

The Hidden Cost of Worker Turnover: Attributing Product Reliability to the Turnover of Factory Workers

Ken Moon, Prashant Loyalka, Patrick Bergemann, Joshua Cohen

Abstract

Product reliability is a long-standing concern among product manufacturers, and researchers have identified a number of factors that significantly affect product reliability. In this paper, we examine a previously under-recognized yet impactful determinant of product reliability: a high rate of turnover on the manufacturing line when the product was assembled. Even when assembly lines are designed to minimize product defect rates through simplified tasks and stringent quality control tests, quality variations may exist in the manufactured units that are only revealed after they have been used repeatedly. If this is the case, then the disruptiveness of high turnover may directly lead to product reliability issues. To evaluate this possibility, we link four post-production years of field failures for tens of millions of consumer mobile devices back to their production lines. After controlling for known factors affecting reliability (workloads, learning, and component quality), we find that the likelihood of field failure increases by 7-8% when a device is produced in the monthly high-turnover weeks following paydays; and that even in other weeks, product reliability responds significantly (2-3%) to the individual assembly lines' weekly turnover rates. Together, we demonstrate that staffing and retaining a stable factory workforce critically underlies product reliability, and we show the value of connected field data in informing manufacturing operations. Ultimately, products are more reliable when workers are more reliable.

Keywords: Data-driven workforce planning, Empirical operations management, Employee turnover, People operations, Product quality, Productivity, Quality management, Supply chain management

Working Paper

April 2020

reap.fsi.stanford.edu



The Hidden Cost of Worker Turnover: Attributing Product Reliability to the Turnover of Factory Workers*

Ken Moon^a, Prashant Loyalka^b, Patrick Bergemann^c, Joshua Cohen^d

a: The Wharton School; b: Stanford Graduate School of Education;

c: Chicago Booth School of Business; d: Apple University

Product reliability is a long-standing concern among product manufacturers, and researchers have identified a number of factors that significantly affect product reliability. In this paper, we examine a previously under-recognized yet impactful determinant of product reliability: a high rate of turnover on the manufacturing line when the product was assembled. Even when assembly lines are designed to minimize product defect rates through simplified tasks and stringent quality control tests, quality variations may exist in the manufactured units that are only revealed after they have been used repeatedly. If this is the case, then the disruptiveness of high turnover may directly lead to product reliability issues. To evaluate this possibility, we link four post-production years of field failures for tens of millions of consumer mobile devices back to their production lines. After controlling for known factors affecting reliability (workloads, learning, and component quality), we find that the likelihood of field failure increases by 7-8% when a device is produced in the monthly high-turnover weeks following paydays; and that even in other weeks, product reliability responds significantly (2-3%) to the individual assembly lines' weekly turnover rates. Together, we demonstrate that staffing and retaining a stable factory workforce critically underlies product reliability, and we show the value of connected field data in informing manufacturing operations. Ultimately, products are more reliable when workers are more reliable.

Key words: Data-driven workforce planning; Empirical operations management; Employee turnover; People operations; Product quality; Productivity; Quality management; Supply chain management.

1. Introduction

Product reliability—the extent to which a product fulfills its intended function over a particular period of time—is a critical concern for product manufacturers. Unreliable products can harm firms in a variety of ways, such as through negative online reviews that decrease demand, warranties or

*We thank Hamsa Bastani, Morris Cohen, Marshall Fisher, Santiago Gallino, Serguei Netessine, Xuanming Su, and our data sponsor's data science team for helpful comments and discussions. We gratefully acknowledge support from our data sponsor and from the Wharton Dean's Research Fund and Wharton Global Initiatives (Moon). Correspondence: kenmoon@wharton.upenn.edu

replacement fees that drain profits, or recalls that damage a firm's reputation and financials. The cost of unreliable products can be enormous: global carmakers spent \$46 billion in warranties in 2018, while the 2016 recall of the Galaxy Note 7 smartphone cost Samsung \$5.3 billion ([Warrantyweek.com 2019](#), [Tsukayama 2018](#)).

Because of these costs, manufacturing firms seek to control product defect rates in a variety of ways, ranging from quality management systems and standards ([Hendricks and Singhal 1997](#), [Corbett et al. 2002, 2005](#)) to the selection of product components ([Ramdas and Randall 2008](#)). On the assembly lines, manufacturers simplify tasks to soften learning curves and install frequent testing stations to ensure product quality. Yet even standardized tasks carried out under stringent testing standards allow some variability in the quality of manufactured units that is revealed only after they are used repeatedly. Such variability becomes visible in field failures, which are the instances in which the product can no longer be used as intended until repaired or replaced.

One factor not previously identified as significantly affecting product reliability is the level of worker turnover in the manufacturing process during a particular unit's production. Worker turnover is prevalent in manufacturing, yet workforce experience and stability may deliver considerable value in maintaining the quality of output, even for units that pass quality control standards. If this is the case, then product reliability may constitute an undetected and previously unrecognized cost of worker turnover.

In this paper, we evaluate this proposition by tracking nearly fifty million consumer mobile electronic devices for field failures (i.e., repairs and replacements) over four years of consumer use, and tracing each device (whether failed or not) back to the factory conditions under which it was produced. Combining field data and production data, we find that the product's field failure rates strikingly synchronize with factory worker turnover rates, as these rates rise and fall over the course of the manufacturer's pay cycle. Empirically, field failure rates peak when devices are produced in the high-turnover week immediately following payday, becoming 10.2% more likely to fail than when they are assembled in the lowest turnover week before payday.

We validate that the product’s field failure rates correspond closely to the recurring episodes of factory worker quits, even after we control for known factors affecting product reliability, including the firm’s organizational learning, the quality of key components, and the assembly lines’ workloads. After controlling for these factors, we estimate that the post-payday week’s peak in worker turnover increases the incidence of field failures by 6.6% to 7.7% (i.e., a device produced then is 6.6% to 7.7% more likely to fail than in other weeks). Even in other weeks, the mean level of worker turnover boosts field failures by an estimated 1.7% to 3.2%. Together, these two effects approach the 10.2% gap in the product’s field failure rate between the pay cycle’s peak and trough.

We view our work as making several contributions. First, despite extensive literatures on both product reliability and worker turnover, to our knowledge no previous study has linked the two together. The strong effects we estimate underscore the importance of staffing and retaining a stable factory workforce in producing reliable products. Second, this work has clear implications for manufacturers seeking to improve the reliability of their products, suggesting that current strategies intended to mitigate the effects of worker turnover may not be as effective as more direct efforts at retention and workforce stability. Finally, we demonstrate the value in combining field data with factory operations. Such data are increasingly available (Porter and Heppelmann 2015), and combining them can help to reveal critical linkages between factory and field.

1.1. Related Literature

Product quality is an important competitive factor in consumer product markets, including for automobiles, consumer technology goods, agricultural goods, and apparel. Consequently, firms adopting quality management systems and standards have reaped significant financial rewards (Hendricks and Singhal 1997, Corbett et al. 2002, 2005). We focus on product reliability, which has grown in salience among quality’s core components. In particular, online technology enables consumers’ critical word of mouth to quickly spread through reviews and social media (e.g., 26% of surveyed consumers used social media to complain about a company or product (Srinivasan and Kurey 2014)) and reduces the costs of collecting data about reliability. As a strategic priority,

product reliability drives firms to respond opportunistically to revelations about the reliability of their competitors' products (Ball et al. 2019) and to calibrate the reliability of their own products to fit their shareholders' financial interests (Kini et al. 2016).

As a result, the empirical operations literature has extensively studied product reliability. By examining product recalls, researchers have shown that manufactured product quality depends on how manufacturers utilize plant capacity and over what product variety (Shah et al. 2016), their proximity-based relationships with key suppliers (Bray et al. 2019), how they trade off components' specific product fit against cross-product familiarity (Ramdas and Randall 2008), and even inspectors' behavioral biases (Ball et al. 2017, Ibanez and Toffel 2019). Delivering on quality extends past the point of sale: consumers significantly value the provision of product maintenance and repair services (Guajardo et al. 2015) and account for timely after-sales services when choosing whether to adopt and recommend technology-intensive goods (Kundu and Ramdas 2019). Appropriately prioritizing across such maintenance services can be critical (Deshpande et al. 2003). Finally, data capturing equipment-level field failures, similar in many respects to our own, have been coupled with maintenance service records to study the impact of performance-based contracting on post-sale maintenance servicing and product reliability (Guajardo et al. 2012, Chan et al. 2018). By uniquely linking a product's field data (over millions of devices) to its factory-based production data, our study complements existing empirical research that examines internal quality metrics and yields (Fisher and Ittner 1999, Moon et al. 2019) or compares field reliability across differently produced products and components (Ramdas and Randall 2008, Ball et al. 2017). To our knowledge, we are the first to empirically associate product reliability with factory workforce turnover.

Our research also adds to prior work addressing the relationship between worker productivity and turnover in operations, including Arthur (1994) and Moon et al. (2019) in manufacturing, Emadi and Staats (2019) and Musalem et al. (2019) in call centers, Ton and Huckman (2008) in retailing, and Song and Huckman (2018) in teaching hospitals. Batt et al. (2019) cover the productivity effects of a phenomenon resembling turnover in healthcare patient handoffs between

work shifts. In estimating turnover’s effect, controlling for workload and utilization effects (e.g., [KC and Terwiesch \(2009\)](#), [Shah et al. \(2016\)](#)) will be important in our study.

Section 2 introduces our data and defines variables. In Section 3, we present our empirical analysis and findings. Section 4 concludes with a brief discussion of implications.

2. Data

Our study is carried out in collaboration with a major consumer technology goods producer (the “firm”) and one of its China-based contract manufacturers. Identifying details are omitted for confidentiality. For the product we study, the firm provided data tracking field failures for nearly 50 million finished devices and data recording weekly production workloads and backlogs and quarterly fallout rates for the product’s key components. The contract manufacturer provided the human resources data required to observe when and from which assembly lines the factory’s workers quit.

Like [Ramdas and Randall \(2008\)](#), [Guajardo et al. \(2012\)](#), and [Chan et al. \(2018\)](#), our reliability data are sourced from the field performance of manufactured products. All the studied devices passed stringent quality control testing (on both their inputs and their station-by-station assembly) before shipping to customers. The firm is able to effectively track device repairs through its program for warranties and authorized repairs. Because the device model is a premium product, customers typically seek repairs using authorized parts through a network of affiliated US retail locations or by shipping defective devices to the firm with round-trip shipping paid. Warranties cover all manufacturing defects and can be extended to cover three years instead of one. Under certain conditions, the firm replaces the defective unit with a new device before the former is shipped by the customer. Generally, we expect that customers do seek to repair defective units, because we study a premium product typically used nearly daily.

Our main production data are obtained from the firm. Production plans provide us with weekly backlogs, which are the production shortfalls measured against the orders the firm generates from its rolling forecasts of end-sales demand, and our production volume data tracks the weekly units passing through individual assembly lines’ benchmark stations. The firm also provides the quarterly

fallout rates for six key components: the device’s battery, display, housing, circuit board, rear camera, and front camera. Lastly, the contract manufacturer’s human resources staffing data track factory worker quits occurring on specific assembly lines and weeks.

We organize our panel dataset around the assembly line-week as the unit of analysis. Critically, we trace consumer-held devices back to their respective line-weeks of origin and thus are able to attribute a field failure rate to each line-week observation. All told, the data cover 40 weeks (September 2014 through June 2015) for 44 assembly lines dedicated to the product and located within three expansive buildings of a common facility in China. The third building operates for only 31 of the 40 weeks, with some assembly lines shutting down sooner. Thus, we are left with 1516 assembly line-weeks in total. Because operations and worker turnover trends are irregular during the Chinese New Year holiday, we omit February 2015 from our regression analyses but include it in descriptive plots. We additionally confirmed with firm experts that no atypical production-related events were recorded for the product during the studied time period.

2.1. Outcome Variable

2.1.1. Field failure rate. As our study’s outcome variable, we examine the fraction of produced devices that failed in the field. A device is deemed to have failed in the field the first time it receives a known service repair or replacement, at which point we call the device a field failure. For our studied devices produced in September 2014 through June 2015, we observe all field failures as of July 6, 2019. Thus, for each assembly line-week, we attribute a field failure rate, defined as the fraction of field failures (out of its units produced) arising through July 6, 2019. Because each device has been observed in the field for at least four years post-production, field failure rates should capture virtually all product reliability variation originating in the factory assembly process.

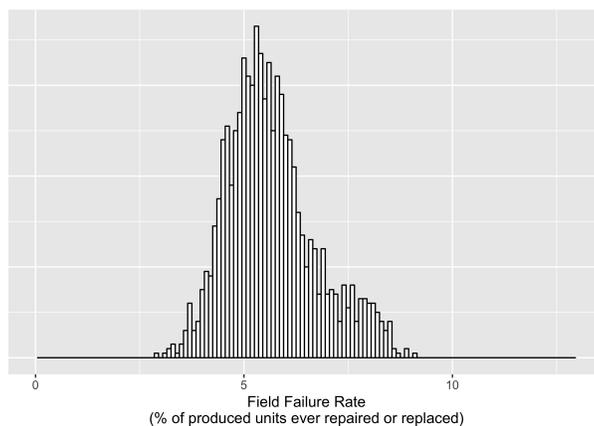
For reasons of confidentiality and data sensitivity, we multiply the field failure rate by a scalar factor between 0.125 and 8, including in the data we provide as the paper’s data supplement. The true field failure rate’s mean is then a percentage between 1% and 45%, and its coefficient of variation is 0.19. The transformed field failure rate’s mean and median are given in Table 1, and its histogram is shown in Figure 2. Our regression results use log field failure rate as the dependent variable, thus the estimated coefficients are not affected by applying the constant factor.

Table 1 Weekly Assembly Line Statistics

	Mean	Median
Field failure rate [†]	5.6%	5.5%
Number of workers staffed	301	259
Worker turnover	5.1%	3.2%
Workload in units assembled	30,152	31,754

Unbalanced panel of 44 assembly lines observed in 1,516 production-active line-weeks in Sept. 2014 – Jun. 2015.

[†] The true field failure rate is multiplied by a scalar between 0.125 and 8.

Figure 2 Histogram of Field Failure Rates

Field failure rates for 1516 assembly line-weeks between September 2014 and June 2015. Field failure rates are the percentages of produced units ever receiving a service repair or replacement through July 6, 2019.

2.2. Study Variables

2.2.1. Measure of worker turnover. We measure worker turnover as the fraction of workers on an assembly line who quit the factory during a given week.¹ The contract manufacturer’s data contain the work assignments of 52,214 workers actively staffing the studied assembly lines during the study. By matching these data with the dates on which workers quit the factory, we attribute quits to assembly line-weeks. To arrive at worker turnover, we divide these quits’ count by the number of workers staffing the assembly line at the start of the same week. Mean and median worker turnover are 5.1% and 3.2% per week, respectively, evidencing an upward skew. The mean assembly line headcount is 301, and staffing is relatively steady as new workers replace quits.

¹ We do not focus on the additional movements of workers between assembly lines or between stations on a given assembly line.

2.2.2. Pay cycle. The manufacturer’s monthly pay cycle strongly shapes worker quits. Because workers value the compensation they receive on payday, assembly line worker turnover averages 2.9% in the week immediately prior to payweek and 8.9% in the week immediately following payweek. We define the weeks to payweek variable, which ranges from a maximum of four (weeks left to the next payweek) down to a minimum of zero (during payweek). Much as with cohort turnover (e.g., [Song and Huckman \(2018\)](#)), managers fully anticipate the pay cycle effect, in which a wave of worker quits causes the replacement of experienced workers by new workers.

2.3. Control Variables

2.3.1. Assembly line workloads. Existing literature supports the existence of workload and utilization effects on worker productivity in a variety of settings (e.g., [KC and Terwiesch \(2009\)](#), [Shah et al. \(2016\)](#)). We measure an assembly line’s weekly workload in device units assembled using the units counted at benchmark stations. The assembled output averages 30,152 units per assembly line-week for a total of 45.7 million units assembled over the data’s 1516 observations. The coefficient of variation in workloads is 0.28.

As a subtlety, because the firm actively adjusts workloads in response to turnover, workloads are not pre-determined controls. In fact, workloads likely mediate worker turnover’s effect on production outcomes. Appendix [A](#) accordingly explains how we use workloads as a mediating control variable. Importantly, our regressions in Section [3.2](#) will consistently estimate the modeled effects of post-payday exit and worker turnover, after excluding any effects arising specifically because of workloads adjusting in response to turnover.

2.3.2. Production backlog. A larger production backlog can affect the urgency of production. The backlog is calculated as the positive shortfall, if any, of the week’s projected production against device orders. As such, it also captures any projected shortages in components supplied.

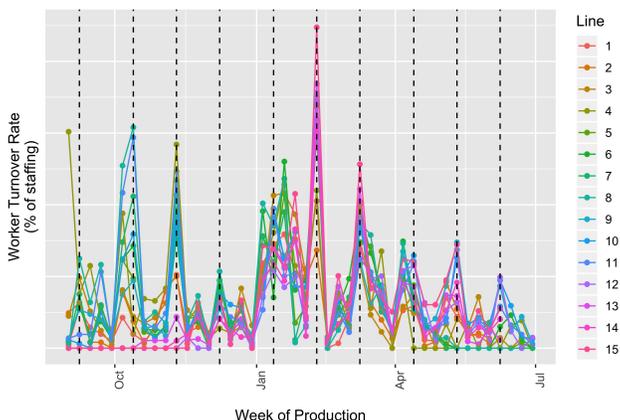
2.3.3. Component fallout rates. Many field failures attribute to the natural depreciation of the product’s functional components (e.g., [Ramdas and Randall \(2008\)](#), [Colak and Bray \(2016\)](#)), where components are designed to last the duration of a product’s intended lifetime or for regular

maintenance and replacement. However, component failures can be accelerated by small problems in installing the component during assembly or by existing defects in the component. To measure pre-existing defect rates for key components, our study tracks their fallout rates, which are their rates of detected failure during manufacturing. We obtain fallout rates for the device’s battery, display, housing, circuit board, front camera, and rear camera for each three-month quarter.

3. Analysis

3.1. Empirical Motivation

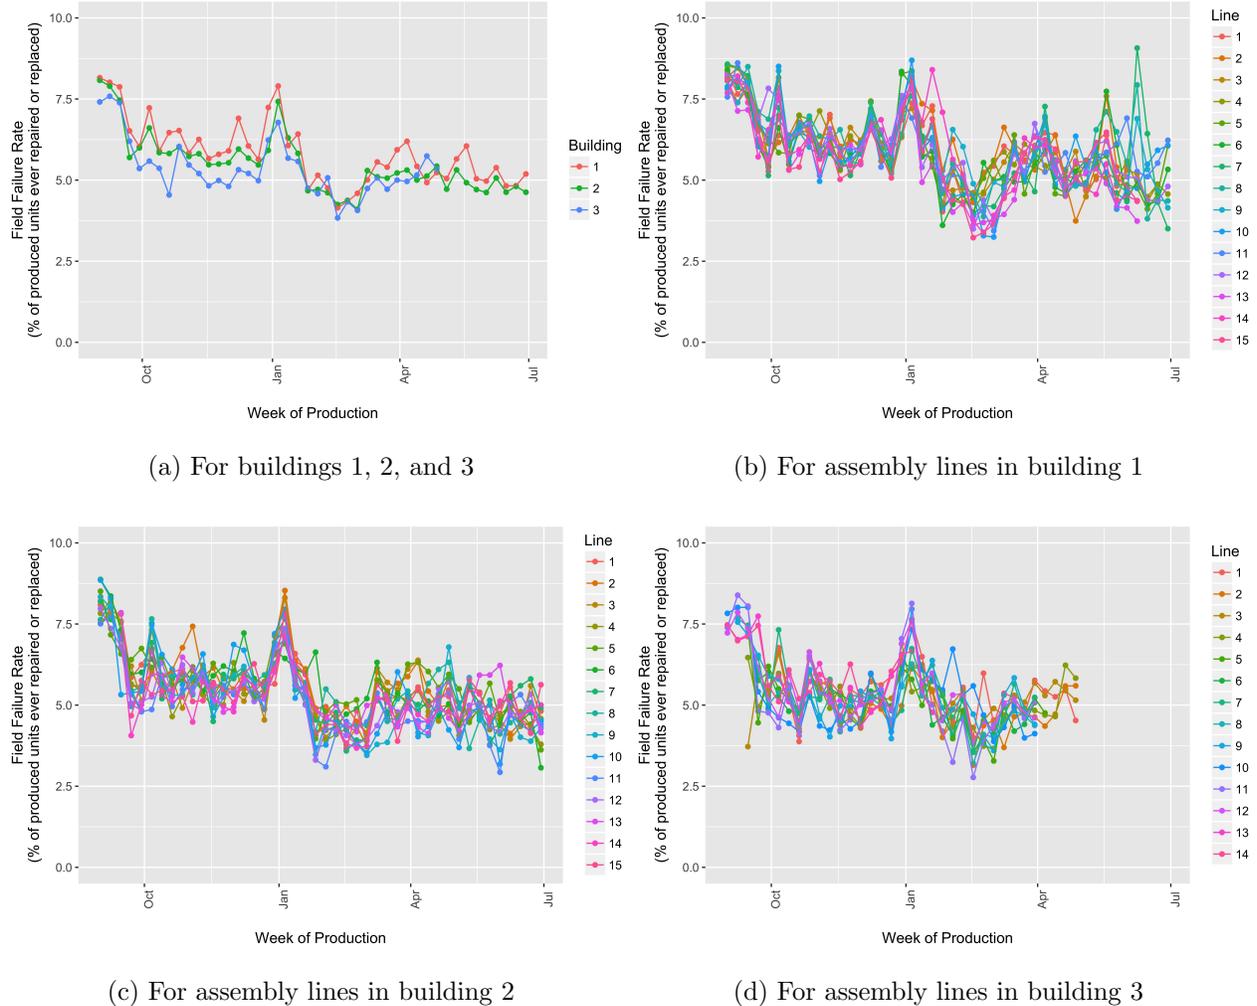
Figure 3 Assembly Line Worker Turnover in Building 1 (Sept. 2014 to June 2015)



Weekly worker turnover rates for assembly lines in Building 1 between September 2014 and June 2015. Dashed vertical lines correspond to the sample’s post-payday weeks.

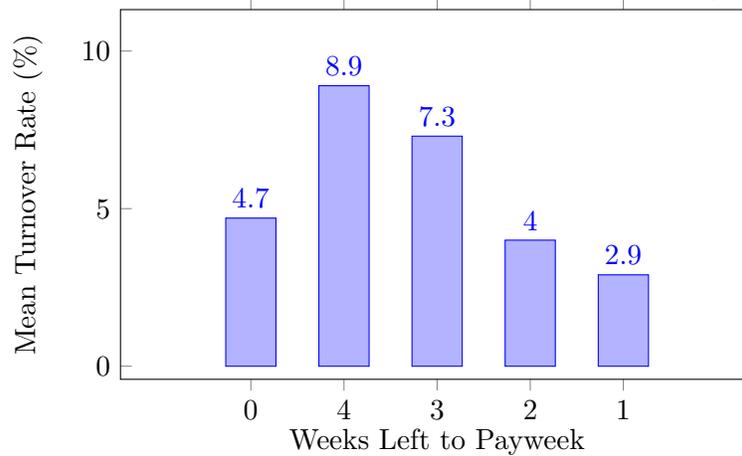
Our analysis considers a chain of events flowing from two key dynamics. First, workers postpone their departures until being compensated, causing workers’ quits to peak significantly following the manufacturer’s payday. Thus, factory worker turnover follows a highly predictable pattern characterized by monthly peaks. Representatively using the assembly lines in Building 1, Figure 3 illustrates worker turnover’s resulting, cyclical variation over time. After each payday, experienced workers exit and new workers are recruited to replenish the ranks.

Second, suppose and consider that worker turnover sharply impacts product reliability. As a result, field failures manifest a stark pattern resembling an electrocardiogram: the product’s field failure rate (when marked by date of production) ricochets through regular, monthly cycles of peaks

Figure 4 Rates of Device Field Failures Attributed to Their Weeks of Production (Sept. 2014 to June 2015)

For the product's 45.7 million units assembled in the studied facility between September 2014 and June 2015, we tallied for each assembly line-week the numbers of units produced that were repaired or replaced and those that were not. All units were tracked for repairs through July 6, 2019, i.e., at least four years post-production. Field failure rates are the percentages of produced units ever receiving a service repair or replacement during the tracked period.

and troughs. Figure 4 reveals the data's actual temporal pattern of field failure rates, shown both for the three buildings (Figure 4a) and for the individual assembly lines in each building (Figures 4b, 4c, 4d). Not only does each building and assembly line transparently exhibit a cyclically peaking trend, the assembly lines' field failure rates co-vary to a striking degree (nearly 41% of variation is explained by the weekly averages), which suggests a common cause behind the pattern of peaking failures.

Figure 5 Field Failure Rates and Worker Turnover Rates over the Manufacturer's Pay Cycle

(a) Assembly lines' mean weekly turnover rates (% headcount) by week in the pay cycle



(b) Assembly lines' mean weekly field failure rates (% units produced) by week in the pay cycle

Descriptive evidence, which we corroborate using statistical tests, is consistent with worker turnover being a common cause. Figure 5 illustrates the empirical correspondence between worker turnover (Figure 5a) and field failure rates (Figure 5b) over the pay cycle. To reiterate, mean worker quits peak at 8.9% post-payweek and bottom out at 2.9% immediately pre-payweek. Correspondingly, mean field failure rates are proportionally 10.2% higher at peak over trough (i.e., a device produced at peak is 10.2% more likely to fail) (Figure 5b). In terms of the underlying statistical relationship between field failure rates and worker turnover rates, their weekly averages exhibit a positive correlation of 0.14.

We formally test whether the post-payday weeks' field failure rates stochastically dominate the failure rates during pre-payday weeks (Table 6). We reject the null hypothesis of a common distribution at the statistically significant, exact p-value of 0.01% (i.e., 0.0001). Thus, a consumer who buys a device produced in a peak-turnover week faces significantly higher odds of needing a service repair or replacement than if she had bought an identical shelf unit *produced* a week earlier.

Table 6 Rank-sum Test for Field Failure Rates Post-payday Stochastically Dominating Pre-payday

Using pre-payday weeks as test statistic's subsample		
Rank sum	Subsample size	P-value
107,673	343 out of 681	0.00010

We carry out a Mann-Whitney-Wilcoxon rank-sum test of the null hypothesis that pre-payday and post-payday weeks' field failure rates share a common distribution, against the one-sided alternative hypothesis that post-payday weeks' field failure rates stochastically dominate those of pre-payday weeks. The p-value is calculated using 100,000 randomly drawn rank permutations. Ranks are assigned by field failure rate from lowest to highest, and ties are randomly broken.

However, we cannot attribute the periodical escalation of field failure rates to worker turnover without ruling out several alternative explanations. Generally, we expect gains in product reliability as the firm accrues assembly experience over time, e.g., about the best way to handle and install a given component (Ramdas and Randall 2008). Thus, a device produced later in time is typically more reliable. Relatedly, components' incoming defect rates tend to fall over time. Visually, there is some indication of organizational learning or component improvements, as field failure rates appear to mildly decline over time in Figure 4. While such factors are influential, they represent largely non-cyclical trends.

Based on the literature regarding worker productivity, we identify two main factors plausibly affecting the quality of work hence of the end product we study. The first factor is assembly line workload, where it is known that elevated workloads and utilization can affect reliability (e.g., Shah et al. (2016)). The second is order backlog, representing the urgency of work. On one hand, it could be argued that when managers adjust workloads in response to post-payday worker turnover, the consequences are part of worker turnover's downstream effects. By controlling for workloads, our

intention is to estimate worker turnover’s effect on field failures occurring outside of a response in workloads. We next validate the empirical relationship between worker turnover and field failures through a series of regression analyses.

3.2. Regression Analyses

We estimate a regression model of worker turnover’s effects on field failure rates. We keep the regression analyses straightforward. The analysis benefits from relatively comprehensive controls for the known factors affecting product reliability. In particular, we closely control for learning, component quality, and workloads. Moreover, both the control variables and the main independent variables (worker turnover and the pay cycle) cleanly pre-date any knowledge of the outcome variable (the subsequent field failures).

3.2.1. Specification. We specify the regression model in terms of causal effects for two reasons. First, one control variable, assembly line workloads, is adjusted by managers *in response* to worker turnover. Adopting a causal framework enables us to more rigorously handle a control variable that also mediates the effect of worker turnover on field failures (see Appendix A). Second, even after controlling for workloads, we expect two types of turnover effects. We expect a general effect: that for any given week, an assembly line that exhibits comparatively higher worker turnover tends to later manifest a higher field failure rate. In contrast, we expect managers to expect and address the waves of high turnover known to arrive in the weeks that follow paydays. Thus, we more flexibly specify worker turnover’s effect in these weeks. A causal model lets us separately explain each effect. However, the model can also be viewed as describing the association between turnover and field failure rates after ruling out controls.

We carry out regressions of the form:

$$\begin{aligned} \log(\text{Field failure rate}_{lt}) = & \alpha_l + \beta_0 \cdot (\text{Post-payday week})_t + \beta_1 \cdot (\text{Worker turnover})_{lt} \\ & + \beta_2 \cdot (\text{Worker turnover if post-payday week})_{lt} + \gamma \cdot X_{lt} + \epsilon_{lt}, \quad (1) \end{aligned}$$

where X_{lt} are the control variables for assembly line l in week t . For weeks other than the post-payday week, the coefficient β_1 measures the general effect of an assembly line’s worker turnover. In

contrast, we devote special attention to post-payday weeks, when managers anticipate that worker quits will peak (Figure 3). For those weeks, we more flexibly specify the effect on the field failure rate: β_0 measures a constant disruptive effect that does not depend on an assembly line’s exact magnitude of worker turnover. Like β_1 for other weeks, β_2 allows the effect’s size to depend on the line’s degree of worker turnover. In particular, we permit the possibility that managers’ mitigating actions result in $\beta_2 < \beta_1$ (that is, a more stable effect) in the post-payday week. We may similarly find that β_0 is large yet $\beta_2 < \beta_1$ if the post-payday week’s disruption results from a general shortage of supervisor attention. We also do not rule out that $\beta_2 = \beta_1$ with β_0 capturing any additional disruption in the post-payday week. Finally, to account for differences in reliability across assembly lines, we present results both with and without the assembly line fixed effects α_i .

In summary, β_0 and β_2 closely characterize high worker turnover’s disruptive effect on product reliability in the post-payday weeks, whereas β_1 characterizes a potentially corroborating effect from assembly lines’ variation in worker turnover in other weeks. For a conservative view on causality, we construe these same coefficients as validating an empirical relationship between worker turnover and field failure rates after controlling for X_{it} .

3.2.2. Results. Table 7 presents our regression results. In Table 7a, we control for a time trend representing the organizational learning that improves the reliability of manufactured units over time. Specification (2) includes assembly line fixed effects, whereas Specification (1) permits a common intercept as the baseline field failure rate. In 7b, we additionally control for assembly line-level workloads and the manufacturer’s backlog for the product. Lastly, Table 7c adds component fallout rates as controls. In Appendix B, we replicate the regression analyses substituting month fixed effects in place of the time trend.

The results are fairly consistent across the regressions. In Tables 7a and 7b, process learning is highly significant and reduces the incidence of field failures proportionally by an estimated 0.7-0.8% weekly (e.g., shaving 0.04% weekly off the median field failure rate of 5.5%). In Table 7c, estimated learning is similar or slightly smaller. Our results are virtually unchanged when we substitute cumulative units assembled to capture process learning.

Table 7 Regressing Field Failure Rates on Post-payday Disruption and Worker Turnover
(a) Log field failure rate as dependent variable

	(1)	(2)
Worker turnover $\widehat{\beta}_1$	0.417*** (0.101)	0.337*** (0.099)
Post-payday week (dummy) $\widehat{\beta}_0$	0.069*** (0.015)	0.064*** (0.014)
Worker turnover in post-payday week $\widehat{\beta}_2$	-0.213* (0.108)	-0.213* (0.104)
Time trend in weeks	-0.007*** (0.000)	-0.008*** (0.000)
Line fixed effects		Y
R^2	0.27	0.33
Significance levels \rightarrow	*** - 0.001	** - 0.01
	* - 0.05	

Unbalanced panel of 1,516 production-active line-weeks in Sept. 2014 – Jun. 2015. Specification (2) includes fixed effects for 44 assembly lines.

(b) Log field failure rate as dependent variable, adding workload-related controls

	(1)	(2)	(3)	(4)
Worker turnover $\widehat{\beta}_1$	0.416*** (0.102)	0.345*** (0.099)	0.434*** (0.102)	0.353*** (0.100)
Post-payday week (dummy) $\widehat{\beta}_0$	0.069*** (0.015)	0.064*** (0.014)	0.068*** (0.015)	0.064*** (0.014)
Worker turnover in post-payday week $\widehat{\beta}_2$	-0.213* (0.108)	-0.211* (0.104)	-0.194 (0.108)	-0.203 (0.104)
Time trend in weeks	-0.007*** (0.000)	-0.008*** (0.000)	-0.008*** (0.000)	-0.008*** (0.000)
Workload (line's log output assembled)	0.001 (0.011)	-0.007 (0.011)	0.001 (0.011)	-0.007 (0.011)
Log backlogged orders+1			-0.002 (0.001)	-0.001 (0.001)
Line fixed effects		Y		Y
R^2	0.27	0.33	0.27	0.33
Significance levels \rightarrow	*** - 0.001	** - 0.01	* - 0.05	

Unbalanced panel of 1,516 production-active line-weeks in Sept. 2014 – Jun. 2015. Specifications (1) and (3) are pooled regressions, whereas (2) and (4) include fixed effects for 44 assembly lines. Backlogged orders are the week's projected backlog calculated as the excess of (A) the firm's reported production targets based on demand forecasts and sell-through (effectively, orders) over (B) clear to build (projected production).

We find robust evidence suggesting that post-payday worker turnover impacts field failure rates. Under all specifications, our estimates of payweek's constant effect are highly significant (p-values 0.001 or smaller). Across the specifications in Tables 7a and 7b, they project that field failure incidence increases proportionally anywhere from 6.6% to 7.1%, and those in Table 7c project

Table 7 Regressing Field Failure Rates on Post-payday Disruption and Worker Turnover (continued)
(c) Log field failure rate as dependent variable, adding component fallout during manufacturing

Control for quarterly fallout rate of:	Battery		Display		Housing	
	(1)	(2)	(3)	(4)	(5)	(6)
Worker turnover $\widehat{\beta}_1$	0.482*** (0.100)	0.405*** (0.098)	0.500*** (0.100)	0.421*** (0.098)	0.617*** (0.104)	0.512*** (0.103)
Post-payday week (dummy) $\widehat{\beta}_0$	0.069*** (0.014)	0.065*** (0.014)	0.070*** (0.014)	0.065*** (0.014)	0.074*** (0.014)	0.068*** (0.014)
Worker turnover in post-payday week $\widehat{\beta}_2$	-0.152 (0.100)	-0.157 (0.102)	-0.148 (0.106)	-0.152 (0.102)	-0.126 (0.107)	-0.133 (0.104)
Time trend in weeks	-0.006*** (0.000)	-0.007*** (0.000)	-0.006*** (0.000)	-0.007*** (0.000)	-0.004*** (0.001)	-0.006*** (0.001)
Workload (line's log output assembled)	0.022 (0.011)	0.012 (0.011)	0.022 (0.011)	0.012 (0.011)	0.018 (0.012)	0.007 (0.011)
Log fallout rate	0.678*** (0.082)	0.558*** (0.078)	0.167*** (0.020)	0.136*** (0.019)	0.336*** (0.047)	0.224*** (0.045)
Line fixed effects		Y		Y		Y
R^2	0.30	0.35	0.30	0.35	0.30	0.34
Control for quarterly fallout rate of:	Circuit board		Rear camera		Front camera	
	(7)	(8)	(9)	(10)	(11)	(12)
Worker turnover $\widehat{\beta}_1$	0.607*** (0.110)	0.470*** (0.110)	0.459*** (0.100)	0.385*** (0.098)	0.588*** (0.103)	0.493*** (0.101)
Post-payday week (dummy) $\widehat{\beta}_0$	0.074*** (0.015)	0.067*** (0.014)	0.069*** (0.014)	0.064*** (0.013)	0.073*** (0.014)	0.067*** (0.014)
Worker turnover in post-payday week $\widehat{\beta}_2$	-0.148 (0.108)	-0.163 (0.105)	-0.159 (0.106)	-0.163 (0.102)	-0.129 (0.107)	-0.135 (0.103)
Time trend in weeks	-0.005*** (0.001)	-0.006*** (0.001)	-0.006*** (0.000)	-0.007*** (0.000)	-0.005*** (0.000)	-0.006*** (0.000)
Workload (line's log output assembled)	0.009 (0.012)	-0.001 (0.011)	0.022 (0.011)	0.012 (0.011)	0.020 (0.012)	0.009 (0.011)
Log fallout rate	0.228*** (0.052)	0.135** (0.051)	0.185*** (0.022)	0.154*** (0.021)	0.742*** (0.099)	0.576*** (0.095)
Line fixed effects		Y		Y		Y
R^2	0.28	0.33	0.30	0.35	0.30	0.34
Significance levels \rightarrow *** - 0.001	** - 0.01	* - 0.05				

Unbalanced panel of 1,516 production-active line-weeks in Sept. 2014 – Jun. 2015. Odd-numbered specifications are pooled regressions, whereas even-numbered specifications include fixed effects for 44 assembly lines.

proportional increases of 6.6% to 7.7%. On the other hand, we do not find that the size of the post-payday week's disruption depends sensitively on the exact magnitude of worker turnover. In Table 7c where we control for learning, workloads, and component fallout, $\widehat{\beta}_2$ is not significant. Across the specifications in Tables 7a and 7b, the mean level of worker turnover (5.1%) reduces field failures proportionally by 0.98% to 1.08%.

In weeks other than post-payday, we estimate a highly statistically significant effect of worker turnover that is moderately significant in size. Over the specifications in Tables 7a and 7b, the effect of mean worker turnover (5.1%) on the incidence of field failures is a proportional increase ranging from 1.7% to 2.2%, while Table 7c's results range from 2.0% to 3.2% for the same. Additionally, we can apply the foregoing estimated effect to the difference between average worker turnover and the lower turnover in pre-payday weeks. Combined with the estimated post-payday spike in field failure rate, we account for between 7.4% to 9.1% (across specifications) of the 10.2% gap that empirically separates the pay cycle's peak from its trough.

Lastly, we find no indication of a significant impact from workloads or backlogs in our setting. In particular, if we were to find significantly positive coefficients on workloads, it would suggest that managers could mitigate the outcome variable (field failures) by reducing workloads, but we find no support for such a capability: no statistical significance is found by any of our regressions. Moreover, even if managers were powerless to mitigate field failures using workloads, given any ability to foresee field failure rates, they could still reduce the overall number of field failures by routing more work to better-performing assembly lines. Such behavior would result in a negative coefficient estimate for workloads, and we find no evidence of such foresight.

4. Concluding Remarks

In this paper, we highlight the role of a stable factory workforce in assuring reliable products. Notwithstanding the manufacturer's stringent, unit-by-unit quality control testing, we find that product reliability at the factory varies significantly from week to week. However, we conclude neither that product reliability depends on the luck of the draw nor that it cannot be controlled precisely. Instead, we argue that products are more reliable when the workforce is more stable, which is a conclusion stronger than any we find in the literature and one that went seemingly unnoticed by a highly sophisticated consumer technology firm and its manufacturer.

Yet the descriptive results of this paper are compelling without sophisticated analysis. According to Figure 4, field failure rates (and hence product reliability) regularly rise and fall. By comparing

the failure rates against worker turnover, we find that the two patterns align. After controlling for a range of factors, our estimates suggest that worker turnover's post-payday peaks make field failures 6.6% to 7.7% more likely, and worker turnover additionally accounts for a 1.7% to 3.2% increase in failure incidence during the rest of the pay cycle.

We discuss a few limitations of the study's findings. First, while we show that worker turnover's effect does not operate by imposing additional workloads or by derailing overall process learning, further research could shed light on the precise mechanisms underlying worker turnover's effect. In particular, this study does not distinguish between the following two theories: (1) Worker turnover disrupts the assembly lines' operations, leading to poorer workmanship and product reliability. Such disruptions may include the loss of the competence of experienced workers, the loss of team familiarity, or the increased demands placed on supervisor attention. (2) By swapping out experienced workers in exchange for new workers, elevated worker turnover may facilitate harmful self-selection. Namely, the firm may lose a greater share of its more capable workers in the post-payday weeks, if such workers self-select into staying longer than their peers.² Second, since the data capture neither variation in manufacturing practices nor evidence of effective mitigation, our study focuses on revealing the novel significance of worker turnover's effect on product reliability rather than on exploring how the effect should be managed and controlled.

References

- Arthur, Jeffrey B. 1994. Effects of human resource systems on manufacturing performance and turnover. *Acad. Management J.* **37**(3) 670–687.
- Ball, George, Enno Siemsen, Rachna Shah. 2017. Do plant inspections predict future quality? The role of investigator experience. *Manufacturing Service Oper. Management* **19**(4) 534–550.
- Ball, George P, Jeffrey T Macher, Ariel Dora Stern. 2019. Recalls, innovation, and competitor response: Evidence from medical device firms. *HBS Working Paper* (19-028) Available at SSRN 3251630.
- Batt, Robert J, Diwas S KC, Bradley R Staats, Brian W Patterson. 2019. The effects of discrete work shifts on a nonterminating service system. *Prod. Oper. Management* **28**(6) 1528–1544.

²It is not clear that self-selection on worker capability can explain our findings in aggregate. In particular, it must explain why worker turnover is also costly in weeks other than the post-payday week, i.e., $\widehat{\beta}_1 > 0$, when presumably the shorter-tenured, poor performers (who must exit at some point) are quitting.

- Bray, Robert L, Juan Camilo Serpa, Ahmet Colak. 2019. Supply chain proximity and product quality. *Management Sci.* **65**(9) 3949–4450.
- Chan, Tian Heong, Francis de Véricourt, Omar Besbes. 2018. Contracting in medical equipment maintenance services: An empirical investigation. *Management Sci.* **65**(3) 1136–1150.
- Colak, Ahmet, Robert Bray. 2016. Why do automakers initiate recalls? A structural econometric game. *Kellogg Working Paper* Available at SSRN 2682645.
- Corbett, Charles J, Maria J Montes, David A Kirsch, Maria Jose Alvarez-Gil. 2002. Does ISO 9000 certification pay? *ISO Management Sys.* **2**(4) 31–40.
- Corbett, Charles J, María J Montes-Sancho, David A Kirsch. 2005. The financial impact of ISO 9000 certification in the United States: An empirical analysis. *Management Sci.* **51**(7) 1046–1059.
- Deshpande, Vinayak, Morris A Cohen, Karen Donohue. 2003. An empirical study of service differentiation for weapon system service parts. *Oper. Res.* **51**(4) 518–530.
- Emadi, Seyed M, Bradley R Staats. 2019. A structural estimation approach to study agent attrition. *Management Sci.* (forthcoming).
- Fisher, Marshall L, Christopher D Ittner. 1999. The impact of product variety on automobile assembly operations: Empirical evidence and simulation analysis. *Management Sci.* **45**(6) 771–786.
- Guajardo, Jose A, Morris A Cohen, Sang-Hyun Kim, Serguei Netessine. 2012. Impact of performance-based contracting on product reliability: An empirical analysis. *Management Sci.* **58**(5) 961–979.
- Guajardo, Jose A, Morris A Cohen, Serguei Netessine. 2015. Service competition and product quality in the US automobile industry. *Management Sci.* **62**(7) 1860–1877.
- Hendricks, Kevin B, Vinod R Singhal. 1997. Does implementing an effective TQM program actually improve operating performance? Empirical evidence from firms that have won quality awards. *Management Sci.* **43**(9) 1258–1274.
- Ibanez, Maria, Michael W Toffel. 2019. How scheduling can bias quality assessment: Evidence from food safety inspections. *Management Sci.* (forthcoming).
- KC, Diwas S, Christian Terwiesch. 2009. Impact of workload on service time and patient safety: An econometric analysis of hospital operations. *Management Sci.* **55**(9) 1486–1498.
- Kini, Omesh, Jaideep Shenoy, Venkat Subramaniam. 2016. Impact of financial leverage on the incidence and severity of product failures: Evidence from product recalls. *Rev. Financial Stud.* **30**(5) 1790–1829.
- Kundu, Amrita, Kamalini Ramdas. 2019. Timely after-sales service and technology adoption: Evidence from the off-grid solar market in Uganda. *LBS Working Paper* Available at SSRN 3477210.
- Moon, Ken, Patrick Bergemann, Daniel Brown, Andrew Chen, James Chu, Ellen Eisen, Gregory Fischer, Prashant Kumar Loyalka, Sungmin Rho, Joshua Cohen. 2019. Manufacturing productivity with worker turnover. *Wharton Working Paper* Available at SSRN 3248075.

- Musalem, Andres, Marcelo Olivares, Daniel Yung. 2019. Balancing agent retention and waiting time in service platforms. *University of Chile Working Paper* Available at SSRN 3502469.
- Porter, Michael E, James E Heppelmann. 2015. How smart, connected products are transforming companies. *Harvard Bus. Rev.* **93**(10) 96–114.
- Ramdas, Kamalini, Taylor Randall. 2008. Does component sharing help or hurt reliability? An empirical study in the automotive industry. *Management Sci.* **54**(5) 922–938.
- Shah, Rachna, George P Ball, Serguei Netessine. 2016. Plant operations and product recalls in the automotive industry: An empirical investigation. *Management Sci.* **63**(8) 2439–2459.
- Song, Hummy, Robert S Huckman. 2018. Cohort turnover and operational performance: The July phenomenon in teaching hospitals. *Wharton Working Paper* Available at SSRN 3037753.
- Srinivasan, Ashwin, Bryan Kurey. 2014. Creating a culture of quality. *Harvard Bus. Rev.* **92**(4) 23–25.
- Ton, Zeynep, Robert S Huckman. 2008. Managing the impact of employee turnover on performance: The role of process conformance. *Organ. Sci.* **19**(1) 56–68.
- Tsukayama, Hayley. 2018. How Samsung moved beyond its exploding phones. *The Washington Post* (Feb. 23, 2018) Accessed: https://www.washingtonpost.com/business/how-samsung-moved-beyond-its-exploding-phones/2018/02/23/5675632c-182f-11e8-b681-2d4d462a1921_story.html.
- Warrantyweek.com. 2019. Worldwide automotive warranty expenses. *Warranty Week* (Aug. 22, 2019) Accessed: <http://www.warrantyweek.com/archive/ww20190822.html>.

Electronic Companion

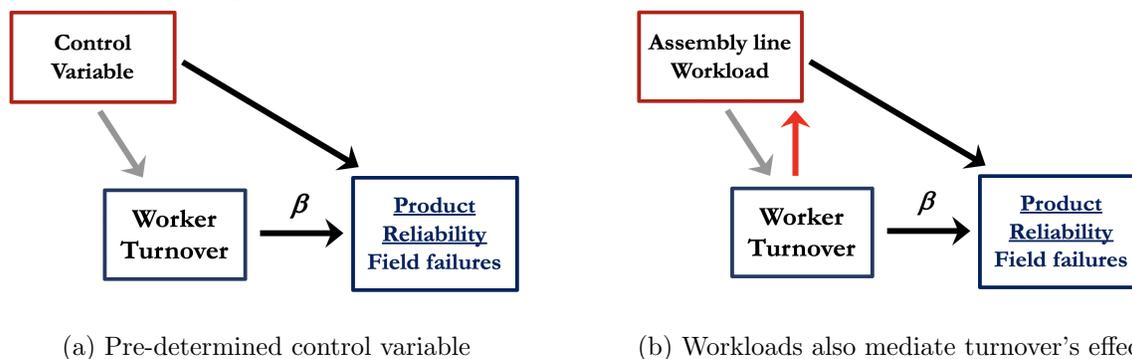
Appendix A: Assembly Line Workload as Mediating Control Variable

Section 3’s regression model specifies and interprets the coefficients $\beta := \{\beta_0, \beta_1, \beta_2\}$ as describing the causal effects of worker turnover. In considering worker turnover as a causal treatment, we recognize that workloads operate as both a control variable and a mediating variable. It is appropriate to control for workloads because they may directly influence product reliability (e.g., Shah et al. (2016)) and through compensation simultaneously affect worker turnover. However, they also mediate worker turnover’s causal treatment effect, because managers adjust assembly lines’ workloads in response to worker turnover.

Below we more extensively explain the role of workloads as a mediating control variable and how it should be interpreted in our regression analysis.

Figure EC.1a illustrates the two types of typical control variables. In the first typical case (dark arrows), the variable’s realized outcome independently affects product reliability. For example, a component’s exogenous defect rate affects the product’s reliability, and including the rate as a control variable yields a more precise estimate of the treatment effect β . In the second typical case (all arrows), the control variable is additionally correlated with the treatment. Excluding such a variable would impose an omitted variable bias on the estimated β . Other than workloads, we control for each of the literature’s known product reliability factors as one of these two types.

Figure EC.1 Separating Worker Turnover’s Effect on Field Failures from Workload Effects



Assembly line workloads are additionally mediators as shown in Figure EC.1b. As with the other control variables, assembly line workloads may independently affect product reliability (Shah et al. 2016). However, factory managers also adjust workloads in response to worker turnover (red arrow). Thus, a workload response is part of worker turnover’s total causal effect on product reliability.

Our regression model estimates the effect β shown in Figure EC.1b. (The coefficient on each regressor is estimated from the regressor’s empirical variation orthogonal to the other regressor.) Importantly, β excludes any part of turnover’s reliability effect that arises because workloads are adjusted for turnover. Economically, β could be interpreted as worker turnover’s effect on product reliability when managers cannot adjust workloads in response.

As Sec. 3.2 and App. B show, we find no workload effects on field failure rates in our setting (the upper dark arrow in Figure EC.1b). We thus believe that the estimated β captures turnover’s total causal effect.

Appendix B: Regressions with Month Fixed Effects

Table EC.2 shows our regression results after replacing the time trend with monthly fixed effects. We expect the fixed effects to additionally handle unobserved, time-varying factors. Table EC.2b replicates both Table 7b and Table 7c, because Table 7c adds component fallout rates on top of Table 7b's specification, but these quarterly rates are absorbed by the monthly fixed effects.

Table EC.2 Replicating Table 7 Including Month Fixed Effects in Place of a Time Trend

(a) Log field failure rate as dependent variable

	(1)	(2)
Worker turnover $\widehat{\beta}_1$	0.466*** (0.109)	0.325** (0.109)
Post-payday week (dummy) $\widehat{\beta}_0$	0.061*** (0.014)	0.059*** (0.013)
Worker turnover in post-payday week $\widehat{\beta}_2$	-0.014 (0.107)	-0.057 (0.104)
Line fixed effects		Y
Month fixed effects	Y	Y
R^2	0.33	0.37
Significance levels \rightarrow	*** - 0.001	** - 0.01 * - 0.05

Unbalanced panel of 1,516 production-active line-weeks in Sept. 2014 – Jun. 2015. Specification (2) includes fixed effects for 44 assembly lines.

(b) Log field failure rate as dependent variable, adding workload-related controls

	(1)	(2)	(3)	(4)
Worker turnover $\widehat{\beta}_1$	0.451*** (0.109)	0.317** (0.109)	0.449*** (0.110)	0.313** (0.109)
Post-payday week (dummy) $\widehat{\beta}_0$	0.061*** (0.014)	0.058*** (0.013)	0.061*** (0.014)	0.058*** (0.013)
Worker turnover in post-payday week $\widehat{\beta}_2$	-0.020 (0.107)	-0.060 (0.104)	-0.019 (0.107)	-0.058 (0.104)
Workload (line's log output assembled)	0.018 (0.012)	0.010 (0.011)	0.018 (0.012)	0.010 (0.011)
Log backlogged orders+1			-0.000 (0.001)	-0.001 (0.001)
Line fixed effects		Y		Y
Month fixed effects	Y	Y	Y	Y
R^2	0.33	0.37	0.33	0.37
Significance levels \rightarrow	*** - 0.001	** - 0.01	* - 0.05	

Unbalanced panel of 1,516 production-active line-weeks in Sept. 2014 – Jun. 2015. Specifications (1) and (3) are pooled regressions, whereas (2) and (4) include fixed effects for 44 assembly lines. Backlogged orders are the week's projected backlog calculated as the excess of (A) the firm's reported production targets based on demand forecasts and sell-through (effectively, orders) over (B) clear to build (projected production).

Across specifications, the estimated effect of post-payday worker turnover on the field failure rate is a constant, proportional increase of 6.0% to 6.3% with high statistical significance (p-value 0.001 or less). In weeks other than the post-payday week, the estimated effect of the mean worker turnover level (5.1%) on field failures incidence ranges from 1.6% to 2.4% in proportional increase. The effect is statistically significant at a p-value of 0.01 or less.