

DISCUSSION PAPER SERIES

IZA DP No. 13898

**Workplace Incentives and Organizational
Learning**

Francesco Amodio
Miguel A. Martinez-Carrasco

NOVEMBER 2020

DISCUSSION PAPER SERIES

IZA DP No. 13898

Workplace Incentives and Organizational Learning

Francesco Amodio
McGill University and IZA

Miguel A. Martinez-Carrasco
Universidad de Los Andes

NOVEMBER 2020

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Workplace Incentives and Organizational Learning*

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, worker's effort choice changes over time as information is disclosed and processed. We also show that changes in workers compensation can trigger such learning process. We test this hypothesis using personnel records from a Peruvian egg production plant. Exploiting a sudden change in the compensation schedule, 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

Corresponding author:

Francesco Amodio
Department of Economics
McGill University
855 Sherbrooke Street West
Montreal, Quebec H3A 2T7
Canada

E-mail: francesco.amodio@mcgill.ca

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

1 Introduction

The creation and retention of knowledge are key features of organizations. A theoretical and empirical literature exists on *organizational learning*, i.e. the process in which information about products, inputs and technologies is disclosed, exchanged, and processed within organizations (Argote 2013). A separate strand of the literature studies the misalignment of interest within organizations and the provision of incentives. In the presence of moral hazard or adverse selection, the provision of incentives or lack thereof matters for efficiency and worker selection (Prendergast 1999). Empirical studies in this domain typically exploit a change in the salary scheme implemented at a particular firm that occurs at a given point in time and investigate its impact on productivity and profits (see for instance Lazear 2000).

This paper brings together these two literatures in studying learning within organizations when incentives change. When information is not perfect, changing incentives brings uncertainty. The objective function of agents change, and so does their optimal decision. The lack of sufficient information on all variables evaluated at the new equilibrium generates the scope for learning among individuals in organizations.

Our analysis proceeds in three steps. First, we develop the intuition above within a simple conceptual framework. We model a principal-agent relationship 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 a signal to update her beliefs over time. Multiple agents observe each other's effort and output, and learn from each other. Importantly, if learning is local, agents only learn about the shape of the production function around a given level of effort. When the contract changes, the optimal effort choice changes, generating scope for learning at the new equilibrium.

Second, we take this prediction to the data. We use personnel records from a Peruvian egg production plant and exploit a sudden change in workers' compensation schedule 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 choice of 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 a higher output. This happens only after the announcement of the new incentive scheme, around the implementation date, and fades away gradually after 4.5 months.

Third, we attempt to 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 derive counterfactual output, revenues, food costs, and wages. We 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 brings together two separate strands of the economic literature on organizations. The first one studies organizational learning. The empirical work in this domain estimates models of learning across firms ([Argote, Beckman, and Epple 1990](#); [Irwin and Klenow 1994](#); [Benkard 2000](#); [Thornton and Thompson 2001](#)). Other studies investigate social learning among farmers over the profitability and use of new production technologies in developing countries ([Foster and Rosenzweig 1995](#); [Munshi 2004](#); [Conley and Udry 2010](#); [BenYishay and Mobarak 2014](#)). There is less evidence of social learning among workers within firms. Two notable exceptions are [Menzel \(2020\)](#), who finds evidence of knowledge spillovers among workers in Bangladeshi garment factories, and [Chan, Li, and Pierce \(2014\)](#), who study peer learning among salespeople in the cosmetics section of a department store.

The second strand of 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](#)). A number of empirical studies provide convincing evidence 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 ([Guiteras and Jack 2018](#)). The most recent empirical evidence shows that the response to workplace incentives changes with the degree of social connectedness ([Bandiera, Barankay, and Rasul 2010](#)), ethnic diversity ([Hjort 2014](#)), and worker's self-control ([Kaur, Kremer, and Mullainathan 2015](#)).

To the best of our knowledge, ours is the first paper showing that changing incentives can trigger learning within organizations and knowledge spillovers among working peers. 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 and Van Reenen 2007](#); [Bloom et al. 2010](#); [Bloom and Van Reenen 2010](#); [Atkin et al. 2017](#)). We provide tangible estimates of such transaction costs.

The remainder of the paper is organized as follows. Section 2 outlines the conceptual framework. Sections 3 and 4 introduce the empirical setting and data respectively. Section 5 illustrates the empirical strategy to detect organizational learning and reports the first set of results. In Section 6, we estimate the transaction cost associated with imperfect information and learning when incentives change. Section 7 concludes.

2 Conceptual Framework

This section illustrates how workers learn over the shape of the production function, how this affects their effort choice over time, and how changes in the compensation schedule can trigger such learning process.

Each worker i in period t independently produces output y_{it} by combining effort a_{it} with an input of heterogeneous quality s_{it} . Output is given by

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

where $f''_{it}(a_{it}) < 0$ for all a_i , so that output is a concave function of worker's effort. Input quality s_i is identically and independently distributed across workers. Workers do not observe input quality, 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_i and $f(\cdot)$ is responsible for the inability of the worker to disentangle the separate contribution of effort and input quality on output, 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 seeks to motivate the worker to provide the optimal amount of effort using all the available information. We assume that both output and worker's effort are observable by the management, that sets the wage equal to

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

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 compensation schedule matches the one we observe in our empirical application, and we take it as given. Notice however that rewarding the worker in both dimensions 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; Amodio and Martinez-Carrasco 2020).²

¹Allowing the worker to partially observe input quality does not change the implications of the model: the worker will discount that information accordingly, but residual uncertainty over input quality will still generate the scope for learning.

²The use of performance measures capturing workers' effort along a particular dimension is a common feature in many working environments, especially when workers make decisions regarding the use of some inputs. In our setting, this is the amount of food distributed among hens. Amodio and Martinez-Carrasco (2020) also show that using performance measures linked to one specific dimension of the effort can be optimal in a multitasking setting.

We assume that workers are risk neutral and have utility

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

The worker 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

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

which implicitly defines the optimal level of effort a_{it}^* . This changes with the shape of incentives as given by α . Applying the implicit function theorem we get

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

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 $y_{it} = s_{it} f(a_{it}^*)$. As mentioned above, since s_{it} is unknown the worker cannot perfectly disentangle the separate contributions of s_i and $f(a_{it}^*)$ to output. Output is therefore an imprecise signal of the shape of the $f(\cdot)$ function. The worker can use this 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 $f(\cdot)$ at 0 and assume without loss of generality $f(0) = 0$. We have

$$y_{it} \approx s_{it} f'(a_{it}^*) a_{it}^* \quad (6)$$

It follows that, given her choice of effort a_{it}^* at time t , when the worker observes a higher than expected output realization $- s_{it} f'(a_{it}^*) a_{it}^* > \mu_s f'(a_{it}^*) a_{it}^*$ – 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 her beliefs upwards. The opposite holds if the worker observes a lower than expected output realization.

The objective of the worker is to maximize utility. If the effort cost parameter θ is low 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 choice in the next period will increase if output is higher than expected, and decrease otherwise. For a given change in beliefs over $f'(\cdot)$, the magnitude

³Qualitatively similar results can be obtained with risk averse agents. If we assume that agents have a CARA utility function and assuming a normal distribution of s_{it} , 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.

of the change in worker's effort choice will depend on the wage contract parameter α .

Upon observing output signals, the worker updates her beliefs over the shape of the production function. If the effort choice and output of coworkers are observable, she will also use this information in her learning process: workers will learn from each other. Specifically, given worker j 's effort choice a_{jt} , worker i has expectation y_{jt}^i on j 's output that is based on i 's beliefs, i.e. $y_{jt}^i = \mu_s f'_{it}(a_{jt})a_{jt}$. 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 our empirical investigation, we use personnel data from a Peruvian egg production plant (Amodio and Martinez-Carrasco 2018). The plant belongs to a company that produces and sells eggs. Production at the plant occurs in different *sectors*. Each sector comprises different *sheds*, long-building facilities containing one to four different *production units*. The production unit is the basic unit of production at this plant.

Each worker is assigned to a given production unit and assigned a batch of laying hens. All hens within a given batch share very similar characteristics. 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 assigned to the same worker for their entire productive life. Workers exert effort along three main dimensions: egg collection and storage, hen feeding, cleaning and maintenance of the unit facilities.

Output is measured by the number of eggs collected during the day. Mapping from our conceptual framework, this is a function of both hen characteristics or input quality and worker's effort. 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, walk along cages in the production unit and distribute food among all hens. Importantly, 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.⁴

Changing Incentives Workers in the firm 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 part of our sampling period, the bonus payment is 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. If this quantity exceeds a given threshold, a piece rate is awarded for each unit above the threshold. On 24 February 2012, the company adopted a new bonus formula. This is now based on the number of boxes of eggs collected only, with no weight attached to the amount of food distributed by the worker. Such quantity is multiplied by two, and a piece rate is awarded for each unit above a given threshold, with the latter being 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 at the firm claims that workers were distributing "too much food" under the earlier incentive scheme. This speaks to the inability of the management to correctly specify from the beginning the contract that maximizes the payoff of the firm. This is hardly surprising in the context of a large firm operating in a developing country setting (Bloom, Eifert, Mahajan, McKenzie, and Roberts 2013). Nonetheless, as we show in what follows, the implementation of the new salary scheme manages to reduce the amount of food distributed by the workers, in line with the management's expectations and goal.

4 Data and Descriptives

For the purpose of this study, 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 in the incentive scheme. We observe 94 production units in total. Across all of them, we identify 211 different hen batches. We also count 127 workers at work in the sector for at least one day.

⁴Production 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.

Online Appendix Table A.1 shows the summary statistics for the main variables that we use in the empirical analysis. It does so separately for the overall sample and for the three subsamples as defined by the dates in which the change was announced and implemented. Across all periods, workers distribute 23.4 sacks of food a day on average. This quantity varies both across and within workers, with a minimum of 0.5 and a maximum of 39.

The total number of hens per batch is heterogenous across production units over time. This is because batches can have a different size to begin with, but also because hens may die as time goes by. Importantly, when hens within a batch die they are not replaced with new ones: only the whole batch is replaced altogether once the remaining hens reach the end of their productive life. As a result, while we observe around 10,000 hens on average per production unit, their number varies considerably from 353 to more than 15,000. Dividing the total amount of food distributed by the number of hens, we derive the amount of food per hen that is distributed by the worker, averaging 116 grams per day.

Output is given by the number of eggs collected. Workers collect an average more than 8,000 eggs per day. This corresponds to 0.8 daily eggs per hen on average, ranging from 0 to 1. As captured by the model, at least part of this variation is attributable to heterogeneity in input quality. Indeed, the productivity of hens in production varies across units over time. This is partially informed by the innate characteristics of the hens, which also determine how their productivity evolves with age. When purchased, each batch comes with detailed information on the average number of eggs per week each hen is expected to produce at every week of its age. This measure is elaborated by the seller, and is therefore exogenous to anything specific of the plant or the worker who ends up being assigned to that batch. These data are stored by the veterinary unit and are not shared with the human resource department. To get a daily measure of input quality, we divide such expected weekly productivity measure by 7. As shown in Table A.1, the measure we obtain varies from 0.02 to 0.93, with an average of 0.81.

Production units are grouped in different sheds. We count 41 of them in our sample. Using information on the location of each production unit within each shed, we can 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. Finally, 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.

5 Empirical Analysis

5.1 Preliminary Evidence

In the model, we assume that output is a concave function of effort. Figure 1 plots the average number of eggs per hen collected by the worker against the amount of food per hen distributed on the same day. The figure plots the smoothed average 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 upon adopting a structural break identification approach that aims 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 R^2 associated with such kinked specification is equal to 0.98. For comparison, the one associated with a quadratic polynomial specification is 0.58. We interpret this evidence as validating our assumption that the production function is concave. It is also well approximated by a locally linear function.

In this setting, the amount of food distributed measures worker's effort. Without further restrictions, our model delivers ambiguous predictions on whether effort falls when the weight attached to output in the bonus formula increases. Yet, Figure 1 shows that the slope of the production function is always lower than one, equal to 0.02 and 0.002 respectively left and right of the kink. Substituting these values to the corresponding term in equation 5 delivers a clear prediction: effort should decrease when the weight α attached to output in the bonus formula increases.

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. Figure 2 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 salary scheme. The amount of food distributed is stable before the announcement, falls discontinuously on announcement and implementation dates, and then seems to stabilize again in the later period at a level that is lower than the initial one.

Online Appendix Table A.1 and Figure A.1 provide additional evidence of the fall in the amount of food distributed by workers. After the implementation of the new incentive scheme, 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. The Table also show that input quality does not change systematically across periods.

Figure A.1 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 scheme 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.

If all workers were fully informed about the shape of the production function, we would observe effort levels to fall only on the implementation date, and stabilize immediately at the new optimum. Figures 2 and A.1 suggest instead that workers do not hold perfect information over the shape of the production function. The announcement of a new salary structure that puts zero weight on the amount of food distributed leads the workers to decrease the amount of effort they exert along this margin. Our hypothesis is that this triggers a learning process over the exact shape of the production function around the new optimum, which could explain the fall and rise in the average effort level, the temporary increase in variance, and the later stabilization.

5.2 Identification Strategy

Our hypothesis is that changing incentives triggers a learning process among workers over the shape of the production function around the new optimal level of effort. This can explain why workers' choice does not stabilize immediately at the new lower level of food distributed when incentives change, but starts decreasing upon announcement before rising again and stabilizing well after implementation.

The spatial arrangement of production units is such that workers in neighboring units can interact and observe each other. In particular, each worker can guess the productivity of peers by monitoring how often they take the boxes of collected eggs to the warehouse in front of their own production unit. If this is the case, we would expect workers to use the available information on food distribution and output of peers to update their prior over the shape of the production function, and inform their own food choice accordingly. This would generate a positive correlation between the choices of neighboring coworkers.

Yet, finding evidence of a positive correlation between the input choices of neighboring coworkers is not necessarily informative of whether workers learn from each other. First, in the absence of social learning, unobserved common factors may independently affect the effort decision of coworkers and tilt it 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 first identified by [Manski \(1993\)](#).

To overcome these issues, we adopt a regression framework that maps from the predictions of our conceptual framework. Our approach builds upon [Conley and Udry \(2010\)](#) and their study

of pineapple growers in Ghana. Rather than looking at the correlation between the choices of neighboring peers, we look at changes in their effort choices over time, and whether they adjust towards their peer choice differentially when the latter achieve higher output. To operationalize this approach, we first define for each worker i operating batch b on day t a variable

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

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 implement the following baseline regression specification

$$\Delta a_{ibt} = \beta M_{ibt} + \gamma Post_t \times M_{ibt} + \mathbf{X}'_{ibt} \kappa + u_{ibt} \quad (8)$$

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 $Post_t$ is a dummy equal to one in the period after the announcement of the new incentive scheme. 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 incentive scheme. The vector \mathbf{X}_{ibt} includes the lagged own and coworkers' input choice and output. It also includes the total number of hens in the batch, an important determinant of the amount of food distributed. Finally, the term u_{ibt} captures any residual determinants of change in the amount of food distributed by the worker. We cluster standard errors along the two dimensions of shed and day to account for unobserved correlation between the residual determinants of food choice across observations belonging to the same shed or day.

5.3 Results

Table 1 shows the corresponding coefficient estimates. Column 1 shows the estimate of 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 include the full set of day fixed effects as regressors. 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. Consistent with our hypothesis, 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 after including progressively worker fixed effects in column 4, and a dummy if the estimated input quality is above the median in column 5. We interpret this as evidence that the announcement of the new incentive scheme

triggers learning among coworkers.

In column 6, we evaluate the robustness of results by excluding those observations belonging to production units and days where the assigned worker was absent, and thus replaced by another one. The estimated coefficients of interest and their significance is not affected. In column 7 and 8, we implement the same regression specification separately for the subsample of production units assigned to workers with lower than average and higher than average tenure respectively. The estimated γ is now significant only for workers with high tenure. 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 by changing their level of effort in the next period.

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

The previous specification pools together all observations, and uses a single identifier for those belonging to the post-announcement period. Yet, we would expect learning to occur only for a limited amount of time: as information is disclosed and processed by workers, their food choice should become more consistent over time and no longer respond to coworkers' effort and output. To test this hypothesis, we augment the regression specification in 8 with a whole set of interactions of M_{ibt} with dummies that identify each two-week pay period. We omit and use as reference the pay period when the change in incentives 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 salary scheme. Estimates are not significantly different from zero for the whole period before the announcement of the new scheme. They become significant at the 5% level shortly before implementation, remain significant for several periods, then return insignificant. When comparing Figure 3 with Figure 2, we can see that the two align remarkably well, with the estimates capturing social learning being insignificant when the food choice is stable, and significant over the adjustment period. We interpret this evidence altogether as showing that the announcement and implementation of the new salary scheme triggers learning among workers over the shape of the production function.

6 The Cost of Learning

In this section, we summarize our attempt to quantify the profit losses associated with the change in incentives 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. In other words, we cannot disentangle the variation in the variables of interest that is driven by learning and experimentation from the one determined by idiosyncratic shocks.

We address this challenge as follows. In the first step, we regress the amount of food per hen distributed by the worker over the full set of day, production unit, and batch fixed effects. We then split the sample in three periods: the one before the announcement of the new incentive scheme, 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 scheme 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 average (equilibrium) level of feeding effort under the old and new incentive scheme 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 scheme respectively. 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.

The top graph in [Figure 4](#) shows the smoothed average of the actual and counterfactual food choices over the sampling period. The two vertical lines define the three periods indicated above. Not surprisingly, given the way we obtain the counterfactual, actual and counterfactual variables coincide in the first and third period. The two are different 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 5](#), 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 change in incentives to calculate the bonuses paid to employees. Finally, we combine all this information to calculate profits. The bottom graph in Figure 4 shows the corresponding results. The area between the two lines measures the daily average profit loss over time.

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.⁵ 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.⁶

7 Conclusions

This paper shows that changing incentives triggers learning among coworkers and knowledge spillovers within firms. We first present a principal-agent framework that illustrates how imperfect information over the true shape of the production function affects worker's effort choice over time. A change in the parameters of the compensation schedule 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 knowledge spillovers. The learning process lasts around 4.5 months, and brings about estimated profit losses of USD 340 to 400K. Our study reveals how imperfect information over the global shape of the production function increases transaction costs for firms associated with changing the shape of incentives at the workplace. It follows that varying degrees of information completeness can in principle explain at least part of the variation we observe across firms within and across countries in the adoption of personnel management practices and incentive schemes. Understanding whether and how this is the case is an open question that we leave for future research.

⁵Online Appendix Figures A.2 and A.3 show the smoothed averages of all actual and counterfactual variables used to calculate profits.

⁶Online Appendix Figure A.4 shows the distribution of absolute and relative profit gains across the 200 repetitions.

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., I. Barankay, and I. Rasul (2010). Social incentives in the workplace. *The Review of Economic Studies* 77(2), 417–458.
- 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 (2014). Social learning and communication. Working Paper 20139, National Bureau of Economic Research.
- 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.
- Hjort, J. (2014). Ethnic divisions and production in firms. *The Quarterly Journal of Economics* 129(4), 1899–1946.
- 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.
- Kaur, S., M. Kremer, and S. Mullainathan (2015). Self-control at work. *Journal of Political Economy* 123(6), 1227–1277.
- 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.

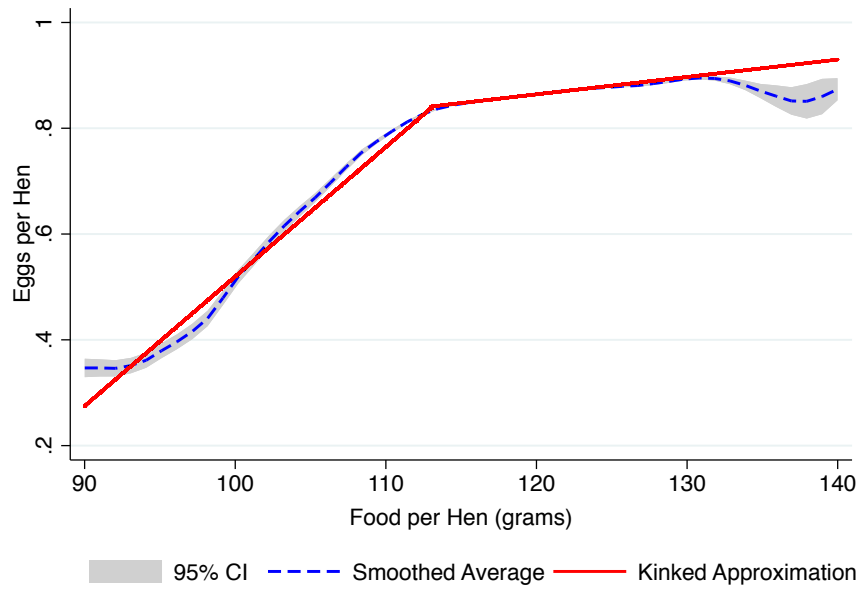
Tables and Figures

Table 1: Incentive Change and Organizational Learning

	Change in Food Distributed							
	(1)	(2)	(3)	(4)	(5)	No Absentees (6)	Low Tenure (7)	High Tenure (8)
M_{ibt}	0.0029*** (0.0009)	0.0046*** (0.0011)	0.0008 (0.0015)	0.0004 (0.0016)	0.0008 (0.0016)	0.0008 (0.0015)	0.0004 (0.0017)	0.0021 (0.0028)
$Post_t \times M_{ibt}$			0.0055** (0.0021)	0.0065*** (0.0022)	0.0064*** (0.0022)	0.0063*** (0.0021)	0.0038 (0.0032)	0.0068** (0.0032)
<i>High Input Quality</i>					-0.0721*** (0.0126)	-0.0691*** (0.0123)	-0.0669*** (0.0178)	-0.0835*** (0.0197)
<i>No. of Hens</i>	0.0003*** (0.0000)	0.0004*** (0.0000)	0.0004*** (0.0000)	0.0004*** (0.0000)	0.0004*** (0.0000)	0.0004*** (0.0000)	0.0003*** (0.0001)	0.0005*** (0.0001)
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	39508	13539	16108
R^2	0.0937	0.1378	0.1388	0.1529	0.1553	0.1477	0.1389	0.2055

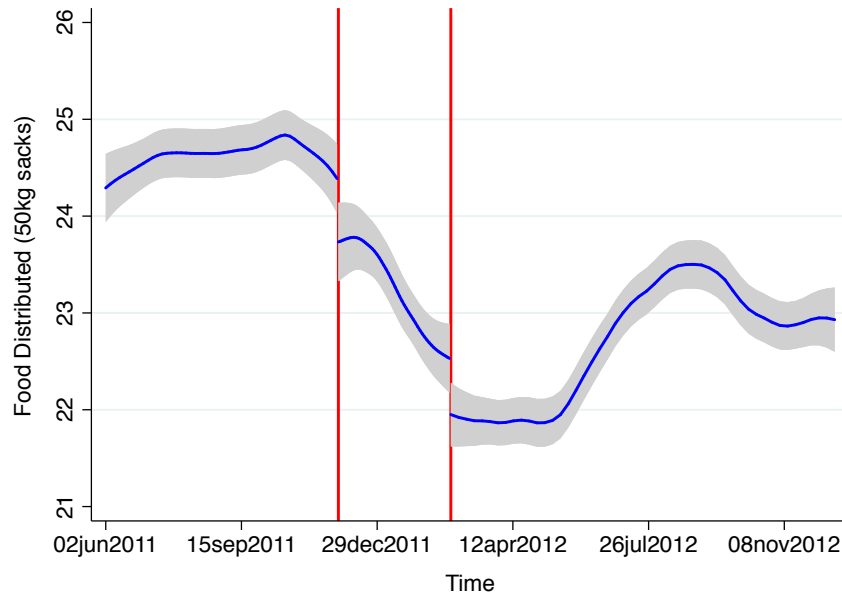
Notes. (* p-value<0.1; ** p-value<0.05; *** p-value<0.01) Two-way clustered standard errors, with residuals grouped along both shed and day. Dependent variable is the change in the amount of food distributed by the worker from the day before to the day of observation, as measured by the change in the 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), $\mathbb{1}\{y_{jbt-1} > y_{ibt-1}\}$. $Post_t$ is a dummy equal to one for all observations belonging to the period after the announcement of the change in incentives. 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), the sample is restricted to those production units and days with no absentee workers. In columns (7) and (8) the sample is split between workers with lower and higher than average tenure respectively, and restricted to those observations that we can merge with the survey of workers that we administered in March 2013.

Figure 1: Output and Feeding Effort



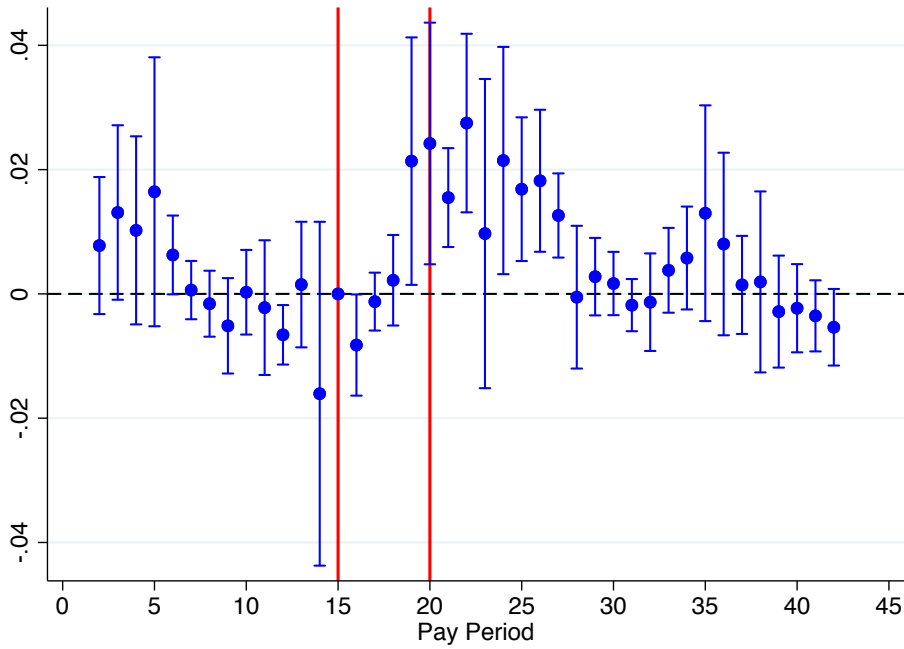
Notes. The figure plots the smoothed average of the number of eggs per hen collected by the worker over the grams of food per hen distributed in the day, together with its 95% confidence interval. It also plots a kinked linear approximation of the production function. The values of amount of food at 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 2: Food Choice Over Time



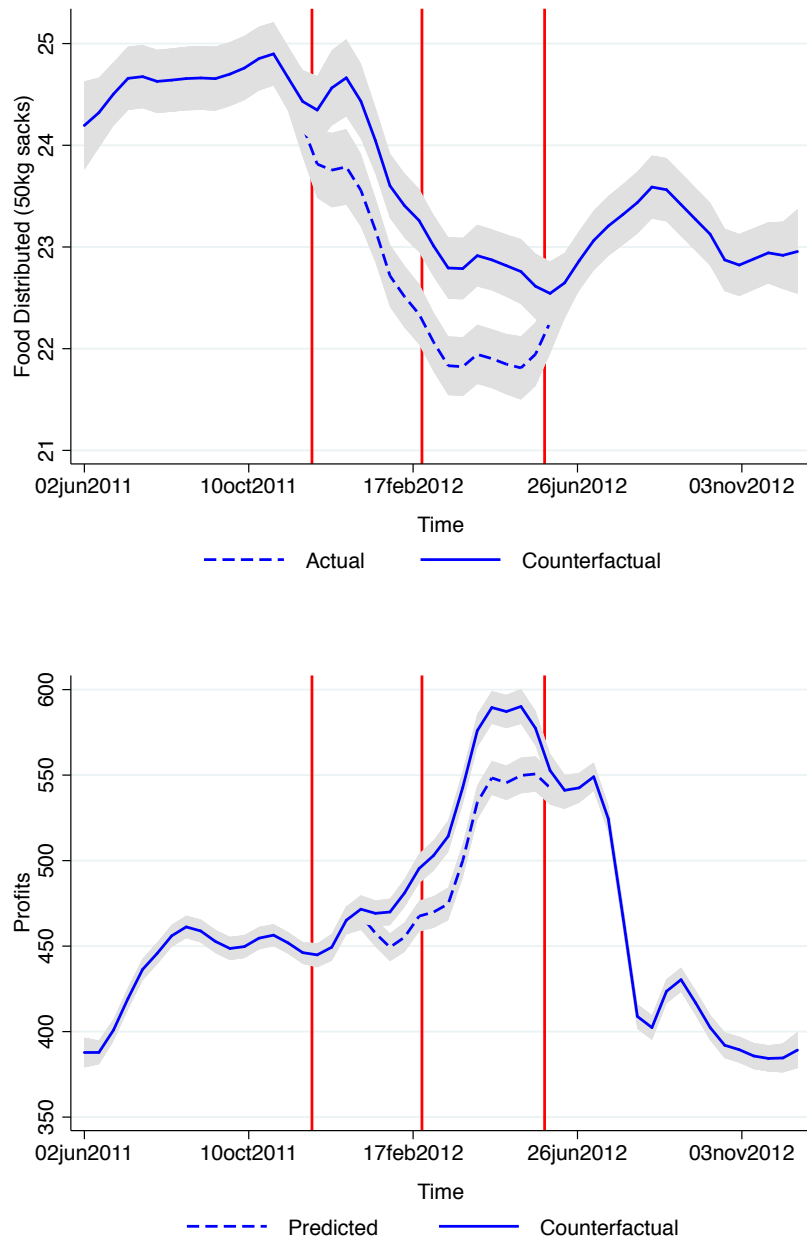
Notes. The figure plots the smoothed average of the total number of 50kgs sacks of food distributed across all production units in a given day, together with its 95% confidence interval. The two vertical lines correspond to the dates of announcement and implementation of the new incentive scheme. The 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 3: Incentive Change and Organizational Learning



Notes. The figure plots the coefficient estimates associated with the whole set of interactions between the M_{ibt} variable specified in Section 5.2 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 incentive scheme. 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 2.

Figure 4: Actual and Counterfactual Food Choice and Profits



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 and counterfactual amount of profits per day. The procedure to construct these counterfactuals is described in Section 6 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
<i>Panel A – All Sample</i>					
Food Distributed (50kg sacks)	23.402	8.705	0.5	39	46049
No. of Hens	10105.576	3672.284	353	15985	46049
Food per Hen (gr)	115.771	9.495	66.774	163.235	46049
Total Eggs Collected	8140.025	3574.481	0	15131	46049
Total Eggs per Hen	0.803	0.19	0	1	46049
Input Quality	0.811	0.147	0.02	0.934	44985
Food Distributed by Coworkers (avg)	24.255	7.302	1	35.5	42281
Experienced	0.498	0.500	0	1	32892
High Schooling	0.537	0.499	0	1	32892
<i>Panel B – Before Announcement</i>					
Food Distributed (50kg sacks)	24.606	8.945	3	39	15358
No. of Hens	10198.691	3654.062	1311	15399	15358
Food per Hen (gr)	120.704	7.079	67.146	163.235	15358
Total Eggs Collected	8550.046	3528.786	0	13830	15358
Total Eggs per Hen	0.84	0.163	0	0.998	15358
Input Quality	0.809	0.155	0.02	0.934	15222
<i>Panel C – Between Announcement and Implementation</i>					
Food Distributed (50kg sacks)	23.188	8.527	1	35	5938
No. of Hens	10025.405	3605.539	406	14963	5938
Food per Hen (gr)	115.619	9.377	68.673	159.795	5938
Total Eggs Collected	8062.368	3395.261	0	13112	5938
Total Eggs per Hen	0.805	0.169	0	1	5938
Input Quality	0.798	0.133	0.02	0.934	5667
<i>Panel D – After Implementation</i>					
Food Distributed (50kg sacks)	22.705	8.513	0.5	35	24753
No. of Hens	10067.036	3698.337	353	15985	24753
Food per Hen (gr)	112.747	9.554	66.774	163.132	24753
Total Eggs Collected	7904.256	3621.822	0	15131	24753
Total Eggs per Hen	0.78	0.206	0	1	24753
Input Quality	0.815	0.145	0.02	0.934	24096

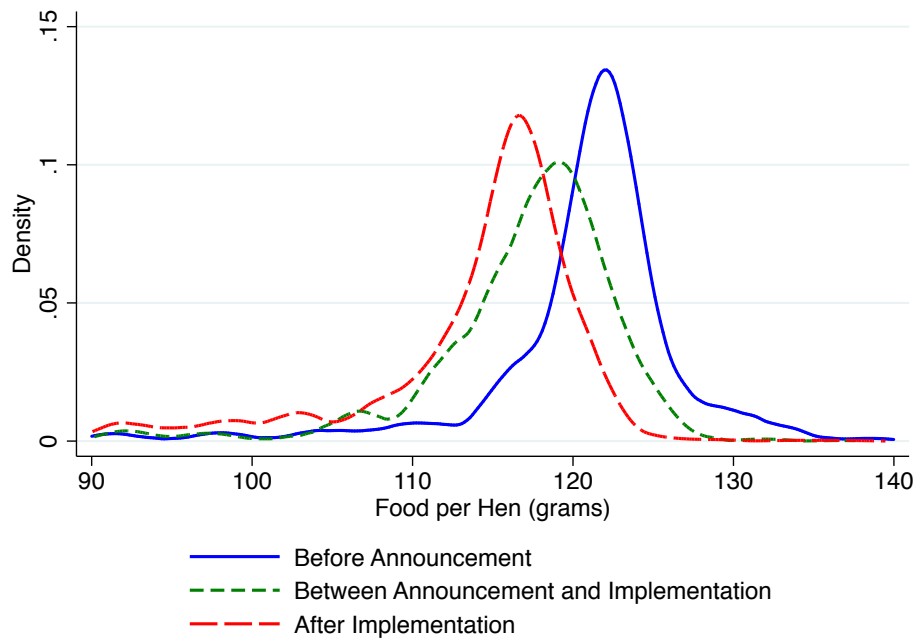
Notes. The table reports the summary statistics of the variable used in the empirical analysis in the overall sample and separately for the period before, during, and after the change in incentives.

Table A.2: Counterfactual Estimates

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

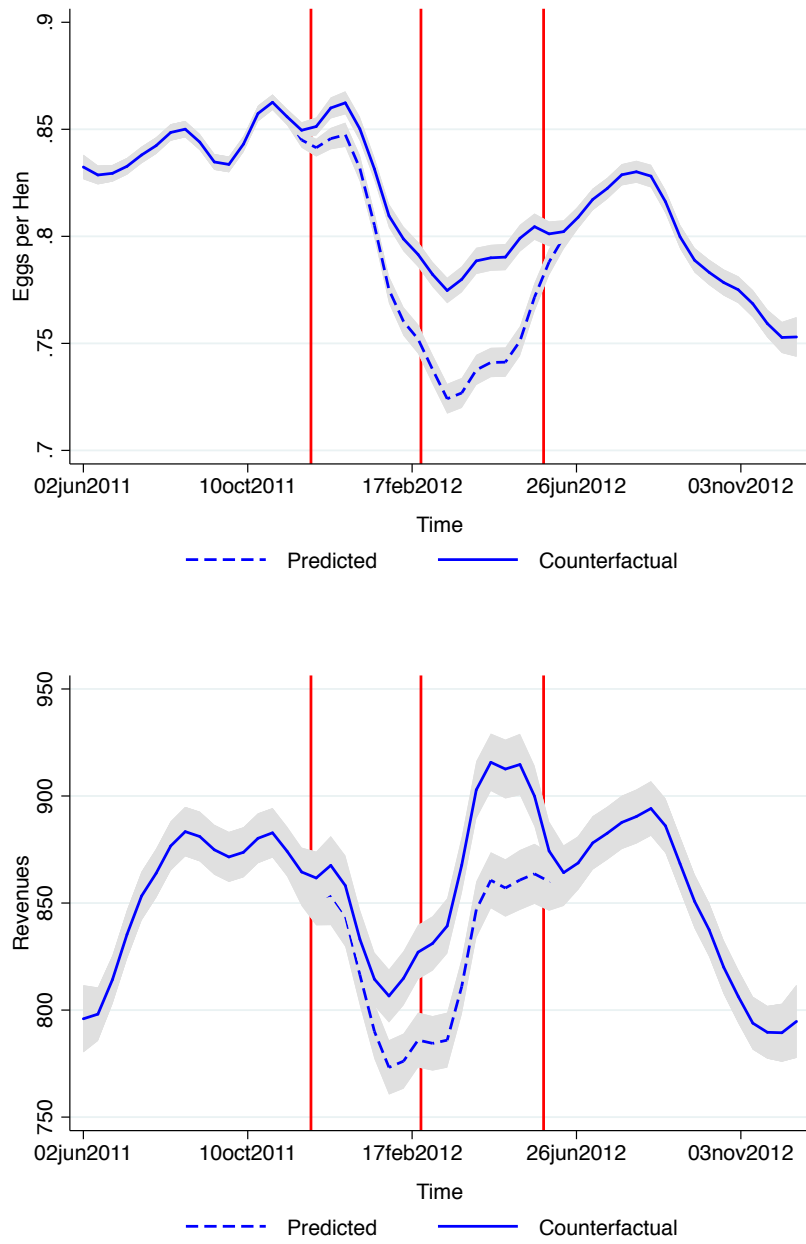
Notes. The table shows the average and standard deviation of predicted and counterfactual variables. Both are estimated with the procedure described in Section 6 and Online Appendix A.2. Distributions are obtained by implementing a bootstrap-type procedure of resampling with replacement in 200 repetitions.

Figure A.1: Distribution of Food Choice Over Time



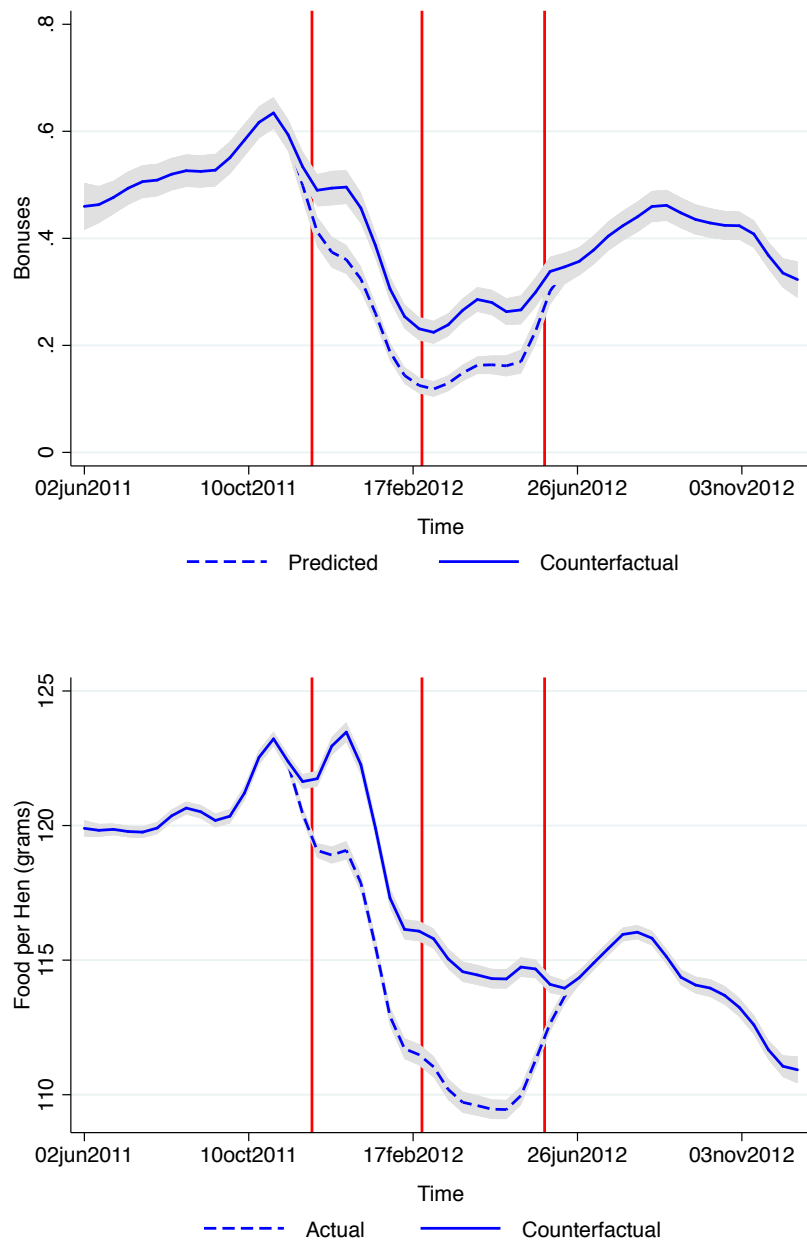
Notes. The figure plots the smoothed kernel density of grams of food per hen distributed in each day across workers and separately in the period before, during, and after the implementation of the change in incentives.

Figure A.2: Actual and Counterfactual Output and Revenues



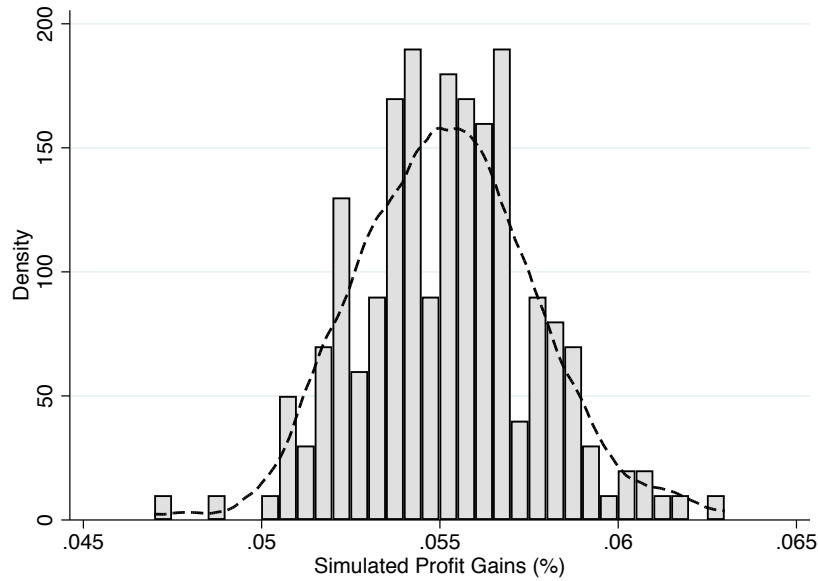
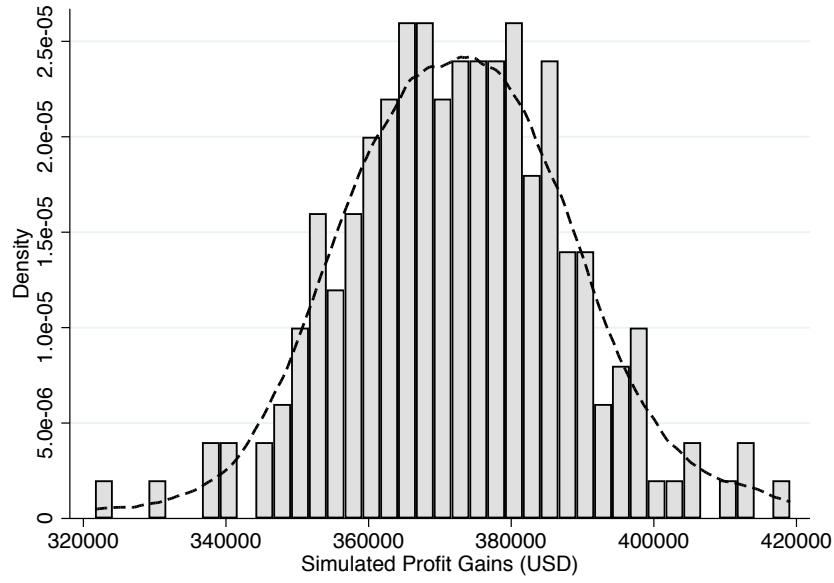
Notes. The top figure shows the predicted smoothed average of the total number of eggs collected, and its counterfactual in a simulated environment with no learning. The bottom figure shows the predicted and counterfactual amount of revenues per day. The procedure to construct these counterfactuals is described in Section 6 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 Wages and Food per Hen



Notes. The top figure shows the predicted smoothed average of bonuses paid, and its counterfactual in a simulated environment with no learning. The bottom figure shows the actual and counterfactual amount of food per hen per day. The procedure to construct these counterfactuals is described in Section 6 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 6 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 6.

In the first step, we implement the following regression specification

$$h_{ibt} = \theta_i + \psi_b + \delta_t + \varepsilon_{ibt} \quad (1)$$

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 incentive scheme, 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 scheme, 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 incentive scheme, $\bar{\hat{\varepsilon}}_D$ 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

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

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

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

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 5, Figure 1 and the specification in equation 1 above, we implement the following kinked regression

specification

$$y_{ibt} = \beta h_{ibt} \times \mathbb{I}\{h_{ibt} < H\} + \gamma h_{ibt} \times \mathbb{I}\{h_{ibt} \geq H\} + \omega_i + \lambda_b + \phi_t + \epsilon_{ibt} \quad (4)$$

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.

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

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 change in incentives 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 $\tilde{\pi}_t = \sum_i \sum_b (\tilde{r}_{ibt} - \tilde{c}_{ibt} - \tilde{b}_{ibt})$ 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.