

IZA DP No. 3960

Sick of Your Colleagues' Absence?

Patrik Hesselius
Per Johansson
Peter Nilsson

January 2009

Sick of Your Colleagues' Absence?

Patrik Hesselius

IFAU and Uppsala University

Per Johansson

IFAU, Uppsala University and IZA

Peter Nilsson

IFAU and Uppsala University

Discussion Paper No. 3960
January 2009

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Sick of Your Colleagues' Absence?*

We utilize a large-scale randomized social experiment to identify how coworkers affect each other's effort as measured by work absence. The experiment altered the work absence incentives for half of *all* employees living in Göteborg, Sweden. Using administrative data we are able to recover the treatment status of all workers in more than 3,000 workplaces. We first document that employees in workplaces with a high proportion treated coworkers increase their own absence level significantly. We then examine the heterogeneity of the treatment effect in order to explore what mechanisms are underlying the peer effect. While a strong effect of having a high proportion of treated coworkers is found for the *nontreated* workers, no significant effects are found for the *treated* workers. These results suggest that pure altruistic social preferences can be ruled out as the main motivator for the behaviour of a nonnegligible proportion of the employees in our sample.

JEL Classification: J24

Keywords: social interactions, employer employee data, work absence, fairness, reciprocal preferences

Corresponding author:

Per Johansson
IFAU
Kykogårdsgatan 6.
Box 513
751 20 Uppsala
Sweden
E-mail: per.johansson@ifau.uu.se

* We have benefited from comments by Gerard van den Berg, Marco Caliendo, Bruno Crepon, Gregory Jolivet, an anonymous referee and by seminar participants at Uppsala University, University of St. Gallen and the Microdata Methods and Practices RTN meeting at CEMFI in Madrid, and the 2008 European Economic Association meeting in Milan. Part of this work was completed when the last author was visiting the Department of Economics at Columbia University whom he thanks for their hospitality. Financial support from the Jan Wallander and Tom Hedelius foundation and The Swedish Council for Working Life and Social Research FAS (dnr 2004-2005) is acknowledged.

1 Introduction

A substantial amount of theoretical work has suggested that social interactions within the workplace are an important determinant of worker effort and firm productivity.¹ Recently a growing empirical literature, using matched employer-employee data, has aimed at identifying to what extent social preferences affect productivity in practice. Due to the lack of consistent and reliable measures of effort/productivity across firms, previous studies have been forced to examine social interactions between coworkers' effort using data from single firms or occupations.² While using high quality data from single firms sometimes enhances the ability to tease out what type of mechanisms are underlying the social interaction effect, clearly, the evidence provided by case studies might be difficult to generalize to other populations.

Therefore, in order to shed further light on how coworkers affect each other's behaviour we take an alternative approach compared to previous studies. We focus on how coworkers affect each other's effort in the form of work absence. Work absence is, of course, intimately related to the productivity of the firm. This relationship is perhaps most obvious in firms with justintime type of production technologies. In such firms the output loss associated with the unexpected absence of one worker is not simply equal to the loss of the single worker's output, but potentially also the lost total value of the entire downstream production line. Moreover, in many jobs both management and co-workers may experience difficulties in observing whether an employee shirks or not. The pattern and frequency of work absence is however arguably more easily monitored and may serve as a proxy for the worker's effort. The relative ease of monitoring co-workers' absence suggests that studying this type of behaviour should be rewarding when trying to investigate whether and through which underlying mechanisms co-workers affect each other's effort.

To address many of the severe identification problems associated with the estimation of social interaction effects we make use of the exogenous variation in work absence incentive induced by a large-scale randomized

¹For example Lazear (1989), Kandel and Lazear (1992), and Rotemberg (1994) incorporate social concerns into the analysis of behavior within firms.

²For example in four recent, interesting and related studies Bandiera, Barankay and Rasul (2005); Mas and Moretti (forthcoming); Guryan, Kroft and Notowidigdo (2007); Kato and Shu (2008) use data from a fruit picking farm in the UK, 6 US supermarket stores, male professional golfers on the PGA tour, and a Chinese textile firm respectively.

experiment.³ During the experiment, formal monitoring during absence was relaxed for half of all employees in the city of Göteborg, Sweden. Instead of having to provide a physicians certificate on the eighth day of an absence spell, the treated individuals did not need to provide a certificate until the 15th day of the spell. Treatment assignment was based on date of birth (even/uneven) and applied to all employees living in Göteborg municipality (pop.~500,000). As treatment was determined at the individual level, the number of treated workers within each workplace varied substantially. Using a rich administrative data set we are able to recover the treatment status of all workers, and hence the proportion treated, within each one of the around 3,000 workplaces in operation in Göteborg during the experiment.

We start off by showing that the experiment strongly affected the short term absence level among the treated workers in our sample.⁴ We then focus on whether social preferences affect short-term absence. Our idea for testing for the prevalence of social preferences among the employees is straightforward. We first show that conditional on treatment status the proportion of treated within each workplace is strongly correlated with the change in individual worker's absence. This result suggests that coworkers indeed have an important influence on employees' behaviour.

We then examine the heterogeneity of the treatment effect to explore what types of social preferences are most likely underlying the social interaction effect. Interestingly, we find stark differences in the effect of having many treated co-workers depending on the worker's own treatment status. While the proportion treated at the workplace is a good predictor for the change in absence level among the *non-treated*, no significant effects can be found among the *treated*. The heterogeneous spill-over effects suggest that the observed co-worker effect is *not* driven by preferences for joint leisure. Neither does it seem in line with a hypothesis suggesting that information sharing among co-workers is causing the observed peer effect. Given the randomized treatment assignment, if either of these two mechanisms would be responsible for the estimated peer effect we would expect that also the treated employees absence level should be correlated with the share of treated at the workplace. The heterogeneous response by treated

³See e.g Manski (1993, 1995) for a description of the difficulties in estimation of social interaction effects.

⁴Hesseliuss, Johansson and Larsson (2005) find a statistically significant and large effect on average duration of absence due to the experiment.

and non-treated workers instead suggest that a non-negligible proportion of the workers have reciprocal type of preferences and/or display fairness concerns.⁵

The prevalence of social preference has previously been documented in several laboratory experiments. The evidence of such preferences outside of the laboratory is however more scarce. Our study complement and extend the evidence on the prevalence of social preferences in the workplace found by Bandiera, Barankay and Rasul (2006), Mas and Moretti (forthcoming), and Ichino and Maggi (2000). These studies all use observational data from single firms. Our main contribution to the literature on social preferences and worker effort is to provide evidence on the prevalence of similar social mechanisms at work using a randomized social experiment and linked employer-employee data from around 3,000 workplaces. These two features significantly improve both the internal as well as the external validity of the results found.

The paper unfolds as follows. In Section 2 we briefly describe the general context of our study, providing details on the Swedish sickness insurance system and the experiment. In Section 3 the data and empirical strategy is presented. Section 4 contains the results and Section 5 concludes.

2 Institutional Background and the Experiment

2.1 Temporary Work Absence in Sweden

In Sweden sickness insurance is compulsory and universal to all employees, students and unemployed. It is financed by a proportional pay-roll tax and replaces individuals earnings lost due to temporary health problems. The benefit level received is related to the loss of earnings during the absence spell.

In an international context the sickness benefit levels are, and have been, generous. For most workers the benefit level was set to 90% of previous earnings. Some workers at the very top of the wage distribution were however excluded from receiving the full 90% due to a benefit cap. Besides the public insurance, most Swedish workers are also covered by extra sickness insurance regulated in agreements between the unions and

⁵See e.g. Rabin (1993), Fehr and Gächter (2000) for references and discussions on the importance of such preferences in the labor market.

the employers' confederations. These top-up insurances generally cover about 10% of the lost earnings but there is considerable variation. Hence the total compensation in case of work absence due to illness could be fully 100%.⁶

The public insurance has no limit for how long and how often sickness benefits are paid. Many spells stretch over a full year and there are examples of even longer durations. The individual have a high degree of discretion when to report sick. The benefit payments are generous, and the monitoring before the eighth absence day is lax. A sickness absence spell starts when the worker calls the public insurance office (and her employer), then within a week (on the eighth day) he/she should confirm her eligibility with the insurance office by presenting a medical certificate proving reduced work capacity due to illness. The public insurance office reviews the certificate and then declines or approves *further* sick-leave. In all but *very* few cases the certificate is approved. In case the insurance office suspects abuse they can make unannounced visits to the claimant's home.

2.2 The Experiment

In the fall of 1988 the regional social insurance board in the municipality of Göteborg, the second largest city in Sweden, and in Jämtland, a large and relatively sparsely populated region in the north of Sweden, agreed on performing a social experiment regarding the timing of the requirement for a physician's certificate. A randomly assigned treatment group was allowed to be absent from work due to illness for 14 days before they needed a physician's certificate in order to continue their absence spell with insurance compensation. The control group faced the ordinary restriction of seven days. Individuals were assigned to the treatment and control group based on their date of birth. Those born on even days ended up in the treatment group, and those born on uneven days in the control group. Hence, due to the universal insurance coverage, everyone who was of working age and lived in the experiment regions was assigned to either the control or the treatment group.⁷

⁶See e.g. Johansson and Palme (1996, 2002, and 2005), Henreksson and Persson (2004) for studies on effects of different compensation schemes on individual work absence.

⁷Government employees were exempted, as they, by law, receive their sick leave compensation from the employer instead of from the social security office. The em-

The insurance agencies had several arguments for running the experiment. All were based on a notion that extending the time-period without monitoring would *decrease* costs and *reduce* sickness absence. The main argument was that with the 15-day restriction unnecessary visits to a physician could be avoided, which would cut costs for both the individual and the public health care system. The insurance agency also believed that physicians by routine prescribed longer absence from work than necessary. With an extended certificate-free period of two weeks many individuals would have time to return to work before a medical certificate was needed, and thus individual and public costs would be reduced.

The experiment started on 1 July 1988 and besides the personnel at the social insurance office, all employers and medical centres were informed before or during the experiment. A massive information campaign also preceded the experiment at the two locations, including mass-media coverage, distribution of pamphlets and posters at workplaces, etc. Short information about the experiment was also written on the form that every insured reporting sick needed to fill in and send to the insurance office to receive sickness benefits. The experiment ended on 1 January 1989, at which point the previous system was resurrected.

Hesselius, Johansson and Larsson (2005) show that absence spell durations increased substantially among the treated group compared to the control group. The characteristic spike in exit rates from absence, which before the experiment typically occurred on day 7, was during the experiment postponed and instead occurred on day 14 of the spell (i.e. just before the need for a physicians certificate). They also report heterogeneous treatment effects between women and men. Men prolonged their work absence spells more than women.

The estimates provided by Hesselius et al. did however not control for potential spill-over effects between the treated and the controls. Whether doing so yields an upward or downward biased estimate of the true treatment effect is *a priori* uncertain. The direction of the bias will depend on if individuals care about e.g. co-workers' behaviour and if so *why* they care. In the following section we provide evidence on the effects of co-workers behaviour on individual absence decision and the reason why.

ployer, in turn, receives the benefit from the social security office. We hence exclude all Governmental workplaces.

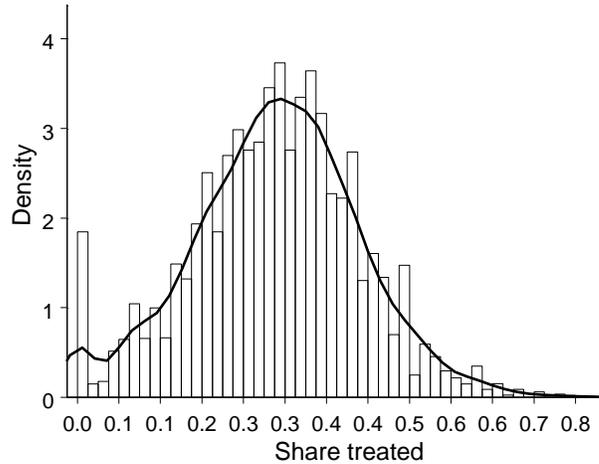


Figure 1: Distribution of treated workers at workplaces in Göteborg

3 Data and Empirical Strategy

We use data from a set of administrative registers compiled by Statistics Sweden. The data contains, besides a set of individual background characteristics, data on start and end date of all absence spells during 1988. We also observe the workplaces where the individual is gainfully employed in November 1988. A few individuals have multiple workplaces, but for simplicity we assume that the workplace from which the highest yearly earning is received is also the main arena for co-worker interaction. The treatment status of each worker was decided by date of birth (even/uneven) and whether the individual is residing in Göteborg municipality or not. Figure 1 show the distribution of treated workers at the workplaces in Göteborg municipality.

As seen in Figure 1 the between workplace variance in proportion of treated is considerable. The average workplace has around 35% of treated workers. The variation in proportion treated workers stem from the random assignment of treatment, but also from the number of commuting co-workers. In the main analysis we focus on workplaces with between 10 and 100 employees as social interactions is probably most prevalent in

small to medium sized workplaces. The workplaces with 10 employees and less are excluded from the sample as alternative rules may apply to these workplaces.

The large variation in proportion treated workers across workplaces provides us with a close to ideal setting to identify the effects of co-worker interactions. Our idea to identify the prevalence of co-worker interactions is straightforward. Given the random assignment of treatment status, if the proportion treated within each workplace affects individual work absence decisions we can be certain that co-workers indeed affect each other’s behavior. The baseline model we estimate by OLS is specified in equation (1)

$$\Delta s_{ij} = \alpha_0 + \beta_0 \textit{treated}_i + \beta_1 \pi_j + \varepsilon_{ij} \quad (1)$$

where Δs_{ij} is the change in number of days of short term absence (shorter than 15 days) between first and second half of 1988 of employee i at workplace j .⁸ *Treated* takes the value 1 if the employee resides in Göteborg and is born on an even date, and 0 otherwise. π_j is the proportion of treated co-workers at employee i ’s workplace (excluding employee i). Given the first differencing and the random assignment of treatment, a significant estimate of β_1 identifies spill-over effects among the employees.⁹ Inference is based on robust standard errors allowing for clustering at the workplace level.

⁸The empirical analysis focuses on the change in work absence to control for the possibility that workplaces with different shares of commuters also are systematically different in some unobserved way. The first differencing approach controls for time-invariant unobserved individual and workplace heterogeneity.

⁹Note that if we only used the directly non-treated individuals when estimating equation (1) could this specification be seen as a reduced form in the estimation of endogenous effects. This second stage structural equation would, among others, require rational expectations of the individuals, that is; the correct prediction of the coworkers response from the more lenient monitoring. Equation (1) does not require that the observed reaction to the fraction of treated is due to an actual increase in absence among treated coworkers. Hence, (non-rational) expectation of an increase in shirking among treated co-workers during the experiment may have an effect on behavior.

Table 1: Direct and indirect effects of experiment

Specification	1	2	3
Sample:	All	All	All: pre- experiment
Treated	.886*** (.057)	.889*** (.058)	.061 (.050)
Proportion treated		.554** (.210)	-.134 (0.193)
<i>R</i> -squared	0.0035	0.0036	0.000
# observations	79,643	79,643	77,647

Note: Dependent variable is the change in number of days in short term absence (spells shorter than 15 days). */**/** denoted statistical significance at 10/5/1 percent level, respectively. The number of co-workers, gender, age and annual earnings is included as control variables. Standard errors are reported in parenthesis and are cluster adjusted at the workplace level. Number of workplaces is 3,008 for the experiment period and 2,910 for the pre-experiment period (1987).

4 Results

4.1 Main Results

In column (1) of Table 1 we start by presenting the “naïve” estimate of the direct treatment effect, β_0 , from a specification using only the treatment status as explanatory variable. The estimated effect suggests that being assigned to treatment increases the short-term work absence with 0.88 days on average. In column (2) of Table 1 both estimates of β_0 and β_1 from equation (1) is reported. The share of treated co-workers indeed increases the short-term absence level substantially.

The finding that the share of treated co-workers within a workplace affects individual worker’s absence decisions provides clear evidence on the importance of peer effects in the workplace. The indirect effect is furthermore sizable in relationship to the direct effect of being treated. An increase in the proportion of treated colleagues from 0 to 1 increases the change in absence by 0.55 days in average. Both the direct and the indirect treatment effects are furthermore substantial given that during the first half of 1988 the average number of short-term absence days was 2.32 in our sample.

We have also performed a number of specification checks in order to

assess the robustness of the results. First we re-ran the same analysis using data from the first and second half of 1987, i.e. the year prior to the experiment. The result from this exercise is presented in column (3) of Table 1. The effects of the share of treated and individual treatment are both statistically insignificant and in the case of share of treated co-workers the sign is changed. This is in sharp contrast with the effect on absence during the experiment. That is, in the year prior to the experiment the share of employees being born on an even date had no significant effect on absence, but a large positive effect during the experiment. This result strengthens our confidence in the validity of the estimation strategy.

We have also re-estimated the same models for different workplace sizes and also used the change in total number of absence days as well as the change in number of absence days in spells shorter than eighth days as the dependent variable.¹⁰ These changes in specification yielded qualitatively similar results. In the smallest workplaces (10-20 employees) considered the estimated spill-over effect is largest and then decreases monotonically as we stepwise enter larger workplaces. We have also checked for nonlinearities by adding a quadratic term in the share of treated to the specification. These estimates suggested that there is a tendency towards a concave relationship (significant at 10% level) between work absence and the share of treated co-workers (maximum when 42% of the co-workers are treated). Using the natural log of the share of treated instead of the share of treated yielded similar results.¹¹ Finally, we also estimated separate models for males and females. The separately estimated models produce slightly higher and more precise estimates for males than for females.

4.2 What mechanisms are underlying the spill-over effect?

Many studies have found that social interactions matter for individual behavior in various settings. Very few studies have however with any certainty been able to tell what type of mechanisms is underlying the effect of peers.¹² This is naturally a much more difficult question to answer empirically. We believe that the spill-over effects found in this setting can be ex-

¹⁰These results are retained due to space limitations but are available upon request from the authors.

¹¹In this specification the fraction of treated was multiplied by 100 and those without any treated were then assigned as having 1% treated before taking the logs.

¹²Two notable exceptions are Mas and Moretti (forthcoming) and Bandiera, Barankay and Rasul (2005).

plained in three ways: 1.) *Joint leisure*: co-workers may form tight bonds and hence enjoy leisure time together. A higher absence level as a result of a higher proportion treated co-workers could therefore be explained by preferences for joint leisure time. 2.) *Information sharing*: Prior to the experiment, information about the absence leave system may have been incomplete. A higher proportion of treated workers may therefore be correlated with average absence level at the workplace simply because more workers now become aware of how the system works. However, given the large information campaign about the setup of the experiment in action prior to and during the experiment, the importance of co-workers as an information channel is probably reduced in this context. We therefore do not expect information sharing between co-workers to have a large effect on work absence. 3.) *Reciprocity/Fairness*: Observing a sudden increased absence level among treated co-workers may induce resentment and lead to ill feelings towards the shirking co-workers. For example, if the workload of the absent worker is shifted to the remaining workers, the absence of a co-worker is costly as the remaining workers may need to increase their effort. If shirking co-workers induces such costs, formal or informal social sanctions of this behaviour might be warranted. While there is a number of imaginable ways through which punishment may take place, a natural retaliation could be to increase one's own absence level for fairness reasons.

The two first hypotheses are more in line with altruistic type of social preferences while the third is more in line with non-altruistic social preferences. Moreover, while the first two hypotheses arguably predict similar (or smaller) effects on non-treated and treated, the third hypothesis suggest that non-treated should respond more strongly to an increasing proportion of treated co-workers.¹³

The results presented in Table 2 helps us distinguish which one of the suggested explanations is most likely driving the observed spill-over effect. Column (1) and (2) reports the estimate of β_1 from equation (1) for the treated and non-treated employees separately. Interestingly, the share of

¹³An additional channel through which an increase in work absence among co-workers might affect individual work absence is through direct negative effects on health; from e.g. the stress of facing an increased workload. Although we cannot completely rule it out, we believe that the short duration of the experiment and our focus on short-term absence most likely diminishes the risk of such effects to have any major influence on the estimates. Also, note that such a mechanism would probably yield a similar effect on both treated and non-treated.

Table 2: Heterogeneous spill-over effects

Specification	1	2	3	4
Dependent variable	#(<8 days)	#(<8 days)	#(<15 days)	#(<15 days)
Sample:	Treated	Non-treated	Treated	Non-treated
Proportion treated	-.023 (.249)	.365*** (.141)	.270 (.457)	.600*** (.223)
# Observations	23,803	55,840	23,803	55,840

Note: Dependent variable is the change in number of days in short-term absence less than 8 and 15 days in columns (1),(2) and (3),(4) respectively. The number of co-workers, gender, age and annual earnings is also included as a control variables. */**/** denotes statistical significance at 10/5/1 percent level respectively. Standard errors presented in parenthesis are clustered at the workplace level. Number of workplaces is 3008.

treated has a negligible and insignificant effect on the treated workers' absence decisions. On the contrary, for the non-treated workers we find a large and significant effect of having many treated co-workers. Columns (3) and (4) display that the same pattern hold when instead the change in number of days in spells shorter than 15 days is used as the dependent variable.

The results in Table 2 clearly point towards the *reciprocity* hypothesis and away from the other two hypotheses suggested; *joint leisure* and *information sharing*. If joint leisure was the main motivator behind the response to the share of treated within the workplace, due to the random assignment of treatment, we would expect that not *only* the non-treated reacted but *also* the treated co-workers. As seen above this is not the case. The same argument holds for the information sharing hypothesis. Information sharing should increase absence not *only* for the non-treated but arguably also the treated. The large information campaign implemented prior to and during the experiment as discussed above probably reduced the importance of treated peers as an information channel. The relatively weak effects among treated and strong effect among non-treated therefore more likely reflect reciprocal type of social preferences or fairness concerns among the workers in our setting.

However, as pointed out by e.g. Sobel (1993), it should be noted that due to the possibility of repeated interactions it is notoriously difficult to distinguish between reciprocal social preferences and pure self-interest

outside of the laboratory. At the very least, as for Mas and Moretti (forthcoming) and Bandiera, Barankay and Rasul (2005), our results suggest that *pure altruistic* social preferences can be ruled out as an explanation for the interaction effect in behaviour among co-workers.

5 Concluding Remarks

We provide preliminary evidence suggesting that social interactions are an important determinant of worker effort as measured by work absence by using data from a social experiment affecting work absence incentives of 70,000 employees in 3,000 workplaces. As previous evidence on the prevalence of social preferences comes from laboratory experiments or from studies using observational data from *single* firms, our study constitutes a significant contribution to the current literature on spill-over effects at work.

It should be noted that the reduced form analysis we employ here prohibits us from drawing definite conclusions about the exact nature of the underlying causes of the observed social interaction effects. In future work we plan to apply a more structural approach, and also use data on the exact timing of the absence spells for treated and non-treated in order to get a better understanding of the nature of the underlying social preferences.

References

- Bandiera, Oriana, Iwan Barankay and Imran Rasul (2005). "Social preferences and the response to incentives: evidence from personnel data", *Quarterly Journal of Economics*, 120, 917-962.
- Fehr, Ernst and Simon Gächter (2000). "Fairness and Retaliation: The Economics of Reciprocity". *Journal of Economic Perspectives* 14, 159-181
- Guryan, Jonathan, Kory Kroft and Matt Notowidigdo (2007). "Peer effects in the workplace: Evidence from random groupings in professional golf tournaments", NBER working paper # 13422.
- Henreksson, Magnus, and Mats Persson (2004). "The Effects on Sick Leave of Changes in the Sickness Insurance System", *Journal of Labor Economics*, 22, 87-113.
- Hesselius, Patrik, Per Johansson, Laura Larsson (2005). "Monitoring sickness insurance claimants: evidence from a social experiment", *IFAU working paper* 2005:15.
- Johansson, Per, and Mårten Palme (1996). "Do Economic Incentives Affect Work Absence? Empirical Evidence Using Swedish Micro Data", *Journal of Public Economics*, 59, 195-218.
- Johansson, Per, and Mårten Palme (2002). "Assessing the Effects of a Compulsory Sickness Insurance on Worker Absenteeism", *Journal of Human Resources*, 37, 381-409.
- Johansson, Per, and Mårten Palme (2005). "Moral hazard and sickness insurance", *Journal of Public Economics*, 89, 1879-1890.
- Kato, Takao, and Pian Shu (2008). "Performance spill-over and social networks in the workplace: Evidence from rural and Urban weavers in a Chinese textile Firm", IZA Discussion paper series no. 3340, February.
- Kandel, Eugene, and Edward Lazear (1992). "Peer Pressure and Partnerships" *Journal of Political Economy* 100: 801-13.
- Lazear, Edward (1989). "Pay Equality and Industrial Politics," *Journal of Political Economy*, LXXXVII, 1261-1284.
- Manski, Charles (1993). "Identification of Exogenous Social Effects: The Reflection Problem," *Review of Economic Studies*, LX, 531-542.
- Manski, Charles (1995). *Identification Problems in the Social Sciences*, Cambridge:Harvard University Press.

Mas, Alexandre, and Enrico Moretti (2008). “Peers at work”, forthcoming *American Economic Review*.

Ichino, Andrea, and Giovanni Maggi (2000). “Work Environment And Individual Background: Explaining Regional Shirking Differentials In A Large Italian Firm”, *Quarterly Journal of Economics* 115, 1057-1090.

Rabin, Matthew (1993). “Incorporating fairness into game theory”, *American Economic Review* 83. 1281–1302.

Rotemberg, Julio (1994). “Human relations in the workplace”, *Journal of Political Economy* 102, 684–718.

Sobel, Joel (2005). “Interdependent preferences and reciprocity,” *Journal of Economic Literature*, XLIII, 392–436.