DISCUSSION PAPER SERIES

DP15498
(v. 2)

Workplace Incentives and Organizational Learning

Francesco Amodio and Miguel Martinez-Carrasco

DEVELOPMENT ECONOMICS

INDUSTRIAL ORGANIZATION

LABOUR ECONOMICS

ISSN 0265-8003

Workplace Incentives and Organizational Learning

Francesco Amodio and Miguel Martinez-Carrasco

Discussion Paper DP15498
First Published 26 November 2020
This Revision 24 February 2021

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

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Development Economics
- Industrial Organization
- Labour Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

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

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

Copyright: Francesco Amodio and Miguel Martinez-Carrasco

Workplace Incentives and Organizational Learning

Abstract

This paper studies learning within organizations when incentives change. We use a simple principal-agent model to show how, in the presence of imperfect information over the shape of the production function, workers' effort choice changes over time as information is disclosed and processed. We also show that changing incentives can trigger this learning process. We test this prediction using personnel data from an egg production plant in Peru. Exploiting a sudden change in the contract parameters, we find that workers learn from each other over the shape of the production function. This adjustment process is costly for the firm.

JEL Classification: D22, D24, J24, J33, M11, M52, M54, O12

Keywords: Organizational Learning, workplace incentives, Inputs

Francesco Amodio - francesco.amodio@mcgill.ca
McGill University and CEPR

Miguel Martinez-Carrasco - ma.martinezc1@uniandes.edu.co
Universidad de Los Andes

Acknowledgements
We are thankful to Oriana Bandiera, Amit Khandelwal, Fabian Lange, Rocco Macchiavello, and Eric Verhoogen for their comments and suggestions. Errors remain our own.

Workplace Incentives and Organizational Learning*

Francesco Amodio† Miguel A. Martinez-Carrasco‡

February 23, 2021

Abstract

This paper studies learning within organizations when incentives change. We use a simple principal-agent model to show how, in the presence of imperfect information over the shape of the production function, workers' effort choice changes over time as information is disclosed and processed. We also show that changing incentives can trigger this learning process. We test this prediction using personnel data from an egg production plant in Peru. Exploiting a sudden change in the contract parameters, we find that workers learn from each other over the shape of the production function. This adjustment process is costly for the firm.

Keywords: organizational learning, workplace incentives, inputs.

JEL Codes: D22, D24, J24, J33, M11, M52, M54, O12.

*We are thankful to Oriana Bandiera, Amit Khandelwal, Fabian Lange, Rocco Macchiavello, and Eric Verhoogen for their comments and suggestions. Errors remain our own.

†francesco.amodio@mcgill.ca, Department of Economics and Institute for the Study of International Development, McGill University, Leacock Building 514, 855 Sherbrooke St. West, Montreal, QC H3A 2T7.

‡ma.martinezc1@uniandes.edu.co, School of Management, Universidad de Los Andes, Calle 21 No. 1-20, Bogotá, Colombia.

1 Introduction

The creation and retention of knowledge are key features of organizations. Information about products, inputs and technologies is continuously disclosed, exchanged, and processed within organizations, a phenomenon known as *organizational learning* (Argote 2013). At the same time, misalignment of interests threatens organizational performance. In the presence of moral hazard or adverse selection, the provision of incentives is crucial to achieve efficiency and hire and retain the most performing agents (Prendergast 1999; Lazear 2000).

This paper studies learning within organizations when incentives change. When information is not perfect, changing incentives brings uncertainty. The objective function of agents changes, and so does their optimal decision. The lack of information on all variables evaluated at the new equilibrium generates the scope for learning. The costs associated with learning may contribute to explain differences in the adoption of personnel management practices and performance across organizations (Bloom and Van Reenen 2007, 2010).

Our analysis proceeds in three steps. First, we develop a principal-agent model where agent's effort maps into output with noise. The agent does not have full information on the global shape of the production function, and uses output as signal to update her beliefs over time. Multiple agents observe each other's effort and output, and learn from each other. If learning is local, agents only learn about the shape of the production function around a given level of effort. When the contract parameters changes, the optimal effort decision changes as well, generating scope for learning at the new equilibrium.

Second, we take this prediction to the data. We use personnel records from an egg production plant in Peru and exploit a change in workers' incentive contract parameters for identification. Workers are assigned batches of hens, exert effort to feed them, and collect eggs as output. Workers get a bonus that depends on both total output and food distributed. The weight attached to these performance measures changes over the sampling period, and the optimal feeding effort changes accordingly. We show that workers change their level of effort towards the one exerted by neighboring peers on the previous day upon observing them achieve higher output. This happens only after the announcement of the new contract, around the implementation date, and fades away gradually after 4.5 months.

Third, we quantify the profit losses associated with the incentive change that result from imperfect information and the need of learning. We estimate the amount of food workers would have distributed in the absence of experimentation, and calculate that profits would have been 5 to 6% or USD 340 to 400K higher during the implementation period in the absence of imperfect information over the production function.

This paper contributes to the literature on organizational learning. The empirical work in this domain estimates models of learning across firms (Argote et al. 1990; Irwin and Klenow 1994; Benkard 2000; Thornton and Thompson 2001). Others focus on uncertainty about the production function and the profitability and use of new production technologies, a prominent feature of low-income countries (Atkin et al. 2017). The majority of these studies investigate social learning among farmers (Foster and Rosenzweig 1995; Munshi 2004; Bandiera and Rasul 2006; Conley and Udry 2010; BenYishay and Mobarak 2018; Beaman et al. 2020). There is less evidence of social learning among workers within firms. Two exceptions are Menzel (2020), who finds evidence of knowledge spillovers among workers in Bangladeshi garment factories, and Chan et al. (2014), who study peer learning among salespeople.

The second related literature is the one on workplace incentives. A large theoretical literature exists on the trade-offs involved in performance pay, and the use of multiple performance measures (e.g., Hölmstrom 1979; Holmstrom and Milgrom 1987; Baker 1992). Starting with Lazear (2000), a number of studies have shown that that performance pay increases output. The most recent empirical literature has devoted increasing attention to working arrangements in developing countries because of the higher prevalence of piece rate pay and the higher labor intensity of the production technology (see for instance Guiteras and Jack 2018).

To the best of our knowledge, ours is the first paper showing that changing incentives can trigger learning within organizations. Imperfect information over the shape of the production function can increase the transaction costs associated with the implementation of new incentive schemes and management practices in general, possibly explaining low levels of adoption (Bloom et al. 2010; Atkin et al. 2017). We provide an estimate of such transaction costs.

The remainder of the paper is organized as follows. Section 2 outlines the conceptual framework. Section 3 introduces the empirical setting and data respectively. Section 4 illustrates the data, empirical strategy, and results. In Section 5, we estimate the transaction cost associated with the contract change. Section 6 concludes.

2 Conceptual Framework

This section illustrates a simple model that formalizes the learning mechanism uncovered by the empirical analysis that follows. Each worker i in period t independently produces output y_{it} combining effort a_{it} with an input of heterogeneous quality s_{it}. Output is given by

$$y_{it}(a_{it}, s_{it}) = s_{it} f(a_{it}) \tag{1}$$

where $f''_{it}(a_{it}) < 0$ for all a_i. Input quality s_{it} is identically and independently distributed across workers. Workers do not observe s_{it}, but know its distribution with mean μ_s and variance σ_s^2. The exact shape of $f(\cdot)$ is also unknown to the worker, who holds in each period beliefs $f_{it}(\cdot)$ over $f(\cdot)$. The combined uncertainty around s_{it} and $f(\cdot)$ is responsible for the inability of the worker to disentangle the separate contribution of effort and input quality to output. It follows that output is only an imprecise signal of the shape of the $f(\cdot)$ function, generating the scope for learning.[1]

The worker's utility cost of effort is given by $C(a_{it}) = \theta a_{it}^2/2$ with $0 < \theta < 1$. The management observes both output and effort, and motivates the worker by setting the wage equal to

$$w(y_{it}, a_{it}) = \kappa + \alpha y_{it} + (1 - \alpha) a_{it} \tag{2}$$

where κ is fixed and α is the weight attached to output relative to effort in compensation. If $\alpha = 0$, the worker is incentivized on effort only. If $\alpha = 1$, the worker is incentivized on output only. If $0 \leq \alpha \leq 1$, the worker is incentivized on both measures. This contract matches the one we observe in our empirical application, and we take it as given. Notice however that using a performance measure that captures worker's effort along a particular dimension is a common feature in many working environments, especially when workers make decisions regarding the use of some inputs.[2]

[1] Allowing the worker to partially observe input quality does not meaningfully change the model and its implications: the worker will discount the known pieces of information accordingly, but residual uncertainty over input quality will still generates scope for learning. For instance, as we show in Amodio and Martinez-Carrasco (2018), we can let $s_{it} = r_{it}\varepsilon_{it}$ with r_{it} being known and ε_{it} being unknown to the worker, who will then use y_{it}/r_{it} as signal in learning about the $f(\cdot)$ function.

[2] Amodio and Martinez-Carrasco (2020) show that using performance measures linked to one specific dimension of the effort can be optimal in a multitasking setting. Moreover, rewarding the worker for output and effort can be optimal if the worker is risk averse. This is because the two metrics are both informative of worker's choice, but vary in the amount of risk they impose on the employee, and enter the principal's payoff in different ways (Hölmstrom 1979; Baker 1992). As explained later, assuming that workers are risk averse does not change the model predictions and comparative statics.

We assume that the worker is risk neutral, has utility

$$u_{it} = \kappa + \alpha y_{it} + (1-\alpha)a_{it} - \theta\frac{a_{it}^2}{2} \tag{3}$$

and chooses the effort level a_{it} that maximizes her expected utility. Given the expected value μ_s of input quality and worker's belief $f_{it}(\cdot)$ on the production function, taking the first order condition we get

$$\alpha\mu_s f'_{it}(a_{it}) + (1-\alpha) = \theta a_{it} \tag{4}$$

which implicitly defines the optimal level of effort a^*_{it}. This changes with the wage contract parameter α. Applying the implicit function theorem we get

$$\frac{\partial a^*_{it}}{\partial\alpha} = \frac{\mu_s f'_{it}(a_{it}) - 1}{\theta - \alpha\mu_s f''_{it}(a_{it})} \tag{5}$$

from which follows that the level of effort may increase or decrease with α depending on whether its expected marginal product is higher or lower than one.[3]

Upon exerting effort, the worker observes the corresponding output realization $y_{it} = s_{it}f(a^*_{it})$. She uses output as signal to update her beliefs over the marginal product of effort in the vicinity of a^*_{it}. In order to see this, consider a Taylor series expansion approximation of $f(.)$ at 0 and assume without loss of generality $f(0) = 0$. We have

$$y_{it} \approx s_{it}f'(a^*_{it})a^*_{it} \tag{6}$$

It follows that, given her choice of effort a^*_{it} at time t, when the worker observes a higher than expected output realization – $s_{it}f'(a^*_{it})a^*_{it} > \mu_s f'_{it}(a^*_{it})a^*_{it}$ – she acknowledges that there is a positive probability that the true marginal product of effort $f'(\cdot)$ in the vicinity of a^*_{it} is higher than her belief $f'_{it}(\cdot)$. This will lead the worker to revise upwards her beliefs on $f'_{it}(\cdot)$. The opposite holds if the worker observes a lower than expected output realization.[4]

The objective of the worker is to maximize utility. If the effort cost parameter θ is low

[3] Assuming risk averse agents yields similar results. Assuming a CARA utility function and s_{it} normally distributed, working in terms of certainty equivalent we obtain the first order condition

$$\alpha\mu_s f'_{it}(a_{it}) + (1-\alpha) = \theta a_{it} + \eta f_{it}(a_{it})\sigma^2$$

where η is the agent's level of risk aversion. The comparative statics with respect to α remains unchanged.

[4] Notice that it is possible to characterize this process as standard Bayesian updating upon taking logs of equation 6 and assuming that s_{it} is log-normally distributed.

enough, higher output maps into higher utility. It follows that the optimal effort choice will change in the same direction of $f'_{it}(\cdot)$: effort will increase in the next period if output is higher than expected, and decrease otherwise. The magnitude of the change will depend on the wage contract parameter α.

If the effort choice and output of coworkers are observable, the worker will also use this information in her learning process. Specifically, given worker j's effort choice a_{jt}, worker i has expectation y^i_{jt} on j's output that is based on i's beliefs, i.e. $y^i_{jt} = \mu_s f_{it}(a_{jt})$. Whenever $a^*_{it} \neq a_{jt}$, such expected output – and corresponding utility – is lower than the one associated with a^*_{it}, as this is the optimal choice of i given her beliefs $f'_{it}(\cdot)$. As a consequence, when worker i observes a realization of coworker's output that is higher than her own, $y_{jt} > y_{it}$, she will update her beliefs over $f'(\cdot)$ and change her level of effort in the next period towards the one exerted by the coworker in the current period.

Notice that workers learn locally over $f(\cdot)$. This implies that learning over $f'(\cdot)$ at one level of effort a_{it} is not necessarily informative of $f'(\cdot)$ at a sufficiently distant level of effort. A change in α changes the optimal choice of effort, and may trigger learning over a different portion of the production function. This is the hypothesis that we take to the data.

3 The Setting

In the empirical analysis, we use personnel data from a Peruvian egg production plant. Production takes place in different *sectors*. Each sector comprises different *sheds*, long-building facilities grouping one to four *production units*.

Each worker is assigned to a given production unit and assigned a batch of laying hens. The batch as a whole is treated as a single input, as all hens within the batch are bought all together from a supplier company, raised in a dedicated sector, and moved to production accordingly. When that happens, they are assigned to a given production unit and to the same worker for their entire productive life. Hen productivity varies over time depending on hens' age and idiosyncratic productivity shocks that materialize throughout their egg-laying span that lasts approximately one year.

Output is measured by the number of eggs collected during the day. This is a function of both hen characteristics and worker's effort. Workers exert effort along three main dimensions: egg collection and storage, hen feeding, cleaning and maintenance of the unit facilities. Hen feeding is observed by the management, which records information on the number of sacks of food distributed by the worker during the day. Effort is costly, as workers need to carry multiple 50kg sacks of food a day and distribute it among all hens. The amount of food distributed is decided by the worker. Each morning, a truck arrives at the production unit and unloads a large (unbinding) number of sacks. The worker decides how many of those to distribute during the day.[5]

Changing Incentives Workers are paid every two weeks. Their salary is equal to a fixed wage plus a bonus component that depends on worker performance as measured in a randomly chosen day within the two-week pay period. The formula to calculate the bonus changed over time. In the first period, the bonus payment was calculated according to the sum of the number of sacks of food distributed by the worker and the total number of boxes of eggs collected, each box containing 360 eggs. If this quantity exceeded a given threshold, a piece rate was awarded for each unit above the threshold. On 24 February 2012, the company adopted a new bonus formula, which has been maintained thereafter. This is now based on the number of boxes of eggs collected only. Such quantity is multiplied by two, and a piece rate is awarded for each unit above a given threshold. Both the piece rate parameter and the threshold were kept the same across the two periods and contracts.

Mapping from our conceptual framework, the total number of boxes of eggs collected is a measure of output y_i, while the number of sacks of food distributed is a measure of worker's effort a_i. The first contract is such that $\alpha = 1/2$, and the second contract is such that $\alpha = 1$. This is the source of variation that we exploit to test the model predictions.

When asked about the reason for changing incentives, the management claims that workers were distributing "too much food" under the earlier incentive scheme. At the same time, managers observed close to a one-to-one relationship between egg boxes and sacks of food distributed, and they expected the contract change not to penalize the workers. We show later that the implementation of the new contract manages to reduce the amount of food distributed by workers, in line with the management's expectations and goal.

[5]Production units are independent from each other and there is no scope for technological spillovers. Egg storage and manipulation is also independent across units, as each one of them is endowed with an independent warehouse for egg and food storage.

4 Empirical Analysis

4.1 Data

We gained access to daily records for all production units in one sector from June 2011 to December 2012. These data cover the period from 8 months prior to 10 months following the change of contract. Overall, we observe 94 production units, 211 different hen batches, and 127 workers present for at least one day.

Online Appendix Table A.1 shows the summary statistics for the main variables that we use. It does so separately for the overall sample and for the three subsamples as defined by the dates in which the contract change was announced and implemented. Across all periods, workers distribute 23 sacks of food a day on average. The total number of hens per batch is heterogenous across production units over time. Dividing the total amount of food distributed by the number of hens, we derive the amount of food per hen distributed by the worker, averaging 116 grams per day.

Output is given by the number of eggs collected, averaging more than 8,000 per day. This corresponds to 0.8 daily eggs per hen on average. Consistent with the model, at least part of this variation is attributable to heterogeneity in input quality as hen productivity varies across and within units and batches over time. Part of this variation is informed by the innate characteristics of the hens. When purchased, each batch comes with detailed information on the average number of eggs per week each hen is expected to produce every week. Divided by 7, this measure of expected daily productivity varies from 0.02 to 0.93, with an average of 0.81.

Production units are grouped in 41 different sheds, 35 of them hosting more than one production unit. We calculate for each production unit the average amount of food and the average number of eggs per hen collected in neighboring production units in the same shed on the same day. We complement all this information with a survey that we administered to all workers in March 2013. We are able to use this information in combination with production data for those workers that were still present on the day of the survey, slightly more than 70% of our study sample. We use this survey to elicit information on worker's tenure at the firm.

4.2 Preliminary Evidence

In the model, we assume that output is a concave function of effort. Online Appendix Figure A.1 shows that this is the case empirically. It plots the average number of eggs per hen collected by the worker against the amount of food per hen distributed on the same day, together with its 95% confidence interval. It also plots a kinked linear approximation of the production function. The amount of food at the kink (113.25g) is set to maximize the R^2 of a kinked regression of number of eggs per hen over the amount of food distributed. Evidence shows that a local linear approximation with only one kink provides a very good approximation of the true shape of the production function.

The model also assumes that input quality is at least partially unknown to the worker. To support this assumption, we regress daily output as measured by eggs per hen collected by the worker over the expected productivity measure provided by the batch supplier. The corresponding coefficient estimate is equal to 0.85 and significant at the 1% level. More importantly, this known measure of expected productivity only explains around 40% of the variation in daily output, up to 65% when including the full sets of worker and batch fixed effects, and their interactions that also capture "match effects." This indicates that idiosyncratic productivity shocks that materialize over the hen life cycle affect their productivity, and that such residual variation in input quality matters.

On 29 November 2011, the firm announced that it would implement a new salary structure, changing the weight α attached to output from $1/2$ to 1. The change was implemented on 24 February 2012. Without further restrictions, our model delivers ambiguous predictions on the impact of such change on effort as measured by the amount of food distributed by the worker. Yet, Figure A.1 shows that the slope of the production function is always lower than one. In this case, equation 5 delivers a clear prediction: effort decreases when the weight α attached to output in the bonus formula increases.

Figure 1 shows the average amount of food distributed daily over time during the sample period. The graph shows the smoothed average value together with its 95% confidence interval. The two vertical red lines correspond to the dates of announcement and implementation of the new contract. The amount of food distributed falls discontinuously on announcement and implementation dates, and then seems to stabilize in the later period at a level that is lower than the initial one. This pattern suggests that the new contract was successful in reducing the amount of food distributed by workers. Yet, variation over time could also be driven by other factors that affect the production process and workers' choices differently on each day. This could explain for instance the slight decrease in the amount of food distributed just before the announcement date. We explain in details in the next section whether this is problematic and, in general, the conditions under which the presence of unobserved determinants of workers' food choice would invalidate our identification strategy.

Online Appendix Table A.1 provides additional evidence of the fall in the amount of food distributed by workers. After the implementation of the new contract, workers distribute on average one sack of food less relative to the period before the announcement. This corresponds to a decrease in food per hen of about 8 grams, or 6.6% of the baseline mean.

Figure 2 plots the distribution of the amount of food per hen distributed by workers in each day and separately for the period before the announcement of the incentive change, between the announcement and implementation date, and after implementation. First, the figure shows how the whole distribution shifts leftwards as the new contract is first announced and then implemented. Second, the distribution is more dispersed in the period between announcement and implementation dates than in the other two periods. This is suggestive of experimentation during that time.

4.3 Identification Strategy

Our hypothesis is that changing incentives triggers learning among workers over the shape of the production function around the new optimal level of effort. The spatial arrangement of production units is such that neighboring peers can interact and observe each other. We would therefore expect workers to use the available information on food distributed and output of peers to update their beliefs and inform their own food choice accordingly. This would generate a positive correlation between the choices of neighboring coworkers. But, finding evidence of such correlation does not necessarily mean that workers learn from each other. First, unobserved common factors may independently affect the effort choice of coworkers and tilt them in the same direction. Second, the simultaneous determination of their decisions makes it difficult to identify causal relationships because of the so-called reflection problem (Manski 1993).

To overcome these issues, we adopt a regression framework that builds upon Conley and Udry (2010) and their study of pineapple growers in Ghana. We look at changes in workers' effort choices over time, and whether they adjust towards their peer choice differentially when the latter achieve higher output. To operationalize this approach, we define for each worker i operating batch b on day t the variable

$$M_{ibt} = (a_{jbt-1} - a_{ibt-1}) \times \mathbb{I}\{y_{jbt-1} > y_{ibt-1}\} \tag{7}$$

where a_{jbt-1} is the average effort choice – sacks of food distributed – of neighboring coworkers on the previous day, a_{ibt-1} is the effort choice of the worker on that same day, and $\mathbb{I}\{y_{jbt-1} > y_{ibt-1}\}$ is an indicator of peer relative success, equal to one if the average output – eggs per hen – of neighboring coworkers was higher than own output.

We then implement the following baseline regression specification

$$\Delta a_{ibt} = \beta\, M_{ibt} + \gamma\, \textit{Post}_t \times M_{ibt} + \mathbf{X}'_{ibt}\kappa + \delta_t + \theta_i + u_{ibt} \tag{8}$$

where $\Delta a_{ibt} = a_{ibt} - a_{ibt-1}$ is the change in the effort choice of worker i from one day to the other, and $\textit{Post}_t$ is a dummy equal to one in the period after the announcement of the new contract. The coefficient β captures whether workers change their level of effort towards the one exerted by neighboring peers on the previous day upon observing them achieve a higher output. γ captures whether this occurs systematically and differentially after the announcement of the new contract. The vector $\mathbf{X}_{ibt}$ includes the lagged own and coworkers' input choice and output as well as the total number of hens in the batch. δ_t captures day fixed effects, which account for and net out all determinants of food choice that vary over time and affect all workers in the same way. Similarly, θ_i captures worker fixed effects, which account for and net out all determinants of food choice that are idiosyncratic to each worker and do not vary over time. Finally, the term u_{ibt} captures any residual determinants of change in the amount of food distributed. We cluster standard errors along the two dimensions of shed and day to account for unobserved correlation between such residuals across observations belonging to the same shed or day.

A first concern with our identification strategy is that the presence of unobserved determinants of food choice could bias the estimated γ. However, this would be the case only insofar as these were systematically related to peer relative success and its indicator nested in the M_{ibt} variable, and if this was differentially the case after the announcement of the new contract. A second concern is related to "match effects," meaning the possibility that batches with particular characteristics are assigned to workers that are particularly good or well-suited to handle batches with those characteristics. Once again, this would be problematic only if the productivity of a match was systematically related to own and coworkers' food choice and peer relative success, and differentially so after the announcement of the new contract.

Before continuing, we investigate the extent of variation in the M_{ibt} variable. Its value is different than zero for around 45% of observations in the sample, and positive (negative) for 25% (20%) of the sample. These relative frequencies are not systematically different across subsamples as defined by whether the observation belongs to before or after the announcement of the new contract, or to workers with lower or higher than median tenure. The R^2 of a regression of M_{ibt} over the full set of worker fixed effects is equal to 0.40, up to 0.45 when also including all their interactions with the post-announcement dummy. This indicates that M_{ibt} varies within workers more than it does between workers.

4.4 Results

Table 1 reports the regression coefficient estimates. Column 1 shows the estimated β from a regression specification that only includes M_{ibt} and the vector of controls $\mathbf{X}_{ibt}$ as independent variables. In column 2, we augment the specification with the full set of day fixed effects. The estimated β is positive and significant at the 1% level. In column 3, we include the interaction between the M_{ibt} variable and the post-announcement dummy.[6] The estimated β is now close to zero and insignificant while the estimate of γ is positive and significant at the 5% level. It becomes significant at the 1% level upon adding worker fixed effects in column 4. In column 5, we include as additional control a dummy that equals one if expected productivity (estimated by the batch supplier) is above the median. We interpret this as evidence that the announcement of the new contract triggers learning among coworkers.

As explained earlier, when hen batches are moved to production they are assigned to a given production unit and to the same worker for their entire productive life. It is reasonable to expect that the extent of unobserved variation in input quality decreases over time. If this is the case, the scope for learning would be highest among workers handling a newly assigned batch. We test this hypothesis by including the triple interaction between the M_{ibt} variable, the post-announcement dummy, and a variable $\textit{MatchDuration}_{ibt}$ capturing the time elapsed since current batch assignment in months.[7] Column 6 of Table 1 reports the corresponding results. The triple interaction coefficient estimate is negative and highly significant while the estimated γ remains positive, larger

[6] Notice that the post-announcement dummy itself is not included as its variation is absorbed by the full set of day fixed effects.

[7] We also include this variable on its own as additional control.

in magnitude, and highly significant. This suggests that the scope for learning is indeed highest among workers handling a newly assigned batch, and decreases thereafter.[8]

In column 7 and 8, we implement the regression specification in column 5 separately for the subsample of production units assigned to workers with lower and higher than median tenure. The estimated γ is significant only for workers with high tenure. While surprising at first, we interpret this as evidence that workers with longer experience are more capable of monitoring their peers, elaborate the information that becomes available, and act accordingly.[9]

To get a sense of the magnitude of the estimated γ, notice that for 99% of our sample the value of Δa_{ibt} is between -1 and 1, with a standard deviation of 0.143. In the post-announcement period, upon observing their peers achieve a higher output, workers change their food choice by 4 to 5% of such standard deviation.

The previous specification pools together all observations belonging to the pre and post-announcement period. Yet, we would expect learning to occur only for a limited amount of time. We thus augment the regression specification in equation 8 with the interactions of M_{ibt} with a set of dummies that identify each two-week pay period. We omit and use as reference the pay period when the contract change was announced. Figure 3 plots the coefficient estimates associated with these interaction terms over time, together with their 95% confidence intervals. The two vertical red lines correspond to the periods of announcement and implementation of the new contract. Estimates are not significantly different from zero for the whole period before the announcement of the new contract. They become positive and significant at the 5% level shortly before implementation, remain significant for several periods, then return insignificant. We interpret this pattern as showing evidence that learning was absent prior to the announcement of the new contract. It only materializes thereafter, spikes around the implementation date, and fades away gradually after 4.5 months.

[8] We also tested if learning is heterogeneous depending on the number of neighboring coworkers. We included the triple interactions between the M_{ibt} variable, the post-announcement dummy, and dummies for whether the worker has two, three or four neighboring coworkers, workers having one being the omitted category. The baseline estimate of γ remains positive and highly significant while none of the interaction coefficients are. Results are available from the authors upon request.

[9] We reach the same conclusion when implementing a specification with the triple interaction between the M_{ibt} variable, the post-announcement dummy, and a dummy for workers with higher than median tenure.

5 The Cost of Learning

In this section, we summarize our attempt to quantify the profit losses associated with the contract change and due to imperfect information and need of learning. We provide the full details of the estimation procedure in Online Appendix A.2.

The fundamental challenge is that we do not observe the counterfactual, i.e. what would have happened to feeding effort, output, and profits in the presence of complete information. That is, we cannot disentangle the variation in the variables of interest that is driven by learning and experimentation from the one determined by idiosyncratic shocks such as changes in output and input prices.

We address this challenge as follows. In the first step, we filter out variation in food choice across days, production units and batches by regressing the amount of food per hen distributed by the worker over the corresponding three sets of fixed effects. We then split the sample in three periods: the one before the announcement of the new contract, the one during which learning occurs, and the one after. The length of the second period is informed by Figure 3 and given by those two-week intervals in which the estimated coefficient capturing knowledge spillovers is positive and significant. Notice that the date of the implementation of the new contract falls within this second period.

In the second step, we use the estimated residuals from the first step to derive the average residual of food distributed per hen in the three periods. We consider the averages in the first and last period as informative of the equilibrium level of feeding effort under the old and new contract respectively.

In the third and last step, we estimate the counterfactual feeding effort choice during the adjustment period by re-centering the distribution of residuals as follows. We subtract the average of the period and add the one of the first period to all observations prior to the implementation date, and do the same using the average of the third period to those after the implementation date. In other words, we re-center the observed distribution of residual food choice in the second period using the averages in the first and third period for the days before and after the implementation of the new contract respectively. Finally, we add to these counterfactual residuals the fixed effects estimated in the first step and obtain the counterfactual choice of food distributed. Not surprisingly given the way we obtain the counterfactual, actual and counterfactual variables differ only in the second period, with the counterfactual amount of food distributed being higher in the absence of experimentation.

Upon obtaining the counterfactual amount of food distributed, we can derive counterfactual output, revenues, food costs, and bonuses paid to the workers. We implement a regression of eggs per hen over the kinked function of food per hen specified in Section 4, and the full set of day, production unit, and batch fixed effects. We then use the estimated coefficients and the counterfactual food per hen to obtain counterfactual output. We calculate revenues using the information on output prices that the firm made available to us. Similarly, we calculate food costs using the information on the price of a sack of food. We use the actual compensation formula before and after the contract change to calculate the bonuses paid to employees. Finally, we combine all this information to calculate profits. Figure 4 shows the corresponding results. The area between the two lines measures the daily average profit loss over time.[10]

To get a sense of the uncertainty surrounding these estimates, we implement a bootstrap-type procedure sampling with replacement from the full dataset and repeating all steps described above 200 times. Online Appendix Table A.2 shows the results from this exercise for each of the variables we use to calculate profits, with standard deviations in parenthesis. We estimate a revenue loss of USD 560K and a profit loss of USD 373K. According to our calculations, profits would have been 5.5% higher over the learning period in the presence of complete information on the global shape of the production function.[11]

6 Conclusions

This paper shows that changing incentives triggers learning among coworkers within firms. We present a principal-agent framework that illustrates how a change in the contract parameters can trigger learning and experimentation over a new, unexplored portion of the production function. We take this hypothesis to the data using personnel records from a Peruvian egg production plant. We show that workers change their level of effort towards the one exerted by neighboring peers on the previous day upon observing them achieve a higher output, which we interpret as evidence of learning among coworkers. The learning process lasts around 4.5 months, and brings about estimated profit losses of USD 340 to 400K. This finding suggests that varying degrees of information completeness and related costs can explain at least part of the variation

[10]Online Appendix Figures A.2 and A.3 show the smoothed averages of all actual and counterfactual variables used to calculate profits.

[11]Online Appendix Figure A.4 shows the distribution of absolute and relative profit gains across the 200 repetitions.

we observe across firms within and across countries in the adoption of personnel management practices. Understanding whether and how this is the case is an open question that we leave for future research.

References

Amodio, F. and M. A. Martinez-Carrasco (2018). Input Allocation, Workforce Management and Productivity Spillovers: Evidence from Personnel Data. *Review of Economic Studies 85*(4), 1937–1970.

Amodio, F. and M. A. Martinez-Carrasco (2020). Inputs, Incentives, and Self-Selection At the Workplace. mimeo.

Argote, L. (2013). *Organizational Learning*. Springer US.

Argote, L., S. L. Beckman, and D. Epple (1990). The persistence and transfer of learning in industrial settings. *Management Science 36*(2), 140–154.

Atkin, D., A. Chaudhry, S. Chaudry, A. K. Khandelwal, and E. Verhoogen (2017). Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan. *The Quarterly Journal of Economics 132*(3), 1101–1164.

Baker, G. P. (1992). Incentive Contracts and Performance Measurement. *Journal of Political Economy 100*(3), 598–614.

Bandiera, O. and I. Rasul (2006). Social networks and technology adoption in northern mozambique*. *The Economic Journal 116*(514), 869–902.

Beaman, L., A. BenYishay, J. Magruder, and A. M. Mobarak (2020). Can Network Theory-based Targeting Increase Technology Adoption? *American Economic Review*. Forthcoming.

Benkard, C. L. (2000). Learning and forgetting: The dynamics of aircraft production. *American Economic Review 90*(4), 1034–1054.

BenYishay, A. and A. M. Mobarak (2018, 07). Social Learning and Incentives for Experimentation and Communication. *The Review of Economic Studies 86*(3), 976–1009.

Bloom, N., B. Eifert, A. Mahajan, D. McKenzie, and J. Roberts (2013). Does Management Matter? Evidence from India. *The Quarterly Journal of Economics 128*(1), 1–51.

Bloom, N., A. Mahajan, D. McKenzie, and J. Roberts (2010). Why Do Firms in Developing Countries Have Low Productivity? *American Economic Review: Papers & Proceedings 100*(2), 619–23.

Bloom, N. and J. Van Reenen (2007). Measuring and Explaining Management Practices

Across Firms and Countries. *The Quarterly Journal of Economics 122*(4), 1351–1408.

Bloom, N. and J. Van Reenen (2010). Why Do Management Practices Differ across Firms and Countries? *Journal of Economic Perspectives 24*(1), 203–24.

Chan, T. Y., J. Li, and L. Pierce (2014). Learning from peers: Knowledge transfer and sales force productivity growth. *Marketing Science 33*(4), 463–484.

Conley, T. G. and C. R. Udry (2010, March). Learning about a New Technology: Pineapple in Ghana. *American Economic Review 100*(1), 35–69.

Foster, A. D. and M. R. Rosenzweig (1995, December). Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture. *Journal of Political Economy 103*(6), 1176–1209.

Guiteras, R. P. and B. K. Jack (2018). Productivity in piece-rate labor markets: Evidence from rural malawi. *Journal of Development Economics 131*, 42 – 61.

Hölmstrom, B. (1979). Moral Hazard and Observability. *The Bell Journal of Economics 10*(1), 74–91.

Holmstrom, B. and P. Milgrom (1987). Aggregation and linearity in the provision of intertemporal incentives. *Econometrica 55*(2), 303–328.

Irwin, D. and P. Klenow (1994). Learning-by-doing spillovers in the semiconductor industry. *Journal of Political Economy 102*(6), 1200–1227.

Lazear, E. P. (2000, December). Performance pay and productivity. *American Economic Review 90*(5), 1346–1361.

Manski, C. F. (1993). Identification of Endogenous Social Effects: The Reflection Problem. *Review of Economic Studies 60*(3), 531–42.

Menzel, A. (2020). Knowledge Exchange and Productivity Spill-overs in Bangladeshi Garment Factories. mimeo.

Munshi, K. (2004). Social learning in a heterogeneous population: technology diffusion in the indian green revolution. *Journal of Development Economics 73*(1), 185 – 213.

Prendergast, C. (1999, March). The provision of incentives in firms. *Journal of Economic Literature 37*(1), 7–63.

Thornton, R. A. and P. Thompson (2001). Learning from experience and learning from others: An exploration of learning and spillovers in wartime shipbuilding. *American Economic Review 91*(5), 1350–1368.

Exhibits

Figure 1: Food Choice Over Time

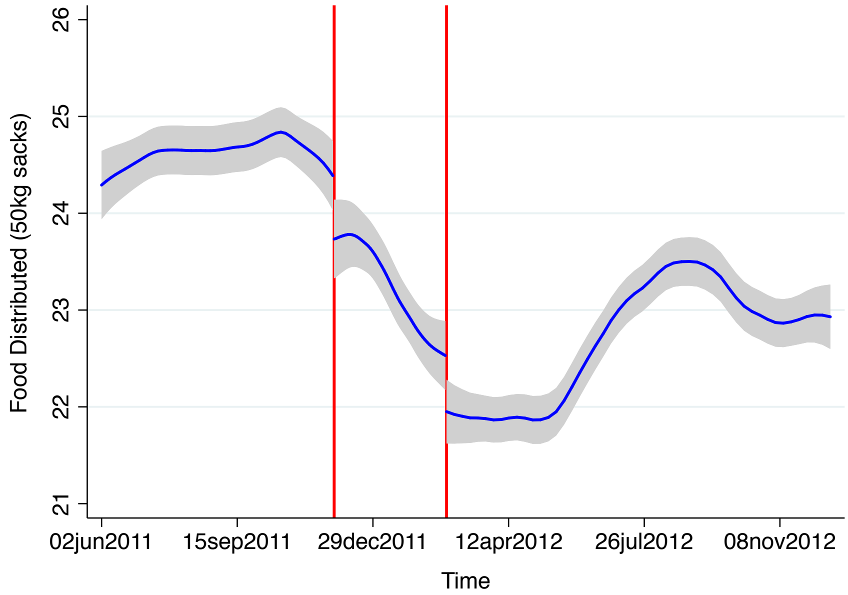

Notes. The figure plots the smoothed average of the total number of 50kgs sacks of food distributed across all production units in a given day, together with its 95% confidence interval. The two vertical lines correspond to the dates of announcement and implementation of the new contract. The amount of food distributed is stable before the announcement, falls discontinuously on announcement and implementation dates, and stabilizes again in the later period at a level that is lower than the initial one.

Figure 2: Distribution of Food Choice Across Periods

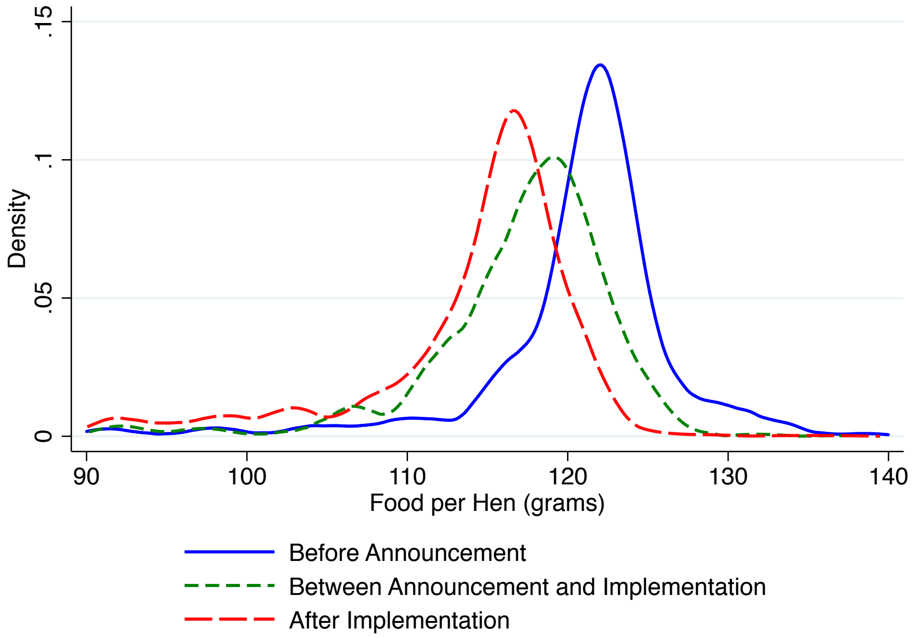

Notes. The figure plots the smoothed kernel density of grams of food per hen distributed in each day across workers and separately in the period before, during, and after the implementation of the contract change.

Table 1: Incentive Change and Learning Estimates

	Change in Food Distributed							
							Low Tenure	High Tenure
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
M_{ibt}	0.0029***	0.0046***	0.0008	0.0004	0.0008	0.0008	0.0004	0.0021
	(0.0009)	(0.0011)	(0.0015)	(0.0016)	(0.0016)	(0.0017)	(0.0017)	(0.0028)
Post$_t \times M_{ibt}$			0.0055**	0.0065***	0.0064***	0.0118***	0.0038	0.0068**
			(0.0021)	(0.0022)	(0.0022)	(0.0026)	(0.0032)	(0.0032)
Post$_t \times M_{ibt} \times$ *MatchDuration*$_{ibt}$						-0.0010***		
						(0.0003)		
High Expected Productivity$_{ibt}$					-0.0721***	-0.0317	-0.0669***	-0.0835***
					(0.0126)	(0.0201)	(0.0178)	(0.0197)
Number of Hens$_{ibt}$	0.0003***	0.0004***	0.0004***	0.0004***	0.0004***	0.0004***	0.0003***	0.0005***
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0001)	(0.0001)
Other Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day FE	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Worker FE	No	No	No	Yes	Yes	Yes	Yes	Yes
Observations	41490	41490	41490	41489	41489	41489	13539	16108
R^2	0.0937	0.1378	0.1388	0.1529	0.1553	0.1589	0.1389	0.2055

Notes. (* p-value<0.1; ** p-value<0.05; *** p-value<0.01) Two-way clustered standard errors, with residuals grouped along both shed and day. Dependent variable is the change in the amount of food distributed by the worker from the day before to the day of observation, as measured by the change in the number of 50kg sacks distributed, $a_{it} - a_{ibt-1}$. M_{ibt} is equal to the interaction of difference between the average amount of food distributed by neighboring coworkers and the one distributed by the worker on the previous day, $a_{jbt-1} - a_{ibt-1}$, with a dummy equal to one if neighboring coworkers achieved higher average output (eggs per hen), $\mathbb{I}\{y_{jbt-1} > y_{ibt-1}\}$. *Post*$_t$ is a dummy equal to one for all observations belonging to the period after the announcement of the contract change. *High Expected productivity*$_{ibt}$ is a dummy equal to one when expected productivity (as estimated by the batch supplier) is higher than the median. The vector of controls includes the amount of food distributed by the worker and the average of neighboring coworkers on the previous day a_{ibt-1} and a_{jbt-1}, and their output y_{ibt-1} and y_{jbt-1}. In columns (6), we include as regressor the triple interaction between M_{ibt}, *Post*$_t$ and the variable *MatchDuration*$_{ibt}$ that captures the time elapsed since the current batch was first assigned to the worker, while also including *MatchDuration*$_{ibt}$ separately as additional control. In columns (7) and (8) the sample is split between workers with lower and higher than median tenure respectively, and restricted to those observations that we can merge with the survey of workers that we administered in March 2013.

Figure 3: Incentive Change and Learning Over Time

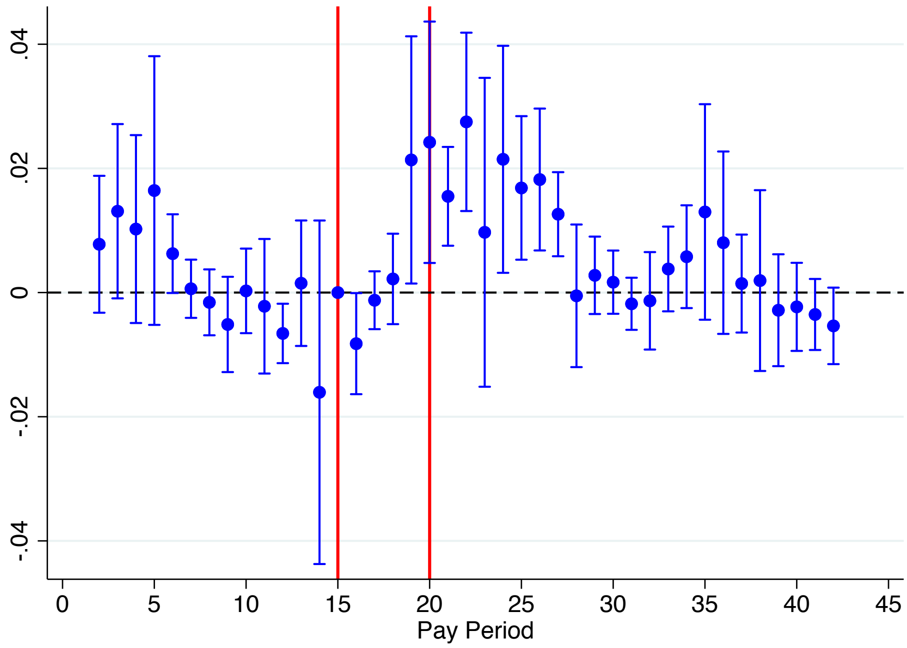

Notes. The figure plots the coefficient estimates associated with the whole set of interactions between the M_{ibt} variable specified in Section 4.3 and a dummy for each two-week pay period. Estimates are obtained from an augmented version of regression specification in equation 8 that includes all these interactions. The two vertical lines correspond to the periods of announcement and implementation of the new contract. The announcement pay period is used as reference. The coefficient estimate that captures learning among coworkers increases after the announcement and becomes positive and significant around and after the implementation date, consistent with Figure 1.

Figure 4: Actual and Counterfactual Profits

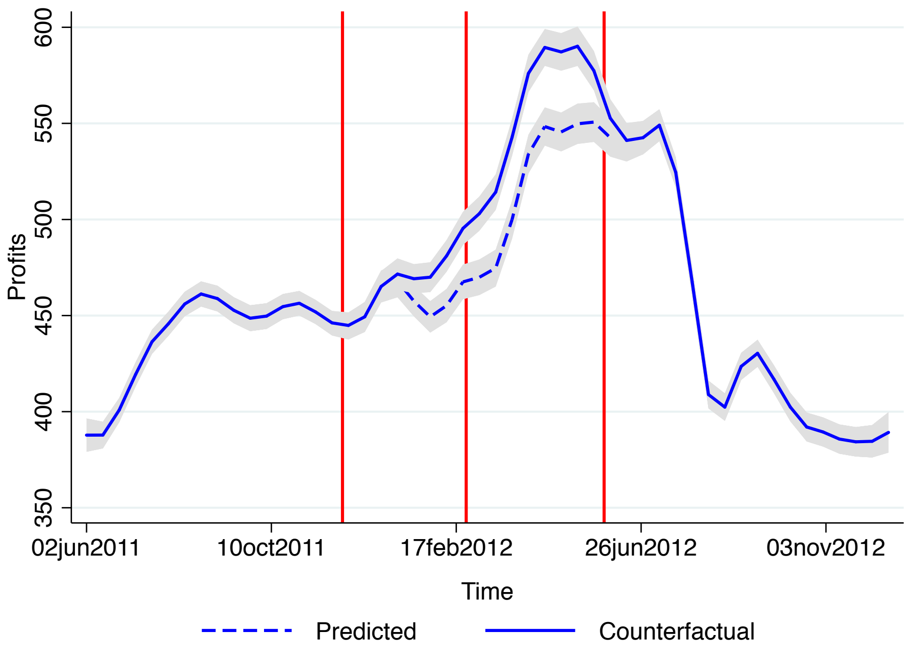

Notes. The figure shows the predicted and counterfactual smoothed average profits per day. The estimation is described in Section 5 and Online Appendix A.2. The first two vertical lines indicate the period of announcement and implementation of the incentive change respectively, while the third one corresponds to the last period in which learning occurs according to the results depicted in Figure 3.

A Appendix for Online Publication

A.1 Additional Tables and Figures

Table A.1: Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
	Panel A – All Sample				
Food Distributed (50kg sacks)	23.339	8.762	0.5	37	41490
No. of Hens	10068.737	3707.413	353	15517	41490
Food per Hen (gr)	115.937	9.145	66.774	163.235	41490
Total Eggs Collected	8154.689	3552.199	0	15131	41490
Total Eggs per Hen	0.809	0.179	0	1	41490
Expected Productivity	0.814	0.138	0.02	0.934	40648
No. of Neighboring Coworkers	2.011	0.87	1	4	41490
Food Distributed by Coworkers (avg)	23.334	8.763	0.5	37	41490
Match Duration (months)	4.265	3.417	0.033	17.567	41490
Experience (years)	5.395	3.531	0.038	15.781	32892
	Panel B – Before Annoucement				
Food Distributed (50kg sacks)	24.443	9.029	3	37	14156
No. of Hens	10134.422	3698.474	1311	15396	14156
Food per Hen (gr)	120.691	6.874	67.146	163.235	14156
Total Eggs Collected	8530.853	3531.786	0	13830	14156
Total Eggs per Hen	0.843	0.154	0	0.993	14156
Expected Productivity	0.809	0.148	0.02	0.934	14037
	Panel C – Between Announcement and Implementation				
Food Distributed (50kg sacks)	23.115	8.541	1	35	4999
No. of Hens	10010.448	3641.729	520	14963	4999
Food per Hen (gr)	115.591	9.122	68.673	159.795	4999
Total Eggs Collected	8026.72	3375.86	60	13112	4999
Total Eggs per Hen	0.806	0.162	0.005	1	4999
Expected Productivity	0.8	0.129	0.02	0.934	4802
	Panel D – After Implementation				
Food Distributed (50kg sacks)	22.689	8.568	0.5	35	22335
No. of Hens	10040.153	3727.173	353	15517	22335
Food per Hen (gr)	113.001	9.154	66.774	159.61	22335
Total Eggs Collected	7944.917	3584.261	0	15131	22335
Total Eggs per Hen	0.788	0.194	0	1	22335
Expected Productivity	0.821	0.133	0.02	0.934	21809

Notes. The table reports the summary statistics of the variable used in the empirical analysis in the overall sample and separately for the period before, during, and after the contract change.

Table A.2: Counterfactual Estimates

	Data	Simulation	Difference	% Difference
Total Eggs	374.084	379.210	5.126	0.014
(Millions)	(0.802)	(0.802)	(0.161)	(0.000)
Revenues	38.860	39.419	0.560	0.014
(USD Millions)	(0.083)	(0.083)	(0.017)	(0.000)
Food	1.077	1.090	0.013	0.012
(Millions of 50kg sacks)	(0.002)	(0.002)	(0.000)	(0.000)
Food Cost	17.816	18.001	0.186	0.010
(USD Millions)	(0.032)	(0.032)	(0.003)	(0.000)
Bonuses	0.018	0.019	0.002	0.089
(USD Millions)	(0.000)	(0.000)	(0.000)	(0.003)
Profits	21.026	21.399	0.373	0.018
(USD Millions)	(0.059)	(0.059)	(0.016)	(0.001)
Profits Adj. Period	6.754	7.126	0.373	0.055
(USD Millions)	(0.062)	(0.061)	(0.016)	(0.002)

Notes. The table shows the average and standard deviation of predicted and counterfactual variables. Both are estimated with the procedure described in Section 5 and Online Appendix A.2. Distributions are obtained by implementing a bootstrap-type procedure of resampling with replacement in 200 repetitions.

Figure A.1: Output and Feeding Effort

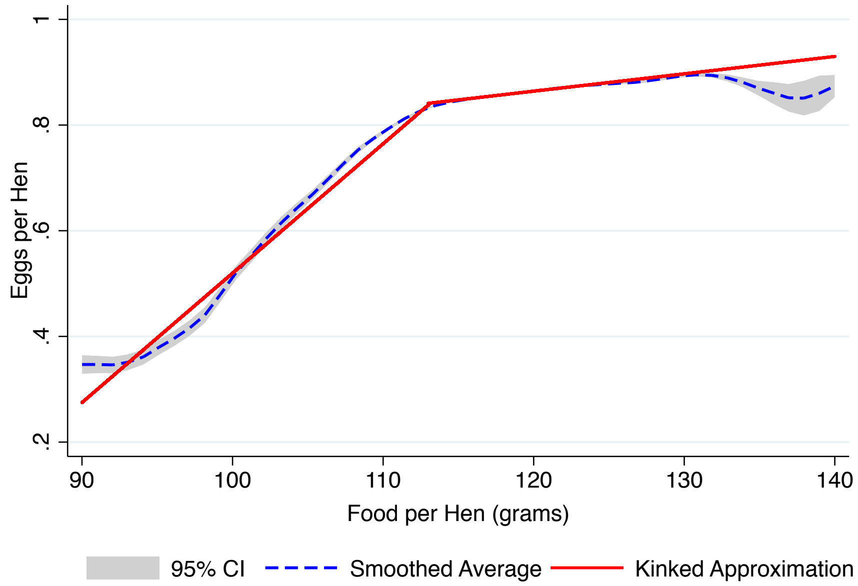

Notes. The figure plots the smoothed average of the number of eggs per hen collected by the worker over the grams of food per hen distributed in the day, together with its 95% confidence interval. It also plots a kinked linear approximation of the production function. The values of amount of food at the kink (113.25g) is chosen in order to maximize the R^2 of a kinked regression of number of eggs per hen over the amount of food distributed.

Figure A.2: Actual and Counterfactual Output and Revenues

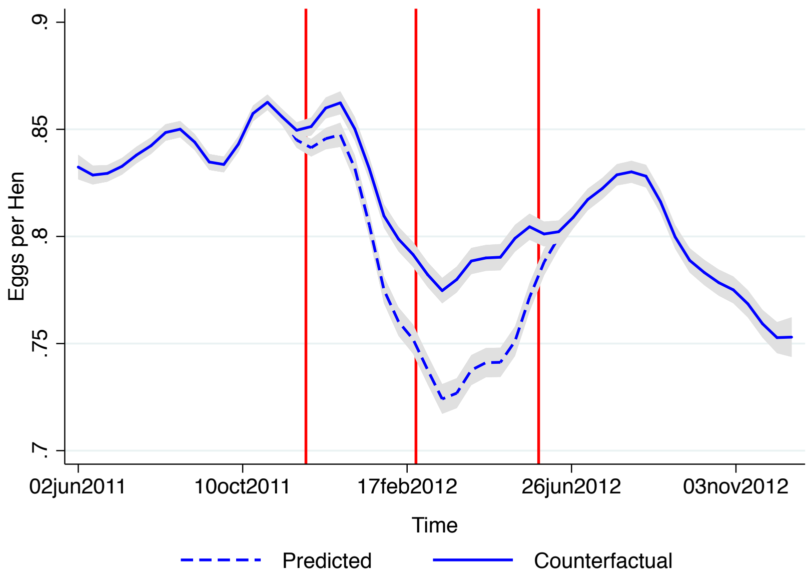

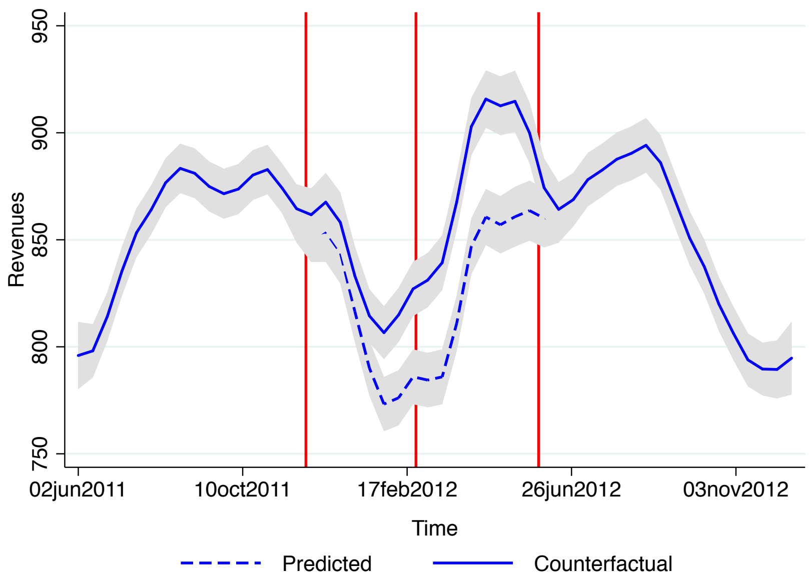

Notes. The top figure shows the predicted smoothed average of the total number of eggs collected, and its counterfactual in a simulated environment with no learning. The bottom figure shows the predicted and counterfactual amount of revenues per day. The procedure to construct these counterfactuals is described in Section 5 and Online Appendix A.2. The first two vertical lines indicate the period of announcement and implementation of the incentive change respectively, while the third one corresponds to the last period in which learning occurs according to the results depicted in Figure 3.

Figure A.3: Actual and Counterfactual Food Choice and Wages

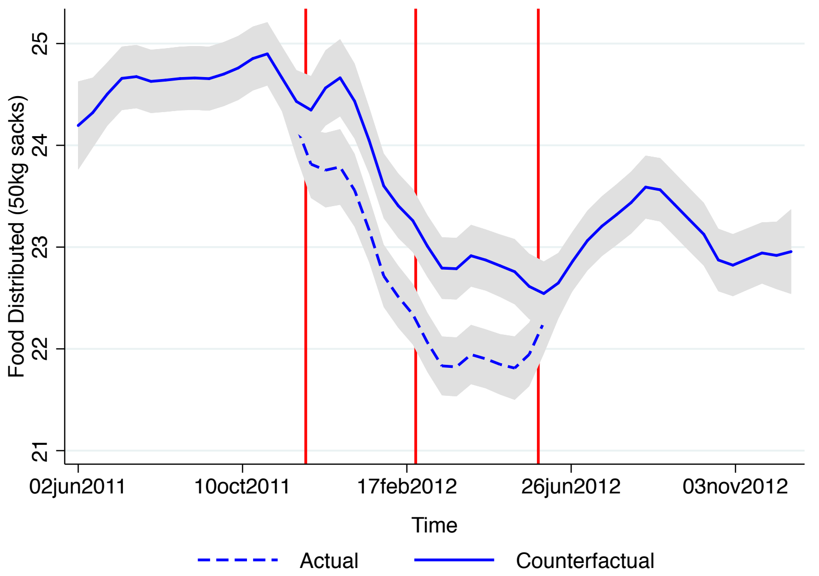

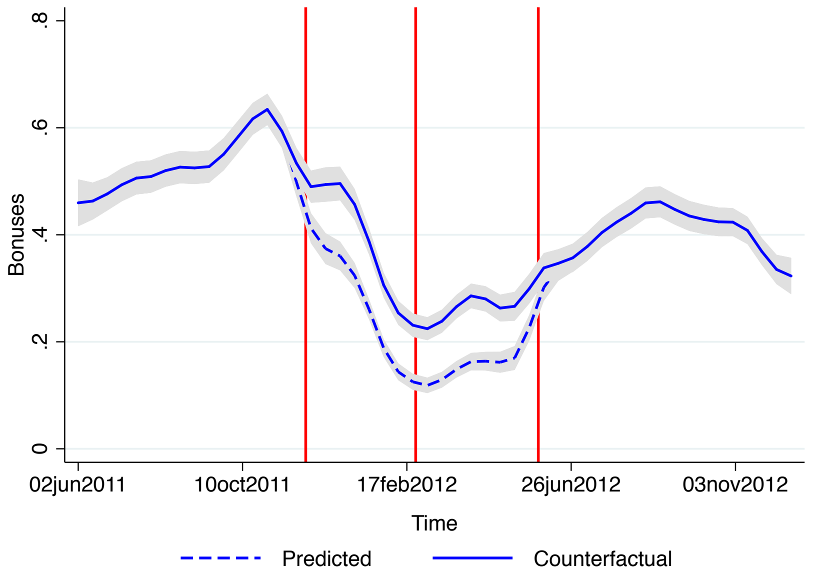

Notes. The top figure shows the actual smoothed average of the amount of food distributed by workers, and its counterfactual in a simulated environment with no learning. The bottom figure shows the predicted smoothed average of bonuses paid and its counterfactual. The procedure to construct these counterfactuals is described in Section 5 and Online Appendix A.2. The first two vertical lines indicate the period of announcement and implementation of the incentive change respectively, while the third one corresponds to the last period in which learning occurs according to the results depicted in Figure 3.

Figure A.4: Distribution of Profit Gains

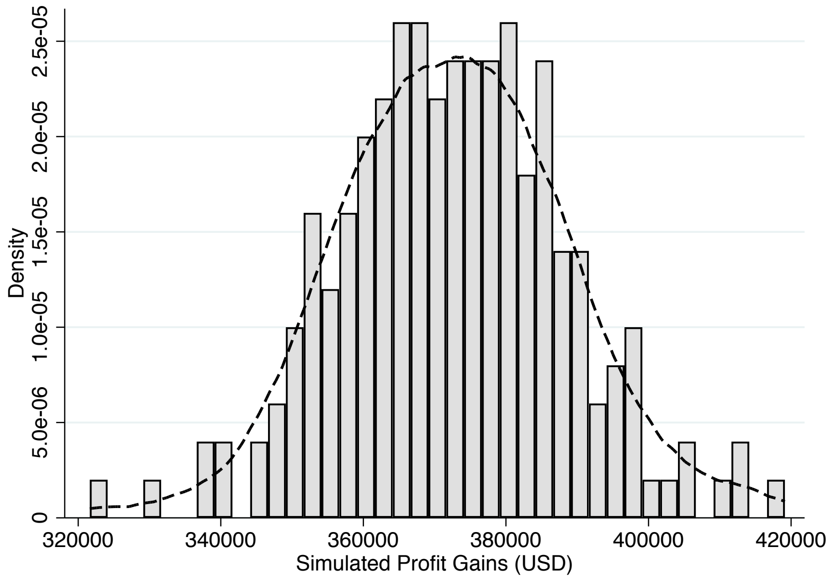

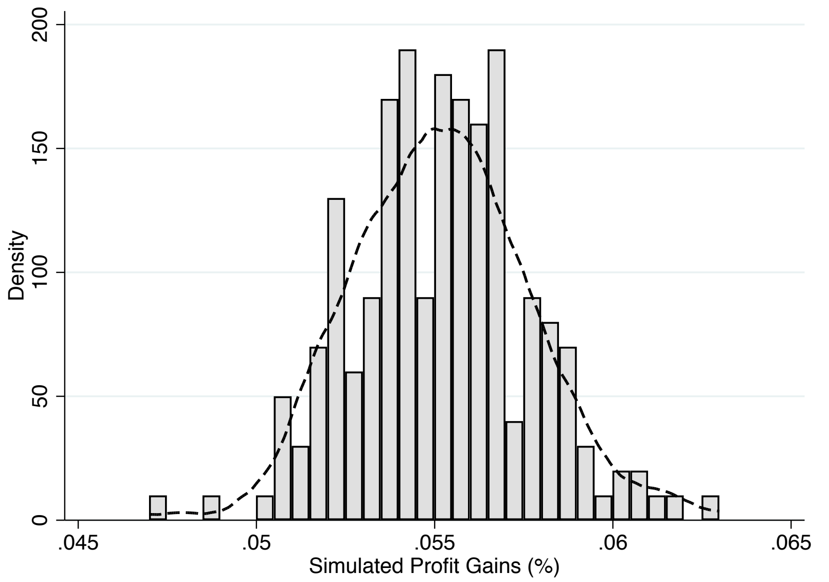

Notes. The top figure shows the distribution of overall profit gains in the absence of learning. The bottom figure shows the percentage change in profits over the adjustment period, between the date of announcement of incentive change and the last period in which learning occurs according to the results depicted in Figure 3. Predictions and counterfactuals are estimated with the procedure described in Section 5 and Online Appendix A.2. Both distributions are obtained after a bootstrap procedure of resampling with replacement in 200 repetitions.

A.2 Counterfactual Estimation

In this section, we provide the full details of the counterfactual estimation procedure summarized in Section 5.

In the first step, we implement the following regression specification

$$h_{ibt} = \theta_i + \psi_b + \delta_t + \varepsilon_{ibt} \tag{1}$$

where h_{ibt} is the amount of food per hen distributed by the worker at production unit i who is assigned batch b on day t while θ_i, ψ_b and δ_t stand for production unit, batch, and day fixed effects.

In the second step, we use the estimated residuals from the first step to derive the average residual of food distributed per hen in the periods before the announcement of the new contract, the one during which learning occurs, and the one after. The length of the second period is informed by Figure 3 and given by those two-week intervals in which the estimated coefficient capturing knowledge spillovers is positive and significant. The date of the implementation of the new contract, which we label as T, falls within this second period.

Specifically, let $\bar{\hat{\varepsilon}}_B$ be the average of $\hat{\varepsilon}_{ibt}$ in the period before the announcement of the new contract, $\bar{\hat{\varepsilon}}_D$ be the average in the period during which learning occurs, and $\bar{\hat{\varepsilon}}_A$ be the average in last period after that.

In the third step, we use these averages to re-center the distribution of residuals. That is, we obtain counterfactual residuals $\tilde{\varepsilon}_{ibt}$ as

$$\begin{aligned} \tilde{\varepsilon}_{ibt} &= \hat{\varepsilon}_{ibt} - \bar{\hat{\varepsilon}}_D + \bar{\hat{\varepsilon}}_B \qquad \text{if} \quad t < T \\ \tilde{\varepsilon}_{ibt} &= \hat{\varepsilon}_{ibt} - \bar{\hat{\varepsilon}}_D + \bar{\hat{\varepsilon}}_A \qquad \text{if} \quad t \geq T \end{aligned} \tag{2}$$

Finally, we add to these counterfactual residuals the fixed effects estimated in the first step and obtain the counterfactual choice of food distributed per hen $\tilde{h}_{ibt}$ as given by

$$\tilde{h}_{ibt} = \hat{\theta}_i + \hat{\psi}_b + \hat{\delta}_t + \tilde{\varepsilon}_{ibt} \tag{3}$$

from which we can derive the counterfactual total amount of food $\tilde{a}_{ibt}$ distributed by the worker, i.e. $\tilde{a}_{ibt} = \tilde{h}_{ibt} n_{ibt}$ where n_{ibt} is the number of assigned hens.

Upon obtaining the counterfactual amount of food distributed, we can derive counterfactual output, revenues, food costs, and bonuses paid to the workers. Consistent with Section 4, Figure A.1 and the specification in equation 1 above, we implement the following kinked regression specification

$$y_{ibt} = \beta\, h_{ibt} \times \mathbb{I}\left\{h_{ibt} < H\right\} + \gamma\, h_{ibt} \times \mathbb{I}\left\{h_{ibt} \geq H\right\} + \omega_i + \lambda_b + \phi_t + \epsilon_{ibt} \tag{4}$$

where y_{ibt} is the number of eggs per hen distributed by the worker at production unit i who is assigned batch b on day t. h_{ibt} is the amount of food per hen distributed by the worker, and H is the kink value, equal to 113.25g. ω_i, λ_b and ϕ_t stand for production unit, batch, and day fixed effects. We use the estimated coefficients and the counterfactual food per hen to obtain counterfactual output $\tilde{y}_{ibt}$, i.e.

$$\tilde{y}_{ibt} = \hat{\beta}\, \tilde{h}_{ibt} \times \mathbb{I}\left\{h_{ibt} < H\right\} + \hat{\gamma}\, \tilde{h}_{ibt} \times \mathbb{I}\left\{h_{ibt} \geq H\right\} + \hat{\omega}_i + \hat{\lambda}_b + \hat{\phi}_t \tag{5}$$

Upon obtaining counterfactual output, we use information on output prices that the firm made available to us to calculate actual and counterfactual revenues, i.e. $r_{ibt} = py_{ibt}$ and $\tilde{r}_{ibt} = p\tilde{y}_{ibt}$ where p is the price per egg. Similarly, we use the information on food price to calculate actual and counterfactual food costs, i.e. $c_{ibt} = qa_{ibt}$ and $\tilde{c}_{ibt} = q\tilde{a}_{ibt}$ where q is the unit price of food. We also use the actual compensation formula before and after the contract change to calculate actual and counterfactual bonuses paid to employees, equal to $b(y_{ibt}, a_{ibt}) = \alpha y_{it} + (1-\alpha)a_{it}$ and $\tilde{b}(\tilde{y}_{ibt}, \tilde{a}_{ibt}) = \alpha\tilde{y}_{it} + (1-\alpha)\tilde{a}_{it}$ respectively with $\alpha = 1/2$ before the change, and $\alpha = 1$ after the change. Finally, we combine all this information to calculate actual and counterfactual profits $\pi_t = \sum_i \sum_b (r_{ibt} - c_{ibt} - b_{ibt})$ and $\tilde{\pi}_t = \sum_i \sum_b \left(\tilde{r}_{ibt} - \tilde{c}_{ibt} - \tilde{b}_{ibt}\right)$ respectively.

To get a sense of the uncertainty surrounding these estimates, we implement a bootstrap-type procedure sampling with replacement from the full dataset and repeating all steps described above 200 times. Online Appendix Table A.2 shows the results from this exercise for each of the variables we use to calculate profits, with standard deviations in parenthesis.