

NBER WORKING PAPER SERIES

PEERS AT WORK

Alexandre Mas
Enrico Moretti

Working Paper 12508
<http://www.nber.org/papers/w12508>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2006

We thank Stefano DellaVigna, Andrea Ichino, Daniel Parent, Jacob Vigdor, and seminar participants at UC Berkeley, IZA/SOLE Transatlantic Meetings, IZA Workshop on Institutions and the Employment Relationship, McGill, MIT Sloan, 2006 NBER Summer Institute, 2006 SITE Conference at Stanford, and University of Torino for useful comments. Sanny Liao and Ashley Langer provided valuable research assistance. E-mail: amas@haas.berkeley.edu and moretti@econ.berkeley.edu. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2006 by Alexandre Mas and Enrico Moretti. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Peers at Work
Alexandre Mas and Enrico Moretti
NBER Working Paper No. 12508
September 2006
JEL No. D01,J0

ABSTRACT

We investigate how and why the productivity of a worker varies as a function of the productivity of her co-workers in a group production process. In theory, the introduction of a high productivity worker could lower the effort of incumbent workers because of free riding; or it could increase the effort of incumbent workers because of peer effects induced by social norms, social pressure, or learning. Using scanner level data, we measure high frequency, worker-level productivity of checkers for a large grocery chain. Because of the firm's scheduling policy, the timing of within-day changes in personnel is unsystematic, a feature for which we find consistent support in the data. We find strong evidence of positive productivity spillovers from the introduction of highly productive personnel into a shift. A 10% increase in average co-worker permanent productivity is associated with 1.7% increase in a worker's effort. Most of this peer effect arises from low productivity workers benefiting from the presence of high productivity workers. Therefore, the optimal mix of workers in a given shift is the one that maximizes skill diversity. In order to explain the mechanism that generates the peer effect, we examine whether effort depends on workers' ability to monitor one another due to their spatial arrangement, and whether effort is affected by the time workers have previously spent working together. We find that a given worker's effort is positively related to the presence and speed of workers who face him, but not the presence and speed of workers whom he faces (and do not face him). In addition, workers respond more to the presence of co-workers with whom they frequently overlap. These patterns indicate that these individuals are motivated by social pressure and mutual monitoring, and suggest that social preferences can play an important role in inducing effort, even when economic incentives are limited.

Alexandre Mas
Haas School of Business
545 Student Services Building #1900
University of California, Berkeley
Berkeley, CA 94720-1900
and NBER
amas@haas.berkeley.edu

Enrico Moretti
University of California, Berkeley
Department of Economics
549 Evans Hall
Berkeley, CA 94720-3880
and NBER
moretti@econ.berkeley.edu

1. Introduction

In many production processes, output is a function not of the effort of a single worker, but of the combined effort of many workers. This kind of group production process is pervasive in modern economies. For example, most white collar jobs, construction, some manufacturing and retail, and co-authored academic research share this characteristic. When it is difficult for an employer to identify and reward the exact contribution made by each employee, free-riding has the potential to be a salient feature of these group work environments. Consider, for example, the case where a person is assigned a partner to complete a project. The employer observes total output perfectly, but individual effort only imperfectly. The effort that the worker puts into the project may depend on the productivity of her partner. If she is assigned a very productive partner, then it may make sense for her to ease her pace, relative to the case where she is assigned a less productive partner. However, if she makes very little effort compared to her partner, she may induce resentment or lose face with her peer. Because of this possibility, it could be optimal for this person to do her “fair share”, and work harder, not slower, in order to reduce the productivity gap with her more productive partner. Kandel and Lazear (1992) make this point theoretically, noting that peer effects can countervail free-riding in partnerships. In theory, peer effects have the potential to internalize some of the externalities that are common in workplaces. Ultimately, the question is an empirical one: do social preferences mitigate the deleterious effects of free-riding in real workplaces that are prone to externalities?¹

In this paper, we empirically investigate how workers influence each other. We explore how and why the productivity of a worker varies as a function of the productivity of her co-workers in a group production process that is particularly prone to free-riding. Our analysis centers on two questions. First, how does the introduction of a high productivity worker affect the productivity of her co-workers? As indicated, this relationship could go in any direction, depending on whether free-riding or peer effects dominate. Having found evidence of positive productivity spillovers, we then investigate the mechanisms that lead to such spillovers. We seek to distinguish between specific forms of peer effects that could be at work, including “social pressure”, “mutual monitoring”, and “contagious enthusiasm.” Economists have long speculated about the existence of productivity spillovers, but few studies have been able to explain the mechanisms that may generate them. This

¹ A different (but related) question is how a worker’s productivity depends on her co-worker productivity in settings with relative pay. In an interesting recent study, Bandiera, Barankay and Rasul (2005) find that workers internalize the negative externality their effort imposes on others under relative incentives—especially when working alongside their friends—but only when mutual monitoring is possible. Their evidence is inconsistent with altruistic behavior.

study is among the first to get inside the black box of productivity spillovers and shed some light on the underlying mechanisms.²

The question of whether there are peer effects in the workplace has significant implications for wage setting when individual output is not contractible. The return to introducing a high-productivity worker into a group is greater than her individual contribution if peer effects are strong. Alternatively, it is lower than her individual contribution if free riding prevails. Moreover, this question is important because peer effects may help explain what motivates workers in jobs with fixed-pay. In many occupations—including the one studied in this paper—career prospects are limited and compensation is not very sensitive to individual output. Economic incentives alone may not be enough to explain what motivates worker to exert effort in these jobs.

We study the productivity of cashiers in a national supermarket chain. In a supermarket, there is potential for negative externalities inherent in the production process. Customers in supermarkets are not committed to a single aisle. Therefore, for a given number of customers, if one checker is working slowly, other checkers will have an additional workload. An attractive feature of this setup is that we can use scanner data to develop a high-quality measure of productivity. Over a two year period, we observe the number of items scanned by each worker in each transaction, and the exact length of the transaction. We define individual productivity as the number of items scanned per second. Unlike much of the previous literature, which has relied on aggregate measures of productivity that vary with low frequency, our measure of productivity is precise, worker-specific and varies instantaneously.³

We relate ten-minute changes in each cashier's productivity to ten-minute changes in the average permanent productivity of the other workers who are active at that time in the same store. Over the course of a given day, the composition of the group of co-workers varies, because workers

² The literature includes many theoretical models that assume the existence of spillovers, and some recent empirical studies. Marshall (1890) is the first to hypothesize that on the job interactions may generate positive externalities across workers. Since then, growth economists have proposed theoretical models where spillovers are the determinant of growth (Lucas, 1988), and empirical studies have tested the existence of spillovers (for example, Moretti, 2004a and 2004b). A related literature investigates the implications of team formation, and their characteristics, on performance. See, for example Hamilton, Nickerson and Owan (2004) and Hayes, Oyer, and Schaefer (2006). Ichino and Maggi (2000) empirically investigate the role of social interactions in shirking behavior. A large literature focuses on peer effects in neighborhoods and schools. Examples include, by are not limited to, Bertrand, Luttmer, and Mullainathan (2000), Katz, Kling and Liebman (2001), Sacerdote (2001 and 2002), Oreopolos (2003 and 2006), Hanushek, Kain, Markman, and Rivkin (2003), Jacob and Lefgren (2003), Zimmerman (2003), Jacob (2004), Figlio (2005), Graham (2005), Vigdor and Nechyba (2004 and 2005), Vigdor (2005), Foster (2006), and Graham, Imbens and Ridder (2006).

³ Because of our need to measure productivity reliably, a study of this kind must necessarily focus on a single occupation. However, this occupation is important: cashiering is the second most common occupation in the United States, after Administrative/Secretarial.

shifts do not perfectly overlap. Therefore, for any given worker, the mix of her co-workers changes throughout the day depending on who enters and who exits. We find that substituting a worker with below average permanent productivity with a worker with above average permanent productivity is associated with a 1.1% *increase* in the effort of other workers on the same shift.⁴ The finding of a positive spillover suggests that positive peer effects dominate free riding. The magnitude of this productivity spillover is not trivial. A 10% increase in co-worker average permanent productivity is associated with 1.7% *increase* in a given worker's effort. The magnitude of this estimate is remarkably similar to recent experimental evidence of productivity spillovers (Falk and Ichino, 2006).⁵

This firm scheduling policy indicates that the timing of within-day changes in the average ability of co-workers can be considered exogenous. The reason is that the firm gives substantial scheduling flexibility to the workers, and there is no attempt by management to assign the best workers to the busiest shifts. Moreover, scheduling is determined two weeks prior to a shift, so that the within-day timing of entry and exit of workers due to shifts changes should largely be predetermined relative to transitory shocks to productivity. We present several empirical tests to verify that the timing of changes in the average ability of co-workers within a day is indeed unsystematic. All the tests confirm that, while shifts are not randomly assigned to workers, the timing of entry and exit of good workers appears uncorrelated with demand shocks or other determinants of individual productivity. For example, the entry of more able workers is not concentrated in the ten minutes prior to large increases in customer volume, as would be the case if managers could anticipate demand changes and bring in fast workers just prior to these increases. Similarly, the timing of exit of more able workers is not concentrated in the ten minutes prior to large declines in customer volume.⁶

Why are there positive spillovers? Two explanations relevant in our context are “social pressure” and “contagious enthusiasm.” By social pressure we mean that a worker experiences

⁴ In our models, the “reflection problem” does not arise, because we regress *current* individual productivity on average *permanent* productivity of co-workers.

⁵ Several laboratory experiments have linked peer pressure to improvements in team performance. In the experiment that most resembles the setting in our study, Falk and Ichino (2006) find that a 10 percent increase in peers' output results in a 1.4% increase in a given individual's effort. Falk, Fischbacher and Gächter (2003) study contributions to a public good and find evidence of strong social interactions. Fehr and Gächter (2000) find that when subjects have the option to sanction other players in the same team based on their contribution to the team output, free-riders are sanctioned and aggregate output is higher relative to cases where sanctions are not possible. See also Falk et al., (2001).

⁶ Moreover, while the mix of co-workers in the current period affects individual productivity in the current period, the mix of co-workers ten minutes into the future has no effect on individual productivity in the current period.

disutility when working less hard than other workers on the shift, but only if co-workers notice. Social pressure could arise from both formal and informal sanctions by co-workers, or through shame [Persico and Silverman 2005]. By contagious enthusiasm we mean that a worker experiences disutility if she knows that she is working slowly relative to other workers on a shift, even if no one else knows that she is working slowly. To distinguish between these mechanisms, we use information on the exact location of each worker's register and its spatial orientation. We estimate models where the effect of co-workers is allowed to vary depending on whether co-workers can monitor each other while working. If social pressure is the dominant channel for the existence of positive spillovers, the spillover effect should be large when a given worker is observed by her co-workers, and small when she cannot be observed by her co-workers. By contrast, if the primary mechanism contagious enthusiasm, the introduction of a new and productive worker should lead to increases in the productivity of workers who can see this new worker.

We find that when more productive workers arrive into shifts, they induce a productivity increase only in workers that are in their line-of-vision. The effect appears to decline with distance between registers: this effect is stronger for workers are in the line-of-vision and closer to the new workers than those that are further away. The effect of introducing more productive personnel on workers that are not in the line-of-vision from their register is not significantly different from zero. Moreover, we find that introducing any worker into a shift, regardless of their typical productivity, increases the productivity of incumbent workers who are in their line-of-vision, while actually decreasing the productivity of incumbent workers who are not in this new worker's line-of-vision.

This evidence is consistent with the notion that the estimated positive spillovers are generated by social pressure arising from workers monitoring each other. As a further test, we look at whether the magnitude of the spillover depends on the degree of schedule overlap of workers on a given shift. If workers on a given shift rarely overlap, they may not be receptive to social pressure due to the limited scope for external enforcement of norms. It is clearly more difficult to exert social pressure and impose sanctions on individuals that we rarely meet than individuals that we see every day. Similarly, the shame a worker endures when working slowly may not be as salient when she does not have to interact with her co-workers frequently in the future. Consistent with these notions, we find that introducing a high-productivity worker into a shift is associated with greater increases in incumbent productivity when the entering worker and the incumbent's schedule have high overlap than when they only coincide infrequently.

Having estimated an average spillover effect, and explored the mechanisms underlying this effect, we then turn to the question of how such effect varies across workers. We find that the

magnitude of the spillover effect varies dramatically depending on the skill level of the relevant worker. The spillover is large for workers with below average productivity, and is small for workers with above average productivity. More general models where an individual-specific spillover is estimated confirm that the effect is a monotonic, decreasing function of a worker skill level, benefiting mostly workers with low skills. Interesting, this relationship is rarely negative, suggesting that the productivity of high skill workers is not hurt by the presence of low skilled co-workers.

This finding is important because it implies that the optimal mix of workers in a given shift is the one that maximizes skill diversity. Overall productivity is higher when high skill and low skill employees work together in the same shift, compared to the case where they work in separate shifts. In the last part of the paper, we compare the mix of workers that is currently adopted by the firm with the random mix of workers and with the mix of workers that maximizes productivity (subject to the skill distribution of the existing set of workers). We find that the mix of workers adopted by the firm is close to the random mix. This finding confirms that shift assignment is not systematic. We also find that the observed mix of workers is different from the one that maximizes productivity. We calculate that that by optimally arranging the mix of workers in each shift, this firm could improve productivity by 0.2%. Taken as a whole, this supermarket chain could generate the same amount of sales with 124,000 fewer hours of work each year. At current wages, this amounts to more than \$2.5 million per year.

Does the presence of spillovers imply that this firm could raise its profit by \$2.5 million per year simply by reallocating current employees across shifts? Not necessarily. Workers currently have freedom of choosing their shift, a job attribute that is presumably valued by workers. The compensating differential associated with this freedom results in lower wages. Limiting this freedom is likely to result in higher equilibrium wages. The productivity gains from optimal mixing could be offset by higher wages.

Overall, the evidence assembled suggests that agents in this firm care about how they are perceived by their peers, either because they feel shame when not exerting effort, or because they are subject to sanctions from reciprocating co-workers.⁷ By seeking to minimize productivity differentials with their faster peers, slow workers display cooperative behavior, even if there are incentives to free ride. However, this behavior does not appear to be motivated by altruism. In fact, these workers display hallmarks of self-interested behavior: their extra-effort only occurs when it can be noticed by their peers, and when numerous future interactions can be expected to occur.

⁷ These findings are related to the experimental work of Charness and Rabin (2002) who find evidence that people wish to be liked (or at least not disliked) by their peers.

Regardless of the underlying preferences of the workers, our results demonstrate that incentives can arise through the social considerations of extrinsically motivated individuals, even when economic incentives are limited. This conclusion is supportive of Kandell and Lazear's [1992] theory, and in-line with a number of laboratory experiments which show that the presence of reciprocally motivated subjects in a labor market can lead to the enforcement of contracts, when there are no formal enforcement mechanisms, and even when there are selfish agents in the market [Fehr, Gächter, and Kirchsteiger 1997].

The paper is organized as follows. Section 2 presents a simple model. In Section 3 we describe the data and a preliminary test of spillovers. In Sections 4, 5 and 6 we describe our main empirical findings. In Section 7 we compare optimal mix with actual mix. Section 8 concludes.

2. Conceptual Framework

In many jobs, employers can observe total output, but can not observe exactly the contribution provided by each individual worker to the production of total output. As indicated in the introduction, this is feature of the workplace is common in most clerical occupations, in many manufacturing jobs, in construction, agriculture and in retail, especially when the number of employees working on a task is large. Consider, for example, the staff of a sale office putting together a presentation for a potential client. The employer can arguably observe the quality of the final presentation, and whether the potential client ends up buying the product. But it may be more difficult for the employer to observe exactly who did what in the project, and consequently how much effort each of the employees provided. It is more likely that the employer observes a noisy signal of their effort. Similarly, consider carpenters building a house. While the quality of the final output is easily observed, the effort provided at each moment in time by each carpenter may not be. As a final example, consider co-authored academic research. While any reader can observe the quality of a paper, the exact contribution of each co-author is not always clear.

In this sense, supermarket cashiers are not an exception. Customers typically choose the shortest line available, so that the length of the line is generally equal for all cashiers working at any given time. While it is always easy for management to observe the length of the line, it may be more difficult to identify which level of effort each cashier is providing at any moment in time. Not only are managers supposed to supervise many workers in the store that are not cashiers, but managers are

also expected to perform many other tasks besides supervising. Moreover, our assumption only requires that individual productivity is observed with some noise, however small.⁸

Our goal in this Section is to investigate how workers in a team will react to an exogenous change in the productivity of their co-workers when peer effects are present and when they are not. Absent peer effects, the basic idea is that a worker will exert less effort following the introduction of a high productivity co-worker to a shift when the worker's marginal benefit of effort declines as the effort of co-workers increases. Peer pressure can potentially mitigate this externality. We present a specific example of how the marginal utility of effort can depend on co-worker effort. The framework described here is intentionally kept very simple, and we note that there are certainly other models, which are in a similar spirit, that will also lead to this dependence. Kandel and Lazear (1992) develop a framework which provides similar insights.

Assume that productivity of worker i at a moment in time is an increasing function of her effort: $y_i = f(e_i)$, where y_i and e_i are unobserved by management, $f' > 0$ and $f'' < 0$. At each moment in time, management observes a noisy signal of each worker's output, $z_i = y_i + u_i$ —where u_i is idiosyncratic noise—and average output of all N workers in a shift, $\bar{y} = \frac{1}{N} \sum_{j=1}^N y_j$. In the context of our application, we can think of \bar{y} as the (inverse of) the length of the customers lines. Following the revelation of worker i 's noisy signal of output, her wage is set to her expected productivity given the signal

$$(1) \quad W_i = E(y_i | z_i) = b [z_i - \bar{z}] + \bar{y}$$

where $b = \frac{\text{var}(y)}{\text{var}(y) + \text{var}(u)}$ and both variances are assumed to be known. Equation (1) simply says that

management imperfectly observes the effort provided by each individual, but perfectly observes the length of the customer lines, and it combines these two pieces of information to infer who is working hard and who is not.⁹

⁸ Although the firm has in theory access to the same data that we use in this study, it has never used the data, mainly for computational difficulties. Indeed, this is the main reason why we were given the data.

⁹ Equation (1) is easily derived from the formula of an hypothetical regression of y_i on z_i , where the OLS intercept is $\bar{y} - \hat{b} \bar{z}$ and the OLS slope is $\text{cov}(y_i, z_i) / \text{var}(z_i) = \hat{b}$. Obviously this regression can not be run by the employer, but the intercept and the slope parameters are known under our assumptions. Note that if the signal has no noise, $\text{var}(u) = 0$ and wage equals true productivity $W_i = y_i$.

Case (a). Consider first the case where there are no social interactions. Workers choose effort to maximize the utility of income, U , minus the cost of effort, C :

$$\max U(W_i(e_i)) - C_i(e_i),$$

where $U' > 0$; $U'' < 0$; $C' > 0$; and $C'' > 0$. As usual, the first order conditions equalize the marginal benefit of effort to its marginal cost:

$$U'(W(e_i^*)) [f'(e_i^*) [b + [1 - b][1/N]]] = C'_i(e_i^*)$$

It is clear that in this context workers have a strong incentive to free ride. Each worker bears the full cost of her effort but gains only a fraction of the benefits in terms of reduced probability of punishment. It is also easy to see that workers' surplus is lower relative to the efficient level because of free riding.¹⁰

In this paper we are interested in what happens to the effort of a worker when the ability of her co-workers changes exogenously. More concretely, we are interested in learning how worker i 's effort changes following an increase in \bar{y} due to the substitution of a co-worker with high cost of effort with an otherwise identical co-worker with low cost of effort. Assume for example that $C_i(e_i) = k_i e_i^2$, where k_i characterizes the individual specific cost of effort. The derivative of the optimal effort of worker i with respect to the cost of effort of co-worker j in the same shift is

$$(2) \quad \frac{de_i^*}{dk_j} = - \frac{[b + [1 - b][1/N]] f'(e_i^*) U''(W(e_i^*)) [[1 - b][1/N]] [f'(e_i^*) \frac{\partial e_j^*}{\partial k_j}]}{[U''(W(e_i^*)) W'(e_i^*) f'(e_i^*) + U'(W(e_i^*)) f''(e_i^*)] [b + [1 - b][1/N]] - C'_i(e_i^*)} > 0$$

This expression is positive. The intuition is that, in the absence of any social considerations, an exogenous increase in co-workers productivity—caused by a decline in their cost of effort—results in more free riding and, therefore, in a decline in i 's effort.

Case (b). The prediction in equation (2) crucially depends on the assumption that mechanisms to internalize the externality generated by free riding are not available. What happens if this assumption is not true? Kandel and Lazear (1992), Huck, Kubler and Weibull (2002), Falk and Ichino (2004), and Falk, Fehr, and Fischbacher (2001) propose alternatives to this assumption. In particular, Kandel

¹⁰ The efficient level of effort is the vector (e_1, e_2, \dots, e_N) that maximizes total surplus, $\sum_i U[W(e_i)] - C_i(e_i)$. Obviously, if effort were observable, the efficient level would be easily achievable.

and Lazear (1992) argue that in team productions, peer pressure is likely to provide a mechanism that helps mitigate the free riding problem. They analyze how peer pressure operates and how factors such as social norms and mutual monitoring create incentives for workers.¹¹ Following their framework and notation, it is easy to incorporate peer pressure in our setting by assuming that workers maximize:

$$\max U[W_i(e_i)] - C_i(e_i) - P(e_i, e_1, e_2, \dots, e_{i-1}, e_{i+1}, \dots, e_N)$$

where $P(e_i, e_1, e_2, \dots, e_{i-1}, e_{i+1}, \dots, e_N)$ is a “peer pressure” function. It differs from the cost of effort function in that $P(\cdot)$ depends on other workers’ effort, not only on the relevant worker effort. The functional form of $P(\cdot)$ is a priori undetermined. In this paper we seek to describe a situation where peer pressure is a function of the distance between workers’ productivity. For example, a reasonable starting point is to assume that the cost of peer pressure is increasing in the distance between a worker’s productivity and the average productivity:

$$\begin{aligned} &P(e_i, e_1, e_2, \dots, e_{i-1}, e_{i+1}, \dots, e_N) \\ &= P \{ (1/(N-1)) [f(e_1)+f(e_2) + \dots + f(e_{i-1})+ f(e_{i+1}) + \dots + f(e_N)] - f(e_i) \}, \end{aligned}$$

where $P' > 0$. If each worker takes others’ effort as given, there is a unique equilibrium. The key implication is that the presence of peer effects may solve, or at least mitigate, the free riding problem.¹² In particular, the introduction of peer effects may change the sign of equation (2). It is possible to show that with strong enough peer effects

$$(3) \quad \frac{de_i^*}{dk_j} < 0.$$

In our empirical analysis we will seek to distinguish between the case described in equation (2) and the case described in equation (3).

For simplicity, we have modeled a worker’s wage as depending on her expected productivity (equation 1). It is easy to see that our results generalize to the case where a firm—like the one in our analysis—pays a fixed wage, provided that the firm can fire workers who are low-performers. The

¹¹ Another mechanism that might generate interdependence is relative compensation. If the probability of promotion or firing depends on relative performance, it is possible that a worker will increase her effort in response to an increase in productivity by her peer.

¹² It is even possible, in theory, that peer pressure result in an efficient equilibrium. However, unless management somehow has the ability to calibrate peer pressure, there is no reason to believe that the efficient level of effort will be achieved.

intuition is that even when a firm pays a fixed wage, there will still be a relationship between productivity and the expected wage if the firm can fire low productivity workers.¹³

3. A New Measure of Productivity and a Preliminary Test of Spillovers

Most existing studies of productivity rely on aggregate output measures, typically defined at the industry or firm level. These aggregate measures make it difficult to empirically identify spillovers. Furthermore, these measures typically vary only yearly. There is an emergent literature that investigates productivity in teams that uses either individual or team-level measures of productivity, for example Bandiera, Barankay, and Rasul (2005), Hamilton, Nickerson, and Owan (2004), and Leonard and Levine (2006). While the data used in these studies are eminently sensible to answer the particular questions they pose, our data is particularly well-suited to investigate the question of productivity spillovers for several reasons. First, we have a real-time measure of productivity, making it possible to identify *instantaneous* changes in individual productivity.¹⁴ Second, we not only know who is working at any moment in time, so that we can identify of a production group, but we also know the exact contribution of each member's output in the group.¹⁵ Third, we have plausibly exogenous variation on the scope of reciprocal monitoring, since we know which workers are able to observe other workers with relative ease based on their register assignment.

We use scanner data from a national supermarket chain to obtain a precise, high-frequency measure of productivity of cashiers. For each transaction, we observe the number of items scanned, and the length of the transaction in seconds. We define individual productivity as the average number of items scanned per second over a ten minute period. We include in our definition of productivity only periods when transactions are occurring.¹⁶

¹³ The whole problem can be recast in terms of the relationship between productivity and firing probabilities, yielding the same conclusions. For example, one could assume that the probability of being fired is a function of expected productivity: $g(E[y_i | z_i])$, with $g' < 0$ and $g'' < 0$. As a worker's productivity declines, the risk of being fired increases, and her expected wage declines provided that the difference between the current wage and the alternative wage (including any unemployment spell) is positive.

¹⁴ We believe that this is the first dataset that provides information on within-day variation in output.

¹⁵ By contrast, Hamilton, Nickerson, and Owan (2004) have individual productivity data prior to team formations, but only total team output once teams are formed.

¹⁶ In particular, for each worker on the shift, we sum the number of items that worker scanned over a ten minute period. We divide this number by the total number of seconds that the worker was in a transaction, where a transaction is defined as the time between when the first item is scanned to when the payment is completed and the receipt for the transaction is produced to the customer. As additional sample restrictions, we exclude any ten minute period where there is only one checker on duty; we only include observations where a worker is at the same register for at least two consecutive ten minutes periods, as we will be estimating first differences models and we wish to hold the registers where workers are stationed constant. We also exclude "10 items or less" registers.

Our sample includes all the transactions that take place in six stores for two years, for dates between 2003 and 2006. The stores are in the same metropolitan area of a state in the Western Census region. In total, we observe 370 cashiers. We know who is working at any moment in time, and the location and spatial orientation of the register operated by each worker within each store. We exclude transactions performed by managers. To minimize dead times, we focus on transactions between 7 AM and 8 PM. In the typical store, there are about 7 registers open with non-managerial workers on average at each moment in time. Table 1 reports descriptive statistics of the sample.

In this supermarket chain, workers are unionized and compensation is a fixed hourly payment. Due to union rules, checkers can only work at the registers. They can not bag items or stock shelves. Discussions with management indicate that the firm gives substantial scheduling flexibility to the workers. Managers have no role in determining which workers are working in particular shifts. Rather, managers submit the shifts schedules to the employees on a bi-weekly basis, and employees submit their scheduling preferences. If there are more workers asking for a particular shift than available slots, shifts are allocated based on seniority. Therefore, while shifts are not randomly assigned to workers, scheduling is quite unsystematic, and there is certainly no attempt by management to assign the best workers to the busiest shifts.

All the workers in our sample perform the same task (scanning items and receiving payment), use the same technology, and are subject to the same incentives. Nevertheless, there is substantial variation in productivity levels across workers, even after controlling for differences in shifts and the presence of co-workers. One way to show the amount of heterogeneity in the skill level of workers is to measure the permanent component of their productivity, holding fixed the day, the time of their shift and the co-workers mix. To do so, we construct a cell-level dataset of workers' productivity over ten minute intervals and other characteristics. We regress log productivity (items per second worked) of worker i at time t in calendar date c in store s , y_{itcs} , on a vector of worker fixed effects, θ_i , dummies for the presence of co-workers on the shift, $\sum_{j \neq i} \pi_j W_{jtcs}$ (with $j \neq i$); a vector of controls, X_{itcs} , that includes number of active workers at a moment in time (in ten minute intervals) as well as dummies for the register where a worker is stationed; a set of dummies for each day of week \times hour of the day \times store combination, γ_{dhs} (d denotes day of week and h denotes hour of the day); and a set dummies for each calendar-date \times store combination, λ_{cs} :

$$(4) \quad y_{itcs} = \theta_i + \sum_{j \neq i} \pi_j W_{jtcs} + \psi X_{itcs} + \gamma_{dhs} + \lambda_{cs} + \epsilon_{itcs}.$$

We interpret the worker fixed effects θ 's as a measure of each worker's permanent productivity. Workers with a high θ are on average more productive than workers with a low θ ,

holding constant the time and location of the shift and the presence of other co-workers. Note that the estimated θ_i 's do not reflect spillovers, because they are estimated controlling for co-workers composition.¹⁷

Figure 1 shows a distribution of the θ_i 's. The Figure indicates that there is a wide variation in workers skill level. The difference between the 10th percentile and the 90th percentile is 0.21, indicating that workers in the top part of the productivity distribution are 21% more productive than workers in the bottom part of the productivity distribution.¹⁸ One way to interpret this finding within the context of our model in Section 2 is that the cost of effort varies significantly across workers.

Before moving to our main results, we present an initial test for the existence of spillovers. The test is general in the sense that we impose very little structure on the form the spillovers can take. We estimate a model where productivity depends just on the presence of co-workers in a ten minute interval, in either a positive or negative direction:

$$(5) \quad \Delta y_{itc} = \sum_{j \neq i} \pi_j \Delta W_{jtc} + \Psi \Delta \mathbf{X}_{tc} + \gamma_{td} + \varepsilon_{itc}$$

where Δy_{itd} is the change in log of average productivity for worker i at time t , over 10 minutes intervals of calendar date c . As before, the W 's are indicators for the presence of each possible co-worker in the same shift in the relevant ten minute period; \mathbf{X}_{itd} is the number of workers at a given time, and γ_{td} is a set of dummies for each day/ten-minute period combination. This regression is run separately for each store.

The F-test of joint statistical significance of the π 's in equation (5) provides a preliminary test of the existence of spillovers. Finding that the presence of other co-workers has no effect on individual productivity would cast doubt on the notion that social pressure or free riding are empirically relevant. An attractive feature of this test is that it does not depend on any assumption on the specific form of the spillover. The test simply asks whether individual productivity depends *in any way* on which co-workers are in the same shift. We strongly reject the absence of interdependency of individuals' output. The average p-value for the F-statistic corresponding to the

¹⁷ One limitation of this specification is that there may be effects specific to combinations of workers. In particular, our model does not control for the interaction of specific group of workers. To address this possibility, we would have to add too many dummies (one for every possible combination of workers). This is computationally not feasible.

¹⁸ This difference is statistically significant. Low productivity workers have a significantly higher job-exit probability than high productivity workers. This type of selection is unlikely to be relevant for our estimates, since we only use variation within a worker and a day. Interestingly, we find limited evidence of an experience gradient in productivity.

null hypothesis that the π 's are jointly zero across the six stores is 0.000. This finding holds when we restrict the sample to changes in the composition of personnel that do not lead to net changes in the number of cashiers. In this case, the π parameters are identified by instances when one worker comes on duty, and a second worker departs.¹⁹ In this specification the average p-value is 0.018. In Section 4.3, we examine in detail possible threats to identification.

4. Estimating the Average Spillover Effect

In this Section, we present our baseline estimates of how the effort of worker i depends on her co-workers' average speed. We also present a series of tests intended to assess possible threats to identification. In the next Section, we allow our estimates of the spillover effect to vary across workers of different quality. In Section 6, we use information on mutual monitoring to investigate possible explanations of the documented spillover effect. Finally, in Section 7, we use our estimates to quantify the potential productivity gains that may be achievable from optimally mixing workers of different quality.

4.1 Baseline Model

We begin by estimating how worker i 's productivity at a given moment in time varies as a function of the average permanent productivity of the other workers who are active at that time in that store. In practice, we proceed in two steps. First, for each worker, store and time period we estimate the permanent productivity of co-workers, θ_i , using equation (4). Recall from equation (4) that the θ_i 's are estimated holding constant co-workers composition and therefore do not reflect spillovers. For each worker, 10 minute period and store, we average the permanent component of productivity of all the co-workers who are active in that period, denoting this quantity $\bar{\theta}_{-its}$, where $-i$ denotes that the average of the permanent productivity component is taken over all workers in store s , working at time t , calendar date c , excluding worker i . (To avoid any possibility of a mechanical relationship, the permanent productivity of co-workers who are active in a given ten-minute period is estimated omitting the data for that period.)

Second, we regress ten minutes *changes* in individual productivity on *changes* in $\bar{\theta}_{-its}$ and controls:

¹⁹ Even though we controlled for changes in the number of workers at the registers in the previous specification, the advantage of using this sub-sample is that it is very much in-line with our identification strategy. Changes in personnel that do not lead to net changes in the number of workers on duty are the result of workers ending shifts and other workers starting shifts, or workers taking breaks and other workers ending breaks, rather than new workers arriving (leaving), potentially because the store has become more or less busy.

$$(6) \quad \Delta y_{itcs} = \beta \Delta \bar{\theta}_{-itcs} + \psi \Delta X_{itcs} + \gamma_{itds} + e_{itcs},$$

where, as before, X_{itds} is the number of active workers in each ten minute interval at the relevant store, and γ_{itds} is a set of dummies for each day of the week \times 10 minute period \times store combination. The dependent variable is the 10-minute change in the log productivity of worker i . The term $\Delta \bar{\theta}_{-itcs}$ is the change in the average permanent productivity of i 's co-workers from $t-1$ to t . This term measures the change in the (inverse of) the average cost of effort of i 's co-workers. The coefficient of interest is β , which represents the derivative in equations (2) and (3). In a world without spillovers, this coefficient should be less than 0. If spillovers are large enough, β should be larger than 0. To account for possible serial correlation, we cluster the standard errors by checker, store and calendar day.²⁰

Four points are worth making. First, because the model is in first differences, we use only variation within a *given day for a given worker* to identify β . For any given worker, the mix of her co-workers changes throughout the day depending on who enters and who exits. Variation in personnel composition mainly comes from the staggered nature of shifts. Shifts overlap because it would be disruptive to change all the cashiers at the same time. Our central assumption is that permanent productivity of workers entering and exiting shifts within a day is orthogonal to changes in the productivity of other workers in the shift, aside from spillovers. This assumption is plausible because scheduling of shifts in the stores in our study is unsystematic, and management's only role in scheduling shifts is to determine how many workers are on duty at every point in time. Moreover, scheduling is determined two weeks prior to a shift, so that the entry and exit of workers due to shifts changes is predetermined relative to transitory shocks to productivity. In the next subsection we

²⁰ Two issues arise because the key independent variable, $\bar{\theta}_{-itcs}$, is estimated and therefore contains some error. First, estimates of the parameter β suffer from some attenuation bias, which will lead us to underestimate the magnitude of productivity spillovers. Second, estimates of the standard errors that we report in the tables underestimate the true standard errors, because they ignore the additional variability that comes from the fact that $\bar{\theta}_{-itcs}$ is an estimate. We do not believe these two problems are empirically serious. The reason is that the magnitude of both problems is a function of the amount of variability of the estimates of the $\bar{\theta}_{-itcs}$, which in turn is mainly a function of the size of the sample used. (In the limit, if an infinite sample was used to estimate the $\bar{\theta}_{-itcs}$, both problems would disappear.) Given that the sample that we use to estimate the $\bar{\theta}_{-itcs}$ parameters is very large, our estimates of the $\bar{\theta}_{-itcs}$ are extremely precise, and the attenuation bias and the underestimate of the standard errors should be negligible. In order to empirically verify whether the reported OLS (clustered) standard errors are far from the theoretically correct standard errors, we estimated the base specification in Table 2, column (1), bootstrapping the standard errors. Specifically, we bootstrapped the entire two-step procedure, where in the first-step we estimate the individual fixed-effects, and in the second step we estimate the base model, as described in the text. Using 500 repetitions, we found that the bootstrapped standard errors are virtually identical to the ones reported.

present a series of empirical tests intended to verify the validity of our assumption. These tests are consistent with our assumption. Among other things, the tests confirm that the timing of entry of good workers is not concentrated in periods of high customer volume and the timing of exit of good workers is not concentrated in periods of low customer volume.

A second point to note is that the parameter β estimates the effect of *permanent* co-workers' productivity on worker i 's current productivity. It does not estimate the effect of *current* co-workers' productivity on worker i 's current productivity. While both effects are potentially interesting, an advantage of the specification in equation (6) is that it does not suffer from the reflection problem.²¹ Because permanent productivity is predetermined, the interpretation of our estimates is not complicated by the fact that the effort of worker i may affect effort of worker j , and vice versa (Manski 1993). For the same reason, a transitory shock to productivity that affects the current effort of all workers at a moment in time will not generate mechanical spurious correlation.

Third, equation (6) assumes that peer effects operate linearly-in-means. While common in the literature on peer effects, this assumption is not the only possible one. For example, it could be that group productivity is affected by the least productive member of a group (Lazear 2001). Alternatively, it may be the case that the most productive person on the shift leads the rest of the group, and makes everyone more productive as well. Below we present some evidence that supports the validity of our specification.

Fourth, the use of a first-differenced model like the one in equation (6) has advantages and disadvantages relative to model in levels. An advantage is that by focusing on first differences we are able to take advantage of the high-frequency nature of the data and hone in on the precise timing of changes of co-worker composition. For example, we are able to identify the effects of a change in the composition of co-workers on the productivity of a worker *at that moment in time*, not ten minutes before or ten minutes after. A disadvantage of the first-differences model arises if the effects of introducing a high productivity worker changes over time. For example, introducing a high-productivity worker may have no effect on the productivity of her co-workers in the first twenty minutes, but lead to positive spillovers after that. In this case, a first-differences model would fail to detect an effect of co-worker composition, but a model in levels may. After analyzing a first-

²¹ The two effects measure two different things. The former measures how a worker's effort at a moment in time depends on her co-workers' permanent quality, irrespective of the effort that the co-workers may exert at that specific time. The latter measures how a worker effort at a moment in time depends on her co-worker effort at the same time. We have experimented with models where we measure co-workers productivity using their current productivity, and we instrument current productivity with permanent productivity. Results from these models are effectively just a rescaled version of the ones reported in Table 2 and are available on request.

differences model with distributed lags and leads, we determined that the first-differences model is well-suited for this study. In particular, in Section 4.2 we find that changes in average co-worker composition are associated with contemporaneous changes in a worker's productivity and, subsequently, this effect does not appear to vary dramatically.

Columns (1) and (2) in Table 2 reports estimates of equation (6). The first column indicates that there is a positive correlation between changes in co-worker permanent productivity and changes in individual productivity. The effect appears to be both statistically and economically significant. A 10% increase in co-worker permanent productivity is associated with a 1.8% increase in reference worker productivity.²² This finding indicates that positive spillovers appear to dominate any free riding effect. The return to introducing a high-productivity worker into a group is greater than her individual contribution.

Including 588 ten minute time interval by day of week dummies as changes in the number of workers on duty slightly lowers the point estimate to 0.16 (column 2). A 10% increase in co-worker permanent productivity is associated with 1.6% increase in reference worker productivity. The magnitude of this estimate is remarkably similar to recent experimental evidence of productivity spillovers. In a laboratory experiment, Falk and Ichino (2006) find that a 10 percent increase in peers' output results in a 1.4 percent increase in individual productivity.

Column (3) of Table 2 shows that workers are more responsive to the entry than the exit of a high productivity worker. This model shows that entry of a worker with above average permanent productivity is associated with a 1.1% increase in the productivity of co-workers. Because we also include a dummy for whether there is any entry of workers into a shift, this estimate should be interpreted as the effect of high-productivity entry above and beyond entry of workers with below average productivity. By contrast, the exit of an above average worker leads to about a 0.5% decline in co-worker productivity relative to the exit of a below average productivity worker. Using the midpoint between the effect of an entry relative to an exit of a high-productivity worker (0.75%), the estimates imply that if in every shift the firm could replace a lower than average productivity worker for a higher than average productivity worker, then the firm's labor inputs would decline by 0.75%, through the effects of the resulting spillovers alone, holding output constant.

In column (4) we consider a specification where we do not use variation in co-workers' quality due to the taking of breaks. The variables of interest are entry and exit of above average productivity personnel due to shift changes alone. In this case, the effect of high-productivity workers starting

²² All estimates are robust to limiting the sample to adjacent periods where there are no net changes in the number of personnel.

and ending their shifts appears to have the same magnitude effect on the productivity of other workers but with opposite signs. We estimate that when a high-productivity worker starts a shift, the productivity of other workers rises by about 0.6%. When a high-productivity worker ends a shift, the productivity of other workers falls by about 0.6%. Both of these estimates are statistically significant at conventional levels. We will come back to the remainder of the table later.

For now, we note that in equation (6) we estimate the spillover using an individual-level regression of worker i current productivity on co-workers average permanent productivity. Alternatively, one can use aggregate models to estimate the spillover. Specifically, one can sum both sides of equation (6) across individuals, and regress average current productivity on average permanent productivity. The level of observation becomes 10 minute period, day and store. In this type of models, the coefficient on average permanent productivity—which is sometimes called “social multiplier” (Glaeser, Sacreote and Scheinman, 2003)—should be equal to $1 + \beta$. In the absence of spillovers, an increase in average permanent productivity by 1% should be associated with an increase in current productivity by 1%, so that the coefficient on average permanent productivity should be 1. In the presence of spillovers, it should be larger than 1. Our estimates of the social multiplier from aggregate models are 1.17 (0.04) (no controls) and 1.14 (0.04) (with controls). The implied β s are therefore 0.17 and 0.14, similar to the β s estimated from the corresponding individual level regressions in columns (1) and (2).

4.2 Testing the Identifying Assumption

Identification of spillovers is typically challenging, because any factor that affects both the productivity and the composition of workers in a store may generate spurious correlation. For example, one might be concerned that shifts of high ability workers tend to coincide with days of the week when the customer volume is high. If higher customer volume causes workers to work harder, we may find a positive association even in the absence of spillovers. In our context, this is unlikely to be a problem. First, scheduling is unsystematic and management does not have control over which workers are working in each day. More importantly, our models are based on very short time-intervals. The parameter β in equation (6) is identified by changes in the composition of co-workers within a given day for a given worker in ten minute windows. Therefore, selection of workers with different ability across shifts does not drive our results.

One may still worry that high productivity workers tend to disproportionately enter (exit) shifts in periods within a day when demand rises (drops), relative to lower productivity workers. Finding that spikes in demand predict entry of high productivity workers may cast doubt on the

causal interpretation of our estimates. For example, this could happen if, in order to shorten the queues, more productive workers are brought at times of the day when demand is elevated and the productivity of all workers is high.²³ In this case, we might find a positive correlation between worker quality and productivity even in the absence of spillovers. Another finding that would cast doubt on the interpretation of our results is if high productivity employees are less likely to exit than low productivity employees during spikes in demand. For example, this could happen if high productivity workers are asked to delay the end of their shift because customer volume is particularly high at a given moment in time.

In general, the identifying assumption in equation (6) for the causal interpretation of β is that the changes in co-workers permanent productivity are orthogonal to changes in unobserved shocks affecting individual effort. We now test this assumption in several different ways.

(1) Entry and Exit. We begin by directly testing whether the entry (exit) of high productivity workers is associated with periods when demand is high (low). Unfortunately, we do not have direct measures of demand, because sales and number of customers passing through the checkout are themselves a function of productivity. Therefore, we devise several alternative measures and tests that are not subject to this problem. These analyses show no evidence of a relationship between changes in demand and the average ability of workers that enter or exit shifts at that time.

(1a) We first investigate whether the speed of entering and exiting workers is associated with *predictable* changes in demand. One way to measure predictable changes in demand is to compute the average quantity sold in each store by day of week and by time of day. A complication that arises is that our measure of customer volume is itself a function of the productivity of personnel. We circumvent this problem by predicting for each worker the average quantity sold omitting observations for which that worker was on duty. In other words, for worker i , we compute average quantity sold in store by day of week by time of day cells using only observations where worker i was not present. Predicting sales in this way allows us to determine the relationship between predictable and exogenous changes in demand across the day and the composition of coworkers in a shift, free of a mechanical relationship between these two measures. Concretely, we estimate the following models:

²³ Although possible in theory, this hypothesis is not very likely to be a serious concern in our context. Due to union rules, the checker's primary job is to work at the registers. Managers tend to work at the registers when demand increases. Managers do not factor into our analysis, however, because we have excluded them from the sample. Therefore it appears unlikely that there are high-productivity workers on the sideline ready to begin working in periods within a day when demand spikes upwards.

$$(7) \quad \text{ENTRY}_{itcs} = \alpha_i + \eta_1 \Delta \ln Q_{idt} + \rho_1 \Delta \ln Q_{idt} \cdot \theta_i + \gamma_{ids} + e_{itcs}$$

$$(8) \quad \text{EXIT}_{itcs} = \alpha_i + \eta_2 \Delta \ln Q_{idt} + \rho_2 \Delta \ln Q_{idt} \cdot \theta_i + \gamma_{ids} + e_{itcs}$$

where ENTRY_{itcs} takes on the value of 1 if worker i is observed on duty at time t , but not at $t-1$, and 0 otherwise; and EXIT_{itcs} takes on the value of 1 if worker i is observed on duty at time t but not at $t+1$, and 0 otherwise. The term $\Delta \ln Q_{idt}$ represents the change in the log of predictable sales, as described above, and $\Delta \ln Q_{idt} \cdot \theta_i$ represents the interaction of the change in the log of predictable sales and the permanent productivity of i . We are interested in ρ_1 and ρ_2 . If high productivity workers tend to enter shifts when demand is on average rising, or exit when demand is on average falling, then $\rho_1 > 0$ and $\rho_2 < 0$.

Table 3 shows that $\hat{\rho}_1$ and $\hat{\rho}_2$ cannot be statistically distinguished from zero and, if anything, the point estimates suggest that entry and exit of high productivity workers are less sensitive than entry and exit of low productivity workers to swings in demand. Columns (1) and (3) present estimates from specifications without worker fixed-effects. Columns (2) and (4) present estimates of the parameters in equations (7) and (8). The estimate of η_1 in column (2) shows that for the average worker, the probability of entry rises with increases in predicted sales. This is to be expected, if shifts are scheduled so that more workers are on duty during periods that are typically busy. The negative sign on the estimate of ρ_1 suggests that, if anything, the entry of fast workers is less affected by changes in predictable demand than the entry of slow workers. In column (4) we examine the determinants of exits. Not surprisingly, we find exit probabilities are lower when predictable demand is rising. But, as with entry, we find no evidence that exit hazards are more sensitive to changes in predictable demand among fast workers than slow workers. Specifically, we cannot distinguish $\hat{\rho}_2$ from 0.

(1b) Finding that entry and exit of good workers is not correlated with predictable demand shocks does not necessarily rule out the possibility that entry and exit of good workers is correlated with unexpected demand shocks. As a second test, we look at the relationship between lagged *actual* changes in sales and the probability of entry and exit of cashiers by productivity-type. Specifically, we test whether 10-minute periods when the change in number of items scanned is large (small) are immediately followed by gains (declines) in the permanent productivity of personnel. Relative to the

test (1a) above, this test has the advantage that it reflects actual demand conditions in a store at a point in time, rather than predicted demand.²⁴

We estimate models that are similar to those in Table 3, but rather than examining the effects of changes in predictable demand from $t-1$ to t on the entry and exit probability in period t , we examine changes in actual demand from $t-2$ to $t-1$ on the entry and exit probability in period t . (t represents a 10 minute interval) We find no evidence that fast workers are more likely to work in busy periods when using entry as the dependent variable. The coefficient on the interaction term in column (2) of Table 4, -0.058, implies that following a positive shock to demand, fast workers are less likely to begin working in the next ten minutes than slow workers. This finding is inconsistent with the view that fast workers tend to enter shifts when demand is rising and when, as a result, employees are working faster than their typical levels. Column (4) shows that the probability that high productivity workers are no more likely to exit than low productivity workers following declines in sales volume. In fact, the coefficient on the interaction of lagged changes in log quantity and individual permanent productivity, which is estimated as 0.01, implies that following positive shocks to sales fast workers are more likely to exit work ten minutes later.

Finally, we note that co-worker productivity matters even when we exclude changes in personnel that involve people coming into and out of breaks (Table 2, column 4). Again, this is inconsistent with the notion that workers end breaks early if there are large changes in demand.

(1c) As a third test, we consider the relationship between the number of personnel on duty and the average quality of workers. A positive relationship between the net change in personnel and the change in average permanent productivity of workers would suggest that the “marginal” worker that enters a shift when customer volume is rising tends to be more able. Relative to test (1b), this measure has the advantage that is based on the relationship between quality of workers at time t and personnel in the same period, rather than a lagged relationship. Figure 2 displays box plots of the change in average permanent productivity of workers across ten minute periods, by net changes in the number of workers on duty and by store. Consistent with our identifying assumption, the plots show that there is virtually no relationship between net changes in the number of personnel and changes in the average permanent productivity of personnel.

(1d) As a fourth test, we examine whether workers with high permanent productivity tend to work faster than their shift-level average in the early parts of their shift relative to the later parts of their shift. Such a relationship would suggest that fast workers tend to join shifts during periods

²⁴ Of course, for this test it is not possible to measure productivity at the same time as the change in customer volume, since such a measure is subject to the effects of a spillover.

when the conditions of the job require them to scan items quickly. In Figure 3 we plot predicted log productivity of workers with above and below average levels of permanent productivity, after controlling for time of day dummies and checker \times calendar date fixed-effects, for the first 3.2 hours of a shift.²⁵ Interestingly, the figure shows that workers with both high and low permanent productivity work below their mean levels in a shift in the early stages of a shift. Productivity tends to rise almost monotonically through the shift. This pattern is consistent with checkers requiring time to warm up. This pattern is inconsistent with the view that high productivity workers are entering during periods when there are positive productivity shocks, and when all workers are working quickly. In fact, it appears that, if anything, above-average workers begin slower than below-average workers early in the shift relative to later on in the shift.

(2) Leads and Lags. We now present a second series of tests intended to investigate whether the mix of co-workers at $t+1$ is correlated with individual effort at t . The rationale for these tests is simple. If our estimates reflect a true productivity spillover, and not spurious correlation, then the mix of workers ten minutes into the future should have no effect on individual productivity at the current period, conditional on the mix of workers in the current period. Finding that the leads are equal to zero means that high productivity workers are not entering in the ten minutes prior to large increase in demand, as would be the case if managers could anticipate demand changes and schedule workers just prior to these increases.

We estimate a model, akin to equation (6), containing the first lead in the average of co-worker productivity, and current average co-worker productivity. Consistent with our assumption, column (5) in Table 2 shows that the coefficient on the first lead is not significantly related to changes in contemporaneous productivity.

We also estimate a more articulated dynamic model, containing seven lags and seven leads in the change in average co-worker productivity. Figure 4 plots the coefficient estimates from the lagged, current, and future changes in average co-worker permanent productivity. Time 0 represents the time of the change in average permanent productivity of co-workers. The estimates show that exactly at the time when the average productivity of co-workers rises, there is an immediate rise in the relevant worker's productivity. Productivity then stabilizes, as inferred from the coefficients on the lagged changes in average co-worker permanent productivity, which are generally statistically

²⁵ Because these estimates come from a model with checker \times calendar date fixed-effects, the points represent log productivity relative to each checker's mean productivity during a shift. For example, a point at -0.02 on the y-axis means that at the time the worker is going approximately 2% slower than her average speed in the shift.

indistinguishable from zero.²⁶ For the purpose of testing the validity of the identifying assumptions, the estimates of interest are the ones on the leads. We would be worried if these estimates were larger than zero, as that would imply that future increases in co-worker quality are positively correlated with the relevant worker's current productivity. Figure 4 shows that all but one of the leads are not statistically different from zero.²⁷ This is inconsistent with the possibility of endogenous turnover of high-productivity checkers.²⁸

(3) Resource Constraints. Up to this point, we have focused on the possibility that our estimates are picking up selective sorting of personnel into and out of shifts that is correlated with demand shocks. We have found no evidence that this kind of sorting is occurring. A different confounder has to do with the possibility that there is a shared productive resource in this production process, which interacts with the productivity of co-workers in such a way as to generate the kinds of patterns we have documented thus far. An obvious candidate is the presence of baggers. Baggers assist the checker in putting groceries into bags. There are often fewer baggers than there are checkers. Therefore, introducing a high productivity worker may have implications on the productivity of other workers for no other reason that there is a constrained resource. While we cannot evaluate the implication of baggers directly from the data at hand, we can seek to assess the implication of baggers from an understanding of what they do. Managers at this firm have indicated to us that their policy is for baggers to “keep busy”. Our observation of baggers at work, when we visited the supermarkets in the sample, confirms this policy. Baggers move from register to register, going specifically to those registers where there are groceries to bag.²⁹ Given this policy, then the introduction of a new high productivity worker should have the effect of lowering, not raising, the productivity of other workers on the shift from the effect of the baggers alone. This is because faster

²⁶ The point estimate on the first lag is negative, implying the spillover effect may diminish somewhat over time. This finding is consistent with the estimates in Table 2 showing that the positive spillover effect of an above average productivity worker entering the checkout stand is twice as large in absolute value as the negative effect associated with a high productivity worker's exit. We note, however, that it is possible that the coefficients on the lags are downward biased, since the introduction of a high productivity worker in period t has a mechanical negative effect on the productivity of other workers in future time periods as the high productivity worker may eventually shorten customer queues. To the extent that workers reduce their speed when there are fewer customers in line, the coefficients on the lags will be biased because our models cannot control for the variable “length of the line”.

²⁷ The lone exception is the second lead, which is *negative* and statistically distinguishable from zero. This coefficient implies that twenty minutes prior to a positive shock in average co-worker productivity, checkers are working at a somewhat *slower* than average pace.

²⁸ Strictly speaking, this conclusion is true unless high-productivity checkers systematically begin (or end) their shifts contemporaneously with lumpy changes in demand. That is, demand is on average increasing between $t-1$ and t , when the high productivity checkers arrive (or depart), and demand is, on average, not changing between t and $t+1$. These conditions are quite special and appear to be rare in our data.

²⁹ We believe that this is the policy in many supermarkets in the United States, not just the ones in our sample. Readers are invited to consult their own experience.

workers on average have more groceries that require bagging, implying that baggers will spend less time with slower workers.

An additional piece of evidence that is relevant to this question, which we discuss in Section 6, is that the addition of high productivity workers has asymmetric effects on the productivity of other workers depending on who is in the new worker’s line-of-sight. We are unable to explain why sharing of a scarce resource would lead to these kinds of asymmetric effects of co-worker productivity when introducing a productive worker to a shift.

4.3 Testing the Linearity Assumption

As we mention above, the functional form of the peer effect function is a priori unknown, and the assumption that peer effects operate linearly-in-means is somewhat *ad hoc*. To shed some light onto this question, in Figure 5 we plot an estimate of the expected influence of an individual worker on her co-workers (the parameter π_i in equation 5), against that worker’s permanent productivity (the parameter θ_i in equation 4), using a local-linear smoother. This plot shows that a linear model actually does a good job in approximating at least the part of the social interactions that are related to the permanent productivity of workers. The plot shows that individuals for whom entry (exit) is associated with gains (losses) in productivity of other workers, tend to have higher levels of permanent productivity. This relationship exhibits some non-linearities, at lower levels of permanent productivity, but is generally linear in the region where most of the data are located. This chart implies that the relationship between an individual’s permanent productivity and her effect on others is essentially linear. Because the mean is a linear operator, we believe it sensible to consider a linear specification to examine the relationship between average co-worker productivity and effort.

5. Estimating Heterogeneity in the Spillover Effect

Having found support for the causal interpretation of the spillovers estimates, we now seek to analyze the characteristics of the spillovers in greater detail. The model described in equation (6) assumes that the spillover effect is the same for all workers. However, it is possible that the spillover effect depends on whether a worker is high ability or low ability. In this Section, we estimate models where we allow for the spillover effect to vary depending on the skill level of the relevant worker:

$$(9) \quad \Delta y_{its} = \beta \Delta \bar{\theta}_{-isct} + \lambda \Delta \bar{\theta}_{-its} L_i + \psi \Delta X_{itds} + \gamma_{itds} + e_{its},$$

where L_i is a dummy equal 1 if worker i permanent productivity is below average in the store. A positive (negative) λ and a positive β imply that low skill workers benefit from the spillover more

(less) than high skill workers. More generally, our longitudinal data allow for models with an individual-specific spillover effect, β_i :

$$(10) \quad \Delta y_{itcs} = \beta_i \Delta \bar{\theta}_{-itcs} + \psi \Delta X_{itcs} + \gamma_{ids} + e_{itcs}.$$

This model is more general than the one in (6) as it does not constraint the coefficient to be the same for all workers in a given skill group.

Notably, we find that the magnitude of the spillover effect varies dramatically depending on the skill level of the relevant worker. In particular, estimates of equation (9) in column (6) of Table 2 indicate that most of the spillover effect benefits the productivity of low productivity workers, but high productivity workers are not reducing their productivity in the presence of low productivity workers. While the coefficient is large and positive for workers who are below average (0.26), it is small but non-negative for workers who are above average (0.26-0.21=0.05).

When we turn to estimates of models with an individual-specific spillover effect, β_i (equation 10), we find that there is substantial variation in β_i . The spillover effect is large for some workers, and small—even negative—for others. Figure 6 plots estimates of the β_i conditional on worker i 's permanent productivity. Consistent with our previous finding for two groups of workers (column 6 of Table 2), the figure confirms that the spillover increases the effort of low productivity workers, and has little effect on the effort of high productivity workers. Notably, this relationship is negative in just a small number of cases, suggesting that the productivity of high skill workers is not hurt by the presence of low skilled co-workers.

The finding that the spillover effect is large for low skilled workers and small for high skilled workers is important because it has implications for the optimal mixing of workers. It implies that productivity is maximized when diversity is a shift is maximum. This finding is consistent with Hamilton, Nickerson and Owan (2004), but not with Costa and Kahn (2005). We will come back to this point in Section 7, where we compute productivity under optimal mixing and we compare it to observed productivity.

6. Inside the Black Box: Exploring the Channels Through Which Spillovers Operate

The results presented thus far indicate that there are significant productivity spillovers. The presence of high productivity workers raises the productivity of other workers, especially the ones who are normally less productive. The magnitude of the spillover is not trivial. What explains these findings? In this subsection we investigate the channels through which the peer effects operate. Kandel and Lazear (1992) propose several theoretical explanations for peer effects, but the existing

empirical evidence is limited. We seek to distinguish between the three (non-mutually exclusive) hypotheses that seem most likely to apply to our context:

- (i) *Social pressure*
- (ii) *Contagious enthusiasm (or Contagious malaise)*
- (iii) *Knowledge spillovers*

By social pressure we mean that a worker experiences disutility when working less hard than other workers on the shift, but only if his co-workers notice. Otherwise, working at a slower pace does not result in disutility. Social pressure could arise from both formal and informal sanctions. For example, if a worker is slow, other workers may impose a cost on her, for example, by reporting her to management, or through informal channels, for example, by ostracizing her socially. Social pressure does not require that workers impose costs on each other though. Viewing social pressure as isomorphic to “shame” in the Kandel and Lazear’s (1992) framework, a worker may care about what her co-workers think of her, even if the co-workers do not have the ability or desire to punish slow work. Irrespective of its specific form, social pressure has the implication that the introduction of a productive worker will lead to increases in the productivity of incumbent workers that are easily observed by the entering worker.

By contagious enthusiasm we mean that a worker experiences disutility if she is not working hard relative to other workers on a shift, even if no one knows that she is working slowly. This disutility could be related to guilt, as in Kandel and Lazear (1992). Alternatively, contagious enthusiasm could be induced by the presence of fast workers who lead-by-example.³⁰ If the spillovers are due to contagious enthusiasm, then the effect of introducing a new and productive worker into a shift should be greatest for incumbent workers who can easily observe the entering worker at their checkout stand.

Knowledge spillovers could occur as information is transmitted from one worker to the next. In our context these spillovers could arise, for example, if productive checkers know the codes for entering the price of fruits and vegetables and are able to transmit that information to other checkers. If there are knowledge spillovers, they should be related to proximity of checkers. While knowledge spillovers may be important in some industries (see, for example, Ichniowski, Shaw and Prennushi 1997), we do not expect them to have a quantitatively important effect in our context. Indeed, knowledge spillovers imply not only that the presence of fast workers makes slow workers more productive, but also that the presence of slow workers should make fast workers less productive. This

³⁰ The opposite type of effect may also exist, whereby slow workers induce other workers to slow down (contagious malaise).

is not consistent with our evidence in Section 5 which shows that faster workers appear to make slower workers go faster, but not vice versa.

To empirically distinguish between these hypotheses, we estimate models where the effect of co-workers is allowed to vary depending on whether co-workers can monitor each other while working (Section 6.1). We also corroborate our findings by examining how the spillover effect varies as a function of the frequency of interactions between workers (Section 6.2).

6.1 Mutual Monitoring

We have information on the location of each checker within a store, her spatial orientation and distance to other checkers. We estimate models that allow for the spillover effect to vary depending on the location of workers relative to their co-workers. The layout of the registers stands is such that when a checker is in position facing the customer, she is facing one set of registers, but not another set of registers. Therefore, co-workers located behind a given worker can easily observe her effort, while co-workers located in front of a given worker can not as easily observe her effort. We estimate the effect of the permanent productivity of a co-worker who enters or exits from a position behind an incumbent worker and the effect of permanent productivity of co-workers entering or exiting positions in front of incumbent workers. In Table 5 we show the hypotheses that are consistent with different combinations of asymmetries (or symmetries) in spillovers. We assume that while it is difficult for a worker to monitor the exact effort level provided at a moment in time by a co-worker located behind, her identity and her average productivity are known.

Table 5: Predictions of Alternative Hypothesis

Hypothesis Supported	Spillovers occur when co-worker is		
	In front	Behind	Both
Contagious enthusiasm	X		
Social Pressure		X	
Knowledge Spillover			X

The results of this exercise are quite stark. The estimates in column (1) of Table 6 show that just about the entire peer effect is operating through changes in workers that are able to monitor other workers due to their positioning. Specifically, as more productive workers are introduced into a shift, only the co-workers that are in their direct line-of-vision seem to become more productive. By contrast, changing composition of workers does not appear to influence the productivity of

incumbent workers when the incumbent workers are not in the line-of-vision of the new workers. This finding is consistent with a social pressure explanation.

The model in column (2) includes only a dummy for the change in the presence of a worker in front and a dummy for the change in the presence of a worker behind. Also consistent with the notion of social pressure, the addition of a worker behind an incumbent worker, regardless of her productivity, results in increased productivity of the incumbent worker of 4%, an estimate that is highly significant. The addition of a worker in front, on the other hand, *decreases* productivity of the incumbent worker by approximately 3%. This finding suggests that there is still scope for free-riding, but only when the free-riding is difficult to observe by other workers. We note that the spatial patterns in column (1) and (2) would be difficult to explain purely based on omitted variables.

In columns (3) and (4) we test whether physical distance matters. The social pressure applied by co-worker who are located behind and are closer appears to have a larger effect than the pressure applied by co-workers who are behind and are farther. For example, the coefficient on permanent productivity of co-workers who are one or two registers away behind the reference worker is 0.16 (column 3). The corresponding coefficient for co-workers who are three or four registers are behind the reference worker is 0.1, or about sixty percent as large. These estimates are significantly different from each other at conventional levels of significance.³¹ We find similar results when we look at dummies for addition of co-workers. We estimate that the change in the presence of a co-worker one or two positions behind worker i increases i 's productivity by 2.5%. A change in the presence of a co-worker three or four positions behind i increases i 's productivity by only 0.7%. The difference in these two estimates is significantly different than zero. Similarly, the coefficient on a dummy for the change in the presence of a co-worker located in front and close is -0.04, much larger in magnitude than the coefficient on a dummy for the change in the presence of a co-worker located behind but farther apart, 0.0001.

6.2 Repeated Interactions

The asymmetric effect by spatial arrangement suggests that the peer effect that we observe is the result of social pressure. Here we seek to provide an additional test of the social interaction hypothesis based on the frequency of interaction between workers. The idea is straightforward. If social interactions are indeed the explanation of the peer effect that we find, we may find that the magnitude of the spillovers depends on whether workers on a given shift have overlapped frequently,

³¹ Considering the effect of permanent productivity of workers that are farther than 4 positions away reduces the sample substantially, too much for the analysis.

or infrequently in the past. If a worker does not overlap often with somebody on a given shift, she may not be as receptive to social pressure because there is not much of a repeated component to the social interaction. It is clearly more difficult to exert social pressure on individuals that we meet rarely than individuals that we see every day. It may also be the case that workers who overlap infrequently may not know each other's ability, and therefore may not be as responsive to each other's permanent productivity.

Suppose that worker i , is on duty with checkers j and k at time t , where time is defined as a ten minute interval. We compute the share of worker i 's work-time that coincides with worker j and k up to time t . We eliminate the first month of the sample because we require a window to calculate these shares.³² We then estimate models where we let the spillover vary depending on the frequency of interaction:

$$(11) \quad \Delta y_{itcs} = \beta_L \Delta \bar{\theta}_{-istc}^L + \beta_M \Delta \bar{\theta}_{-istc}^M + \beta_H \Delta \bar{\theta}_{-istc}^H + \psi \Delta X_{itcs} + \gamma_{itds} + e_{itds},$$

where L denotes workers that have previously had low schedule overlap with i , M denotes workers that have previously had medium schedule overlap with i , and H denotes workers that have previously had high schedule overlap with i . Equation (11) models changes in worker productivity as depending on changes in the average levels of permanent worker productivity, where the average co-worker productivity is taken over co-workers in three categories of schedule overlap: high, medium and low. For example, the term $\Delta \bar{\theta}_{-i,s,t}^M$ denotes the change in the average permanent productivity of i 's co-workers who have previously had medium overlap with i .³³ The vector of covariates includes the change in the number of workers on duty in a ten minute interval, as well as store by day of week by time of day dummies.

Equation (12) further breaks down the permanent productivity averages by spatial orientation of co-workers:

$$(12) \quad \begin{aligned} \Delta y_{itcs} = & \beta_{B,L} \Delta \bar{\theta}_{-i,s,t}^{B,L} + \beta_{F,L} \Delta \bar{\theta}_{-i,s,t}^{F,L} + \beta_{B,M} \Delta \bar{\theta}_{-i,s,t}^{B,M} + \\ & + \beta_{F,M} \Delta \bar{\theta}_{-i,s,t}^{F,M} + \beta_{B,H} \Delta \bar{\theta}_{-i,s,t}^{B,H} + \beta_{F,H} \Delta \bar{\theta}_{-i,s,t}^{F,H} + \psi \Delta X_{itcs} + \gamma_{itds} + e_{itcs}, \end{aligned}$$

³² In principle, there are at least two ways to measure how exposed workers are to each other. For each pair of workers, one can count the number of interactions workers have had with each other in the past, or one can compute the share of all time spent working that workers overlapped. Data limitations prevent us from computing the first measure for all but the workers that began working in the sample period, due to censoring. A more attractive measure, that circumvents the censoring problem, is to estimate the share of the total time working that overlap with each of the other workers at every point in time.

³³ We define low schedule overlap as cases where co-workers have previously coincided with i between 0% and 5% of i 's schedule. Medium overlap co-workers have coincided with 5%-20% of i 's schedule. High overlap workers are those that have coincided with 20%-100% of i 's schedule.

where B denotes workers that are positioned behind i , and F denotes workers positioned in front of i .

In column (1) of Table 7, we show estimates of equation (11). The point-estimate of β_L is virtually zero, meaning that changes in the permanent productivity of co-workers who have had little previous overlap with i , have no affect on the change in i 's productivity from $t-1$ to t . By contrast, changes in the permanent productivity of co-workers with medium and high previous overlap with i have a positive and statistically significant relationship with changes in i 's contemporaneous productivity. Specifically, β_M and β_H are both estimated as approximately 0.08. The table shows that both $\hat{\beta}_M$ and $\hat{\beta}_H$ are statistically distinguishable from $\hat{\beta}_L$. In column (2) we present estimates from equation (12), which break-out the permanent productivity component of co-workers by previous overlap and spatial orientation in relation to i . As before, the entire spillover effect is coming from changes in the composition of workers for whom i is in the line-of-sight. Changes in the permanent productivity of these workers, however, do not appear to affect i 's productivity if exit and entry corresponds to workers that i has had relatively little previous overlap. We note that while the point estimate on $\hat{\beta}_{B,L}$ is close to zero, the large standard error corresponding to this estimate results in imprecise estimates of $\hat{\beta}_{B,H} - \hat{\beta}_{B,L}$ and $\hat{\beta}_{B,L} - \hat{\beta}_{B,M}$, which have p-values of 0.09 and 0.11 respectively against the null of zero. These findings are qualitatively consistent with results in Vigdor and Nechyba (2004) and Jacob and Vigdor (2005).

We have already seen in Table 6 that not only is there a relationship between changes in the permanent productivity of i 's co-workers and changes in i 's contemporaneous productivity, depending on the spatial orientation of the co-workers in relation to i , but there is also a relationship between changes in the presence of any co-worker, irrespective of their permanent productivity, depending on whether the workers who are entering or exiting are facing i . In column (3) of Table 7 we estimate a similar model. We test whether changes in the presence of co-workers, irrespective of their permanent productivity levels, affect i 's productivity depending on whether these co-workers have high, medium, and low previous overlap with i , as well as the spatial orientation of these co-workers in relation to i . Consistent with the findings for spatial orientation alone, we find that changes in the presence of a co-worker in front of i has either no effect, or a negative effect on changes i 's productivity. Changes in the presence of workers who face i are associated with positive changes in i 's productivity, again depending on whether the entering or exiting workers have previously had high or low schedule overlap with i . If in period $t-1$ there is no worker on duty that is facing i and with whom i has previously had high schedule overlap, and a worker meeting this

criteria enters, then i 's productivity increases by approximately 1.5% (t-ratio = 8). A similar conclusion is reached in the case of workers with medium previous schedule overlap with i . However, in the case of workers with low schedule overlap with i , their entry is not associated with a statistically significant change in i 's productivity. We can rule out positive changes in i 's productivity larger than 0.49% at the 5% level of significance.

In sum, the body of evidence in subsection 6.1 suggests that social pressure and monitoring of co-workers is important in understanding the source of the spillovers. When we further examine the spillovers by previous schedule overlap in subsection 6.2, we find results that are consistent with this explanation. Workers do not appear to be as affected by the presence of co-workers when their work schedules are different and therefore when they interact infrequently.³⁴ The results demonstrate that social pressure can partially internalize the negative externalities associated with a group production process. In such an environment, social preferences provide incentives for work when they are not provided in the marketplace. However, the findings in subsections 6.1 and 6.2 do not suggest that these workers display prosocial behavior, at least in the dimensions of effort that we have examined. While these individuals work harder in the presence of their more productive peers, they do so only when they are being observed and when future interactions are likely to occur.

7. Productivity Gains from Optimal Mix

In Section 5, we have shown that low productivity workers benefit from the spillover more than high productivity workers. This finding is important, because it implies that the mix of workers that maximizes productivity is the one that maximizes the diversity of skills in each shift. In other words, overall productivity is higher when high skill workers and low skill workers are employed in the same shift, compared to the case where some shifts are made of only of high skill workers and other shifts are made only of low skill workers. In this Section, we quantify the productivity gains that can be achieved by optimally mixing workers given the existing skill distribution. In particular, we compare workers' productivity under optimal mixing with workers' productivity under the mixing that is implemented by this employer. The difference is a measure of how far this employer is from optimal allocation. We then compute the possible savings in labor inputs that can be achieved by optimally mixing workers, holding constant sales. We also compare the observed workers mix with the mix that is obtained by randomly mixing workers.

³⁴ It is possible that our inability to find a spillover effect when workers have low overlap may be because workers must have information about their co-worker's abilities, which is less likely to be true if they have overlapped infrequently. We are unable to distinguish between these two explanations for the findings, but both are consistent with social pressure.

For computational simplicity, we divide the workers in our sample in two groups: workers with permanent productivity above the median (H), and workers with permanent productivity below the median (L). As in equation (9), output of high and low productivity workers in shift j can be written as $Y_{Hj} = a_H + b [N_{Hj} / (N_{Hj} + N_{Lj})]$ and $Y_{Lj} = a_L + c [N_{Lj} / (N_{Hj} + N_{Lj})]$, respectively, where N_{Hj} and N_{Lj} are the number of hours of work of H and L workers in shift j , and $[N_{Hj} / (N_{Hj} + N_{Lj})]$ is the share of hours of high productivity workers in that shift. Based on our findings in Section 5, we know that $c > b$. Indeed, our estimates suggest that $b = 0.0064$ and $c = 0.046$.³⁵

A first interesting comparison is between the observed mix of workers and the mix of workers that is obtained by randomly mixing high and low types in each shift, constraining the overall proportion of high and low types to be the same to the observed proportion. We find that the difference between the productivity obtained under the observed mix and the productivity obtained under the random mix is virtually zero (-0.0001%). This finding confirms that the mix currently employed by the firm is not systematic. Indeed it is close to be random. This is consistent with the fact that workers are free to choose the shift that they prefer and there is no attempt by management to affect workers mix.

An even more interesting comparison is between productivity obtained under the observed mix and productivity obtained under an ideal mix. This comparison allows us to calculate possible saving in labor inputs obtainable by optimally mixing the existing set of workers. To keep things simple, we begin by considering the optimal mix of high and low productivity workers in each shift, ignoring the issue of the optimal spatial location of workers within a store. Practically, we choose the fraction $[N_{Hj} / (N_{Hj} + N_{Lj})]$ for each shift j that minimizes the labor inputs necessary to produce the current level of output:

$$\min \sum_{j=1}^K (N_{Lj} + N_{Hj})$$

subject to

$$\sum_{j=1}^K N_{Lj} / \left(\sum_{j=1}^K N_{Lj} + \sum_{j=1}^K N_{Hj} \right) = \sum_{j=1}^K M_{Lj} / \left(\sum_{j=1}^K M_{Lj} + \sum_{j=1}^K M_{Hj} \right)$$

$$\sum_{j=1}^K N_{Hj} / \left(\sum_{j=1}^K N_{Lj} + \sum_{j=1}^K N_{Hj} \right) = \sum_{j=1}^K M_{Hj} / \left(\sum_{j=1}^K M_{Lj} + \sum_{j=1}^K M_{Hj} \right)$$

³⁵ These estimates are obtained from first differenced models similar to equation 10, controlling for all the covariates that are included in equation 10.

$$\sum_{j=1}^K \{N_{Lj} [a_L + c (N_{Hj} / (N_{Hj} + N_{Lj}))] + N_{Hj} [a_H + b (N_{Hj} / (N_{Hj} + N_{Lj}))]\} = \sum_{j=1}^K \{M_{Lj} [a_L + c (M_{Hj} / (M_{Hj} + M_{Lj}))] + M_{Hj} [a_H + b (M_{Hj} / (M_{Hj} + M_{Lj}))]\}$$

where M_{Lj} and M_{Hj} are the numbers of hours actually worked by low and high productivity workers in shift j , and K is the total number of shifts in the data. The first two constraints simply say that the optimal allocation needs to have the same overall fraction of H and L types that we observe in the data. In our case, this fraction is $\frac{1}{2}$. The third constraint simply says that the overall output under optimal allocation needs to be equal to the overall output that we observe in the data.

With this simple linear technology, it is easy to see that productivity is maximized—and labor inputs are minimized—when the share of high productivity workers $[N_{Hj} / (N_{Hj} + N_{Lj})]$ is set equal to $\frac{1}{2}$ in each shift. The skill mix that is actually adopted by this firm is not identical to the optimal skill composition. While the *average* share of high productivity workers across all shifts is (by construction) $\frac{1}{2}$, there is wide variation across shifts. Indeed, the share of high productivity workers in the typical shift is far from $\frac{1}{2}$. For example, in 73% of shifts, the fraction of high skill workers is below $\frac{1}{4}$ or above $\frac{3}{4}$.

Overall, our estimates indicate that the difference in productivity between the current mix of workers and the optimal mix of workers is 0.02%. This implies that by rearranging the mix of workers in each shift to maximize skill variance in each shift, this supermarket could produce the same amount of sales with 0.2% fewer hours worked each year. For this firm as a whole, the difference in labor inputs between the optimal mix of workers and the current mix of workers amounts to 123,529 hours worked per year. At an hourly labor cost (wage + health benefits + other costs) of \$20 an hour, this difference amounts to a wage bill that is approximately \$2.5 million per year higher than under the optimal mix.³⁶

Does the fact that the firm could use 123,529 fewer hours of work by simply rearranging the shifts of its employees mean that the firm is not maximizing profits? Not necessarily. The reason why optimizing shifts may not result in higher profits is that optimizing shifts may be costly for the firm because it may involve paying higher wages. Workers currently have the freedom of choosing their shift, a job attribute that is presumably valued by workers. The compensating differential associated with this freedom results in lower wages. Limiting this freedom could ultimately result in higher wages. This suggests that the additional productivity that the firm could obtain by imposing optimal

³⁶ This calculation includes all the stores of this chain, and assumes that the 6 stores in our data are representative.

mixing of workers may be equal, in equilibrium, to the total value of the compensating differential for flexibility in shift scheduling.³⁷

Based on our spatial results in Section 6, one might argue that the firm may be able to improve productivity by allowing workers to pick their preferred shift, but optimally locating workers across registers within each shift. The optimal allocation would be the one where the number of high productivity workers who can observe low productivity workers is maximized. In practice, however, the set of active cashiers changes quite frequently throughout the day, as cashiers start shifts, end shifts, take lunch breaks, etc. Moving cashiers location across registers is costly, because it takes time and disrupts consumer lines. For this reason, it is probably hard to significantly improve efficiency by constantly moving cashiers around.

8. Conclusion

We find that there is a social multiplier associated with the introduction of high productivity workers into work-groups: a 10% increase in co-workers productivity results in a 1.7% increase in individual productivity. The finding of a positive effect is particularly surprising, given that our data come from a group production process that is particularly prone to free-riding. In the absence of peer effects, we would expect to see that the introduction of a high productivity worker has a negative effect on other workers effort. Furthermore, while it may not be too surprising to find productivity spillovers in creative professions (for example, R&D or scientific research), it is surprising to find them in a low skill occupation where the tasks performed by workers are highly standardized. The magnitude of the estimated effect is remarkably similar to recent experimental evidence (Falk and Ichino, 2006). We consider several possible confounds that would threaten the validity of our estimates. Crucially, we find no evidence that workers are sorted into or out of shifts in such a way that turnover is systematically correlated with large changes in demand in a particular direction.

The average positive effect masks substantial heterogeneity in the spillover. Low productivity workers benefit from the spillover substantially more than high productivity workers. An important implication is that the mix of workers that maximizes productivity is the one where skill diversity in each shift is maximized. Our estimates indicate that by rearranging the mix of workers in each shift to maximize productivity, this firm could generate the same amount of sales with 123,529 fewer hours worked each year, although the effect on profits is unclear.

³⁷ This point was confirmed in a conversation with an executive of this firm. Informed of our results and the possibility of increasing productivity under optimal mixing, he immediately raised the point that wages might increase if workers can not pick their preferred shifts.

Our findings have important implications for wage setting. Hiring of a high productivity worker raises total output directly because of the worker has higher productivity, but also indirectly because the spillover raises the productivity of other workers. The return to a high productivity worker is therefore greater than her individual direct contribution.

While the literature on productivity spillovers is extensive, very little is known about the mechanisms that might generate such spillovers. Our findings indicate that peer pressure and mutual monitoring play an important role in inducing effort in peers. Consistent with this interpretation, we find that the introduction of highly productive workers into a shift boosts the productivity of incumbent workers that are easily observed by the new workers, but not incumbent workers that are not easily observed. Moreover, the introduction of new personnel into a shift does not have as much of an influence on the productivity of incumbent workers that they rarely overlap with.

The evidence we have assembled implies that social factors can partially internalize externalities that are built into many workplaces. While workers in this setting do not appear to be particularly altruistic, our findings, that people appear to care about how others perceive them, may be viewed optimistically. When workers hold themselves accountable to their peers, workplaces have the potential to be cooperative environments. Under this model, self-interest does not necessarily dictate that impulse towards motivation has its counterpart in inertia.

References

- Bandiera, Oriana; Barankay, Iwan, and Imran Rasul, "Social Preferences and the Response to Incentives: Evidence from Personnel Data," *Quarterly Journal of Economics* 120: 3 (August 2005): 917-962.
- Bertrand, Marianne; Erzo F.P., Luttmer, and Sendhil Mullainathan, "Network Effects and Welfare Cultures," *Quarterly Journal of Economics* 115 (2000): 1019-1055.
- Charness, Gary, and Mathew Rabin, "Understanding Social Preferences with Simple Tests," *Quarterly Journal of Economics*, 117(3) (2002): 817-869.
- Costa, Dora L. And Matthew E. Kahn "Surviving Andersonville: The Benefits of Social Networks in POW Camps" NBER Working Paper No. 11825
- Falk, Armin, Ernst Fehr and Urs Fischbacher, "Driving Forces of Informal Sanctions", IEW WP N. 59, University of Zurich (2001)
- Falk, Armin and Andrea Ichino, "Clean Evidence on Peer Effects," *Journal of Labor Economics* 24:1 (January 2006): 39-57.
- Falk Armin, Urs Fischbacher and Simon Gächter, "Living in Two Neighborhoods: Social Interactions in the Lab" IZA DP 1381 (2003).
- Fehr, Ernst and Simon Gächter, "Cooperation and Punishment in Public Goods Experiments," *American Economic Review* 90, (2000): 980-994.
- Fehr, Ernst; Gächter, Simon, and Georg Kirchsteiger, "Reciprocity as a Contract Enforcement Device: Experimental Evidence," *Econometrica* 65(4) (1997): 833-860.
- Figlio, David, "Boys Named Sue: Disruptive Children and Their Peers," NBER Working Paper No. 11277 (2005).
- Foster, Gigi, "It's not your Peers, and it's not your Friends: Some Progress Towards Understanding Educational Peer Effects", *Journal of Public Economics* (Forthcoming).
- Glaeser, Edward; Sacerdote, Bruce, and Jose Scheinkman, "The Social Multiplier," *Journal of the European Economic Association* 1(2) (2003): 345-353.
- Graham, Bryan. "Identifying Social Interactions through Excess Variance Contrasts," University of California, Berkeley mimeo (2005).
- Graham, Bryan; Imbens, Guido, and Geert Ridder, "Measuring the Average Outcome and Inequality Effects of Segregation in the Presence of Social Spillovers," University of California, Berkeley, Mimeo (2006).
- Hanushek, Eric; Kain, John; Markman, Jacob, and Steven Rivkin, "Does Peer Ability Affect Student Achievement?" *Journal of Applied Econometrics* 18(5) (September/October 2003)
- Hamilton, Barton; Nickerson, Jack, and Hideo Owan, "Team Incentives and Worker Heterogeneity: An Empirical Analysis of the Impact of Team on Production and Participation," *Journal of Political Economy* (June 2004): 465-497.
- Hayes, Rachel; Oyer, Paul, and Scott Schaefer, "Co-Worker Complementarity and the Stability of Top Management Teams" *Journal of Law Economics & Organization*, 22, (Spring 2006): 184-212.

- Huck, Steffen; Kubler, Dorothea, and Jorgen Weibull, "Social Norms and Incentives in Firms", mimeo, University College London (2002)
- Jacob, Brian, "Public Housing, Housing Vouchers and Student Achievement: Evidence from Public Housing Demolitions in Chicago" *American Economic Review* 94(1), 2004.
- Jacob, Brian and Lars Lefgren, "Are Idle Hands the Devil's Workshop? Incapacitation, Concentration and Juvenile Crime" *American Economic Review* 93(5), 2003
- Ichino, Andrea "Clean Evidence on Peer Effects," *Journal of Labor Economics* 24:1 (January 2006): 39-57.
- Ichino, Andrea and Giovanni Maggi, "Work Environment And Individual Background: Explaining Regional Shirking Differentials In A Large Italian Firm," *Quarterly Journal of Economics* 115 (2000): 1057-1090
- Ichniowski, Casey; Shaw, Kathryn, and Giovanna Prennushi, "The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines," *American Economic Review* 86 (June 1997): 291-313.
- Kandel, Edward and Edward Lazear, "Peer Pressure and Partnerships," *Journal of Political Economy* 100:4 (August 1992): 801-817
- Katz L.,J.R. Kling and J.B. Liebman, "Moving to Opportunities in Boston: Early results of a Randomized Mobility Experiment," *Quarterly Journal of Economics* 116(2) (2001): 607-654.
- Lazear, Edward, "Educational Production," *Quarterly Journal of Economics* 116:3 (August 2001): 777-803.
- Leonard, Jonathan and David Levine, "The Effect of Diversity on Turnover: Evidence from a Very Large Case Study," *Industrial Relations and Labor Review* (Forthcoming 2006).
- Lucas, R.E., On the mechanics of economic development. *Journal of Monetary Economics* (1988) 22, 3-42,.
- Marshall, A. Principles of Economics. Macmillan, New York (1890).
- Manski, Charles, "Identification of Endogenous Social Effects: The Reflection Problem," *Review of Economic Studies* (1993), 60: 531-542.
- Moretti, Enrico "Estimating the Social Return to Higher Education: Evidence From Longitudinal and Repeated Cross-Sectional Data" *Journal of Econometrics* 121(1-2) (2004).
- Moretti, Enrico, "Workers' Education, Spillovers and Productivity: Evidence from Plant-Level Production Functions," *American Economic Review* 94(3) (2004).
- Oreopoulos, Philip "The Long-run Consequences of Living in a Poor Neighborhood" *Quarterly Journal of Economics*, Vol. 118, No. 4, November 2003: 1533 – 1575.
- Oreopoulos, Philip "Neighbourhood Effects in Canada" mimeo, Toronto University, 2006.
- Persico, Nicola, and Dan Silverman "Other-Regarding Behavior: Theories and Evidence", University of Pennsylvania, 2006.

Sacerdote, Bruce, "Peer Effects With Random Assignment: Results for Dartmouth Roommates." *Quarterly Journal of Economics*, 116 (2001).

Sacerdote, Bruce "Peer Effects in Occupational Choice for Dartmouth Students," *European Economic Review*, Vol. 46 (May 2002).

Vigdor Jacob and Thomas Nechyba "Peer Effects in North Carolina Public Schools" in *Schools and the Equal Opportunity Problem*, P.E. Peterson and L. Woessmann, eds., MIT Press (2004)

Vigdor Jacob and Thomas Nechyba "Peer Effects in Neighborhoods and Housing" in *Deviant Peer Influences in Programs for Youth: Problems and Solutions*, K.A. Dodge, T.J. Dishion, and J.E. Lansford, eds., Guilford Press, forthcoming.

Vigdor Jacob and Thomas Nechyba "Peer Effects in Elementary School: Learning from 'Apparent' Random Assignment", mimeo, 2004.

Zimmerman, D.J.. "Peer Effects in Academic Outcomes: Evidence From a Natural Experiment. *The Review of Economics and Statistics* (2003) 85,1: 9–23

Table 1: Descriptive statistics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Store # 1	Store # 2	Store # 3	Store # 4	Store # 5	Store # 6	All stores
Share of 10 minute interval that checkers are transacting	0.66 [0.27]	0.61 [0.25]	0.60 [0.27]	0.69 [0.26]	0.66 [0.25]	0.60 [0.26]	0.64 [0.26]
Minutes per customer	1.33 [0.81]	1.18 [1.07]	1.74 [1.08]	1.27 [1.14]	1.40 [0.85]	1.38 [0.90]	1.36 [0.99]
Productivity in ten minute intervals	0.18 [0.07]	0.16 [0.07]	0.17 [0.08]	0.16 [0.07]	0.18 [0.07]	0.20 [0.08]	0.18 [0.07]
Checkers on duty in ten minute intervals	5.77 [1.94]	5.84 [1.60]	4.87 [1.75]	8.19 [2.38]	8.89 [2.67]	6.92 [2.30]	7.07 [2.62]
Estimated individual Fixed-Effects	0.00 [0.07]	0.00 [0.10]	-0.00 [0.07]	0.00 [0.06]	-0.00 [0.06]	-0.00 [0.07]	0.00 [0.07]
Average co-worker permanent productivity	0.00 [0.04]	0.00 [0.05]	-0.00 [0.04]	0.00 [0.02]	-0.00 [0.03]	-0.00 [0.03]	0.0 [0.03]
Change in co-worker permanent productivity	0.00 [0.04]	0.00 [0.02]	0.00 [0.02]	-0.00 [0.01]	0.00 [0.01]	-0.00 [0.01]	0.00 [0.02]

Notes: Standard deviations are in brackets. The units of observation are checker \times ten minute cells. Individual productivity is defined as the number of items scanned per second over a ten minute period. Specifically, for each worker on the shift, we sum the number of items that worker scanned over a ten minute period. We divide this number by the total number of seconds that the worker was in a transaction, where a transaction is defined as the time between when the first item is scanned to when the payment is completed and the receipt for the transaction is produced to the customer. We include in our definition of productivity only periods when transactions are occurring. The sample excludes any observations that do not occur in the 7:00 AM – 8:00 PM interval. The analysis excludes managers.

Table 2: The effect of changes of average co-worker permanent productivity on reference person current productivity

	(1)	(2)	(3)	(4)	(5)	(6)
Δ Co-worker permanent Productivity	0.176 (0.023)	0.159 (0.023)			0.160 (0.026)	0.261 (0.033)
Δ Co-worker permanent productivity _{t+1}					-0.010 (0.026)	
Δ Co-worker permanent prod. \times Above average worker						-0.214 (0.046)
Entry of above average productivity worker			0.011 (0.001)			
Exit of an above average productivity worker			-0.005 (0.001)			
Shift entry of above average productivity worker				0.006 (0.002)		
Shift exit of an above average productivity worker				-0.006 (0.002)		
Observations	1,734,140	1,734,140	1,734,164	1,734,164	1,356,643	1,734,140
Controls	No	Yes	Yes	Yes	Yes	Yes

Notes: Standard errors clustered by store \times calendar date \times checker are in parentheses. The units of observation are checker \times ten minute cells. The dependent variable is the change in the log productivity of a checker across ten minute periods. Individual productivity is defined as the number of items scanned per second over a ten minute period. The change in co-worker permanent productivity is computed as the simple average of co-worker permanent productivity components estimated by fitting equation (4) to the data. Controls are the change in the number of workers on duty in ten minute intervals (except columns 3-4), and 588 10 minute time interval by day of week dummies. In column (3)-(4), we also include in the models dummies for whether there was any entry or exit into or out of the shift, irrespective of the productivity of the worker entering or departing. Entry is 1 if in a given day a checker was on duty at time t but not on duty at time $t-1$, and 0 otherwise. Exit is 1 if in a given day a checker was on duty at time t but not on duty at time $t+1$, and 0 otherwise. Shift entry is 1 if an observation is a checker's first in a calendar date, and 0 otherwise. Shift exit is 1 if an observation is a checker's last in a calendar date, and 0 otherwise.

Table 3: The relationship between changes in predictable sales volume and exit and entry probabilities of personnel

	Entry at t (1)	Entry at t (2)	Exit at t (3)	Exit at t (4)
Δ Predicted log average Transactions between t-1 and t	0.026 (0.004)	0.022 (0.004)	-0.08 (0.01)	-0.09 (0.01)
Δ Predicted log average transactions between t-1 and t \times Permanent Prod.	-0.12 (0.05)	-0.08 (0.05)	-0.01 (0.09)	0.04 (0.09)
Permanent Prod.	0.11 (0.01)		0.10 (0.01)	
Checker fixed effects	No	Yes	No	Yes
Observations	2,045,091	2,045,091	2,045,091	2,045,091

Note: Standard errors clustered at the store \times day of week \times time of day level are in parentheses. Entries are estimates of equations (7) and (8). The units of observation are checker \times ten minute cells. “Exit at t” means that in a given day the checker is last observed working in period t. “Entry at t” means that in a given day the checker is first observed in period t. To compute predicted transactions for worker i , we compute the average number of items scanned in store \times day of week \times time of day cells, excluding any observations where worker i was on duty (to prevent a mechanical correlation between a worker’s permanent productivity and predicted demand).

Table 4: Relationship between lagged changes in sales volume and entry and exit probabilities of personnel

	Entry at t (1)	Entry at t (2)	Exit at t (3)	Exit at t (4)
Δ log items sold between t-2 and t-1	-0.029 (0.001)	-0.031 (0.001)	0.005 (0.001)	0.001 (0.001)
Δ log items sold between t-2 and t-1 \times Permanent Prod.	-0.051 (0.009)	-0.058 (0.009)	0.015 (0.008)	0.009 (0.008)
Permanent Prod.	0.113 (0.003)		0.106 (0.003)	
Checker fixed effects	No	Yes	No	Yes
Observations	2,062,431	2,062,431	2,062,431	2,062,431

Notes: Standard errors clustered at the store \times calendar date \times time of day level are in parentheses. The units of observation are checker \times ten minute cells. “Exit at t” means that in a given day the checker is last observed working in period t. “Entry at t” means that in a given day the checker is first observed in period t.

Table 6: The effect of changes of average co-worker permanent productivity on reference person current productivity; Models by spatial orientation and proximity

	(1)	(2)	(3)	(4)	(5)
Δ Co-worker permanent productivity behind	0.233 (0.019)				
Δ Co-worker permanent productivity in front	0.007 (0.018)				
Δ Co-worker permanent productivity behind & closer			0.162 (0.016)		
Δ Co-worker permanent productivity in front & closer			0.016 (0.015)		
Δ Co-worker permanent productivity behind & farther			0.100 (0.018)		
Δ Co-worker permanent productivity in front & farther			0.003 (0.018)		
Δ Worker behind		0.040 (0.003)			0.040 (0.003)
Δ Worker in front		-0.033 (0.003)			-0.033 (0.003)
Δ Worker behind & closer				0.025 (0.002)	
Δ Worker in front & closer				-0.042 (0.002)	
Δ Worker behind & farther				0.007 (0.002)	
Δ Worker in front & farther				0.0001 (0.002)	
Δ Worker behind \times Average FE					0.159 (0.039)
Δ Worker in front \times Average FE					-0.045 (0.037)
Observations	1,660,312	1,734,164	1,501,555	1,734,164	1,734,164

Notes: Standard errors clustered by store \times calendar date \times checker are in parentheses. The units of observation are checker \times ten minute cells. The dependent variable is the change in the log productivity of a checker across ten minute periods. See note for Table 2 for a description of the dependent variable. All models include controls for the change in the number of workers on duty in a ten minute interval, and 588 ten minute time interval by day of week dummies. “Behind” refers to workers that are facing worker i . “In front” refers to workers that i is facing. “Closer” refers to workers that are one or two positions away from i . “Farther” denotes workers that are three or four positions away from i . The permanent productivity averages are taken over the indicated sub-groups. For example, Δ co-worker permanent productivity in front & closer denotes the change in the average permanent productivity of co-workers that worker i is facing are who are one or two registers away from i . In this example, if there are no workers that are positioned one or two positions in front of i in both $t-1$ and t , then this change is coded as 0. The variable “ Δ Worker” behind denotes the change in the presence of a worker for whom i is in their line-of-sight. In column (5), “Average FE” denotes the average permanent productivity of the workers who have entered or exited the checkout stands.

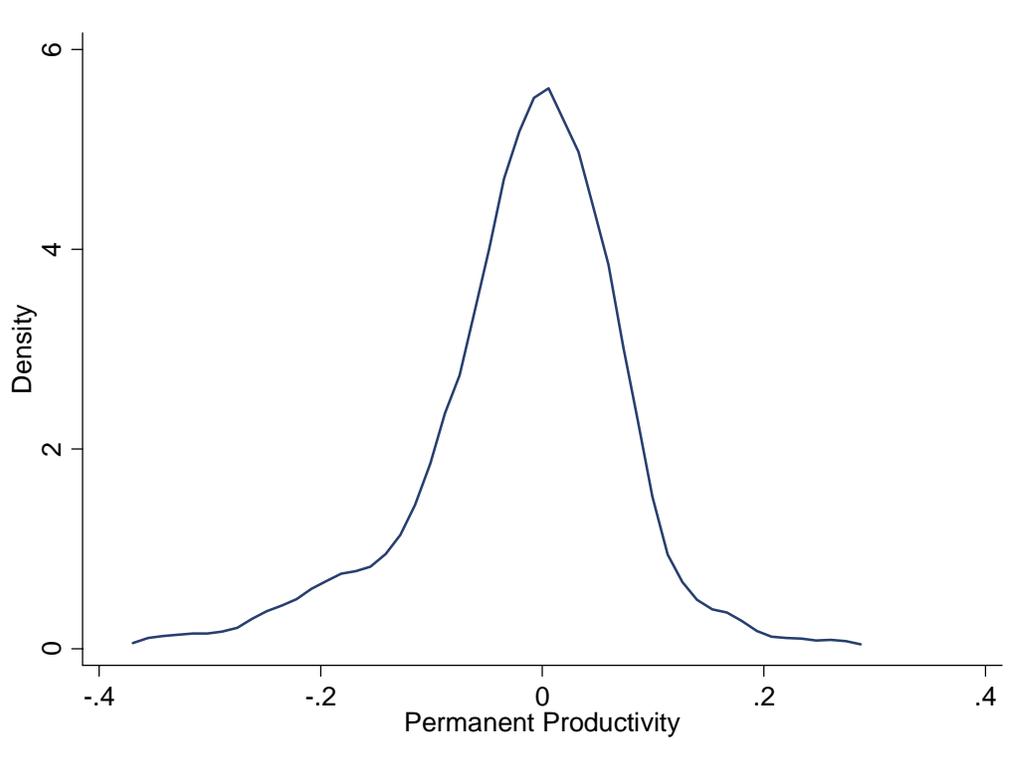
P-value for H_0 : Δ Co-worker permanent productivity behind (closer) = Δ Co-worker permanent productivity behind (farther) is 0.028

Table 7: The effect of changes of average co-worker permanent productivity on reference person current productivity; Models by previous exposure to co-workers and spatial orientation

	(1)	(2)	(3)
(I) Δ Co-worker permanent prod: low exposure	0.013 (0.012)	(I) Δ Co-worker permanent prod: behind & low exposure	0.018 (0.056)
(II) Δ Co-worker permanent prod: medium exposure	0.084 (0.014)	(II) Δ Co-worker permanent prod: behind & medium exposure	0.116 (0.023)
(III) Δ Co-worker permanent prod: high exposure	0.075 (0.017)	(III) Δ Co-worker permanent prod: behind & high exposure	0.122 (0.025)
		(IV) Δ Co-worker permanent prod: in front & low exposure	0.002 (0.016)
		(V) Δ Co-worker permanent prod: in front & medium exposure	-0.007 (0.006)
		(VI) Δ Co-worker permanent prod: in front & high exposure	0.000 (0.005)
p-value: $H_0: (I) = (II)$	0.000	p-value: $H_0: (I) = (II)$	0.105
$H_0: (I) = (III)$	0.003	$H_0: (I) = (III)$	0.090
$H_0: (II) = (III)$	0.655	$H_0: (II) = (III)$	0.852
		$H_0: (IV) = (V)$	0.562
		$H_0: (IV) = (VI)$	0.891
		$H_0: (V) = (VI)$	0.328
Observations	1,659,450	Observations	1,659,450

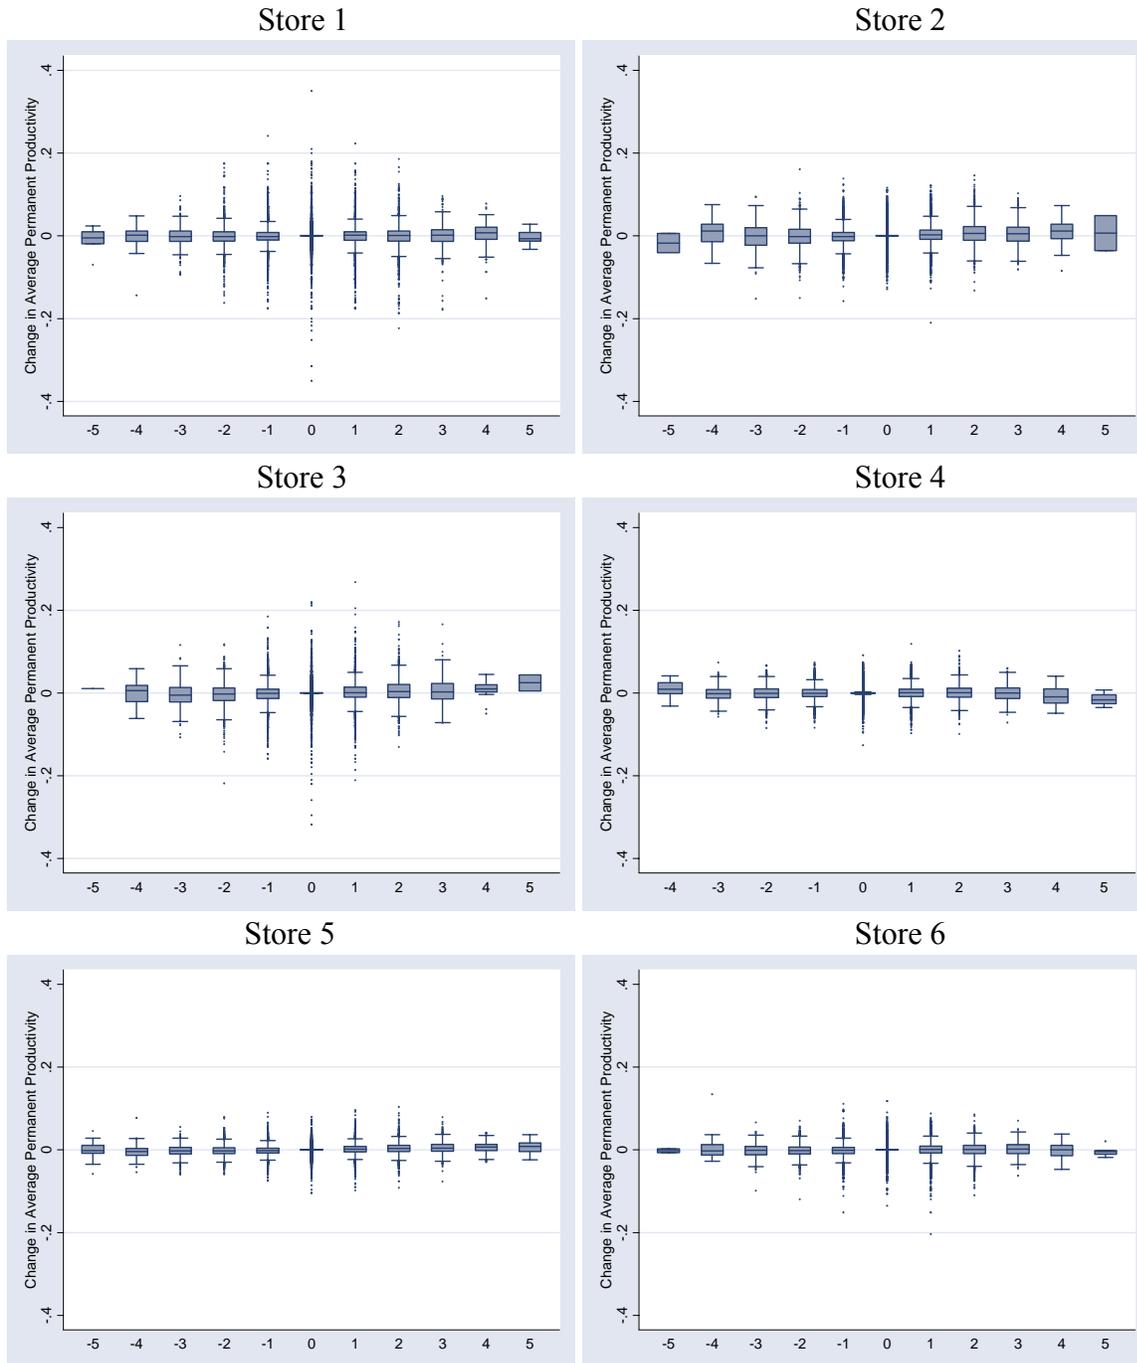
Notes: Standard errors clustered by store \times calendar date \times checker are in parentheses. The units of observation are checker \times ten minute cells. The dependent variable is the change in the log productivity of a checker across ten minute periods. See note to Table 2 for a description of the dependent variable. The first two months of the sample in each store are dropped because we require an initial period to estimate the degree of schedule overlap. In the case of worker i , “Low exposure” denotes the set of workers who have previously overlapped between 0 and 5 percent of i ’s schedule. “Medium exposure” denotes workers who have overlapped between 5 and 20 percent of i ’s schedule. “High exposure” denotes workers who have overlapped between 20 and 100 percent of i ’s schedule. “Behind” refers to workers that are facing worker i . “In front” refers to workers that are facing i . The permanent productivity averages are taken over the indicated sub-groups. For example, variable (II) in column (2) is the change in the productivity of co-workers that are positioned behind the reference worker and have medium previous exposure to worker i . In this example, if there are no workers that are positioned behind worker i who have medium previous exposure to i in $t-1$ and t , this change is coded as 0. If in $t-1$ there is no worker with medium previous exposure to i that is positioned behind i , but in period t there is, this variable is coded as missing. However, to prevent the sample size from dropping substantially due to the introduction of missing values, we code missing values as 0, and then include indicator variables for whether the value in each of the explanatory variables of interest is missing. All models include controls for the change in the number of workers on duty in a store over ten minute intervals, and 588 ten minute time interval by day of week dummies.

Figure 1: Distribution of Workers Permanent Productivity



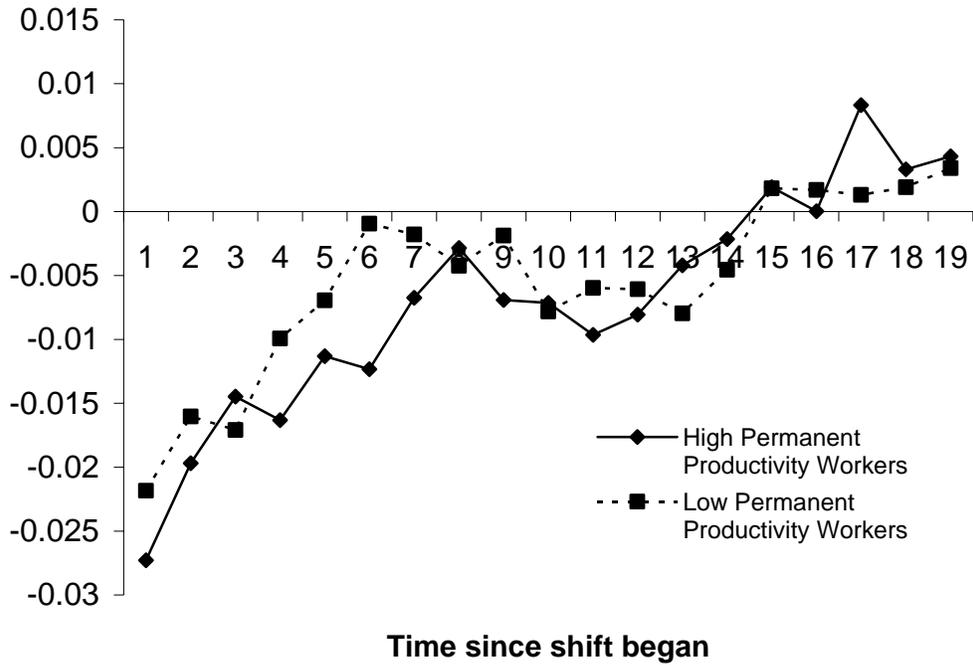
Note: This figure shows the distribution of worker permanent productivity (the parameter θ_i), obtained by estimating equation (4). Specifically, we regress log productivity (items per second worked) of worker i at time t in calendar date c in store s , on a vector of worker fixed effects, dummies for the presence of co-workers on the shift; a vector of controls that includes number of active workers at a moment in time (in ten minute intervals) as well as dummies for the register where a worker is stationed; a set of dummies for each day of week \times hour of the day \times store combination; and a set dummies for each calendar-date \times store combination. The Figure shows the kernel density estimate of the fixed-effects.

Figure 2: Relationship between net changes in personnel on duty and changes in the average permanent productivity of workers



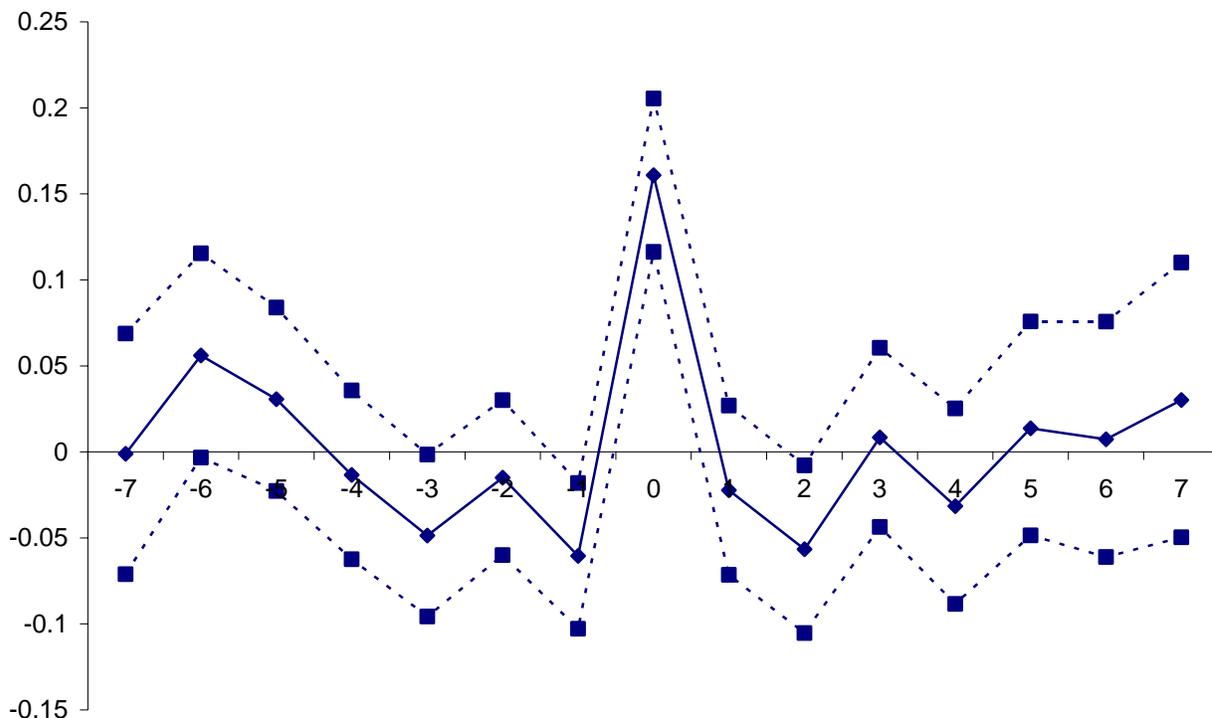
Note: The x-axis is the net change in the number of personnel across ten minute periods. The y-axis is the corresponding change in the average permanent productivity of workers on duty. The units of observation are store \times ten minute period cells.

Figure 3: Productivity by Time-in-Shift



Note: Each unit of time is ten minutes. To construct this chart, the log of productivity is regressed on time of day dummies (in ten minute intervals), checker \times calendar date fixed effects, and indicators for each time in shift, using the sample of Table 2. This regression was run separately for above average and below average permanent productivity workers. The first nineteen estimated time-in-shift indicators, by high/low permanent productivity type, are plotted above. Because these estimates come from a model with checker \times calendar date fixed-effects, the points represent log productivity relative to each checker's mean productivity during a shift. For example, a point at -0.02 on the y-axis means that at the time the worker is going approximately 2% slower than her average speed in the shift.

Figure 4: Distributed lags and leads for the effect of changes of average co-worker permanent productivity on reference person current productivity

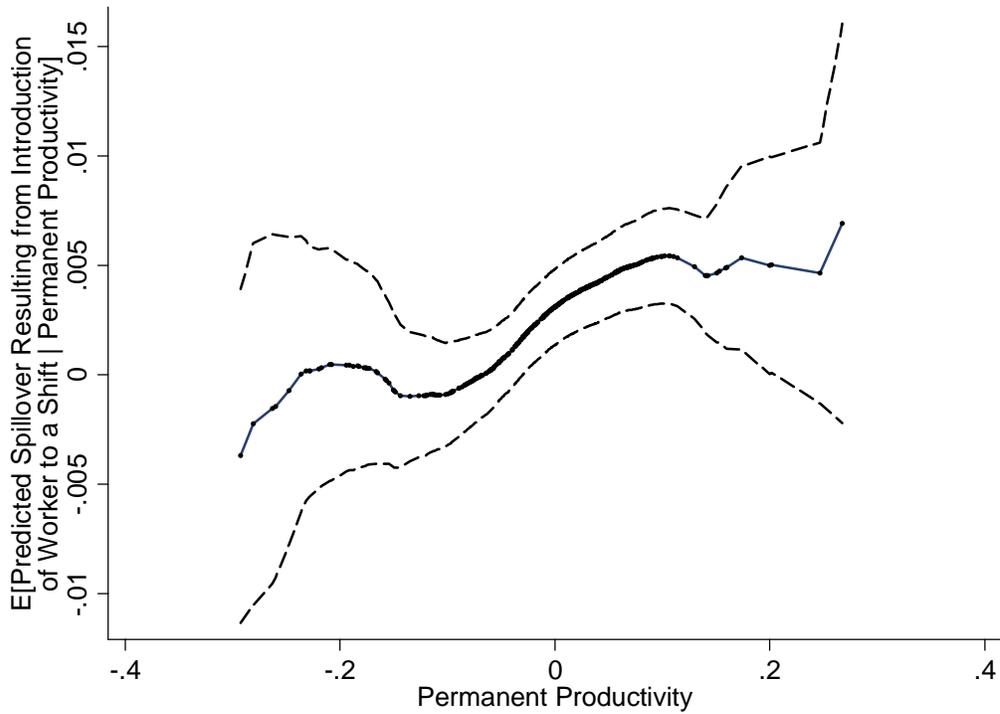


Note: The x-axis denotes the distributed lag and lead coefficients. For example, “-4” denotes the fourth lag, “0” is the contemporaneous change, and “4” is the fourth period lead. The unit is 10 minutes. Points are the estimated β 's from the following model:

$$\begin{aligned} \Delta y_{itcs} = & \beta_{-7} \Delta \bar{\theta}_{-i(t-7)cs} + \beta_{-6} \Delta \bar{\theta}_{-i(t-6)cs} + \beta_{-5} \Delta \bar{\theta}_{-i(t-5)cs} + \beta_{-4} \Delta \bar{\theta}_{-i(t-4)cs} + \beta_{-3} \Delta \bar{\theta}_{-i(t-3)cs} + \beta_{-2} \Delta \bar{\theta}_{-i(t-2)cs} \\ & + \beta_{-1} \Delta \bar{\theta}_{-i(t-1)cs} + \beta_0 \Delta \bar{\theta}_{-i(t)cs} + \beta_1 \Delta \bar{\theta}_{-i(t+1)cs} + \beta_2 \Delta \bar{\theta}_{-i(t+2)cs} + \beta_3 \Delta \bar{\theta}_{-i(t+3)cs} + \beta_4 \Delta \bar{\theta}_{-i(t+4)cs} + \beta_5 \Delta \bar{\theta}_{-i(t+5)cs} \\ & + \beta_6 \Delta \bar{\theta}_{-i(t+6)cs} + \beta_7 \Delta \bar{\theta}_{-i(t+7)cs} + \zeta \mathbf{M} + e_{itcs}, \end{aligned}$$

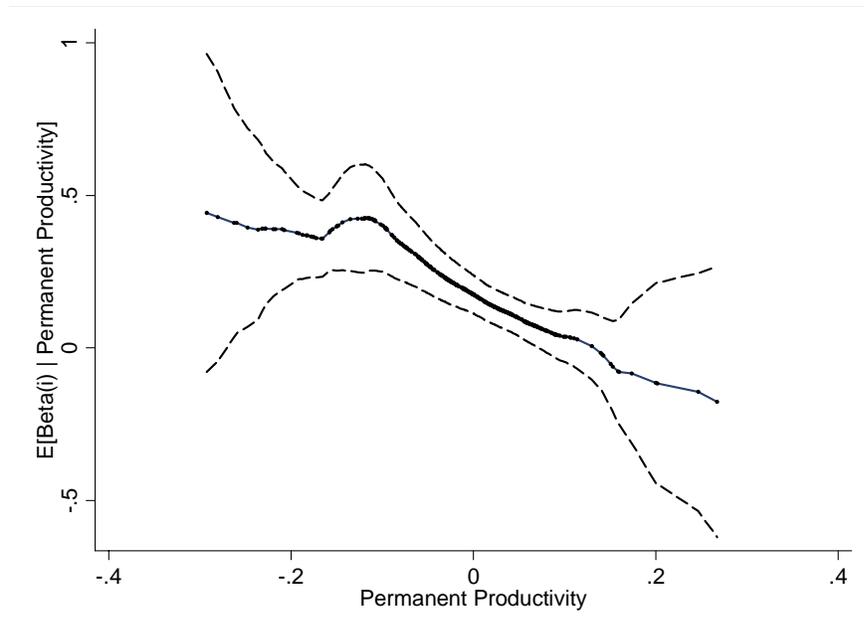
using the sample of Table 2. Dotted line is the 95% confidence interval. In order to prevent the sample size from falling substantially with the introduction of the lagged and lead changes in average co-worker productivity, we assign lags and leads with missing values “0” and include in the model fifteen dummy variables—one for each variable—taking on the value of “1” if its corresponding lag or lead contained a missing value. These indicators are represented by the matrix \mathbf{M} in the above equation.

Figure 5: The relationship between the permanent productivity of worker i and the effect that the presence of worker i has on the productivity of her co-workers



Notes: This figure plots an estimate of the influence of an individual worker on her co-workers (the parameter π_i in equation 5), against that worker's permanent productivity (the parameter θ_i in equation 4), using a local-linear smoother. The lower and upper lines are the 95% confidence interval.

Figure 6: The relationship between individual permanent productivity and worker specific spillover effect



Notes: This figure plots estimates of a worker specific spillover effects (the parameter β_i in equation 10) against that worker permanent productivity (the parameter θ_i in equation 4). The lower and upper lines are the 95% confidence interval.