

Advances in Economic Analysis & Policy

Volume 6, Issue 2

2006

Article 4

FIELD EXPERIMENTS

The Evolution of Cooperative Norms: Evidence from a Natural Field Experiment

Oriana Bandiera*

Iwan Barankay†

Imran Rasul‡

*LSE, o.bandiera@lse.ac.uk

†University of Essex, itbara@essex.ac.uk

‡University College London, i.rasul@ucl.ac.uk

The Evolution of Cooperative Norms: Evidence from a Natural Field Experiment*

Oriana Bandiera, Iwan Barankay, and Imran Rasul

Abstract

We document the establishment and evolution of a cooperative norm among workers using evidence from a natural field experiment on a leading UK farm. Workers are paid according to a relative incentive scheme under which increasing individual effort raises a worker's own pay but imposes a negative externality on the pay of all co-workers, thus creating a rationale for cooperation. As a counterfactual, we analyze worker behavior when workers are paid piece rates and thus have no incentive to cooperate.

We find that workers cooperate more as their exposure to the relative incentive scheme increases. We also find that individual and group exposure are substitutes, namely workers who work alongside colleagues with higher exposure cooperate more. Shocks to the workforce in the form of new worker arrivals disrupt cooperation in the short term but are then quickly integrated into the norm. Individual exposure, group exposure, and the arrival of new workers have no effect on productivity when workers are paid piece rates and there is no incentive to cooperate.

KEYWORDS: cooperation, group exposure, individual exposure

*We thank the editor, John List, and an anonymous referee for providing useful comments. Financial support from the ESRC under grant RES-000-22-0785 is gratefully acknowledged. We thank all those involved in providing the data. This paper has been screened to ensure no confidential information is revealed. All errors remain our own.

The ability to cooperate, namely to abstain from individually profitable actions for the sake of the common good, is a key determinant of economic performance in settings where individual and social optima do not coincide. In this paper we present evidence from a natural field experiment to document the establishment and evolution of a cooperative norm.¹

The experiment was run in collaboration with the management of a leading farm in the United Kingdom. Our subjects are farm workers, whose main task is to pick fruit. Workers are paid according to a relative incentive scheme that provides a rationale for cooperation. Under the relative scheme, each worker's pay depends on the ratio of her individual productivity to the average productivity of the group of her co-workers. Increasing individual effort therefore increases a worker's own pay but imposes a negative externality on all co-workers by raising average productivity and lowering co-workers' pay, other things equal. The welfare of the group is maximized when workers fully internalize the negative externality their effort places on others and cooperate to exert the minimum feasible level of effort.

In previous work using data from the same experiment we have shown that, on average, workers managed to cooperate to some extent under the relative incentive scheme (Bandiera et al 2005a). In other words, the productivity of the average worker lies between those predicted in two benchmark models of worker behavior – the individualistic Nash equilibrium, and the Pareto optimum among workers.

In this paper we analyze how cooperation evolves with time, namely how the behavior of a worker is affected by her exposure to the relative scheme, by the exposure of her co-workers and finally how workers react to the arrival of new individuals who are unaware of the norm. It is important to stress that cooperation can arise either because of altruism or collusion; workers might cooperate either because they truly care about colleagues' payoffs, or because they have established an implicit collusive agreement enforced by credible threats of punishment. In this paper we focus on how cooperation evolves with time, regardless of its underlying motives.²

To provide a counterfactual, workers on the same farm at the same point of the year in the following season were paid piece rates. Under this compensation

1. Following the taxonomy developed in Harrison and List (2004), the experiment falls into the "natural field" category as the subjects naturally undertake the task in the environment under study and they are not aware of being involved in an experiment.

2. The economic environment we study has a number of features that facilitate both collusion and altruism. For example, workers live and work together, interacting repeatedly both inside and outside the work environment. This makes it relatively easy for them to build social ties, and provides a variety of mechanisms to provide transfers and enforce punishment. In Bandiera et al (2005a) we show that workers cooperate only when they can monitor each other actions, which suggests that cooperation is not driven by pure altruism.

scheme each worker's pay depends only on their own productivity and hence workers have no incentives to cooperate. The analysis of how individual behavior changes with time and with the characteristics of co-workers in this counterfactual scenario allows us to separate the effects of individual and group exposure on cooperation from their effect on productivity per se.

We address three questions. First, we analyze whether workers learn to cooperate as their exposure to the scheme increases. To identify the effect of exposure on cooperation we use daily data on each worker's productivity. Identification of the parameters of interest thus arises from the comparison of a given worker to herself at different points in time, implying that time invariant sources of unobservable individual heterogeneity, such as worker's ability and motivation, are controlled for throughout. Importantly, the organization of the farm is such that workers do not pick fruit everyday, which creates a wedge between the workers' exposure to the relative incentive scheme and their experience with fruit picking. This wedge allows us to separate the effect of time on learning how to reduce productivity to "game" the incentive scheme from the effect of time on learning how to pick fruit and thus increase productivity. Finally, we exploit the fact that different workers arrive on the farm at different points of the season to establish whether individuals in early arrival cohorts learn to cooperate at the same rate as later cohorts that arrive after the norm is established.

Second, we investigate whether group exposure is a substitute for individual exposure, namely whether workers with low exposure to the scheme cooperate more when they work alongside workers who are familiar with the norm, and whether individual exposure loses relevance once other workers in the group are familiar with the norm. To this purpose we exploit the fact that the group of co-workers an individual is assigned to changes on a daily basis.

Third, we present evidence on whether shocks to the group, such as the arrival of new workers who are unaware of the cooperative norm, disrupts cooperation and, if so, whether the effect is long lasting.

Our main results are follows. First, individuals cooperate more, namely their productivity is significantly lower, as their exposure to the relative incentive scheme increases. This effect is significantly larger for the cohort of early worker arrivals, namely individuals who started working at the beginning of the peak season when the scheme was first introduced.

Second, individuals cooperate more when they work with co-workers who have been exposed to the scheme for longer and hence are more familiar with the norm. This effect is larger for the cohort of late worker arrivals, namely individuals who arrive after the cooperative norm has been established. We find that individual and group exposure are substitutes, so the marginal effect on productivity of individual exposure becomes weaker as group exposure increases and vice versa.

Third, the arrival of new workers who are unaware of the norm disrupts cooperation, namely workers who work alongside new arrivals significantly increase their productivity. The disruption is however short lived -- on average new workers learn to behave according to the cooperative norm within a week of their arrival. Moreover, the disruptive effect of new arrivals becomes weaker as the average exposure of the group they join increases.

None of these effects are present in the counterfactual strategic environment when workers are paid piece rates and hence have no reason to cooperate. In particular, neither individual nor group exposure affect productivity, and the arrival of new workers does not change the behavior of workers who arrived earlier in the season.

In summary, the findings indicate that under a relative incentive scheme, individuals quickly learn how to cooperate in the workplace, that individual and group exposure are equally important in determining levels of cooperation, and that individuals are able to quickly transmit the cooperative norm to new arrivals.

Our findings have direct relevance for the experimental literature on public goods and common pool resources games as the strategic environment individual workers face under the relative incentive scheme shares elements of both types of game. Since the seminal works of Isaac et al (1985) and Andreoni (1985), a well established finding is that cooperation in laboratory experiments decays with time (Ledyard 1995). Our results can be reconciled with the laboratory evidence by noting that our context differs from the standard laboratory setting in two important respects – individuals can communicate and punish each other throughout the duration of the experiment. Indeed, laboratory evidence suggests that cooperation increases when subjects are allowed to punish, express disapproval and communicate (Fehr and Gächter 2000, Ostrom et al 1992, Masclet et al 2003, Carpenter and Seki 2005).

More generally, our results are consistent with the finding that learning and experience have a large effect on individual behavior in both laboratory and field experiments. Slonim and Roth (1998) show that people learn how to play the ultimatum game as the pattern of offers converge to the equilibrium predictions when the same subjects play repeatedly. In a series of field experiments, List (2003, 2004) shows that market experience eliminates market anomalies, namely individuals with substantial trading experience are more likely to overcome the endowment effect.

Our paper contributes to this literature as it provides evidence from a field setting on how individuals learn to cooperate over time, both from their experience and the experience of others, and it sheds light on the robustness of cooperative norms to shocks caused by the arrival of new inexperienced individuals.

The remainder of the paper is organized as follows. Section 2 describes the context for the study, the design of the experiment, and the data. Section 3 presents the empirical results. Section 4 concludes.

1. Experiment and Data Description

1.1. The Experiment

To provide evidence on the establishment and evolution of cooperation in the workplace we exploit a natural field experiment run in collaboration with a leading producer of soft fruit in the United Kingdom. Our subjects are university students from Eastern Europe, hired as seasonal farm workers for three to six months.³

The workers' main task is to pick fruit on a number of fields each day. Workers work alongside each other but they are assigned their own row of fruit and pick independently from others; namely, each worker's productivity depends solely on her own effort and field conditions.

Workers live on the farm, nearby the fields, and interact repeatedly both inside and outside the work environment. The organization of the workplace thus provides opportunities for workers to build social ties with others and provides workers a variety of mechanisms by which to make transfers or hand out punishments to enforce cooperative norms. The magnitude of such transfers across workers, as well as the individual incentive to deviate from a cooperative norm, are expected to be large because the real value of earnings for workers on the farm is high. Gross monthly earnings at the UK minimum wage (€ 1105) are 5 times as high as at the minimum wage in Poland (€ 201) and almost 20 times higher than in Bulgaria (€ 56).⁴

The rationale for cooperation derives from the fact that workers are paid according to a relative incentive scheme. For two consecutive months during the first half of the 2002 peak picking season, workers face a compensation schedule of the form;

$$(1) \quad \textit{compensation} = \beta K_i$$

3. There are ten nationalities represented in the data, both genders are equally represented, and individuals are aged 20 to 25 years. In order to be recruited, individuals must be full-time university students, have at least one year before graduation, and must return to the same university in the Fall. Only a handful of workers are hired for two consecutive seasons.

4. This monthly minimum wage data is from Eurostat, January 2003. These differences remain even if PPP adjustments are made. Not surprisingly, three quarters of workers in our sample report coming to the farm for financial reasons.

where β is the unit wage and K_i is the total kilograms of fruit picked by worker i on the field-day. The productivity of worker i , y_i , is defined as the amount of kilograms of fruit she picks per hour. The defining feature of the incentive scheme is that the unit wage β is endogenously determined by the average productivity of all workers on the field-day. Hence β is set according to;

$$(2) \quad \beta = w/\mu_y,$$

where w takes the same value throughout the season, and μ_y is the average hourly productivity of all workers in the same field on the same day. There are, on average, forty workers on each field-day. At the start of each field-day the farm manager announces an ex ante picking rate based on her expectations of worker productivity. This picking rate is revised at the end of each field-day to ensure a worker with productivity μ_y earns the pre-established hourly wage w .

Under relative incentives an increase in worker i 's effort increases her own pay, but also increases the average productivity on the field-day and thus imposes a negative externality on her co-workers by reducing the unit wage β in (2). The relative incentive scheme creates a wedge between individual and group optima, thus providing a rationale for cooperation. To be clear, in this setting higher cooperation corresponds to lower productivity. We then employ daily data on individual productivity to analyze how cooperation evolves with time, as a function of individual and group exposure to the relative scheme.

To disentangle the effects of cooperation on productivity from other factors that can lower productivity through time, we use daily data on individual productivity during the first half of the 2003 peak picking season, when workers are paid piece rates and thus have no incentive to cooperate. The piece rate compensation schedule is the same as in (1) but the unit wage β is set ex ante and is not revised according to average productivity. An increase in worker i 's effort then does not affect the unit wage received by her co-workers, namely there is no externality and the individual optimum coincides with the group optimum.⁵

5. In a dynamic framework workers may under perform if they believe that working hard will result in management setting lower piece rates in the future. There are two reasons why in this setting, there are unlikely to be such large ratchet effects. First, given the stochastic nature of agricultural production, it is difficult for workers to disentangle changes in the piece rate due to changing conditions and those due to management learning about workers' true ability. Such ratchet concerns have been documented in firms where productivity shocks are less common such as shoe making (Freeman and Kleiner 2005) and bricklaying (Roy 1952). Second, the effect of a worker's current performance on the unit wage she faces in the future is weak as the unit wage is field-day specific and workers are reallocated to different fields in different days. In particular, workers face uncertainty over which fields they will be assigned to in the future and about the identity of their future co-workers.

The 2003 season provides an appropriate counterfactual as the farm uses the same technology on the same set of fields and employs workers from the same pool of individuals in both years. Workers originate from the same set of countries, attend the same set of universities, and so the age, gender and nationality distributions are very similar across the two years.

Our previous work compares the average behavior of workers under relative incentives to the average behavior of the same workers later in the 2002 season when they were paid piece rates and thus had no incentive to cooperate (Bandiera *et al* 2005a). We show that, notwithstanding the large group size, the productivity of the average worker was at least 50 percent higher under piece rates than under relative incentives and that this was due to workers partially internalizing the negative externality. In this study we build on these previous results to understand the establishment of the cooperative norm and the evolution of cooperation over time.

For this purpose, the first half of the 2003 season provides a better counterfactual than the second half of 2002. This is so because in the second half of the same season workers have, for example, higher levels of picking experience, more established social ties with co-workers, and greater familiarity with the organization of the farm. These factors cloud identification of the parameters of interest in this study. In contrast, in the first half of 2003 workers are more comparable to the workers in the first half of 2002 along each of these dimensions. Moreover, characteristics at the field-day level such as field conditions and group exposure are also more comparable across treatment and control groups.

1.2. The Data

We retrieve workers' productivity on each field and each day on which they pick fruit from personnel records. Productivity is defined as kilograms picked per hour, and is measured electronically by assigning a unique bar code to each worker. The personnel records also contain information on the identity of all co-workers on each field-day, and on the dates of arrival and departure for each worker on the farm.

The relative incentive scheme was in place for the first half of the peak season in 2002, from mid-May until the first week of July. This represents our treatment season. The control season is the first half of the peak season in 2003, from early May until the end of June.⁶

6. Thus the peak season started 10 days earlier in 2003. Since overall the peak season has the same length in the two years, the comparable sample for 2003 ends at the end of June, that is 10 days earlier than in 2002.

The fruit type that is picked is the same in both seasons. The 2002 sample covers 289 workers, 15 fields, 150 field-days, and provides a total of 6177 worker-field-day level observations. The 2003 sample covers 349 workers, 14 fields, 152 field-days and provides 9858 worker-field-day level observations.⁷

Table 1 presents descriptive statistics for the variables used in the analysis. The table highlights that on average, worker productivity is about 60% higher in 2003 under piece rates, compared to 2002 under the relative incentive scheme.⁸

To analyze the evolution of cooperation with time we assess whether individual productivity depends on the individual's exposure to the relative scheme. If workers learn how to cooperate, they should work less hard as their exposure, defined as the cumulative number of days the worker has been present on the farm, increases. Importantly, while workers are present on the farm for every day of their stay in the UK, they do not pick fruit every day.⁹

This is so for two reasons. First, as workers are hired on a casual basis the employment contract provides no guarantee of being employed each day. Hence on days in which there is less fruit to be picked some workers may be left unemployed. Second, on some days workers will be engaged in non-picking tasks only, such as planting or weeding. The allocation of workers between picking, non-picking tasks, and unemployment is made by management on the basis of the demand for labor on each task.

These factors create a wedge between the picking experience of each worker, that is the cumulative number of days the worker has been picking fruit, and the workers' exposure to the scheme. The fact that workers accumulate picking experience at an exogenously different rate to their exposure to the scheme allows us to separately identify the effects of experience and exposure. This is important because as time passes workers might learn how to cooperate with others and how to "game" the relative incentive scheme by exerting less effort. This is so if the cooperative norm can be communicated to all workers, even if those workers themselves are not picking fruit and hence not being paid according to the relative incentive scheme. At the same time, however, workers naturally become more productive as they accumulate picking experience.

Table 1 shows that there is no difference in the amount of picking experience among workers across the two seasons; workers are exposed to the

7. While the farm is physically the same in the two years, yields were higher in 2003 due to a combination of weather conditions and the maturity of plants. The number of hired workers was accordingly higher in 2003 compared to 2002.

8. It is also worth noting that productivity in the first half of the 2003 season is not significantly different from productivity in the second half of the 2002 season (not reported) when workers' were paid piece rates. This is consistent with workers cooperating to lower productivity under the relative incentive scheme in 2002.

9. In our sample, workers pick fruit on average every other day. The ratio of picking days to total days on the farm is 48% for 2002 and 51% for 2003.

farm a week more in 2002 than in 2003 but this difference is not statistically significant.

An second important source of quasi-random variation that we exploit in the analysis is in the time of arrival of workers. Workers arrive on the farm at different points of the season partly because of variations in term dates across universities in their home countries. In addition, farm management aims to keep constant the supply of labor over the peak picking season.¹⁰

As each individual's work permit places an upper bound on how long they can work for, a stable supply of labor is ensured by staggering the issue of work permits. As shown in Figure 1 and in the first panel of Table 1, some workers arrive in April before the start of the peak picking season, while others arrive during the peak season. Two points are of note. First, the pattern of arrivals is very similar in the treatment and control years. Second, workers arrive almost every week during the season and this creates variation in the average exposure of the group of workers at any given point in time. Most workers stay throughout the peak season and, on average, they depart some nine more weeks after the end of the peak picking period in either year. The personnel records also indicate that no worker was fired in either year.¹¹

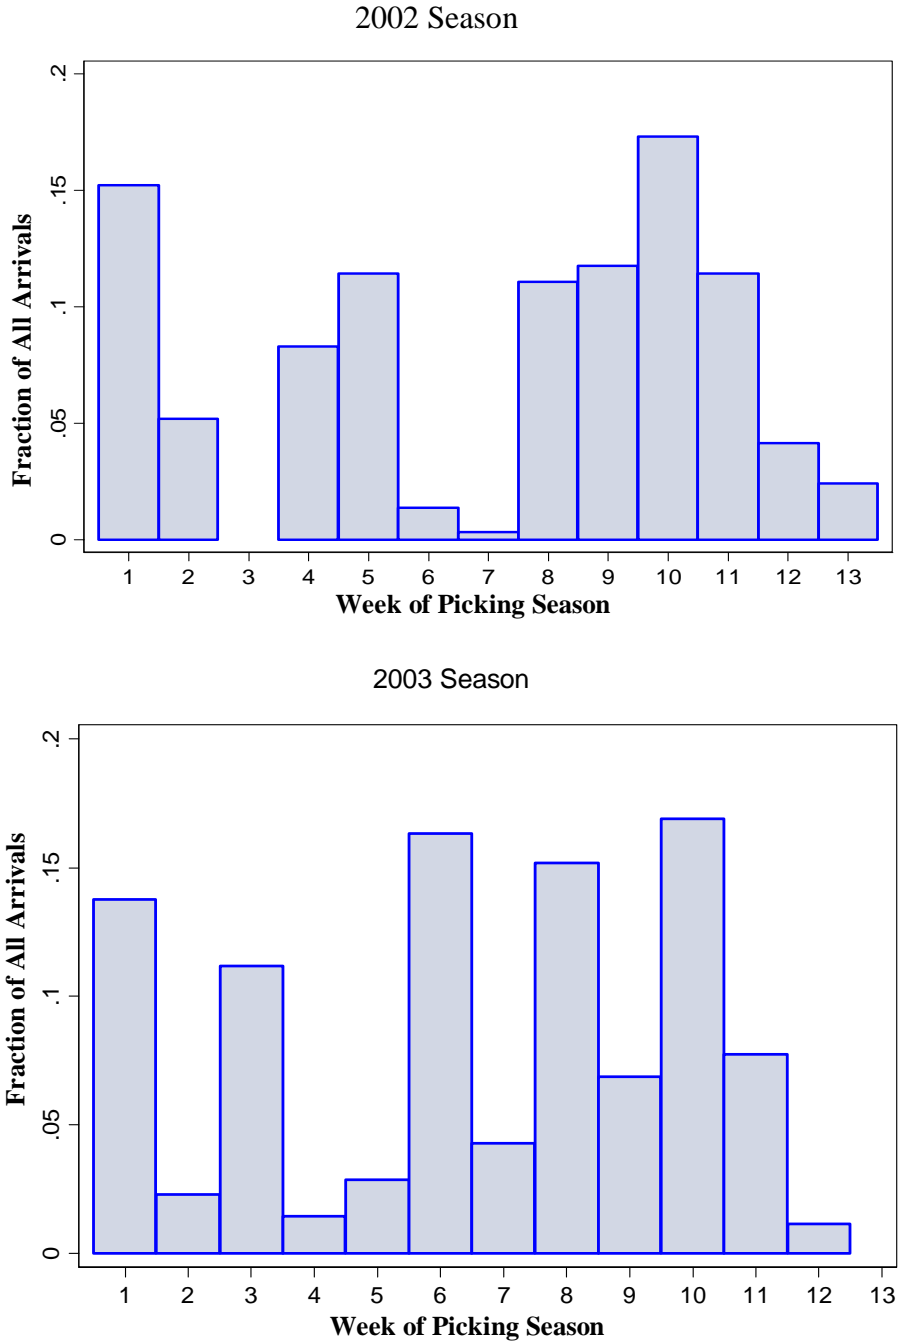
Table 1 reports two characteristics of the group of workers on the field-day that are used later in the analysis. The first is average group exposure, defined as the mean of exposure of each worker in a given field day. The second is the share of new arrivals, defined as the number of workers on the field-day who have been on the farm for less than one week divided by the total number of workers on the field-day. In line with individual exposure being higher in 2002, we also find that group exposure is higher in 2002. The share of new arrivals, defined as the number of workers with less than one week exposure divided by the total number of workers on the field-day, is also slightly higher in the 2003 season as more workers arrive later.

Finally, Table 1 reports data on the average field life cycle as a measure of field conditions. The life cycle is defined as the number of days the field has been picked until day t divided by the total number of days that the field is picked over the season. This captures a key feature of the fruit growing technology, that is, the quantity of fruit available in a field depletes over time. Average field conditions are not significantly different across the two samples.

10. Fruit is planted some years in advance to ensure a near constant supply of fruit over the peak picking season. Hence it is optimal to have a near constant supply of labor over the peak season.

11. A few workers, accounting for 3% of total observations, leave before the end of our sample in either year. The results are not affected if we drop these "early leavers" from the sample.

Figure 1: Distribution of Worker Arrivals by Season



Notes: The first week of the picking season is defined to be the week beginning from the first Monday of April in each season.

Table 1: Descriptive Statistics

	2002 Season	2003 Season
Worker Characteristics		
April arrival cohort	83	100
May arrival cohort	70	135
June arrival cohort	136	114
Total	289	349
Worker Characteristics (means, standard deviations in parentheses)		
Worker productivity (kg/hr)	4.92 (3.13)	8.06 (4.10)
Picking experience (days)	11.3 (8.27)	11.3 (8.32)
Individual exposure (days)	27.1 (21.0)	20.9 (16.0)
Field-Day Characteristics (means, standard deviations in parentheses)		
Field life cycle (0-1)	.504 (.108)	.520 (.105)
Group exposure (days)	29.6 (11.7)	20.0 (7.45)
Share of new arrivals on field-day (0-1)	.164 (.162)	.239 (.201)

Notes: Productivity is defined as kilograms of fruit picked per hour. Individual exposure equals the cumulative number of days the worker has spent on the farm since fruit picking started. Picking experience is the cumulative number of days the worker has been picking fruit. The field life cycle is defined as the cumulative number of days a field has been operated divided by the total number of days the field is operated during the season. Group exposure is the mean exposure of workers on the same field on the same day. The share of new arrivals is the number of workers with less than one week exposure divided by the total number of workers on the field-day. The sample from the 2002 season covers 289 workers, 15 fields, 150 field-days, and provides a total of 6177 worker-field-day level observations. The sample from the 2003 season covers 349 workers, 14 fields, 152 field-days and provides 9858 worker-field-day level observations.

2. Empirical Analysis

The empirical analysis proceeds in three stages. First, we analyze how the behavior of a given worker changes as a function of her exposure to the relative incentive scheme. That is, do workers become more or less cooperative over time? Second, we analyze how the behavior of a given worker changes as a function of the average exposure of her co-workers. That is, do workers cooperate more when they work in groups that are more familiar with the norm? Third, we ask whether cooperation is disrupted by the arrival of new workers that are unaware of the norm, and whether workers learn to deal with new arrivals over time.

Throughout we use 2003 as a control group to separately identify the effect of the variables of interest on cooperation from their effect on individual productivity per se. Indeed, while the allocation of workers to a given field-day might depend on information that is available to the farm management and not to the econometrician, we are able to isolate the effect of the variables of interest on cooperation to the extent that the omitted variables are orthogonal to the incentive scheme in place. Observing workers under relative incentives and piece rates allows us to separate the effect of, say, group exposure on productivity through cooperation, which is only relevant under the relative incentive scheme, from the spurious effect of unobservables that determine group composition and productivity under both schemes.

2.1. Individual Exposure

We estimate the following panel data regression on the stacked 2002-2003 data, where all continuous variables are in logarithms;

$$(3) \quad y_{ift} = \gamma^{02} d^{02} X_{it} + \gamma^{03} d^{03} X_{it} + \delta T_{it} + \eta Z_{ft} + \alpha_i + \varphi_f + u_{ift}$$

where y_{ift} denotes the productivity of worker i on field f on day t . X_{it} measures the worker's exposure, namely the number of days she has spent on the farm, independent of whether she has been picking or not. The dummy variables d^{02} and d^{03} take a value of one in 2002 and 2003 respectively, and zero otherwise. The coefficients of interest are γ^{02} and γ^{03} which capture the effect of individual exposure on productivity in 2002 and 2003 respectively.¹²

The null hypothesis is that individual exposure has no effect on cooperation and hence productivity, that is $\gamma^{02} = \gamma^{03} = 0$. If individual exposure

12. We checked for robustness to functional form here and in the two following models using a linear, log-linear and quadratic specification for the RHS variables and a linear specification for the LHS. Findings are robust to these alternative specifications.

affects productivity only through its effects on cooperation we expect $\gamma^{02} \neq 0$, $\gamma^{03} = 0$, that is the effect of exposure is zero only when there is no rationale for cooperation. If workers learn to cooperate as their exposure to the scheme increases we expect $\gamma^{02} < 0$, while if repetition destroys cooperation the opposite occurs so that $\gamma^{02} > 0$.

T_{it} measures the worker's picking experience, that is the cumulative number of days the worker has picked fruit until day t . This captures the fact that workers become more productive as they accumulate experience and allows us to separate the effect of exposure on cooperation from the natural increase in productivity due to the positive returns from picking experience. Z_{ft} measures the life cycle of field f on day t . This is defined as the number of days the field has been picked until day t divided by the total number of days that the field is picked over the season. As the technology and characteristics of hired workers are the same in the two years, we assume the effects on productivity of picking experience and field life cycle are the same across years.¹³

We include workers' fixed effects, α_i , to capture time invariant worker level determinants of productivity such as the value of their outside option, innate ability, and intrinsic motivation. The effect of individual exposure and experience is thus identified by comparing a worker to herself at different points in time. This ensures that individual exposure and experience do not proxy for unobservable time invariant worker characteristics that drive productivity. Therefore the parameters of interest are consistently estimated even if, for example, more motivated workers arrive earlier and thus have higher exposure, or if more able workers are selected to pick more frequently and thus have more experience.

A further concern arises if ability or other unobservables affect the probability of being selected differently at different points in time. For instance this could happen if it takes time for managers to learn how to identify good pickers so that unobserved ability has a stronger effect on the probability of being selected later in the season. Our identification relies on the assumption that any spurious time varying effects are the same in the two seasons. Namely, since we identify the effect of exposure on cooperation by comparing the effect of exposure on productivity in the 2002 and 2003 seasons, inferences about the difference in workers' behavior between the two years can still be made as long the bias is the same in the two seasons. In the example above this entails assuming that managers learn at the same rate in the two seasons, which is sensible given that the identity of managers and workers differ each season, managers and workers are chosen from a similar pool of applicants each year, and all other aspects of farm operations are unchanged.

13. A test of the hypotheses that the coefficients of picking experience and of the field life cycle are the same in 2002 and 2003 fails to reject the null once individual exposure is controlled for.

Finally, we include field fixed effects, ϕ_f , to capture time invariant field level determinants of productivity such as soil quality or plant spacing. Note that while some of the fields are the same in the two years we treat them as different in (3). The reason for this is that the fruit type grows on a three year cycle, implying that the crop on a given field in 2003 is one year older than the crop on the same field in 2002. The field fixed effects thus also absorb any level differences between the two years.

The disturbance term, u_{ijt} , captures unobserved determinants of productivity at the worker-field-day level. Observations within the same field-day are unlikely to be independent since workers face similar field conditions. This is accounted for by clustering standard errors at the field-day level in all the regressions.

Table 2 presents estimates of (3). Column 1 shows the baseline estimate of (3) without controlling for individual exposure. As is intuitive, the results show that productivity increases as workers accumulate picking experience, and decreases as fields get later into their life cycle.¹⁴

We introduce our variable of interest, individual exposure, in Column 2. We find that conditional on the worker's picking experience, individual exposure has a significantly negative effect on productivity in 2002 ($\gamma^{02} < 0$). In comparison to the specification in Column 1, the coefficient on picking experience rises, which given that exposure has a negative effect on productivity, is as expected because experience is positively correlated with exposure.¹⁵

Quantitatively, the coefficients imply that a one standard deviation increase in worker's picking experience, evaluated from the mean, increases productivity by 7.8%. A one standard deviation increase in individual exposure decreases productivity by 17.3%. Finally, a one standard deviation increase in the field life cycle decreases productivity by 14.3%.

Individual exposure has no effect on productivity in 2003 ($\gamma^{03} = 0$). Importantly, this is not due to the coefficient being imprecisely estimated. Since the key difference between the two years is that in 2002 workers have incentives to cooperate while in 2003 they do not, we interpret the result as indicating that the longer a worker is exposed to the relative incentive scheme, the more she cooperates by reducing her productivity when she actually picks.

In Column 3 we test the hypothesis that the effect of exposure depends on the worker's month of arrival. We divide workers into three month of arrival

14. The worker and field fixed effects are also jointly significant. By themselves, worker fixed effects explain 36% of the variation in productivity, highlighting the importance of controlling for unobserved worker heterogeneity throughout.

15. In contrast, the coefficient on the field life cycle is unchanged over the two columns suggesting that the experience and exposure of workers is uncorrelated to the stage of the life cycle on fields to which they are allocated.

Table 2: Individual Exposure

Dependent variable = Log worker productivity (kg/hr)

Robust standard errors in parentheses, clustered by field-day

	(1) Baseline	(2) Exposure	(3) Cohort
Picking experience	.047* (.027)	.083** (.034)	.115*** (.034)
Field life cycle	-.822*** (.187)	-.805*** (.185)	-.584*** (.215)
Individual exposure 2002		-.193*** (.053)	
Individual exposure 2003		.022 (.036)	
April arrival cohort x individual exposure 2002			-.847*** (.184)
April arrival cohort x individual exposure 2003			-.098 (.119)
May arrival cohort x individual exposure 2002			-.366*** (.081)
May arrival cohort x individual exposure 2003			-.055 (.045)
June arrival cohort x individual exposure 2002			-.168** (.066)
June arrival cohort x individual exposure 2003			.025 (.047)
Worker fixed effects	Yes	Yes	Yes
Field fixed effects	Yes	Yes	Yes
Number of observations (worker-field-day)	16035	16035	16035
Adjusted R-squared	.4284	.4375	.4438

Notes: All continuous variables are in logarithms. Individual exposure equals the cumulative number of days the worker has spent on the farm since fruit picking started. Picking experience is the cumulative number of days the worker has been picking fruit. The field life cycle is defined as the cumulative number of days a field has been operated divided by the total number of days the field is operated during the season.

cohorts – April, May and June – and allow the coefficients of interest, γ^{02} and γ^{03} , to differ across cohorts. Intuitively, early arrivals have to establish the cooperative norm while late arrivals find the norm already in place. This implies the effect of exposure should be stronger for early arrivals than for later arrivals who might learn how to cooperate from the existing workers.

The pattern of coefficients are largely in line with this reasoning. For the cohort of workers who arrive in April, evaluated from the mean, a fifteen day increase in individual exposure reduces productivity by 22.2%, by 18.7% for the May cohort, and by 13.8% for the June cohort.

In contrast, in the 2003 season, month of arrival has no effect on worker productivity. The interactions of exposure with month of arrival are actually more precisely estimated for 2003 than 2002, the point estimates are orders of magnitude smaller than for 2002, and close to zero.

These results can also be represented graphically. To do this we estimate (3) without controlling for individual exposure. In Figure 2 we then plot the residuals from this regression against individual exposure, for each month of arrival cohort. This shows how the unexplained component of worker productivity – after worker fixed effects, field fixed effects, picking experience and field life cycle are controlled for – relates to individual exposure, and how this varies by month of arrival. The Figure shows that earlier cohorts cooperate more as their exposure to the relative incentive scheme increases, whereas workers in the last cohort cooperate almost immediately from when they arrive on the farm. Both Column 3 of Table 3 and the lower panel of Figure 2 show that exposure does not matter for any cohort in 2003.

A natural explanation for the differential effect of exposure by month of arrival is that workers who arrive in June find that their co-workers have already established a cooperative norm under relative incentives, and that this norm can be imposed onto later arrivals. The next two subsections explore this hypothesis in more detail.


2.2. Group Exposure

We estimate the following panel data regression;

$$(4) \quad y_{ift} = \psi^{02} d^{02} G_{ift} + \psi^{03} d^{03} G_{ift} + \delta T_{it} + \eta Z_{ft} + \alpha_i + \phi_f + u_{ift}$$

where G_{ift} is the mean exposure of worker i 's co-workers in field f on day t . If worker i learns from her co-workers, she should cooperate more, and hence have lower productivity, when she works alongside others who are more familiar with the norm. Given that the identity of co-workers changes each field-day, we

Figure 2: Time Series of Productivity Residual, by Month of Arrival



Notes: The productivity residual is the residual from the productivity regression in specification (3) in the main text, where individual exposure, X_{it} , is not controlled for. Each cohort comprises workers that arrived during that month. Individual exposure equals the cumulative number of days the worker has spent on the farm since fruit picking started.

identify the effect of group exposure on individual productivity by comparing the worker to herself as she works alongside different co-workers.

A necessary condition for the parameters of interest, ψ^{02} and ψ^{03} , to be consistently estimated in (4) is that there is no systematic year-specific correlation between G_{ift} and u_{ift} . Namely, the identifying assumption is that any unobservable that creates a spurious correlation between G_{ift} and y_{ift} does so regardless of the incentive scheme in place. For instance, if workers who have been on the farm for longer are more likely to be assigned to a low productivity field (so that G_{ift} and y_{ift} are negatively correlated), identification requires this to hold true in both seasons. In other words, we assume that the rule according to which managers allocate workers to fields is orthogonal to the incentive scheme in place.

While managers might want to alter group composition to reduce cooperation under relative incentives, data on group composition suggests otherwise. To keep the comparison as clean as possible we look at group composition in the first and second half of 2002, when the same managers choose from the same pool of workers and the only difference is the incentive scheme in place. We find that the probability of working in the same group as one's self-reported friends or in the same group as people who live in the same caravan is the same under both schemes (55% and 60% respectively). The mean share of workers on the field who are friends of a given worker is .043 under relative incentives and .037 under piece rates, the difference not being statistically significant. Finally, the mean share of workers on the field who are of the same nationality of a given worker is .175 under relative incentives and .156 under piece rates, the difference again is not significant. We also note that if managers were to allocate workers to groups to discourage cooperation under relative incentives, we should observe less variation in group composition across field-days under relative incentives, leading to less precise estimates.¹⁶

The result in Column 1 of Table 3 shows that under relative incentives, a given worker cooperates more when she works in a group that has higher mean exposure. Evaluated from the mean, a one standard deviation increase in mean exposure of the group of co-workers decreases productivity by 12.4% under relative incentives.

16. One possible explanation for managers' behavior is that they do not internalise the effect of group composition on the farm's profits because, being paid a fixed wage, they have no stake in it. Data from a different experiment (Bandiera et al 2005b) indeed shows that when managers' pay is conditional on workers' performance, managers significantly change how they allocate workers to fields.

Table 3: Group Exposure

Dependent variable = Log worker productivity (kg/hr)

Robust standard errors in parentheses, clustered by field-day

	(1) Exposure	(2) Cohort	(3) Interactions
Group exposure 2002	-.422** (.218)		-.904** (.383)
Group exposure 2003	-.048 (.080)		-.068 (.123)
April arrival cohort x group exposure 2002		-.206 (.205)	
April arrival cohort x group exposure 2003		-.085 (.083)	
May arrival cohort x group exposure 2002		-.558*** (.206)	
May arrival cohort x group exposure 2003		-.058 (.088)	
June arrival cohort x group exposure 2002		-.455 (.322)	
June arrival cohort x group exposure 2003		.024 (.107)	
Individual exposure 2002			-.743*** (.289)
Individual exposure 2003			.158 (.129)
Group exposure x individual exposure 2002			.186** (.094)
Group exposure x individual exposure 2003			-.039 (.048)
Worker fixed effects	Yes	Yes	Yes
Field fixed effects	Yes	Yes	Yes
Other controls	Yes	Yes	Yes
Number of observations (worker-field-day)	16035	16035	16035
Adjusted R-squared	.4347	.4358	.4445

Notes: All continuous variables are in logarithms. Individual exposure equals the cumulative number of days the worker has spent on the farm since fruit picking started. Group exposure is the mean exposure of workers on the same field on the same day. Other controls include worker's picking experience and the field life cycle.

In comparison, a one standard deviation increase in picking experience increases productivity by 4.8%.¹⁷

Column 2 repeats the analysis allowing the coefficients of interest to vary by month of arrival cohort. Although the results by cohort are not as pronounced as for individual exposure, we still find that group mean exposure significantly reduces the productivity of May and June arrivals relative to those workers that arrive in April. This is consistent with the idea that workers who arrive when the norm is already in place are more affected by the exposure of their colleagues. In other words, early arrivals learn how to cooperate, establish the norm and transmit their knowledge to new arrivals.


The final specification sheds light on whether individual and group exposure are substitutes – namely whether individuals learn to cooperate both from by being individually exposed, and from working alongside co-workers who have been exposed. To do this we re-estimate (4) and allow for an interaction between group and individual exposure. The result in Column 3 shows that both individual and group exposure favor cooperation under the relative incentive scheme. Their interaction is positive, suggesting that the marginal effect of individual exposure is smaller in absolute value on field-days when an individual works alongside high-exposure co-workers. Similarly, the exposure of co-workers has a smaller effect on workers that have been individually more exposed to the relative incentive scheme.

Figure 3 illustrates this finding. We plot the marginal effect of individual exposure as a function of group exposure and vice versa. The marginal effect of individual exposure varies from -.2 log points in a group that has 20 days of exposure on average, to zero when the average group exposure rises to 50 days. Similarly, the marginal effect of group exposure is eight times larger for a worker with five days individual exposure compared to a worker who has been exposed to the scheme for 50 days. To provide an indication of the relative importance of individual versus group exposure we evaluate the marginal effect of each at the mean level of the other. We then find that the marginal effect of group exposure is to reduce productivity by .33 log points, and the marginal effect of individual exposure is to reduce productivity by .23 log points.

The fact that, conditional on individual exposure and experience, group exposure still has a significant effect on workers' behavior under relative incentives, implies that individuals learn the cooperative norm from others. In our setting this is not surprising as individuals can communicate to each other,

17. One concern may be that an individual that intends to break the cooperative norm prefers to work alongside co-workers with higher levels of group exposure. This is because such a group has lower productivity, hence a higher unit wage β , and therefore the returns to breaking the cooperative norm are higher. However, such an endogenous allocation of workers to field by group exposure would bias ψ^{02} downwards in absolute value and so provide a lower bound on the true effect.

Figure 3: Marginal Effects of Individual and Group Exposure



Notes: All Figures are for the 2002 season. Individual exposure is equal to the cumulative number of days the worker has spent on the farm since fruit picking started. Group exposure is the mean exposure of workers on the same field on the same day. The share of new arrivals is the number of workers with less than one week exposure divided by the total number of workers on the field-day.

deviators from the cooperative norm can be costlessly identified, and there are a variety of mechanisms through which they can be punished.¹⁸

Finally, the last specification in Table 3 reaffirms that the effects of exposure on productivity go through cooperation as we find no evidence that group or individual exposure affect productivity in 2003. The 2003 coefficients are estimated precisely throughout and are close to zero.

2.3. Shocks: New Arrivals

In this subsection we analyze the impact of shocks, namely what happens to cooperation when new workers who are unaware of the norm join workers on the field-day. We estimate the following specification;

$$(5) \quad y_{ift} = \lambda^{02} d^{02} N_{ft} + \lambda^{03} d^{03} N_{ft} + \delta T_{it} + \eta Z_{ft} + \alpha_i + \varphi_f + u_{ift}$$

where N_{ft} is the share of workers on field-day ft who have less than one weeks exposure, and we restrict the sample to workers who have more than one week exposure. The coefficients of interest, λ^{02} and λ^{03} , measure the effect of new arrivals on the productivity of existing workers. If new arrivals disrupt cooperation, that is they cause productivity to rise, then $\lambda^{02} > 0$ and $\lambda^{03} = 0$. The second condition is necessary to ensure that new arrivals affect productivity solely through their effect on cooperation. If the share of new arrivals were correlated with the productivity of every worker on the field-day for spurious reasons, for example if new arrivals were allocated to fields that are easier to pick, we would find $\lambda^{03} \neq 0$. As with previous estimates, the comparison between λ^{02} and λ^{03} allows us to identify the presence of any spurious correlation that might arise for similar reasons over the two years.

Estimates of (5) are presented in Table 4. The result in Column 1 shows that new arrivals disrupt the cooperation norm under the relative incentive scheme. The presence of workers with less than one weeks exposure to the scheme causes the productivity of other workers on the field-day to significantly increase. Quantitatively, an increase in the share of new arrivals by one standard deviation from its mean, increases the productivity of other workers by 11.6% in 2002, while new arrivals have no significant effect in 2003 when workers are paid according to piece rates.

18. Huck et al (2004) show the conditions under which, in the context of a game of Cournot competition, trial and error learning leads players to converge to the joint profit maximizing levels of output. While such a process of learning does not fit our data, we cannot rule out workers imitating the behavior of others. Learning processes with some element of imitation have, in the Cournot setting, been shown to converge to the collusive levels of output (Apesteguia et al 2002).

Table 4: New Arrivals

Dependent variable = Log worker productivity (kg/hr)
 Robust standard errors in parentheses, clustered by field-day

	(1) Arrivals	(2) Arrivals	(3) Exposure
Share of new arrivals 2002	1.20*** (.414)		5.74 (3.64)
Share of new arrivals 2003	.075 (.175)		-.286 (1.52)
Share of new arrivals (1 to 4 days) 2002		1.14** (.486)	
Share of new arrivals (5 to 7 days) 2002		1.36** (.638)	
Share of new arrivals (8 to 11 days) 2002		-.640 (.466)	
Share of new arrivals (12 to 14 days) 2002		-.087 (.429)	
Share of new arrivals x individual exposure 2002			.243 (.246)
Share of new arrivals x individual exposure 2003			-.050 (.229)
Share of new arrivals x group exposure 2002			-1.77* (1.06)
Share of new arrivals x group exposure 2003			.101 (.483)
Individual exposure 2002			-.479*** (.133)
Individual exposure 2003			-.079 (.103)
Group exposure 2002			.283 (.274)
Group exposure 2003			-.248
experience	.040 (-.038)	-.077 (.078)	.035 (.038)
field life cycle	-.326 (.228)	-.285 (.561)	-.245 (.248) (.168)
Worker fixed effects	Yes	Yes	Yes
Field fixed effects	Yes	Yes	Yes
Other controls	Yes	Yes	Yes
Test 1: 1 to 4 days = 5 to 7 days (p-value)		.7546	
Test 2: 8 to 11 days = 12 to 14 days (p-value)		.3692	
Number of observations (worker-field-day)	12858	4089	12858
Adjusted R-squared	.4441	.3299	.4529

Notes: All variables are in logarithms. Individual exposure is the cumulative number of days a worker has spent on the farm since fruit picking started. Group exposure is the mean exposure of workers on the same field on the same day. Share of new arrivals is the number of workers with less than one week exposure divided by the total number of workers on the field-day. The samples in Columns 1 and 3 exclude workers with less than seven days of exposure. The sample in Column 2 excludes all workers with less than fourteen days of exposure. Other controls include worker's picking experience and the field life cycle.

Column 2 investigates further precisely how long recently new arrivals have to have been exposed to the relative incentive scheme to cause this effect on the pre-existing cooperative norm. To ease exposition we focus on the 2002 sample, and on workers that have at least two weeks of exposure. We then separate new arrivals into four groups according to whether they have between 1 and 4 days exposure, between 5 and 7, between 8 and 11 and between 12 and 14. We find that the effect of the first two groups is the same, namely workers with any level of exposure between 1 and 7 days are equally disruptive. Furthermore, seven days is the threshold level of individual exposure beyond which new arrivals do not disrupt cooperation.

One interpretation is that new workers learn the norm when they socialize with others on their first weekend. The data does not however allow us to explore this hypothesis any further since most workers begin employment on Mondays, so that their first weekend occurs after seven days.

Finally, Column 3 investigates whether existing workers learn to cope with new arrivals over time, namely whether the marginal impact of new arrivals decreases with individual and group exposure. For this purpose, we control for individual and group exposure in (5), and for their interactions with the share of new arrivals.

We find that new arrivals are less disruptive when they first work in a group comprising of workers who are more familiar with the cooperative norm -- the interaction between group exposure and the share of new arrivals is negative and significant at the 10% level in 2002. In contrast, there is no interaction effect between individual exposure and the share of new arrivals. This suggests, as is intuitive, that it is the group of co-workers as a whole, and not specific individual workers, that learn to cope with the presence of new arrivals.

In the control group in 2003, new arrivals on the field-day have no direct effect on the productivity of their co-workers, nor do they have differential effects as individual or group exposure vary. This lends support to the hypothesis that the presence of new arrivals has an effect through the ability of the group to sustain a cooperative norm, rather than some other mechanism that affects productivity.

3. Conclusions

In this paper we exploit a natural field experiment to document the establishment and evolution of a cooperative norm among workers. The rationale for cooperation derives from the fact that workers are paid according to a relative incentive scheme. Under this scheme, increasing individual effort increases a worker's own pay but imposes a negative externality on all co-workers by raising average productivity and lowering co-workers' pay. The welfare of the group is

maximized when workers fully internalize the negative externality their effort places on others and cooperate to exert the minimum feasible level of effort.

To identify how cooperation evolves through time we use panel data on each worker's productivity from personnel files and exploit three quasi-random sources of variation in the data – that workers begin employment at different times, that there is a divergence between how long a worker has been exposed to the relative incentive scheme and how long they have actually worked under it, and that workers are re-allocated to different groups of co-workers each day.

We find that individuals learn how to cooperate, namely a given individual works less hard as her exposure to the relative incentive scheme increases. This effect is strongest for workers who begin employment when the scheme is first introduced. Exposure to the scheme has a significantly smaller effect for workers who arrive when the cooperative norm is already well established.

Individuals cooperate more when they work alongside co-workers that have been exposed to the scheme for longer. This effect is strongest for workers who start employment later, when the relative incentive scheme has been in place for at least a month. Moreover, individual and group exposure are substitutes -- the marginal effect of individual exposure falls as group exposure rises and vice versa.

Finally, the arrival of new workers who are unaware of the cooperative norm disrupts cooperation. In other words, existing workers cooperate less when they work alongside new arrivals. The disruption is however only temporary as new arrivals conform within a week of starting employment. Finally, workers learn how to cope with new arrivals -- the effect on the productivity of existing workers of these new arrivals is smaller when the group of co-workers have greater levels of exposure to the relative incentive scheme.

While in our setting the rationale for cooperation stems from the incentive contract individuals are offered, our results apply more generally to games of public goods provision and common resource management. A first important implication of our results is to provide field evidence to corroborate experimental results in such games. We show that individuals cooperate outside of the laboratory, and in line with existing experimental evidence, if individuals can communicate and socially sanction other players, cooperative norms tend to strengthen over time, and become more robust to the arrival of new players.

Second, under the relative incentive scheme workers have incentives to cooperate because each worker's effort imposes a negative externality on co-workers. Of course, in many strategic environments the actions of individuals place positive externalities on others. Understanding the establishment, evolution, and robustness of cooperative norms in such strategic environments, remains open findingsto future research.

References

- Andreoni, James (1988) "Why Free Ride? Strategies and Learning in Public Goods Experiments," *Journal of Public Economics* 37: 291-304.
- Apestequia, J. S., S. Huck and J. Oechssler (2002) "Imitation – Theory and Evidence," mimeo University of Bonn.
- Bandiera, O., I. Barankay, and I. Rasul (2005a) "Social Preferences and the Response to Incentives: Evidence from Personnel Data," *Quarterly Journal of Economics* 120: 917-62
- , ——— and ——— (2005b) "Incentives for Managers and Inequality Among Workers: Evidence from a Firm Level Experiment," mimeo, London School of Economics.
- Carpenter, J. and E. Seki (2005) "Competitive Work Environments and Social Preferences: Field Experimental Evidence from a Japanese Fishing Community," mimeo Middlebury College.
- Fehr, E. and S. Gächter (2000) "Cooperation and Punishment in Public Goods Experiments", *American Economic Review*, 90: 980-94.
- Freeman, R. B. and M. M. Kleiner (2005) "The Last American Shoe Manufacturers: Decreasing Productivity and Increasing Profits in the Shift from Piece Rates to Continuous Flow Production," *Industrial Relations*, 44: 307-30.
- Harrison, G. and J. List (2004) "Field Experiments," *Journal of Economic Literature*, 42: 1009-55.
- Huck, S., H-T. Normann and J. Oechssler (2004) "Through Trial and Error to Collusion," *International Economic Review*, 45: 205-24.
- Isaac, M., K. McCue and C. Plott (1985) "Public Good Provision in an Experimental Environment," *Journal of Public Economics*, 26: 51-74.
- Ledyard J. (1995) "Public Goods: A Survey of Experimental Research" in Kagel.J, Roth.A (eds.), *The Handbook of Experimental Economics*, Princeton: Princeton University Press, 111-94.

- List, J. (2003) "Does Market Experience Eliminate Market Anomalies?" *Quarterly Journal of Economics*, 118: 41-71.
- (2004) "Neoclassical Theory Versus Prospect Theory: Evidence from the Marketplace," *Econometrica*, 72: 615-25.
- Masclet, D., C. Noussair, S. Tucker and M-V. Villeval (2003) "Monetary and Nonmonetary Punishment in the Voluntary Contributions Mechanism," *American Economic Review*, 93: 366-80.
- Ostrom, E., J. M. Walker, and R. Gardner (1992). "Covenants With and Without a Sword: Self-governance is Possible," *American Political Science Review*, 86: 404-17.
- Roy, D. (1952) "Quota Restriction and Goldbricking in a Machine Shop," *American Journal of Sociology*, 62: 427-42.
- Slonim, R. and A. Roth (1998) "Learning in High Stakes Ultimatum Games: An Experiment in the Slovak Republic," *Econometrica*, 66: 569-96.