

Clean Evidence on Peer Effects*

Armin Falk

Andrea Ichino

IZA AND UNIVERSITY OF BONN

EUROPEAN UNIVERSITY INSTITUTE

May 19, 2005

Abstract

While confounding factors typically jeopardize the possibility of using observational data to measure peer effects, field experiments offer the potential for obtaining clean evidence. In this paper we measure the output of subjects who were asked to stuff letters into envelopes, with a remuneration completely independent of output. We study two treatments. In the “pair” treatment two subjects work at the same time in the same room. Peer effects are possible in this situation and imply that outputs within pairs should be similar. In the “single” treatment, which serves as a control, subjects work alone in a room and peer effects are ruled out by design. Our main results are as follows: First, we find clear and unambiguous evidence for the existence of peer effects in the pair treatment. The standard deviations of output are significantly smaller within pairs than between pairs. Second, average output in the pair treatment largely exceeds output in the single treatment, i.e., peer effects raise productivity. Third, low productivity workers are significantly more sensitive to the behavior of peers than are high productivity workers. Our findings yield important implications for the design of the workplace.

JEL Classification: D2, J2, K4.

Keywords: Peer effects, field experiments, incentives.

*Both authors are also affiliated with CEPR, CESifo and IZA. We would like to thank Urs Fischbacher, Simon Gächter, Lorenz Goette, Sally Gschwend, Soren Johansen, Michael Kosfeld, Ed Lazear, Alan Manning, Imran Rasul, Gianfranco Pieretti, Kathryn Shaw, Eric Verhoogen, one anonymous referee and seminar participants at Toulouse, CESifo, EUI, ESSLE, LSE, Siena and Pisa for insightful comments and suggestions. Jasmin Gulden and Barbara Nabold provided excellent research assistantship. *Corresponding author:* Armin Falk, IZA, P.O. Box 7240, D-53072 Bonn, Germany, e-mail: falk@iza.org

1 Introduction

Scholars in many disciplines have long tried to estimate empirically the extent to which individual behavior is modified by peer effects. The reason why doing this is difficult, despite the apparent wealth of evidence from daily experience, is that observational data do not allow us to easily separate the pure effect of peer behavior from the effect of confounding factors. Using data from a controlled field experiment where randomly selected subjects were paid independently of their work output, we show in this paper that the productivity of a worker is systematically influenced by the productivity of peers in the absence of confounding factors. These results provide clean evidence for the existence of peer effects on work behavior.

In order to understand the nature of our experiment, consider two individuals working on separate tasks, where one is in sight of the other. Suppose that we observe them behaving in a similar way, which we suspect could be generated by peer effects. To be precise, we say that peer effects exist if the output of individual i increases when the output of j increases and nothing else changes. Following Manski [1993], a first set of confounding factors is generated by the possibility that local attributes of the environments in which the two individuals operate determine their behavior. If observational data do not allow us to fully control for these local attributes, we could observe the behavior of i and j changing simultaneously even in the absence of true peer effects simply because some unobserved local attributes have changed. Second, it is possible that the two individuals have similar characteristics, which would make them behave similarly even if they were not working one in sight of the other. With respect to both of these possibilities, it could also happen that i and j decide to work near each other because they like the

same local attribute, which in turn affects their behavior, or because they both like to be near individuals with similar characteristics. In these cases, the supposed effect of peers would instead be the result of sorting according to local or personal attributes.

The most recent generation of studies, which try to measure peer effects with observational data, has made several important steps towards solving these problems.¹ However, even if the setting offers an almost perfect opportunity to identify peer effects in many of these studies, the impossibility of controlling for all local or personal confounding factors and for endogenous sorting makes the identification strategy not fully convincing. The most significant recent steps forward in this literature are offered by Sacerdote [2001] and Katz, Kling and Liebman [2001] who use data based on randomized assignments of individuals to peer groups. However, both of these papers are confronted with the consequences of local confounding factors. More specifically, Sacerdote [2001] finds evidence of peer effects among Dartmouth students randomly assigned to the same dorm but cannot completely exclude the possibility that these effects might be due to local time varying shocks. This is less of a problem in Katz, Kling and Liebman [2001], who analyze the consequences of randomly changing the residential neighborhood of families residing in high-poverty public housing projects and, therefore, are not primarily interested in isolating pure peer effects from local effects. A further important difference with respect to our setting is that neither of these papers focuses on a work environment.

In contrast, we focus explicitly on a real work environment in our study

¹See, among others, Wilson [1987], Case and Katz [1991], Crane [1991], Glaeser et al. [1996], Topa [1997], Encinosa et al. [1998], Aaronson [1998], Van Den Berg [1998], Bertrand et al. [2000], Ichino and Maggi [2000], Katz, Kling and Liebman [2001] and Sacerdote [2001]. See also the literature based on the classic Hawthorne experiments (e.g., Whitehead [1938] and more recently Jones [1990]).

and we aim to assess the existence of peer effects in a fully controlled setting where no possible confounding factor can hinder this assessment. As in any other controlled experiment, the possibility of obtaining clean evidence complements the evidence generated by observational studies in an informative way.²

Our subjects were recruited randomly and asked to perform a typical short term job, which was paid independently of individual or team output. The work task was to stuff letters into envelopes. We study two treatments. In the ‘pair’ treatment, which is our main treatment, two subjects work simultaneously in the same room. This setting allows for the possibility that the behavior of a subject is affected by the behavior of the other member of the pair. Given two subjects i and j in a pair, we speak of positive peer effects if the output of i systematically raises the output of j , and vice versa, leading to similar output levels within the pair. A formal characterization of this definition will be given in Section 3. In the second treatment (the ‘single’ treatment), which serves as control, peer effects are ruled out by design because subjects work alone in a room. Output in this treatment reveals the level of productivity in the absence of any peer influence. The comparison of the outputs arising in the pair treatment with those from the single treatment permits the assessment of the effects of peers on individual productivity.

Our main results are the following: First, we find strong and unambiguous evidence for the existence of positive peer effects in the pair treatment. This can be inferred from the fact that output within pairs is very similar, while

²For related literature on laboratory experiments aimed at measuring peer effects see Falk and Fischbacher [2002] and Falk, Fischbacher, and Gächter [2002]. Nagin et al. [2002] provide instead an example of controlled experimentation in a real labor setting, although their focus is on a different issue.

differing substantially between pairs. This difference is particularly striking when compared to what happens in random allocations of subjects from the pair and the single treatment in simulated pairs. By comparing the standard deviation of output within and between true and simulated pairs, we show that peer effects are large and highly significant. Second, even though economic incentives are identical, average output in the pair treatment is higher than that in the single treatment. Thus, peer effects significantly increase output. Third, we show that peer influence affects subjects differently. In particular we find that it mainly improves the output of less productive subjects. Finally we derive an implicit estimate for the strength of peer effects. Interestingly, the estimated coefficient is very similar to a comparable estimate, which was derived by Maggi and Ichino (2000) with observational data.

The paper is organized as follows. In the next Section we present the design of our experiment. Section 3 discusses the behavioral hypotheses. Section 4 contains our results. Section 5 concludes.

2 Design of the field experiment

The goal of this paper is to study potential peer effects on work behavior. We therefore conducted a field experiment where subjects who performed a simple task in a highly controlled environment were exogenously sorted into two different treatments. Before discussing our treatments in detail, we describe the recruitment process, the work task and the procedures.

2.1 Recruitment

All our subjects were high-school students who were recruited from different schools in the area of Winterthur, a city in the canton of Zurich (Switzerland).

Students were asked in announcements posted on blackboards whether they wanted to do a simple short term job requiring no previous knowledge. In the announcement it was stated that the job was a one-time four hour job, which was paid 90 Swiss Francs (1 Swiss Franc \approx .70 US or \approx .70 EURO). The payment was obviously attractive as we were able to recruit the number of subjects we had planned to recruit within 24 hours.

Students applied by email. After receiving their applications, we informed them of the precise date and location where they were expected for the job. The experiment took place during the 2002 spring vacations, which cover two weeks. It was performed in a high-school building in Winterthur.

2.2 Procedure and task

Upon arrival, subjects were welcomed and informed about the task and the procedural details. In particular, they were told that they had to work for four hours without a break and that at the end of this time, they would receive their payment.

We chose a work task, which is simple, requires no previous knowledge and is easy to measure. In particular, students had to prepare the mailing of a questionnaire study for the University of Zurich. This job basically involved stuffing letters into envelopes. First, subjects had to fold two sheets of paper (one sheet contained the description of the questionnaire, the other was to be filled out by the recipients of the study). After placing the two sheets into the envelope, subjects had to seal the envelope and to put an A-priority sticker on it. When a set of 25 envelopes had been completed the set had to be bundled with a rubber band and put in a box. The work environment was exactly the same for each subject, including, e.g., the same type of desk and chair and the same large number of envelopes and sheets (Figure 1 displays

a picture of a subject’s desk). Payment was independent of output and paid in cash. In both treatments the procedure was exactly the same.

Note that we determined a subject’s output only after they had been paid. This implies that we have only one output observation per subject. An interesting question, which we cannot address with this procedure concerns the importance of peer effects over time. For example, peer effects may be stronger at the beginning than at the end of the experiment. To address differences over time it would have been necessary to document output not only at the end of the experiment but also during the experiment. This would have required, e.g., to install a camera and to videotape subjects’ behavior or to have an experimenter counting output during the experiment. It is likely that this type of intervention would have been experienced as quite “unnatural” by or subjects. Since one of our primary goals was to implement a rather natural work environment we decided against this option.

2.3 Treatments

We study two treatments, the “pair” and the “single” treatment. In the pair treatment two subjects did the task described above at the same time in the same room. The two desks were situated in such a way that a subject could easily realize the output of the other subject (the position of the second desk can be seen in the background of Figure 1). Subjects were free to communicate but instructed that they had to perform the task described above independently. Hence, they were not allowed to engage in teamwork or division of labor. We invited only students from different high-schools to participate in this treatment in order to minimize the possibility that two subjects in the pair treatment knew each other.

In the pair treatment peer effects were possible. In contrast, peer effects

were ruled out by design in the single treatment. In this control treatment everything was exactly the same as in the pair treatment except that in this case each subject worked alone in a room. Since subjects did not have any contact to another subject and were not informed about other subjects' output in this treatment, the single treatment rules out any potential peer effect stemming from a co-worker. Therefore, a comparison of output arising in the single treatment with that of the pair treatment, indicates the potential effects of peers on productivity.

A total of 24 subjects participated in our study, eight in the single treatment and 16 (eight pairs) in the pair treatment. Treatment assignment was random. No subject participated in more than one treatment.

From a methodological point of view some aspects about the design are worth pointing out: Unlike most lab experiments that study work behavior, our subjects performed a 'real' task. In a typical lab experiment the choice of work effort is represented by an increasing monetary function, i.e., instead of choosing real effort subjects choose a costly number. This procedure has been used in tournament experiments, e.g., Bull, Schotter and Weigelt [1987], or in efficiency wage experiments, e.g., Fehr, Kirchsteiger and Riedl [1993]. Some authors have recently conducted so-called 'real effort' experiments to study incentive mechanisms and efficiency wages. In Fahr and Irlenbusch [2000] subjects had to crack walnuts, in van Dijk et al. [2001] subjects performed cognitively demanding tasks on the computer (two-variable optimization problems) and in Gneezy [2003] and Gneezy, Niederle, and Rustichini [2003] subjects had to solve mazes at the computer. While these real effort tasks resemble regular work more than just choosing a number, subjects probably do not perceive these tasks as economically valuable. This means that an important dimension of regular work is missing. In contrast, subjects

in our study performed a regular and economically valuable job.

3 Behavioral hypotheses

To illustrate what we expect to happen in our experiment in the presence of peer effects we present in this section a slightly modified version of the model proposed by Ichino and Maggi (2000). Consider a generic subject i of our experiment who chooses a level of output denoted by $X_i \in [0, X^{\max}]$. We assume that the gain from producing X_i is given by $G(X_i, Y^e, \theta_i)$, with $G_1 > 0$ and $G_{11} < 0$, where θ_i is a preference parameter (the subject's "type") and Y^e is a vector of characteristics describing the local environment e in which i operates. A higher value of θ_i indicates a worker with a higher marginal gain from producing X_i . This amounts to assuming $G_{13} > 0$.

We introduce the possibility of peer effects by modelling the cost of producing X_i as given by $L(X_i, \bar{X}_i^e)$, where \bar{X}_i^e is the average output of peers in the environment in which i operates. In the absence of peer effects $L_{12} = 0$, while if instead $L_{12} \leq 0$ there are peer effects in the sense that the cost of producing X_i is lower when average production is higher. Note that our goal is not to identify the determinants of peer effects, but just to understand what we should see in the data generated by our experiment if peer effects are at work. By allowing L to depend on \bar{X}_i^e we suggest the possibility that, for whatever reason, it may be costly for a subject not to keep effort in line with what the others do³. Note that the way we model peer effects differs from the way peer effects are modeled, e.g., in Kandel and Lazear [1992] or Huck, Kübler and Weibull [2002]. The latter models focus on the question whether peer pressure improves performance in situations where payments

³For a discussions of possible determinants of peer effects see, among others, Kandel and Lazear [1992], Akerlof [1997], Spagnolo [1999] and Huck, Kübler and Weibull [2002].

are based on team incentives, i.e., where a subject's effort has a positive externality on other team members' profits.

Considering a team of subjects like i we can characterize the Nash equilibria of this game. The first step is to derive an individual subject's optimal choice given the other subjects' choices. Each subject chooses X_i to maximize the individual utility of production,

$$U^i = G(X_i, Y^e, \theta_i) - L(X_i, \bar{X}_i^e).$$

Therefore, the optimal output level will be a function of \bar{X}_i^e , Y^e and θ_i :

$$X_i = g(\bar{X}_i^e, Y^e, \theta_i). \quad (1)$$

Given our assumptions, we have $\partial X_i / \partial \bar{X}_i^e \geq 0$ with strict inequality if positive peer effects exist, and $\partial X_i / \partial \theta_i > 0$. Note that equation (1) is a structural condition, because \bar{X}_i^e is endogenous.

Using (1) and denoting with $f^e(\theta)$ the distribution of types in the environment e , we can write

$$\bar{X}_i^e = \int g(\bar{X}_i^e, Y^e, \theta) df^e(\theta). \quad (2)$$

The solutions of this equation in \bar{X}_i^e represent the equilibrium average output levels. Note that, if g is linear, there is a unique equilibrium, but if g is nonlinear, multiple equilibria are possible.

Coming closer to the setup of our experiment, note that in the pair treatment $\bar{X}_i^e = X_j$ where i and j denote the two subjects in a pair. Moreover, since the environment is controlled by the experimenter and kept constant for all subjects and treatments, $Y_i^e = Y$. Thus, linearizing for simplicity equation (1) we obtain:

$$X_i = Y + \beta X_j + \theta_i \quad (3)$$

and symmetrically for j . In this equation β measures how the output of i depends on the output of j when they work in a pair. Within this context we say that peer effects exist and are positive if the output of i increases with the output of j , which formally means:

Definition 1 *If $\beta > 0$, positive peer effects exist in a pair. $\beta = 0$ implies absence of peer effects, while these effects are negative if $\beta < 0$.*

In the equilibrium of the pair treatment, the output of subject i is given by

$$X_i^p = \frac{Y}{(1 - \beta)} + \frac{\theta_i + \beta\theta_j}{(1 - \beta^2)} \quad (4)$$

while the same subject in the single treatment would produce

$$X_i^s = Y + \theta_i \quad (5)$$

since in this treatment no other subject exercises any pressure on i . Symmetrically, we can derive analogous expressions for j . It is important to note that random assignment ensures that types θ are randomly distributed in the two treatments.

Points P and S in Figure 2 describe the respective equilibria of the pair and single treatments. The figure also plots the reaction curves described by equation (3) for the pair treatment, which cross at P , and by equation (5) for the single treatment, which cross at S .

It is immediately obvious that the difference between the output levels of the two subjects within each pair is equal to

$$|X_i^P - X_j^P| = \frac{|\theta_i - \theta_j|}{1 + \beta}. \quad (6)$$

As a result, positive peer effects can be detected in the pair treatment according to the following proposition, which will be tested in Section 4.

Proposition 1 *If positive peer effects exist, i.e., $\beta > 0$, the absolute value of the difference between output levels within pairs should be smaller than if there were no peer effects.*

An illustration of Proposition 1 is given with the help of Figure 2, where P shows an equilibrium with $\beta > 0$ and S shows an equilibrium with $\beta = 0$. Since P is closer to the 45-degree line than S , output levels are more similar in P in comparison to S . Moreover, it is obvious that a higher β implies output levels which are increasingly similar in the P equilibrium.

The setting of our experiment offers the possibility for testing further implications of peer effects. In the absence of these effects, the distributions of output should be the same in the pair and in the single treatment. This is so because the economic incentives are identical in both conditions. Of course, there might be individual differences because some subjects are, e.g., more talented than others or feel more obliged to perform well than others do. Since subjects are randomly allocated to the treatment conditions, however, individual differences should cancel out.

On the contrary, if peer effects do exist, it is easy to show that the average output in the two treatments should differ. Using equation (3), the average output of i and j when they work in a pair is

$$\frac{X_i^P + X_j^P}{2} = \frac{Y}{(1 - \beta)} + \frac{\frac{\theta_i + \theta_j}{2}}{1 - \beta} \quad (7)$$

while the average output of the same two subjects working alone in the single treatment would be

$$\frac{X_i^S + X_j^S}{2} = Y + \frac{\theta_i + \theta_j}{2} \quad (8)$$

A comparison of equations (7) and (8) shows that, in the presence of positive peer effects such that $0 < \beta < 1$, average output is higher in the pair treatment than in the single treatment. This can also be inferred from Figure 2

where output in the P equilibrium is clearly higher compared to output in the S equilibrium. If instead $\beta > 1$ the output level of the two subjects would still be higher in the pair treatment but it would be equal to infinity. On the contrary, in the case of negative effects ($\beta < 0$) the output of a subject reduces the output of the other, in which case the output of the pair treatment would be lower than the output of the single treatment. Our model therefore suggests a second proposition, which will be tested in Section 4.

Proposition 2 *In the presence of positive peer effects, the average output of the pair treatment exceeds that of the single treatment.*

Note that Proposition 2 states a behavioral consequence of peer effects, which is similar to the so-called ‘social facilitation’ paradigm in social psychology. According to this paradigm even the mere presence of another person improves one’s performance. Numerous studies have supported evidence for this type of behavior.⁴

Our final proposition derives immediately from the first two given the setting, but it is worth stating it explicitly. Let’s focus on a pair of subjects i and j and consider the differences $\Delta X_i = X_i^P - X_i^S$ and $\Delta X_j = X_j^P - X_j^S$, which measure the change between the two potential output levels in the pair and in the single treatment respectively for i and j . Using 4 and 5 note that $\Delta X_j - \Delta X_i = (X_i^S - X_j^S) - \frac{(X_i^S - X_j^S)}{1+\beta}$. It is easy to see that, if $\beta > 0$, this quantity is strictly positive only when $X_i^S > X_j^S$, i.e., only when i is the member of the pair who would produce more if both subjects were working

⁴See for example Zajonc [1965], Cottrell et al. [1968] and Hunt and Hillery [1973]. In Allport [1920], performance of subjects doing simple tasks (like chain word association) was much better in groups than if subjects did the tasks alone. In a more recent study, Towler [1986] takes the time cars need to reach a 100-yards mark from a standing start at traffic lights. He reports that the if there are two cars at the traffic light the time to travel the 100 yards is significantly shorter than if there is just one car.

alone.⁵ Thus, the following proposition holds.

Proposition 3 *Given a pair of subjects working together, the one who would be less productive working alone is the one whose output increases by more when joining the pair as opposed to working alone.*

Hence, our simple model suggests three propositions, which describe the implications of peer effects in our treatments. We test these propositions in the next section where we also show how our data can be used, in the light of the model described above, to derive an implicit estimate of β .

4 Results

In this section we present our results and test our behavioral predictions. Our main interest concerns the existence of positive peer effects, which are revealed by the observation that output levels within pairs are similar in the pair treatment.

In order to test Proposition 1, consider the standard deviation of output within and between pairs.⁶ In the absence of peer effects (i.e., $\beta = 0$), working in a particular pair has no effect on individual behavior. In this case, therefore, the standard deviations of output within pairs should be identical to those generated by any simulated configuration of pairs constructed from the subjects who worked alone in the single treatment. Given our data there are 105 possible configurations of 4 pairs with 8 individuals⁷. Only *one* of

⁵This is also graphically evident in Figure 2.

⁶We use standard deviations instead of differences to facilitate the computation and the comparison of within and between statistics. This, however, does not change the substance of our results because, in our specific case, the standard deviation within a pair is equal to the absolute value of the difference between the output levels of the pair divided by the square root of 2.

⁷ This number of configurations is in general equal to $\prod_{i=0}^{(N-2)/2} (N - 2i - 1)$, where N is the (even) number of individuals, i.e. 8 in our case.

these 105 configurations originates a hypothetical within standard deviation lower than that obtained with the true pairs of the pair treatment. The likelihood that this finding is just pure coincidence in the absence of peer effects is below 1 percent.

A second test of the same proposition is possible by comparing the true pairs and simulated pairs using only the subjects involved in the pair treatment. In the absence of peer effects (i.e., if $\beta = 0$), the standard deviation of output within the true pairs should be identical to the within standard deviation generated by any simulated configuration of pairs constructed from the same group of people. Moreover, there should be no reason to expect that the between and within standard deviations obtained with the true pairs should differ in any specific direction. These comparisons are shown in Figures 3, 4 and 5.

The first of these figures plots the kernel density of the simulated within pairs standard deviations computed for 20,271 randomly chosen different configurations of pairs of the 16 individuals involved in the pair treatment. To be more precise, we generated all 2,027,025 possible configurations of 8 pairs with these 16 individuals,⁸ and for one out of every 100 of these configurations we computed the within pairs standard deviation.⁹

The variation of these simulated within standard deviations ranges from 9.6 to 34.8 letters. The vertical line in Figure 3 identifies the standard deviation within true pairs, i.e., that computed for the pairs who actually worked together in our experiment. This standard deviation is equal to 14.6 letters and only 1.17% of the simulated configurations originated a lower value. This

⁸See footnote 7

⁹We would have liked computing the within pairs standard deviations for all the 2,027,025 configurations but this calculation would have required a substantial amount of computer time without any major gain from the viewpoint of the reliability of our results.

suggests that on average the output levels of two individuals working in the same room on separate tasks, are significantly more similar than the output levels of two individuals working separately. In other words, in the absence of any peer effect, the probability of observing a within-pairs deviation as low as 14.6 is on average less than 1.17%.¹⁰ Hence, we can reject the hypothesis of the absence of peer effects with a high level of confidence.

In line with Figure 3, we find in Figure 4, that the observed standard deviation between the true pairs in the experiment (which is equal to 33.7 letters) is higher than 98.85% of the between standard deviations generated by the simulated configurations of pairs. The chance that such a high between standard deviation could be generated in the absence of peer effects is extremely low (in particular smaller than 1.15%). Moreover, Figure 5 plots the kernel density of the between minus within difference for each hypothetical configuration of pairs. It is evident that this difference is not systematically positive or negative since it is approximately symmetric around zero. Note that this is exactly what one would expect in the absence of peer effects, while in the presence of these effects, the between standard deviation should be larger than the within. This is indeed what we find for the true pairs of our experiment: the between minus within difference is equal to 19.0 letters, as indicated by the vertical line in the figure. For only less than 1.17% of the simulated configurations the analogous difference reaches a higher value. Hence, while in the absence of peer effects there would be no reason to expect the within standard deviation to be smaller than the between standard deviation or vice versa, Figure 5 suggests that when individuals are paired

¹⁰Note that the standard deviations computed for the simulated configurations are identically but not independently distributed random variables. Because of stochastic variation, the true probability of observing a within standard deviation smaller than 14.6 in a simulated configuration might be larger or smaller than 1.17%. However, it will be equal to this value on average, since the random variables are identically distributed.

in the same room the between pairs deviation is significantly larger than the within pairs deviation. This implies that, *ceteris paribus*, working in pairs induces more similar output levels than working separately.

We now turn to our second proposition. Remember that according to standard economic theory average output levels in the pair and the single treatment should be similar because incentives are identical in both treatments. In the presence of peer effects, however, output should be higher in the pair compared to the single treatment. This is in fact what we find. The average output in the single treatment is 190 envelopes while average output in the pair treatment is 221 envelopes. The difference is not only sizeable in percentage terms (16.3 percent) but also statistically significant despite the small sample size. To show this we regress outputs in both treatments on a treatment dummy for the pair treatment. The respective *p*-value of this dummy is 0.068. This is confirmed by the non parametric Wilcoxon ranksum test ($p = 0.049$, one sided). Thus peer effects lead to higher average output as hypothesized in Proposition 2.

Finally we provide evidence in favour of our third proposition which suggests that in the presence of positive peer effects the subject of a pair who would be less productive in the single treatment is the one who would increase the productivity by more when joining the pair as opposed to working alone. With the data at our disposal we cannot perform a direct test of Proposition 3 because for each subject we observe the output level only in the treatment she is assigned to but not the output level in the counterfactual treatment. However, the comparison of the quantiles of the distributions for the single and the pair treatment provide evidence which is consistent with this proposition. These quantiles are shown in columns 1 and 2 of Table 1. The output of the 10th quantile in the single treatment is 133 while the output of the

same quantile in the pair treatment is 175, which implies a difference of 42 letters. For the 90th quantile instead the output levels of the two treatments are respectively 256 and 265, with a difference of only 9 letters. Thus, at low productivity levels peer effects determine large differences in output between the pair and the single treatment, while at high productivity levels the differences are small. In fact the Spearman Rank correlation between the differences and the corresponding productivity levels is negative and significant (Spearman's rho = -0.900, p= 0.018, one sided). This is what one would expect to see in the data given Proposition 3.

We conclude this section by showing how, in the light of our simple model of Section 3, the data generated by our experiment can be used to estimate β . Remember that this parameter measures how the output of i influences the output of j in a pair and vice versa. Equations (4) and (5) say that a subject i 's outputs in the pair and the single treatments are given by $X_i^p = \frac{Y}{(1-\beta)} + \frac{\theta_i + \beta\theta_j}{(1-\beta^2)}$ and $X_i^s = Y + \theta_i$, respectively. Substituting the sample averages \bar{X}^p for X_i^p and \bar{X}^s for θ_i and θ_j , after simplifying we can compute the average β solving $\bar{X}^p = \frac{\bar{X}^s + \beta\bar{X}^s}{(1-\beta^2)}$ or $221 = \frac{190 + \beta 190}{1-\beta^2}$. This gives an implicit estimate of $\beta = 0.14$, which implies that when the output of j increases by one unit, the output of i increases by 0.14 units, on average. Of course, we do not claim that 0.14 is a universal number. Yet, it is interesting and reassuring to see that Maggi and Ichino (2000), who derive a comparable estimate of β with observational data, get very similar numbers. Their study exploits personnel longitudinal data on workers in different branches of a large Italian firm to analyse peer effects in absenteeism behaviour. On the basis of a first difference version of equation (3) they look at movers between branches and estimate how much their absenteeism changes as a function of the difference in the absenteeism of stayers of the departure and arrival branch. This

procedure allows them to control for individual and branch fixed effects as well as for several individual and local time varying covariates. Depending on the specification, their estimates of β are $\beta = 0.14$, $\beta = 0.18$ and $\beta = 0.15$.

5 Concluding remarks

In this paper we have presented clean evidence in favor of the existence of peer effects. We show in a controlled field experiment that the behavior of subjects working in pairs is significantly different than the behavior of subjects working alone. The standard deviations within pairs are significantly smaller than between pairs. As a second result, peer effects work in the direction of raising the overall average productivity significantly. We also show that the less productive workers are more affected by peer effects than high productivity workers. In other words, “bad apples”, far from damaging “good apples”, seem instead to gain in quality when paired with these latter. The presence of peer effects is robust and quantitatively important even though subjects interacted only once and did not know each other. This suggests that the effects measured in our study are a lower boundary for the effects that prevail in actual labor relations.

In our experiment we implemented a fixed pay regime, i.e., payments were not conditioned on output. This was done in order to keep things as simple as possible. In many firms, however, more sophisticated pay regimes are common. Often pay is based on individual performance as in piece rate schemes. Alternatively pay often depends on group performance as in profit sharing, team based incentives or stock option plans. It would be interesting to study the interaction of these incentive schemes and peer effects in our real effort set-up.

If pay is based on group incentives it may well be that peer effects play

a decisive role in enhancing performance. This holds in particular if group members can directly exhibit peer pressure in the form of social exclusion, sanctions or by pushing poor performers out of the team. This type of peer pressure is not modeled in our experiment where peer effects are operative through comparisons only. There is, however, evidence from public goods experiments that peer pressure enhances team performance. In these experiments subjects are given the opportunity to sanction other players after observing their contributions to the team. In the absence of such sanction possibilities cooperation and team performance is typically rapidly declining over time. In contrast if sanctions are possible, free-riders are severely sanctioned and contribution levels are high and stable (Fehr and Gächter 2000). The driving forces behind this type of sanctions are studied in much detail in Falk et al. [2001].

In case of individual pay-for-performance peer effects may play a less important role simply because the situation is much more structured and individual output produces no externality on the earnings of other employees. However, even if pay is based on individual output, employees may care about their own output and pay relative to the output and pay of their co-workers. If this holds, peer effects of the kind studied in this paper may affect effort even in pay-for-performance schemes.

The existence of peer effects raises an interesting question concerning the efficient design of the workplace. Should employees work in groups or alone? If they work in groups, how should low and high productive workers be optimally grouped? In light of our results, the output maximizing strategy is to let people work in groups rather than alone. In addition since high productive workers seem to elevate low productive workers it may be optimal to group low and high productivity workers instead of grouping workers of

similar productivity. These conclusions, however, are based on peer effects observed in a fixed pay regime and for simple job tasks. Whether they remain valid in a more general setting and specifically if employees are paid according to group incentives or individual pay-for-performance remains an empirical question for future research.

References

- Aaronson, Daniel, “Using Sibling Data to Estimate the Impact of Neighborhoods on Children’s Educational Outcomes”, *Journal of Human Resources*, XXIII (1998), 915–946.
- Akerlof, George, “Social Distance and Social Decisions”, *Econometrica* 65(4), (1997), 1005-1027.
- Allport, F.H., “The Influence of the Group upon Association and Thought”, *Journal of Experimental Psychology* 3 (1920), 159-182.
- Bertrand, Marianne, Erzo F.P. Luttmer and Sendhil Mullainathan, “Network Effects and Welfare Cultures”, *Quarterly Journal of Economics* 115 (2000), 1019-1055.
- Bull, C., Schotter, A. and K. Weigelt, “Tournaments and Piece Rates: An Experimental Study”, *Journal of Political Economy* 95 (1987), 1-33.
- Case, Anne, and Lawrence Katz, “The Company You Keep: The Effect of Family and Neighborhood on Disadvantaged Youth”, NBER Working Paper 3705 (1991).
- Cotrell, N. B., Wack, D. L., Sekerak, F. J. and R. H. Rittle, “Social Facilitation of Dominant Responses by the Presence of an Audience and the Mere Presence of Others”, *Journal of Personality and Social Psychology* 10 (1968), 245-250.
- Crane, Jonathan, “The Epidemic Theory of Ghettos and Neighborhood Effects on Dropping Out and Teenage Childbearing”, *American Journal of Sociology* XCVI (1991), 1226–1259.
- Encinosa, William E., Martin Gaynor and James B. Rebitzer, “The Sociology of Groups and the Economics of Incentives: Theory and Evidence on Compensation Systems”, Case Western Reserve University (1998).
- Fahr, René, and Bernd Irlenbusch, “Fairness as a Constraint on Trust in Reciprocity: Earned Property Rights in a Reciprocal Exchange Experiment”, *Economics Letters* 66 (2000), 275-282.
- Falk, Armin, Ernst Fehr and Urs Fischbacher, “Driving Forces of Informal Sanctions”, IEW Working Paper No. 59 University of Zurich (2001).

- Falk Armin and Urs Fischbacher, "Crime in the Lab-detecting Social Interaction", *European Economic Review* 46 (2000), 859-869.
- Falk, Armin, Fischbacher, Urs and Simon Gächter, "Isolating Social Interaction Effects - An Experimental Investigation", IEW Working Paper 150, University of Zurich (2003).
- Fehr, Ernst and Simon Gächter, "Cooperation and Punishment in Public Goods Experiments, *American Economic Review* 90, (2000), 980-994.
- Fehr, E., Kirchsteiger, G. and A. Riedl, "Does Fairness Prevent Market Clearing? An Experimental Investigation", *Quarterly Journal of Economics* 108 (2) (1993), 437-459.
- Glaeser, E. L., B. Sacerdote and Jose A. Scheinkman, "Crime and Social Interactions", *Quarterly Journal of Economics* CXI (1996), 507-548.
- Gneezy, U., "Do High Wages Lead to High Profits? An Experimental Study of Reciprocity Using Real Effort", Working paper, University of Chicago, GSB 2003.
- Gneezy, U., M. Niederle, and A. Rustichini, "Performance in competitive environments: Gender differences", *Quarterly Journal of Economics*, CXVIII, (2003), 1049-1074.
- Huck, Steffen, Dorothea Kubler and Jörgen Weibull, "Social Norms and Incentives in Firms", mimeo, University College London (2002).
- Hunt, Peter J., and Joseph M. Hillery, "Social Facilitation in a Coaction Setting: An Examination of the Effects Over Learning Trials", *Journal of Experimental Social Psychology* 9 (1973), 563-571.
- Ichino A. and G. Maggi, "Work Environment and Individual Background: Explaining Regional Shirking Differentials in a Large Italian Firm" *Quarterly Journal of Economics* 115 (2000), 1057-1090.
- Jones, R. G. Stephen, "Worker Interdependence and Output: The Hawthorne Studies Reevaluated", *American Sociological Review* LV (1990), 176-190.
- Kandel, E. and E. P. Lazear, "Peer Pressure and Partnerships", *Journal of Political Economy* 100 (4) (1992), 801-817.

- Katz L., J. R. Kling and J. B. Liebman, “Moving to Opportunities in Boston: Early results of a Randomized Mobility Experiment”, *Quarterly Journal of Economics* 116(2) (2001) , 607-654.
- Manski, Charles F., “Identification of Endogenous Social Effects: The Reflection Problem”, *Review of Economic Studies* LX (1993), 531–542.
- Nagin, D. S., Rebitzer J. B., Sanders S. and J. L. Taylor, “Monitoring, Motivation and Management: The Determinants of Opportunistic Behavior in a Field Experiment”, *American Economic Review*, 92 (4) (2002), 850-873.
- Sacerdote, B., “Peer Effects with Random Assignment: Results for Dartmouth Roommates”, *Quarterly Journal of Economics* (2001) 116(2), 681-704.
- Spagnolo, Giancarlo, “Social Relations and Cooperation in Organizations”, *Journal of Economic Behavior and Organization* (1999) 38, 1-25.
- Topa, Giorgio., “Social Interactions, Local Spill–overs and Unemployment”, Working Paper, New York University (1997).
- Towler, Gary, “From Zero to One Hundred: Coaction in a Natural Setting”, *Perceptual and Motor Skills* 62 (1986), 377-378.
- Van den Berg, Gerard, Bas Van der Klaauw and Jan van Ours, “The Effect of Neighborhood Characteristics on the Labor Supply of Welfare Recipients”, CentER, Tilburg, mimeo (1998).
- Van Dijk, Frans, Joep Sonnemans and Frans van Winden, “Incentive systems in a real effort experiment”, *European Economic Review* 45 (2001), 187-214.
- Whitehead, Thomas N., *The Industrial Worker: A Statistical Study of Human Relations in a Group of Manual Workers* (Cambridge, MA: Harvard University Press, 1938).
- Wilson, William J., *The Truly Disadvantaged* (Chicago, IL: The University of Chicago Press, 1987).
- Zajonc, Robert B., “Social Facilitation”, *Science* 149 (1965), 269-274

Table 1: Quantiles of the output distribution in each treatment

Quantile	single treatment	pair treatment	difference
10th	133	175	42
25th	173	207	34
50th	194	212	18
75th	213	236	23
90th	256	265	9

Note: columns 1 and 2 of the table report the quantiles of the output distribution for the single and the pair treatments, estimated using a quantile regression of output on a dummy for the pair treatment plus a constant. Column 3 reports the absolute value of the difference between the quantiles estimated for the two treatments.



Fig. 1: The desk




Fig. 2: Reaction curves and equilibria in the pair and in the single treatment




Fig. 3: St. dev. within true and hypothetical pairs in pair sample




Fig. 4: St. dev. between true and hypothetical pairs in pair sample

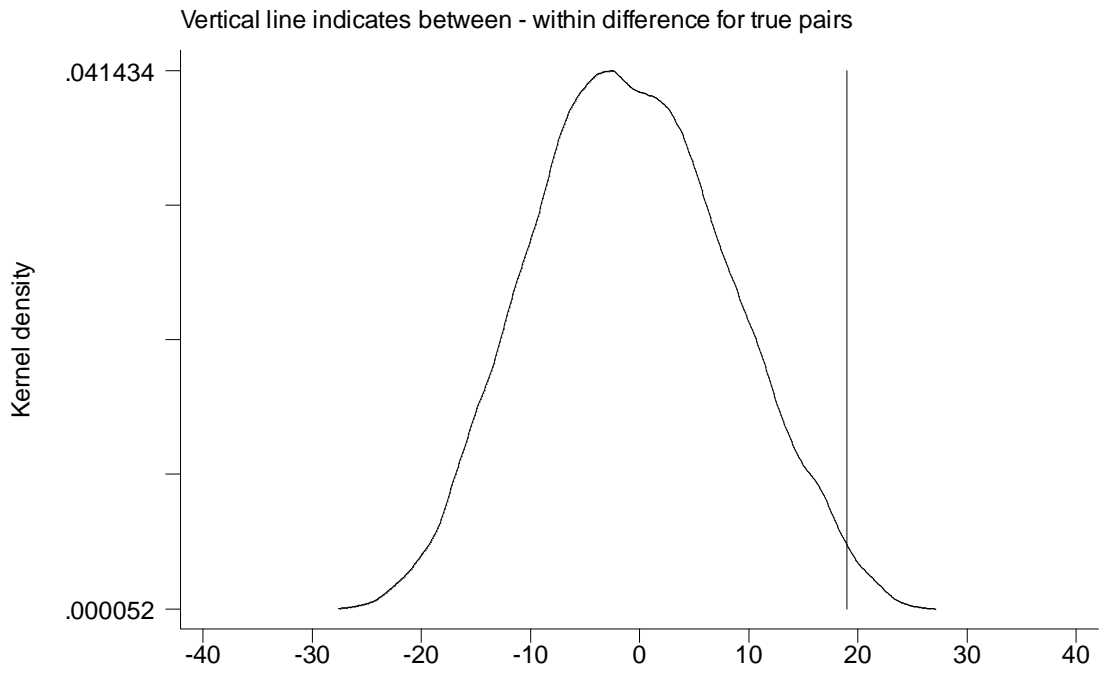


Fig. 5: Between - within st. dev. for true and hypothetical pairs