at the moment of their initial assignment. Throughout the experiment, information about the randomisation of the order of the teams was given only to the head of the training department and the department manager. No information was given to the team leaders or the agents about the

⁵ The production technology, however, is not necessarily without externalities. Agents may follow a strategy of providing low quality to score well on the main performance criteria (average handling time). If low-quality calls lead to a higher rate of customers calling repeatedly, other agents can be affected if they have to take more care of those customers. In practice, however, agents cannot easily shirk by providing low quality with a low average handling time because the quality of agents is regularly monitored by their team leaders, who can listen in on the calls. Moreover, if there was any shirking by providing low quality, these agents should be similarly distributed to the treatment and control groups.

⁶ Both the t-test as well as the Wilcoxon-rank test do not show significant differences between the groups. It should be noted, however, that the sample size for this comparison is rather small.

Table 2
*Descriptive Statistics of Teams' First and Second Training Groups
 (Treatment Group Only)*

Agents	(1)	(2)	(3)
	First group	Second group	Diff (2)–(3)
Gender	0.3846	0.3810	0.0037
(share male agents)	(0.5064)	(0.4976)	
Age	35.659	33.085	2.575
	(10.759)	(9.799)	
Tenure	3.707	4.297	–0.590
(in years)	(3.978)	(3.784)	
Performance (y_{it})	0.3687	0.3539	0.0148
	(0.0685)	(0.0843)	
Call quality (y_{it}^q)	0.4526	0.4972	–0.0446
	(0.0798)	(0.2269)	
Composite productivity (y_{it}')	0.2844	0.2706	0.0139
	(0.0538)	(0.0616)	
Average working hours	18.539	13.952	4.586
	(10.674)	(7.743)	
Share peak hours	0.5181	0.5783	–0.0602
	(0.2242)	(0.1720)	
Number of agents	13	21	

Notes. Significance levels: difference significant at * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$; two-sided t-test. Standard deviations are shown in parentheses. Descriptive statistics are calculated before agents were assigned to the training (week 50/2008).

randomisation or the evaluation of the training programme. There was no need to do this, since all employees know that the size of the training centre is restricted. For this reason, the training department always makes a training schedule in collaboration with the department managers when they start a new training programme. The training centre is occupied most of the time. During the post-experiment period, agents from other departments of the call centre were trained in the training centre.

1.4. Measuring Productivity

We measure agent productivity by means of the call centre's key performance indicator: the average handling time, defined as the average time an agent needs to handle a customer call (Liu and Batt, 2007; Breuer *et al.*, 2010). The handling time includes the time needed to talk to the customer, as well as the time needed to log the call in the customer database. Information on the average handling time of each individual agent i is available for each working week t . This performance measure does not contain information on the quality of calls. Section 2 introduces additional indicators of call quality and discusses the quality–quantity relation.

The department's aim is to improve performance by decreasing average handling time. Throughout this study, we therefore use the inverse of the average handling time (aht_{it}) multiplied by 100: $y_{it} = (1/aht_{it}) \cdot 100$. Since lower values of aht_{it} are interpreted as higher performance, this transformation allows us to interpret improvements in y_{it} as increases in individual productivity. The average performance of all agents assigned to training is 0.3307, which relates to 5.6 min for an average call (see column 1 of Table 1).

There is substantial heterogeneity in individual performance within and between agents. This suggests that not only individual-specific characteristics but also other, department-specific effects, such as technical problems, affect the individual productivity of all agents working in the department. Compared to individual heterogeneity in productivity, however, variation over time is less important: while period fixed effects explain only about 12% of the overall variation in individual performance, worker fixed effects alone explain 43% of the overall variation in performance.

It is essential that our measure of productivity is comparable within and between agents. Calls are randomly assigned to agents. Agents have no direct influence on the types of calls they receive or the types of customers put through to them, and, therefore have the same probability of exceptionally long or short calls. Before talking to agents, however, customers must first state the purpose of their call. Based on this information, calls are routed to agents who have sufficient knowledge to resolve the call. The assignments of agents to types of calls can be changed at any time by management. Agents are often reassigned *ad hoc* if the structure of customer calls changes.

This assignment of calls to agents does not violate the assumption that our performance measure is comparable across time and agents. First, because agents are exogenously assigned to the treatment and control groups, the skill distribution of agents and thus the types of calls should, on average, be the same in the treatment and control groups. Second, we compare an agent's performance before and after the training intervention. Because the training did not focus on resolving different types of customer requests, calls assigned to agents after the treatment do not systematically differ from those assigned prior to the training.

2. Empirical Analysis

2.1. Causal Effect of Training on Individual Performance

We observe an agent i 's performance y_i and training participation d_i . The observed outcome can thus be written as

$$y_i = y_i(d_i) = d_i y_i(1) + (1 - d_i) y_i(0), \quad (1)$$

where $y_i(1)$ and $y_i(0)$ denote performance in the treated and untreated states, respectively. The randomised assignment of agents to the treatment and control groups ensures the independence of treatment status and potential outcomes $E[y_i(j) \mid d_i = 0] = E[y_i(j) \mid d_i = 1]$ for $j = 0, 1$. The ATE τ (Rosenbaum and Rubin, 1983) is thus identified by $\tau = E[y_i(1) - y_i(0)]$. The ATE can be estimated by performing a linear regression of individual i 's performance y_{it} in week t on a treatment dummy d_{it} which is defined as being one in each after-training period, and zero otherwise:

$$y_{it} = \alpha_i + \tau d_{it} + \beta_1 t + \beta_2 \mathbf{X}_{it} + \beta_3 X_t + u_{it}, \quad (2)$$

where α_i are individual fixed effects to account for individual heterogeneity remaining despite the experimental design, and $\mathbf{X}_{it}$ are covariates that are assumed to be independent from the treatment status d_{it} , such as working hours in week t and the type

of shift the agent is working.⁷ The variable t_t is a linear time trend that controls for trends in aggregate performance affecting all agents, and u_{it} is an idiosyncratic error term. To control for aggregate effects on performance (X_t), we include the overall number of calls divided by the total number of full-time equivalent agents. Throughout this analysis, standard errors are clustered at the agent level.⁸

2.2. Baseline Results

We provide first evidence of the treatment effect in Figure 2, which shows the average performance of agents in the treatment group. The treatment week is denoted week 0, with positive (negative) values of the x axis showing the t th calendar week after (before) the training. Agent performance appears to be, on average, higher in the weeks after the training than before the training. When the performance of untreated agents is not controlled for, treated agents perform significantly better after the training. Given the random assignment of agents to the treatment and control groups, this can also be shown by comparing the mean performance of agents in the treatment group with that of agents in the control group. Table 3 shows that agents in the treatment group are significantly more productive after the training, while there are no significant pre-treatment differences between the two groups.

Table 4 shows the results when estimating (2). The treatment dummy is defined as being one in all weeks after an agent has been trained and zero otherwise. While agents from the control group thus always have a treatment dummy equal to zero, the share of treated agents increases in time with the growing number of groups that have been trained. Column (1) of Table 4 shows that agent performance after participation in the training is 11.3% higher than before the training, controlling for untreated agents' performance. When, in addition, individual heterogeneity is controlled for by including worker fixed effects, the effect increases slightly to 12.4% (column 2). Figure 2, however, suggests that aggregate trends matter. When controlling for a linear time trend, the effect decreases by about 3.6% to an estimated effect of 8.8% (column 3 of Table 4).⁹ The time trend itself is significant, suggesting that there is an overall positive trend in performance during the sample period.

The estimates in column (3) of Table 4 show that the number of working hours do not have a significant impact on the performance y_{it} . To control for possible differences in the nature and amount of calls during peak hours, we use a variable that contains the share of peak hours in the total number of hours agent i worked in week t . The estimates show that a 10-percentage point increase in the share of peak hours leads to a decrease in performance of 3.4%. This suggests that the calls during peak hours are

⁷ We do not use agent tenure as a covariate in our regressions for two reasons. First, because the agents in our sample are all relatively experienced, there is not sufficient variation in individual tenure to control for tenure *and* for common time trends. Second, previous research shows that experience effects matter only for new agents (De Grip *et al.*, 2011).

⁸ In panels increasing in t with fixed N , the appropriate assumptions about standard errors can be crucial to the significance of the results (Bertrand *et al.*, 2004). The results here were re-estimated, allowing for serial correlation in the standard errors. This does not change the size or significance of the results.

⁹ To allow for more flexibility in time trends common to all agents, we also estimated (2) including linear and squared time trends. The estimated treatment effect $\hat{\tau}$ then decreases slightly. The difference between the two estimates is not significantly different from zero.

Fig. 2. *Average Performance of the Treatment Group Before and After the Training*
Notes. This Figure shows the average performance of agents in the treatment group in the weeks before and after the training week (solid line). The dashed lines show the 95% confidence interval. The training week is denoted as week 0.

Table 3
Average Performance Post-treatment Period

	(1)	(2)	(3)
	Treatment group	Control group	Difference TG–CG
Average performance (predicted)	0.4148 (0.0143)	0.4027 (0.0157)	0.0122***
Number of agents	28	35	
Number of observations	296	409	

Notes. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$; two-sided t-test. Standard deviations in parentheses. Predicted performance is generated from a regression of performance on working hours of agent i in week t plus a linear time trend.

more time consuming than the calls at other times of the day.¹⁰ To control for overall pressure in week t , we furthermore control for the effect of the overall number of incoming calls in week t , divided by the number of full-time equivalent agents. This variable, however, does not significantly affect the performance of agents.

The preceding analysis makes the standard assumption that participation in training leads to a persistent shift in performance by including a dummy for participation in training. Given the weekly performance information at hand, we can exploit dynamic patterns of the treatment effect. Figure 3 shows the estimated shape of the treatment effect in the weeks after the training. Panel (a) of Figure 3 is based on an estimation of (2), where the treatment variable d_{it} has been replaced by a set of dummies for each post-training week ($\sum_{t'=1}^T d_{it'}$), where $t' = 1$ denotes the first week after the training. The untreated state $d_{it} = 0$ serves as the reference. It shows that the increase in performance in the first week after the training is not significantly different from zero.

¹⁰ We also estimated regressions with similar variables to control for the timing of working hours, such as dummies for different working hours. These estimates show very similar results.

Table 4
Direct Treatment Effect

	(1)	(2)	(3)
Treatment dummy	0.1127*** (0.0295)	0.1244*** (0.0152)	0.0882*** (0.0206)
Working hours	0.0017 (0.0018)	-0.0012 (0.0010)	-0.0015 (0.0010)
Share peak hours	-0.2564*** (0.0930)	-0.3486*** (0.0739)	-0.3416*** (0.0749)
Calls per FTE	0.0002*** (0.0001)	0.0002** (0.0001)	0.0001 (0.0001)
Time trend			0.0023** (0.0009)
Constant	-1.0008*** (0.0782)	-0.8770*** (0.0582)	-1.0000*** (0.0888)
Individual fixed effects	No	Yes	Yes
Observations	1,859	1,859	1,859
Number of agents	74	74	74
R ²	0.0689	0.6113	0.6165

Notes. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable: $\log(y_{it})$. All standard errors are clustered at the agent level.

In the weeks thereafter, agent performance is significantly higher than before the training. Soon after, however, the estimated treatment effect gradually decreases and becomes insignificant from about the eighth week after the training. This decrease suggests that the estimate $\hat{\tau}$ can only be interpreted as an average effect over the whole post-treatment period.

While this result seems to suggest that the treatment effect declines over time, there may be other mechanisms at work. From a human capital perspective, it is less reasonable that human capital acquired during the training depreciates within a short number of weeks. Moreover, the declining treatment effect may be due to a decline in the agents' motivation. From this perspective, however, it is difficult to understand that workers are least productive in the first week after the training. When the treatment effect would be driven by motivation, there should be a high positive performance effect in the first week. It is more likely that the decreasing treatment effect is driven by training externalities. If (yet) untreated agents are affected by the training of their peers, either due to knowledge spillover or peer pressure, their performance will increase when their co-workers have been treated. In this case, the treatment effect measured should be highest in the weeks shortly after the training and then constantly diminish, because agents from the control group are not actually fully untreated but indirectly affected by the training as well. Panel (b) of Figure 3 takes account of possible training externalities by showing the estimated shape of the treatment effects in the weeks after the training for the agents in the first training groups only, with the agents who will not be trained throughout the observation period as the control group. This Figure shows a similar pattern as Panel (a). The treatment effect, however, remains significant for a longer period. Both the treatment and control groups in this setting are not affected by within-team externalities suggesting that Panel (a) exhibits contamination through externalities leading to a decreasing treatment effect over time.

Fig. 3. *Development of the Treatment Effect on Performance Over Time*

Notes. This Figure shows the estimated treatment effect on performance for each week after the training, controlling for a linear time trend. Panel (a) shows the estimated treatment effect from a regression including all agents. Panel (b) shows the treatment effect for first groups only. Week 0 denotes the training week. The dashed lines show the 95% confidence interval estimated from a regression of y_{it} on after-training week dummies.

2.3. *The Role of Social Interaction*

The preceding subsection estimates the ATE from participation in the training programme in the way that is standard in the training literature. An unbiased estimate $\hat{\tau}$, however, requires the stable unit treatment value assumption, hereafter SUTVA (Angrist *et al.*, 1996). The most important implication of the SUTVA is that there are no externalities from treated workers on untreated workers; that is, the control group's performance [$y(0)$] is not affected by the treatment. In settings where individuals potentially interact, the SUTVA is violated. In the presence of externalities on untreated agents, the observed outcome changes to

$$y_i = y_i(d_i) = d_i y_i(1) + (1 - d_i) (1 + d_i^* \tau_s) y_i(0), \quad (3)$$

where d_i , $y_i(0)$ and $y_i(1)$ are defined as in (1). In addition, we allow the observed performance in the untreated state to be affected by the treatment of workers in the same team. The indicator d_i^* is defined as being one if a worker is untreated but exposed to treated workers from the same team. If an agent is untreated ($d_i = 0$) and working in a team in which all agents are untreated ($d_i^* = 0$), the potential outcome from (3) simplifies to that from (1). The same holds for treated agents who participated in the training ($d_i = 1$ and $d_i^* = 0$). The indirect treatment effect τ_s scales the size of the externality. If we assume that externalities are non-negative and affect untrained agents only, (3) implies that, unless $\tau_s = 0$, the estimated treatment effect ($\hat{\tau}$) is underestimated in (2). If direct and indirect treatment effects are stable, the true treatment effect is underestimated by τ_s .

In practice, the identification of social effects is cumbersome due to endogeneity (Manski, 1993). In our setting, however, we can exploit the fact that team agents were randomly assigned to two separate training groups. In between the training of two groups belonging to the same team, treated and untreated agents of the same team worked together for one or two weeks. During these weeks, a randomly selected group of the team that was treated worked with the remainder of the team that was untreated. Untreated agents of these teams were thus exposed to treated peers. Assuming that there are no peer effects between teams, the random assignment enables us to identify within-team externalities. We can do this in both a direct and an indirect way. The first column of Table 5 shows the results of a direct test of externalities. Here, we test whether the share of treated peers affects the performance of agents who were not yet treated. We therefore use precise information about an individual agent's shift to calculate the share of treated agents among all the agents with whom the agent worked during that shift. The share of treated peers is calculated for peers from the same team only. The results show that a 10% increase in the share of treated peers leads to a performance increase of 0.51%. This effect, however, is only statistically significant at the 10% level of significance.¹¹

The second training groups having already increased their productivity through the externality effect of the first group being trained, we can expect a smaller additional effect when the second training groups finally receive the training themselves. We therefore test whether $\hat{\tau}$ is lower for agents who were exposed to treated agents before their own training. Column (2) of Table 5 shows the estimated treatment effect for agents who comprise the first group of their team to be treated, compared to that for the control group of agents who were in teams in which none of the agents was treated throughout the observation period. The estimated effect is slightly higher (9.9%) than when estimating (2) for all agents (8.8%; see column 3 of Table 4). This suggests that the causal effect from training is underestimated because some agents were exposed to

¹¹ The regression presented in column (1) of Table 5 does not contain the variable *share peak hours*. This is because both variables *share of peers treated* and *share peak hours* are based on the same hourly information. This causes a high between-agent within week correlation of the two variables and makes it impossible to identify the effect of *share peak hours* in this regression. This is supported by the fact that agents have only limited possibilities to choose shifts and that the assignment to the first and second training group is random (Table 2).

Table 5
Estimation of Externalities

	(1)	(2)	(3)
Share of peers treated	0.0513* (0.0275)		
Treatment dummy		0.0992*** (0.0213)	0.0675** (0.0287)
Constant	-1.1772*** (0.0702)	-1.0548*** (0.1042)	-0.9924*** (0.1018)
Individual fixed effects	Yes	Yes	Yes
Control variables	Yes	Yes	Yes
Observations	1,544	1,352	1,510
Number of agents	74	53	61
Number of agents (treatment group)	34	13	21
Number of agents (control group)	40	40	40
R ²	0.5612	0.6173	0.5894

Notes. Significance levels: *p < 0.10, **p < 0.05, ***p < 0.01. Dependent variable: log(*y_{it}*). All standard errors are clustered at the agent level. *Share of peers treated* is calculated as the average share using information on the actual working hours in a specific week. The sample used for column (1) is restricted to agents in an untreated state only (treatment dummy = 0). Control variables include *working hours*, *share peak hours*, a linear time trend and *calls per FTE*. The variable *share peak hours* is not used in column (1).

treated agents in their teams. Accordingly, the same estimation for second-group agents results in a lower estimate. Column (3) of Table 5 shows that their treatment effect is 6.8%. The difference between the two estimates, however, is not significantly different at the 5% significance level. Despite this, the results shown in Table 5 strongly suggest the existence of peer effects from treated agents on untreated peers in their team.

2.4. Returns to Training

Standard theory predicts that firms invest in human capital if the expected returns from training investments exceed costs. Personnel information with specific information on the length and costs of training, as well as information on performance, allows to estimate the returns to training (Liu and Batt, 2007). For the training analysed in this study, the costs of training per agent are defined as

$$C = \frac{(n_g w_a + w_l + w_c)h}{n_g},$$

(4)

where *n_g* is the size of the training group; *w_a*, *w_l* and *w_c* are the hourly wages for the agents, team leader and training coach, respectively; and *h* is the length of the training in hours.¹² Because our performance measure is the time an agent needs, on average, to handle a call, and there are no other tasks in which the agent is involved, benefits are defined as the percentage increase in performance times the weekly workload and the agent's wage. The benefit *B_d* can be calculated for each week after the training *d*:

¹² The wage costs of the team leader and training coach who supervise the training are included as direct training costs.

$$B_d = (\hat{\tau}_d aht_{id} n_{c,id}) w_a, \quad (5)$$

where $n_{c,id}$ is the number of calls an agent completes in after-training week d and $\hat{\tau}_d$ is the estimated treatment effect for this week as shown in panel (b) of Figure 3. The term B_d is thus the decrease in the time a trained agent needs to handle a certain workload compared to an untrained agent in week d , multiplied by his wage.

Given the decreasing treatment effect over time, we use the actual average performance and actual number of calls made to calculate the returns to training for the firm. The firm will *de facto* retain the full benefits of the performance increase because agents' wages are almost entirely related to their years of tenure. Figure 4 shows the estimated curve of cumulated returns $\sum_{d=1}^D B_d$ until week D . The solid line shows that, for first-group agents, the cumulated benefits do not exceed the costs of the training in our sample period. When considering within-team externalities, however, the training reaches its break-even point in the eighth week after the training.¹³ Figure 4 shows that in week 11 after the training, the overall return to the training is 37.7% higher when including externalities in the calculation of the benefits.¹⁴

Fig. 4. *Estimated Payoff (in Euros)*

Notes. Cumulative payoff for first-group agents (solid line) and first-group agents including externalities on teammates (dashed line). The horizontal axis shows the individual training costs based on an agent's wages and the trainer's and the team leader's wage costs (see Section 2.4).

¹³ Because half of the training consisted of actual customer calls, we calculated the costs of the training using 50% of the agents' wage costs. Considering the full wage costs of agents roughly doubles the training's costs.

¹⁴ Externalities were calculated based on the estimates of the treatment effect for the first- and the second-group agents ($\hat{\tau}_1$ and $\hat{\tau}_2$) from Columns (2) and (3) of Table 5. We use the difference between the treatment effects as a measure of externalities. The externality effect on second-group agents is thus calculated as $(\hat{\tau}_2 - \hat{\tau}_1) aht_{id} n_{c,id} w_a$.

3. Additional Evidence

3.1. First-week Effects

One finding of the preceding analysis of the dynamics of the treatment effect is that agents perform better only starting from the second week after the treatment (see Figure 3). We can rule out the argument that this is due to selectivity, since over 90% of the agents who participated in the training also worked the following week. Another, yet untestable, explanation for the low effect is that the low performance in the first week after the training is due to passing knowledge on to teammates.

A testable explanation, however, is fatigue. Most of the agents work part-time, with an average of about 20 hours per week. The training programme is one full-time week, that is, 38 hours. When using additional data on previous training programmes, we find that after rather short-training programmes (<15 hours of training in one week), the difference between pre- and post-training is about the same, irrespective of the call agent’s usual working hours (see column 1 of Table 6). This is different for long-training programmes where agents with relatively low numbers of working hours perform worse in the week after the programme than before. In contrast, agents with long-working hours do not exhibit this decrease in performance in the first week following the training. This result suggests that fatigue is an explanation for the slight decrease in performance in the first week after the training.

3.2. Personnel Turnover and Training Effects

As in other call centres, the call centre analysed in this study experiences a high agent turnover. In total, we assigned 86 agents to the treatment and control groups, of which 12 were reassigned by the management or dropped out (see Section 2). Table 1 shows that the treatment and control groups do not differ with respect to observable characteristics such as age, gender and tenure. This holds for the initial assignment, as well as for agents who eventually participated in the training.

Table 6
Pre-/Post-differences in Performance for Other Training Programmes

	(1)	(2)
Length of courses	Short	Long
Agents with shortworking hours	−0.037 (0.080)	−0.047*** (0.012)
Number of observations	42	526
Agents with longworking hours	−0.006 (0.012)	−0.040 (0.034)
Number of observations	742	64

Notes. Significance levels: *p < 0.10, **p < 0.05, ***p < 0.01. The numbers show the differences in y_{it} between the first week after the training and the week before the training. The calculations are based on a larger data set from the same call centre. Short courses are defined as courses with less than 15 hours per week. Long courses are defined as courses with more than 25 hours per week. Short and longworking hours are relative to the course length. Standard errors are in parentheses.

Ordinary least squares estimation of (2) results in an unbiased estimator of τ only if the probability of belonging to either the treatment or control group does not depend on unobservable characteristics that are correlated with both the likelihood to participate in the training programme and our performance measure, average handling time. This type of selectivity occurs when agents with these characteristics are more likely to drop out of the treatment group than from out of the control group.

To test this hypothesis, we limit the estimation sample to agents that eventually participated in the training programme and stayed in the department throughout the observation period. This decreases the sample size from 74 agents in our main estimation sample to 55 agents. Of the 34 agents of the treatment group, 26 (77%) were trained and did not leave the call centre. For the control group, this limits the number of agents from 40 to 29 (73%). The main reason for the higher attrition in the control group is that those agents followed the training several weeks later, which increases the probability of dropping out before actually participating in the training. Calculating the attrition rate for treatment group agents who were staying at least until the last week of our full sample which contains longer panel information, we find there is hardly any difference in attrition between the treatment and control groups. In addition, we do not find a systematic pattern in the exit dates that reveals that agents may have left the department due to the upcoming training.¹⁵ Table 7 shows that for this sample the estimated treatment effect decreases to 7.7%, which is lower but not significantly different from the estimated effect for the main sample (see column 3 Table 4).

3.3. Effects on Call Quality

Throughout this study, we used a transformation of the average length of calls as a measure of performance. Average handling time provides a clear measure of perform-

Table 7
Estimated Treatment Effect and Estimation Samples

Treatment dummy	0.0773*** (0.0214)
Constant	-0.9934*** (0.0993)
Individual fixed effects	Yes
Control variables	Yes
Observations	1,532
Number of agents	55
Number of agents (treatment group)	26
Number of agents (control group)	29
R ²	0.6178

Notes. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable: $\log(y_{it})$. All standard errors are clustered at the agent level. The samples is defined as all agents minus agents dropping out due to selection. Control variables include *working hours*, *share peak hours*, a linear time trend and *calls per FTE*.

¹⁵ There are potentially two dates on which selective agents may be more likely to leave due to the training. First, several months ahead of the first training, agents were informed that the whole department was going to be trained. Second, each agent received her schedule four weeks ahead. At neither of the two moments did exit rates appear to be higher than usual.

ance that allows us to quantify the returns to training precisely. One can, however, argue that individual productivity is a function of two dimensions. Calls can differ not only by quantity, but also by quality provided by the agent. In this case, 'actual' productivity $y'_{it} = f(y_{it}, y^q_{it})$ is a function of the quantity-performance measure y_{it} and a performance measure for quality y^q_{it} . If y'_{it} is a non-monotonic function in y_{it} , improvements in average handling time will not necessarily translate into higher productivity. Since call quality is much more difficult to monitor than work speed, agents may aim to improve their handling times by providing lower call quality. However, agents cannot easily shirk by delivering low quality with a low average handling time because their calls are regularly monitored by their team leaders.

To test whether agents substitute higher performance y_{it} by providing lower quality y^q_{it} , we employ an alternative indicator for which we have data for each worker and week: the share of repeat calls. This measure is defined as the share of customers an agent talked to who called the call centre again within seven days. The firm uses this measure as an indicator of call quality because low customer satisfaction with the first call is the main reason for repeat calls.¹⁶ On average, 23.6% of the customers called back within seven days after the first phone call (see column 1 of Table 1). Because low values of repeat calls (rc_{it}) indicate higher performance, we define call quality as $y^q_{it} = 1/(rc_{it} \times 10)$, with an average of 0.454.

The first column of Table 8 shows that training participation does not significantly affect call quality. The argument that the shorter average handling time after the training may be associated with lower call quality is therefore not supported by the data.

Having information on the quality of calls allows us to create a composite measure of productivity that incorporates both quantity and quality dimensions. We define y'_{it} as

Table 8
The Effect of Training Participation on Call Quality

	$\log(y^q_{it})$	$\log(y'_{it})$	Grade 1	Grade 2	Grade 3
Outcome	(1)	(2)	(3)	(4)	(5)
Treatment dummy	0.0425 (0.0391)	0.0985*** (0.0228)	0.8701** (0.3721)	1.5742*** (0.4439)	-0.4701 (0.9417)
Constant	-0.5157*** (0.1374)	-1.2132*** (0.0876)	6.0604 (3.8802)	3.2397 (3.6075)	5.2156 (5.4162)
Individual fixed effects	Yes	Yes	Yes	Yes	Yes
Control variables	Yes	Yes	Yes	Yes	Yes
Observations	1,730	1,730	112	115	107
Number of agents	74	74	14	15	15
R ²	0.3384	0.5763	0.0283	0.0474	0.0405

Notes. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All standard errors are clustered at the agent level. Call quality y^q_{it} is measured as $1/(\text{share of repeat calls} \times 10)$. The composite productivity measure y'_{it} is defined as $y'_{it} = aht_{it} \times (1 - \text{share of repeat calls})$. Grades are given on 'knowledge of agent' (Grade 1), 'agent understood question' (Grade 2) and 'solution of the problem' (Grade 3). Grades are given on a scale of 1 (very bad) to 10 (very good). Control variables include *working hours*, *share peak hours*, a linear time trend and *calls per FTE*.

¹⁶ Note that the type (and potentially also the length) of repeat calls can be different from that of initial calls. Repeat calls are, however, not linked to certain agents. This implies that every agent has a given probability of receiving repeat calls.

$y'_{it} = y_{it} (1 - rc_{it})$. Higher shares of repeat calls rc_{it} can thus be interpreted as a penalty in performance, since each additional percentage point in the share of repeat calls (lower quality) relates to a lower composite productivity y'_{it} . The second column of Table 8 shows that the estimated treatment effect on the composite measure of productivity is 9.85% and statistically significant at the 1% level.

In addition to the information on the share of repeat calls, we use data gathered in a customer survey. Individual calls made by call agents were randomly selected and evaluated by customers on a scale from 1 (very bad) to 10 (very good). Customers gave grades on three different dimensions: the 'knowledge of the agent' (*Grade 1*), whether the 'agent understood the question' (*Grade 2*) and whether the agent had a 'solution to the problem' (*Grade 3*). Because the week t in which the call was made and the corresponding agent i are known, we are able to match this information to whether the agent was treated yet or not. The unit of observation for this analysis is thus a single call, and not agent-week information. Columns (3) to (5) of Table 8 show the estimation results. They indicate that the training had a positive effect on the agent's knowledge and understanding capabilities. However, according to column (5), the training had no effect on the agent's ability to provide a solution to the problem. Despite the fact that the analysis using customer-survey data is based on a rather small sample, the results are in line with the previous results showing that quality is not substituted at the cost of average handling time.

3.4. *Heterogeneous Treatment Effects*

The preceding analysis did not consider interaction effects between worker characteristics and treatment. This subsection analyses whether an agent's tenure and number of working hours affect the size of the estimated treatment effect.

First, more experienced agents may exhibit a different effect from the treatment compared to less experienced agents. On average, more experienced workers perform better than less experienced workers. When a ceiling exists in the potential productivity of call centre agents, less experienced agents have higher potential gains from attending the training. In contrast, complementarity in human capital acquisition can lead to greater effects from participation in training for more experienced agents. Column (1) of Table 9 shows the treatment effect for agents with a tenure below the median tenure in our sample. Compared to column (2), which shows regression results for agents with a tenure above the median, the estimated treatment effect is slightly higher for experienced agents (by 1.7 percentage points). The point estimates, though, are not significantly different from each other.

Second, agents with more working hours may experience a different treatment effect compared to agents with fewer working hours. Column (3) of Table 9 shows that the interaction effect between the number of working hours and the treatment dummy is significantly negative. This implies that the training participation of agents with more working hours has a lower effect on their productivity than the training of agents who work fewer hours. This result can be explained by either the greater fatigue of agents with more working hours or the selection of individuals into contracts with lower working hours based on unobservable characteristics such as ability and motivation. Students, for instance, may have contracts with shorter average working hours.

Table 9
Direct Treatment Effect with Interaction Terms

	Below-median tenure	Above-median tenure	All
Sample	(1)	(2)	(3)
Treatment dummy	0.0811*** (0.0243)	0.0977*** (0.0325)	0.1640*** (0.0378)
Working hours			-0.0011 (0.0011)
Working hours $\times$ treatment			-0.0037*** (0.0014)
Constant	-1.0476*** (0.0988)	-1.0467*** (0.1265)	-1.0027*** (0.0872)
Individual fixed effects	Yes	Yes	Yes
Control variables	Yes	Yes	Yes
Observations	908	951	1,859
Number of agents	36	38	74
R ²	0.1277	0.1779	0.1585

Notes. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable: $\log(y_{it})$. Columns (1) and (2) show the baseline regression for agents with below-median (1) and above-median (2) tenure, respectively. All standard errors are clustered at the agent level. Control variables include *share peak hours*, a linear time trend and *calls per FTE*.

4. Conclusions

This study analyses the effect of training participation on worker performance by means of a field experiment held at a telephone company call centre. Agents had to participate in a compulsory five-day training programme. We randomly assigned agents to training groups, thereby generating exogenous variation in training participation. Regression results show that the performance of agents was about 10% better after having followed the training. This return is much higher than the returns to training found in studies that attempt to correct for selectivity. This may be explained by the fact that the effects of training on productivity are much greater than the returns measured by workers' wages (Dearden *et al.*, 2006; Konings and Vanormelingen, 2010). However, as mentioned in the Introduction, we also expect that returns to training in the call centre sector will be relatively high because the sector lacks proper vocational education for their core jobs. Furthermore, we find that the performance increase is not due to lower quality provided by the worker as indicated by a higher rate of repeat calls. Instead, training participation has a positive effect on customer ratings on agents' knowledge and understanding of the customers' question.

Moreover, we show that agents who did not participate in the training also improved their performance through externalities. These indirect effects may arise due to either knowledge spillover from trained to untrained workers or social pressure. We find that increasing the share of treated peers by 10 percentage points leads to a 0.51% increase in performance. We show that the effects measured are indeed caused by the training and not other effects, such as selective labour turnover.

Our finding of externalities in the workplace has important implications for the estimation of the effects of work-related training on wages using individual data. Even if

estimates properly control for selective training participation, the estimated treatment effect is likely to be an underestimation of the true treatment effect, because externalities of training participation increase the productivity of those who have not been trained. Since the tasks carried out by call agents follow an individual production technology without team production, these externalities are probably even higher in other sectors of industry in which team work is more important than in the call centres.

The finding of externalities has important implications for firms' training strategies. When the externalities of training participation increase the productivity of workers who have not been trained, it may be more efficient to design training and the assignment by internalising the treatment externalities. This does, however, involve important questions for further research. Are there non-linearities in the externalities and what is the optimal share of workers being treated? How and to what extent do externalities propagate further to other workers?

Maastricht University and Institute for the Study of Labour (IZA) Bonn

References

- Angrist, J.D., Imbens, G.W. and Rubin, D.B. (1996). 'Identification and causal effects using instrumental variables', *Journal of the American Statistical Association*, vol. 91, pp. 444–55.
- Bandiera, O., Barankay, I. and Rasul, I. (2007). 'Incentives for managers and inequality among workers: evidence from a firm-level experiment', *Quarterly Journal of Economics*, vol. 122(2), pp. 729–73.
- Bartel, A.P. (1995). 'Training, wage growth, and job performance: evidence from a company database', *Journal of Labor Economics*, vol. 13(3), pp. 401–25.
- Bartel, A.P. (2000). 'Measuring the employer's return on investments in training: evidence from the literature', *Industrial Relations*, vol. 39(3), pp. 502–24.
- Bassanini, A., Booth, A., Brunello, G., De Paola, M. and Leuven, E. (2007). 'Workplace training in Europe', in (G. Brunello, P. Garibaldi and E. Wasmer, eds.), *Education and Training in Europe*, pp. 143–289, Oxford: Oxford University Press.
- Batt, R., Doellgast, V. and Kwon, H. (2005). 'Service management and employment systems in U.S. and Indian call centers', in (S.M. Collins and L. Brainard, eds.), *Brookings Trade Forum 2005: Offshoring White-Collar Work—The Issues and Implications*, pp. 335–72, Washington, D.C.: Brookings Institution Press.
- Batt, R., Holman, D. and Holtgrewe, U. (2009). 'The globalization of service work: comparative institutional perspectives on call centers', *Industrial and Labor Relations Review*, vol. 62(4), pp. 453–88.
- Bertrand, M., Duflo, E. and Mullainathan, S. (2004). 'How much should we trust differences-in-differences estimates?' *Quarterly Journal of Economics*, vol. 119(1), pp. 249–75.
- Black, S.E. and Lynch, L.M. (2001). 'How to compete: the impact of workplace practices and information technology on productivity', *Review of Economics and Statistics*, vol. 83(3), pp. 434–45.
- Breuer, K., Nieken, P. and Sliwka, D. (2010). 'Social ties and subjective performance evaluations: an empirical investigation', IZA Discussion Paper No. 4913, Bonn: Institute for the Study of Labor (IZA).
- De Grip, A., Sauermann, J. and Sieben, I. (2011). 'Tenure-performance profiles and the role of peers: evidence from personnel data', ROA Research Memorandum no. 14, Maastricht: Research Centre for Education and the Labour Market (ROA).
- Dearden, L., Reed, H. and Van Reenen, J. (2006). 'The impact of training on productivity and wages: evidence from british panel data', *Oxford Bulletin of Economics and Statistics*, vol. 68(4), pp. 397–421.
- Falk, A. and Ichino, A. (2006). 'Clean evidence on peer effects', *Journal of Labor Economics*, vol. 24(1), pp. 39–58.
- Goux, D. and Maurin, E. (2000). 'Returns to firm-provided training: evidence from French worker-firm matched data', *Labour Economics*, vol. 7(1), pp. 1–19.
- Guryan, J., Kroft, K. and Notowidigdo, M. (2009). 'Peer effects in the workplace: evidence from random groupings in professional golf tournaments', *American Economic Journal: Applied Economics*, vol. 1(4), pp. 34–68.
- Holzer, H.J., Block, R.N., Cheatham, M. and Knott, J.H. (1993). 'Are training subsidies for firms effective? The Michigan experience', *Industrial and Labor Relations Review*, vol. 46(4), pp. 625–36.
- Ichniowski, C. and Shaw, K. (2009). 'Insider econometrics: empirical studies of how management matters', NBER Working Paper No. 15618, Cambridge, MA: National Bureau of Economic Research.

- Konings, J. and Vanormelingen, S. (2010). 'The impact of training on productivity and wages: firm level evidence', IZA Discussion Paper No. 4731, Bonn: Institute for the Study of Labor (IZA).
- Krueger, A. and Rouse, C. (1998). 'The effect of workplace education on earnings, turnover, and job performance', *Journal of Labor Economics*, vol. 16(1), pp. 61–94.
- Leuven, E. and Oosterbeek, H. (2008). 'An alternative approach to estimate the wage returns to private-sector training', *Journal of Applied Econometrics*, vol. 23(4), pp. 423–34.
- Liu, X. and Batt, R. (2007). 'The economic pay-offs to informal training: evidence from routine service work', *Industrial and Labor Relations Review*, vol. 61(1), pp. 75–89.
- Manski, C.F. (1993). 'Identification of endogenous social effects: the reflection problem', *Review of Economic Studies*, vol. 60, pp. 531–42.
- Mas, A. and Moretti, E. (2009). 'Peers at work', *American Economic Review*, vol. 99(1), pp. 112–45.
- Moretti, E. (2004). 'Estimating the social returns to higher education: evidence from longitudinal and repeated cross-section data', *Journal of Econometrics*, vol. 121, pp. 175–212.
- Parent, D. (2003). 'Employer-supported training in Canada and its impact on mobility and wages', *Empirical Economics*, vol. 28(3), pp. 431–59.
- Rosenbaum, P.R. and Rubin, D.B. (1983). 'The central role of the propensity score in observational studies for causal effects', *Biometrika*, vol. 70(1), pp. 41–55.
- Schwerdt, G., Messer, D., Woessmann, L. and Wolter, S. (2011). 'Effects of adult education vouchers on the labor market: evidence from a randomized field experiment', IZA Discussion Papers No. 5431, Bonn: Institute for the Study of Labor (IZA).
- Shaw, K. and Lazear, E.P. (2008). 'Tenure and output', *Labour Economics*, vol. 15(4), pp. 705–24.
- Sieben, I., De Grip, A., Van Jaarsveld, D. and Sørensen, O. (2009). 'Technology, selection, and training in call centers', *Industrial and Labor Relations Review*, vol. 62(4), pp. 553–72.
- Wilcoxon, F. (1945). 'Individual comparisons by ranking methods', *Biometrics Bulletin*, vol. 1(6), pp. 80–83.
- Zwick, T. (2006). 'The impact of training intensity on establishment productivity', *Industrial Relations*, vol. 45(1), pp. 26–46.