

Recover Overnight?

Work Interruption and Worker Productivity

This Version: February 2017

Abstract

This paper investigates the effect of work interruption on workers' subsequent productivity. We employ a data set of individual productivity and machine conditions, in which each worker faces the chance, on a daily basis, that her machine will break down randomly. Our analysis finds that compared to a workday with smooth production, experiencing a machine breakdown is associated with a 3.3 percentage point decline in the worker's productivity the following day. We discuss possible explanations for the observed effect, including negative emotions, increased cautiousness in operating the machine and proficiency loss. Our findings shed light on the importance of understanding and managing interruptions in the workplace, and contribute to a growing literature on the determinants of productivity at the micro level.

Keywords: Productivity, Labor supply, Work interruption

1 Introduction

Interruptions are common in the workplace and costly to individuals, firms, and organizations. From equipment breakdowns to unscheduled meetings and communication requests, unplanned breaks from a smooth and continuous production process trigger losses in work hours and reductions in worker productivity. For instance, according to Spira and Feintuch (2005), the direct costs of unnecessary interruptions have been estimated at an average of 28% of daily time for knowledge workers in the U.S. Nowadays, with firms adopting modern organizational and communication technologies—open-plan offices, e-mail, instant messaging, etc.—managing interruptions has become increasingly important and challenging for businesses and their workers.

In addition to impairing direct productivity, work interruption may have spillover effects on subsequent production. On the one hand, it has been widely recognized that interruptions can have negative spillover: Workers may need extra time to warm up and regain full engagement and concentration; productivity may also decline due to negative emotions, such as stress from time pressure or frustration about failing to meet targets (Mandler 1990).¹ On the other hand, interruptions may be beneficial for subsequent productivity, as occasional breaks could help workers alleviate fatigue and boredom (Roy 1959). Moreover, workers might try to make up for losses caused by interruptions and exert greater effort leading to increased productivity (Camerer et al. 1997). In the field of management and organization sciences, both positive and negative effects of interruptions have been extensively examined (see Jett and George 2003 for a review). Nevertheless, scant empirical work has estimated the consequences of work interruption, partly because of the difficulty of identifying each incident and establishing a causal relationship.

This paper is the first to estimate the effect of interruption on workers' subsequent productivity using exogenous incidents of work interruption. To do so, we examine the consequences of machine breakdown in a plastics-printing company in China. Manufacturing companies continuously face the possibility of machine breakdown, and must consider how much resources to expend on machine maintenance. For manufacturers in developing countries in transition from labor-intensive to technology-intensive structures, the cost of inadequate maintenance is particularly important yet also little understood.

Specifically, we collect a data set of worker-level daily output and machine conditions. Our sample spans a period of 473 days, during which a total of 273 workers labored on 75

¹In the workplace, it has been observed that negative shocks, such as failed salary negotiations (Mas 2006), bonus payments that fall short of individually assigned bonus targets (Ockenfels, Sliwka and Werner 2014), and bereavement and family illness (Oswald, Proto and Sgroggi 2014), adversely affect workers' subsequent productivity.

machines for 25 different products. As workers are paid by piece rate, we have an accurate record of workers' daily output levels. On a daily basis, team managers assign each worker to a machine and a product type, and each worker faces the chance that the machine will break down. On average, machines have a breakdown frequency of 14%, and each breakdown takes about 4 hours to repair.²

Our empirical strategy compares workers with and without the experience of machine breakdown the previous day. We find that the incidence of machine breakdown leads to a 3.3 percentage point decline in the worker's productivity on the following day; this result is consistent across alternative estimation methods and several robustness checks.

Our findings document a hidden cost of work interruption: The cost is not limited to the hours and productivity lost to interruption, but may also spill over to subsequent production and cast significant cost for the firm and the workers. For the firm, the 3.3 percentage point decline in productivity can be translated to a drop in annual output with estimated sales value of 334,000 to 668,000 RMB. For workers, our calculation suggests that machine breakdown would reduce a worker's daily income by 9.85 RMB, or 5.28%. In light of these back-of-the-envelope calculations, the costs that arise from negative spillover are economically meaningful and should be accounted for in firms' management decisions.

We further discuss possible interpretations of the observed negative effect. First, it may arise from a negative emotional response to the shock of interruption. Second, workers may become more cautious in operating the machine after machine breakdown, leading to a drop in productivity. Third, workers may become less proficient and less engaged after the interruption; for example, they must start over with the techniques and procedures. While we can not completely disentangle different channels, our evidence is generally consistent with the emotion hypothesis. We further explore worker heterogeneity, and find that the effect is similar for male versus female, young versus old, and local versus migrant workers. However, the productivity decline is smaller when more peers also experience the shock of machine breakdown. This result indicates that social comparison may moderate the adverse effect, which also lends support to the emotion hypothesis.

Our study contributes to the understanding of productivity differentials across firms and countries and over time. In a comprehensive review, Syverson (2011) survey micro-level determinants within the firm such as information technology, capital input, and management practice, as well as macro-level influences external to the firm including industry or market environment. For example, Hall and Jones (1999) show that productivity differences across countries can only be partially explained by physical capital and educational attainment,

²Our data contain information on the duration of machine breakdown, but not its exact time, which precludes us from examining the effect of breakdown on productivity in the same day.

and suggest the importance of institutions and government policies. Surveying managers over 700 firms in the United States, Bloom and Van Reenen (2007) document that the quality of management practices is positively correlated with the productivity of the firm. In a randomized controlled field experiment in India, Bloom et al. (2013) show that provision of free consultation on management practices increases firm-level productivity. Combining survey data and longitudinal earnings records of Germany firms, Bender et al. (2016) show that firms with better management practices systematically recruit workers with higher human capital. Following this literature, our study suggests that managing interruptions in the workplace can be yet another aspect of management practices in determining productivity.

Our study also contributes to a growing literature on the economic consequences of work interruption. Researchers have investigated various forms of interruptions in the workplace, including worker absence, menstruation, weather, pollution, and multitasking. Herrmann and Rockoff (2012b) use unique data on worker absence to estimate the effect of work absences on productivity. They find that expected loss in daily productivity from employing a temporary substitute is on par with replacing a regular worker of average productivity with a worker of 10 – 20% lower productivity. Ichino and Moretti (2009) observe that menstruation as a form of interruption contributes to gender gaps in absenteeism and earnings, while Herrmann and Rockoff (2012a, 2013) find little support for its role in explaining the gender gap in earnings. Connolly (2008) explores the effect of exogenous variation in daily weather on labor supply and finds that on average men shift 30 minutes from leisure to work on rainy days, giving rise to a rough estimate of the intertemporal elasticity of labor supply at around 0.01. Zivin and Neidell (2012) use a novel panel dataset with farm workers to investigate the impact of pollution on labor supply, and find that a 10 ppb decrease in ozone concentrations increases worker productivity by 4.2%. Zivin and Neidell (2014) estimate the influence of temperature on time allocation, and find that temperature increases at the higher end of the distribution reduce hours worked in industries with high exposure to climate, while temperature increases at the lower end of the distribution do not show such an effect. Coviello et al. (2014) propose a model of task juggling, in which a worker is interrupted and switches from one project to another too frequently and Coviello et al. (2015) estimate the causal effect of an exogenously induced increase in parallel working of some judges in Italy and show that juggling causally and substantially lowers their productivity.

Our findings are also related to income targeting behavior in labor supply, which has been extensively investigated in the recent literature. In a seminal study by Camerer et al. (1997), the authors find that cabdrivers work more on days when the transient salary is low, which suggests that they may have a daily income target and quit working once they reach that target. Several follow-up studies use field evidence and observational data to investigate the

daily income-targeting hypothesis. The evidence appears to be mixed. Several studies find positive evidence for the daily income-targeting hypothesis (Fehr and Goette 2007; Crawford and Meng 2011), while others suggest otherwise (Farber 2005, 2008; Andersen et al. 2014). As our data do not contain same-day productivity before and after interruption (see data section for details), we cannot directly test daily income targeting. Alternatively, our setting could be interpreted as a test of cross-day income-targeting behavior, in which workers would work harder on the second day to make up for losses on the first day. However, our results do not support this prediction, suggesting that income targeting behavior is more likely to be “one day at a time.”

The rest of the paper proceeds as follows. In Section 2 we describe our research setting, including the institutional background and a stylized model. In Section 3 we discuss the data, main variables, and identification strategy. Section 4 presents our empirical estimates, robustness checks and interpretation of the results. Section 5 concludes.

2 Context

2.1 Workplace

We collect a data set of worker-level daily output from a leading company that produces double-wall paper cups in the city of Xiamen, Fujian Province of China. Established in 2003, the company has become a major supplier of paper cups for the food and beverage industry; in 2011, the company accounted for around 15% of the national market. By 2013, it had \$7 million in total assets, \$5 million in annual revenue, and about 300 employees.

Our data focus on production workers in the molding and wrapping divisions. Their main task is to operate the molding machines to seal the cup body and bottom, and then use wrapping machines to bond sleeves for the paper cups. Workers are hired from both within and outside Fujian Province. Each day, workers are divided into day and night shifts, and typically work for 12 hours. Specifically, the day shift is from 6 a.m. to 6 p.m., and the night shift from 6 p.m. to 6 a.m. the next day. In our sample, 49% of the observations are from the day shift, and the other 51% from the night shift. Within a shift, each worker is assigned a machine and a product type. A machine can generally be configured for multiple types of products, and production efficiency may differ across products.

Notably, workers do not select their product or machine. Each day, a general production manager determines the composition of the production teams, which consist of a team manager and four or five workers. The composition varies daily, and is largely unknown to team managers and workers until they arrive at the site. Team managers assign their

member workers to machines and product types, and once assigned, workers cannot switch to a different machine or product. The only choice variable a worker has is how much effort to exert in operating the machine and preparing materials for the next stage of production. Moreover, workers rotate tasks and machines, which reflects (1) the firm’s desire to ensure fairness, as machine efficiency varies across products, and (2) the fact that there is little room for human capital accumulation in these low-skilled jobs.

Once production begins, team managers prepare packaging boxes labeled with workers’ names and products, and provide assistance to smooth operations if necessary. A worker’s output depends on the normal operation and functioning of the machine. If a machine breaks down, the worker who operates it calls maintenance staff for repairs and the maintenance staff records the duration of the breakdown. The worker can neither change to a different machine nor leave the site, which means that the worker’s production is completely interrupted until the machine is fixed.³

Team managers are paid a flat rate, while workers are paid a two-tier piece rate based on their daily output and the specified target level. Specifically, on a typical day and for each product-machine pair, the general production manager specifies a target, denoted by \bar{Q}_{pm} , which depends on the features of product p and machine m . When workers meet the target, they are paid a piece rate denoted by α_1 per unit of output. When output exceeds the target, a higher rate, denoted by α_2 (where $\alpha_2 > \alpha_1$), applies to each unit above the target. When a worker fails to meet the benchmark, the unfulfilled units are also deducted by the rate α_2 . Formally, given the worker’s daily production Q_{ipmd} , the salary is determined as $\alpha_1 \bar{Q}_{pm} + \alpha_2 (Q_{ipmd} - \bar{Q}_{pm})$.⁴ The target will be proportionally adjusted if machine breakdown occurs; i.e., given a breakdown of x hours, the worker’s new target becomes $\tilde{Q}_{pm} = \bar{Q}_{pm} \frac{12-x}{12}$.

2.2 A Stylized Model

We construct a parsimonious and stylized model to illustrate our research setting. Consider that the company has N workers indexed by i . The production function is denoted as $Q_{id} = Q(e_{id}, x_{id})$, where e_{id} is the effort chosen by worker i on day d ; and x_{id} are exogenous variables that affect daily output, such as worker proficiency, product features, machine conditions, and workplace environment (e.g., weather and pollution). $Q(\cdot)$ has the features of a standard production function: Output is an increasing and concave function of worker effort, and there is complementarity between worker effort and external productive factors—i.e., $Q_e > 0$,

³Whether the machine is fixed or not by the end of the workday, the next day the worker always rotates machines and products, as other workers do.

⁴In principle, the daily salary could be negative if output level is very low, i.e., $Q_{ipmd} < \frac{\alpha_2 - \alpha_1}{\alpha_2} \bar{Q}_{pm}$. However, since there is no such case in our data, presumably the specified target output level \bar{Q}_{pm} is set at a reasonable level and workers exert sufficient effort.

$Q_{ee} < 0$, $Q_x > 0$, and $Q_{ex} > 0$. A worker i 's cost function is $C_{id} = C(e_{id}, w_{id})$, where w_{id} are external factors, such as the worker's mental status, and the regular assumptions for a cost function apply, namely, $C_e > 0$, $C_{ee} > 0$, $C_w < 0$, and $C_{ew} < 0$.

The optimal effort chosen by worker i on day d is given by

$$\begin{aligned} \max_{e_{id}} Q(e_{id}, x_{id}) - C(e_{id}, w_{id}) \\ \Rightarrow e_{id}^* = e(x_{id}, w_{id}). \end{aligned}$$

Linking the optimal output Q_{id}^* to the external variables (x_{id} and w_{id}), we have

$$\frac{\partial Q_{id}^*}{\partial x_{id}} = Q_e \frac{\partial e_{id}^*}{\partial x_{id}} + Q_x = Q_e \frac{Q_{ex}}{C_{ee} - Q_{ee}} + Q_x > 0, \quad (1)$$

and

$$\frac{\partial Q_{id}^*}{\partial w_{id}} = Q_e \frac{\partial e_{id}^*}{\partial w_{id}} = Q_e \frac{C_{ew}}{Q_{ee} - C_{ee}} > 0. \quad (2)$$

Next, we use equations (1) and (2) to illustrate how machine breakdown on day $d - 1$ affects Q_{id} and the potential bias from omitted variables. We also extend the framework to incorporate heterogeneous effects across workers.

Effect of Machine Breakdown on Day $d - 1$. Machine breakdown on day $d - 1$ can affect output Q_{id}^* on day d through production factor x_{id} or cost factor w_{id} or both. Consider first the possibility that breakdown affects x_{id} . For example, after a substantially long interruption, workers may become less proficient or engaged in the task, and must start over with the techniques and procedures (Jett and George, 2003). The process of reaching maximum proficiency imposes a fixed cost on productivity; once it is interrupted, workers have to start over and incur such a cost. Formally, let B_{id-1} denote that worker i experienced a machine breakdown on day $d - 1$. The proficiency effect means that $\frac{\partial x_{id}}{\partial B_{id-1}} < 0$, which causes $\frac{\partial Q_{id}^*}{\partial B_{id-1}} < 0$.

The other possibility is that machine breakdown affects the external factor of workers' cost of effort w_{id} . For example, workers may feel frustrated by machine breakdown because it is unpleasant, it clashes with their expectations for the production process, or they are disappointed because their daily salary will be compromised due to the lost hours. Such negative emotions render working hard more costly, or $\frac{\partial w_{id}}{\partial B_{id-1}} < 0$, which causes $\frac{\partial Q_{id}^*}{\partial B_{id-1}} < 0$.

Heterogeneous Effects. In the baseline setting, the production function $Q(\cdot)$ and cost function $C(\cdot)$ are the same for all workers on the site. This implies that our estimate is the average effect across all workers. However, worker heterogeneity is commonly observed in

the workplace. We can extend the framework and allow for heterogeneity in both the output and cost function—i.e., $Q_g(\cdot)$ and $C_g(\cdot)$, where g indexes the worker type; this allows us to estimate potentially heterogeneous effects from machine breakdown. In further analyses, we consider differential effects across several demographic characteristics—gender, age, and local versus migrant workers.

3 Data, Variables, and Estimation Strategy

Data. We use data on worker output and machine conditions for each workday from October 2012 to April 2014. The sample covers 473 workdays, 273 workers (all full-time), 75 machines, and 25 product categories. Due to workforce turnover, this is not a balanced panel data set. Specifically, 14 workers (5.1%) were observed throughout our sample period, 53 incumbent workers (19.4%) left the firm, and 206 (75.5%) were newly hired during this period. We plot the density distribution of workers’ tenure during the 473 workdays, with the unit of observation at worker level in Appendix Figure A1. On average, workers’ tenure during the 473 workdays is about 132 days.⁵ A few workers had notably short tenure with the firm (e.g., less than 1 month), yet constitute only a small amount of worker-day observations for our empirical analysis. As a result, excluding them does not affect our main results (see the estimation results in Appendix Table A1). We also examine whether machine breakdown (our regressor of interest) affects workforce turnover, and find no significant impact (see Table 3, column 1).

In the end, we have 24,081 worker-day observations. Similarly, on a given day, depending on the production task—i.e., the assigned product’s features and the specific machines required to produce them—workers operate the assigned machines while other machines remain idle. We plot the density of machine utilization during the sample period in Appendix Figure A2. On average, 29 machines were operating per day.

Baseline Specification. The effect of machine breakdown on a worker’s subsequent productivity comes from comparison of output at day d between workers with and without machine breakdown at day $d - 1$. Specifically, we implement the following baseline regression model:

$$y_{impd} = \beta * Treatment_{im'p'd-1} + \lambda_i + \lambda_d + \lambda_{pm} + \eta Shift_{impd} + \varepsilon_{impd}, \quad (3)$$

⁵During our sample period, the firm had a monthly turnover rate between 9% and 25%, comparable to the industry average in the province. The Fujian Provincial Bureau of Statistics reports that in 2015, manufacturing companies had turnover rates between 25% and 30%. Source: http://www.stats-fj.gov.cn/xxgk/gzdt/wshjyxc/201505/t20150515_37733.htm.

where i indicates worker; d indicates day; and $m(m')$ and $p(p')$ indicate machine and product for $d(d-1)$, respectively. We cluster standard errors at the machine level to control for any heteroskedasticity and serial correlation (Bertrand et al. 2004).

To account for variations in machine efficiency across product types, we define worker productivity y_i as the over-target percentage of output; specifically, $y_{impd} \equiv (Q_{impd} - \bar{Q}_{pm}) / Q_{impd}$, where \bar{Q}_{pm} is the targeted output. As workers are paid by daily output, both the firm and the workers make an effort to record output precisely, so measurement errors are not a large concern. One possible concern, however, is that using the over-target percentage may add a certain degree of persistency in the outcome variable. As a robustness check, we use the log of absolute output, i.e., $\ln Q_{impd}$, as the outcome measure and obtain similar results (see Section 4.1).

The regressor of interest, $Treatment_{im'p'd-1}$, is a dummy variable that indicates whether on day $d-1$, worker i 's machine broke down. Note that since a worker rotates among products and machines on a daily basis, the worker may not use the same machine on two consecutive days. In other words, the machine used by a worker on day $d-1$ may not be the same machine the worker operates on day d , i.e., $m' \neq m$. Similarly, it is possible that workers are assigned to different products across two days, i.e., $p' \neq p$.

We control for confounders through a battery of fixed effects: (1) worker fixed effects λ_i , so that the identification comes from within-worker comparison, which helps control for possible nonrandom rotation after a breakdown episode; (2) product-machine fixed effects λ_{pm} , which helps control for nonrandom machine effects in a flexible way; (3) day fixed effects λ_d , which helps control for daily variations (such as temperature) that are common to all workers; and (4) job shift fixed effects, which allows comparison of workers from the same shift (i.e., day shift or night shift).

Given the potential effect of breakdown on productivity, we estimate equation (3) for a sample of workers without machine breakdown at day d . We also cluster standard errors at the worker level to control for any heteroskedasticity and serial correlation (Bertrand et al. 2004).

To further check the concern about omitted variables in our baseline specification, we conduct a placebo test. Specifically, we replace $Treatment_{im'p'd-1}$ with a dummy variable $Treatment_{im'p'd-1}^{other}$ that indicates whether worker i 's counterpart on the other shift, who also operates machine m at day d , experienced machine breakdown on day $d-1$. Unless there was a perfect assortative matching (that is, always pair two workers whose machines broke down on the previous day to work on the same machine the following day), $Treatment_{im'p'd-1}^{other}$ should have no effect on worker i 's performance on day d (i.e., $\beta^{other} = 0$), as the two workers had different trajectories; otherwise, it indicates a misspecification of our estimation model.

Alternative Specification. In a robustness check, we use an alternative identification framework; that is, a single difference design in which we compare a pair of workers with similar job assignments except for the experience of machine breakdown. Specifically, there are two job shifts every day and during each shift one worker operates the machine. We therefore focus on the two workers who operates the same machine during different shifts on $d-1$, in which the machine breaks down during one shift (treatment group) but not the other (control group). If there had been any manipulation by team managers in worker-machine assignments, the two workers assigned to the same machine on the same day—albeit during two different shifts—should be reasonably similar.⁶

We estimate the effect of machine breakdown on a worker’s subsequent productivity based on the following single difference framework:

$$y_i = \beta * Treatment_i + \varepsilon_i, \tag{4}$$

where y_i denotes the productivity of worker i ; and $Treatment_i$ is a dummy variable that indicates whether worker i ’s machine broke down the previous day. Standard errors are clustered at the shift-pair level to adjust for any heteroskedasticity and serial correlation.

To further address the concern that workers in the treatment and control groups might differ systematically in some unobservable characteristics, we use their productivity for the two previous days as a further control and conduct a difference-in-differences (DD) analysis. Specifically, our DD estimation specification is

$$y_{id} = \beta * Treatment_i \times Post_d + \lambda_i + \lambda_d + \varepsilon_{id}, \tag{5}$$

where $Post_d$ is a post-breakdown day indicator, and λ_i and λ_d are worker and day fixed effects, respectively. Standard errors are clustered at the machine level to control for any heteroskedasticity and serial correlation.

Such within-worker comparisons can eliminate all differences across workers—including the differences between the treatment and control groups—that do not change across two days. This could include, for example, worker ability, experience, social connections with the managers or within the firm, etc.

⁶To verify whether there was nonrandom assignment of machine at day d after a breakdown episode on day $d-1$, we examine whether the machines assigned at day d to workers with and without machine breakdown in the previous day had similar average productivity. Specifically, we examine the machine’s average productivity over the last day, over the last 7 days, and over the last 30 days. As reported in Appendix Table A2, there are no statistically and economically significant differences in the machines’ average productivity, lending support to the plausibly random assignment of machines after a breakdown episode.

4 Empirical Findings

Table 1 reports the regression results of specification (3). In column 1, we only include three sets of fixed effects—worker fixed effect, product-machine fixed effect, and day fixed effect—and find that the coefficient of interest, *Treatment*, is negative and statistically significant. In column 2, we further add the job shift control. Despite a significant drop in observations (information on job shift is only available for a subset of the observations), we continue to find a negative and statistically significant coefficient of similar magnitude. Using the more conservative estimates, we find that machine breakdown on the previous day leads to a 3.3 percentage point drop in over-target output.

[Insert Table 1 here]

In the next subsections, we test the robustness of the results by conducting several validity checks, calculate the economic magnitude of the effects, and examine possible interpretations.

4.1 Validity Checks

In this subsection, we adopt the methods detailed in Section 3 and report the results of several validity checks—namely, using an alternative estimation framework, measuring productivity using absolute output, conducting a placebo test to rule out omitted variables and testing for longer-term dynamic effects.

Alternative Estimation Framework. We use the alternative estimation strategy described in Section 3—that is, the difference between a pair of workers on the same machine but different shifts at day $d - 1$, with one having a machine breakdown and the other having no breakdown. Appendix Table A3, columns 1 and 2, report the results of the single difference equation (4) with and without worker characteristics controls, respectively, and column 3 reports the results of DD specification (5). Across all specifications, we consistently find a negative and statistically significant effect of machine breakdown on worker productivity on the following day. Meanwhile, the magnitude from the DD specification is close to that in the baseline specification, suggesting that our estimated effect is robust and stable.

The DD estimation requires that treatment and control groups would have followed the same time trend if the treatment group did not receive the treatment shock. However, it is difficult to test the parallel trend assumption in our setting, as a machine may have multiple breakdowns. One standard solution to check the possible violation of the parallel trend assumption is to include machine-specific linear time trend; that is, to include the interaction between machine fixed effects (λ_m) and day variable (d), in equation (5). Table

A3, column 4, reports the results with machine-specific linear trends. The estimated effect of machine breakdown remains robust without any changes in the statistical significance and magnitude, lending support to the validity of our DD estimation.

Alternative Measurement of Productivity. To address the concern that our measurement of worker productivity (i.e., over-target percentage of output) may increase the persistency of outcome variables over time, we use the logarithm of the daily output as an alternative measure. Estimation results are reported in panel A of Appendix Table A4, column 1. Again, we find a negative and statistically significant coefficient: Machine breakdown leads to a 4% decline in output on the following day. This effect size is comparable to the magnitude of baseline estimates, suggesting that our findings are not driven by a particular measure of worker productivity. We conduct similar analysis using the DD framework in panel B, column 1, and continue to find a negative, although marginally insignificant effect (p -value=0.102).

Placebo Test. To further ensure that our estimates are not caused by omitted variables, we conduct a placebo test as described in Section 3. Specifically, we replace $Treatment_{im'p'd-1}$ with $Treatment_{im'p'd-1}^{other}$, a dummy variable that indicates whether worker i 's counterpart from the other shift—who used the same machine as the concerned worker on day d —experienced machine breakdown on day $d - 1$. Unless the two workers were assigned to the same machine for two consecutive days, $Treatment_{im'p'd-1}^{other}$ should have no effect on worker i 's productivity on day d . Thus, a significant estimate of $Treatment_{im'p'd-1}^{other}$ would indicate that there exist omitted variables in our main specification. As shown in panel A of Appendix Table A4, column 2, the coefficient of $Treatment_{im'p'd-1}^{other}$ is insignificant and small in magnitude, suggesting that there are no severe omitted variables in our baseline specification. Similar results are found in the DD estimation in panel B, column 2.

Persistent Effects. We further investigate how long this negative effect persists by including dummies that indicate machine breakdown day $d - 2$ and day $d - 3$. As shown in panel A of Appendix Table A4, column 3, results suggest that while the negative effects persist even after three days, the magnitude drops by around 30% on the third day (with an estimated coefficient of -0.019) and further drops by 47% on the fourth day (with an estimated coefficient of -0.010). This strengthens our observation, and suggests that the negative spillover could last for a number of days after the interruption. The persistent effects are also found in the DD estimation (panel B, column 3).

4.2 Economic Magnitude

Our estimate suggests that the adverse effect is about a 3.3 percentage point drop in the over-target percentage of output (Table 1, column 2). Given that the control group—i.e., workers without machine breakdown on day $d - 1$ —averages 19.95% over-target output on day d , machine breakdown reduces over-target output by $3.3/19.95 = 16.54\%$, or gross daily output by 3.96%.⁷

Understanding the economic magnitude of adverse breakdown effects is also valuable. To understand their impact on the firm, we conduct a back-of-the-envelope calculation to assess the cost to the firm of a breakdown and the accompanying productivity loss. On average, 29 machines operate every day, with a 14.4% chance of a breakdown that will take 4 hours to fix. Given that the average output per machine per hour is 2,309.68, annual losses in output due to suspended production are about $2,309.68 * 4 * 14.4\% * 29 * 365 = 14$ million units. Less tangible is the cost of lost productivity on the following day. Using our benchmark estimates (i.e., 3.96% loss of gross daily output and the control group’s average output of 27,626.82), annual losses due to our documented effect are $3.96\% * 27,626.82 * 14.4\% * 29 * 365 = 1.67$ million units, or 334,000 to 668,000 RMB in revenue (evaluated at the sales price of 0.2 to 0.4 RMB per unit).⁸ In practice, firms are aware of the lost hours and maintenance costs that arise from machine breakdown, but tend to ignore the cost incurred by lower subsequent productivity. Our calculation suggests that this hidden cost is indeed economically significant, and is therefore relevant to how firms manage interruptions.⁹

We then calculate the worker’s economic loss. Recall that a worker’s income is determined by his or her actual output Q_{ipmd} and target \bar{Q}_{pm} : $w_{id} = \alpha_1 \bar{Q}_{pm} + \alpha_2 (Q_{ipmd} - \bar{Q}_{pm})$. Our estimates show that machine breakdown lowers subsequent gross daily output by 3.96%. Given that the average output of the control group on day d is 27,626.82 and $\alpha_2 = 0.009$, a worker earns $\Delta w_{id} = \alpha_2 \Delta Q_{ipmd} = 9.85$ RMB less on the following day. In our sample, a worker’s average daily earnings are 186.47 RMB, machine breakdown therefore leads to a fall in daily earnings of $9.85/186.47 = 5.28\%$. On any given day, machine break down with an average probability of 14.4%. Losses to a worker’s annual income, therefore, are around $14.4\% * 365 * 9.85 = 517.72$ RMB. Assuming that the adverse effect applies to the 1.05 million manufacturing workers in Xiamen, total losses would be about 543.61 million RMB in earnings (equivalent to US\$80.42 million at an exchange of 6.76) due to reduced

⁷We define over-target percentage of output as $y_i \equiv (Q_i - \bar{Q}) / Q_i$. The average gross output of the control group is then $Q_i^c = \bar{Q} / (1 - 0.1995)$, while the average output of the treatment group is $Q_i^t = \bar{Q} / (1 - 0.1665)$. The treatment effect is equivalent to $(Q_i^t - Q_i^c) / Q_i^c = -3.96\%$ of the gross output.

⁸Because the firm’s profit margin is confidential information, we use the sale price instead to calculate losses in revenue.

⁹As we don’t have information on the cost of repairs, this is a lower bound of estimated costs to the firm.

output following machine breakdown. Viewed in this light, the adverse effect is economically substantial.

4.3 Interpretation

Our analyses demonstrate that machine breakdown has a negative and significant effect on workers’ productivity the following day; we now discuss possible interpretations of this effect. In particular, we test three possible channels of negative spillover: negative emotional reactions, increased cautiousness, and production proficiency loss.

Negative Emotional Reactions. Machine breakdown may cause emotional reactions—e.g., frustration or disappointment—and, therefore, reduce output on day d . As explained in Section 2.1, workers’ payment is entirely dependent on their daily output. Machine breakdown causes direct losses in working hours and payment, so it is plausible that workers are disappointed about failing to earn as expected for the day.

There is evidence that the emotions triggered by external shocks affect workers’ output. For example, Ockenfels et al. (2014) show that when bonus payments fall short of individually assigned bonus targets, workers are disappointed, which leads to lower work satisfaction and performance. Mas (2006) finds that after New Jersey police officers lose in final-offer arbitration over salary demands, relative to when they win, arrest rates and average sentence length decline and crime reports rise. Oswald et al. (2014) show that increased happiness leads to higher productivity, and decreased happiness caused by major real-world shocks, including bereavement and family illness, leads to lower productivity.¹⁰

To check the relevance of the potential channel through negative emotions, we conduct several exercises. First, given that machine conditions are plausibly common knowledge, workers who are assigned to “bad” machines (i.e., those more likely to break down) would be unhappy, which might affect their performance even in the absence of an actual breakdown. To examine this emotional response, we first proxy for a machine’s condition using its historical frequency of breakdown (i.e., the previous day, in the past 7 days, and in the past 30 days), and then for observations where there was no breakdown event, regress worker productivity on this machine condition measurement along with a full set of fixed effects as in the baseline equation (3). Estimation results are reported in Table 2. We find that all

¹⁰Relatedly, negative shocks in the workplaces may induce angry and even harmful behavior by the worker—e.g., drinking, violence, or deliberate sabotage—which in turn, lowers their second-day productivity. Card and Dahl (2011) observe that losses in professional football matches increase the rate of at-home violence by men against their wives or girlfriends. Lowenstein (2000) suggests that emotions, including a wide range of visceral factors, underpin daily economic behavior. For example, angry negotiators could become obsessed with causing harm to the other party, even at the cost of their own interests.

coefficients are negative and statistically significant, which suggests that being assigned to a machine with a history of breaking down lowers worker productivity on that day (columns 1-3) and the following day (columns 4-6) even if the machine does not actually break down, consistent with a hypothesis of negative emotional reactions.

[Insert Table 2 here]

Second, after experiencing machine breakdown, workers could get too frustrated to attend work on the following day. To examine how interruption affects work participation, we replace the productivity outcome in equation (3) with a dummy variable to indicate absence from work at day d . Estimation results are reported in Table 2, column 7. We find higher probability of being absent from work on the day following machine breakdown. The increased absenteeism after a breakdown episode lends further support to the notion that machine breakdown triggers negative emotional reactions, which discourage labor supply.¹¹

Increased Cautiousness. Another possible channel is that workers become more cautious in handling their machines after they experience breakdowns—e.g., they spend more time fine-tuning, slow down, or take more breaks; this would in turn reduce their output. Such behavioral change can be justified even when the breakdown is completely exogenous, because machine damage incurs costs for the company (including repair expense and lost production time), and workers may worry about being fired if they break the machine again in a small interval of time after the previous breakdown.

To access this possibility, we first check whether a worker is more likely to be fired or quit after his or her machine breaks down. As shown in Table 3, column 1, we find a small and statistically insignificant coefficient, which suggests that machine breakdown does not result in workers leaving the firm.

[Insert Table 3 here]

¹¹The increased absenteeism may raise the concern of sample selection: If more productive workers are more likely to be absent from work after a breakdown episode, then our estimates simply reflect the sample selection instead of change in worker productivity. We address this potential bias using the methodology in Lee (2009), the premise of which is to consider the best and worst scenarios caused by sample selection and bound the estimated effect. Specifically, in our context, consider the case in which the 2.2% higher absence in the treatment group arises from the most productive workers, i.e., the largest bias. Then we drop the top 2.2% productive workers in the control group to construct a balanced sample, and the resulting estimates constitute the lower bound of the true effect. Similarly, if we assume that in the treatment group, workers who are absent the following day are the least productive ones, then we exclude the bottom 2.2% productive workers in the control group and obtain the upper bound of the true effect. Appendix Table A4, columns 5 and 6 report the lower and upper bounds of the effects for both baseline and DD specifications, respectively. We find that both are significant and close to our baseline estimates, which suggests that the adverse effect is not driven by sample selection.

Nevertheless, workers may still try to avoid another breakdown; after all, they have to stop and wait for maintenance staff to fix the machine, which lowers output, and might be pressed by team managers and peers to be cautious. One possible implication of being more cautious is that they work fewer hours the next day because they spend more time fine-tuning their machine. To check this possibility, we examine whether machine breakdown influences the worker’s production hours in the following day. As shown in Table 3, column 2, the coefficient is statistically insignificant and the magnitude is quite small, suggesting that workers do not lose production time the day after machine breakdown.

Another testable implication of increased cautiousness is to see whether working on a previously repaired machine affects workers’ productivity. If past experience of machine breakdown makes a worker more cautious, the effect would be stronger when the worker operates a machine that was repaired on the previous day, because he or she is concerned about the machine’s breaking down again. Indeed, we show in Appendix Table A4, column 4, that a previously broken machine, after being fixed, is more likely to break down again. Accordingly, we include an additional control, $Breakdown_{md-1}$, in baseline regression (3) to indicate whether machine m used on day d by worker i broke down on day $d - 1$ (not necessarily used by the same worker). Estimation results are reported in Table 3, column 3. While our regressor of interest ($Treatment_{im'p'd-1}$) remains negative and statistically significant, its magnitude drops by 36%. Meanwhile, the new control, $Breakdown_{md-1}$, is also statistically significant. These results suggest that after a machine breaks down, the worker’s subsequent output declines, especially when assigned to a machine that was repaired on the previous day. The increased cautiousness hypothesis, therefore, can partly explain our findings.

Production Proficiency Loss. The third channel through which machine breakdown may negatively affect workers’ subsequent production is the proficiency lost by being interrupted. It is possible that after a substantial period of pausing and waiting—especially when the interruption is unscheduled—workers become less proficient, and need extra time to regain their momentum with techniques and procedures before returning to full engagement and concentration.

To test this potential channel, we investigate whether the effect of machine breakdown on day $d - 1$ varies by the worker’s previous working hours. Specifically, we look at the working hours at day $d - 2$, the cumulative working hours in the past 7 days (up to day $d - 1$), and the cumulative working hours in the past 30 days (up to day $d - 1$). Estimation results are reported in Table 4. None of the three interactions has any statistical and economical significance. These results largely reject the proficiency hypothesis—that is, there is no evidence

that workers who previously worked more intensively are less affected by the interruption.

[Insert Table 4 here]

4.4 Heterogeneity and Peer Effects

We now explore heterogeneous effects across worker characteristics. The literature suggests that there is substantial heterogeneity in economic behavior, which can be partially accounted for by demographic information—gender, age, socioeconomic background, etc. For instance, Dohmen et al. (2011) conduct a study with a representative sample of roughly 22,000 individuals in Germany, and find that willingness to take risks is negatively related to age and gender, and positively related to height and parental education. Similarly, it has been suggested that gender plays an important role in economic preference (Croson and Gneezy 2009), and that for some decision behaviors, elderly individuals are less biased than younger individuals (Kovalchik et al. 2005).

Motivated by these studies, we examine whether the effect of machine breakdown differs by workers’ gender, age, or place of residence. To do this, we interact each worker characteristic with our regressor of interest (i.e., $Treatment_{im'p'd-1}$); hence, coefficients of the interactions represent heterogeneous effects across these characteristics. Table 5 reports the results. Column 1 shows results for heterogeneous effects across gender. We find that despite the magnitude’s not being small, the interaction term is statistically insignificant. These results suggest that there is no significant difference in machine breakdown effects between males and females.

[Insert Table 5 here]

Column 2 investigates heterogeneous effects for young and old workers. The coefficient of the interaction term is statistically insignificant and small in magnitude, which suggests that young and old workers experience similar negative effects of machine breakdown.

Column 3 examines heterogeneous effects for local and immigrant workers—i.e., those who are not registered as Fujian residents. The interaction is statistically insignificant and small in magnitude, which suggests that local residency does not play an important role in coping with the shocks of work interruption.

Finally, we check whether the effect is reinforced or mediated by peer effect. Peer effects in the workplace have been extensively explored in the literature; for instance, Mas and Moretti (2009) observe that a 10% increase in coworker productivity is associated with a 1.5% increase in a worker’s productivity. De Grip and Sauermann (2012) exploit a field

experiment and observe that a 10 percentage points increase in the share of treated peers improves an individual’s performance by 0.51%. Overall, these studies suggest that workers are motivated by social pressure and mutual monitoring.

Here, we include an interaction term between a dummy that indicates high frequency of machine breakdown at the firm level on day $d - 1$ (i.e., above sample median) and the treatment dummy. Table 5, column 4 presents the results. The coefficient on the interaction term is positive and statistically significant, which suggests that when more peers experience interruption, the negative effect on a worker’s productivity tends to be smaller. This moderating effect further supports the channel through emotional reactions: Interruption may damage a worker’s morale and lower his productivity, but the damage will be milder if more peers experience a similar shock. To that extent, social comparison in the workplace can moderate the emotional reactions to negative production shocks.

4.5 Discussion

Overall, our analyses show that machine breakdown leads to a 3.3 percentage point decline in the worker’s over-target percentage of output the following day. In this subsection, we examine the results in light of recent discussions on statistical power and false positive of empirical studies (see, for example, Levitt and List 2009; List, Sadoff and Wagner 2011; Maniadis, Tufano and List 2014). In particular, Maniadis, Tufano and List (2014) suggest that statistical significance as the sole criterion can lead to an excessive number of false positives. Moreover, the authors propose the notion of the Post-Study Probability (PSP), the probability that a research finding reveals a true effect, which depends not only on statistical significance, but also on the prior assigned to the observation and the statistical power of the design.

[Insert Table 6 here]

Using the method proposed in Maniadis, Tufano and List (2014) and a significant level 0.01, we compute the post-study probability with three levels of statistical power and three levels of priors of a true effect (see Table 6). As the analysis employs a panel data of 273 workers over 473 work days, our study can be regarded as having high statistical power. For instance, if we take our statistical power as 0.8, our result can substantially change the prior probability of 1%, 25% and 50% to the post-study probability of 44.7%, 96.4% and 98.8%, respectively. Moreover, an important implication drawn from Maniadis, Tufano and List (2014) is that successful replication can substantially increase the post-study probability and reduce the likelihood of reporting false positives. We therefore encourage subsequent

replications on our study.

5 Conclusion

This paper investigates how machine breakdown affects workers' subsequent productivity. Using daily output data, we show that individual productivity declines following a workday with machine breakdown. This adverse effect cannot be explained by either traditional theories of labor supply or income targeting behavior. In the sense that the marginal return on effort is larger when production is restored, workers would be expected to work harder the following day. Alternatively, if the worker has a targeted level of income that guides his labor supply, we would also observe higher subsequent output from the worker to compensate for his income losses during the interruption. By investigating possible channels, we find evidence that interruption may result in negative emotions for the workers, which lowers their next-day productivity.

Our findings document a hidden cost of work interruption: The cost is not limited to the hours and productivity lost to interruption, but may also spill over to subsequent production and can persist for days. This hidden but economically significant cost is relevant to firms' decisions about how much resources to expend on managing interruptions, including the costs of maintaining equipment and arranging for standbys. A further implication lies in the remedies for interruptions. We find evidence that interruption causes subsequent output loss because of negative emotions. One implication, therefore, is that managers may want to motivate the affected individuals and restore their morale.

More generally, our study highlights work interruption as yet another determinant of productivity along with external factors such as incentive schemes, peer effects, weather and pollution, etc. It would be interesting to compare the effects of other work interruptions, or examine how the effect of machine breakdown can be generalized to other types of interruptions—and, in particular, those caused by communication devices and social media, as they are increasingly common and have become a concern for modern firms and organizations. These would be fruitful avenues for future research and contribute to better understanding of the causes and consequences of, and remedies for, interruptions in the workplace.

References

Andersen, Steffen, Alec Brandon, Uri Gneezy, and John A List, “Toward an Understanding of Reference-Dependent Labor Supply: Theory and Evidence from a Field

- Experiment,” Technical Report, National Bureau of Economic Research 2014.
- Bender, Stefan, Nicholas Bloom, David Card, John Van Reenen, and Stefanie Wolter**, “Management Practices, Workforce Selection and Productivity,” Working Paper, National Bureau of Economic Research 2016.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan**, “How Much Should We Trust Differences-in-Differences Estimates?,” *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.
- Bloom, Nicholas and John Van Reenen**, “Measuring and Explaining Management Practices Across Firms and Countries,” *The Quarterly Journal of Economics*, 2007, *122* (4), 1351–1408.
- , **Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts**, “Does Management Matter? Evidence from India,” *The Quarterly Journal of Economics*, 2013, *128* (1), 1–51.
- Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler**, “Labor supply of New York City cabdrivers: One day at a time,” *The Quarterly Journal of Economics*, 1997, pp. 407–441.
- Connolly, Marie**, “Here Comes the Rain Again: Weather and the Intertemporal Substitution of Leisure,” *Journal of Labor Economics*, 2008, *26* (1), 73–100.
- Coviello, Decio, Andrea Ichino, and Nicola Persico**, “Time Allocation and Task Juggling,” *The American Economic Review*, 2014, *104* (2), 609–623.
- , – , and – , “The Inefficiency of Worker Time Use,” *Journal of the European Economic Association*, 2015, *13* (5), 906–947.
- Crawford, Vincent P and Juanjuan Meng**, “New York City Cab Drivers’ Labor Supply Revisited: Reference-Dependent Preferences with Rational-Expectations Targets for Hours and Income,” *The American Economic Review*, 2011, pp. 1912–1932.
- Crosen, Rachel and Uri Gneezy**, “Gender Differences in Preferences,” *Journal of Economic literature*, 2009, pp. 448–474.
- Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G Wagner**, “Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences,” *Journal of the European Economic Association*, 2011, *9* (3), 522–550.

- Farber, Henry S**, “Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers,” *Journal of Political Economy*, 2005, 113 (1).
- , “Reference-Dependent Preferences and Labor Supply: The Case of New York City Taxi Drivers,” *The American Economic Review*, 2008, 98 (3), 1069–1082.
- Fehr, Ernst and Lorenz Goette**, “Do Workers Work More if Wages Are High? Evidence from a Randomized Field Experiment,” *The American Economic Review*, 2007, pp. 298–317.
- Grip, Andries De and Jan Sauermann**, “The Effects of Training on Own and Co-worker Productivity: Evidence from a Field Experiment*,” *The Economic Journal*, 2012, 122 (560), 376–399.
- Hall, Robert E and Charles I Jones**, “Why Do Some Countries Produce So Much More Output Per Worker Than Others?,” *The Quarterly Journal of Economics*, 1999, 114 (1), 83–116.
- Herrmann, Mariesa A and Jonah E Rockoff**, “Does Menstruation Explain Gender Gaps in Work Absenteeism?,” *Journal of Human Resources*, 2012, 47 (2), 493–508.
- **and** – , “Worker Absence and Productivity: Evidence from Teaching,” *Journal of Labor Economics*, 2012, 30 (4), 749–782.
- **and** – , “Do Menstrual Problems Explain Gender Gaps in Absenteeism and Earnings?: Evidence from the National Health Interview Survey,” *Labour Economics*, 2013, 24, 12–22.
- Ichino, Andrea and Enrico Moretti**, “Biological Gender Differences, Absenteeism, and the Earnings Gap,” *American Economic Journal: Applied Economics*, 2009, 1 (1), 183–218.
- Jett, Quintus R and Jennifer M George**, “Work Interrupted: A Closer Look At the Role of Interruptions in Organizational Life,” *Academy of Management Review*, 2003, 28 (3), 494–507.
- Kovalchik, Stephanie, Colin F Camerer, David M Grether, Charles R Plott, and John M Allman**, “Aging and Decision Making: A Comparison between Neurologically Healthy Elderly and Young Individuals,” *Journal of Economic Behavior & Organization*, 2005, 58 (1), 79–94.
- Lee, David S**, “Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects,” *The Review of Economic Studies*, 2009, 76 (3), 1071–1102.

- Levitt, Steven D and John A List**, “Field experiments in economics: The past, the present, and the future,” *European Economic Review*, 2009, 53 (1), 1–18.
- List, John A, Sally Sadoff, and Mathis Wagner**, “So you want to run an experiment, now what? Some simple rules of thumb for optimal experimental design,” *Experimental Economics*, 2011, 14 (4), 439.
- Mandler, George**, “Interruption (Discrepancy) Theory: Review and Extensions,” *On the move: The psychology of change and transition*, 1990, 13, 32.
- Maniadis, Zacharias, Fabio Tufano, and John A List**, “One swallow doesn’t make a summer: New evidence on anchoring effects,” *The American Economic Review*, 2014, 104 (1), 277–290.
- Mas, Alexandre**, “Pay, Reference Points, and Police Performance,” *Quarterly Journal of Economics*, 2006, 121 (3).
- **and Enrico Moretti**, “Peers at Work,” *American Economic Review*, 2009, 99 (1), 112–145.
- Ockenfels, Axel, Dirk Sliwka, and Peter Werner**, “Bonus Payments and Reference Point Violations,” *Management Science*, 2014.
- Oswald, Andrew J, Eugenio Proto, and Daniel Sgroi**, “Happiness and Productivity,” *Journal of Labor Economics*, 2014. forthcoming.
- Roy, Donald F**, “Banana Time: Job Satisfaction and Informal Interaction,” *Human Organization*, 1959, 18 (4), 158–168.
- Spira, Jonathan B and Joshua B Feintuch**, *The Cost of Not Paying Attention: How Interruptions Impact Knowledge Worker Productivity*, Basex New York, NY, 2005.
- Syverson, Chad**, “What Determines Productivity?,” *Journal of Economic literature*, 2011, 49 (2), 326–365.
- Zivin, Joshua Graff and Matthew Neidell**, “The Impact of Pollution on Worker Productivity,” *The American Economic Review*, 2012, 102 (7), 3652–3673.
- **and –**, “Temperature and the allocation of time: Implications for climate change,” *Journal of Labor Economics*, 2014, 32 (1), 1–26.

Table 1: Baseline Estimates

	(1)	(2)
	Baseline	With Day shift control
Treatment	-0.036*** (0.009)	-0.033*** (0.006)
Day Shift		-0.001 (0.002)
Controls:		
Individual FE	YES	YES
Product*Machine		
FE	YES	YES
Day FE	YES	YES
Observations	17,587	11,285

Notes: The dependent variable is over-target output percentage. The main independent variable is treatment, namely, whether the worker experiences machine breakdown the previous day. All standard errors are clustered at the machine level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 2: Interpretation I, Negative Emotions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Machine condition on same-day productivity			Machine condition on next-day productivity			
	Yesterday	Past 7 days	Past 30 days	Yesterday	Past 7 days	Past 30 days	Absent
Yesterday's breakdown frequency	-0.049*** (0.009)			-0.020*** (0.004)			
Breakdown frequency in past 7 days		-0.072*** (0.013)			-0.035*** (0.006)		
Breakdown frequency in past 30 days			-0.071*** (0.014)			-0.034*** (0.006)	
Treatment							0.022*** (0.006)
Day Shift	-0.001 (0.002)	-0.001 (0.002)	-0.001 (0.002)	0.002 (0.002)	0.002 (0.002)	0.002 (0.002)	
Controls:							
Individual FE	YES	YES	YES	YES	YES	YES	YES
Product*Machine FE	YES	YES	YES	YES	YES	YES	YES
Day FE	YES	YES	YES	YES	YES	YES	YES
Observations	10,865	10,865	10,865	9,161	9,161	9,161	24,081

Notes: The dependent variable in columns (1) to (6) is over-target output percentage. Columns (1) to (3) report how output depends on machine's condition, measured by its breakdown frequency the previous day, in the past 7 days and past 30 days, respectively. Columns (4) to (6) report how output depends on yesterday's machine condition. In column (7), the dependent variable is a dummy variable indicating absence from work, and the main independent variable is treatment, namely, whether the worker experiences machine breakdown the previous day. All standard errors are clustered at the machine level. *** p<0.01, ** p<0.05, * p<0.1.

Table 3: Interpretation II, Increased Cautiousness

	(1)	(2)	(3)
	Quit	Hours	Machine <i>D</i>
Treatment	0.003 (0.002)	-0.028 (0.029)	-0.021*** (0.006)
Day Shift		-0.012 (0.017)	-0.001 (0.002)
Machine <i>D</i> breakdown in day <i>d-1</i>			-0.018*** (0.003)
Controls:			
Individual FE	YES	YES	YES
Product*Machine FE	YES	YES	YES
Day FE	YES	YES	YES
Observations	11,285	10,865	24,081

Notes: The dependent variables are: whether the worker quits the job (column 1), his working hours at day *d* (column 2) and over-target output percentage (column 3). The main independent variable is treatment, namely, whether the worker experiences machine breakdown the previous day. All standard errors are clustered at the machine level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4: Interpretation III, Proficiency Loss

	(1)	(2)	(3)
	Day d-2	Past 7 days	Past 30 days
Treatment	-0.029 (0.043)	-0.034*** (0.010)	-0.036*** (0.008)
Working hours in day d-2	0.004*** (0.001)		
Treatment*(work hour) _{d-2}	-0.000 (0.004)		
Working hours in past 7 days		0.000*** (0.000)	
Treatment*Working hours in past 7 days		0.000 (0.000)	
Working hours in past 30 days			0.000*** (0.000)
Treatment*Working hours in past 30 days			0.000 (0.000)
Day Shift	0.000 (0.002)	-0.000 (0.002)	-0.001 (0.002)
Controls:			
Individual FE	YES	YES	YES
Product*Machine FE	YES	YES	YES
Day FE	YES	YES	YES
Observations	9,660	11,285	11,285

Notes: The dependent variable is over-target output percentage. The main independent variable is treatment, namely, whether the worker experiences machine breakdown the previous day. Each column reports the estimates of the main treatment effect and its interaction term with the worker's cumulative working hours. All standard errors are clustered at the machine level. *** p<0.01, ** p<0.05, * p<0.1.

Table 5: Heterogeneity effect and Social Comparison

	(1)	(2)	(3)	(4)
	Male	Older worker	Local (Fujian)	Peer Effect
Treatment	-0.025*** (0.004)	-0.029*** (0.005)	-0.037*** (0.010)	-0.045*** (0.012)
Day Shift	-0.001 (0.002)	-0.001 (0.002)	-0.001 (0.002)	-0.001 (0.002)
Male*Treatment	-0.020 (0.015)			
Older worker*Treatment		-0.009 (0.011)		
Local (Fujian)*Treatment			0.010 (0.011)	
(Peer: High breakdown ratio)*Treatment				0.023* (0.013)
Controls:				
Individual FE	YES	YES	YES	YES
Product*Machine FE	YES	YES	YES	YES
Day FE	YES	YES	YES	YES
Observations	11,285	11,285	11,285	11,285

Notes: This table reports the results of heterogeneous effects. The dependent variable is over-target output percentage. The main independent variable is treatment, namely, whether the worker experiences machine breakdown the previous day. Columns 1 to 4 report the estimates of the main treatment effect and its interaction with gender, age, residency status and whether more peers experience machine breakdown, respectively. All standard errors are clustered at the machine level. *** p<0.01, ** p<0.05, * p<0.1.

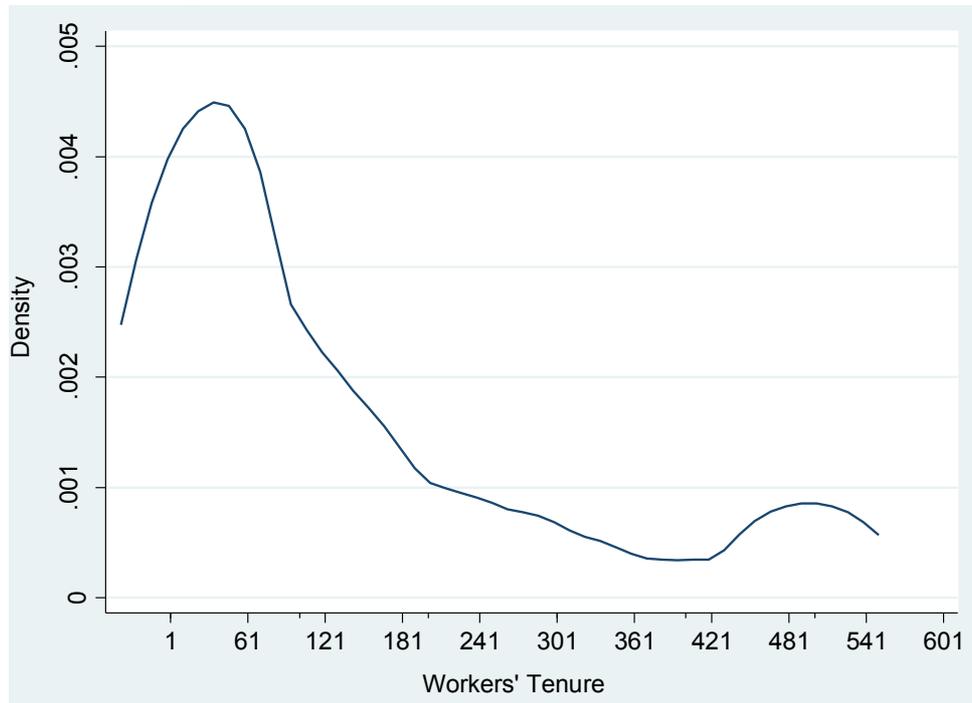
Table 6: Statistical Power, Prior Probability and Post-Study Probability

Statistical Power	Prior	Post-Study Probability (PSP)
0.8	0.01	0.447
0.8	0.25	0.964
0.8	0.5	0.988
0.5	0.01	0.336
0.5	0.25	0.943
0.5	0.5	0.980
0.2	0.01	0.168
0.2	0.25	0.870
0.2	0.5	0.952

Notes: The PSP is computed with the significant level of 0.01 using the method in Maniadis, Tufano, and List (2014).

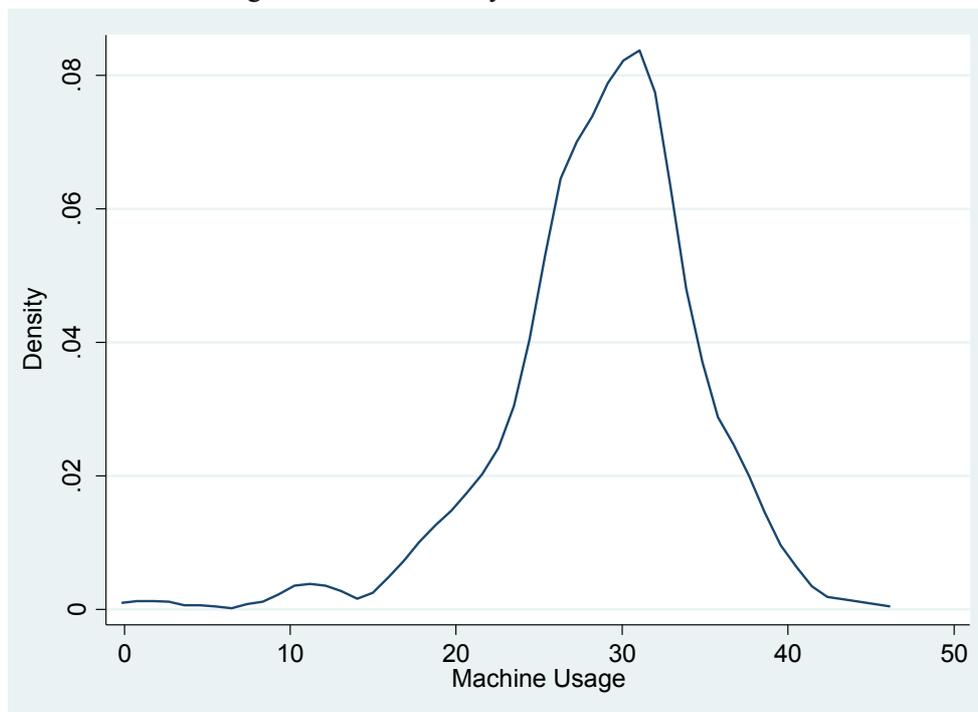
Appendix

Figure A1: The Density Distribution of Workers' Tenure



Notes: This figure plots the density distribution of workers' tenure (in days) during our sample period. For each worker, we calculate the length of time from the first to the last time that s/he is observed in the data, and therefore may include those non-workdays.

Figure A2: The Density of Machine Utilization



Notes: This figure plots the density distribution of the number of machines utilized for each workday during our sample period.

Table A1: Exclusion of the workers who work less than one month

	(1) Baseline	(2) With Day shift
Treatment	-0.036*** (0.009)	-0.033*** (0.006)
Day Shift		-0.001 (0.002)
Controls:		
Individual FE	YES	YES
Product*Machine FE	YES	YES
Day FE	YES	YES
Observations	17,221	11,285

Notes: The dependent variable is a worker's daily over-target output percentage. The main independent variable is treatment, namely, whether the worker experiences machine breakdown the previous day. All standard errors are clustered at the machine level. *** p<0.01, ** p<0.05, * p<0.1.

Table A2: Checks on the Nonrandom Assignment of Machine

	(1)	(2)	(3)
	Productivity on the previous day	Average Productivity in past 7 days	Average Productivity in past 30 days
Treatment	-0.009 (0.014)	-0.007 (0.014)	-0.008 (0.014)
Constant	0.108*** (0.011)	0.091*** (0.015)	0.089*** (0.016)
Observations	2,146	2,146	2,146

Notes: This table reports the estimates of the effect of machine breakdown on the quality of the machine that is assigned to the worker the following day. The dependent variables are the machine (m)'s average productivity on the previous day (column 1), its average productivity in past 7 days (column 2), and average productivity in past 30 days (column 3). The main independent variable is treatment, namely, whether the worker experiences machine breakdown the previous day. All standard errors are clustered at the machine level. *** p<0.01, ** p<0.05, * p<0.1.

Table A3: Alternative Estimation Framework

	(1)	(2)	(3)	(4)
	Single Difference	Single Difference	Difference-in- differences	Difference-in- differences
Treatment	-0.020** (0.010)	-0.020*** (0.005)	0.016 (0.015)	0.018 (0.015)
Age		0.001** (0.000)		
Local (Fujian) Resident		0.001 (0.008)		
Male		-0.016** (0.006)		
Day shift		-0.003 (0.006)		
Post*Treatment			-0.036* (0.018)	-0.036** (0.018)
Controls:				
Machine-specific Linear Trend	NOT	NOT	NOT	YES
Individual FE	NOT	NOT	YES	YES
Day FE	NOT	NOT	YES	YES
Observations	2,943	1,758	5,551	5,551

Notes: The dependent variable for column (1) and (2) is over-target output percentage. The dependent variable for column (3) and (4) is the difference in output between the considered workday and two days prior. The main independent variable is treatment, namely, whether the worker experiences machine breakdown the previous day. All standard errors are clustered at the machine level. *** p<0.01, ** p<0.05, * p<0.1.

Table A4: Further Robustness Checks

	(1)	(2)	(3)	(4)	(5)	(6)
	Log(Output)	Placebo	Persistent Effect	Breakdown probability	Correct for selection: lower bound	Correct for selection: upper bound
Panel A: Baseline analysis						
Treatment	-0.040*** (0.006)	-0.009 (0.005)	-0.027*** (0.003)	0.156*** (0.013)	-0.031*** (0.006)	-0.038*** (0.006)
Treatment in d-2			-0.019*** (0.003)			
Treatment in d-3			-0.010*** (0.002)			
Day Shift	-0.001 (0.003)	-0.001 (0.004)	0.001 (0.002)		-0.000 (0.002)	0.000 (0.002)
Controls:						
Individual FE	YES	YES	YES	NOT	YES	YES
Product*Machine FE	YES	YES	YES	YES	YES	YES
Day FE	YES	YES	YES	YES	YES	YES
Machine FE	NOT	NOT	NOT	YES	NOT	NOT
Observations	11,285	8,761	8,264	11,712	11,085	11,202
Panel B: Difference-in-Difference analysis						
Post*Treatment	-0.028 (0.017)	0.010 (0.011)	-0.036** (0.018)		-0.035* (0.019)	-0.044** (0.020)
Postd-2*Treatment			-0.056*** (0.020)			
Postd-3*Treatment			-0.019 (0.019)			
Controls:						
Individual FE	YES	YES	YES		YES	YES
Day FE	YES	YES	YES		YES	YES
Observations	5,551	4,070	10,087		5,525	5,442

Notes: The dependent variables are: the logarithm of output (column 1); over-target output percentage of the worker who operates the same machine in another shift on the same day as the concerned (treated) worker (column 2); over-target output percentage (columns 3, 5 and 6); a dummy variable indicating whether the machine m breaks down in day d (column 4). Panel A reports the estimates using baseline specification. The main independent variable is treatment, namely, whether the worker experiences machine breakdown the previous day. Panel B reports the estimates using DD specification. The main independent variable is Post*Treatment. All standard errors are clustered at the machine level. *** p<0.01, ** p<0.05, * p<0.1.