

Peers at Work

By ALEXANDRE MAS AND ENRICO MORETTI*

We study peer effects in the workplace. Specifically, we investigate whether, how, and why the productivity of a worker depends on the productivity of coworkers in the same team. Using high-frequency data on worker productivity from a large supermarket chain, we find strong evidence of positive productivity spillovers from the introduction of highly productive personnel into a shift. Worker effort is positively related to the productivity of workers who see him, but not workers who do not see him. Additionally, workers respond more to the presence of coworkers with whom they frequently interact. We conclude that social pressure can partially internalize free-riding externalities that are built into many workplaces. (JEL J24, L81, M54)

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 coauthored academic research share this characteristic to some degree. 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 devotes to 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 face sanctions from her peer. Because of this possibility, it could be optimal for this person to do her “fair share,” and work harder, in order to reduce the productivity gap with her more productive partner. Edward Kandel and Edward 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 considerations mitigate the deleterious effects of free-riding in real workplaces that are prone to externalities?¹

* Mas: Haas School of Business, 545 Student Services Building #1900, University of California, Berkeley, CA 94720-1900 (e-mail: amas@haas.berkeley.edu); Moretti: Department of Economics, 637 Evans Hall #3880, University of California, Berkeley, CA 94720-3880 (e-mail: moretti@econ.berkeley.edu). We thank Oriana Bandiera, Stefano DellaVigna, Andrea Ichino, Alan Krueger, Lars Lefgren, Bentley Macleod, Daniel Parent, Jacob Vigdor, Giulio Zanella, and seminar participants at Berkeley, Brigham Young University, Case Western Reserve University, Columbia University, University of Chicago GSB, Duke University, Boston Federal Reserve Bank, Georgetown University, IZA/SOLE Transatlantic Meetings, IZA Workshop on Institutions and the Employment Relationship, Harvard University, Kellogg School of Management, McGill University, MIT, NBER Summer Institute, 2006 SITE Conference, Stanford University, University of Torino, UCLA, UC Irvine, UC Santa Barbara, University of Virginia, Yale University, and the Wharton School for useful comments. Sanny Liao and Ashley Langer provided valuable research assistance. This research was partially financed by a Labor and Employment Research Fund grant.

¹ The question of the role of social relations in the workplace in motivating effort has been an important theme of the literature on organizations at least as far back as Chester Barnard (1938). A different but related question is whether workers display social preferences in team production settings. In an interesting recent study, Oriana Bandiera, Iwan Barankay, and Imran Rasul (2005) find that workers internalize the negative externality their effort imposes on others

In this paper, we empirically investigate how workers influence each other in the context of a retail firm. We explore how and why the productivity of a worker varies as a function of the productivity of her coworkers in a group production process that is particularly prone to free-riding. Our analysis centers on two questions. First we ask how the introduction of a high-productivity worker affects the productivity of her coworkers. As indicated, this relationship could go in any direction, depending on whether free-riding or positive spillovers dominate. Having found evidence of positive productivity spillovers, we then investigate their underlying mechanisms. We seek to distinguish between specific forms of peer effects that could be at work, including social pressure and prosocial behavior. 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 study is among the first to get inside the black box of productivity spillovers and to 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. Monetary incentives alone may not be enough to explain what motivates workers 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 environment 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 checkers who are working at that time in the same store. Over the course of a given day, the composition of the group of coworkers varies, because worker shifts do not perfectly overlap. Therefore, for any given worker, the mix of her coworkers

under relative incentives. Another related phenomenon involves workers colluding with each other in defiance of management. This collusion, often seen in manufacturing, can involve restricting output. See Donald Roy (1954) for a sociological account of this kind of collusion in a piece-rate machine shop.

² Alfred Marshall (1890) is the first to hypothesize that on-the-job interactions may generate positive externalities across workers. Growth and urban economists have long proposed models where spillovers are the determinant of growth (Robert E. Lucas 1988), and empirical studies have tested the existence of spillovers (for example, Moretti 2004a, b). Recent papers that investigate the role of moral hazard in teamwork environments include Martin Gaynor, James Rebitzer, and Lowell Taylor (2004); William Encinosa, Gaynor, and Rebitzer (2007); and Marc Knez and Duncan Simester (2001). Also related are papers on diversity and performance (Barton Hamilton, Jack Nickerson, and Hideo Owan 2005; Jonathan Leonard and David Levine 2006), friendship networks in the workplace (Bandiera, Barankay, and Rasul 2007), and complementarities among coworkers in management teams (Rachel Hayes, Paul Oyer, and Scott Schaefer 2006).

³ Because we need to measure productivity reliably, a study of this kind must focus on a single occupation or industry where the data on effort are reliably measured, and where the institutional features of the work environment are understood. In this respect, our approach is similar in spirit and design to other empirical studies in personnel economics, for example Casey Ichniowski, Kathryn Shaw, and Giovanna Prennushi (1997), Harry Paarsch and Bruce Shearer (1999), Andrea Ichino and Giovanni Maggi (2000), Lazear (2000), Hamilton, Nickerson, and Owan (2003), and Bandiera, Barankay, and Rasul (2005).

changes throughout the day. The firm gives substantial scheduling flexibility to the workers, and management does not have the ability 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 shift changes should largely be predetermined relative to transitory shocks to productivity. Because of this scheduling policy, the timing of within-day changes in the average ability of coworkers can be considered plausibly exogenous.

We find strong evidence of productivity spillovers. Substituting a worker with below average permanent productivity with a worker with above average permanent productivity is associated with a 1 percent *increase* in the effort of other workers on the same shift. The finding of a positive spillover suggests that positive peer effects dominate free-riding. The magnitudes of our estimates are in line with related laboratory experiments.⁴ We consider a variety of empirical tests examining whether the timing of changes in the average ability of coworkers within a day is indeed exogenous. The tests confirm that the patterns we see in the data are likely not coming about from selective timing of checker entry and exit.

The magnitude of the spillover effect appears to vary dramatically depending on the skill level of the relevant worker. Low-productivity workers are far more responsive to the composition of coworkers than high-productivity workers. Interestingly, while low-productivity workers benefit from the presence of more capable workers, the productivity of high-skill workers is not hurt by the presence of low-skill coworkers. This finding is important because it implies that the optimal mix of workers in a given shift is the one that maximizes skill diversity. By rearranging the mix of workers to maximize skill variance in each shift, this supermarket could produce the same amount of sales with fewer hours worked each year.

Why are there positive spillovers? Two explanations relevant in our context are social pressure and prosocial behavior. We define social pressure as encompassing cases where workers experience disutility if they are *observed* behaving selfishly by their peers. This utility loss may be due to formal or informal sanctions by coworkers, or shame. By prosocial behavior we mean a broad class of altruistic behavior whereby a worker experiences disutility if she is acting noncooperatively, even if no one notices. A novel feature of our setting is that we can identify the spatial orientation of workers in this production process based on their register assignment. To distinguish between the underlying mechanisms, we use this information to estimate models where the effect of coworkers is allowed to vary depending on whether coworkers can monitor each other.

We find that when more productive workers arrive at shifts, they induce a productivity increase only in workers who are in their line of vision. We find no effect on workers who are not in their line of vision. Moreover, the effect appears to decline with distance. It is stronger for workers who are in the line of vision and close to the more productive peer than those who are farther away. There is also evidence that workers free-ride when coworkers cannot observe them. We interpret these findings as consistent with the notion that productivity spillovers are due to social pressure. These findings appear less consistent with altruistic behavior.

As an additional test, we consider whether the magnitude of the spillover depends on frequency of interaction in the workplace. If workers rarely interact, they may not be receptive to social pressure due to the limited scope for sanctions. Consistent with this hypothesis we find that introducing a high-productivity worker into a shift is associated with greater increases in

⁴ In the experiment that most resembles the setting in our study, Armin Falk and Ichino (2006) find that a 10 percent increase in peers' output results in a 1.4 percent increase in a given individual's effort. Falk, Urs Fischbacher, and Simon Gächter (2003) study contributions to a public good and find evidence of social interactions. Ernst 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.

incumbent productivity when the entering worker and the incumbent's schedule have high overlap than when they coincide only infrequently.

Overall, the evidence assembled suggests that social pressure can partially internalize externalities that are built into many workplaces. 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 occurs only when it can be noticed by their peers, and when numerous future interactions can be expected to occur. Our results demonstrate that social considerations can motivate workers, and even offset limited monetary incentives. This conclusion is supportive of Kandel and Lazear's (1992) theory, and is 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 I presents a simple model. In Section II we describe our measure of productivity. In Sections III, IV, and V we describe our empirical specification and our empirical findings. Section VI concludes.

I. Conceptual Framework

In many jobs, employers can observe total output, but cannot observe the exact contribution provided by each worker to the production of total output. As indicated in the introduction, this feature of the workplace is common in most clerical occupations, many manufacturing jobs, construction, agriculture, and retail, especially when the number of employees working on a task is large. Consider, for example, a sales team writing a marketing presentation. 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 exact individual contributions. It is more likely that the employer observes a noisy signal of each worker's effort. A similar story could be told about carpenters building a house, or even coauthored academic research, where any reader can observe the quality of a paper but the contribution of each coauthor 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. Managers are supposed to supervise many workers in the store who are not cashiers, and they must also perform many other tasks. Moreover, our assumption requires only that individual productivity is observed with some noise, however small.⁵

Our goal in this section is to investigate how workers in a team react to an exogenous change in the productivity of their coworkers 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 coworker to a shift when the worker's marginal benefit of effort declines as the effort of coworkers increases. Peer pressure can potentially mitigate this externality. We present a specific example of how the marginal utility of effort can depend on coworker effort. The framework described here is intentionally kept very simple, and we note

⁵ Although the firm has in theory access to the same data that we use in this study, it has never used them for this purpose. Indeed, this is one of the reasons we were allowed access to the data.

that there are certainly other models, in a similar spirit, that will also lead to this dependence. Kandel and Lazear (1992) develop a framework that 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, $\bar{y} = (1/N) \sum_{j=1}^N y_j$. In the context of our application, we can think of \bar{y} as the (inverse of) the length of the customer lines. Following the revelation of worker i 's noisy signal of output, management's best guess of a worker's productivity given the signal is

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

where $b = \text{var}(y)/(\text{var}(y) + \text{var}(u))$ and both variances are assumed to be known. Equation (1) simply says that management imperfectly observes the effort provided by each worker, but perfectly observes the length of the lines, and it combines these two pieces of information to infer who is working hard and who is not.⁶ In our data, workers are unionized and compensation is a fixed hourly payment, but workers can be fired if they are perceived by management as underperforming. We assume that the probability of keeping the job, Q_i , is an increasing function of management's best guess of a worker's productivity, and that if a worker is fired, she receives no utility.⁷

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

$$\max [Q_i U(w)] - C_i(e_i),$$

where w is the wage, $Q_i = q(E(y_i | z_i))$, with $q' > 0$; $q'' < 0$; and $C' > 0$; $C'' > 0$. Given that Q_i depends implicitly on effort, the first-order conditions are

$$U(w) q' [f' [b + [1 - b] [1/N]]] = C'_i.$$

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.⁸

In this paper we are interested in what happens to the effort of a worker when the productivity of her coworkers 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 coworker with high cost of effort with an otherwise identical coworker with low cost of effort. Assume for example that $C_i(e_i) = (1/\theta_i)e_i^2$, where $1/\theta_i$ is the individual specific cost of effort. In other words, the parameter

⁶ Equation (1) is easily derived from the formula of a hypothetical regression of y_i on z_i , where the OLS intercept is $\bar{y} - b \bar{z}$ and the OLS slope is $\text{cov}(y_i, z_i)/\text{var}(z_i) = b$. Obviously this regression cannot be run by the employer, but the intercept and the slope parameters are known under our assumptions. If the signal has no noise, $\text{var}(u) = 0$ and $E(y_i | z_i)$ equals true productivity: $E(y_i | z_i) = y_i$.

⁷ We have modeled a worker's wage as fixed, and the probability of being fired as a function of productivity. This approach is consistent with our empirical application. It is easy to see that our results generalize to the case where the wage is not fixed. In this case, the problem can be recast in terms of the relationship between productivity and wage, yielding the same conclusions.

⁸ The efficient level of effort is the vector (e_1, e_2, \dots, e_N) , which maximizes total surplus, $\sum_i U - C_i$. Obviously, if effort were observable, the efficient level would be easily achievable.

θ_i characterizes the worker i 's typical productivity, or her ability. High-productivity types have a high θ_i (or low cost of effort); low-productivity types have a low θ_i (or high cost of effort).

The derivative of the optimal effort of worker i with respect to the productivity of coworker j in the same shift is

$$(2) \quad \frac{de_i^*}{d\theta_j} = \frac{U(w)[b + [1 - b][1/N]]f'q''[[1 - b][1/N][(f')\partial e_j^*/\partial \theta_j]]}{2\theta_i - U(w)[q''f'^2 [b + [1 - b][1/N]] + q'f''][b + [1 - b][1/N]]} < 0.$$

Expression (2) is negative. The intuition is that, in the absence of any social considerations, an exogenous increase in coworker productivity—caused by a decline in their cost of effort—results in more free-riding and, therefore, a decline in i 's effort.

Case (b). The prediction in equation (2) crucially depends on the assumption that there are no mechanisms to internalize the externality generated by free-riding. Kandel and Lazear (1992), Steffen Huck, Dorothea Kubler, and Jörgen Weibull (2002), John M. Barron and Kathy Paulson Gjerde (1997), Armin Falk and Ichino (2006), and Falk, Fehr, and Fischbacher (2001) propose alternatives to this assumption. In particular, Kandel 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.⁹ Following their notation, it is easy to incorporate peer pressure in our setting by assuming that workers maximize:

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

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 focal worker's effort. The functional form of $P(\cdot)$ is a priori undetermined. In this paper we seek to describe a situation where peer pressure is parameterized as a function of the average of coworker 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 mitigate the free-riding problem. 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,

$$(3) \quad de_i^*/d\theta_j > 0.$$

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

⁹ 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.

II. The Setting

Most existing studies of productivity rely on low-frequency aggregate output measures, typically defined at the industry or firm level. Such measures are not well suited to empirically identify spillovers. There is an emergent literature that investigates productivity in teams that use either individual or team-level measures of productivity, for example Bandiera, Barankay, and Rasul (2005), Hamilton, Nickerson, and Owan (2003), and Leonard and Levine (2006). While the data used in these studies are eminently sensible to answer the particular questions they pose, our data are particularly well suited to investigate the question of productivity spillovers. 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.¹⁰

There are several reasons that these data are attractive for our purposes. First, we have a close to continuous time measure of productivity, making it possible to identify *instantaneous* changes in individual productivity. Second, we know not only who is working at any moment in time, so that we can identify the production group, but also the exact contribution of each member's output in the group. Third, we have information on the scope of reciprocal monitoring based on worker register assignments, which is useful when examining the underlying mechanisms. On the other hand, this measure is not without flaws. While it captures very accurately a cashier's speed, it completely abstracts from the quality of her service (friendliness, care in handling items, etc.), and from other relevant measures of performance, for example, absenteeism.

In this supermarket chain, workers are unionized and compensation is a fixed hourly payment. Because of union rules, checkers can only work at the registers, as opposed to, for example, stocking shelves. Discussions with management indicate that the firm gives substantial scheduling flexibility to the workers. Managers have no role in determining which workers are assigned to particular shifts. Rather, managers determine the number of workers in a given shift. They then provide the schedules to the employees on a biweekly 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.

III. Econometric Specification

We begin by specifying a model describing productivity determinants in an environment where spillovers may operate. We assume that productivity of worker i , working in store s , in calendar date c , and time t (where t is measured in ten-minute intervals) can be written as

¹⁰ Specifically, for each worker on the shift, we sum the number of items that a 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 given to the customer. We exclude any ten-minute period where there is only one checker on duty; and we include only observations where a worker is at the same register for at least two consecutive ten-minute periods, as we will be estimating first-differences models and we wish to hold the registers where workers are stationed constant. An alternative measure of productivity is the waiting time between customers. Unfortunately we cannot compute this measure because, while we know the number of seconds elapsed in the transaction, we know the end time of the transactions only in intervals of minutes (not seconds). As a result, any gap between customers less than a minute long is impossible to measure.

$$(4) \quad y_{itcs} = \theta_i + P_{itcs}(\theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_k) + \pi N_{itcs} + \tau R_{itcs} + \gamma_{itds} + e_{itcs},$$

where θ_i denotes worker fixed effects; $P_{itcs}(\theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_k)$ is a peer pressure function that affects an individual's productivity depending on her coworker ability-type, as explained in Section I; N_{itcs} denotes the number of workers on duty at the given time of day; R_{itcs} is an indicator for checker i 's register location; and γ_{itds} is a vector of interactions for all possible combinations of hour of the day, day of the week, and store. As in the theoretical model in Section I, we interpret the parameter θ as a measure of a worker's permanent productivity, or ability. In particular, in the model in Section I, we assume that the parameter θ is the inverse of her cost of effort. For this reason, throughout the paper, effort and productivity are synonymous. Workers with a low cost of effort (high θ) are on average more productive than workers with a low cost of effort (low θ).

We consider a parameterization of the peer pressure function whereby it consists of the average of coworker productivity, so that $P_{itcs}(\theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_k) = \bar{\theta}_{-itcs} = 1/n \sum_{j \neq i} \theta_{jics}$. $\bar{\theta}_{-itcs}$ is the average permanent productivity of all the coworkers who are active in period t , where $-i$ denotes that the average of the permanent productivity component is taken over all workers in store s , working at time t , and calendar date c , excluding worker i .¹¹ From (4), we then obtain

$$(5) \quad y_{itcs} = \theta_i + \beta \bar{\theta}_{-itcs} + \pi N_{itcs} + \tau R_{itcs} + \gamma_{itds} + e_{itcs}.$$

Taking first differences of (5) gives our baseline estimating equation:

$$(6) \quad \Delta y_{itcs} = \alpha + \beta \Delta \bar{\theta}_{-itcs} + \pi \Delta N_{itcs} + e_{itcs}.$$

The coefficient of interest is β . If there are no spillovers, a checker who is exposed to speedier workers may free-ride, thus lowering her effort. In this case β should be negative (see equation (2)). In the presence of spillovers, a worker will increase her effort when exposed to faster peers. If spillovers are large enough, β should be positive (see equation (3)). In the absence of spillovers or free riding, β should be zero.

We estimate (6) in two steps. In the first step we seek to estimate the θ_i terms. To estimate these terms it is necessary to take into account the fact that an individual's productivity may be affected by coworker composition, as suggested by econometric model (5). The purpose of the first step is therefore to estimate θ_i in a model that is consistent with (5). To accomplish this, we estimate

$$(7) \quad y_{itcs} = \theta_i + \mathbf{M}' \boldsymbol{\varphi}_{Ci} + \pi N_{itcs} + \tau R_{itcs} + \gamma_{itds} + e_{itcs}.$$

The term $\boldsymbol{\varphi}_{Ci}$ is a vector of all possible interactions from the set $\{C_{il}, \dots, C_{ik}\}$, where

$$C_{il} = \begin{cases} 1 & \text{if worker } l \text{ is on duty at time } t, \text{ on calendar date } c, \text{ and store } s, \\ 0 & \text{if } i = l, \\ 0 & \text{if worker } l \text{ is not on duty at time } t, \text{ on calendar date } c, \text{ and store } s. \end{cases}$$

¹¹ While common in the literature on peer effects, this linearity 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.

In words, the term \wp_{Ci} is the set of dummy variables, one for every possible combination of coworker composition. For example, there is a dummy variable in \wp_{Ci} for every instance when a checker is working with workers 2, 3, and 4, and another dummy variable for every instance when a checker is working with checkers 2, 9, and 12. Note that, in a given shift, all checkers are working with a different set of coworkers. For example, if checkers 1, 2, and 3 are working together, 2 and 3 are checker 1's coworkers, while checkers 1 and 2 are checker 3's coworkers. The vector \mathbf{M} contains parameters, once for every possible member of \wp_{Ci} .¹² Equation (7) is consistent with (5) because the peer pressure function is absorbed by \wp_{Ci} . The term \wp_{Ci} accounts for the focal worker's productivity response to working in a shift with a specific composition of coworkers. This term will allow us to estimate the worker fixed effects purged of possible influences of arbitrary social interactions, like those we are seeking to estimate.

We use our estimated fixed-effect estimates to construct a measure of average coworker productivity in every shift, denoted $\bar{\theta}_{-iics}$. In the second step, we use $\bar{\theta}_{-iics}$ to estimate equation (6). Because $\bar{\theta}_{-iics}$ is derived from estimated quantities, we adjust standard errors in (6) to take into account the sampling variability of this term. We implement a simple, transparent version of the Bayesian parametric bootstrap in which we use the estimated variability from simulated draws of the estimated fixed effects to adjust the standard errors. This approach is very similar in practice to computing standard errors with multiple imputed data (Donald Rubin 1987).¹³ We describe this procedure in detail in the Appendix. An important feature of this procedure is that it permits for arbitrary covariances in the error term between every pair of time periods for a given checker in a given day. In this way we are allowing for possible serial correlation.¹⁴

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 coworkers changes throughout the day depending on who enters and who exits. Variation in personnel composition comes primarily 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 behavioral response of workers to their peers. 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 shift changes is predetermined relative to transitory shocks to productivity. In Section IVD we present a series of empirical tests intended to verify the validity of this assumption.

The parameter β represents the effect of *permanent* coworkers' productivity on worker i 's current productivity. An alternative model that we could specify would be to have the peer pressure function take the form $P_{iics}(y_{1cts}, \dots, y_{i-1cts}, y_{i+1cts}, \dots, y_{kcts})$. That is, the focal worker's effort would depend on the *contemporaneous* effort of coworkers, rather than their *permanent* productivity.¹⁵

¹² In practice, we include only dummies for observed combinations of workers. Estimating individual fixed effects and coworker composition fixed effects is similar in spirit to the problem of estimating individual fixed effects and firm fixed effects in longitudinal data that contain firm and worker identifiers. See, for example, John Abowd, Francis Kramarz, and David Margolis (1999). Equation (7) is separately estimated for each of the six stores.

¹³ An alternative version would be to block bootstrap the entire two-step procedure. This approach is prohibitive computationally. Analyses conducted employing a block bootstrap on a subsample of the data yield standard errors that are almost identical to those employed when employing the parametric version.

¹⁴ Because the estimated individual fixed-effects are estimated, the parameter estimate of β might also be prone to attenuation bias. However, because the regressor of interest $\bar{\theta}_{-iics}$ is an average of these estimated fixed effects, the average of these errors will converge toward zero. Moreover, as our estimates of the fixed effects are extremely precise, attenuation bias should be negligible.

¹⁵ Charles F. Manski (1993, 2000), William Brock and Steven Durlauf (2001a, b, 2002) and Robert A. Moffitt (2001) discuss the differences between contextual effects (how the characteristics of the group affects its members) and

These two models have different interpretations. If the peer pressure function is assumed to depend on the permanent productivity of coworkers, the focal worker is assumed to know who is typically a fast worker and who is typically a slow worker, and adjusts her productivity accordingly, irrespective of their speed at a particular point in time. On the other hand, if the peer pressure function is assumed to depend on contemporaneous coworker productivity, workers are assumed to influence each other only through their point-in-time speed, irrespective of their permanent speed. The two models are *ex ante* equally plausible. Ultimately, we are unable to distinguish empirically between them.¹⁶ Of course, both forces could be at play at the same time. Therefore, a possible interpretation of our estimates is that they capture some combination of a true effect of permanent productivity and a true effect of contemporary coworker effort/productivity.

As far as estimation is concerned, a regression of changes in worker productivity on changes in coworker contemporaneous productivity would be more problematic than (6). First, it would be subject to the reflection problem (Manski 1993), in that the effort of worker i may affect effort of worker j , and vice versa. Second, productivity shocks that affect all workers at a given point in time can lead to a spurious relationship between worker and coworker productivity. In principle, one could estimate such a model by instrumental variables, using permanent productivity as an instrument for contemporaneous productivity. Note, however, that the parameter of interest in this approach would simply be the rescaled estimate of β in equation (6), with the rescaling factor equal to the first-stage coefficient from a regression of changes in coworker contemporaneous productivity on changes in coworker permanent productivity.

IV. Estimates of Productivity Spillovers

In this section, we begin by presenting our data and baseline estimates of how the effort of worker i depends on her coworkers' permanent productivity (Section IVA). We then investigate whether the estimated spillover is persistent over time (Section IVB). Third, we estimate more general models where we allow the spillover to vary across workers of different ability (Section IVC). Finally, we present a series of tests intended to assess possible threats to identification (Section IVD).

A. Data and Baseline Estimates

Our sample includes all the transactions that take place in six stores of a national supermarket chain for two years, for dates between 2003 and 2006.¹⁷ The stores are in the same metropolitan area of a state in the Western Census region. In total, we observe 394 cashiers. 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 approximately seven registers open with nonmanagerial workers, on average. Table 1 reports descriptive statistics of the sample.

endogenous effects (how the behavior of others in the group affects the individual), and how changes in group membership can aid identification.

¹⁶ To see why this is, consider the case where the true model is the one that has contemporaneous productivity of coworkers as a key independent variable, but the estimated model uses permanent productivity of coworkers instead (as in equation (6)). The estimated model would yield a positive regression coefficient on permanent productivity of coworkers because coworkers who have higher permanent productivity tend by definition to have higher contemporaneous productivity.

¹⁷ While we have only two years of data for each store, the dataset spans four calendar years because the starting dates differ by store.

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 general time patterns in shift productivity and the presence of coworkers. Figure 1 shows a distribution of the θ_i 's. The figure indicates that there is a wide variation in worker skill levels. The average 90–10 percentile differential in the estimated fixed effects across the six stores is 0.30, indicating that the top part of the productivity distribution is 30 percent more productive than the bottom part. One way to interpret this finding within the context of our model in Section I is that the cost of effort varies significantly across workers. Of course, dispersion in the estimated fixed effects will be overstated due to sampling variability. The estimated fixed effects are estimated quite precisely, however, and their standard deviation is hardly affected by adjusting for sampling variability. In particular, the standard deviation in the estimated fixed effects is 0.0901 (Table 1, column 7), and 0.086 after adjusting for sampling variability.¹⁸

Column 1 of Table 2 presents our baseline estimate of β from fitting (6) to the data. This estimate indicates a positive relationship between changes in average coworker permanent productivity and changes in individual productivity. The effect appears to be both statistically and economically significant. A 10 percent increase in coworker permanent productivity is associated with a 1.5 percent increase in reference worker productivity. This estimate is robust to the inclusion of a variety of controls, including dummies for store by hour of day and by day of week and a quadratic polynomial in the number of minutes a checker has been on duty (column 2).^{19, 20} The estimated β indicates that positive spillovers appear to dominate any free-riding effect. In other words, the return to introducing a high-productivity worker into a group is greater than her individual contribution. The magnitude of this estimate is remarkably similar to recent experimental evidence of productivity spillovers. For example, in a laboratory experiment, Falk and Ichino (2006) find that a 10 percent increase in peer output results in a 1.4 percent increase in individual productivity.²¹

In column 3 we limit the sample to periods when the number of workers on duty did not change. In this subsample, variability in $\bar{\theta}_{-itcs}$ results from workers of differing abilities replacing each other. This cut of the data is interesting because changes in personnel that do not involve changes in the number of workers on duty are less likely to result from changes in staffing levels due to some shock that could affect group-level productivity, for example, increases in the number of customers. Using this sample, we estimate β as 0.13, which is close to the full-sample estimate. In Section IVD we discuss possible threats to validity in greater detail and present results from associated tests.

In column 4 we test for whether the spillover effect is symmetric for positive and negative changes in coworker quality. Specifically, we allow the effects of $\Delta\bar{\theta}_{-itcs}$ on the focal worker's

¹⁸ To compute the variance of the fixed effects adjusted for sampling variability, we use $\text{Var}(\theta_i) = \text{Var}(\hat{\theta}_i)/A$, where A is the Wald statistic corresponding to the estimated fixed effects divided by 393 (number of workers – 1). See Cory Koedel and Julian Betts (2007).

¹⁹ Estimates are robust to the use of store \times ten-minute time interval \times day-of-week dummies. Using dummy variables for the number of ten-minute periods a worker has been on duty (in place of the polynomial) yields the same point estimate and standard error. Estimates are also robust to inclusion of the lagged log productivity level of worker i in period $t-1$. When adding this control to Table 2, column 2, β is estimated as 0.12 (0.02). To control for potential demand shifts in past periods, we also estimate models with four lagged variables in the change in the number of workers on duty. The estimated β is 0.16 (0.02) when including these lags to the specification in Table 2, column 2.

²⁰ By adding these controls, the model we are estimating is no longer strictly consistent with the one specified in equation (4).

²¹ An alternative specification would be a model where the primary regressor of interest is the change in the average *contemporaneous* productivity of coworkers. One could estimate this model by 2SLS, using the change in the average *permanent* productivity of coworkers as instrument. As explained above, this approach would simply generate the estimates in Table 2 rescaled by the first-stage coefficient, which in this case is equal to 0.99.

TABLE 1—DESCRIPTIVE STATISTICS

	Store # 1 (1)	Store # 2 (2)	Store # 3 (3)	Store # 4 (4)	Store # 5 (5)	Store # 6 (6)	All stores (7)
Share of ten-minute interval that checkers are transacting	0.67 [0.32]	0.61 [0.25]	0.64 [0.28]	0.69 [0.26]	0.68 [0.24]	0.60 [0.26]	0.65 [0.27]
Minutes per customer	1.4 [1.0]	1.2 [1.1]	1.6 [1.1]	1.3 [1.1]	1.4 [0.86]	1.4 [0.91]	1.4 [1.0]
Productivity in ten-minute intervals	0.18 [0.09]	0.16 [0.07]	0.17 [0.08]	0.16 [0.07]	0.18 [0.07]	0.20 [0.08]	0.17 [0.08]
Checkers on duty in ten-minute intervals	5.8 [1.9]	5.9 [1.6]	4.7 [1.7]	7.7 [2.1]	8.3 [2.4]	7.0 [2.3]	6.9 [2.4]
Estimated individual fixed effects	[0.07]	[0.12]	[0.08]	[0.08]	[0.09]	[0.09]	[0.09]
Average coworker permanent productivity	[0.04]	[0.06]	[0.04]	[0.03]	[0.04]	[0.04]	[0.04]
Change in coworker permanent productivity	[0.02]	[0.03]	[0.03]	[0.02]	[0.02]	[0.02]	[0.02]

Notes: Main entries are means. Figures in brackets are standard deviations. The estimated fixed effects are mean 0 by construction. The units of observation are checker × ten-minute cells. Individual productivity is defined as the number of items scanned per seconds spent in transactions over a ten-minute period. Specifically, for each worker on the shift, we sum the number of items that a 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 and when the payment is completed and the receipt for the transaction is given 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 AM–8 PM interval. The sample excludes transactions involving managers. Average coworker permanent productivity is the average permanent productivity of all workers in a shift, excluding permanent productivity of the focal worker i .




FIGURE 1. DISTRIBUTION OF WORKERS' PERMANENT PRODUCTIVITY

Note: This figure shows a kernel density estimate of estimated worker permanent productivity (the parameter θ_i), obtained by fitting equation (7). The sample is 394 workers. We used an Epanechnikov kernel and “optimal” bandwidth.

TABLE 2—THE EFFECT OF CHANGES IN AVERAGE COWORKER PERMANENT PRODUCTIVITY ON FOCAL WORKER PRODUCTIVITY
(Dependent variable is the difference in log productivity of the focal worker between t and $t - 1$)

	(1)	(2)	(3)	(4)	(5)
Δ Average coworker permanent productivity	0.15 (0.02)	0.15 (0.02)	0.13 (0.03)	-0.03 (0.03)	
Δ Average coworker permanent productivity \times positive Δ indicator				0.24 (0.05)	
Positive Δ indicator				0.004 (0.001)	
Entry of above average productivity worker					0.010 (0.001)
Exit of an above average productivity worker					-0.005 (0.001)
Observations	1,718,052	1,718,052	823,274	1,718,052	1,732,941
Additional controls?		Yes			
No net change in number of workers from $t - 1$ to t ?			Yes		

Notes: OLS estimates. Bootstrapped standard errors that are robust to serial correlation are in parentheses. The parametric bootstrap procedure is described in Section III and in the Appendix. Observations are checker \times ten-minute cells. See Table 1 notes for additional details on the sample. The symbol Δ indicates the difference in the relevant variable between periods t and $t - 1$. The dependent variable is the change in the log productivity of a checker across consecutive ten-minute periods. Individual productivity is defined as the number of items scanned per second transacting over a ten-minute period. Average coworker permanent productivity is computed as the simple average of coworker permanent productivity components (estimated by fitting equation (7)) in a given shift, where the average excludes focal worker i . All specifications other than (3) control for the change in the number of workers on duty across consecutive ten-minute periods. In column 3 the sample consists of time periods where there was no change in the number of workers on duty. The additional controls in column 2 are dummies for store by hour of day and by day of week and a quadratic polynomial in the number of ten-minute periods a worker has been on duty. In column 5 we also include in the model dummies whether there was any entry or exit into or out of the shift, irrespective of the productivity of the worker entering or departing (not reported). “Entry” is defined as one if in a given day a checker was on duty at time t but not on duty at time $t - 1$, and zero otherwise. “Exit” is defined as one if in a given day a checker was on duty at time t but not on duty at time $t + 1$, and zero otherwise. “Positive Δ indicator” is one if “ Δ average coworker permanent productivity” is positive and zero otherwise.

productivity to vary depending on whether $\Delta\bar{\theta}_{-itcs}$ is positive or negative. To do so, we interact $\Delta\bar{\theta}_{-itcs}$ with an indicator variable that is one if $\Delta\bar{\theta}_{-itcs}$ is positive and zero otherwise, and control for a main-effect variable indicating whether there was an increase in the average productivity of coworkers. Interestingly, we find that workers are primarily responsive to positive shocks to coworker productivity.

In the model in column 5, the key independent variables are indicators for the entry and the exit of a worker with above average permanent productivity. 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. The entry of a worker with above average permanent productivity is associated with a 1 percent increase in the productivity of coworkers, relative to entry of a below average worker. By contrast, the exit of an above average worker is associated with approximately a 0.5 percent decline in coworker productivity relative to the exit of a below average productivity worker. Using the mid-point between the effect of an entry relative to an exit of a high-productivity worker (0.75 percent), the estimates imply that if in every shift the firm could replace a lower than average productivity worker with a higher than average productivity worker, the firm’s labor inputs would decline by 0.75 percent, through the effects of the resulting spillovers alone, holding output constant.

B. Is the Spillover Effect Persistent over Time?

We have found that a worker's effort is increasing in the permanent productivity of coworkers. The implications of this effect are very different depending on whether the effect is short lived or permanent, both for the firm and for our understanding of productivity spillovers. In order to determine the persistence of this effect, we estimate a version of equation (6) that includes both the current change in average coworker permanent productivity and seven leads and seven lags (the unit of time is ten minutes, as before):

$$(8) \quad \Delta y_{itcs} = \beta_{-7} \Delta \bar{\theta}_{-i(t-7)cs} + \beta_{-6} \Delta \bar{\theta}_{-i(t-6)cs} + \dots + \beta_0 \Delta \bar{\theta}_{-i(t)cs} + \dots + \beta_6 \Delta \bar{\theta}_{-i(t+6)cs} \\ + \beta_7 \Delta \bar{\theta}_{-i(t+7)cs} + \pi \Delta N_{ict} + e_{itcs}$$

The coefficients on the lag terms from this model, β_{-7} through β_{-1} , allow us to examine how a shock to the composition of coworkers propagates over time. The coefficients on the lead terms, β_1 through β_7 , allow us to determine how a checker's productivity responds to a future shock to coworker composition.


In panel A of Figure 2, we plot the coefficients $\hat{\beta}_7$ through $\hat{\beta}_{-7}$. The leftmost coefficient, $\hat{\beta}_7$, represents the *change* in the focal worker's productivity in response to a shock to the permanent productivity of coworkers 70 minutes into the future. The rightmost coefficient, $\hat{\beta}_{-7}$, represents the focal worker's *change* in productivity in response to a shock to the permanent productivity of coworkers 70 minutes in the past. The dotted lines are a 95 percent confidence band. In panel B, we present the cumulative estimates corresponding to (8). Specifically, the figure displays $\hat{\sigma}_u = \hat{\beta}_7 + \hat{\beta}_6 + \dots + \hat{\beta}_u$, for $u = -7$ through 7, along with an accompanying error band. To interpret this figure, suppose that there is a shock to the average coworker permanent productivity at time $t = 0$. The point that is farthest to the left, $\hat{\sigma}_7$, represents the productivity response 70 minutes prior to the shock. The next point moving to the right, $\hat{\sigma}_6$, is the estimated cumulative productivity response up to 60 minutes before the shock ($\hat{\beta}_7 + \hat{\beta}_6$). The point $\hat{\sigma}_0$ is the cumulative productivity response from 70 minutes before the shock through 70 minutes after the shock.

Two features of Figure 2 are striking. First, estimates reveal that exactly at the time when there is a positive shock to the average coworker productivity, there is an immediate rise in the focal worker's productivity. Productivity then declines somewhat over the next 40 minutes, as inferred from the negative coefficients on the first four lagged terms in panel A, and from the downward drifting cumulative response in panel B.²² The spillover is nevertheless highly persistent: the sum of the estimated coefficients on the contemporaneous and lagged terms ($\hat{\beta}_{-7} + \hat{\beta}_{-6} + \dots + \hat{\beta}_0$) is 0.13, with a t -ratio of 2.6. This quantity implies that approximately 73 percent of the effect at $t = 0$ is still present 70 minutes after the shock.

A second important feature of Figure 2 is that the lead terms in (8) provide a test of the validity of our identifying assumptions. 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 in the current period, conditional on the mix of workers in the current period. Such a finding would be inconsistent with the possibility that high-productivity workers begin shifts prior to a large increase in demand, as would be the case if managers could anticipate demand and schedule workers just prior to increases. Figure 2 shows that we cannot reject that any of the coefficients on the lead terms are zero. This evidence is inconsistent with the pos-

²² 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.

Panel A: Estimated coefficients on the lag and lead terms for changes in the average permanent productivity of coworkers, where the dependent variable is change in focal worker productivity. See equation (8).



Panel B: Cumulative response.




FIGURE 2: DYNAMIC RESPONSE FOLLOWING A SHOCK TO THE AVERAGE PERMANENT PRODUCTIVITY OF COWORKERS

Notes: Figures based on equation (8) in the text. In panel A, we plot the estimated coefficients on the lagged and lead terms in (8) along with accompanying 95 percent confidence intervals (dotted line). For example, $\hat{\beta}_7$ is the coefficient on the seventh lead term. In panel B we plot the cumulative response: $\hat{o}_u = \hat{\beta}_7 + \hat{\beta}_6 + \dots + \hat{\beta}_u$, for $u = -7$ through 7. For example, $\hat{o}_{-2} = \hat{\beta}_7 + \hat{\beta}_6 + \hat{\beta}_5 + \hat{\beta}_4 + \hat{\beta}_3 + \hat{\beta}_2 + \hat{\beta}_1 + \hat{\beta} + \hat{\beta}_{-1} + \hat{\beta}_{-2}$. In estimating (8), standard errors were computed using the bootstrapped procedure described in Section III and in the Appendix.

sibility of endogenous turnover of high-productivity workers generating the observed patterns in the data.²³ We return to this point in Section IVD, where we discuss the validity of our design in more detail.

C. Heterogeneity in the Spillover Effect

The baseline 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 focal worker. Specifically, we estimate models of the form:

$$(9) \quad \Delta y_{itcs} = \alpha + \beta \Delta \bar{\theta}_{-itcs} + \lambda \Delta \bar{\theta}_{-itcs} L_i + \pi \Delta N_{itcs} + e_{itcs},$$

where L_i is a dummy equal to one if worker i 's permanent productivity is above average in the store. A negative λ and a positive β imply that low-skill workers are more responsive to changes in coworker composition than are high-skill workers.

Estimates of λ and β coefficients are presented in column 1 of Table 3. Notably, we find that the magnitude of the spillover effect varies dramatically depending on the skill level of the relevant worker. In particular, the negative estimate of λ indicates that workers with low permanent productivity are substantially more responsive to changes in the average permanent productivity of coworkers than workers with high permanent productivity. While the spillover coefficient is large and positive for workers who are below average (0.24), it is small but nonnegative for workers who are above average (0.24 – 0.19 = 0.05). These estimates are robust to the inclusion of the same additional controls used in Table 2 (Table 3, column 2).

Our longitudinal data allow for a more general model than the one in equation (9). Specifically, we estimate a random coefficient model where the spillover effect is allowed to vary by individual:

$$(10) \quad \Delta y_{itcs} = \beta_i \Delta \bar{\theta}_{-itcs} + \pi \Delta N_{itcs} + e_{itcs},$$

where β_i is an individual-specific spillover. Unlike equation (9), this model does not constrain the coefficient to be the same for all workers in a given skill group. As before, we find that there is substantial heterogeneity in how workers respond to peers. The spillover effect is large for some workers, and small—even negative—for others. Figure 3 presents estimates of the average β_i , conditional on the worker's permanent productivity using a local-linear smoother.²⁴ Consistent with our previous finding for two groups of workers in Table 3, 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 typically not hurt by the presence of low-skill coworkers.²⁵

²³ 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.

²⁴ Specifically, we estimate a local-linear regression model, weighting by the inverse variance of the estimated β_i from (10), using an Epechnikov kernel with a bandwidth of 0.02.

²⁵ One possible reason for the differential response of high and low permanent productivity checkers to coworker composition is that fast checkers are always working at their potential, while the slower ones are not. This is consistent

TABLE 3—HETEROGENEITY IN RESPONSIVENESS BY ABILITY TYPE
(Dependent variable is the difference in log productivity of the focal worker between t and $t - 1$)

	(1)	(2)
Δ Average coworker permanent productivity	0.24 (0.03)	0.24 (0.03)
Δ Coworker permanent productivity × above average worker	-0.19 (0.04)	-0.18 (0.04)
Observations	1,718,052	1,718,052
Additional controls?		Yes

Notes: Specifications present estimates from OLS models. The sample is the same as in Table 2, column 1. See Table 2 notes for more details on sample. Bootstrapped standard errors that are robust to serial correlation are in parentheses. The parametric bootstrap procedure is described in Section III and in the Appendix. The units of observation are checker × ten-minute cells. The symbol Δ indicates the difference in the relevant variable between periods t and $t - 1$. 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 transacting over a ten-minute period. “Above average worker” refers to a worker with above average permanent productivity in his or her store. Average coworker permanent productivity is computed as the simple average of worker permanent productivity components (estimated by fitting equation (7)) in a given shift, where the average excludes the focal worker i . The additional controls in column 2 are dummies for store by hour of day and by day of week and a quadratic polynomial in the number of ten-minute periods a worker has been on duty.

The finding that the spillover effect is large for low-skilled workers and small for high-skilled workers is important because it implies that the mix of workers that maximizes productivity is the one that maximizes skill diversity within shifts. 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 only of high-skill workers and other shifts are made only of low-skill workers.

How large are the possible savings in labor inputs obtainable by optimally mixing the existing set of workers? When we compare productivity obtained under the observed mix and productivity obtained under an ideal mix (maximizing skill diversity), we find 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 percent fewer hours worked each year.²⁶ 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 current labor costs, this difference amounts to a wage bill that is approximately \$2.5 million per year higher than under the optimal mix. Because optimizing shifts may result in higher labor costs, this finding does not necessarily imply that the firm is not maximizing profits.²⁷

D. Tests of 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 induce spurious correlation. For

with the earlier finding that workers are primarily responding to positive coworker productivity shocks.

²⁶ When we compare the observed mix with the mix that is obtained by *randomly* mixing workers, we find virtually no difference. This finding is inconsistent with a systematic effort by management to group workers to maximize productivity.

²⁷ 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. It is unclear how the additional productivity that the firm could obtain by imposing optimal mixing of workers relates to this compensating differential.




FIGURE 3. RELATIONSHIP BETWEEN INDIVIDUAL PERMANENT PRODUCTIVITY AND WORKER-SPECIFIC SPILLOVER EFFECT

Notes: This figure plots the local-linear regression fit for the relationship between the estimated worker specific spillover effects (the parameter β_i in equation (10)) against that worker's permanent productivity (the parameter θ_i in equation (7)). The dashed lines are the 95 percent confidence intervals. We use an Epanechnikov kernel with a bandwidth of 0.02. The regression is weighted by inverse variance of the estimated worker-specific spillover effects. The unit of observation is checker. There are 394 checkers in the sample.

example, one might be concerned that high-ability workers are scheduled during busy days. If higher customer volume causes workers to speed up, we may estimate positive spillovers where none exists. In our context, we believe that this issue is unlikely to be a serious problem. First, scheduling is unsystematic and management does not have control over which workers are on duty in any given 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 coworkers *within a given day for a given worker in ten-minute windows*. Differences across days do not contribute to identification.

One might still be concerned about the possibility of within-day personnel changes, whereby high-productivity workers differentially enter or exit shifts when demand is changing. For example, we would worry if, in order to shorten the queues, more productive workers were brought at times of the day when demand is elevated and the productivity of all workers is high.²⁸ However, we have seen in Figure 2 that a worker's productivity is not affected by future changes in the average permanent productivity of coworkers. If managers are scheduling high-ability workers

²⁸ This hypothesis is not very likely to be a serious concern in our context. Due to union rules, the checkers work only at the registers. Managers and assistant 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 upward.

during periods when shift-level productivity is high for other reasons, they are doing so in a very remarkable way—these workers are being assigned two weeks in advance to shifts that begin at exactly the moment when productivity spikes, not ten minutes before or ten minutes after.

In general, the identifying assumption in equation (6) for the causal interpretation of β is that the changes in coworker 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). 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 that better workers are more likely to enter (or less likely to exit) when there is a positive demand shock.

(1a) We first investigate whether the speed of entering and exiting workers is associated with *predictable* changes in demand. One way to measure predictable demand is to compute the average quantity sold in each store by day of week and by time of day. This measure reflects a “typical” quantity at each time of day and day of week, something that the manager plausibly knows in advance. Because our measure of customer volume is itself a function of the productivity of personnel, for each worker we use the average quantity sold by all other workers, omitting observations for which that worker was on duty. This approach allows us to determine the relationship between exogenous and predictable changes in demand (depending on the day of week and the time of day) and changes in composition of coworkers in a shift free of a possible mechanical relationship between these two measures. For worker i , the average quantity sold for a given day of week and time of day is denoted Q_{idt} . This term is subscripted with an i because for every individual we compute average quantity in cells for day of week and time of day, excluding that individual’s observations. Using this measure, we estimate the following linear probability models:

$$(11) \quad \text{ENTRY}_{ics} = \alpha_i + \eta_1 \Delta \ln Q_{idt} + \bar{\omega}_1 \hat{\theta}_i + \rho_1 \Delta \ln Q_{idt} \hat{\theta}_i + e_{ics},$$

$$(12) \quad \text{EXIT}_{ics} = \alpha_i + \eta_2 \Delta \ln Q_{idt} + \bar{\omega}_2 \hat{\theta}_i + \rho_2 \Delta \ln Q_{idt} \hat{\theta}_i + e_{ics}.$$

ENTRY_{ics} takes on the value of one if worker i is observed on duty at time t , but not at $t - 1$; EXIT_{ics} takes on the value of one if worker i is observed on duty at time t but not at $t + 1$ and; $\hat{\theta}_i$ is the estimated fixed effect for individual i obtained by estimating (7). The term $\Delta \ln Q_{idt}$ represents the change in the log of predictable sales, as described above, and $\Delta \ln Q_{idt} \hat{\theta}_i$ 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 are more likely to enter shifts when demand is on average rising, or exit when demand is on average falling, then $\rho_1 > 0$ and $\rho_2 < 0$.

Columns 1 and 2 of Table 4 present estimates of the parameters in equations (11) and (12), respectively. The estimate of η_1 presented in column 1 suggests that, for the average worker, the probability of entry rises with increases in predicted demand, although the estimate is not precise enough to be significant. A positive relationship is to be expected if shifts are scheduled so that more workers are on duty during typically busy periods. 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. However, we cannot reject that ρ_1 is zero. In column 2 we examine the determinants of exits. Not surprisingly, we find exit probabilities are significantly lower when predictable demand is rising. But we find no evidence that more productive workers delay exiting the registers

TABLE 4—RELATIONSHIP BETWEEN CHANGES IN PREDICTABLE SALES VOLUME AND EXIT AND ENTRY PROBABILITIES OF PERSONNEL
(Linear probability models)

	Dependent variable:	
	Entry at t (1)	Exit at t (2)
Δ Predicted log average transactions between $t - 1$ and t	0.029 (0.048)	-0.086 (0.009)
Permanent productivity	0.090 (0.007)	0.077 (0.007)
Δ Predicted log average transactions between $t - 1$ and $t \times$ permanent productivity	-0.078 (0.048)	0.21 (0.07)
Observations	2,029,970	2,029,970

Notes: 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 . Bootstrapped standard errors that are robust to serial correlation are in parentheses. The parametric bootstrap procedure is described in Section III and in the Appendix. Column 1 corresponds to equation (11) in the text and column 2 corresponds to equation (12). The symbol Δ indicates the difference in the relevant variable between periods t and $t - 1$. 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). See text for additional details.

when demand is rising. In fact, the positive point estimate of ρ_2 suggests that high-productivity workers are somewhat *more* likely to exit when predictable demand is rising.

(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 ten-minute periods, during which the change in number of items scanned is large (small), are immediately followed by gains (declines) in the permanent productivity of personnel. 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. Relative to 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.

In practice, we estimate models that are similar to those in Table 4, 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 ten-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 1 of Table 5, -0.054 , 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 may be working faster than their typical levels. Column 2 shows the probability that high-productivity workers are no more or less likely to exit than low-productivity workers following increases in sales volume.²⁹

²⁹ A puzzling feature of this table, however, is that positive demand shocks between $t - 2$ and $t - 1$ are associated with lower entry probabilities at time t .

TABLE 5—RELATIONSHIP BETWEEN LAGGED CHANGES IN SALES VOLUME AND ENTRY AND EXIT PROBABILITIES OF PERSONNEL
(Linear probability models)

	Dependent variable	
	Entry at t (1)	Exit at t (2)
Δ log items sold between $t - 2$ and $t - 1$	-0.032 (0.0002)	0.006 (0.0007)
Permanent productivity	0.090 (0.005)	0.081 (0.005)
Δ log items sold between $t - 2$ and $t - 1$ × permanent productivity	-0.054 (0.008)	-0.0005 (0.007)
Observations	2,050,558	2,050,558

Notes: “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 . Bootstrapped standard errors that are robust to serial correlation are in parentheses. The parametric bootstrap procedure is described in Section IIIA and in the Appendix. The units of observation are checker × ten-minute cells. The symbol Δ indicates the difference in the relevant variable between periods t and $t - 1$.

(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 who enters a shift when customer volume is rising tends to be more able. Relative to test (1b), this measure has the advantage that it is based on the relationship between quality of workers at time t and personnel in the same period, rather than a lagged period. Figure 4 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 discernible relationship between net changes in the number of personnel and changes in the average permanent productivity of personnel. The observed variability in the median change in the average permanent productivity of personnel is small relative to its standard deviation of 0.021.

(2) *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, which 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 coworkers in such a way as to generate the kinds of patterns we have documented thus far. An obvious candidate is the presence of baggers, who 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 for the productivity of other workers for no reason other than 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. Given this policy, 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.




FIGURE 4. RELATIONSHIP BETWEEN NET CHANGES IN PERSONNEL ON DUTY AND CHANGES IN THE AVERAGE PERMANENT PRODUCTIVITY OF WORKERS

Notes: The x-axis is the net change in the number of personnel across consecutive ten-minute periods. Absolute values in changes larger than five are excluded due to small sample sizes. 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. Change in average permanent productivity is computed as the difference in the average fixed effects of workers in a store at time t and the average fixed effects of workers in a store at time $t - 1$. Box plots exclude outside values.

This is because faster workers, on average, have more groceries that require bagging, implying that baggers will spend less time with slower workers. This observation would suggest that our estimates of the spillover are a lower bound.

Another piece of evidence that is relevant to this issue, which we discuss in the next section, 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 coworker productivity when introducing a productive worker to a shift.

V. 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. What explains these findings? In this section, we consider three explanations for these peer effects. We emphasize that these explanations are to be interpreted as broad classes of possible mechanisms. While broad, these hypotheses are a first step to understanding the possible channels responsible for the documented productivity spillovers.

Social Pressure.—Because of the features of the production process, when a checker is working more slowly than her coworkers, she is essentially giving a larger share of the workload to her peers. Social pressure may be relevant in this situation. We define social pressure as encompassing cases where workers have preferences over how they are perceived by their coworkers. Specifically, under social pressure, workers are subject to a loss of utility if they are *observed* behaving noncooperatively by their peers. Workers might care about how they are perceived for a variety of reasons, including shame, sanctions, or reputational concerns which could arise in repeated interactions. For example, if a worker is slow, other workers may impose a cost on her, for example, by reporting her to management or by ostracizing her socially.³⁰ The social pressure explanation has the implication that the introduction of a productive worker will lead to increases in the productivity of incumbent workers who are easily observed by the entering worker.

Prosocial Preferences.—We define prosocial preferences as encompassing cases where workers lose utility if they act uncooperatively, regardless of whether the worker is being observed acting in this manner. These kinds of preferences would be embodied by altruistic workers, or those subjected to feelings of guilt. Alternatively, they could be due to competitive spirit, leading by example, or some kind of contagious enthusiasm (or contagious malaise). Irrespective of the reason, if prosocial behavior is the primary underlying mechanism, 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 the checkout stand.

Knowledge Spillovers.—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 nearby. If knowledge spillovers are operating, they should be highly localized; it is implausible that workers could consistently communicate more than two registers away. We have no a priori reason to think that the spatial orientation of workers matters for potential knowledge spillovers. That is, worker A can communicate with worker B if she is positioned in front of B or behind him. While knowledge spillovers may be important in some settings (see, for example, Moretti 2004a, b; Ichimowski, Shaw, and Prenzushi 1997), based on what we already know from Table 3, we do not expect them to have a quantitatively important

³⁰ Jeffrey Carpenter, Samuel Bowles, and Herbert Gintis (2006) provide experimental evidence from a laboratory that reciprocal behavior is relevant in a team production setting with repeated interactions.

effect in our context. 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 prediction is inconsistent with Table 3, which shows that faster workers appear to make slower workers go faster, but the presence of slow workers does not reduce the speed of the faster workers.

A. *Distinguishing between Alternative Hypotheses*

To empirically distinguish between these hypotheses, we present two pieces of evidence. First, in Section VA(1) we estimate models where the effect of coworkers is allowed to vary depending on whether coworkers can observe each other while working. The idea is simple: 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 coworkers, and small when she cannot easily be observed. By contrast, if the primary mechanism is a form of prosocial preferences, the introduction of a new and productive worker should lead to increases in the productivity of workers who can see this new worker, but have a more limited effect on the workers who cannot see her. On the other hand, simple models of knowledge spillovers do not predict an asymmetry depending on spatial orientation. In the presence of such spillovers, we expect to see highly localized effects that are not sensitive to whether coworkers are observing or observable by the focal worker.³¹

Second, in Section VA(2), we examine how the spillover effect varies as a function of the frequency of interactions between workers. If social pressure is behind the peer effect that we find, we should find that the magnitude of the spillovers depends on whether workers on a given shift have overlapped frequently 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 we meet rarely than on individuals 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.³²

(I) Spatial Orientation.—We have information on the location of each checker within a store, her spatial orientation, and her distance to other checkers. The layout of the registers 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. A given worker can more easily observe the set of coworkers positioned in front of her than the coworkers positioned behind her. For example, suppose that in a hypothetical store there are ten registers. In each of these registers the checker is facing to the right, from the point of view of the customers in line. In this case, the checker in the register all the way to the left (from the point of view of the customer) has a direct view of all of the other checkers. The checker all the way to the right cannot easily observe any checkers, but is directly observed by all. For a given worker i , we will call the set of coworkers who can see i directly as her “observing” set of coworkers. We will call the set of workers i can see directly as her “observable” set of coworkers.³³

³¹ Observability of actions is a natural starting point for our analysis, as it has been emphasized by both the theoretical literature and experimental literature as playing an important role in determining economic behavior (e.g., Kandel and Lazear 1992; Steve Tadelis 2007). While not possible to analyze with our data, the role of communication between workers in eliciting effort would also be of interest given the findings of Gary Charness and Martin Dufwenberg (2006).

³² On the other hand, it is also possible that interacting infrequently may raise effort for precautionary reasons.

³³ As an illustration, consider the case the registers are positioned as follows:

1→2→3→4→5→6→

We estimate models that allow for the spillover effect to vary depending on the location of workers relative to their coworkers. Specifically, we estimate the effect of the permanent productivity of a coworker who enters or exits from a position where it is easy to observe an incumbent worker, as well as the effect of permanent productivity of coworkers entering or exiting positions in the line of sight of the incumbent worker. We assume that while it may be difficult for a worker to monitor the exact effort level provided at a moment in time by a coworker located behind, her identity and her average productivity are known. In practice, we estimate the following model:

$$(13) \quad \Delta y_{itcs} = \alpha + \beta^G \Delta \bar{\theta}_{-itcs}^G + \beta^E \Delta \theta_{-itcs}^E + \pi \Delta N_{itcs} + e_{itcs},$$

where $\Delta \bar{\theta}_{-itcs}^G$ denotes the change in the average permanent productivity of checker i 's observing set of coworkers, while $\Delta \theta_{-itcs}^E$ denotes the change in the average permanent productivity of checker i 's observable set of coworkers.

Baseline Estimates by Spatial Orientation. Column 1 of Table 6 presents estimates of β^G and β^E from (13). The findings are quite stark. Just about the entire peer effect documented in Table 2 is operating through changes in workers from the observing set. Changes in the permanent productivity of the observable set has virtually no effect on the focal worker's productivity. Specifically, we estimate β^G as 0.17 (t -ratio = 8.5) and β^E as 0.01 (t -ratio = 0.5). Notably, this pattern is not being driven by any single store. We observe this pattern in five out of the six stores in our sample.³⁴

In Figure 5 we present estimates from models with seven lags and seven leads in the change in permanent productivity of average coworker in the observing and observable sets. Specifically, we estimate the following two models:

$$(14) \quad \begin{aligned} \Delta y_{itcs} = & \beta_{-7}^G \Delta \bar{\theta}_{-i(t-7)cs}^G + \dots + \beta_0^G \Delta \bar{\theta}_{-i(t)cs}^G + \beta_1^G \Delta \bar{\theta}_{-i(t+1)cs}^G + \dots \\ & + \beta_7^G \Delta \bar{\theta}_{-i(t+7)cs}^G + e_{itcs}, \end{aligned}$$

$$(15) \quad \begin{aligned} \Delta y_{itcs} = & \beta_{-7}^E \Delta \bar{\theta}_{-i(t-7)cs}^E + \dots + \beta_0^E \Delta \bar{\theta}_{-i(t)cs}^E + \beta_1^E \Delta \bar{\theta}_{-i(t+1)cs}^E + \dots \\ & + \beta_7^E \Delta \bar{\theta}_{-i(t+7)cs}^E + e_{itcs}. \end{aligned}$$

These models are just like (8), but rather than using the lags and leads of the change in average coworker productivity as the regressors of interest, we use the lags and leads of the change in the average coworker productivity for workers in the observing and observable sets ((14) and (15), respectively).


Here, each number represents a register, and the arrow points toward the direction that the checker on that register is facing. Register 1 is directed toward register 2, which is directed toward register 3, and so on. Suppose that a worker is stationed at register 4. Then workers who are on registers 1, 2, and 3 are able to directly observe the worker on 4. The checker on registers 5 and 6 are positioned to be easily seen by 4; that is, 5 and 6 are in 4's direct line of sight. Alternatively, consider the case where registers positioned as follows:

$$1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow \leftarrow 5 \quad 6 \rightarrow$$


In this case, the observing set for the worker on register 4 includes coworkers stationed on registers 1, 2, 3, and 5. The observable set includes coworkers on registers 5 and 6. These two cases are just an illustration. We cannot reveal the exact details of the spatial arrangements of the registers in the stores for reasons of confidentiality.

³⁴ In the store where we do not observe this pattern, we do not estimate any spillovers on average (β in equation (6) is not statistically different from 0).


Panel A: Dynamic response to a positive average coworker permanent productivity shock in the observing set



Panel C: Dynamic response to a positive average coworker permanent productivity shock in the observable set



Panel B: Cumulative response to a positive average coworker permanent productivity shock in the observing set



Panel D: Cumulative response to a positive average coworker permanent productivity shock in the observable set




FIGURE 5. DYNAMIC RESPONSE MODELS BROKEN DOWN BY SPATIAL ORIENTATION

Note: In panel A we plot the estimated coefficients on the lagged and lead terms from estimating (14) along with accompanying 95 percent confidence intervals (dotted line). For example, $\hat{\beta}_7$ is the coefficient on the seventh lead term. In panel B we plot the cumulative response: $\hat{\delta} = \hat{\beta}_7 + \hat{\beta}_6 + \dots + \hat{\beta}_u$, for $u = -7$ through 7 . Panels C and D correspond to the estimated lag and lead coefficients and the cumulative response for equation (15), respectively. Standard errors were computed using the bootstrapped procedure described in Section III and in the Appendix.

In panels A and B of Figure 5, we consider the focal worker’s dynamic response to a change in average coworker permanent productivity in the observing set. Panel A presents the estimates $\hat{\beta}_7^G$ through $\hat{\beta}_{-7}^G$ and panel B presents the cumulative response, as in Figure 2. There is a sharp and highly persistent increase in the focal worker’s productivity following a positive shock in the observing set. This pattern is similar to the pattern uncovered in Figure 2. Notably, this increase is not observed for shocks in the observable set, as seen in panels C and D. In fact, there appears to be a moderate downward drift in the focal worker’s productivity following a productivity shock in the observable set.

Overall, the estimates imply that as the more productive workers are introduced into a shift, only the coworkers who are in their direct line of sight (i.e., those they can observe), seem to become more productive. This effect manifests itself exactly at the time of entry of more productive workers and appears to be largely permanent. 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 sight of the new workers. The estimates suggest that the workers have preferences over how they are perceived by others, providing support for the social pressure explanation.

Testing for Threats to Validity. One may worry that our estimates are capturing location-specific effects. Special cases can be constructed whereby heterogeneity in responsiveness to peers

TABLE 6—EFFECT OF CHANGES OF AVERAGE COWORKER PERMANENT PRODUCTIVITY ON FOCAL WORKER'S PRODUCTIVITY
(Models by spatial orientation and proximity; dependent variable is the difference in log productivity of the focal worker between t and $t - 1$)

	(1)	(2)	(3)	(4)	(5)	(6)
Δ Avg. coworker permanent productivity in observing set	0.17 (0.02)	0.18 (0.02)				
Δ Avg. coworker permanent productivity in observable set	0.01 (0.02)	0.00 (0.02)				
Δ Avg. coworker permanent productivity observing set & closer					0.11 (0.01)	
Δ Avg. coworker permanent productivity in observable set & closer					0.01 (0.01)	
Δ Avg. coworker permanent productivity in observing set & farther					0.06 (0.01)	
Δ Avg. coworker permanent productivity in observable set & farther					0.01 (0.01)	
Δ Presence of worker in observing set			0.031 (0.003)	0.031 (0.004)		
Δ Presence of worker in observable set			-0.03 (0.003)	-0.022 (0.004)		
Δ Presence of worker in observing set & closer						0.027 (0.002)
Δ Presence of worker in observable set & closer						-0.026 (0.002)
Δ Presence of worker in observing set & farther						0.013 (0.002)
Δ Presence of worker in observable set & farther						0.008 (0.002)
Using only middle registers?		Yes		Yes		
Observations	1,649,916	683,933	1,732,941	721,176	1,487,086	1,732,941

Notes: Specifications present estimates from OLS models. Column 1 corresponds to estimating equation (13) in the text. The remaining columns are variants. The original sample is the same as the one used in Table 2. The sample sizes differ from those in Table 2 because, for some individuals, there are no observations in the relevant subgroup (e.g., average coworker permanent productivity in the observable set). Bootstrapped standard errors robust to serial correlation are in parentheses. The parametric bootstrap procedure is described in Section IIIA and in the Appendix. 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. "Observing set" refers to workers who are facing worker i . "Observable set" refers to workers whom i is facing. "Closer" refers to workers who are one or two positions away from i . "Farther" denotes workers who are three or four positions away from i . The permanent productivity averages are taken over the indicated subgroups. The symbol Δ indicates the difference in the relevant variable between periods t and $t - 1$. For example, " Δ avg. coworker permanent productivity in the observable set & closer" denotes the change in the average permanent productivity of coworkers in i 's observable set who are one or two registers away from i . In this example, if there are no workers positioned one or two positions in front of i in both $t - 1$ and t , this change is coded as 0. The variable " Δ presence of worker in observing set" denotes the change in the presence of a worker for whom i is in the line of sight. As in the baseline specifications from Tables 2 and 3, models presented in columns 1 and 3 control for the change in the number of workers on duty. The models presented in columns 2 and 4 do not control for this variable.

depending on register assignment, and not the spatial orientation of the checker, produces the patterns we observe in the data. To address this potential problem, we estimate a model where the slope coefficients β^G and β^E are allowed to vary depending on the register assignment of the checker. We employ such a specification by estimating (13) separately for every register across all stores. The average estimated β^G and β^E across all of these regressions, weighted by the number of observations corresponding to a given register, is 0.13 and 0.04, respectively. These estimates

are close to the ones presented in column 1 of Table 6, and we therefore do not believe that the observed asymmetries are coming about from this kind of register heterogeneity. To further reinforce this point, in column 2 of Table 6 we present estimates for β^G and β^E for the subsample of registers closer to the center of the store.³⁵ These registers are different from the registers at the extreme ends of the store, in the sense that they have a roughly equal number of registers in the observing versus observable set, whereas the registers at the extremes are imbalanced in this respect. The estimated β^G and β^E coefficients using this cut of the data are very close in magnitude to those in column 1, implying that there is not a lot of heterogeneity in these effects depending on register location.

A second possible concern has to do with sorting across registers. It may be the case that when customer volume increases, and when the shift-level productivity is high, fast workers are more likely to be positioned at registers at the extreme ends of the store, where they are better positioned to observe other workers. This reallocation of workers across registers may lead us to conclude that productivity increases when a worker is being observed, when in fact it is the case that productivity has increased for other reasons, and it is simply that fast workers are more likely to observe other workers at this time. To test for this kind of sorting, for every register we constructed a “visibility index” by counting how many registers are directly visible from that location. We then examined whether more productive workers were more likely to switch to registers with better visibility following shocks to customer volume.³⁶ We found no evidence that this kind of sorting occurred, on average.³⁷

Effect of Adding a New Worker, by Spatial Orientation. Up to this point we have focused on whether the speed of coworkers affects the speed of the focal worker. We now consider a different question: does the *presence* of a coworker in the observing or observable set influence productivity? The model in column 3 includes only a dummy for the change in the presence of a worker in the line of sight and a dummy for the change in the presence of a worker who can monitor. Unlike the other models we have estimated, in this model we do not control for the change in the total number of coworkers on duty. Consistent with the notion of social pressure, we find that the addition of a worker who can monitor an incumbent worker, regardless of her productivity, results in increased productivity of the incumbent worker of 3 percent, an estimate that is highly significant. On the other hand, the addition of a worker in the line of sight of the incumbent workers *decreases* productivity of the incumbent worker by 3 percent. As before, this estimate is robust to the exclusion of the registers at the extreme ends of the store (column 4). We interpret this finding as evidence that, for a given number of customers in line, increasing the number of active cashiers has a different effect depending on where the new entrant is located. When a worker goes from a situation where no coworker can easily observe her to one where she is easily observed, she works significantly faster. When a worker goes from a situation where she cannot directly observe any coworkers to one where she can, she works significantly more slowly. This last finding suggests that there is still scope for free-riding, but only when the free-riding is difficult to observe by other workers.

Effects by Distance to Coworkers. In columns 5 and 6 we test whether physical distance matters. The social pressure applied by coworkers who are located behind and are closer appears to

³⁵ In practice, we eliminate from the sample the three registers on both ends of the store.

³⁶ In practice, this involves estimating models exactly as those presented in Tables 4 and 5, but using the change in the visibility index as the dependent variable.

³⁷ Tables are available upon request. We also investigated whether more productive workers enter registers with better views of other registers following shocks to demand. We do not find evidence that this kind of sorting is occurring.

have a larger effect than the pressure applied by coworkers who are behind and are farther away. For example, the coefficient on permanent productivity of coworkers who are one or two registers behind the reference worker is 0.11 (column 5). The corresponding coefficient for coworkers who are three or four registers behind the reference worker is 0.06, or about 55 percent as large. These estimates are significantly different from each other at conventional levels. We find similar results when we consider dummies for addition of coworkers. We estimate that the change in the presence of a coworker one or two positions behind worker i increases i 's productivity by 2.7 percent. A change in the presence of a coworker three or four positions behind i increases i 's productivity by only 1.3 percent. The difference in these two estimates is significant. Similarly, the coefficient on a dummy for the change in the presence of a coworker located in front and close is -0.026 , much larger in magnitude than the coefficient on a dummy for the change in the presence of a coworker located behind but farther apart (0.008).

(2) *Repeated Interactions.*—The asymmetric effect by spatial arrangement suggests that the peer effect we observe is the result of social pressure. Here we provide an additional test of the social interaction hypothesis by testing whether the estimated spillover between workers is larger for workers who interact often relative to workers who rarely interact. Specifically, we model changes in worker productivity as depending on changes in the average levels of permanent worker productivity, where the effect of average coworker productivity is allowed to vary depending on the amount of schedule overlap.

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 an initial window to calculate these shares.³⁸ We define low schedule overlap as cases where coworkers have previously coincided with i between 0 percent and 5 percent of i 's schedule. Medium-overlap coworkers have coincided with 5 percent to 20 percent of i 's schedule. High-overlap workers are those who have coincided with 20 percent to 100 percent of i 's schedule. We then estimate models where we let the spillover vary depending on the degree of previous scheduling overlap:

$$(16) \quad \Delta y_{itcs} = \beta_L \Delta \bar{\theta}_{-itcs}^L + \beta_M \Delta \bar{\theta}_{-itcs}^M + \beta_H \Delta \bar{\theta}_{-itcs}^H + \pi \Delta N_{itcs} + e_{itds},$$

where L denotes workers who have had low schedule overlap with i , M denotes workers who have had medium schedule overlap with i , and H denotes workers who have had high schedule overlap with i . For example, the term $\Delta \bar{\theta}_{-itcs}^M$ denotes the change in the average permanent productivity of i 's coworkers who have had medium overlap with i .

Equation (17) further breaks down the permanent productivity averages by spatial orientation of coworkers:

$$(17) \quad \Delta y_{itcs} = \beta_{G,L} \Delta \bar{\theta}_{-itcs}^{G,L} + \beta_{E,L} \Delta \bar{\theta}_{-itcs}^{E,L} + \beta_{G,M} \Delta \bar{\theta}_{-itcs}^{G,M} + \beta_{E,M} \Delta \bar{\theta}_{-itcs}^{E,M} + \beta_{G,H} \Delta \bar{\theta}_{-itcs}^{G,H} \\ + \beta_{E,H} \Delta \bar{\theta}_{-itcs}^{E,H} + \pi \Delta N_{itcs} + e_{itcs},$$

where the superscript G denotes workers in the observing set of i , and E denotes workers positioned in the observable set of i .

³⁸ Another way to measure how exposed workers are to each other is to count the number of interactions each pair of workers have had. Data limitations prevent us from computing this measure for all but the workers who began working in the sample period, due to censoring. The measure described in the text circumvents this censoring problem.

In column 1 of Table 7, we present estimates of equation (16). The point estimate of β_L is 0.01 and insignificant, meaning that changes in the permanent productivity of coworkers who have had little overlap with i have at most a small effect on the change in i 's productivity from $t - 1$ to t . By contrast, changes in the permanent productivity of coworkers with medium and high previous overlap with i have a larger and statistically significant relationship with changes in i 's contemporaneous productivity. Specifically, β_M and β_H are estimated as 0.06 and 0.07, respectively. The p -values reported in the table show that both $\hat{\beta}_M$ and $\hat{\beta}_H$ are statistically distinguishable from $\hat{\beta}_L$. In column 2 we present estimates from equation (17), which break out the permanent productivity component of coworkers by their previous overlap and spatial orientation in relation to i . As before, the entire spillover effect comes from changes in the composition of workers for whom i is in their line of sight. Regardless of spatial orientation, changes in the permanent productivity do not appear to affect i 's productivity if workers' schedules have had only a small degree of overlap in the past.

We have already seen in Table 6 that not only is there a relationship between changes in the permanent productivity of i 's coworkers and changes in i 's contemporaneous productivity, depending on the spatial orientation of the coworkers in relation to i , but there is also a relationship between changes in the presence of any coworker, 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 explore this result further by examining whether this effect is mediated by previous scheduling overlap. Again, consistent with the social pressure hypothesis, we find that changes in the presence of a coworker in front of i has either no effect, or a negative effect, on changes in i 's productivity. As before, changes in the presence of workers in i 's observing set are associated with positive changes in i 's productivity. Column 3 shows, however, that this effect is limited if the entering or exiting worker has had limited contact with i in the past. Specifically, if there is a change in the presence of a worker in i 's observing set—those workers who are looking at i directly— i 's productivity increases by 2 to 3 percent if the worker exiting or entering has had medium or high previous overlap with i . However, if the entering or exiting worker has had little overlap with the focal worker, the productivity response is only 0.002 percent, and we can reject an effect greater than 0.004 percent at the 5 percent level.

The evidence from this analysis implies that a worker responds to changes in the presence of her peers, but only if the peers who are entering or exiting are positioned to observe this worker, and there has been sufficient previous interaction. This finding is consistent with the social pressure explanation because workers may not care very much about staying on good terms with coworkers with whom they rarely interact.³⁹

B. Interpretation

In sum, the body of evidence in Section VA suggests that mutual monitoring of coworkers and social pressure are important for understanding the source of the spillovers. Notably, a worker is affected by the presence of a skilled coworker when that coworker can observe her, but the same worker is not affected by the presence of a skilled coworker when the coworker cannot observe her. When we further examine the spillovers by previous schedule overlap, we find results that are consistent with this explanation. A worker is affected by the presence of a skilled coworker only when their work schedules are similar and therefore future interactions are likely to occur.

³⁹ This finding is also consistent with workers who rarely overlap not being aware of each other's abilities, thus being unresponsive to each other's presence.

TABLE 7—THE EFFECT OF CHANGES OF AVERAGE COWORKER PERMANENT PRODUCTIVITY ON FOCAL WORKER'S PRODUCTIVITY

(Models by previous scheduling overlap between coworkers and spatial orientation; dependent variable is the difference in log productivity of the focal worker between t and $t - 1$)

	(1)		(2)		(3)
(I) Δ Avg. coworker permanent productivity: low overlap	0.01 (0.01)	(I) Δ Avg. coworker permanent productivity: observing set & low overlap	0.01 (0.02)	(I) Δ Presence of low overlap worker in observing set	0.002 (0.002)
(II) Δ Avg. coworker permanent productivity: medium overlap	0.06 (0.01)	(II) Δ Avg. coworker permanent productivity: observing set & medium overlap	0.11 (0.02)	(II) Δ Presence of medium overlap worker in observing set	0.020 (0.002)
(III) Δ Avg. coworker permanent productivity: high overlap	0.07 (0.02)	(III) Δ Avg. coworker permanent productivity: observing set & high overlap	0.09 (0.03)	(III) Δ Presence of high overlap worker in observing set	0.027 (0.002)
		(IV) Δ Avg. coworker permanent productivity: observable set & low overlap	-0.02 (0.02)	(IV) Δ Presence of low overlap worker in observable set	0.007 (0.002)
		(V) Δ Avg. coworker permanent productivity: observable set & medium overlap	0.01 (0.02)	(V) Δ Presence of medium overlap worker in observable set	-0.004 (0.002)
		(VI) Δ Avg. coworker permanent productivity: observable set & high overlap	-0.06 (0.03)	(VI) Δ Presence of high overlap worker in observable set	-0.009 (0.002)
p -value: H_0 : (I) = (II)	0.002	p -value: H_0 : (I) = (II)	0.001	p -value: H_0 : (I) = (II)	0.000
H_0 : (I) = (III)	0.001	H_0 : (I) = (III)	0.010	H_0 : (I) = (III)	0.000
H_0 : (II) = (III)	0.673	H_0 : (II) = (III)	0.710	H_0 : (II) = (III)	0.003
		H_0 : (IV) = (V)	0.266	H_0 : (IV) = (V)	0.000
		H_0 : (IV) = (VI)	0.234	H_0 : (IV) = (VI)	0.000
		H_0 : (V) = (VI)	0.049	H_0 : (V) = (VI)	0.022
Observations	1,658,491		1,169,175		1,658,491

Notes: Each column refers to a different OLS regression. Bootstrapped standard errors that are robust to serial correlation are in parentheses. The parametric bootstrap procedure is described in Section III and in the Appendix. The sample is based on the one used in Table 2. The units of observation are checker \times ten-minute cells. The sample sizes differ from those in Table 2 because for some individuals there are no observations in the relevant subgroup (e.g., average coworker permanent productivity in the observable set). Also, the first 31 days of the sample in each store are dropped because we require an initial period to estimate the degree of schedule overlap among workers. The dependent variable is the change in the log productivity of a checker across ten-minute periods. In the case of worker i , "low overlap" denotes the set of workers who have previously overlapped between 0 and 5 percent of i 's schedule. "Medium overlap" denotes workers who have overlapped between 5 and 20 percent of i 's schedule. "High overlap" denotes workers who have overlapped between 20 and 100 percent of i 's schedule. "Observing set" refers to workers facing worker i . "Observable set" refers to workers whom i is facing. The permanent productivity averages are taken over the indicated subgroups. The symbol Δ indicates the difference in the relevant variable between periods t and $t - 1$. For example, variable II in column 2 is the change in the average productivity of coworkers who are in the focal worker's observing set and have medium previous overlap with worker i . In this example, if there are no workers who are positioned in worker i 's observing set (i.e., no workers who are directly looking at i) who have medium previous overlap with i in $t - 1$ and t , this change is coded as zero. If there are no workers positioned in worker i 's observing set in just t or $t - 1$, then this variable is coded as missing. As in the baseline model in Table 2, all models include controls for the change in the number of workers on duty in a store over ten-minute intervals.

The results demonstrate that social pressure can partially internalize the negative externalities associated with a group production process. However, the findings in Sections VA(1) and VA(2) do not suggest that these workers display prosocial behavior or altruism, at least in the dimensions of effort 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.

VI. Conclusion

We find evidence of strong peer effects associated with the introduction of high-productivity workers into work groups: a 10 percent increase in coworker productivity results in a 1.5 percent increase in individual productivity. This effect manifests itself at precisely the time of entry of more productive coworkers, and is persistent. The finding of positive productivity spillovers is particularly surprising given that our data come from a group production process that is particularly prone to free-riding. Furthermore, while one may expect to see productivity spillovers in creative professions (for example, R&D or scientific research), it is more surprising to find them in a low-skill occupation where the tasks performed by workers are highly standardized.

These findings have important implications for wage setting. The hiring of a high-productivity worker raises total output directly because 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, and efficient compensation should take it into account. Furthermore, the finding that low-productivity workers are substantially more sensitive to coworker composition than high-productivity workers has implications for the optimal workplace organization. The mix of workers that maximizes productivity is the one where skill diversity in each shift is maximized.

The evidence we have assembled based on the spatial orientation of workers and the frequency of their interactions supports the hypothesis that the workers in our sample are responding to social pressure. Workers exhibit cooperative behavior only when they are observed by coworkers and when they are likely to interact with them again in the future. There is also evidence that workers free-ride when coworkers cannot easily notice. These findings imply that workers in this firm care about how they are perceived by their peers, either because they are subject to sanctions from reciprocating coworkers or because they feel shame when not exerting effort. While workers in this setting do not appear to be particularly altruistic, the 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 toward motivation has its counterpart in inertia.

APPENDIX: CONSTRUCTION OF STANDARD ERRORS

This appendix describes how we constructed the standard errors used in this paper. Our procedure involves creating ten datasets in addition to the original one. Each of these datasets is identical to the main dataset used in the analysis, except that rather than using the estimated fixed effects θ , for each dataset we draw a new vector of fixed effects from the distribution $N(\hat{\theta}, \hat{\Sigma})$, where $\hat{\theta}$ is the vector of estimated fixed effects and $\hat{\Sigma}$ is the variance-covariance matrix for $\hat{\theta}$, both obtained by fitting (7). For each of these ten datasets, every variable that depends on the estimated fixed effects will instead be constructed using the simulated version. In order to compute standard errors, we estimate every model 11 times, once for the original dataset, clustering on checker \times calendar date, and once for each of the simulated datasets. The variance of $\hat{\beta}$ in (6) is the sum of the between- and within-simulation variability. Specifically, standard errors are computed as $\sqrt{s_{\beta}^2 + \sigma_{\beta}^2}$, where s_{β}^2 is the sampling variance of the estimated fixed effects and σ_{β}^2 denotes the between-simulation variance, which is just the variance of the estimates across the ten simulated datasets, $\text{var}(\hat{\beta}_1, \dots, \hat{\beta}_{10})$. The resulting standard errors are not sensitive to increasing the number of simulated datasets beyond ten.

APPENDIX TABLE 1—ESTIMATED COEFFICIENTS FROM DYNAMIC RESPONSE MODELS

Estimated coefficients from equation (8)			Estimated coefficients from equation (14)			Estimated coefficients from equation (15)		
(1)			(2)			(3)		
$\hat{\beta}_7$	-0.008	(0.021)	$\hat{\beta}_7^G$	0.003	(0.017)	$\hat{\beta}_7^E$	-0.010	(0.017)
$\hat{\beta}_6$	0.000	(0.020)	$\hat{\beta}_6^G$	-0.002	(0.016)	$\hat{\beta}_6^E$	0.001	(0.016)
$\hat{\beta}_5$	-0.012	(0.020)	$\hat{\beta}_5^G$	-0.034	(0.016)	$\hat{\beta}_5^E$	0.015	(0.016)
$\hat{\beta}_4$	-0.014	(0.019)	$\hat{\beta}_4^G$	0.001	(0.015)	$\hat{\beta}_4^E$	-0.005	(0.016)
$\hat{\beta}_3$	-0.011	(0.019)	$\hat{\beta}_3^G$	-0.005	(0.016)	$\hat{\beta}_3^E$	-0.029	(0.016)
$\hat{\beta}_2$	-0.014	(0.020)	$\hat{\beta}_2^G$	-0.008	(0.016)	$\hat{\beta}_2^E$	-0.008	(0.015)
$\hat{\beta}_1$	0.007	(0.019)	$\hat{\beta}_1^G$	0.007	(0.015)	$\hat{\beta}_1^E$	0.025	(0.016)
$\hat{\beta}_0$	0.178	(0.019)	$\hat{\beta}_0^G$	0.168	(0.017)	$\hat{\beta}_0^E$	-0.001	(0.016)
$\hat{\beta}_{-1}$	-0.041	(0.020)	$\hat{\beta}_{-1}^G$	-0.032	(0.015)	$\hat{\beta}_{-1}^E$	-0.032	(0.015)
$\hat{\beta}_{-2}$	-0.007	(0.019)	$\hat{\beta}_{-2}^G$	0.003	(0.015)	$\hat{\beta}_{-2}^E$	-0.009	(0.016)
$\hat{\beta}_{-3}$	-0.020	(0.019)	$\hat{\beta}_{-3}^G$	0.005	(0.015)	$\hat{\beta}_{-3}^E$	-0.021	(0.016)
$\hat{\beta}_{-4}$	-0.021	(0.020)	$\hat{\beta}_{-4}^G$	-0.021	(0.015)	$\hat{\beta}_{-4}^E$	-0.021	(0.016)
$\hat{\beta}_{-5}$	0.018	(0.019)	$\hat{\beta}_{-5}^G$	0.034	(0.015)	$\hat{\beta}_{-5}^E$	0.014	(0.016)
$\hat{\beta}_{-6}$	0.017	(0.019)	$\hat{\beta}_{-6}^G$	0.005	(0.015)	$\hat{\beta}_{-6}^E$	0.003	(0.016)
$\hat{\beta}_{-7}$	0.006	(0.019)	$\hat{\beta}_{-7}^G$	-0.007	(0.016)	$\hat{\beta}_{-7}^E$	0.016	(0.017)

Notes: Coefficients in column 1 are plotted in Figure 2. Coefficients in columns 2 and 3 are plotted in Figure 5. See notes to Figures 2 and 5 for further details.

REFERENCES

- Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica*, 67(2): 251–333.
- Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2005. "Social Preferences and the Response to Incentives: Evidence from Personnel Data." *Quarterly Journal of Economics*, 120(3): 917–62.
- Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2007. "Social Incentives in the Workplace." Unpublished.
- Barnard, Chester. 1938. *Functions of the Executive*. Cambridge, MA: Harvard University Press.
- Barron, John M., and Kathy Paulson Gjerde. 1997. "Peer Pressure in an Agency Relationship." *Journal of Labor Economics*, 15(2): 234–54.
- Brock, William A., and Steven N. Durlauf. 2001a. "Discrete Choice with Social Interactions." *Review of Economic Studies*, 68(2): 235–60.
- Brock, William A., and Steven N. Durlauf. 2001b. "Interactions-Based Models." In *Handbook of Econometrics Volume 5*, ed. James J. Heckman and Edward Leamer, 3297–3380. Amsterdam: Elsevier Science, North-Holland.
- Brock, William A., and Steven N. Durlauf. 2002. "A Multinomial-Choice Model of Neighborhood Effects." *American Economic Review*, 92(2): 298–303.
- Charness, Gary, and Martin Dufwenberg. 2006. "Promises and Partnership." *Econometrica*, 74(6): 1579–1601.
- Carpenter, Jeffrey, Samuel Bowles, and Herbert Gintis. 2006. "Mutual Monitoring in Teams: Theory and Experimental Evidence on the Importance of Reciprocity." Institute for the Study of Labor Discussion Paper 2106.
- Encinosa, William E., III, Martin Gaynor, and James B. Rebitzer. 2007. "The Sociology of Groups and the Economics of Incentives: Theory and Evidence on Compensation Systems." *Journal of Economic Behavior and Organization*, 62(2): 187–214.
- Falk, Armin, Ernst Fehr, and Urs Fischbacher. 2001. "Driving Forces of Informal Sanctions." Institute for Empirical Research in Economics Working Paper 59.
- Falk, Armin, Urs Fischbacher, and Simon Gächter. 2003. "Living in Two Neighborhoods: Social Interactions in the Lab." Institute for the Study of Labor Working Paper 1381.

- Falk, Armin, and Andrea Ichino.** 2006. "Clean Evidence on Peer Effects." *Journal of Labor Economics*, 24(1): 39–57.
- Fehr, Ernst, and Simon Gächter.** 2000. "Cooperation and Punishment in Public Goods Experiments." *American Economic Review*, 90(4): 980–94.
- Fehr, Ernst, Simon Gächter, and Georg Kirchsteiger.** 1997. "Reciprocity as a Contract Enforcement Device: Experimental Evidence." *Econometrica*, 65(4): 833–60.
- Gaynor, Martin, James Rebitzer, and Lowell Taylor.** 2004. "Physician Incentives in Health Maintenance Organizations." *Journal of Political Economy*, 112(4): 915–31.
- Hamilton, Barton H., Jack A. Nickerson, and Hideo Owan.** 2003. "Team Incentives and Worker Heterogeneity: An Empirical Analysis of the Impact of Teams on Productivity and Participation." *Journal of Political Economy*, 111(3): 465–97.
- Hamilton, Barton, Jack Nickerson, and Hideo Owan.** 2005. "Diversity and Productivity in Production Teams." Unpublished.
- Hayes, Rachel M., Paul Oyer, and Scott Schaefer.** 2006. "Coworker Complementarity and the Stability of Top-Management Teams." *Journal of Law, Economics, and Organization*, 22(1): 184–212.
- Huck, Steffen, Dorothea Kubler, and Jörgen Weibull.** 2002. "Social Norms and Incentives in Firms." Unpublished.
- Ichino, Andrea, and Giovanni Maggi.** 2000. "Work Environment and Individual Background: Explaining Regional Shirking Differentials in a Large Italian Firm." *Quarterly Journal of Economics*, 115(3): 1057–90.
- Ichniowski, Casey, Kathryn Shaw, and Giovanna Prennushi.** 1997. "The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines." *American Economic Review*, 87(3): 291–313.
- Kandel, Eugene, and Edward P. Lazear.** 1992. "Peer Pressure and Partnerships." *Journal of Political Economy*, 100(4): 801–17.
- Knez, Marc, and Duncan Simester.** 2001. "Firm-Wide Incentives and Mutual Monitoring at Continental Airlines." *Journal of Labor Economics*, 19(4): 743–72.
- Koedel, Cory, and Julian Betts.** 2007. "Re-Examining the Role of Teacher Quality in the Educational Production Function." University of Missouri Department of Economics Working Paper 0708.
- Lazear, Edward P.** 2000. "Performance Pay and Productivity." *American Economic Review*, 90(5): 1346–61.
- Leonard, Jonathan S., and David I. Levine.** 2006. "The Effect of Diversity on Turnover: A Large Case Study." *Industrial and Labor Relations Review*, 59(4): 547–72.
- Lucas, Robert E., Jr.** 1988. "On the Mechanics of Economic Development." *Journal of Monetary Economics*, 22(1): 3–42.
- Manski, Charles F.** 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *Review of Economic Studies*, 60(3): 531–42.
- Manski, Charles F.** 2000. "Economic Analysis of Social Interactions." *Journal of Economic Perspectives*, 14(3): 115–36.
- Marshall, Alfred.** 1890. *Principles of Economics*. New York: Macmillan.
- Moffitt, Robert A.** 2001. "Policy Interventions, Low-Level Equilibria, and Social Interactions." In *Social Dynamics*, ed. S. N. Durlauf and H. P. Young, 45–82. Washington, DC: Brookings Institution Press.
- Moretti, Enrico.** 2004a. "Estimating the Social Return to Higher Education: Evidence from Longitudinal and Repeated Cross-Sectional Data." *Journal of Econometrics*, 121(1–2): 175–212.
- Moretti, Enrico.** 2004b. "Workers' Education, Spillovers, and Productivity: Evidence from Plant-Level Production Functions." *American Economic Review*, 94(3): 656–90.
- Paarsch, Harry J., and Bruce S. Shearer.** 1999. "The Response of Worker Effort to Piece Rates: Evidence from the British Columbia Tree-Planting Industry." *Journal of Human Resources*, 34(4): 643–67.
- Roy, Donald.** 1954. "Efficiency and the Fix: Informal Intergroup Relations in a Piecework Machine Shop." *American Journal of Sociology*, 60: 255–66.
- Rubin, Donald.** 1987. *Multiple Imputation for Nonresponse in Surveys*. New York: John Wiley.
- Tadelis, Steve.** 2007. "The Power of Shame and the Rationality of Trust." Unpublished.