ANNUAL REVIEWS

Social Incentives in Organizations

Nava Ashraf[1,2] and Oriana Bandiera[1,2]

[1]Department of Economics, London School of Economics, London WC2A 2AE, United Kingdom; email: n.ashraf1@lse.ac.uk, o.bandiera@lse.ac.uk

[2]Suntory and Toyota International Centres for Economics and Related Disciplines, London WC2A 2AE, United Kingdom

ANNUAL REVIEWS CONNECT

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

Keywords

Abstract

We review the evidence on social incentives, namely on how social interactions with colleagues, subordinates, bosses, customers, and others shape agents' effort choices in organizations. We propose a two-way taxonomy based on (*a*) whether the social group is horizontal (peers at the same level of the hierarchy) or vertical (individuals at different levels within or outside of the organization) and (*b*) whether the agent's effort creates externalities for the other members of their social group. We show settings in which social incentives improve productivity and settings in which they reduce it. In most cases, the size of the effect is approximately 10%, which is half of the typical effect of performance pay. We also show that social incentives can interfere with financial incentives, making them ineffective or even detrimental. We conclude that social incentives are a powerful motivator that must be taken into account in the design of organizational policies and that more research is needed to understand how policies can shape the preferences that underpin these incentives.

1. INTRODUCTION

Understanding how to motivate employees is the key challenge of every organization. The principal–agent model, a workhorse of agency theory, makes clear how this challenge can be met by offering financial incentives that align the agent's preference for pay with the organization's preference for performance.

Looking at organizations through the lens of the principal–agent model, however, obscures their raison d'être, which is to bring together the labor inputs of several individuals who interact with one another. The fact that these interactions shape workplace behavior has long been acknowledged by sociologists (Barnard 1938, Mayo 1933, Roethlisberger & Dickson 1939, Roy 1952) and, more recently, by economic theorists (Kandel & Lazear 1992, Rotemberg 1994). These interactions and the underpinning social preferences can create social incentives; that is, they can affect the marginal benefit or cost of effort and therefore shape individuals' motivation. They can also interact with financial incentives and shape individuals' responses to the latter.

This article reviews the empirical evidence on social incentives in organizations and their interaction with monetary incentives. We define a social incentive as any factor that (*a*) affects the marginal benefit or marginal cost of effort and (*b*) stems from interactions with others. Thus, social incentives can be underpinned by either nonstandard preferences (altruism, reciprocity, etc.) or social interactions between selfish agents. To guide our review, we develop a simple framework that allows us to organize and draw lessons from the evidence on social incentives in organizations. We begin by setting up the principal–agent model and derive the conditions under which social incentives emerge. The standard model assumes that agents' utility is increasing in pay and decreasing in effort. To generate social incentives, we need to augment preferences to include the outcomes of others within the organization (colleagues, bosses) or outside of it (clients, beneficiaries, potential hires). We follow the literature on social preferences and assume that the agent can care about others' outcomes either relative to the agent's own or for their own sake. The first type of preference includes preferences for fairness, inequity aversion, and dominance (Charness & Rabin 2002, Fehr & Schmidt 1999, Kranton et al. 2016). The second type includes positive preferences for others' welfare due to pure altruism, warm glow, or cooperation, as well as negative preferences due to spite or strategic punishment (Andreoni 1989, Falk et al. 2005, Fehr & Gachter 2000). In both cases, social preferences can be a structural parameter of the utility function or a reduced-form representation of the equilibrium of a repeated game among selfish agents. Social incentives can arise regardless of the deeper nature of social preferences.

The framework illustrates how the effect of social incentives depends on the structure of social preferences and how responses to social incentives can be used to infer the structure of preferences. The framework suggests a two-way taxonomy that we use to organize our review. The two dimensions are the identity of the social group and the link between agents' effort and the welfare of their social group. We show that, in all cases, the effect of social incentives on effort and the performance of the organization is ambiguous. In particular, agents with relative preferences can get discouraged if they perform worse than their comparison group, and agents whose effort affects the welfare of their subordinate can be led to benefit their in group at the expense of more deserving individuals.

We also show that social incentives interact with financial incentives, sometimes making the latter ineffective or even detrimental. For instance, incentives that create inequality within the peer group might demotivate the least-productive workers and further reduce their productivity. One notable exception is that financial incentives can ameliorate the targeting bias of bosses who favor their in group at the expense of others and, by doing so, harm the organization.

The review is organized as follows. Section 2 develops the theoretical framework, and discusses methods and scope. We focus on the analysis of personnel data and field experiments within organizations. Social incentives in other domains, for instance, charitable giving or consumers'

choices, and laboratory tests of social preferences are beyond the scope of this review. Section 3 reviews the evidence from contexts where the social group is made up of peers, with and without effort externalities. Section 4 focuses on contexts where the peer group is vertical or outside the firm. Section 5 discusses open questions on the origin of social preferences and whether these can be shaped by the policies of an organization.

2. THEORY AND METHODS

2.1. Theory

In the standard principal–agent model, there are two actors: a principal (the organization) and an agent (the employee). The principal hires the agent to produce output, which is increasing in the effort e exerted by the agent according to $f(.)$, where $f' > 0$ and $f'' < 0$. Effort entails a cost $d(e)$ for the agent, where $d' > 0$ and $d'' > 0$. The moral hazard problem arises because effort is unobservable, and output also depends on a shock ε with mean 0. Output is thus a noisy signal for effort equal to $f(e) + \varepsilon$.

The model provides a stark illustration of the misalignment of interests between the organization and its employees, as effort is desirable for the former but costly for the latter.

The principal offers a contract of the form $y(e) = w + bf(e)$, where w is a fixed wage and $bf(e)$ is a performance bonus. The principal chooses (w, b) to maximize

$$(1 - b)f[e^*(w, b)] - w,$$

where e^* solves the agent's utility maximization problem

$$max_e\, U = w + bf(e) - d(e),$$

and thus

$$b\frac{\partial f(e)}{\partial e} = \frac{\partial d(e)}{\partial e}.$$

In this stylized world, the agent is solely motivated by performance pay, that is, $e^* = 0$ (or the minimum feasible level) whenever $b = 0$. This follows from the assumptions that the effort of the agent only benefits the organization and that the agent places no value on this benefit.

Relaxing this assumption requires modifying the utility function of the agent U_i to capture social interactions, that is, how i's effort affects the utility of others and how the others' effort, actions, or traits affect i's own marginal return to effort. We do so by introducing a social interaction function $S(e_i)$:

$$U_i = w + bf(e_i) - d(e_i) + S(e_i).$$

Social interactions change the first-order condition for i's choice of effort as in

$$b\frac{\partial f(e_i)}{\partial e_i} + \frac{\partial S(.)}{\partial e_i} = \frac{\partial d(e_i)}{\partial e_i}.$$

The second term on the left-hand side captures social incentives, and it can be either negative or positive depending on the nature of the social interaction $S(.)$ and whether this increases or decreases the marginal benefit of effort.

This simple extension of the principal–agent model suggests a taxonomy of social incentives as a function of the identity of the social group and the nature of the interaction $S(.)$. The social group includes colleagues at different rungs of the hierarchy, other social relations outside the organization (family and friends), and downstream beneficiaries. $S_i(.)$ can take many forms depending on whether i's effort affects others and on their social preferences. For instance, if i's effort affects

j's welfare, then $S_i(.) = \sigma_{ij} U_j(e_i)$, where $\partial U_j / \partial e_i > 0$, and $\sigma_{ij} > 0$, commonly called the social preference parameter, measures the weight that agent i puts on the utility of j. The specification $S_i(.) = \sigma_{ij} U_j(e_i)$ also captures the case in which the good produced by the organization affects the welfare of beneficiaries outside it and in which i puts a positive weight on this benefit. This mirrors Besley & Ghatak's (2005) mission preferences, where the agent shares the interest of the organization, in this case improving the welfare of the beneficiaries of the service. In both of these examples, social interactions affect i's effort if i cares about others and if i's effort affects others' welfare.

Social interactions can affect effort, even if i's effort does not directly affect others, if social preferences are positional, that is, if i cares about how their effort, pay, or performance compares to j's. This can occur for multiple reasons, including if j's effort, pay, or performance creates a benchmark for i's progress. The literature has focused on three classes of preferences: inequality aversion, status seeking, and envy (for an overview, see Sobel 2005). Inequality aversion or conformity can be modeled as i's utility decreasing in the absolute difference between i and j:

$$S(.) = \sigma_{ij} |y_i - y_j|, \ \sigma_{ij} < 0.$$

In contrast, i has status-seeking preferences if their utility increases in the difference between i and j, but only if this is positive:

$$S(.) = \begin{cases} \widehat{\sigma_{ij}}(y_i - y_j) & \text{with} \quad \widehat{\sigma_{ij}} > 0 & \text{if} \quad (y_i - y_j) > 0 \\ 0 & & \text{if} \quad (y_i - y_j) < 0 \end{cases}.$$

In contrast, i has envy if their utility decreases in the difference between i and j and if this difference is negative:

$$S(.) = \begin{cases} 0 & & \text{if} \quad (y_i - y_j) > 0 \\ \widehat{\sigma_{ij}}(y_i - y_j) & \text{with} \quad \widehat{\sigma_{ij}} < 0 & \text{if} \quad (y_i - y_j) < 0 \end{cases}.$$

This simple setup highlights two useful facts. First, if social preferences are absolute, then social incentives can arise only if there is a link or externality between i's effort and j's outcomes. Indeed, when $S_i(.) = \sigma_{ij} U_j$, we have $\partial S / \partial e_i \neq 0$ only if $\partial U_j / \partial e_i \neq 0$. On the contrary, if preferences are relative, then social incentives can arise regardless of the existence of such externalities.

Second, the effect of social incentives on the marginal benefit of effort is ambiguous. When preferences are relative, the sign of the effect depends on the agent's position relative to others. For instance, agents who are inequality averse will reduce their effort if their outcome is above the norm but increase it if it is below the norm. When preferences are absolute, its sign depends on the sign of the externality that the agent creates and the sign of the weight that the agent puts on the welfare of others. If these signs are equal, then social incentives increase effort; if they are opposite, then they reduce it.

2.2. Methods

Our aim is to review field evidence on social incentives in organizations. We thus exclude two literatures: field evidence outside of organizations—e.g., charitable donations or health-seeking behavior—and evidence from the laboratory. Whilst social preferences play an important role in both of these literatures, our aim is to study how social preferences shape effort and organizational performance rather than to provide evidence on social preferences in general.

Research on social incentives in organizations uses three methods. The first relies on detailed personnel data from real-world organizations combined with a naturally occurring source of variation in social incentives. The second is based on collaborations between researchers and organizations to evaluate a commonly agreed policy via experimental methods. The third has

researchers set up an organization for the purpose of the experiment and hire workers for short-term jobs through either university boards or online platforms.

The three methods lie at different points on the continuum between realism and control. Realism is highest in naturally occurring data and lowest in purpose-built organizations, while the opposite is true for control; experiments with real-world organizations fall in between. Three dimensions of realism are particularly relevant for the study of social incentives. The first is the nature of the task: In real-world organizations, workers perform tasks that are appropriate for their skill level and that constitute their regular day-to-day job, while in purpose-built organizations, subjects, typically university students, are hired for occasional, low-skilled tasks such as cataloguing books, entering data, or stuffing envelopes. Performing tasks that do not require much attention might make employees more sensitive to social preferences, but the fact that the job is only temporary might make them respond less. The second dimension is the time horizon. Purpose-built organizations tend to be short lived, employees are subject to different treatments over a short period of time, and only short-run responses can be evaluated. This is relevant for social incentives because employees do not have time to establish relationships with other employees or with the employer, and thus, the effect of social incentives might be underestimated. The third dimension is the fact that the stakes are much higher in real-world organizations, as the employees' responses to social incentives affect their main source of income and possibly their career in the organization. This might heighten or dampen responses depending on whether the incentives are beneficial or costly in terms of income and career prospects.

The main advantage of creating organizations for the sole purpose of the experiment is that it allows full control: When researchers own the organization, they can implement any treatment (subject to ethical constraints) and can closely match the design of the experiment to theory, which allows for a more granular study of mechanisms.

The identification of social incentives is related to but distinct from the identification of peer effects. The latter aims to uncover the effect of a group outcome on the same outcome of a member of the group, for instance, the effect of the class average scores on the test score of a student in that class. Manski (1993) shows that the correlation captures the effect of interest, known as the endogenous peer effect; other students' traits, known as exogenous peer effects; and common shocks such as teacher quality, known as correlated effects, on individual student performance.

Most of the papers that we review aim to identify the equivalent of exogenous peer effects in Manski's taxonomy. The challenge is to separate these from correlated effects. This is achieved either by finding a plausibly exogenous source of variation or by creating it by means of a field experiment.

3. HORIZONTAL SOCIAL RELATIONSHIPS

3.1. Settings Without Externalities

When the agent's effort does directly affect the outcomes of their peers, social incentives can arise only if the agent has relative preferences, either because they care directly about their relative performance or because they can infer information on the marginal benefit of effort from their comparison with others. Field experiments on this topic rely on two sources of variation—either creating actual differences or informing subjects of naturally occurring differences.

Breza et al. (2018) and Cohn et al. (2014) provide recent examples of the first type. Breza et al. (2018) collaborate with an Indian manufacturing firm to introduce variation in pay within groups of workers who work together but do not have complementarities. They randomize teams to receive either the same wages for all workers or heterogeneous wages such that each team

member is paid a different wage. This allows them to identify the effect of relative pay on effort by comparing workers who receive the same absolute wage but work with coworkers who receive different wages. Workers reduce output by 12% when their coworkers are paid more than they are, and they are also 13.5 percentage points less likely to go to work. In contrast, Breza et al. find no effect when coworkers are paid less. Thus, pay inequality demoralizes the weakest workers without boosting the strongest, and as such, it is unambiguously bad for the firm. Interestingly, when workers are given a plausible reason for the pay differences (the fact that they are based on baseline productivity levels), the negative impact on the lowest-paid workers disappears.

Cohn et al. (2014) collaborate with a firm that hires workers for one-off sales promotions. Workers worked in groups of two performing identical individual tasks, again without complementarities. The experiment has two stages. In the first stage, workers in the same group are paid the same. In the second stage, the authors cut by 25% either the wages of both workers or the wage of just one of the two. They find that cutting both wages caused a decrease in performance by 15% relative to the control group, whose wages were not cut. In line with the existence of relative preferences, cutting only one group member's wage by 25% caused a decrease in the performance of the affected worker by 34%.

Blanes i Vidal & Nossol (2011) use the second type of identification; that is, they inform workers of their position in the distribution of performance and pay. They find that employees respond positively to the provision of rank information: Productivity increased by 7% once the feedback policy was announced, and productivity remained at that higher level.

Taken together, these papers suggest that social incentives are a significant factor that shapes the choice of effort and productivity. The fact that social incentives arise in settings without externalities suggests that workers have relative preferences. Two out of the three papers find a negative effect for the workers at the bottom of the pay for performance distribution, while one finds a positive effect throughout. The discrepancy might be due to the fact that workers in the experiments of Breza et al. (2018) and Cohn et al. (2014) were hired for a short-term job, and therefore, all repeated games considerations that arise in long-term relationships were muted. These effects might undo any negative effects on pay differentials, as shown by Blanes i Vidal & Nossol (2011). Comparing across contexts is, of course, of limited use, as the time horizon is only one of the possible differences driving different responses. Understanding the conditions under which social incentives underpinned by relative preferences can benefit the organization remains an open area for future research.

3.2. Settings with Externalities

The earliest evidence on social incentives in economics came as a by-product of economists' interest in the employee stock options (ESOPs) and profit sharing plans that became popular in the 1980s. Using panel data on US and Japanese firms, respectively, both Jones & Kato (1995) and Kruse (1993) show that the adoption of these plans was associated with a 4–5% increase in productivity. For comparison, most of the field and lab experiments that identify the causal impact of individual bonuses on productivity find a 20–25% increase (Bandiera et al. 2017a). The evidence is suggestive of social incentives because, in their absence, profit sharing and ESOPs should be ineffective given that each agent bears the entire cost of their effort while receiving a trivial fraction of the marginal benefit, that is, $1/N$, where N often equals several tens of thousands. The fact that ESOPs and profit sharing plans are correlated with productivity is consistent with the assumption that agents internalize the externality, namely, that the numerator is larger than 1 because they value the benefits accruing to others. The nature of the data, however, does not allow for establishing causality.

Following from this literature, several papers using more detailed data from individual firms uncover widespread use of collective or team incentives that cannot be rationalized without social incentives (Bartel et al. 2011, Boning et al. 2007, Gaynor et al. 2004, Griffith & Neely 2009, Hansen 1997). Hamilton et al. (2003) identify the effect of shifting from individual production and incentives to team production with team incentives using individual-level personnel data from a textile factory. The standard model with selfish agents predicts that team incentives lead to free riding and lower productivity. In contrast, Hamilton et al. (2003) show that productivity increases. This is consistent with social incentives but could equally be driven by the improvement in technology associated with the shift to team production.

3.2.1. Social incentives in personnel data. Two papers exploit precise, high-frequency productivity data from personnel records to isolate the causal effect of social incentives on productivity. Like many authors of earlier studies, Bandiera et al. (2005) analyze a setting where the incentive scheme generates an externality between workers; in the case of their setting, the externality is negative. Their setting is a farm where workers pick fruit in a field with approximately 40 other workers, and their pay equals a piece rate times the quantity of fruit picked (quality adjusted). The piece rate is relative; namely, it is set by dividing a fixed wage by the average productivity of workers in that field on that day. This scheme creates a negative externality because, by exerting effort, a worker increases average productivity and lowers the piece rate for everyone in the field. To identify social incentives, Bandiera et al. (2005) collaborated with senior management to replace the relative scheme with a standard piece rate set independently of workers' productivity. This allows them to identify social incentives by comparing the productivity of the same workers under the two schemes. Without social incentives, the difference should be nil because workers do not internalize the effect of their effort on others and thus take the piece rate as fixed in both cases. In contrast, Bandiera et al. (2005) find that the productivity of the average worker was 50% higher under the absolute scheme, suggesting that workers were withholding effort when it damaged their colleagues. To quantify the effect of social incentives, Bandiera et al. (2005) calibrate the social preference parameter σ for each worker. They find that $\sigma > 0$ for 98% of the workers and that the average worker weighs his colleagues' pay at two-thirds of his own. This implies that, while social incentives affect effort choices, they are not strong enough for workers to reach the Pareto optimum.

Mas & Moretti (2009) analyze a setting where workers' effort generates a positive externality on their colleagues, this time through the production function rather than through the pay scheme. Their setting is a supermarket chain where workers operate cash tills and receive fixed wages, and their utility decreases as more customers form a queue at their till. Workers who take on more customers benefit their colleagues because they shorten their queues. Supermarkets stagger the change of shifts to ensure that a large number of tills stay open. This generates groups that differ in innate ability, which Mas & Moretti (2009) exploit to identify the effect of increasing colleagues' ability on workers' productivity. They find that the arrival of a fast worker increases the effort of other workers on the same shift by 1%. This cannot be explained without some form of social incentive. If workers chose their effort without taking into account the externality on others, then the size of the externality should not matter. The fact that workers speed up when the externality that they impose on their colleagues is larger (because faster workers would attract more customers, other things equal) suggests that they must internalize the externality to some extent.

3.2.2. Implications for performance and personnel policies. The theoretical framework makes clear that social incentives increase effort and benefit the organization when the

externality is positive. While none of the studies discussed above compares the behavior of the same workers with different types of externalities, the cross-study comparison is consistent with the theoretical prediction.

Positive externalities are much more common, and all but one of the studies reviewed above examine settings with positive externalities through the pay scheme (Jones & Kato 1995, Knez & Simester 2001, Kruse 1993), the production function (Mas & Moretti 2009), or both (Hamilton et al. 2003). In all of these cases, the fact that workers internalize the positive externality, at least to some extent, leads to more effort and higher performance for the organization. The only study with a negative externality (Bandiera et al. 2005) provides a useful falsification test, as workers internalize the externality by reducing effort, and this reduces aggregate performance.

These results have implications for the choice of incentive schemes and other personnel policies, highlighting the fact that policies' effectiveness depends on the extent to which they leverage social incentives. Bandiera et al. (2005) use their estimates to calibrate the scheme that maximizes profits and show that team pay would lead to 30% higher effort at the same cost to the principal; this is entirely due to social incentives because there are no complementarities in production.

3.2.3. Implications for social preferences.
Besides identifying social incentives, these findings can be used to infer the nature of the social preferences that underpin them.

In a case study of profit sharing at Continental Airlines, Knez & Simester (2001) argue that the relevant in group is the team of workers at the same airport and that the introduction of a profit sharing scheme improved performance because workers internalized the effect of their performance on the pay of their immediate colleagues in the same airport.

Bandiera et al. (2005) collect data on friendship networks among workers. They find that, when the pay scheme generates a negative externality, the average worker slows down by 21% if all of their friends work on the same field, relative to when none of their friends are there. In contrast, the presence of friends has no effect when the pay scheme does not generate an externality. The fact that the extent to which workers internalize the externality depends on the size of the in group indicates that social preferences are stronger for the members of the in group. Mas & Moretti (2009) define the in group on the basis of familiarity, measured by the frequency with which workers work on the same shift. They find that the effects only materialize when familiarity is high, which, again, is consistent with heterogeneous preferences.

Both sets of results are consistent with different motives for social preferences. Workers might have pure altruism and care about the welfare of their colleagues, or helping others might give warm glow. Otherwise, they might be selfish but able to sustain cooperation through repeated interaction. Mas & Moretti (2009) and Bandiera et al. (2005) exploit variation in visibility to disentangle these alternatives. Bandiera et al. (2005) observe the same workers picking fruit that grows in short plants and fruit that grows in tall shrubs. The cash tills in Mas & Moretti (2009) face sideways so workers can see those in front of them but can only be seen by those behind. Bandiera et al. (2005) and Mas & Moretti (2009) find that workers internalize the externality only when their colleagues can see them. This rules out pure altruism and warm glow because visibility does not affect the externality. It is consistent with impure altruism based on reputation concerns and also with the ability to sustain cooperation through repeated interactions.[1]

---

[1]It is important to note, however, that these findings do not imply that visibility is a necessary condition for cooperation. In both the setting analyzed by Bandiera et al. (2005) and that analyzed by Mas & Moretti (2009), the effect of visibility is identified using within-worker variation, so that each worker is exposed to both cases with and without visibility and only cooperates in the former. If there were no visibility at all, it is possible that workers would find ways to monitor each other and cooperate nonetheless.

4. VERTICAL SOCIAL RELATIONSHIPS

4.1. Vertical Social Groups: Bottom Up

The core assumption of the principal–agent model is that agents do not internalize the welfare of the organization unless they are paid to do so. This assumption might fail if the agent has social preferences. The literature has focused on two violations of this assumption: the agent having reciprocal altruism toward the employer, as in models of gift exchange, or deriving utility from fulfilling the mission of the organization. Contrary to the literature on social preferences among peers, there has been no attempt to test whether relative comparisons across the hierarchy affect effort.

4.1.1. Gift exchange.
A well-established theoretical literature starting with the work of Akerlof (1982) models the employer–employee relationship as a gift exchange, where the employers buy the employees' goodwill by offering good conditions, and the latter reciprocate by providing high effort.

The theory of gift exchange has been subject to several tests in the laboratory, but field evidence remains slim, especially in organizations that exist independently of the research. Gneezy & List (2006), DellaVigna & Pope (2018), and DellaVigna et al. (2016) set up their own organizations to test the relevance of gift exchange in the field. They hire workers for tasks ranging from fundraising to stuffing envelopes and generate exogenous variation in wages by randomly offering some workers more than originally agreed. Of the three studies, only that of Gneezy & List (2006) finds an effect, but this effect turns out to be only temporary.

It is key to note, however, that these ad hoc organizations are short-lived, so that employers and employees do not have the time to establish a reciprocal agreement or to develop social preferences as a consequence of many acts of good will, as postulated by Akerlof (1982). We need evidence from long-term employment relationships to establish whether the effect of gift exchange is indeed short-lived.

4.1.2. Identity and mission.
In an influential line of work, Akerlof & Kranton (2010) highlight the motivating power of identity and show how many commonly used management practices can be ascribed to the organization's desire to make the employees feel that they belong and, therefore, care about the success of the organization.

Whilst Akerlof & Kranton (2010) assume that identity is created once an individual joins an organization, Besley & Ghatak (2005) suggest that it emerges as an equilibrium of the matching process between individuals and organizations, as individuals who care about a given mission sort into organizations that pursue that mission.

To provide evidence on mission incentives, researchers have followed two strategies. The first consists of setting up an organization for the purpose of the experiment and randomly varying whether workers are offered jobs with a mission, typically by changing the identity of the employer from a for-profit organization to a charity or nongovernmental organization (NGO). The evidence suggests that social incentives motivate workers; that is, workers exert more effort when their effort benefits a mission that they care about (DellaVigna & Pope 2018; Tonin & Vlassopoulos 2010, 2015).

The second strategy relies on collaboration between researchers and real-world organizations that have a social mission. In this case, the nature of the job is fixed, but researchers can test whether agents who care more about the mission exert more effort. For instance, Ashraf et al. (2014) study the motivation and performance of agents hired by a public health organization to sell condoms in Lusaka, Zambia. They measure how much each agent cares about the mission

through a lab game and then show that this experimental measure of social preferences is strongly correlated with sales performance.

4.1.3. Implications for performance and personnel policies. Social incentives unambiguously increase effort and the performance of the organization when the employees care about the mission. However, the implications for the effectiveness of financial incentives are not clear cut. A hypothesis that has received considerable attention in the social sciences is that the two motives are substitutes, so that financial incentives reduce or crowd out the effect of social incentives. In contrast to the large theoretical and experimental literature (see Bowles & Polania-Reyes 2012), tests of crowding out in organizations are rare. An exception is the work of Ashraf et al. (2014), who randomly allocate agents to incentive treatments and test whether the effects of these interact with the agents' social preferences, measured as described in Section 4.1.2. Ashraf et al. (2014) find evidence of crowding in, that is, that both financial and nonfinancial rewards are more effective for agents with stronger social preferences.

Three recent papers test for crowding out on the extensive margin, that is, whether offering financial incentives attracts agents with weaker social preferences. Dal Bó et al. (2013) and Ashraf et al. (2016) randomize the offer of incentives when recruiting civil servants in Mexico and health workers in Zambia, respectively. In both cases, the evidence indicates that stronger incentives attract higher-quality applicants without displacing social preferences. Deserranno (2017) employs a similar design to hire NGO workers for a job that consists of two tasks: basic goods sales (soap, oil, etc.) and health promotion. She finds that stronger financial incentives signal that the sales component is more important and thus attract agents who are more interested in this component and have weaker social preferences.[2]

4.2. Vertical Social Groups: Top Down

The study of social incentives of individuals at higher levels of the hierarchy toward individuals at lower levels raises the key issue of effort allocation. Indeed, in most cases, agents at the top need to allocate effort or other resources to agents below them, and social preferences can affect the allocation in addition to affecting the level. In the **Supplemental Appendix**, we model the level and allocation choice of a manager who manages two workers of different ability and for whom the manager has different social preferences. The model makes clear that social incentives have an effort boost effect that is always positive, that caring about at least one worker makes the manager work harder, and that the targeting effect can be negative if social incentives increase the effort devoted to the worker who is socially connected to the manager at the expense of the more able worker.

The model also makes clear that financial incentives that give the manager a stake in the performance of the organization can ameliorate the negative effects of social incentives.

Three recent papers study whether social incentives affect the behavior of agents who can benefit others within the organization via the allocation of either effort or other resources. The three papers all exploit natural variation in group composition that allows them to observe the same agent when they belong to the in group of the agent making the allocation and when they do not.

4.2.1. Social incentives and the allocation problem. Bandiera et al. (2009) study how managers allocate effort across workers in a fruit farm. In their setting, managerial effort makes workers

---

[2]The fact that, with incomplete information, incentives can act as signals also explains many instances of crowding out in laboratory experiments (see Bowles & Polania-Reyes 2012).

more productive and, as workers are paid piece rates, increases their pay. To identify the effect of social connections, they exploit the fact that the same worker is observed working with different managers. They find strong evidence that social incentives are at play: A worker's productivity and pay are 9% higher on days when the worker is managed by a manager that they are socially connected to, relative to days when they are managed by someone that they are not connected to.

Hjort (2014) analyzes social incentives in a flower packing firm in Kenya, where employees belong to rival ethnic groups. Employees work in teams of three, where an upstream manager needs to allocate flowers to one of two downstream workers. As in the study of Bandiera et al. (2009), targeting implies higher pay, and social incentives are identified from naturally occurring differences in group composition from day to day. Hjort (2014) finds that, when downstream workers have different ethnicities, those who share the same ethnicity as the upstream worker earn 24% more.

Xu (2017) studies how the Secretary of State of the British Empire allocated colonies to governors. Salary and amenities varied considerably across colonies and were positively correlated, so that some colonies were strictly preferable to others. These colonies also yielded larger revenues, and there was thus a complementarity between the governor's skill and the colony's revenue potential. Xu (2017) finds that social incentives affected allocation: A given governor got a more profitable or desirable colony when they were connected to the Secretary of State in charge in that year.

The three papers find evidence that social incentives shape the allocation of effort and resources in three very different contexts. In all three cases, social incentives are underpinned by partial altruism, as subordinates are favored when they belong to the in group of the agent in charge of allocating resources.

4.2.2. Implications for performance and personnel policies.
The outstanding question is whether this form of social incentive is good for the organization. The theoretical framework makes clear that the answer depends on the balance of three factors: Social incentives (*a*) increase effort toward the in group, (*b*) increase the effectiveness of effort targeted to the in group, and (*c*) decrease effort toward the out group. The aggregate effect is positive if the first two dominate the third.

To provide evidence, Bandiera et al. (2009) engineer an exogenous change in managerial incentives that gives managers a bigger stake in the success of the firm by offering them a bonus based on the productivity of the managed workers. The theoretical framework makes precise how changes in incentive power *b* can be used to sign the effect of social incentives on the performance of the organization. The intuition is that an increase in *b* aligns the interest of the manager with the organization; thus, if devoting more effort to worker 1 is beneficial to the organization as a whole, increasing *b* will increase that effort and vice versa. Bandiera et al. (2009) find that, when managers are offered steeper incentives, they reallocate effort from their socially connected workers to high-ability workers, which indicates that their original allocation did not maximize the firm's productivity.

Hjort (2014) exploits the fact that he observes both teams made up entirely of workers of the same ethnicity and teams made up of different ethnicities. Since social incentives are at play only in the latter, he can identify their effect by comparing the average productivity of ethnically homogeneous groups to that of ethnically heterogeneous groups. He finds that the latter are 8% slower, suggesting that social incentives reduce productivity. The fact that the difference disappears when the firm changes the incentive scheme so that the upstream worker cannot help their friends downstream suggests that the difference was driven by social incentives rather than differences in unobservables between heterogeneous and homogeneous groups.

Xu (2017) collects measures of colonies' performance to test whether being connected to the Secretary of State makes governors more productive. Since the goal of the governor was to maximize revenues for the Crown, Xu (2017) collects data on revenues that reveal that governors raise less revenue when connected to the Secretary of State, therefore ruling out the idea that appointments driven by social incentives benefitted the organization. To corroborate this interpretation, Xu (2017) also shows that, when allocation rights are shifted from the Secretary of State to an independent commission, social connections to the Secretary of State play no role, and the matching between colonies and governors becomes more assortative; that is, more effective governors are assigned to better colonies.

Taken together, these three papers show that social incentives distort the allocation of effort and that this damages productivity. Yet these are settings in which the tasks that workers undertake are relatively simple, and thus the benefits of social connections in terms of better communication or enforcement are likely to be small. In settings where these benefits are more relevant, the allocative effect of social connections on managerial effort might be beneficial. We present some such evidence below.

5. SOCIAL GROUPS OUTSIDE THE ORGANIZATION

5.1. Potential New Hires

Besides affecting the welfare of colleagues within the same organization, agents' effort also affects individuals outside the organization, including potential new hires, customers, and downstream beneficiaries.

One of the main channels through which employees' effort affects the welfare of agents outside the organization is through hiring and referrals. The effect of social incentives on hiring and referrals is ambiguous for reasons similar to those discussed above. When individuals' preferences display limited altruism, social incentives lead them to hire or refer members of their in group. Social incentives might be beneficial for the organization if the agent has better information about their in group or more effective tools to solve the moral hazard problem, but they might be detrimental if they crowd out a more deserving or productive out group.

Giuliano et al. (2009) use personnel data from a large US retail firm to test whether the race of the hiring manager affects the racial composition of new hires. Exploiting manager changes at several stores, they find that managers tend to hire workers of their same race. The main cleavage is between Black and non-Black managers; that is, Black managers are more likely to hire Black workers, while non-Black (White, Hispanic, and Asian) managers are more likely to hire non-Black workers. While this is consistent with social incentives underpinned by limited altruism toward a racial in group, Giuliano et al. (2009) were not able to test whether this is beneficial for the firm.

A closely related context is that of job referrals. To test whether social incentives shape referrals, Beaman & Magruder (2012) create short-term jobs in a laboratory in the field in Kolkata, India and ask their employees to refer others. The relevant in group in this setting is the extended family, and indeed, Beaman & Magruder (2012) find that employees often refer relatives. To test whether social incentives benefit the firm, Beaman & Magruder (2012) vary the power of financial incentives: Some employees are randomly offered a flat finder's fee, while others are paid a fee that is positively linked to the performance of the referral. As in the work of Bandiera et al. (2009), strengthening financial incentives aligns the interests of the referee with that of the firm. Beaman & Magruder (2012) find that this leads employees to refer coworkers rather than family members and that the former perform better. The findings thus suggest that social incentives distort the choice of referrals from high-ability workers to family members and that high-powered financial

incentives counteract this effect. This, however, is a setting in which, aside from the financial incentives, the employee who makes the referral does not work with the referred employee in the long run, so, by definition, the fact that in group members might be better able to address the moral hazard problem plays no role in this case.

We are not aware of any study that directly tests whether the effect of social incentives on referrals benefits the organization, but two recent papers show that referred workers are more profitable for the firm than workers hired without referrals, even when the referrers are not offered financial incentives, and thus, presumably, their choices are shaped by social incentives.

The first paper is by Burks et al. (2015), who use personnel data from nine large firms in three sectors: call centers, trucking firms, and high-tech firms. They find that, compared to nonreferred workers, referred workers are more profitable because they are less likely to drop out, and thus, they reduce turnover and affiliated costs. The second paper is by Pallais & Sands (2016), who set up a series of field experiments in an online labor market to identify the individual channels through which social incentives can benefit the firm. They hire workers on an online platform and invite them to provide referrals without offering any reward. They find evidence that employees hired through referrals are of better quality (referrals contain information about the workers' quality that is not otherwise observable by the employer) and that referred workers perform substantially better when paired with their own referrers.

The evidence suggests that the effect of social incentives goes beyond the borders of the firm as it shapes its hiring policy. There is evidence of limited altruism, as hiring managers and referrers seem to prefer members of their in group. However, while there is evidence that workers hired through referrals are more productive, we have no direct evidence that this is due to social incentives. What we do know, from one experiment (Beaman & Magruder 2012), is that financial incentives that reward the referrer for the productivity of the referred make them refer more productive workers at the expense of family members. More research is needed given the importance of referrals in the job market: In the United States, one in two jobs is found through referrals (Burks et al. 2015), and the rate is likely to be higher in poorer countries where labor markets are thinner.

5.2. Customers

Besides affecting hiring decisions and, therefore, the welfare of potential hires, agents in organizations directly affect the welfare of customers or, in the case of service delivery, beneficiaries. Because of time or stock constraints, agents might favor those in their in group. This can be beneficial to the organization if, as before, members of the same group can better solve problems deriving from asymmetric information. It can be detrimental if the out group has a higher willingness to pay or a greater need for the service.

Leonard et al. (2010) use data from 800 outlets of the same US retail chain to test whether the race match between the sales force and customer base affects sales. In principle, sales representatives might devote more time to customers of their same race, but customers might also be more likely to buy from someone of their same race. Using variation in sales force composition in the same store over time, Leonard et al. (2010) find very modest effects.

Fisman et al. (2017) use data on individual loans issued by over 4,000 employees of large Indian state banks over 6 years. Fisman et al. (2017) exploit the bank's rotation policy to test whether loan officers give out more loans when they work in branches whose customers belong to the same religion or caste group. They find strong evidence of social incentives: The total amount of new loans to borrowers in a religion or caste group increases by 6.5% when the officer assigned to the branch belongs to the same group. In-group officers also increase the number of new loan recipients by 5.7% and the probability that a member of the group receives any credit by 2.5%. To

test whether social incentives benefit the organization, Fisman et al. (2017) track the performance of loans during and after the term of the loan officer. They find that both improve, suggesting that targeting the in group improves both selection and moral hazard.

The difference in the effect of social incentives in retail and banking is consistent with our simple framework. In retail, adverse selection and moral hazard are irrelevant, and there is no constraint on the quantity of goods that can be sold. Thus, there is no underpinning for social incentives, as there are no benefits or costs to targeting the in group. The opposite is true for banking. The comparison between the two sectors thus illustrates that social incentives affect the allocation of goods and services when there is an advantage to the agent or the beneficiaries.

5.3. Beneficiaries

Social incentives can play an important role in the allocation of public services, especially in low-income countries where delivery is typically delegated to local agents who are embedded in a network of social relationships.

Bandiera et al. (2017a) study how social incentives affect the choices of agents who deliver agricultural extension services (training and improved seeds) in Uganda. The aim of these services is to improve the productivity of the poorest and least-productive farmers. To study social incentives, Bandiera et al. (2017a) collaborate with the implementing NGO to randomly select one out of two eligible candidates for the role of delivery, thus generating random variation in the in-group membership of the farmers.[3]

Bandiera et al. (2017a) find strong evidence that social incentives lead to a redistribution of resources from the out group to the in group: As expected, agents put a positive weight on the welfare of their in group. To investigate what drives social preferences in this context, Bandiera et al. (2017a) test the hypothesis that conflict generates parochial altruism, namely, that it increases altruism toward the in group while at the same time increasing discrimination toward the out group (Bauer et al. 2016). They show that delivery agents favor their in group only in villages with political cleavages. In the same villages, they discriminate against the out group. This distorts the allocation of services from poor farmers to farmers connected to the delivery agent, regardless of their wealth level. Calibration of a simple model reveals that this behavior is consistent with negative altruism or spite toward the out group.

6. DISCUSSION

The study of social relations in the workplace has produced a wealth of evidence that social preferences, whether in structural or reduced form, underpin social incentives and that these shape agents' choice of effort and the performance of the organizations. The studies we review, summarized in **Tables 1–4**, suggest that social incentives are a key determinant of employee motivation and that they can be used to the benefit of the organization. Three patterns emerge clearly from the tables. First is that the effect of social incentives on productivity is economically significant. Most studies find effects between 7% and 16%. To put these numbers into context, it

---

[3]Compared to the within-agent variation used in all other studies, this design allows one to quantify the redistributive effect of social incentives by comparing farmers connected to the winning candidate to those connected to the losing candidate. This comparison is not informative with observational data because connected and nonconnected agents differ on unobservables that are likely to be correlated with the outcome of interest. The experimental design overcomes this obstacle by creating a valid counterfactual for connected agents in the absence of connections.

Table 1   Studies of social groups made up of peers

| | Study details | | | | | | Findings | Implications | | |
Reference	Organization	Principal and time horizon	Agents	Social group	Externality	Source of variation		Social preferences	Performance	Financial incentives
Breza et al. 2018	Manufacturing firm, India	Researchers, 35 days	Seasonal workers, hired daily	Three team members	None	Randomize workers into equal pay teams (three wage levels) and unequal pay teams	Workers with lower relative pay are less likely to show up for work, and, conditional on working, they are 12% slower. Workers with higher pay are unaffected. The effect disappears when pay differentials are justified.	Envy	Negative	Incentives that create unfair inequality can demotivate low-productivity workers
Cohn et al. 2014	Sales agency, Germany	Researchers, 4 days	Individuals hired for a one-off sales promotion	Two team members	None	Randomize workers between three groups: no pay cut, uniform pay cut, uneven pay cut	The uniform pay cut reduces all workers' productivity by 15%; the uneven pay cut reduces the productivity of the affected worker by 34%. Thus, envy reduces productivity by 15%.	Envy	Negative	Incentives that create unfair inequality can demotivate low-productivity workers
Blanes i Vidal & Nossol 2011	Retail and wholesale organization, Germany	Organization, 3 years	Full-time warehouse workers	None	None	Natural variation in feedback policy on productivity and pay rank	Workers increase productivity by 7% as soon as the feedback policy is announced; the increase is sustained for at least 3 months.	Rat race and status seeking	Positive	Incentives that create competition among peers can increase productivity

(Continued)

Table 1 *(Continued)*

	Study details						Findings	Implications			
Reference	Organization	Principal and time horizon	Agents	Social group	Externality	Source of variation	Findings	Social preferences	Performance	Financial incentives
Bandiera et al. 2005	Fruit farm, United Kingdom	Organization, 3.5 months	Seasonal workers, hired daily over 4 months	Approximately 40 pickers on the same field	Negative under the relative incentive regime; none under the piece rate regime	Natural variation in group composition and experimental variation in incentive scheme	Worker productivity is 50% lower under relative incentives relative to piece rates and are even lower if the group contains several friends.	Altruism and cooperation	Negative	Incentives that create negative externalities backfire if agents internalize the externalities
Mas & Moretti 2009	Supermarket chain, United States	Organization, 2 years per store	Cashiers paid hourly by number of shifts	Approximately 7 cashiers on duty per 10-minute interval	Positive: faster cashiers shorten the queues at their colleagues' tills	Natural variation in group composition by ability and familiarity	Worker productivity is 1% higher when a worker with above-average productivity enters the shift relative to times when a below-average worker enters the shift. The effect is 6–7% when a worker already on duty is more familiar with the entering cashier.	Altruism and cooperation	Positive	Incentives that create positive externalities on team members might be more effective than individual incentives

The header spans: "Study details" spans Organization through Source of variation. "Implications" spans Social preferences, Performance, Financial incentives.

Table 2 Studies of social groups at other layers of the hierarchy: subordinates

| | Study details | | | | | | | Implications | | |
Reference	Organization	Principal and time horizon	Agents	Social group	Externality	Source of variation	Findings	Social preferences	Performance	Financial incentives
DellaVigna et al. 2016	Charities and a grocery store, United States	Researchers, 1 day	Individuals hired for one-off lab task	Beneficiaries of organization	Positive: effort translates into charity revenue	Randomized order of rounds across sessions; within session, randomized between four gift exchange treatments: control, positive monetary, negative monetary, and in-kind	When the piece rate is held constant, participants stuff 10% more envelopes when working for charity than when envelopes are discarded; productivity does not change when the return to the charity is increased. Overall, none of the gift treatments have a significant effect.	Preferences for the mission	Positive	NA
Tonin & Vlassopoulos 2010	University lab, United Kingdom	Researchers, 2 days over 2 weeks	Students hired for two sessions of a lab task	Beneficiaries of organization	Positive: effort translates into charity donation	Randomized three pay schemes at second session: no change, baseline pay plus lump sum donation, or baseline pay plus performance-based donation to charity	Productivity increases by 7–8% with lump-sum donation relative to control; no difference is observed between lump-sum and performance-based pay.	Preferences for the mission	Positive	None
Tonin & Vlassopoulos 2015	University lab, United Kingdom	Researchers, 4 days over 1 week	Students hired for four sessions of a lab task	Beneficiaries of organization	Positive: effort translates into charity donation	Randomized four pay schemes: constant pay over sessions, varying piece rate, and piece rate with one of two varying conditions of donation to charity	Productivity increases by 13%, relative to control, with exogenously set charity donation and 26% when participants can choose the division of the piece rate.	Preferences for the mission	Positive	None

(Continued)

Table 2   (Continued)

	Study details							Implications		
Reference	Organization	Principal and time horizon	Agents	Social group	Externality	Source of variation	Findings	Social preferences	Performance	Financial incentives
DellaVigna & Pope 2018	Online labor market, United States	Researchers, 1 day	Individuals hired for one-off task	Beneficiaries of organization	Positive in charity treatment: effort translates into charity donation	Randomized into one of 18 treatments that span varying piece rate, behavioral treatments (e.g., appeal to social preferences), and psychology-based treatments	Social comparison increases performance by 16–21.5% relative to the benchmark of no pay. Task significance treatment increases performance by 14%. Introducing charitable giving increases performance by 25%, but the return to the charity does not matter. An unexpected bonus gift increases performance by 5%.	Preferences for the mission and status seeking	Positive	None
Ashraf et al. 2014	Public health NGO, Zambia	Organization, 1 year	Hairdressers hired to sell condoms indefinitely	Beneficiaries of organization	Positive: effort translates into reduction in HIV transmission	Randomized four treatments: no incentives, large financial margin, small financial margin, and nonfinancial reward (stars) with cross-sectional variation in motivation	Agents who have prosocial motivation above the median sell 48% more condoms.	Preferences for the mission	Positive	The effect of both the financial and nonfinancial incentives is stronger for these prosocially motivated agents

Abbreviations: NA, not applicable; NGO, nongovernmental organization.

Table 3   Studies of social groups at other layers of the hierarchy: bosses

	Study details						Implications		
Reference	Organization	Principal and time horizon	Agents	Social group	Externality	Source of variation	Findings	Social preferences	Performance	Financial incentives
Bandiera et al. 2009	Fruit farm, United Kingdom	Organization, 4 months	Seasonal fruit pickers hired for one season (approximately 100 days)	Approximately 20 other pickers on the same field and a manager, who may be a neighbor of or the same nationality as the worker	Positive: managerial effort complements worker effort	Natural variation in days with in-group managers due to rotation policy; experimental variation in manager pay between fixed wage and performance-based pay	When managers receive a fixed wage, workers are 9% more productive on days when they are connected to the manager relative to days when they are not. The effect disappears when managers are paid performance bonuses, as they reallocate effort to high-ability, unconnected workers.	Altruism toward in-group	Negative	Eliminate favoritism and increase firm performance
Hjort 2014	Flower packaging plant, Kenya	Organization, 2 years	Residential flower packagers	Three team members, who may be co-ethnic with the worker	Positive: managerial effort complements worker effort	Natural variation in ethnic composition of teams due to rotation policy; experiment also exploits exogenous political conflict due to election and change in pay scheme to piece rate for team output	Relative to homogeneous teams, output is 8% lower when the upstream worker is not co-ethnic with downstream workers; in horizontally mixed teams, the output of a non-co-ethnic worker is 18% lower, but the output of a co-ethnic worker is 7% higher. Political conflict decreases the output of mixed teams by 4%–7%. Team pay postconflict increases output in horizontally mixed teams by 4%.	Altruism toward in-group and spite toward out-group	Negative	NA

Table 3 (Continued)

Reference	Study details						Implications		
	Organization	Principal and time horizon	Agents	Social group	Externality	Source of variation	Findings	Social preferences	Performance	Financial incentives
Xu 2017	Colonial Office of the British Empire, United Kingdom	Secretary of State, approximately 110 years	Governors appointed to lead colonies	Secretary of State, who may share schooling or family with the Governor	None	Electoral turnover in the Secretary of State position creates variation in social connections of sitting governors; law passed during study period limits discretionary appointments	During the period when patronage is allowed, a well-connected Governor in a given position generates 4% less annual revenue and spends 9% less on tax collection and 11% less on public works; he receives 10% higher salary, driven by transfers to higher-paid governorships. These effects disappear after the law preventing patronage is passed.	Altruism toward in-group	Negative	NA

Abbreviation: NA, not applicable.

Table 4 Studies of social groups outside the organization

| | Study details | | | | | Implications | | | |
Reference	Organization	Principal and time horizon	Agents	Social group	Source of variation	Findings	Social preferences	Performance	Financial incentives
Giuliano et al. 2009	Retail chain, United States	Organization, 30 months	New hires for entry-level part-time sales positions	Manager, who may be the same race or ethnicity as the worker	Natural turnover in management creates within-store variation in manager race	When a Black manager is replaced by a non-Black manager, the share of Black new hires falls by 19%, and the share that is White increases by 8%. When >30% of the local population is Hispanic, replacing a Hispanic manager with a White manager decreases the share of Hispanic new hires by 17% and increases the share of White new hires by 45%.	Altruism toward in-group	NA	NA
Beaman & Magruder 2012	Lab, India	Researcher, 2 days	Individuals hired for one day of a lab task; may return with a referral who does the same task	Participant and their referral	Randomizing participants between five payment schemes for bringing a referral, varying the amount of payment and whether it is performance based or fixed	Offering high-stakes performance pay increases the likelihood of referring a coworker, rather than a relative, by 66.5% and induces higher-ability participants to recruit better-performing referrals.	Altruism toward in-group	Negative	Undo the negative effect of financial incentives
Pallais & Sands 2016	Online labor market, Philippines	Researcher, 5–6 days per experiment	Individuals hired for short-term lab tasks	Participant and their referral or partner	Inviting random sample of workers and asking them for a referral, then inviting all referred workers and a random sample of nonreferred workers for several experiments	Nonmonitored referred workers were 15.5% more accurate than nonreferrals and performed no differently from monitored referred workers. In a team task, working with one's referrer improved performance from 58% to 110%, depending on the metric, relative to working with someone else's referrer.	Altruism toward in-group	Positive	NA
Burks et al. 2015	Call centers, trucking firms, and technology firms, United States	Organization, 4–8 years, depending on the industry	Applicants and new hires for full-time jobs	None	None	While productivity is similar between nonreferred and referred workers, in trucking, referrals have a 6% lower risk of accidents each week, and in tech, referrals produce 24% more patents. Referred workers are 11–26% less likely to quit and earn higher profits.	Altruism toward in-group	Positive	NA

(Continued)

Table 4   (Continued)

| Reference | Study details | | | | | Findings | Implications | | |
	Organization	Principal and time horizon	Agents	Social group	Source of variation		Social preferences	Performance	Financial incentives
Leonard et al. 2010	Retail chain, United States	Organization, 30 months	Part-time sales employees	Customer base, which may be predominantly the same ethnicity or race as the worker	Natural variation in sales force demographics	When the Hispanic employment share increases by 1 SD (6%) above the mean, sales increase by 2% through the White customer base; when Black employment share increases by 1 SD (7%), sales decrease by 2% through the White customer base. These changes are small, as the SD of changes in sales is 18%. Raising the share of Asian American workers in areas where many people only speak Asian languages increases sales, but no similar effect is found for Hispanic workers.	NA	Positive	NA
Fisman et al. 2017	State-owned bank, India	Organization, 5 years	Borrowers	Branch officers, who may be of the same religion or caste as the borrower	Natural variation in head officers through rotation policy	When the borrowers belong to the same group as the branch officer, this increases the number of new loans by 6.5%, the number of new loan recipients by 6%, and the probability that a member of the group receives credit by 2.5%. Loans made to in-group borrowers are 7% less likely to be defaulted on.	Altruism toward in-group	Positive	NA
Bandiera et al. 2017a	BRAC, Uganda	Organization	Delivery agents of an agricultural extension program (seeds and training)	Connected farmers in the village	Random choice of delivery agent creates random connection status among farmers	The delivery agent is 8.5 percentage points more likely to target their friends; program coverage increases by 12 percentage points for each connected farmer.	Altruism toward in-group and spite towards out-group	Positive if the number of poor farmers connected to the delivery agent is large	NA

Abbreviations: NA, not applicable; SD, standard deviation.

is useful to note that financial incentives typically increase productivity by 20% in a wide range of settings (Bandiera et al. 2017b).

Second, the effect of social incentives can be positive or negative overall, but within cells of our taxonomy, one sign generally prevails. In particular, when the social group is made of peers, social incentives tend to reduce productivity through positional preferences unless individual effort generates a positive externality for other members of the group (**Table 1**). When the social group is vertical, the sign of social incentives depends on whether the externality is excludable, namely, whether the agent can target their effort to some and exclude others. When they cannot, social incentives typically increase productivity (**Table 2**). When they can, they might target their in group at the expense of a more deserving out group, thereby reducing aggregate performance (**Tables 3** and **4**). To be clear, we do not imply that there is a relationship between group composition and the sign of the social incentive effect. Rather, we see these observed patterns as interesting correlations that warrant further study.

Third, the evidence suggests that social and financial incentives interact both positively and negatively, depending on the type of social preferences underpinning social incentives. For instance, when social incentives push the agent to target their in group at the expense of a more deserving out group, financial incentives can increase the cost of doing so and therefore reduce the net benefit, as found by Bandiera et al. (2009).

To make further progress in research on social incentives, we need more evidence on the root causes of social preferences. Comparing across studies reveals some patterns, but, of course, these comparisons are mostly speculative, as contexts differ on many dimensions. With this caveat in mind, we can tentatively say that envy is more likely to be found in short-term anonymous workplaces, whereas cooperation seems to prevail in longer-term work relationships. Whether cooperation subsequently benefits the organization depends on whether individuals can cooperate with some at the expense of others.

All of the evidence is consistent with social incentives stemming from either structural parameters of the utility function or the equilibrium outcome of a repeated game among selfish agents, which results in heuristics that look like social preferences in reduced form. Both cases lead to the conclusion that social interactions in the workplace lead to social incentives that shape effort and performance. One piece of evidence that is missing is how employees respond to inequality between layers of the hierarchy. One of the most striking patterns in organizations is the exponential growth of the CEO–worker pay ratio, which, in the United States, has increased more than tenfold over the past 50 years, from approximately 20 in the 1960s to over 300 in 2015. Yet little is known about whether this has an impact on workers' productivity through social incentives.

At the core, what is missing is an understanding of how social preferences, in structural or reduced form, are formed and whether they can be shaped by policy. There is some evidence that preferences are somewhat malleable. For instance, the favoritism toward employees of the same ethnicity documented by Hjort (2014) became much stronger following contested elections and the violence that followed. Along the same lines, Bandiera et al. (2017a) find that favoritism toward socially connected agents only emerges if there are political cleavages.

Research on the determinants of social preferences is in its infancy. One possibility is that preferences are state dependent, so the same person can display different preferences in different settings. Laboratory experiments provide considerable support for this idea, starting with Andreoni's (1995) ground-breaking finding that individuals behave very differently in the same public good game depending on how the game is framed. Whether this holds outside the lab and whether organizations can design policies to leverage social preferences, as suggested by Ashraf & Bandiera (2017), are fruitful avenues for future research.

DISCLOSURE STATEMENT

ACKNOWLEDGMENTS

LITERATURE CITED

Akerlof G. 1982. Labor contracts as partial gift exchange. *Q. J. Econ.* 97:543–69

Akerlof G, Kranton R. 2010. *Identity Economics: How Our Identities Shape Our Work, Wages, and Well-Being*. Princeton, NJ: Princeton Univ. Press

Andreoni J. 1989. Giving with impure altruism: applications to charity and Ricardian equivalence. *J. Political Econ.* 97:1447–58

Andreoni J. 1995. Warm-glow versus cold-prickle: the effects of positive and negative framing on cooperation in experiments. *Q. J. Econ.* 110:1–21

Ashraf N, Bandiera O. 2017. Altruistic capital. *Am. Econ. Rev.* 107:70–75

Ashraf N, Bandiera O, Jack BK. 2014. No margin, no mission? A field experiment on incentives for public service delivery. *J. Public Econ.* 120:1–17

Ashraf N, Bandiera O, Lee S. 2016. *Losing prosociality in the quest for talent? Sorting, selection and productivity in the delivery of public services*. Work. Pap., London School Econ.

Bandiera O, Barankay I, Rasul I. 2005. Social preferences and the response to incentives: evidence from personnel data. *Q. J. Econ.* 120:917–62

Bandiera O, Barankay I, Rasul I. 2009. Social connections and incentives in the workplace: evidence from personnel data. *Econometrica* 77:1047–94

Bandiera O, Burgess R, Deserranno E, Morel R, Rasul I, Sulaiman M. 2017a. *Social connections and the delivery of development programs*. Work. Pap., London School Econ.

Bandiera O, Fischer G, Prat A, Ytsma E. 2017b. *Do women respond less to performance pay? Building evidence from multiple experiments*. Work. Pap., London School Econ.

Barnard CI. 1938. *The Functions of the Executive*. Cambridge, MA: Harvard Univ. Press

Bartel A, Freeman R, Ichniowski C, Kleiner MM. 2011. Can a workplace have an attitude problem? Workplace effects on employee attitudes and organizational performance. *Labour Econ.* 18:411–562

Bauer M, Blattman C, Chytilova J, Henrich J, Miguel E, Mitts T. 2016. Can war foster cooperation? *J. Econ. Perspect.* 30:249–74

Beaman L, Magruder J. 2012. Who gets the job referral? Evidence from a social networks experiment. *Am. Econ. Rev.* 102:3574–93

Besley T, Ghatak M. 2005. Competition and incentives with motivated agents. *Am. Econ. Rev.* 95:616–36

Blanes i Vidal J, Nossol M. 2011. Tournaments without prizes: evidence from personnel records. *Manag. Sci.* 57:1721–36

Boning B, Ichniowski C, Shaw K. 2007. Opportunity counts: teams and the effectiveness of production incentives. *J. Labor Econ.* 25:613–50

Bowles S, Polania-Reyes S. 2012. Economic incentives and social preferences: substitutes or complements? *J. Econ. Lit.* 50:368–425

Breza E, Kaur S, Shamdasani Y. 2018. The morale effects of pay inequality. *Q. J. Econ.* 133:611–63

Burks SV, Cowgill B, Hoffman M, Housman M. 2015. The value of hiring through employee referrals. *Q. J. Econ.* 130:805–39

Charness G, Rabin M. 2002. Understanding social preferences with simple tests. *Q. J. Econ.* 117:817–69

Choi JK, Bowles S. 2007. The coevolution of parochial altruism and war. *Science* 318:636–40

Cohn A, Fehr E, Herrmann B, Schneider F. 2014. Social comparison and effort provision: evidence from a field experiment. *J. Eur. Econ. Assoc.* 12:877–98

Dal Bó E, Finan F, Rossi MA. 2013. Strengthening state capabilities: the role of financial incentives in the call to public service. *Q. J. Econ.* 128:1169–218

DellaVigna S, List JA, Malmendier U, Rao G. 2016. *Estimating social preferences and gift exchange at work*. NBER Work. Pap. 22043

DellaVigna S, Pope D. 2018. What motivates effort? Evidence and expert forecasts. *Rev. Econ. Stud.* 85(2):1029–69

Deserranno E. 2017. *Financial incentives as signals: experimental evidence from the recruitment of village promoters in Uganda*. Work. Pap. 17-105, Glob. Pov. Res. Lab, Buffett Inst. Northwestern Univ., Evanston, IL

Falk A, Fehr E, Fischbacher U. 2005. Driving forces behind informal sanctions. *Econometrica* 73:2017–30

Fehr E, Gachter S. 2000. Cooperation and punishment in public goods experiments. *Am. Econ. Rev.* 90:980–94

Fehr E, Schmidt KM. 1999. A theory of fairness, competition, and cooperation. *Q. J. Econ.* 114:817–68

Fisman R, Paravisini D, Vig V. 2017. Cultural proximity and loan outcomes. *Am. Econ. Rev.* 107:457–92

Gaynor M, Rebitzer JB, Taylor LJ. 2004. Physician incentives in health maintenance organizations. *J. Political Econ.* 112:915–31

Giuliano L, Levine DI, Leonard J. 2009. Manager race and the race of new hires. *J. Labor Econ.* 27:589–631

Gneezy U, List JA. 2006. Putting behavioral economics to work: testing for gift exchange in labor markets using field experiments. *Econometrica* 74:1365–84

Griffith R, Neely A. 2009. Performance pay and managerial experience in multitask teams: evidence from within a firm. *J. Labor Econ.* 27:49–82

Hamilton BH, Nickerson JA, Owan H. 2003. Team incentives and worker heterogeneity: an empirical analysis of the impact of teams on productivity and participation. *J. Political Econ.* 111:465–97

Hansen DG. 1997. Worker performance and group incentives: a case study. *Ind. Labor Relat. Rev.* 51:37–49

Hjort J. 2014. Ethnic divisions and production in firms. *Q. J. Econ.* 129:1899–946

Jones DC, Kato T. 1995. The productivity effects of employee stock-ownership plans and bonuses: evidence from Japanese panel data. *Am. Econ. Rev.* 85:391–414

Kandel E, Lazear EP. 1992. Peer pressure and partnerships. *J. Political Econ.* 100:801–17

Knez M, Simester D. 2001. Firm-wide incentives and mutual monitoring at Continental Airlines. *J. Labor Econ.* 19:743–72

Kranton R, Pease M, Sanders S, Huettel S. 2016. *Groupy and non-groupy behavior: deconstructing bias in social preferences*. Work. Pap., Duke Univ., Durham NC

Kruse DL. 1993. *Does profit sharing affect productivity?* NBER Work. Pap. 4542

Leonard JS, Levine DI, Giuliano L. 2010. Customer discrimination. *Rev. Econ. Stat.* 92:670–78

Manski C. 1993. Identification of endogenous social effects: the reflection problem. *Rev. Econ. Stud.* 60:531–42

Mas A, Moretti E. 2009. Peers at work. *Am. Econ. Rev.* 99:112–45

Mayo E. 1933. *The Human Problems of an Industrial Civilization*. New York: Macmillan

Pallais A, Sands EG. 2016. Why the referential treatment: evidence from field experiments on referrals. *J. Political Econ.* 124:1793–828

Roethlisberger F, Dickson WJ. 1939. *Management and the Worker*. Cambridge, MA: Harvard Univ. Press

Rotemberg JJ. 1994. Human relations in the workplace. *J. Political Econ.* 102:684–717

Roy D. 1952. Quota restriction and goldbricking in a machine shop. *Am. J. Sociol.* 57:427–42

Sobel J. 2005. Interdependent preferences and reciprocity. *J. Econ. Lit.* 43:392–436

Tabellini G. 2008. The scope of cooperation: values and incentives. *Q. J. Econ.* 123:905–50

Tonin M, Vlassopoulos M. 2010. Disentangling the sources of pro-socially motivated effort: a field experiment. *J. Public Econ.* 94:1086–92

Tonin M, Vlassopoulos M. 2015. Corporate philanthropy and productivity: evidence from an online real effort experiment. *Manag. Sci.* 61:1795–811

Xu G. 2017. *The costs of patronage: evidence from the British Empire*. Work. Pap., Haas School Bus., Univ. Calif., Berkeley