# Workplace Incentives and Organizational Learning*

Francesco Amodio† Miguel A. Martinez-Carrasco‡

March 12, 2019

**Abstract**

This paper studies organizational learning when incentives change. We first illustrate how imperfect information over the true shape of the production function affects worker's effort choice over time as information is disclosed and processed. We then show that changes in the compensation schedule can trigger such learning process. We take this hypothesis to the data using personnel records from a Peruvian egg production plant. Exploiting a sudden change in the worker salary structure, we find evidence that workers learn from each other over the shape of the production function, and change their effort accordingly. This adjustment process is costly for the firm. Our study shows that lack of information over the global shape of the production function increases the cost for firms associated with changing the shape of incentives at the workplace.

**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, 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. A large theoretical and empirical literature exists on "organizational learning," the process in which information about products, inputs and technologies is disclosed, exchanged, and processed within organizations (Argote 2013). A separate literature studies the misalignment of goals within organizations and the provision of incentives. The shape of incentives or their absence determine the extent of moral hazard and adverse selection issues, and thus matter for efficiency and worker selection (Prendergast 1999). Empirical studies in this domain typically exploit changes in the salary structure of a given firm that occur at a given point in time (see for instance Lazear 2000).

This paper studies organizational learning when incentives change. When information is not perfect, changing incentives can bring 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 optimum can trigger learning among individuals in organizations.

First, we develop this intuition 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 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 optimum.

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. At this company, workers are assigned batches of hens, exert effort to feed them, and collect eggs as output (Amodio and Martinez-Carrasco 2018). 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. 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 for a period of 4.5 months.

Third, we attempt to quantify the profit losses associated with the incentive change and due to imperfect information and 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 this counterfactual scenario.

Our paper contributes to two strands of the literature. 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 exceptions are Menzel (2017), 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). Existing studies show that 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 knowledge spillovers among peers at the workplace. Imperfect information over the global shape of the production function can increase the transaction costs associated with the implementation of new incentive schemes and management practices in general (Bloom and Van Reenen 2007; Bloom et al. 2010; Bloom and Van Reenen 2010). We provide empirical evidence that this is the case, and provide an estimate of the transaction costs associated with imperfect information and learning.

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

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

where $\partial^2 f(\cdot)/\partial a_i^2 \leq 0$ 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 $\mu_s$ and variance $\sigma_s^2$.[1]

The exact shape of $f(\cdot)$ is unknown to the worker, who holds in each period beliefs $f_{it}(\cdot)$ over $f(\cdot)$. This determines the worker's effort choice. The worker is risk-neutral, and her cost of effort is given by $C(a_{it}) = \theta a_{it}^2/2$, with $0 < \theta < 1$. Both output and worker's effort are observable by the management, that sets the wage equal to

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

where $\kappa$ is fixed and $\alpha$ 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.[2]

Worker's utility is given by

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

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

$$\theta a_{it}^* = \alpha \mu_s f_{it}'(a_{it}) + (1-\alpha) \tag{4}$$

We assume that $f(\cdot)$ is locally linear. This means that the marginal product of effort is constant within intervals of $a_i$, but changes discontinuously across such intervals. This also implies that the previous equation defines the optimal level of effort $a_{it}^*(\alpha, f_{it}')$ in closed form.

[1] Allowing 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 generates scope for learning.

[2] 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).

The optimal level of effort changes with the shape of incentives as given by $\alpha$. Applying the implicit function theorem we get

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

so that the level of effort may increase or decrease with $\alpha$ depending on whether its expected marginal product is higher or lower than one.

Upon exerting effort, the worker observes the corresponding output realization

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

Output is an imprecise signal of the shape of the production function. Since $s_{it}$ is unknown, the worker cannot perfectly disentangle the separate contributions of $s_i$ and $f(a^*_{it})$ to output. Yet, the worker can use this output signal to update her beliefs over the marginal product of effort. 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

$$y_{it} \approx s_{it} f'(a^*_{it}) a^*_{it} \tag{7}$$

It follows that, given her choice of effort $a^*_{it}$ at time $t$, when the workers observe a higher than expected output realization – $s_{it} f'(a^*_{it}) a^*_{it} > \mu_s f'_{it}(a^*_{it}) a^*_{it}$ – there is a positive probability that the true marginal product of effort $f'(\cdot)$ 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 $\theta$ 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. Importantly, for a given change in beliefs over $f'(\cdot)$, the magnitude of the change in worker's effort choice will depend on the contract structure and its parameter $\alpha$.

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 $y^i_{jt}$ on $j$'s output that is based on $i$'s beliefs, i.e. $y^i_{jt} = \mu_s f'_{it}(a_{jt}) a_{jt}$. 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.

The assumption that $f(\cdot)$ is locally linear also implies that learning over the marginal product $f'(\cdot)$ at one level of effort $a_{it}$ is not informative of the marginal product at another level of effort. A change in $\alpha$ changes the optimal choice of effort, and therefore triggers learning over a different segment of the production function. This is the hypothesis that we subject to empirical scrutiny.

# 3 The Setting

In our empirical application, 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, and 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 observable 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, walking within the production unit along cages and distributing it 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.[3]

**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. Importantly, the formula to calculate the bonus has 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

[3] 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.

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 refers to the workers distributing "too much food" under the earlier incentive scheme. This speaks to the inability of the management to correctly specify 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 our empirical analysis, 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.

Table 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 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 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 merge this information with those for workers that were still present on the day of the survey, which amounts to 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 assumed that output is a concave function of effort, and that the production function is locally linear. 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 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. Evidence shows that even 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 prediction is equal to 0.98. For comparison, the one associated with a quadratic approximation is 0.58. We interpret these results as validating our assumption that the production function is concave and locally linear in effort.

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 $\alpha$ 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 $\alpha$ 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 our sampling period. The graph shows the smoothed average of such residual 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.

Table 1 and Figure 3 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. Importantly, the Table also shows that input quality does not change systematically across periods. Figure 3 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 3 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. That 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. However, testing for the presence of such positive correlation is not necessarily informative of whether workers learn from each other. First, in the absence of social learning, unobserved third factors may independently affect their effort decision 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 matches 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 achieves higher output. To operationalize this approach, we first define for each worker $i$ operating batch $b$ on day $t$ a variable

$$M_{ibt} = (a_{jbt-1} - a_{ibt-1}) \times \mathbb{I}\{y_{jbt-1} > y_{ibt-1}\} \tag{8}$$

where $a_{jbt-1}$ is the effort (food) choice 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 success, equal to one if the output (eggs per hen) of neighboring coworkers was higher than own output.

We implement the following baseline regression specification

$$\Delta a_{ibt} = \beta\, M_{ibt} + \gamma\, \mathit{Post}_t \times M_{ibt} + \mathbf{X}'_{ibt}\kappa + \nu_i \tag{9}$$

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 $\mathit{Post}_t$ is a dummy equal to one in the period after the announcement of the new incentive scheme. The coefficient $\beta$ 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. $\gamma$ captures whether this is differentially the case 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 obvious source of variation for the amount of food distributed. Finally, the term $\nu_i$ 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.

## 5.3 Results

Table 2 shows the corresponding coefficient estimates. Column 1 shows the estimate of $\beta$ 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 $\beta$ 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 $\beta$ is now close to zero and insignificant while the estimate of $\gamma$ 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 $\gamma$ 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 $\Delta a_{ibt}$ 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 will become more stable and will no longer respond to coworkers' effort and output. To test this hypothesis, we augment the regression specification in 9 with the 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 4 plots the coefficient estimates of all these interaction terms, 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 4 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 attempt to quantify the profit losses associated with the change in incentives, due to imperfect information and learning. 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 observed choice of food distributed per hen 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 4 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 the second period. In the second step, we derive the average of residual food distributed per hen in the three periods. We consider the averages in the first and last period as informative of the average optimal level of feeding effort under the old and new incentive scheme respectively. We thus define the counterfactual feeding effort choice by re-centering the distribution of residuals in the second period 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. In the third step, 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 5 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.

With such counterfactual measure in hand, we can derive output, revenues, food costs, and bonuses paid to the workers. We derive output by running a regression of eggs per hen over the kinked function of food per hen estimated 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. 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 5 shows the corresponding results. The area between the two lines measures the average profit loss.

To get a sense of the precision of the estimates, we implement a bootstrap procedure. We sample with replacement from the full dataset and repeat the whole procedure described above 200 times. Table 6 shows the results from this exercise, with standard deviations in parenthesis.[4] Figure A.3 in Appendix A.1 shows the distribution of absolute and relative profit gains across the 200 repetitions. The average profit loss is estimated to be equal to USD 373K. According to our exercise, profits would have been 5.5% higher over the learning period in the presence of complete information.

# 7 Conclusions

This paper studies organizational learning when incentives change. 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 such learning process. We take this hypothesis to the data using personnel records from a Peruvian egg production plant. We show that workers change their

[4] Figures A.1 and A.2 in Appendix A.1 show the smoothed averages of all actual and counterfactual variables used to calculate profits.

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.

# 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.

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. (2017). 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: Summary Statistics

| Variable | Mean | Std. Dev. | Min. | Max. | N |
|---|---|---|---|---|---|
| | *Panel A – All Sample* | | | | |
| 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 |
| | *Panel B – Before Annoucement* | | | | |
| 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 |
| | *Panel C – Between Announcement and Implementation* | | | | |
| 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 |
| | *Panel D – After Implementation* | | | | |
| 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 2: Incentive Change and Organizational Learning

| | Change in Food Distributed | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | | | | | | No Absentees | 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) | (0.0015) | (0.0017) | (0.0028) |
| $Post_t \times M_{ibt}$ | | | 0.0055** | 0.0065*** | 0.0064*** | 0.0063*** | 0.0038 | 0.0068** |
| | | | (0.0021) | (0.0022) | (0.0022) | (0.0021) | (0.0032) | (0.0032) |
| *High Input Quality* | | | | | -0.0721*** | -0.0691*** | -0.0669*** | -0.0835*** |
| | | | | | (0.0126) | (0.0123) | (0.0178) | (0.0197) |
| *No. of Hens* | 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) |
| 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 current day, 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 in the amount of food distributed between the worker and her neighbor on the previous day, $a_{jbt-1} - a_{ibt-1}$, with a dummy equal to one if the neighbor achieved a higher output, $\mathbb{I}\left\{y_{jbt-1} > y_{ibt-1}\right\}$, where output is measured as eggs per hen. *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 her neighbor 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.

Table 3: Estimated Cost of Learning

| | 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 6. Distributions are obtained after a bootstrap procedure of resampling with replacement in 200 repetitions.

Figure 1: Output and Feeding Effort

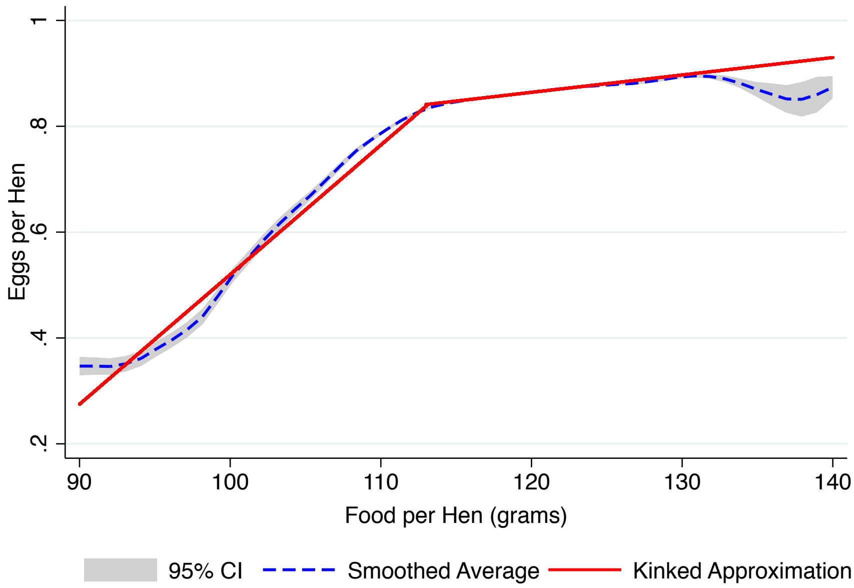


*Notes.* The figure plots the smoothed average of the number of eggs per hen collected by the worker over the grams of food per hen distributed in the day, together with its 95% confidence interval. It also plots a kinked linear approximation of the production function. The values of amount of food at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

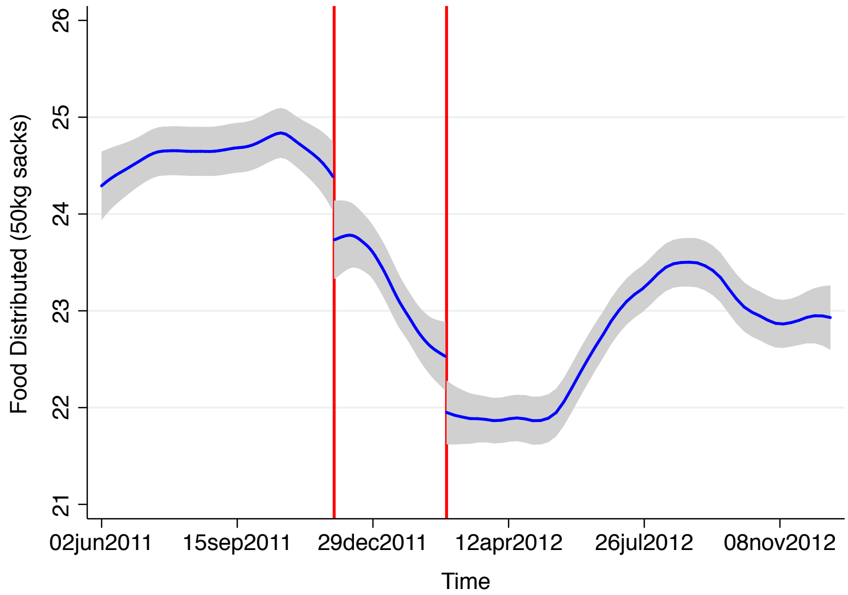


*Notes.* The figure plots the smoothed average of the total number of 50kgs sacks of food distributed across all production units in a given day, together with its 95% confidence interval. The two vertical lines correspond to the dates of announcement and implementation of the new incentive scheme. The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

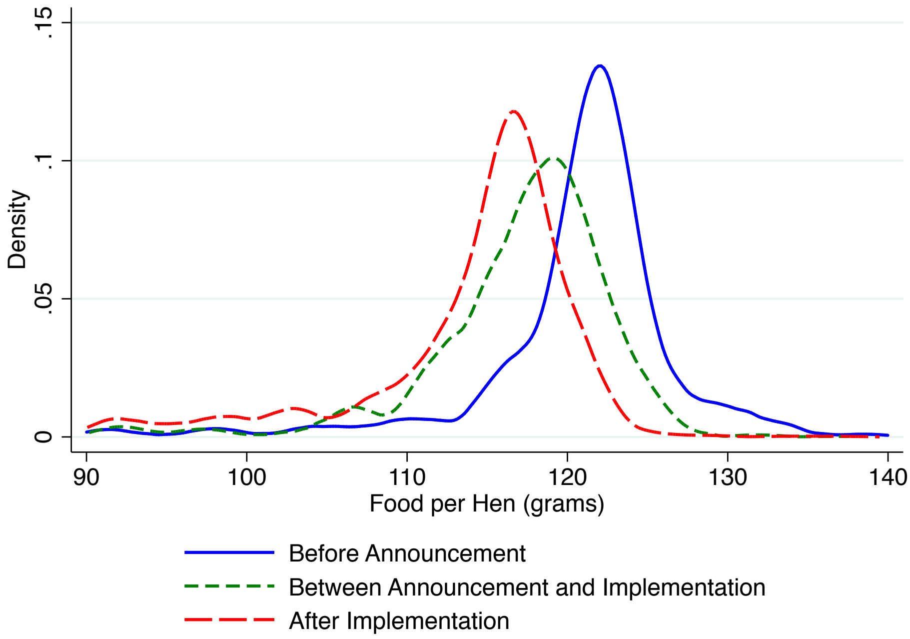


*Notes.* The figure plots the smoothed kernel density of grams of food per hen distributed in each day across workers and separately in the period before, during, and after the implementation of the change in incentives.

Figure 4: Incentive Change and Organizational Learning

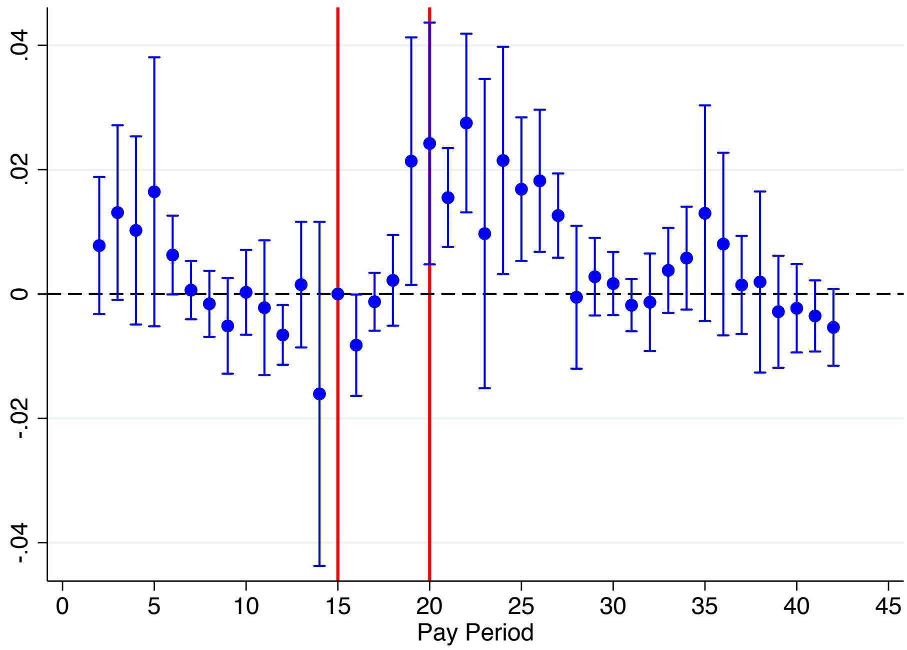


*Notes.* The figure plots the coefficient of the interaction between the $M_{ibt}$ variable specified in Section 5.2 and a dummy for each two-week pay period as estimated from an augmented version of regression specification in equation 9 that includes these interactions. The two vertical lines correspond 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 organizational learning increases after the announcement and becomes positive and significant around and after the implementation date, consistent with Figure 2.

Figure 5: Counterfactual Food Choice and Profits

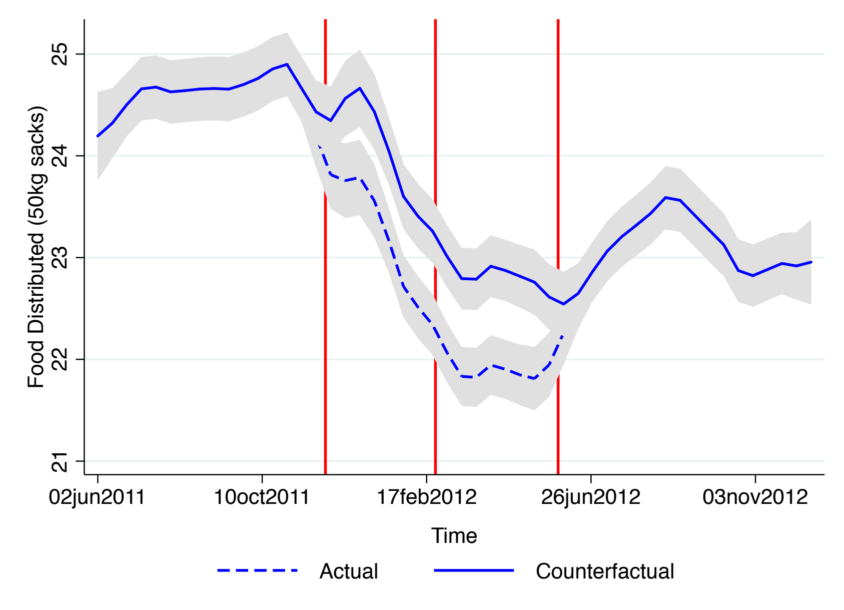


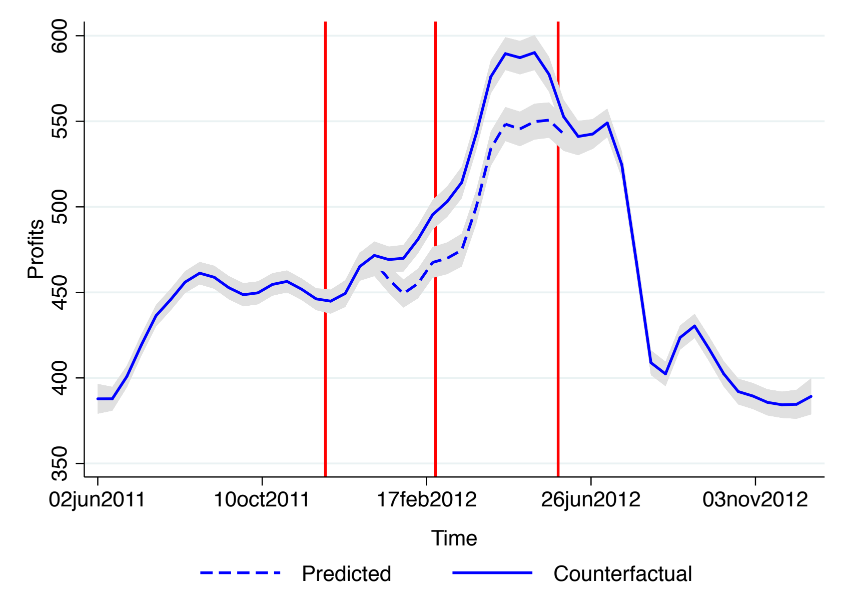


*Notes.* The top figure shows the actual smoothed average of the amount of food distributed by workers, and its counterfactual in a simulated environment with no learning. The bottom figure shows the predicted and counterfactual amount of profits per day. The procedure to construct these counterfactuals is described in Section 6. 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 4.

# A Appendix

## A.1 Additional Tables and Figures

Figure A.1: Counterfactual Output and Revenues

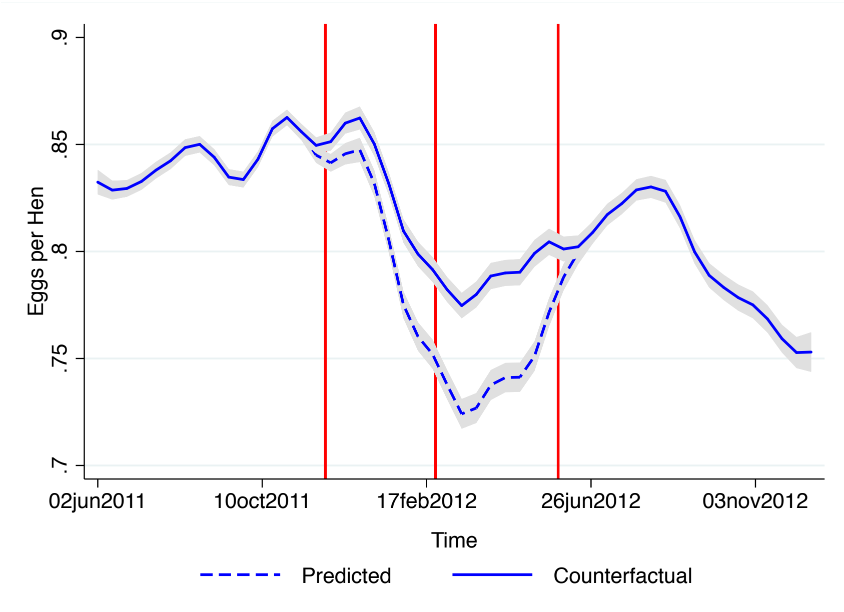


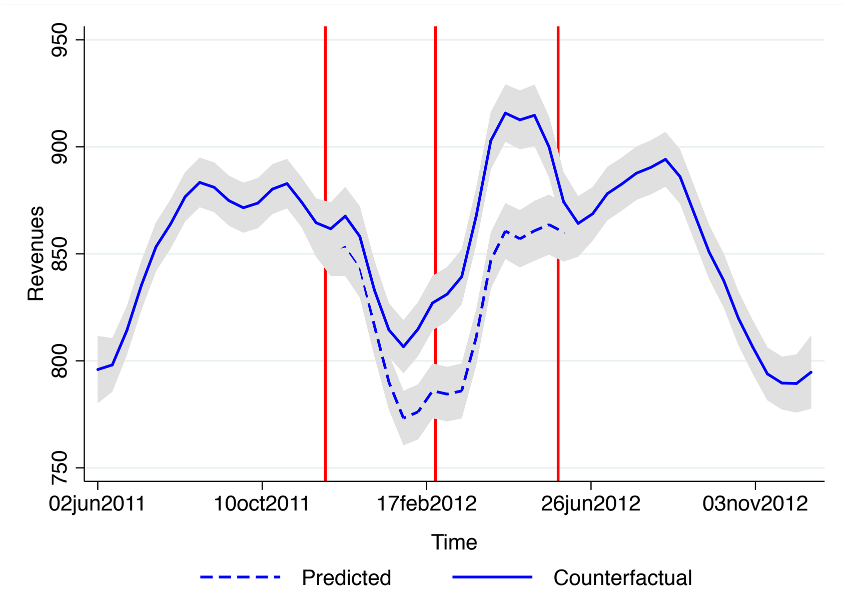


*Notes.* The top figure shows the predicted smoothed average of the total number of eggs collected, and its counterfactual in a simulated environment with no learning. The bottom figure shows the predicted and counterfactual amount of revenues per day. The procedure to construct these counterfactuals is described in Section 6. 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 4.

Figure A.2: Counterfactual Wages and Food per Hen

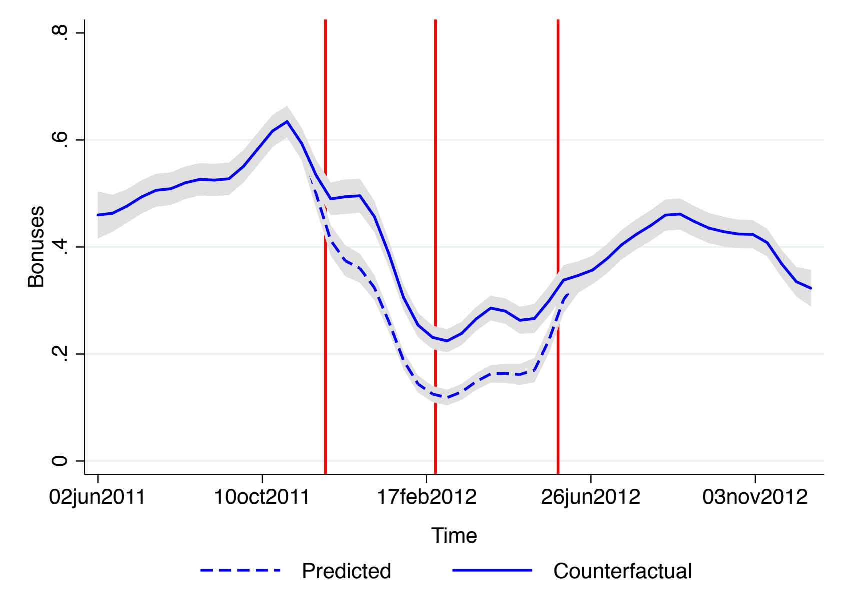


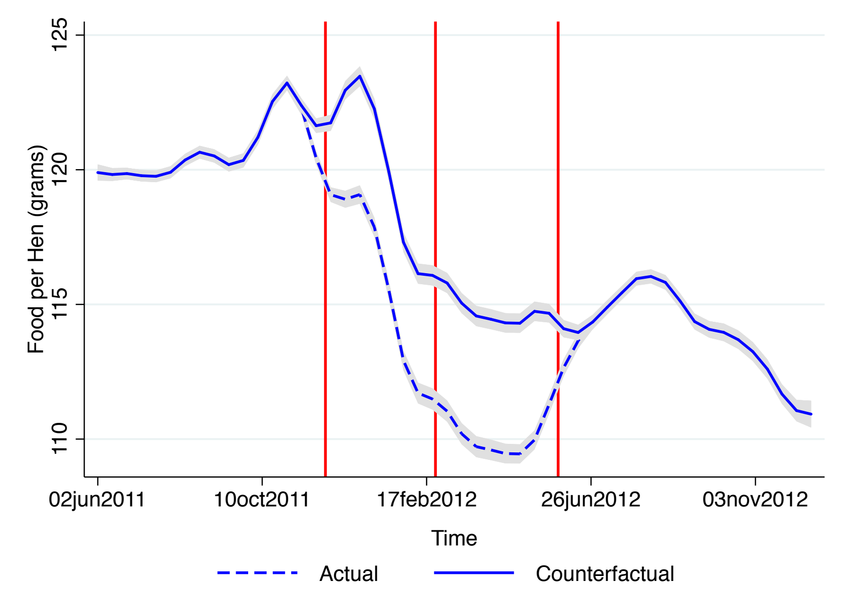


*Notes.* The top figure shows the predicted smoothed average of bonuses paid, and its counterfactual in a simulated environment with no learning. The bottom figure shows the actual and counterfactual amount of food per hen per day. The procedure to construct these counterfactuals is described in Section 6. 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 4.

Figure A.3: Distribution of Profit Gains

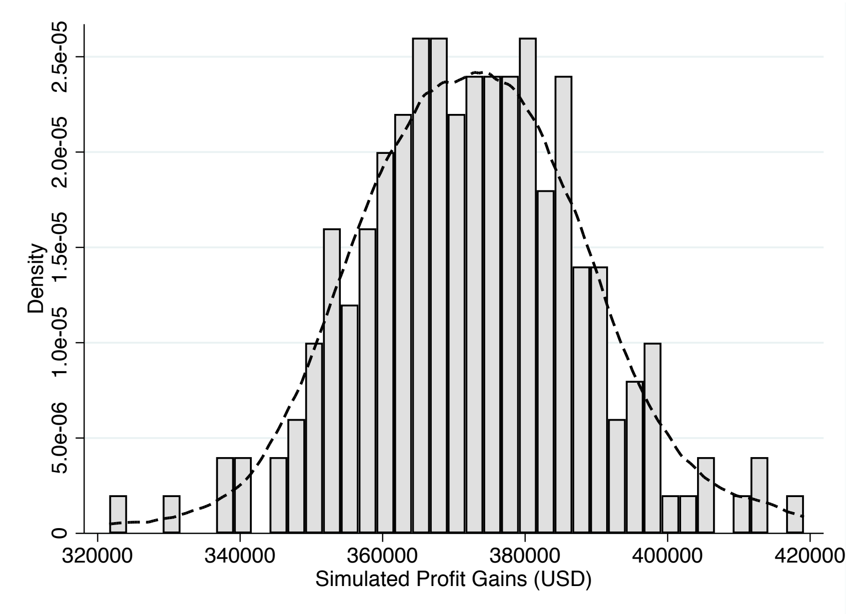


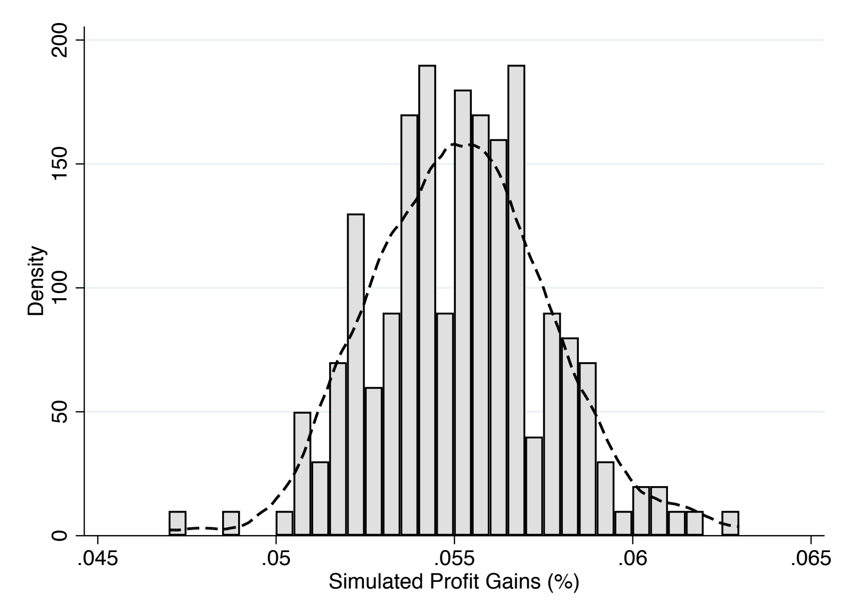


*Notes.* The top figure shows the distribution of overall profit gains in the absence of learning. The bottom figure shows the percentage change in profits over the adjustment period, between the date of announcement of incentive change and the last period in which learning occurs according to the results depicted in Figure 4. Predictions and counterfactuals are estimated with the procedure described in Section 6. Both distributions are obtained after a bootstrap procedure of resampling with replacement in 200 repetitions.