

DISCUSSION PAPER SERIES

No. 10928

NETWORK EFFECTS ON WORKER PRODUCTIVITY

Matthew Lindquist, Jan Sauermann
and Yves Zenou

LABOUR ECONOMICS



NETWORK EFFECTS ON WORKER PRODUCTIVITY

Matthew Lindquist, Jan Sauermann and Yves Zenou

Discussion Paper No. 10928

November 2015

Submitted 29 October 2015

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: (44 20) 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **LABOUR ECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Matthew Lindquist, Jan Sauermann and Yves Zenou

NETWORK EFFECTS ON WORKER PRODUCTIVITY

Abstract

We use data from an in-house call center of a multi-national mobile network operator to study how co-worker productivity affects worker productivity via network effects. We also exploit data from a field experiment to analyze how exogenous changes in worker productivity due to on-the-job training affect co-worker productivity, including non-trained workers. We show that there are strong network effects in co-worker productivity. This effect is driven by conformist behavior. We also show that exposure to trained workers increases the productivity of non-trained workers. This effect works through strategic complementarities (knowledge spillovers). We demonstrate how our network model of worker productivity can be used to inform a variety of practical decisions faced by personnel managers including the design of optimal training policy.

JEL Classification: J24, M53 and Z13

Keywords: on-the-job training, peer effects, social networks and worker productivity

Matthew Lindquist Matthew.Lindquist@sofi.su.se
SOFI, Stockholm University

Jan Sauermann jan.sauermann@sofi.su.se
SOFI, Stockholm University, IZA Bonn and ROA Maastricht

Yves Zenou yves.zenou@ne.su.se
Stockholm University, IFN, IZA and CEPR

1 Introduction

There is a growing empirical literature showing that workers are sensitive to the output choices of their peers. Falk and Ichino (2006), Mas and Moretti (2009), and Bandiera, Barankay and Rasul (2010) all demonstrate that co-workers can exert economically significant effects on their peers, via channels often not explicitly created by their firms. Yet it is not always clear what are the exact economic mechanisms behind these effects and what are the policy consequences of such results. The aim of this paper is to shed more light on these questions by studying how co-worker productivity affects worker productivity via network effects. In particular, we would like to answer the following questions: Does co-worker productivity affect worker productivity? If so, then what are the mechanisms? Do productivity increases from on-the-job training spread from trained workers to their untrained co-workers? How does the structure of a co-worker network enhance or impede the spread of productivity through the network? Is it possible to identify key workers in the firm? If so, then how can such information be used to aide personnel managers when making important decisions about who to train and who to retain? Can we use information about co-worker networks to help firms organize work hours and teams in a more optimal fashion or to improve upon the design of their training programs?

We study these questions using a high quality data set from an in-house call center of a multi-national mobile network operator that covers a period of two years. The data include detailed information on the performance of individual workers, their characteristics, their team affiliation, and the exact times that they punch in and out of work. These data allow us to create co-worker networks for each week, where a weighted link between two agents indicates the amount of time they have been working together on the same team during a week. That is, we have very precise knowledge of the co-workers that a worker is exposed to during the week together with the intensity of this exposure. We show that these co-worker networks are as good as random after conditioning on team, week and individual fixed effects. Unlike most network papers, we do not identify exposure to peers off of the stable part of the network, since the stable part may be prone to non-random sorting and to common shocks. In our setup, every worker receives a unique, exogenously varying dose of co-worker productivity each week due to worker turnover, due to changes in the scheduling needs of the company, and due to idiosyncratic changes in one's own work schedule and the work schedules of teammates.

We present a formal model of worker productivity that includes both local average network effects and local aggregate network effects. The local average effect represents the role

of social work norms (e.g. conformist behavior or peer pressure), while the local aggregate effect represents strategic complementarities (e.g. knowledge spillovers). As discussed in our literature review, these are the two main mechanisms that are typically put forward to explain productivity spillovers among workers.

We use this model to guide the empirical part of our paper. Our estimation equation is the best reply function from the Nash equilibrium of our model. We test to see which type of network effect is most relevant in our particular setting. But we allow for the inclusion of both mechanisms if needed.

We run two different regression experiments. We first use our exogenous network exposure matrix to study the effect of co-worker productivity on worker productivity. We show that there are indeed strong network effects in worker productivity. A 10% increase in the current productivity of a worker’s co-worker network leads to a 1.7% increase in own current productivity. We attribute this productivity spillover to conformist behavior. Low tenure workers react particularly strong to the work norm of their co-worker network.

In our second regression experiment, we exploit data from a field experiment with random assignment of workers to a one week on-the-job training program to analyze how exogenous changes in worker productivity due to on-the-job training affect co-worker productivity, including non-trained workers. Previous work by De Grip and Sauermann (2012) showed that this particular training program had a positive causal impact on the productivity of the call center agents analyzed in our paper.¹ We show that there are strong network effects in co-worker productivity and that being exposed to trained workers increases the productivity of non-trained workers. Adding one additional trained co-worker to a worker’s network increases the worker’s own productivity by 0.7%. We show that this effect is driven by knowledge spillovers (strategic complementarities) and not by conformist behavior.

In the policy section of our paper, we show how these empirical findings can be used in conjunction with our model-based network approach to address several important personnel management questions. The answers to these questions hinge crucially upon the presence of externalities in worker productivity from network effects, the underlying mechanism, and on the structure of co-worker networks.

The rest of the paper unfolds as follows. In the next section, we review related studies and highlight our contribution with respect to these studies. Section 3 exposes our theoretical

¹The data used in this study were originally collected for the purpose of evaluating a randomized on-the-job training experiment. De Grip and Sauermann (2012) only use a small share of the data that they collected ($N = 74$). In this paper, we make use of the entire data set, including data from the field experiment ($N = 425$).

framework, while Section 4 describes the data used in each regression experiment. In Section 5, we describe how we define and construct our co-worker networks. Section 6 is devoted to our first regression experiment concerning the effect of co-worker productivity on worker productivity and its results. Section 7 focuses on our second regression experiment, which studies network effects in on-the-job training. In Section 8, we demonstrate the relevance of our findings by using them to answer a number of practical policy questions faced by a typical personnel manager. Section 9 concludes.

2 Related literature

In this section, we discuss a number of related studies. We highlight important similarities and differences between our work and previous work, emphasizing what we believe to be our most important original contributions to the literature.

2.1 Peer effects in worker productivity

A number of important studies have been published that convincingly show that there are meaningful peer effects in worker productivity, at least in some settings.² Falk and Ichino (2006) provide evidence on peer effects in worker productivity by randomly varying whether an experimental subject worked on a simple task alone or together with another subject in the same room. Although the task (stuffing envelopes) was individual, the presence of a second subject in the same room increased a workers' productivity significantly. A 10% increase in co-worker productivity increased worker productivity by 1.4%. Low productivity workers respond more strongly to the productivity of their peers.

In their study of supermarket cashiers, Mas and Moretti (2009) analyze peer effects in worker productivity by exploiting the quasi-random placement of cashiers who just started

²We use the term *network effect* rather than *peer effect* to describe what we do in this paper, since we want to make it clear that we are studying productivity spillovers among a network of inter-linked co-workers. Peer effects, on the other hand, are usually conceived as an average intra-group externality that affects all members of a given group identically. For example, if we look at two prominent papers in this literature (Sacerdote 2001 and Carrell et al. 2009), all individuals who are members of the same “unit” (dormitories for the former and squadrons for the latter) are considered peers. This means that the group boundaries for such a homogeneous effect are often determined at a rather aggregate level and that all people from the same “unit” are assumed to interact with each other equally. In this paper, teams of workers are clearly not fully and equally interlinked. “Peers” in our setting are those co-workers that each worker is linked to either directly or indirectly. Also, each link is weighted by the amount of time two workers spend working together. As such, the “peer” relationship is quite heterogenous within each team.

their shift to registers at which they can observe some of the other cashiers and at which they can be observed by some of the other cashiers. They show that cashiers work faster when being observed by highly productive workers, suggesting that co-workers may face sanctions from their peers for low output. They find that a 10% increase in co-worker productivity raises worker productivity by 1.5%. Like Falk and Ichino (2006), they also find that is the less productive workers who respond to the more productive work norm set by their high productivity peers. In line with the arguments made by Kandel and Lazear (1992), their findings indicate that peer pressure in the work place may help overcome the problem of free-riding. Their findings also imply that the firm could increase overall productivity by striving to have skill level diversification across shifts.

Bandiera, Barankay and Rasul (2010) combine self-reported information on friendships among fruit pickers in the UK with random assignment of workers to rows on a field. They show that, even in the absence of externalities from either the production technology or the compensation scheme, workers adjust their effort when working with friends. More productive workers are willing to forgo up to 10% of their own earnings when working with less productive friends and less productive friends are willing to exert 10% more effort when working with more productive friends. The net effect on aggregate performance among fruit pickers in this firm is positive, which suggests that firms could harness these types of social incentives to boost aggregate production.

The actual expression of potential peer effects in the work place and the role played by social connectedness may depend on the type of payment system in place in the firm. Chan, Li and Pierce (2014), for example, study peer effects among salespersons under individual-based and team-based pay: the presence of high ability peers improves performance for low ability peers under team-based compensation, but not under the individual-based pay system. Babcock et al. (2015) show that peer effects can arise from monetary team incentives and that these social incentives can be quite effective³ in motivating effort intensive tasks.

All of the above mentioned studies examine peer effects in worker productivity among workers performing relatively low skilled tasks. One advantage of studying low skilled tasks is that they can often be quantified and, hence, allow for rather precise measures of individual worker productivity. As discussed above, this strand of research typically finds that peer effects are best explained by peer pressure (and work norms) in the workplace. Hamilton et

³It is possible, however, that social connectedness between managers and workers could lead to favoritism and/or nepotism that could potentially harm firm productivity. Bandiera, Barankay and Rasul (2009) show that managers favor workers to which they are socially connected if managers are paid fixed wages. But when manager compensation is performance-related, they change their behavior and favor more productive workers.

al. (2003), however, demonstrate that peer effects among low skilled workers can also arise when production is team based and when team members have collaborative skills or when team members can specialize in different steps of the production process.

The evidence concerning the existence of peer effects among high skilled workers is more mixed. Also, the mechanism that most authors point towards when discussing peer effects among high skilled workers is knowledge spillovers (i.e. strategic complementarities) as opposed to work norms and peer pressure.

Jackson and Bruegemann (2009) show that teacher output, measured by students' grades, is higher when the teacher has more effective colleagues. The effect is particularly large for less experienced teachers and it appears to persist over time. These results imply that such peer effects are likely driven by peer-to-peer learning (knowledge spillovers).

Azoulay et al. (2010) find that researchers collaborating with "super star" scientists experienced a lasting and significant decline in their quality adjusted publication rate after the unexpected death of their super star colleague. Waldinger (2010) provides evidence that the expulsion of high quality Jewish scientists from Germany by the Nazi government resulted in negative effects on the productivity of Ph.D. students that were left behind. But he finds no such negative peer effects among faculty members that were left behind (Waldinger 2012).

Recent work by Cornelissen, Dustmann and Schönberg (2013) study peer effects in wages using worker-firm matched data for an entire local labor market in Germany. They find very small average effects of peer productivity on own wages. A 10% increase in peer quality increases own wage by only 0.1%. Thus, they reject the notion that there are large productivity spillovers in wages for a representative set of occupations and firms. When looking at the most repetitive and predefined occupations, those that are likely to be most susceptible to peer pressure, they find that a 10% increase in peer quality increases own wages by 0.84%, which is just over half the number reported in Mas and Moretti (2009). When looking at high skilled an innovative occupations, they find spillover effects that are as small as those for the economy as a whole.

2.1.1 Our contribution

To the best of our knowledge, our paper is the first to use an explicit network approach to study the effect of co-worker productivity on worker productivity. This allows us to make several important contributions to the literature on peer effects in the workplace.

First, our network approach allows us to estimate causal network effects of contemporaneous co-worker productivity on contemporaneous worker productivity. We are measuring the

total effect that arises from taking a worker’s full network of co-workers into consideration, as opposed to only considering the effect of her immediate peers. Second, our model based approach allows us to distinguish between different mechanisms underlying these network effects. Third, our model based network approach allows us to run a set of well defined policy experiments. We show how important it is for personnel managers to take the existence of network effects and the structure of work place networks into consideration when designing firm policy.

Our first regression experiment, looking at the effect of co-worker productivity on worker productivity, is very similar to Mas and Morreti’s (2009) experiment using data on grocery store cashiers. Both studies use detailed time clock data to measure exposure to peers and both have very precise, automated measures of worker productivity. The main difference between our first experiment and their experiment is that we allow a worker to be influenced by her entire co-worker network, while the cashiers in Mas and Morreti’s (2009) paper can only be influenced by persons working on the same shift. It is quite likely that this is, in fact, the most relevant specification in their context. In our context, however, we see that workers can be influenced by workers on their team whom they don’t actually see during the week. That is, the co-workers of co-workers also matter. They influence worker productivity indirectly by influencing the productivity of a worker’s co-workers.

2.2 The effect of training on worker productivity

There is an extensive literature on firms’ incentives to sponsor training, starting with Becker (1962). In a competitive environment, Becker conjectured that firms should only pay for specific training since it increases workers’ productivity only in their current job while general training should be wholly financed by workers themselves as it improves their productivity in future jobs as well. The contributions in the late 1990’s (Acemoglu, 1997; Acemoglu and Pischke, 1998, 1999), which consider frictional markets, have questioned Becker’s argument: firms do have incentives to finance general training and firms’ investments in specific training are subjected to hold-up issues. Other papers put the emphasis on firms’ incentives to invest in training. Acemoglu (1997), Fella (2005) and Lechthaler (2009) show that firing costs help raise firms’ investment in general training, so optimal training choices can be achieved.⁴

Most studies analyzing the returns to training focused on estimating wage or establishment productivity (Bartel, 1995; Loewenstein and Spletzer, 1999; Goux and Maurin, 2000; Dearden, Reed and Van Reenen, 2006; Konings and Vanormelingen 2015). A few studies

⁴See Leuven (2005) for an overview on these issues.

in economics exploit personnel data and experimental variation to provide causal evidence on the returns to training. These studies examine the returns to training using data from the manufacturing sector (Bartel, 1995; Krueger and Rouse, 1998; Breuer and Kampkötter, 2013; Pfeifer et al., 2013), the public sector (Gelderblom and de Koning, 1996), and the service sector (Krueger and Rouse, 1998; Liu and Batt, 2007), and use either supervisor ratings or more objective measures of performance. Because training incidence is likely to be endogenous, some of these studies have used fixed effects regressions.

To the best of our knowledge, De Grip and Sauermann (2012) is the only study that exploits randomized on-the-job training to study the causal effect of training on individual worker productivity. Using a sample of 74 call center workers (a subsample of the call center workers used in this paper), they report an 8.8% increase in productivity among the 34 workers who were randomized into the week long training program relative to the 40 control workers.

De Grip and Sauermann (2012) also investigate potential spillover effects from trained workers onto untrained workers. Since the training of some of the treated teams was split into two sessions (due to capacity constraints of the training facility), they were also able to estimate a positive externality from the early trained ($N = 13$) on the late trained ($N = 21$). An increase of 10 percentage points in the share of early treated peers (during the one to two weeks that they were waiting to be trained) improved the performance of their late trained peers' performance by 0.51%. Although these results are only marginally significant (at the 10% level), they do give us a first indication of the potential presence of spillover effects from on-the-job training.⁵

2.2.1 Our contribution

Our contributions to the on-the-job training literature are threefold. First, we propose to study spillover effects using two new identification strategies, both of which use much larger samples and longer evaluation periods than the approach used in De Grip and Sauermann (2012). Our first strategy relies on the randomization of the field experiment. We study spillover effects among the large group of never treated co-workers of the treated group and compare their outcomes to the outcomes of the never treated co-workers of the control group. Our second strategy uses all workers in our data and relies upon our exogenous exposure matrix to achieve identification. Larger samples and alternative methods allow us to present more precise and reliable estimate of spillover effects that are relevant for the entire set of workers in this firm. We can study if and how these spillover effects differ across workers

⁵This imprecision is likely due to the small sample size and the short evaluation period.

with high versus low tenure. We also avoid measuring any pre-treatment effects, since we look at spillover effects among workers who are not scheduled for training.

Our second contribution is that we can use these training data together with our network model to pin down the mechanism more precisely. Lastly, we can use our estimates together with our model to answer important questions concerning the costs and benefits of training and the role played by the structure of the co-worker network in enhancing knowledge spillovers. We can also talk more precisely about who the firm should train. These types of policy experiments concerning optimal training have not been done before. More generally, we believe that our results clearly show that a network approach can provide valuable new information to the on-the-job training literature.

2.3 Networks

There is a growing literature on the economics of networks, both theoretical and empirical.⁶ The main empirical challenge faced by the network literature is that of obtaining a causal relationship between the outcome of an individual and the outcomes of her peers. The two main threats to identification of peer effects in the context of social networks are *non-random sorting* into networks (endogenous network formation) and *correlated shocks*. These phenomena imply that the estimation of peer effects might be flawed because of the presence of peer-group specific unobservable factors affecting both individual and peer behavior.⁷

The network literature has dealt with these two difficulties in a number of ways. Some researchers have simultaneously modeled the network formation process and the outcome peer-effect equation and estimate both equations together (Goldsmith-Pinkham and Imbens, 2013; Badev, 2014; Hsieh and Lee, 2015; Del Bello, Patacchini and Zenou, 2015; Boucher, 2015; Patacchini, Rainone and Zenou, 2015). A second approach has been to use instrumental variables to identify the exogenous part of an otherwise endogenous network and to correct for correlated shocks (Bifulco, Fletcher and Ross, 2011; Patacchini and Zenou, 2014).

The network literature has also seen a rapid growth in the number of controlled experiments which provide identification. Experiments have been implemented by either (*i*) fully

⁶For overviews, see Jackson (2008, 2011, 2014), Blume et al. (2011), Ioannides (2012), Aral (2016), Boucher and Fortin (2016), de Paula (2015), Graham (2015), Jackson and Zenou (2015), Jackson et al. (2015), Topa and Zenou (2015).

⁷The *reflection problem* (Manski, 1993) is usually an issue when one studies peer effects in linear-in-means models. However, when using explicit network data, the reflection problem can be readily solved using the structure of the network; see Bramoullé, Djebbari and Fortin (2009). See Blume et al. (2011 and 2015) for a complete discussion of the identification of peer effects in social networks.

controlling for the network of relationships in the laboratory (Choi et al., 2012; Kearns et al., 2009; Charness et al., 2013; Aral, 2016) or (ii) assigning subjects in the field to specific positions in a network through which they must communicate (Centola, 2010, 2011; Goeree et al., 2010; Babcock and Hartman, 2010). Some field experiments match subjects together, but do not control for who communicates with whom. Examples include Cai et al. (2015), Breza and Chandrasekhar (2015), Paluck, Shepherd and Aronow (2015) and Beaman et al. (2015). In these papers, it is the intervention that is randomized, not the network. There are also a few recent field experiments where agents are randomly allocated to networks (Algan et al., 2015; Hahn et al., 2015).⁸

2.3.1 Our contribution

Our contributions to the network literature are threefold. First, our paper appears to be the first to apply an explicit network approach to individual level personnel data to study the effect of co-worker productivity on worker productivity. We are also the first to apply a network model and methods to study productivity spillover effects from on-the-job training. Our hope is that the literature on social networks will expand more vigorously into the field of personnel economics; a field that we believe is particularly suited for network methods and models.

Our second contribution to the network literature is our solution to the problem of non-random sorting and correlated shocks. We adapt the identification strategies of Bayer, Hjalmarsson and Pozen (2009) and Mas and Moretti (2009) to an explicit network framework.⁹ We also combine our “exogenous exposure matrix” with data from a field experiment to study the causal effect of on-the-job training on worker productivity. We believe that this “exposure matrix” approach could be used quite fruitfully to study various types of interactions in social networks.

Third, we demonstrate the value of positing a hybrid model of network effects (as opposed to simply positing a model based on peer averages as is often times done). With

⁸There are, of course, many papers that use a random allocation of individuals to assess peer effects but very few that look explicitly at network effects.

⁹Bayer et al. (2009) study peer effects in recidivism among juvenile offenders. To control for non-random assignment of juveniles to correctional facilities, they include facility and facility-by-prior offense fixed effects in their regressions. This ensures that the estimated peer effect is identified using only the variation in the length of time that any two juvenile offenders who are placed in the same facility happen to overlap. As discussed above, Mas and Moretti (2009) use random shift assignments to measure overlap in hours worked by grocery store cashiers. We also use overlap in hours worked to define peer exposure, although we extend this approach to the full network of a worker’s co-workers (more on this below).

such a model in hand, researchers can then run a series of step-wise tests to arrive at the most empirically relevant specification of their model. It is this model that should be used (along with a set of unbiased parameter estimates) to run policy relevant experiments. In particular, by defining “key workers” using our model, we are able to answer questions such as: Who should the firm strive to retain? Who should the firm let go? Who should the firm train?

3 Theoretical Framework

3.1 Notations and preferences

A co-worker network, g , is a collection of $N = \{1, \dots, n\}$ workers and the links between them. The adjacency matrix $\mathbf{G} = \{g_{ij}\}$ keeps track of these links, where $g_{ij} = 1$ if i and j are co-workers, and $g_{ij} = 0$, otherwise. In this paper, co-worker links are defined as those who work the same shift on the same team in the same company.¹⁰ Links are reciprocal so that $g_{ij} = g_{ji}$. We also set $g_{ii} = 0$ so that individuals are not linked to themselves. The adjacency matrix is thus a 0 – 1 symmetric matrix describing the architecture of a co-worker network. Denote by $\mathbf{G}^* = \{g_{ij}^*\}$ the row-normalized matrix of \mathbf{G} where $g_{ij}^* = g_{ij}/g_i$, where $g_i = \sum_{j=1}^n g_{ij}$ is the number of links (co-workers) of individual i .

Individuals decide how much productive effort to exert on the job. We denote the effort level of individual i by y_i and the population effort profile by $\mathbf{y} = (y_1, \dots, y_n)'$. Each agent i selects an effort $y_i \geq 0$, and obtains a payoff $u_i(\mathbf{y}, g)$ that depends on the effort profile \mathbf{y} and the underlying network g , in the following way:

$$u_i(\mathbf{y}, g) = (a_i + \eta + \epsilon_i) y_i - \frac{1}{2} y_i^2 + \lambda_1 \sum_{j=1}^n g_{ij} y_j y_i - \frac{1}{2} \lambda_2 \left(y_i - \sum_{j=1}^n g_{ij}^* y_j \right)^2 \quad (1)$$

where $\lambda_1 > 0$, $\lambda_2 > 0$. The structure of this utility function is an extension of the one usually used in games on networks (Ballester, Calvó-Armengol and Zenou, 2006; Calvó-Armengol, Patacchini and Zenou, 2009; Patacchini and Zenou, 2012; Bramoullé, Kranton and d’Amours, 2014; Jackson and Zenou, 2015) where both *local-aggregate* and *local-average effects* are incorporated in (1). This utility function has been introduced by Liu, Patacchini and Zenou (2014) and referred to as the *hybrid utility function*. Indeed, there are two network effects in (1). The first network term $\sum_{j=1}^n g_{ij} y_j y_i$ represents the *aggregate* effort of i ’s co-workers with the *social-multiplier* coefficient λ_1 . As individuals may have different locations

¹⁰In Section 5, we define the co-worker networks in our data more precisely.

in the network, $\sum_{j=1}^n g_{ij}y_jy_i$ is heterogeneous in i even if every individual in the network chooses the same effort level. The second network term $\left(y_i - \sum_{j=1}^n g_{ij}^*y_j\right)^2$ represents the moral cost due to deviation from the *social norm* of the reference group (i.e., the average effort of the peers) with the *social-conformity* coefficient λ_2 . Thus, an individual's utility is positively affected by the total effort of her co-workers *and* negatively affected by the distance from the average effort of her co-workers. If $\lambda_1 = 0$, we obtain the *local-average model* while, if $\lambda_2 = 0$, we have the *local-aggregate model*.

In (1), there is also an idiosyncratic exogenous part, $(a_i + \eta + \epsilon_i)y_i - \frac{1}{2}y_i^2$, where a_i represents the ex ante individual *observable* heterogeneity in the return to effort, ϵ_i captured the *unobservable* individual heterogeneity, η , the *network fixed effect*, and $-\frac{1}{2}y_i^2$ is a quadratic effort cost. To be more precise, a_i , the observable individual heterogeneity in productive ability, is assumed to be deterministic, perfectly observable by all individuals in the network, and corresponds to the observable characteristics of individual i (e.g., age, sex, participation in on-the-job-training, etc.) and to the observable average characteristics of individual i 's immediate co-workers. It can thus be written as:

$$a_i = \sum_{m=1}^M \beta_{1m}x_i^m + \frac{1}{g_i} \sum_{m=1}^M \sum_{j=1}^n \beta_{2m}g_{ij}x_j^m \quad (2)$$

where x_i^m belongs to a set of M variables accounting for observable differences in individual characteristics of individual i . β_{1m} and β_{2m} are parameters and $g_i = \sum_{j=1}^n g_{ij}$ constitute the total number of immediate co-workers of individual i .

3.2 Nash Equilibrium

We now characterize the Nash equilibrium of the game where agents choose their effort level $y_i \geq 0$ simultaneously. In equilibrium, each agent maximizes her utility (1). We obtain the following best-reply function for each $i = 1, \dots, n$:

$$y_i = \phi_1 \sum_{j=1}^n g_{ij}y_j + \phi_2 \sum_{j=1}^n g_{ij}^*y_j + \alpha_i \quad (3)$$

where $\phi_1 = \lambda_1/(1 + \lambda_2)$, $\phi_2 = \lambda_2/(1 + \lambda_2)$, and $\alpha_i = (a_i + \eta + \epsilon_i)/(1 + \lambda_2)$. As $\lambda_1 \geq 0$ and $\lambda_2 \geq 0$, we have $\phi_1 \geq 0$ and $0 \leq \phi_2 < 1$. The coefficient ϕ_1 is called the *local-aggregate* endogenous peer effect. As $\phi_1 \geq 0$, this coefficient reflects *strategic complementarity* in efforts. The coefficient ϕ_2 is called the *local-average* endogenous peer effect, which captures the *taste for conformity*. Note that, $\phi_1/\phi_2 = \lambda_1/\lambda_2$. That is, the relative magnitude of ϕ_1

and ϕ_2 is the same as that of the social-multiplier coefficient λ_1 and the social-conformity coefficient λ_2 .

3.3 Key players

The concept of the key player in economics has been introduced by Ballester, Calvó-Armengol and Zenou (2006) and was initially defined for criminal activities. The key player is the agent that should be targeted by the planner so that, once removed, she will generate the highest level of reduction in total activity. It has been tested empirically and applied to other activities than crime, such as financial networks, R&D networks, wars, etc. (see Zenou, 2016, for an overview of this literature). Here, the key player will be the worker that the firm would most like to retain because, if removed, total productivity will be reduced the most. In some sense, the key player(s) is (are) the critical worker(s) in a company.

Formally, a *key player* is the agent whose removal from the network leads to the largest reduction in the aggregate effort level in a network. Let $\mathbf{M}(g, \phi_1, \phi_2) = (\mathbf{I} - \phi_1 \mathbf{G} - \phi_2 \mathbf{G}^*)^{-1}$, with its (i, j) -th entry denoted by $m_{ij}(g, \phi_1, \phi_2)$. Let

$$\mathbf{b}(g, \phi_1, \phi_2, \boldsymbol{\alpha}) = \mathbf{M}(g, \phi_1, \phi_2) \boldsymbol{\alpha}$$

with its i -th entry denoted by $b_i(g, \phi_1, \phi_2, \boldsymbol{\alpha}) = \sum_{j=1}^n m_{ij}(g, \phi_1, \phi_2) \alpha_j$. Let $B(g, \phi_1, \phi_2, \boldsymbol{\alpha}) = \sum_{i=1}^n b_i(g, \phi_1, \phi_2, \boldsymbol{\alpha}) = \mathbf{1}'_n \mathbf{M}(g, \phi_1, \phi_2) \boldsymbol{\alpha}$ denote the aggregate effort level in network g , where $\mathbf{1}_n$ is an $n \times 1$ vector of ones. Let $g^{[-i]}$ denote the network with agent i removed. Let $\mathbf{G}^{[-i]}$ and $\boldsymbol{\alpha}^{[-i]}$ denote the adjacency matrix and vector of covariates corresponding to the remaining agents in network $g^{[-i]}$. Then, the key player i^* in network g is given by $i^* = \arg \max_i d_i(g, \phi_1, \phi_2, \boldsymbol{\alpha})$, where

$$d_i(g, \phi_1, \phi_2, \boldsymbol{\alpha}) = B(g, \phi_1, \phi_2, \boldsymbol{\alpha}) - B(g^{[-i]}, \phi_1, \phi_2, \boldsymbol{\alpha}^{[-i]}). \quad (4)$$

3.3.1 The key player in the local-aggregate network game

Ballester, Calvó-Armengol and Zenou (2006, 2010) and Ballester and Zenou (2014) have studied the key-player policy for the *local-aggregate* network game when $\lambda_2 = 0$ in (1). Observe that, in this case, $\phi_2 = 0$ and $b_i(g, \phi_1, 0, \boldsymbol{\alpha})$ is the well-known *Katz-Bonacich centrality* of node i (Katz, 1953; Bonacich, 1987). Let $\boldsymbol{\alpha}^{[i]}$ denote the vector of covariates calculated based on the network consisting $g^{[-i]}$ and the isolated i . It follows from Ballester and Zenou (2014) that the key player can be determined by the *generalized intercentrality*.

Proposition 1 For network g , let the generalized intercentrality of node i be denoted by

$$d_i(g, \phi_1, 0, \boldsymbol{\alpha}) = \underbrace{\mathbf{1}'_n \mathbf{M}(g, \phi_1, 0)(\boldsymbol{\alpha} - \boldsymbol{\alpha}^{[i]})}_{\text{contextual variable change effect}} + \underbrace{\frac{b_i(g, \phi_1, 0, \boldsymbol{\alpha}^{[i]}) \sum_{j=1}^n m_{ji}(g, \phi_1, 0)}{m_{ii}(g, \phi_1, 0)}}_{\text{network structure change effect}}. \quad (5)$$

Then, agent i^* is the key player of the local-aggregate network game if and only if i^* has the highest generalized intercentrality in network g .

The generalized intercentrality (5) highlights the fact that when an agent is removed from a network, two effects are at work. The first one is the *contextual variable change effect*, which is due to the change in α after the removal of an agent. The second effect is the *network structure change effect*, which captures the change in \mathbf{G} when an agent is removed. More generally, the generalized intercentrality measure accounts both for one’s exposure to the rest of the group and for one’s contribution to every other exposure.

3.3.2 The key player in network games with the local-average peer effect

To the best of our knowledge, nobody has studied the key-player policy for the local-average model. Liu, Patacchini and Zenou (2014) discuss this issue by providing some examples. In general, when the agents are *ex ante* homogeneous in terms of observable characteristics, which agent to remove from the network does not matter in terms of the aggregate effort level reduction, unless the agent holds a very special position in the network such that removing this agent generates isolated nodes in the network.

If the agents have different values of \mathbf{x}_i , the key-player problem for the local-average network game and the general network game with utility (1) does not have an analytical solution. Yet we can still determine the key player numerically using its definition given by (4) if we can estimate the unknown parameters in the best-response function (3).

4 The Call Center Data and Institutional Setting

We use data from an in-house call center of a multi-national mobile network operator to estimate the best reply function given by equation (3). We construct two different samples of data from this call center. We label the first data set our “full sample”, which are taken from observational data on all call center workers from week 1/2008 until week 10/2010. We label the second data set our “experimental sample”. These data are a subset of our full sample. They include information on workers that were involved in a randomized on-the-job training experiment and also information on the peers of these workers, even those

not directly involved in the training experiment. We describe each of these data sets below. We also provide descriptive statistics and important information concerning the operational setting.

The call center provides services for current and prospective customers and is divided into 5 departments, which are segmented by customer group. In this study, we use data from the largest department, which handles calls from private customers with fixed mobile phone contracts. Customers contact customer services in case of problems, complaints or questions. Our data contain information on 439 call center workers. Because we are missing important information on gender, age and tenure for 14 agents, we are dropping them from our sample. This leaves us with 425 workers and 14,079 worker-week observations.¹¹ On average, 124 agents work in this department each week.

Call center agents working in this department answer customer calls and make notes in their customer database for documentation. Agents are recruited by an external recruiting company through which they are employed in the beginning. After hiring, agents receive an initial training of three weeks. Throughout their career, agents receive further training, e.g. for specific campaigns or computer skills.

All agents are placed in teams which are led by a team leader. The main purpose of grouping agents into teams is that it facilitates monitoring, evaluation, and coaching by the team leader. Teams are not specialized for specific types of calls, or specific customer groups. There are also no team-based incentives. We observe an average of 10 teams per week. Average team size is 12 workers.

Our panel data include weekly information on agents' individual performance. Workers' performance as well as other indicators are continuously, and automatically measured by the IT system of the call center. Throughout the sample period used in this paper, average handling time, i.e. the time an agent needs on average to handle a customer call, is used as the main key performance indicator to evaluate agents. The management's aim is to reduce costs by reducing average handling time aht_{it} without a loss in quality. We define worker performance as $y_{it} = \frac{100}{aht_{it}}$. A decrease in average handling time can thus be interpreted as an increase in worker performance. Team coaches receive weekly scorecards for each worker on this and other key performance indicators.

¹¹We are also missing information on the age of an additional 29 workers. But instead of dropping them, we assign the mean age to them and create a dummy variable indicating that we are missing information about their age.

4.1 The Full Sample

To study how networks affect performance in the workplace, we link these performance data to data on the exact time agents are present at their workplace. This information is gathered from the turnstiles where agents need to log in when entering and log out when leaving the call centre. Agents also need to log in and out when they have breaks. Since agents who belong to the same team sit next to each other, two agents who are present at the same time are exposed to each other. Thus, we will use team membership and overlap in hours worked to define links between co-workers and co-worker networks (more on this below).

The panel data consisting of performance and network information is complemented with information on agents' gender, age, tenure, the number of hours they work each week, which days of the week they work, and when they work (morning, midday or evening). Information on the overall work load (i.e. the volume of incoming calls) during a specific week is also available. Importantly, we also know if and when a worker has received one week of on-the-job training in a newly introduced training program (more on this below).

Descriptive statistics for the full sample are shown in Table 1. The full sample consists of 425 different workers and 14,079 worker \times week observations. Two-thirds of the workers are females. The average age of a worker is 29 years old and most work part-time (around 22 hours per week). Most workers work Monday through Friday during the middle of the day. Some work on Saturdays, but very few work on Sundays. Worker performance y_{it} varies quite substantially. One standard deviation is equal to 26% of the mean in average worker productivity (see the right hand panel of Table 1) and 32% of the mean productivity in our worker \times week panel (see the left hand panel of Table 1).

[Insert Table 1 here]

4.2 The Experimental Sample

In 2008, the firm decided to introduce a new on-the-job training program. To assess the effectiveness of this new program, it was first introduced in the form of a randomized treatment and control experiment (De Grip and Sauermann, 2012). The experimental sample is a subset of the workers in our full sample.

In week 50/2008, agents were selected for participation in the training program. The training program was focused on more tenured agents. This was done for two reasons. First, the company's management believed that the new training program would be more suitable for workers with some experience on the job. Second, management wanted to reduce the risk of losing investments in on-the-job training through turnover.

The training program took place over the course of 27 weeks, starting in week 10/2009. The training took place in an in-house training center and consisted of 10 half-day sessions that were held from Monday to Friday. Half of these sessions contained group discussions led by the training coach and the team leader. These discussions were about which skills the agents were missing when executing their task, how these could be improved, and how agents could help each other on the work floor. Agents were also trained in conversational techniques designed to decrease average handling time. During the other half of the sessions, agents handled incoming customer calls that were routed to the training center. Training coaches and team leaders assisted these calls and gave feedback.

De Grip and Sauermann (2012) evaluate the causal effect of this training program using data on the 74 agents who were chosen for the experiment and then randomly assigned to treatment and control groups. They found that participation in the training programme produced a 8.8% increase in performance. Since the training of some teams was split into two sessions (due to capacity constraints of the training facility), they were also able to estimate a positive externality from the early trained ($N = 13$) on the late trained ($N = 21$): An increase of 10 percentage points in the share of treated peers improved a worker's performance by 0.51%.

One important aspect of this experiment is that workers were randomized by teams and not as individuals, which is an important fact that we will make use of when constructing our own experimental sample.

We construct our experimental data set using data from week 45/2008 to week 24/2009 (cf. De Grip and Sauermann, 2012). This includes a *pre-experimental period*, weeks 45/2008-9/2009, the *training period*, weeks 10-14/2009, and a 10 week *follow up period*, weeks 15-24/2009. Once the experiment was concluded, the control group was also trained and following this, the program was expanded to include other workers.

In our first experiment, we use data from the full sample (described above) to test for network effects by estimating workers' best reply function given by equation (3). In our second experiment, we test for network effects that arise from a very specific increase to workers' productivity, namely the causal effect on performance of this experimental training program. In essence, if there are spillover effects on non-trained workers or multiplier effects among trained workers (or both), then these phenomena will give rise to a positive network effect from on-the-job training. Measuring such effects is necessary to get an accurate picture of the costs and benefits of on-the-job training.

To explore this hypothesis, we use our new experimental sample, which is comprised of 4 groups of workers: (1) the treated ($N = 29$), (2) the control ($N = 41$), (3) the untrained

peers of the treated ($N = 24$), and (4) the untrained peers of the controls ($N = 43$). In the presence of network effects, the post-experiment performance of group 3 will be higher than that of group 4.

Descriptive statistics for the experimental sample are shown in Table 2. Worker characteristics balance quite well across the treated and controls and across the peers of the treated and the peers of the controls. As mentioned above, the training program was tailored for workers with longer tenure, so the peers of the treated and controls are younger, have lower tenure, and lower average productivity than those who were selected to participate in the experiment.¹²

[Insert Table 2 here]

5 Defining Co-Worker Networks

Let us now describe how we define co-worker networks. As we saw above, most call agents are part-timers who work on average 22 hours per week. Call agents are organized into approximately 10 teams at any given moment. There are 20 different teams in our full sample. All teams work on the same floor of the building. The physical workspace is organized into work islands, with up to eight agents of a team sitting next to each other. There are two levels of co-worker interactions. First, each worker is assigned to a team. Then, each week, individuals work in shifts and interact with different persons within the team that they are allocated to.

For each week, we have information about the exact time when entering and leaving the call centre. This information is used to identify networks (which can, and will, be weighted by joint working hours between worker i and j). As a result, one can reconstruct the whole geometric structure of a co-worker network, which is summarized by the adjacency matrix \mathbf{G} . We define each network component r (henceforth network) such that all individuals belonging to a network are path-connected. We define time periods t as weeks to make the problem tractable.

¹²Note that De Grip and Sauermann (2012) only analyze the data for the 74 workers originally chosen to participate in the training experiment. They do not analyze the 67 peers of these 74 workers, nor do they make use of the full data sample on all 425 workers that we use. Note also that our experimental design requires strict adherence to the original randomization. We, therefore, drop four of the individuals used in De Grip and Sauermann’s (2012) original experiment.

Two employees are defined as co-workers if they both come from the same team τ and their work hours during week t overlap. In total, we have 114 weeks of data for 20 different teams. Over time, new teams are created and old teams are dissolved so that we observe roughly 10 teams each week. In total we observe 1188 team by week networks. To keep the notation simple, we label networks by r , leaving the team and time aspect $r(\tau, t)$ implicit.

As in Section 3, we first define an unweighted adjacency matrix \mathbf{G}_r , where each cell $g_{ij,r} \in \{0, 1\}$ keeps track of whether team members i and j have worked together during week t or not. We can also define matrix \mathbf{G}_r^* , which is the row-normalized matrix of \mathbf{G}_r where each cell $g_{ij,r}^* = g_{ij,r}/g_{i,r} \in [0, 1]$, where $g_{i,r} = \sum_j g_{ij,r}$ is the total number of team members individual i has worked with during week t . We also define a matrix \mathbf{H}_r , in which each cell $h_{ij,r} \geq 0$ keeps track of the number of hours team members i and j have worked together during week t . The weighted adjacency matrix \mathbf{H}_r^w is such that each cell $h_{ij,r}^w = h_{ij,r}/h_{i,r}$, where $h_{i,r} = \max_j[h_{ij,r}]$. This normalizes the weights so that the weight on the link between worker i and the co-worker j that she works the most with is equal to one.

To illustrate this, consider the following network g_r as shown in Figure 1. There are three agents i in team τ . Agent 1 holds a central position whereas agents 2 and 3 are peripherals.

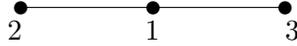


Figure 1: A network of 3 agents

The unweighted adjacency matrix \mathbf{G}_r for this network is:

$$\mathbf{G}_r = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}$$

This means that, during the week t for which the network is observed, agents 1 and 2 as well as agents 1 and 3 have worked together while agents 2 and 3 have not. We can also define \mathbf{G}_r^* , which is the *row-normalized matrix* of \mathbf{G}_r and defined as:

$$\mathbf{G}_r^* = \begin{pmatrix} 0 & 1/2 & 1/2 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}$$

so that $g_{ij,r}^* = g_{ij,r}/g_{i,r} \in [0, 1]$, where $g_{i,r} = \sum_j g_{ij,r}$ is the total number of persons individual i has worked with during week t .

Imagine that, during week t , agents 1 and 2 have worked 6 hours together while agents 1 and 3 have worked 10 hours together. We have:

$$\mathbf{H}_r = \begin{pmatrix} 0 & 6 & 10 \\ 6 & 0 & 0 \\ 10 & 0 & 0 \end{pmatrix}$$

Each link is given a weight by dividing through by the maximum value in each row. This means that we wait each link relative to the strongest link and the strongest link is normalized to 1. The weighted adjacency matrix \mathbf{H}_r^w is given by:

$$\mathbf{H}_r^w = \begin{pmatrix} 0 & 6/10 & 10/10 \\ 6/6 & 0 & 0 \\ 10/10 & 0 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0.6 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}$$

The row-normalized and weighted adjacency matrix \mathbf{H}_r^{w*} is given by:

$$\mathbf{H}_r^{w*} = \begin{pmatrix} 0 & 0.6/1.6 & 1/1.6 \\ 1/1 & 0 & 0 \\ 1/1 & 0 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0.375 & 0.625 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}$$

The weighted matrices reflect the fact that worker 1 works 2/3 more hours with worker 3 than she does with worker 2. Henceforth, we drop the superscript w and refer to our weighted adjacency matrix as \mathbf{H}_r and our row-normalized weighted adjacency matrix as \mathbf{H}_r^* .¹³

Let $\bar{r}(\tau, t)$ be the total number of networks in the sample, n_r the number of individuals in the r th network, and $n = \sum_{r=1}^{\bar{r}} n_r$ the total number of sample observations. In our full data set, $n = 14,079$ and $\bar{r}(\tau, t) = 1188$. The minimum network size is one. The average network size is 12. The median size is 13 and the maximum is 36.

In Figure 2, we plot the graph of one such (randomly chosen) co-worker network. There are 17 workers who work together as a team for a given week. Each node represents a worker. The size of the node reflects how many co-workers from the same team a worker has worked on the same shift with during that particular week. The thickness of the lines represents the number of hours each pair of workers worked side by side during this particular week. This network is not complete. All workers are not directly connected to each other. Note also that there are large differences in the amount of time each pair of workers is exposed to each other.

¹³Observe that all the results obtained in Section 3 hold true if we use the matrices \mathbf{H}^w and \mathbf{H}^{w*} in (3) instead of \mathbf{G} and \mathbf{G}^* .

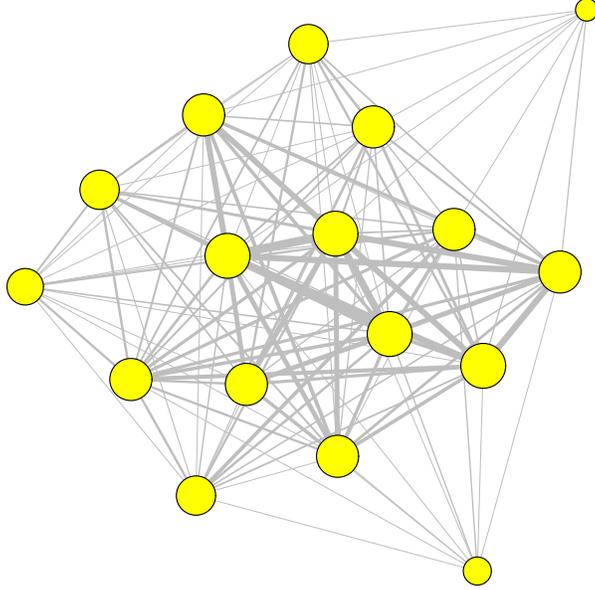


Figure 2: A real-world co-worker network.

6 Experiment #1: Network Effects on Worker Productivity

In this experiment, we use our full panel data set (shown in Table 1 and discussed above) to estimate the best reply function of workers given by equation (3) where we replace the matrices \mathbf{G} and \mathbf{G}^* by \mathbf{H} and \mathbf{H}^* and add the subscript r to all variables to indicate to which network each individual belongs. We want to see if $y_{i,r}$, the productivity of individual i belonging to network r (measured by the average time needed to handle inbound customer calls), is positively influenced by the productivity of the team members who work the same shift as individual i during the week weighted by the number of hours worked together during that week.

Similar to Bayer, Hjalmarsson and Pozen (2009) and Mas and Moretti (2009), the identification strategy in our first experiment relies on the exogenous exposure matrices \mathbf{H}_r and \mathbf{H}_r^* . As we will demonstrate in Section 6.2 (below), \mathbf{H}_r and \mathbf{H}_r^* can be treated as exogenous after first conditioning on individual, team and week fixed effects. Each week, an individual

is exposed to a different dose of both aggregate- and average peer productivity, since each week she faces a somewhat different group of co-workers who, in turn, vary in their productive capacity. It is this exogenous variation in co-worker productivity that we use to identify causal network effects.

We further strengthen this identification strategy and increase precision by including observable characteristics of individuals and the average characteristics of their co-workers. Importantly, we also include individual fixed effects, ϵ_i , team fixed effects, τ , and week fixed effects, t .

The econometric model corresponding to the best-reply function (3) of agent i belonging to network $r(\tau, t)$ can be written as:¹⁴

$$y_{i,r} = \phi_1 \sum_{j=1}^n h_{ij,r} y_{j,r} + \phi_2 \sum_{j=1}^n h_{ij,r}^* y_{j,r} + \sum_{m=1}^M \beta_{1m} x_{i,r}^m + \sum_{m=1}^M \sum_{j=1}^n \beta_{2m} h_{ij,r}^* x_{j,r}^m + \epsilon_i + \tau + t + \varepsilon_{i,r}. \quad (6)$$

Recall that ϕ_1 and ϕ_2 represent the local aggregate and the local average network effects (respectively), which are the main objects of interest in this study; $\varepsilon_{i,r}$ represents i.i.d. innovations with zero mean and variance σ^2 for all i and r . The characteristics $x_{i,r}^m$ and $x_{j,r}^m$ are gender, age, tenure, total work hours during the week, day(s) of the week worked and the time of the day worked (morning, midday, evening). Inference will be made using standard errors that are clustered on individual workers.

6.1 Threats to Identification

It is well-known that when estimating peer effects using linear-in-means model endogenous peer effects (ϕ_1 and ϕ_2) and contextual effects (β_{2m}) cannot always be separately identified due to the *reflection problem* (Manski 1993). When individuals are influenced by the members of their own group, but not by individuals outside their group, there arises a simultaneity in the behavior of individuals within the group that introduces a perfect collinearity between the endogenous peer effect and the contextual effects. In this very special case, one cannot disentangle these two effects.

Using the terminology of social networks, the reflection problem arises when networks are complete. That is, when all agents are connected to (and influenced by) all other agents in the network. However, most networks (such as those studied in this paper) are not complete; everyone is not connected to everyone else. Bramoullé, Djebbari and Fortin (2009), Lee, Liu

¹⁴where, as stated above, we replace the matrices \mathbf{G} and \mathbf{G}^* by \mathbf{H} and \mathbf{H}^* .

and Lin (2010), Liu and Lee (2010), Liu et al. (2012) and others have shown us how the architecture of social networks can be used to identify endogenous peer effects.¹⁵ Loosely speaking, endogenous peer effects and exogenous contextual effects are identified if at least two individuals in the same network have different links (Bramoullé, Djebbari and Fortin, 2009). This condition is generally satisfied in any real-world network. In practice, a priori knowledge of \mathbf{H}_r and \mathbf{H}_r^* provides us with a set of restrictions on the coefficients in the reduced form equation that are used to identify the structural model.

While our network approach does allow us to separately identify endogenous effects and contextual effects, it does not necessarily identify the causal effect of peers' influence on individual behavior. In our context, we face two sources of potential bias arising from *correlated effects* and *endogenous network formation* (non-random sorting).

Individuals within the same network who share the same environment and face the same set of incentives and/or shocks are likely to behave in a similar manner. We control for these types of correlated effects by adding team and week fixed effects. Team members share the same physical environment and answer to the same team leader who may have her own personal management style, may be more or less experienced, etc. Shift members may share typical workloads and/or shift-specific shocks. We, therefore, control for week fixed effects and also for the day(s) of the week worked and the time of day worked (morning, midday, evening). These controls should deal with correlated effects.

Unlike the networks in most applications, our networks are not formed by individuals who self-select into them. Workers do choose to work for the firm and they also state the shifts that they would be willing and able to work. For example, homemakers may want to work in the middle of the day, while students may only be available evenings and weekends. But it is the firm that places these workers into teams and sets the weekly work schedule.

The firm, however, clearly has the power to place like with like if they so desire. The firm could choose to place all homemakers into one team and all students into another. They could also choose to take workers' requests to work together into account when forming teams if they thought it to be in the best interests of the firm. If teams are formed through a process of assortative matching, then we would find ourselves in a situation with endogenous network formation. This would result in a positive bias to our estimated peer effect, since similar people would be placed into the same groups and are likely to have positively correlated outcomes even in the absence of true peer effects.

¹⁵See Blume et al. (2011 and 2015) for recent overviews of the literature on the identification of social interactions and for a set of original and important contributions on the topic. See also our discussion in Section 2.3.

We have four explicit ways of dealing with the potential issue of endogenous team and shift formation (i.e., network formation). First, we control for the observable characteristics of co-workers, $\sum_{m=1}^M \sum_{j=1}^n \beta_{2m} h_{ij,r}^* x_{j,r}^m$. Second, we control for both team and week fixed effects. Third, we control for both observable, $\sum_{m=1}^M \beta_{1m} x_{i,r}^m$, and unobservable, ϵ_i , characteristics of each individual worker using an important set of observables controls and individual fixed effects. Fourth, we rely on our exogenous exposure matrices \mathbf{H} and \mathbf{H}^* to provide exogenous variation in co-worker productivity. We are not using the stable part of the co-worker network (that is potentially endogenous) to identify network effects. We are identifying network effects off of changes in the dose of co-worker productivity that each individual worker faces from week to week due to non-systematic changes in the make-up her network of co-workers.

There is a considerable amount of week-to-week variation in the dose of co-worker productivity that each worker is exposed to even after controlling for team-, week, and individual fixed effects. The standard deviation of the residualized variation in $\mathbf{H}y$ is 0.86, while the mean of $\mathbf{H}y$ is 2.14. The standard deviation of the residualized variation in \mathbf{H}^*y is 0.07. The mean of \mathbf{H}^*y is 0.34.

6.2 Diagnostic Test of Identifying Assumption

In our setting, the main threat to the identification of a causal network effect on worker productivity is the potential for the firm to construct co-worker networks (i.e., teams and shifts) in such a way that the innate productivities of team- and shift-mates are correlated and that firms use information that is unobservable to us when forming teams and shifts. Our main identifying assumption is that the within-worker, -team and -week variation in co-worker networks is essentially random, i.e. that \mathbf{H}_r and \mathbf{H}_r^* are conditionally exogenous after controlling for worker, week and team fixed effects. Our argument is that this variation is as good as random since it is based on variation in the overlap of hours worked (shifts) within teams of co-workers, due to non-systematic events (e.g., idiosyncratic changes in availability, own illness, sick children, holidays, school schedule changes – affecting both mothers with school children and college students, exam periods, school holidays, etc.) and due to new hires and quits (i.e., co-worker turnover). If this assumption holds, then we should see no correlation between this variation and the average observable characteristics of a worker’s peers. Nor should we be able to detect any correlation between a worker’s own observable characteristics and her co-workers’ characteristics after controlling for worker, team and week fixed effects.

But before we test this assumption, we would like to examine whether or not we actually see evidence of non-random link formation between co-workers in networks. To do this, we estimate a logistic model of link formation using the variables age, gender and tenure as covariates.¹⁶ What we find is that men only have a 0.05 higher odds of linking with another man (as opposed to linking with a woman), and people aged +/- 5 years apart only have a 0.04 higher odds of being linked with each other.^{17,18} In contrast to these very small amounts of non-randomness, we see that people with similar tenure (+/- 12 weeks) have a 2.6 higher odds of being linked with each other.

What we see in the data is that firms tend to hire more than one new person at a time. These new people tend to be placed in the same team (or teams) for training and then stay closely linked to each other for many months to come. Thus, links are significantly non-random in tenure. Since productivity is strongly increasing in tenure (we show this below; see also De Grip, Sauermann, and Sieben, 2011), those who are linked together will tend to have correlated productivities generated by this correlation in tenure. It is exactly this type of threat to identification that we need to be wary of and motivates our use of control variables together with the use of individual, team and week fixed effects.

Thus, a test of our main identifying assumption must demonstrated that variation in \mathbf{H}_r and \mathbf{H}_r^* is unrelated with average co-worker characteristics (age, gender and, in particular, tenure) after conditioning on our set of fixed effects. We run several versions of this basic test.

In our first test, we use age, tenure and gender along with individual, team and week fixed effects to predict work productivity, \hat{y} . We then regress this measure (or index) of the productive characteristics of workers on to our two measures of endogenous network effects: the local aggregate network effect $\sum_{j=1}^n h_{ij,r} y_{j,r}$ and the local average network effect $\sum_{j=1}^n h_{ij,r}^* y_{j,r}$. In Panel A of Table 3, we see that our index of worker characteristics, \hat{y} , is correlated with our two measures of endogenous network effects. However, in Panel B, we see that these correlations completely disappear when we include individual, team and week fixed effects. This result speaks in favor of the conditional exogeneity of \mathbf{H}_r and \mathbf{H}_r^* and of our main identifying assumption.

[Insert Table 3 here]

¹⁶We estimate this model using the *ergm* package in R (Hunter et al. 2008). Our model is one with dyadic independence. Hence, we avoid the well-known problem of model degeneracy.

¹⁷These results are available upon request.

¹⁸Note that 0.05 and 0.04 are extremely small values of assortative matching relative to what is typically seen in the literature (e.g. Currarini, Jackson and Pin, 2010).

In Column (2) of Table 3, we relate our index of worker characteristics, \widehat{y} , to the average observable characteristics of her co-workers, $\sum_{m=1}^M \sum_{j=1}^n \beta_{2m} h_{ij,r}^* x_{j,r}^m$. Once again, we see that the correlations that appear to exist (see Panel A) disappear once we include individual, team and week fixed effects (see Panel B).

In a similar fashion, we can test whether or not the number of predicted (weighted) links an individual has, $\sum_{j=1}^n \widehat{h}_{ij,r} \mathbf{1}_{j,r}$, is related to the average observable characteristics of her co-workers. We also test to see if the predicted values of our two measures of endogenous network effects, $\sum_{j=1}^n \widehat{h}_{ij,r} y_{j,r}$ and $\sum_{j=1}^n \widehat{h}_{ij,r}^* y_{j,r}$, are correlated with the average observable characteristics of a worker's co-workers. In Panel B of Table 3, we clearly see that the variation in co-worker productivity that is used to identify a causal network effect is uncorrelated with the average characteristics of co-workers within networks after conditioning on individual, team and week fixed effects. Once again, these results all speak in favor of our main identifying assumption.

6.3 Results

Estimation results of equation (6) are reported in Table 4. In Column (1), we report the raw associations between our two different network effects and worker productivity. The local average network effect has a strong positive association with worker productivity equal to 0.69 (0.040), while the local aggregate network effect has a negative association -0.01 (0.003). After including individual, team and week fixed effects in Column (2), the large positive association between the average network effect and worker productivity is reduced to 0.20 (0.057), while the local aggregate effect changes only slightly. In Column (3), we add controls for individual characteristics (including workday dummies) and average co-worker characteristics. These additional controls reduce our estimate of the local aggregate network effect and render it insignificant. Our estimate of the local average effect is only slightly reduced.

We conclude that the local average network model fits the data better than the local aggregate model does. This is our first important empirical finding. It implies that a worker's current productivity is affected by her co-workers' current productivity through her desire to conform to the local work norm and not through strategic complementarities. In Column (4) of Table 4, we present our preferred baseline model of worker productivity, which includes the local average network effect only.¹⁹

¹⁹When we estimate a similar model including the local aggregate effect only, the aggregate effect is a precisely estimated zero equal to 0.0007 (0.0017).

[Insert Table 4 here]

Our point estimate of 0.17 (0.060) implies that a 10% increase in average co-worker productivity produces a 1.7% increase in a worker’s own productivity. This is an economically meaningful effect and is quite similar to the effects reported in previous studies by Falk and Ichino (2006) and Mas and Moretti (2009). These earlier studies report effects of 1.4% and 1.5%, respectively. Importantly, both of these studies also present evidence that these productivity spillovers are driven by social norms.

The main threat to the identification of a casual effect in our case is not the reflection problem. Our network approach eliminates this worry. Nor is it endogenous network formation, since we are identifying our effect off of the non-stable part of the network. The remaining potential threat to identification is that of correlated shocks. More specifically, one might worry that each team faces a new, team-specific shock each week and that such team- and week-specific shocks occur on a regular basis and affect worker productivity in a meaningful way. Such correlated shocks could lead to positive associations in co-worker productivity and, thus, bias our estimates upwards.

To address this potential problem, we use $\mathbf{H}^*\mathbf{H}^*X$ as an instrument for \mathbf{H}^*y in a two-stage least squares (2SLS) regression. This has been shown to be a valid instrument given that \mathbf{H}^* is conditionally exogenous (Bramoullé, Djebbari and B. Fortin, 2009). This instrument entails using the characteristics of a reference worker’s co-workers’ co-workers. These co-workers are two steps away from the reference worker, i.e. the reference worker is not directly connected to them. This provides us with a valid exclusion restriction. We include age, gender and tenure in X . These characteristics are fixed at any point in time and cannot be correlated with potential team-week shocks. In Column (5) of Table 4, we see that our 2SLS point estimate is not significantly lower than our baseline estimate. We, therefore, conclude that our results are not likely driven by correlated shocks.

6.4 Heterogeneous Effects by Tenure

In Table 5, we examine the importance of network effects for high- versus low tenure workers. The average network effect appears to affect both high- and low tenure workers. But it is clearly more important for low tenure workers than for high tenure workers. This seems

reasonable if it is the more tenured workers who set the work norm, while newer workers strive to live up to this norm.²⁰

[Insert Table 5 here]

7 Experiment #2: Network Effects on Worker Productivity from On-the-Job Training

In our second regression experiment, we make use of the random variation in on-the-job training that is available in our experimental data set to identify causal network (spillover) effects from such training. Recall that our experimental data set is a subset of our full data set that includes teams of workers that participated in a randomized training experiment. See Table 2 and our discussion above.

We begin by replicating the main finding reported in Column (3) of Table 4 in De Grip and Sauermann (2012). We use the subset of our data (70 workers) who were chosen to take part in an on-the-job training experiment and who complied with their randomly assigned treatment status. Following De Grip and Sauermann (2012), we estimate the following fixed effects regression:

$$\log y_{it} = \alpha_i + \delta ojt_{it} + \beta_1 working_hours_{it} + \beta_2 share_peak_hours_{it} + \beta_3 trend_t + u_{it} \quad (7)$$

As before, y_{it} is agent’s performance in week t , which we now take the log of. In this equation, ojt_{it} is an indicator variable for on-the-job training. Treatment status is randomized by team so that δ represents the average treatment effect. Their baseline regression also includes individual fixed effects, α_i , the number of hours worked during a particular week, $working_hours_{it}$, whether or not a worker’s shift coincides with the peak workload hours, $share_peak_hours_{it}$, and a linear time trend, $trend_t$.

In Column (1) of our Table 6, we see that treated workers are 7.8% more productive during the 10 weeks following their training (relative to the control group). Using a sample of 74 workers, De Grip and Sauermann (2012) report an 8.8% increase in productivity. Thus, on-the-job training does, in fact, have a direct effect on a worker’s own productivity and our basic estimate is in line with previous estimates reported in De Grip and Sauermann (2012).

²⁰We have also run similar regressions that include both the local average network effect and the local aggregate network effect, and also regressions that include only the aggregate effect. The local aggregate network effect is never significant.

[Insert Table 6 here]

Does this increase in worker productivity due to on-the-job training spill over to non-treated co-workers? To study this question, we look at our sample of 67 workers who did *not* take part in the original training experiment but who were teammates of those who did take part. Since the original randomization was done at the team level, we label the teammates of those in the treatment group from the original experiment our “treated” peers, while the teammates of the control group are our “control” peers (see Table 2). In Column (2) of Table 6, we see that the productivity of treated peers is 8.5% higher than that of the control peers during the 10 weeks following the training experiment.

But how is it possible for the indirect peer effect to be as large as the direct treatment effect? What we see in Column (3) of Table 6 is that the indirect peer effect is concentrated among co-workers with low tenure. Recall that the actual training experiment was conducted using a sample of workers with relatively high tenure. Low tenure workers were not included in the original experiment. A likely explanation is that high tenure workers take their new training back to their teams and either consciously or otherwise pass their newly gained skills onto their low tenure teammates. Since low tenure workers are on the steep part of their learning curve, small doses of indirect training via network effects can affect their productivity substantially. Thus, indirect treatment effects on low tenure workers (who have much to learn) can increase the productivity of those workers by as much as the direct training effect on high tenure workers who are on the relatively flat part of their learning curve.

This explanation of the observed spillover effects relies on the idea of learning-on-the-job through strategic complementarities. In our model framework, strategic complementarities are represented by the local aggregate network effect. We can test the hypothesis that on-the-job training leads to productivity spillovers via strategic complementarities by estimating the best reply function from our model. The estimation equation is modified to include a dummy variable in $x_{j,r}^m$ equal to one if worker j has received training and zero if not. We also replace the terms $\sum_{j=1}^n h_{ij,r} y_{j,r}$ and $\sum_{j=1}^n h_{ij,r}^* y_{j,r}$ with $\sum_{j=1}^n h_{ij,r} ojt_{j,r}$ and $\sum_{j=1}^n h_{ij,r}^* ojt_{j,r}$, respectively.

$$y_{i,r} = \phi_3 \sum_{j=1}^n h_{ij,r} ojt_{j,r} + \phi_4 \sum_{j=1}^n h_{ij,r}^* ojt_{j,r} + \sum_{m=1}^M \beta_{1m} x_{i,r}^m + \sum_{m=1}^M \sum_{j=1}^n \beta_{2m} h_{ij,r}^* x_{j,r}^m + \epsilon_i + \tau + t + \varepsilon_{i,r}. \quad (8)$$

We could have instead used the number of trained and the share trained **workers?** as instruments for aggregate- and average productivity in a 2SLS framework. But studying the reduced form directly makes the results easier to interpret. Note that the sample size is reduced to 264 workers and 5,572 worker \times week observations, since there was no training during the first period of our sample. Training begins with the above mentioned field experiment and then continued on a regular basis after the field experiment was concluded.

Our estimates of ϕ_3 and ϕ_4 are shown in Table 7. The aggregate number of trained workers present in a worker’s network has a significantly positive affect on own productivity. Once again, we see that there are positive spillover effects from on-the-job training. These effects work through the existence of strategic complementarities. We see no evidence that the average number of trained workers in a network has any significant effect on worker productivity.

Our estimate of ϕ_3 is equal to 0.003 (0.0011). This estimate tells us that if we increase the number of trained workers in a network by one worker, then this increase leads to a 0.7% increase in worker productivity.²¹ The standard deviation of $\sum_{j=1}^n h_{ij,r} ojt_{j,r}$ in the estimation sample is 2.22. If we increase the number of trained workers by one standard deviation, then this increases worker productivity by 1.6%.

[Insert Table 7 here]

8 Personnel Policy as Seen Through the Lens of Our Network Model

In this section, we demonstrate how our network model of worker productivity can be put to use to inform personnel policy and to increase firm productivity. We start by asking our model three questions. Who should the firm strive to retain? Who should the firm let go? Who should the firm train? We contrast our answers to the answers from a standard model of worker productivity without network effects. We then continue with a brief discussion of the implications of our model for the design of shift schedules and work teams.

8.1 Worker retention and dismissal

To begin, imagine that we ask a personnel manager to pick out 10 workers that she feels the company should work the hardest *to retain*. One reasonable strategy would be to pick out

²¹ $\widehat{\phi}_3/\bar{y} = 0.0025/0.3589 = 0.0069$, where 0.3589 is the mean productivity level in the estimation sample.

the 10 workers with the highest average productivity. Our (model-based) strategy would be instead to pick out the 10 workers with the highest average *intercentrality measures*. These are our “key workers”. Recall that our measure of intercentrality (defined by Equation (4)) positively depends on a worker’s own productivity. But it also depends on the network effects that this worker generates. There are two sources of such co-worker peer effects in our local average model.²² There are contextual effects and local average endogenous peer effects. The size of these effects depend not only on workers’ own characteristics (including own productivity), but also on the unique position in the structure of the network occupied by each worker.

After picking our 10 key workers, we compare our workers to the 10 most individually productive workers chosen by the firms personnel manager. Despite a strong correlation between our measure of intercentrality and individual productivity (0.78), there are only 3 workers that are on both of our lists. According to our model, the total productivity loss incurred by losing our 10 key workers is 22% higher than the loss incurred by the 10 workers picked using the naive strategy based solely on a worker’s own productivity. In other words, our 10 key workers have a much higher overall value to the firm than the 10 workers with the highest average individual productivity. This is because these particular workers generate large positive externalities.

In the specific network depicted in Figure 3 (which is the network described in Figure 2 with 17 workers), we see that the agent with the highest productivity is not the agent with the highest intercentrality measure. In this particular network, the productivity loss from losing the key worker in this network is 26% larger than the productivity loss incurred by losing the working with the highest own productivity. This is due to the fact that the most productive worker in this network has a more peripheral position than the key worker.

If we instead are forced *to layoff* 10 workers and we want to minimize the aggregate productivity loss associated with these dismissals, then the firm should layoff the 10 workers with the lowest measures of intercentrality and not necessarily those with the lowest own productivity. Firing the 10 workers with the lowest average productivity leads to a productivity loss that is 12% larger than the loss that would be incurred if the firm had instead laid off the 10 workers with the lowest average measures of intercentrality. There is, however, a significant overlap between our list and the naive list. They have 5 workers in common.

In Figure 3, we see that the agent with the lowest productivity is not the agent with the lowest intercentrality measure. In this specific case, the productivity loss from laying off the

²²We here focus on the local-average model because we have shown that it was the right model for the impact of peers’ productivity on own productivity. See Table 4.

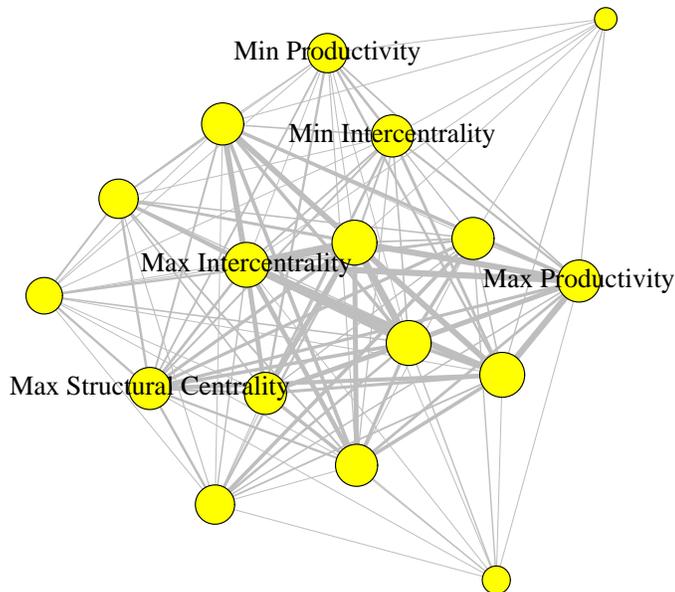


Figure 3: A co-worker network with productivity and centrality.

worker with the lowest productivity is 15% larger than the loss incurred by dismissing the worker with the lowest measure of intercentrality instead.

8.2 Which workers should be trained?

To analyze which workers should be trained, we use our local aggregate model with strategic complementarities (knowledge spillovers).²³ Our measure of intercentrality (given by Equation (5) for the local-aggregate model) allows us to identify individuals who lead to the largest drop in firm productivity if permanently removed from the firm. These are our “key workers”. On the-job-training, however, is a very different type of policy. The goal of on-the-job training is to raise a worker’s productivity and then place this worker back into the job or position in the network that she occupied before being trained. Training alters this worker’s productivity but not her personal characteristics and, hence, there are

²³We here focus on the local-aggregate model because we have shown that it was the right model for evaluating the impact of the training policy on the productivity of non-trained workers. See Table 7.

no changes in contextual effects. Furthermore, training does not alter the structure of the network. We must, therefore, adapt our measure of intercentrality so that it only measures the spillover effect that comes from the fact that each worker occupies a unique position in the existing network structure. This measure corresponds in fact to the key-player formula given by Ballester, Calvó-Armengol and Zenou (2006), which measures the total knowledge spillover effect that a worker’s productivity has on her co-workers’ productivity generated solely by the local aggregate peer effect. As long as everyone is not linked to everyone else, then this measure will differ across agents and depend only on the unique position in the network that a worker occupies and not on her characteristics. It is a structural property of the graph. This intercentrality measure can be calculated using formula (5) with no contextual effects. It is therefore defined by: $d_i(g, \alpha) = [b_i(g, \alpha)]^2 / m_{ii}(g, \alpha)$, where we set all individual α s equal to 1.

Between week 50/2008 and week 36/2009, the firm picked out 88 workers to be trained. We choose 88 workers using our index of structural intercentrality. The overlap in these two lists is not large; only 27 people are on both. According to our model, the productivity gain from training our 88 workers is 3.8% larger than the productivity gain achieved when training the 88 workers chosen by the firm. This gain is solely do to the unique positions in the network occupied by those we choose to train.

If we restrict ourselves to training only those workers who have at least 2 years of tenure (which is an exaggeration of the firm’s policy), then the overlap between who we choose to train and who the firm chooses to train increases to 37 workers (out of 50) and the productivity gain of our choice falls to 3.3%, which is a quite large effect in this context. In this example, gains from training could be increased from 7.8% to 11.1% by simply choosing who to train more strategically.²⁴

8.3 Optimal network design

In the *local average model*, the optimal network structure (from a worker’s perspective) is the empty network.²⁵ This, however, is not allowed by the firm. The second-best, again from the worker’s perspective, would be to have a small network in which a single worker can influence average productivity. Employers, on the other hand, may want to have larger teams in order to rationalize organization and monitoring. In the local average model, employers have no

²⁴7.8% is the coefficient on on-the-job training reported in Column (1) of Table 6.

²⁵Indeed, for the local-average model, the utility function is given by (1) for which $\lambda_1 = 0$. Since social interactions with others only involve a cost (the cost of conforming to the norm), it should be clear that the optimal network that maximizes the sum of the utilities of all workers is the empty network.

incentive to manipulate the structure of the network given a fixed productivity level. They will, however, have incentives to try and maximize average productivity across networks, perhaps by moving key workers to new teams.

The optimal network in the *local aggregate model* is either the complete network or a nested split graph (König, Liu and Zenou, 2014; Billand et al., 2015; Belhaj, Bervoets and Deroïan, 2015).²⁶ Once again, such an extreme result is not likely to be of practical use for the firm.

There are no such analytical results available for a hybrid model (that includes both average and aggregate network effects). Instead, there is a built-in tension between the complete network and the empty network, subject to the constraints of the firm (e.g., total hours worked, the timing of customer demand, etc.). The optimal network structure in the hybrid model will have some “smoothing” properties (network effects spread more rapidly in more complete networks) and workers who work more than one shift during the day will play an important role. These “bridge” workers are necessary to facilitate the spread of network effects across shifts.²⁷ More generally, larger more well connected networks should dominate, but the optimal size and degree of connectedness should fall well short of the complete network.

To illustrate these ideas, we run several descriptive regressions at the network level using our aggregate model (i.e. we focus on the role of network size and structure for the propagation of knowledge spillovers). First, we regress predicted productivity (from our aggregate model) onto indicators variables for the share of workers in each network who work more than one type of shift during the week (i.e., morning and midday or midday and evening). We also control directly for the share of workers needed to man each shift. According to this descriptive regression, increasing the share of morning and midday “bridge” workers in a network by 10%, while keeping the total number of workers on each shift, hours worked, and worker quality constant, leads to a rise in aggregate productivity of 3.8%.²⁸

We then examine the role played by network connectedness as measured by the average betweenness centrality of a co-worker network.²⁹ An increase of one standard deviation in

²⁶A nested split graph is a hierarchical structure such that the neighborhood of an agent with low centrality is a subset of the neighborhood of another agent with higher centrality, i.e. neighborhoods are nested. See König, Tessone and Zenou (2014) for a precise definition.

²⁷In the data, we define the bridge workers as the ones who, for a given week, work either in the morning *and* in the afternoon or in the afternoon *and* in the evening so that productivity can flow from morning workers to evening workers who would not be “in contact” otherwise.

²⁸The mean share of these bridge workers in the data is 0.3 with a standard deviation of 0.18.

²⁹The betweenness centrality of a given worker is equal to the number of shortest paths between all pairs of workers that pass through the given agent. In other words, a worker is central if she lies on several shortest

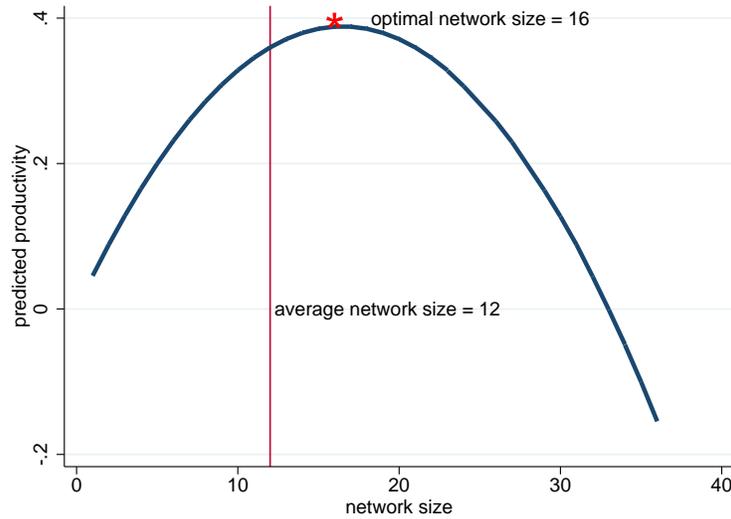


Figure 4: Optimal network size.

the average betweenness of a network is associated with a 2.8% higher average predicted productivity.³⁰

Finally, we examine the optimal network size. We do this by regressing average network productivity from our aggregate model on a quadratic function of network size. The resulting function is presented in Figure 4. Currently, the mean network size is 12, while the optimal network size is 16. Increasing the mean network size by four individuals is associated with an increase in average network productivity of 7.9%. The local-average model produces similar results.

Our policy conclusions concerning the optimal structure of co-worker networks can be summarized as follows: (i) The firm should increase team size, (ii) they should hire more “bridge” workers, and (ii) they should increase the average betweenness in each network.

9 Conclusion

We present evidence that co-workers can exert economically significant effects on their peers through network effects. A 10% increase in the current productivity of a worker’s co-worker paths among other pairs of workers. See Jackson (2008) for definitions and discussions of the different centrality measures.

³⁰The average betweenness centrality in the data is 0.09 with a standard deviation of 0.06.

network leads to a 1.7% increase in own current productivity. This productivity spillover can be attributed to conformist behavior. We see that low tenure workers react particularly strong to the work norm of their co-worker network.

We also find evidence of significant knowledge spillover effects from trained workers to their non-trained co-workers. Adding one additional trained co-worker to a worker's network increases that worker's own productivity by 0.7%.

The presence of such network effects affect the answer to a wide variety of policy questions faced by personnel managers on a daily basis. Our hope is that the literature on social networks will expand more vigorously into the field of personnel economics; a field that we believe is particular suited for network methods and models.

References

- [1] Acemoglu, D. (1997), "Training and innovation in an imperfect labour market," *Review of Economic Studies* 64, 445-464.
- [2] Acemoglu, D. and J.-S. Pischke (1998), "Why do firms train? Theory and evidence," *Quarterly Journal of Economics* 113, 79-119.
- [3] Acemoglu, D. and J.-S. Pischke (1999), "The structure of wages and investment in general training," *Journal of Political Economy* 107, 539-571.
- [4] Algan, Y., Dalvit, N., Do, Q.-A., Le Chapelain, A. and Y. Zenou (2015), "How social networks shape our beliefs: A natural experiment among future French politicians," Unpublished manuscript, Sciences Po, Paris.
- [5] Aral, S. (2016), "Networked experiments," In: Y. Bramoullé, B.W. Rogers and A. Galeotti (Eds.), *Oxford Handbook on the Economics of Networks*, Oxford: Oxford University Press, forthcoming.
- [6] Azoulay, P., Graff Zivin, J. and J. Wang (2010) "Superstar Extinction," *Quarterly Journal of Economics* 125, 549-589.
- [7] Babcock, P.S. and J.L. Hartman (2010), "Networks and workouts: Treatment size and status specific peer effects in a randomized field experiment," NBER Working Paper No. 16581.

- [8] Babcock, P.S., Bedard, K., Charness, G., Hartman, J.L. and H. Royer (2015), “Letting down the team? Social effects of team incentives,” *Journal of the European Economic Association* 13, 841-870.
- [9] Badev, A. (2014), “Discrete Games in endogenous networks: Theory and policy,” Unpublished manuscript, Federal Reserve Board, Washington D.C.
- [10] Ballester, C., Calvó-Armengol, A. and Y. Zenou (2006), “Who’s who in networks. Wanted: the key player,” *Econometrica* 74, 1403-1417.
- [11] Ballester, C., Calvó-Armengol, A. and Y. Zenou (2010), “Delinquent networks,” *Journal of the European Economic Association* 8, 34-61.
- [12] Ballester, C. and Y. Zenou (2014), “Key player policies when contextual effects matter,” *Journal of Mathematical Sociology* 38, 233-248.
- [13] Bandiera, O., Barankay, I. and I. Rasul (2009), “Social Connections and Incentives in the Workplace: Evidence From Personnel Data,” *Econometrica* 77, 1047-1094.
- [14] Bandiera, O., Barankay, I. and I. Rasul (2010), “Social Incentives in the Workplace,” *Review of Economic Studies* 77, 417-459.
- [15] Bartel, A.P. (1995) “Wage Growth, and Job Performance: Evidence from a Company Database,” *Journal of Labor Economics* 13(3), 401-25
- [16] Bayer, P., Hjalmarsson, R. and D. Pozen (2009), “Building criminal capital behind bars: Peer effects in juvenile corrections,” *Quarterly Journal of Economics* 124, 105-147.
- [17] Beaman, L., BenYishay, A., Magruder, J. and M. Mobarak (2015), “Can network theory-based targeting increase technology adoption?” Unpublished manuscript, Northwestern University.
- [18] Becker, G. (1962), “Investment in human capital: A theoretical analysis,” *Journal of Political Economy* 70, 9-49.
- [19] Belhaj, M., Bervoets, S. and F. Deroïan (2015), “Efficient networks in games with local complementarities,” *Theoretical Economics*, forthcoming.

- [20] Bifulco, R., Fletcher, J.M. and S.L. Ross (2011), “The effect of classmate characteristics on post-secondary outcomes: Evidence from the Add Health,” *American Economic Journal: Economic Policy* 3, 25-53.
- [21] Billand, P., Bravard, C., Durieu, J. and S. Sarangi (2015), “Efficient networks for a class of games with global spillovers,” *Journal of Mathematical Economics*, forthcoming.
- [22] Blume, L.E., Brock, W.A., Durlauf, S.N. and Y.M. Ioannides (2011), “Identification of social interactions,” In: J. Benhabib, A. Bisin, and M.O. Jackson (Eds.), *Handbook of Social Economics, Vol. 1B*, Amsterdam: Elsevier Science, pp. 853-964.
- [23] Blume, L.E., Brock, W.A., Durlauf, S.N. and R. Jayaraman (2015), “Linear social interactions models,” *Journal of Political Economy* 123, 444-496.
- [24] Bonacich, P. (1987), “Power and centrality: A family of measures,” *American Journal of Sociology* 92, 1170-1182.
- [25] Boucher, V. and B. Fortin (2016), “Some challenges in the empirics of the effects of networks,” In: Y. Bramoullé, B.W. Rogers and A. Galeotti (Eds.), *Oxford Handbook on the Economics of Networks*, Oxford: Oxford University Press, forthcoming.
- [26] Boucher, V. (2015), “Conformism and self-selection in social networks,” Unpublished manuscript, Laval University.
- [27] Bramoullé, Y., Djebbari, H. and B. Fortin (2009), “Identification of peer effects through social networks,” *Journal of Econometrics* 150, 41-55.
- [28] Bramoullé, Y., Kranton, R. and M. d’Amours (2014), “Strategic interaction and networks,” *American Economic Review* 104, 898-930.
- [29] Breuer, K. and P. Kampkötter (2013), “Determinants and effects of intra-firm trainings: evidence from a large German company,” *Journal of Business Economics* 83(2), 145-169.
- [30] Breza, E. and A.G. Chandrasekhar (2015), “Social networks, reputation and commitment: Evidence from a savings monitors experiment,” NBER Working Paper No. 21169.
- [31] Cai, J., de Janvry, A. and E. Sadoulet (2015), “Social networks and the decision to insure,” *American Economic Journal: Applied Economics* 7, 81-108.

- [32] Calvó-Armengol, A., Patacchini, E. and Y. Zenou (2009), “Peer effects and social networks in education,” *Review of Economic Studies* 76, 1239-1267.
- [33] Carrell, S.E., Fullerton, R.L. and J.E. West (2009), “Does your cohort matter? Estimating peer effects in college achievement,” *Journal of Labor Economics* 27, 439-464.
- [34] Centola, D. (2010), “The spread of behavior in an online social network experiment,” *Science* 329, 1194-1197.
- [35] Centola, D. (2011), “An experimental study of homophily in the adoption of health behavior,” *Science* 334, 1269-1272.
- [36] Chan, T.Y., Li, J. and L. Pierce (2014) “Compensation and Peer Effects in Competing Sales Teams,” *Management Science* 60, 1965-1984
- [37] Charness, G., F. Feri, M.A. Melendez-Jimenez and M. Sutter (2013), “Experimental games on networks: Underpinnings of behavior and equilibrium selection,” *Econometrica* 82, 1615-1670.
- [38] Choi, S., Gale, D., and S. Kariv (2012), “Social learning in networks: A quantal response equilibrium analysis of experimental data,” *Review of Economic Design* 16,93-118.
- [39] Cornelissen, T., Dustmann, C. and U. Schönberg (2013), “Peer effects in the workplace,” IZA Discussion Paper No. 7617.
- [40] Currarini, S., Jackson, M.O., and P. Pin (2010), “Identifying the roles of choice and chance in network formation: Racial biases in high school friendships,” *Proceedings of the National Academic of Sciences of the USA* 107, 4857-4861.
- [41] De Grip, A. and J. Sauermann (2012) “The effects of training on own and co-worker productivity: Evidence from a field experiment,” *Economic Journal* 122, 376-399.
- [42] De Grip, A., Sauermann, J. and I. Sieben (2011) “Tenure-performance profiles and the role of peers: Evidence from personnel data,” IZA Discussion Paper No. 6164.
- [43] Dearden, L., Reed, H. and J. Van Reenen (2006) “The impact of training on productivity and wages: Evidence from British panel data,” *Oxford Bulletin of Economics and Statistics* 68(4), 397-421.

- [44] Del Bello, C.L., Patacchini, E. and Y. Zenou (2015), “Neighborhood effects in education,” IZA Discussion Paper No. 8956.
- [45] De Paula, A. (2015), “Econometrics of network models,” Unpublished manuscript, University College London.
- [46] Falk, A. and A. Ichino (2006), “Clean evidence on peer effects,” *Journal of Labor Economics* 24, 39-57.
- [47] Fella, G. (2005), “Termination restrictions and investment in general training,” *European Economic Review* 6, 1479-1499.
- [48] Gelderblom, A. and J. de Koning (1996) “Evaluating Effects of Training within a Company: Methods, Problems and One Application,” *LABOUR* 10(2), 319-337
- [49] Goeree, J.K., McConnell, M.A., Mitchell, T., Tromp, T. and L. Yariv (2010), “The 1/d law of giving,” *American Economic Journal: Microeconomics* 2, 183-203.
- [50] Goldsmith-Pinkham, P. and G.W. Imbens (2013), “Social networks and the identification of peer effects,” *Journal of Business and Economic Statistics* 31, 253-264.
- [51] Goux, Dominique, and Eric Maurin (2000) “Returns to firm-provided training: evidence from French worker-firm matched data,” *Labour Economics* 7(1), 1-19
- [52] Graham, B.S. (2015), “Methods of identification in social networks,” *Annual Review of Economics* 7, 465-485.
- [53] Hahn, Y., Islam, A., Patacchini, E. and Y. Zenou (2015), “Teams, organization and education outcomes: Evidence from a field experiment in Bangladesh,” CEPR Discussion Paper No. 10631.
- [54] Hamilton, B.H., Nickerson, J.A. and H. Owan (2003), “Team incentives and worker heterogeneity: An empirical analysis of the impact of teams on productivity and participation,” *Journal of Political Economy* 111, 465-497.
- [55] Hsieh, C.-S. and L.-F. Lee (2015), “A social interaction model with endogenous friendship formation and selectivity,” *Journal of Applied Econometrics*, forthcoming.
- [56] Hunter, D.R., M.S. Handcock, C.T. Butts, S.M. Goodreau, and M. Morris (2008) “ergm: A Package to Fit, Simulate and Diagnose Exponential-Family Models for Networks,” *Journal of Statistical Software* 24(3), 1-29.

- [57] Ioannides, Y.M. (2012), *From Neighborhoods to Nations: The Economics of Social Interactions*, Princeton: Princeton University Press.
- [58] Jackson, C. Kirabo, and Elias Bruegmann (2009) “Teaching Students and Teaching Each Other: The Importance of Peer Learning for Teachers,” *American Economic Journal: Applied Economics* 1, 85-108.
- [59] Jackson, M.O. (2008), *Social and Economic Networks*, Princeton: Princeton University Press.
- [60] Jackson, M.O. (2011), “An overview of social networks and economic applications,” In: J. Benhabib, A. Bisin and M.O. Jackson (Eds.), *Handbook of Social Economics Vol. 1A*, Amsterdam: Elsevier Science, pp. 511-585.
- [61] Jackson, M.O. (2014), “Networks in the understanding of economic behaviors,” *Journal of Economic Perspectives* 28, 3-22.
- [62] Jackson, M.O., Rogers, B.W. and Y. Zenou (2015), “The economic consequences of social network structure,” CEPR Discussion Paper 10406.
- [63] Jackson, M.O. and Y. Zenou (2015), “Games on networks,” In: P. Young and S. Zamir (Eds.), *Handbook of Game Theory, Vol. 4*, Amsterdam: Elsevier, pp. 91-157.
- [64] Kandel, E. and E. P. Lazear (1992), “Peer pressure and partnerships,” *Journal of Political Economy* 100, 801-817.
- [65] Katz, L. (1953), “A new status index derived from sociometric analysis,” *Psychometrika* 18, 39-43.
- [66] Kearns, M.J., Judd, S., Tan, J. and J. Wortman (2009), “Behavioral experiments on biased voting in networks,” *Proceedings of the National Academy of Sciences* 106, 1347-1352.
- [67] Konings, Jozef and Stijn Vanormelingen (2015), “The Impact of Training on Productivity and Wages: Firm-Level Evidence,” *Review of Economics and Statistics* 97(2), 485-497.
- [68] König, M.D., Liu, X. and Y. Zenou (2014), “R&D networks: Theory, empirics and policy implications,” CEPR Discussion Paper No. 9872.

- [69] König, M.D., Tessone, C. and Y. Zenou (2014), “Nestedness in networks: A theoretical model and some applications,” *Theoretical Economics* 9, 695-752
- [70] Krueger, A.B. and C. Rouse (1998) “The Effect of Workplace Education on Earnings, Turnover, and Job Performance,” *Journal of Labor Economics* 16, 61-94.
- [71] Lawshe, C.H., Bolda, R.A. and R.L. Brune (1959) “Studies in management training evaluation: II. The effects of exposures to role playing,” *Journal of Applied Psychology* 43, 287-292.
- [72] Lechthaler, W. (2009), “The interaction of firing costs and firm training,” *Empirica* 36, 331-350.
- [73] Lee, L-F., Liu, X. and X. Lin (2010), “Specification and estimation of social interaction models with network structures,” *Econometrics Journal* 13, 145-176.
- [74] Leuven, E. (2005), “The economics of private sector training: A survey of the literature,” *Journal of Economic Surveys* 19, 91-111.
- [75] Liu, X. and L-F. Lee (2010), “GMM estimation of social interaction models with centrality,” *Journal of Econometrics* 159, 99-115.
- [76] Liu, X., Patacchini, E. and Y. Zenou (2014), “Endogenous peer effects: Local aggregate or local average?” *Journal of Economic Behavior and Organization* 103, 39-59.
- [77] Liu, X., Patacchini, E., Zenou, Y. and L-F. Lee (2012), “Criminal networks: Who is the key player?” CEPR Discussion Paper No. 8772.
- [78] Liu, X. and R. Batt (2007) “The Economic Pay-Offs to Informal Training: Evidence from Routine Service Work,” *Industrial and Labor Relations Review* 61, 75-89.
- [79] Loewenstein, Mark A., and James R. Spletzer (1999) “General and Specific Training: Evidence and Implications,” *Journal of Human Resources* 34(4), 710-733.
- [80] Manski, C.F. (1993), “Identification of endogenous effects: The reflection problem,” *Review of Economic Studies* 60, 531-542.
- [81] Mas, A. and E. Moretti (2009), “Peers at work,” *American Economic Review* 99, 112-145.

- [82] Paluck, E., Shepherd, H. and P.M. Aronow (2015), “Changing climates of conflict: A social network experiment in 56 schools,” Unpublished manuscript, Princeton University.
- [83] Patacchini, E., Rainone, E. and Y. Zenou (2015), “Heterogeneous peer effects in education,” Unpublished manuscript, Stockholm University.
- [84] Patacchini, E. and Y. Zenou (2012), “Juvenile delinquency and conformism,” *Journal of Law, Economic, and Organization* 28, 1-31.
- [85] Patacchini, E. and Y. Zenou (2014), “Social networks and parental behavior in the intergenerational transmission of religion,” CEPR Discussion Paper No. 8443.
- [86] Pfeifer, C., Janssen, S., Yang, P. and U. Backes-Gellner (2013) “Effects of Training on Employee Suggestions and Promotions: Evidence from Personnel Records,” *Schmalenbach Business Review* 65, 270-287.
- [87] Sacerdote, B. (2001), “Peer effects with random assignment: Results from Dartmouth roommates,” *Quarterly Journal of Economics* 116, 681-704.
- [88] Topa, G. and Y. Zenou (2015), “Neighborhood versus network effects,” In: G. Duranton, V. Henderson and W. Strange (Eds.), *Handbook of Regional and Urban Economics, Vol. 5A*, Amsterdam: Elsevier Publisher, pp. 561-624.
- [89] Waldinger, F. (2010), “Quality matters: The expulsion of professors and the consequences for PhD students outcomes in a nazi Germany,” *Journal of Political Economy*, 118, 787-831
- [90] Waldinger, F. (2012) “Peer effects in science - Evidence from the dismissal of scientists in nazi Germany,” *Review of Economic Studies* 79, 838-861
- [91] Zenou, Y. (2016), “Key players,” In: Y. Bramoullé, B.W. Rogers and A. Galeotti (Eds.), *Oxford Handbook on the Economics of Networks*, Oxford: Oxford University Press, forthcoming.

Table 1. Descriptive Statistics for the Full Sample

	worker \times week observations				within worker averages			
	N = 14,079				N = 425			
	mean	s.d.	min	max	mean	s.d.	min	max
performance, y_{it}	0.34	0.11	0.06	2.61	0.31	0.08	0.10	0.60
age_{it}	32.3	11.1	17	65	29.3	9.74	17	65
$male_i$	0.29	0.46	0	1	0.34	0.48	0	1
$tenure_{it}$	144	188	1	701	81.4	150	1.5	646
weekly hours worked, $hours_{it}$	21.3	10.1	0.05	82.7	22.3	7.69	2.82	47.65
on-the-job training, ojt_{it}	0.52	0.50	0	1	0.23	0.42	0	1
$monday_{it}$	0.76	0.43	0	1	0.74	0.22	0	1
$tuesday_{it}$	0.76	0.43	0	1	0.74	0.22	0	1
$wednesday_{it}$	0.74	0.44	0	1	0.74	0.23	0	1
$thursday_{it}$	0.73	0.44	0	1	0.74	0.23	0	1
$friday_{it}$	0.71	0.45	0	1	0.72	0.22	0	1
$saturday_{it}$	0.32	0.47	0	1	0.30	0.22	0	1
$sunday_{it}$	0.04	0.20	0	1	0.04	0.08	0	0.46
$morning_{it}$	0.32	0.47	0	1	0.30	0.24	0	1
$midday_{it}$	0.93	0.25	0	1	0.95	0.12	0	1
$evening_{it}$	0.28	0.45	0	1	0.28	0.27	0	1

Table 2. Pre-Treatment Descriptive Statistics for the Experimental Sample

	Treated, N = 29				Control, N = 41			
	mean	s.d.	min	max	mean	s.d.	min	max
performance, y_{it}	0.36	0.06	0.23	0.51	0.37	0.07	0.27	0.51
age_{it}	34.9	10.2	21	56	36.1	11.5	20	59
$male_i$	0.38	0.49	0	1	0.27	0.45	0	1
$tenure_{it}$	230	206	22	604	207	205	23	642
weekly hours worked, $hours_{it}$	19.8	6.05	6.42	29.6	20.2	5.70	8.33	31.9
$morning_{it}$	0.41	0.28	0	0.94	0.34	0.25	0	0.86
$midday_{it}$	0.93	0.13	0.53	1	0.91	0.20	0.21	1
$evening_{it}$	0.25	0.29	0	1	0.32	0.38	0	1
	Peers of Treated, N = 24				Peers of Control, N = 43			
performance, y_{it}	0.28	0.11	0.17	0.73	0.30	0.08	0.17	0.50
age_{it}	27.6	8.95	19	48	29.6	9.70	19	61
$male_i$	0.21	0.41	0	1	0.40	0.49	0	1
$tenure_{it}$	30.3	32.0	3.50	113	49.2	117	2	583
weekly hours worked, $hours_{it}$	17.9	7.89	4.57	28.6	19.3	6.91	7.23	29.2
$morning_{it}$	0.25	0.30	0	1	0.30	0.31	0	1
$midday_{it}$	0.98	0.05	0.82	1	0.96	0.10	0.50	1
$evening_{it}$	0.32	0.34	0	1	0.30	0.25	0	0.88

Table 3. Diagnostic Tests of Main Identifying Assumption

Dependent variables:	(1)	(2)	(3)	(4)	(5)
	\widehat{y}_i	\widehat{y}_i	$\widehat{h_{ij,r}^* \mathbf{1}_{j,r}}$	$\sum_{j=1}^n \widehat{h_{ij,r}^*} y_{j,r}$	$\frac{1}{g_{i,r}} \sum_{j=1}^n \widehat{h_{ij,r}^*} y_{j,r}$
Panel A: No fixed effects					
$\sum_{j=1}^n h_{ij,r} y_{j,r}$	-0.003*** (0.0014)				
$\sum_{j=1}^n h_{ij,r}^* y_{j,r}$	0.541*** (0.0271)				
$\sum_{j=1}^n h_{ij,r}^* tenure_{j,r}$		0.000*** (0.0000)	-0.004*** (0.0007)	-0.000** (0.0002)	0.000*** (0.0000)
$\sum_{j=1}^n h_{ij,r}^* age_{j,r}$		0.001*** (0.0004)	-0.025* (0.0140)	-0.007 (0.0041)	0.000 (0.0003)
$\sum_{j=1}^n h_{ij,r}^* male_{j,r}$		-0.008 (0.0091)	0.106 (0.2570)	-0.089 (0.0800)	-0.025*** (0.0059)
Panel B: With fixed effects					
$\sum_{j=1}^n h_{ij,r} y_{j,r}$	-0.000 (0.0000)				
$\sum_{j=1}^n h_{ij,r}^* y_{j,r}$	0.000 (0.0000)				
$\sum_{j=1}^n h_{ij,r}^* tenure_{j,r}$		-0.000 (0.0000)	0.000** (0.0000)	0.000 (0.0000)	0.000 (0.0000)
$\sum_{j=1}^n h_{ij,r}^* age_{j,r}$		0.000 (0.0000)	0.000 (0.0000)	0.000 (0.0000)	-0.000 (0.0000)
$\sum_{j=1}^n h_{ij,r}^* male_{j,r}$		0.000 (0.0000)	0.000 (0.0000)	-0.000 (0.0000)	0.000** (0.0000)

Dependent variables are the predicted values from linear regressions that include age, tenure, gender and individual, team and week fixed effects. Standard errors (in parentheses) are clustered on individuals; *** indicates significance at 1% level, ** at 5% level, * at 10% level.

Table 4. Estimation Results

	(1)	(2)	(3)	(4) Baseline	(5) 2SLS
Local aggregate network effect	-0.006** (0.0025)	-0.005** (0.0019)	-0.002 (0.0018)		
Local average network effect	0.692*** (0.0397)	0.203*** (0.0572)	0.182*** (0.0617)	0.174*** (0.0597)	0.261** (0.1281)
Weekly work hours			-0.001*** (0.0002)	-0.001*** (0.0002)	-0.001*** (0.0002)
Tenure			0.001*** (0.0002)	0.001*** (0.0002)	0.001*** (0.0002)
Morning			0.001 (0.0021)	0.001 (0.0021)	0.001 (0.0021)
Midday			-0.025*** (0.0041)	-0.025*** (0.0041)	-0.025*** (0.0041)
Evening			0.000 (0.0023)	0.000 (0.0023)	0.000 (0.0023)
Workday dummies			YES	YES	YES
Co-worker \bar{X}_j			YES	YES	YES
Individual fixed effects		YES	YES	YES	YES
Team fixed effects		YES	YES	YES	YES
Week fixed effects		YES	YES	YES	YES
H^*H^*X as instrument for H^*y					YES
Observations	14,079	14,079	14,079	14,079	14,070
F -stat					38
Individuals		425	425	425	416

Standard errors (in parentheses) are clustered on individuals; *** indicates significance at 1% level, ** at 5% level, * at 10% level. We include age, gender, and tenure in X when constructing H^*H^*X .

Table 5. Estimation Results for Workers with High versus Low Tenure

	(1)	(2)	(3)	(4)
	Baseline	Low tenure <= 104 weeks	Low tenure <= 52 weeks	Low tenure <= 26 weeks
Low tenure		-0.023 (0.0305)	-0.059** (0.0262)	-0.101*** (0.0248)
Local average network effect	0.174*** (0.0597)	0.083 (0.1090)	0.091 (0.0898)	0.056 (0.0813)
Low tenure \times average network effect		0.124 (0.0873)	0.148** (0.0716)	0.200*** (0.0700)
Observations	14,079	14,079	14,079	14,079
Number of person id	425	425	425	425

All regressions include our full set of controls and fixed effects. Standard errors (in parentheses) are clustered on individuals; *** indicates significance at 1% level, ** 5% level, * 10% level.

Table 6. Estimates of the Direct Experimental Treatment Effect and the Indirect Experimental Peer Effect

	Replication of original	Our peer effects experiment	
	training experiment		Low tenure <= 26 weeks
	(1)	(2)	(3)
On-the-Job Training	0.078*** (0.0200)	0.085** (0.0372)	0.028 (0.0437)
OJT*low tenure			0.134*** (0.0429)
Weekly work hours	-0.004*** (0.0011)	-0.001 (0.0014)	-0.001 (0.0014)
Share peak hours	-0.189*** (0.0542)	0.011 (0.0633)	0.010 (0.0639)
Time Trend	0.003*** (0.0008)	0.012*** (0.0021)	0.012*** (0.0026)
Low tenure			-0.021 (0.0356)
Individual fixed effect	Yes	Yes	Yes
Observations	1,781	1,116	1,116
Individuals	70	67	67

Dependent variable is $\log y_{it}$. Standard errors (in parentheses) are clustered on individuals; *** indicates significance at 1% level, ** at 5% level, * at 10% level.

Table 7. Network Effects from On-the-Job Training

	(4)	(5)	(6)
Local aggregate network effect, $\mathbf{H}_{trained}$	0.003** (0.0016)		0.003** (0.0011)
Local average network effect, $\mathbf{H}^*_{trained}$	-0.010 (0.0131)	0.006 (0.0090)	
Worker X_i (including $trained_i$)	YES	YES	YES
Workday dummies	YES	YES	YES
Co-worker \bar{X}_j	YES	YES	YES
Individual fixed effects	YES	YES	YES
Team fixed effects	YES	YES	YES
Week fixed effects	YES	YES	YES
Observations	5,572	5,572	5,572
Number of person id	264	264	264

Standard errors (in parentheses) are clustered on individuals;

*** indicates significance at 1% level, ** at 5% level, * at 10% level.