

Das, Sanghamitra; Krishna, Kala M.; Lychagin, Sergey; Somanathan, Rohini

Working Paper

Lifting the veil: The face of TFP in an Indian rail mill

CESifo Working Paper, No. 3515

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Das, Sanghamitra; Krishna, Kala M.; Lychagin, Sergey; Somanathan, Rohini (2011) : Lifting the veil: The face of TFP in an Indian rail mill, CESifo Working Paper, No. 3515, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/49492>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Lifting the Veil:
The Face of TFP in an Indian Rail Mill

Sanghamitra Das
Kala Krishna
Sergey Lychagin
Rohini Somanathan

CESIFO WORKING PAPER NO. 3515
CATEGORY 11: INDUSTRIAL ORGANISATION
JULY 2011

PRESENTED AT CESIFO AREA CONFERENCE ON GLOBAL ECONOMY, FEBRUARY 2011

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

Lifting the Veil: The Face of TFP in an Indian Rail Mill

Abstract

We use a proprietary data set on the floor-level operations at the Bhilai Rail and Structural Mill (RSM) in India to understand how output rose sharply in response to competitive pressures. Output increases came predominantly from reductions in production delays of various kinds. We model interruptions to the production process as a function of worker characteristics and training and find that a large part of the avoidable delay reductions are attributable to a particular form of training, suggesting that such investments can have very high returns. Our work suggests very high returns to knowledge-enhancing investments in emerging economies.

JEL-Code: D240, J240, L230, L610, M530.

Keywords: Total Factor Productivity (TFP), plant level data, competitiveness and trade.

Kala Krishna
Department of Economics
Pennsylvania State University
523 Kern Graduate Building
USA - University Park, PA 16802
kmk4@psu.edu

Sergey Lychagin
Department of Economics
Pennsylvania State University
307 Kern Building
USA - University Park PA 16802
lychagin@gmail.com

Rohini Somanathan
Department of Economics
Delhi School of Economics
Delhi University
India – Delhi 110007
rohini@econse.org

July 6, 2011

Sanghamitra Das passed away unexpectedly in December 2008 and is greatly missed. The paper has benefited from comments by the participants at the Conference on the Global Economy (CESifo Munich), Fifth Annual Conference on Economic Growth and Development 2009 (ISI Delhi), and a seminar at the University of New South Wales. We are grateful to George Alessandria, Ana Cecilia Fieler, Morita Hodaka, Michael Keane, James A. Schmitz, E. Somanathan and two anonymous referees for comments.

1 Introduction

Productivity differences have been found to be important in explaining differences in income. Hall and Jones (1999) find that of the 35-fold difference in output per worker between the United States and Niger, Total Factor Productivity (TFP) differences explain about twice as much as differences in physical and human capital. Syverson (2004) estimates within-industry TFP distributions in the United States and finds that firms in the 90th percentile are twice as productive as those in the 10th percentile. Ratios for emerging economies appear to be even larger. Based on plant level data from India and China, Hsieh and Klenow (2009) argue that a reallocation of resources that would bring the dispersion of productivity down to U.S. levels could raise output by as much as 50%.

Much of the literature treats productivity as a black box and evaluates its response to policy through its correlation with a handful of observed variables, more recent work tries to go deeper, though in narrower settings. These studies have been labeled “insider econometrics” since they rely either on data generated inside plants and firms as part of their normal production activities or on surveys with managers Shaw (2009) and such data is usually only available to “insiders”. These are often floor-level data on details of production process and work practices that allow researchers to zoom in on the micro determinants of productivity. This “bottom-up” approach has been at the frontier of research in both economics and management and has emphasized the role of contractual and organizational factors in explaining productivity.

Our paper is very much a part of this emerging literature. We use data from over 3,000 production shifts in the Rail and Structural Mill at the Bhilai Steel Plant in Central India. The mill has historically supplied all rails used by the Indian Railways and there was an implicit agreement with the state that this would continue unless the rails were found to be of substandard quality or the plant was unable to meet orders. Both of these factors led to the threat of entry of private players in the 1990s. A series of railway accidents brought rail quality into question and it was also unclear whether the plant could increase production in line with the plans for expanding the rail network.

The mill responded to these competitive pressures with a 28% increase in output between the first quarters of 2000 and 2003 in spite of no significant changes in employment or wage contracts and higher quality standards. The number of defects were cut in half and avoidable interruptions to production went down by 43 percent. We attribute much of this productivity surge to a brief and low cost episode of training. The use of detailed factory floor and administrative data on inputs, outputs and personnel allow us to control for the composition of the work teams, narrowly specified product types, unobservable worker fixed effects and the other episodes of training that occurred during this period.

Our work is closely related to that of Ichniowski, Shaw and Prennushi (1997) who use longi-

tudinal data on productivity and human resource management practices for 36 steel finishing lines across the United States and find that innovative practices have a productivity premium. A major source of higher productivity in their case is also increased uptime. Bloom and Van Reenen (2007) relate good management practices and productivity using detailed data on firms in different industries and countries. On the specific role of training in improving productivity, Das and Sengupta (2007) study steel-making in India and find that managers contribute to productivity only when trained. A recent randomized experiment that provided consulting services to textile firms in India shows that the resulting improvement in management practices reduced defects levels by over 50% and that many managers were simply unaware of low cost changes that could raise productivity (Bloom et al. 2010, 2011).

Also related is the literature that examines the response of productivity to competition. Galdon-Sanchez and Schmitz (2005) show that when competitive pressures mounted in the market for iron ore in the early 1980s, countries with mines that were close to becoming non-competitive increased efficiency, while others did not. Schmitz (2005) suggests this efficiency increase came about from improved work practices. Holmes, Levine and Schmitz (2008) model disruptions caused by technology shifts and argue that monopolies are less likely to innovate because they are under less pressure to tolerate the production losses that occur while the firm adjusts to new technology. Other studies on competition and productivity includes Caves and Christensen (1980), Rodriguez and Rodrik (1999), Tybout and Corbo (1991), Tybout and Westbrook (1991), Trefler (2004) and Nickell (1996).

Given the literature discussed above, it is worth emphasizing what we believe to be our main contributions. First, the factory floor data that we use, though restricted to a particular plant, is more detailed than any used in previous work. The rolling of rails is a continuous process with 3 shifts a day and our data allows us to capture this process in *real-time*. As management in the mill struggled to increase productivity, a variety of training programs were introduced and we can therefore compare their relative influence in a manner that is difficult with other data sets. Second, we consider a public sector undertaking where employment is secure. Such institutions are believed to be characterized by weak incentives and soft budget constraints. Although workers and managers in our study experienced no changes in marginal incentives for performance, perceived competition brought about changes in output. Organized sector workers in poor countries typically earn premium wages with very little monitoring. Our work suggests that competitive pressures may operate quite effectively with public ownership and the case for privatization in emerging economies that is so often made may need to be qualified. In fact, competition in these settings may play a much more important role than in industries dominated by private firms precisely because earnings are often above market wages and there is a gap between the actual and potential productivity of workers. Finally, in modeling the production process in the mill, we provide a methodological approach which is generalizable to other settings.

The rest of the paper is organized as follows. Section 2 provides background on the plant and an overview of the data. Section 3 estimates a reduced form model which supports our hypothesis that the productivity surge is explained by employee training. This is a precursor to the more structural approach in Section 4 which explicitly models the production process, fits the model to our data, and runs some counterfactual experiments to help identify the role of major contributors to productivity. Section 5 summarizes the lessons learned and concludes.

2 The Background and Data

2.1 The Bhilai Rail and Structural Mill

The Rail and Structural Mill in Bhilai is a part of a large integrated steel plant owned by the Steel Authority of India and built with Soviet cooperation in the 1950s. The location of the plant was part of the state planning strategy to bring jobs to remote areas and the plant successfully transformed Bhilai and ninety six of its surrounding villages into a company town. Local residents were given preference in employment. The jobs of regular workers are secure, with excellent fringe benefits including schooling, health care, housing and paid leave. There is a strong correlation between seniority and pay and a very weak one between pay and performance.¹These jobs have always been highly valued and workers here, in comparison with other industrial workers in India, have been referred to as the *aristocracy of labor* Parry (1999).

The Rail and Structural Mill (RSM) was commissioned in 1960 and has since then been the sole supplier of rails for Indian Railways. In addition to producing rails, the mill produces a variety of different products (beams, channels, angles) that are collectively called *structurals* and are either used directly in major infrastructural projects or as intermediate inputs in industries producing heavy machinery. The production technology in the mill is outlined in detail in Appendix A. In a nutshell, blocks of steel called blooms are fed into the furnace, rolled into rails or structurals, and then cut and moved to the cooling bed.

Each shift at the plant is generally devoted to either rails or structurals of a particular specification since switching between products results requires adjustments to the machinery and some downtime. Production takes place continuously in three eight-hour shifts. A total of 330 workers are employed at the mill during the period we study. Each worker can be

¹In spite of this, Bhilai is regarded as one of the more successful of the public sector undertakings. In our own field visits we found groups of extremely committed workers operating under physical conditions, particularly in the summer when floor temperatures can exceed 50 degrees centigrade in parts of the plant. In addition, managers seem to put in long hours and pitch in where needed.

tracked through a unique *personal number* that they are allotted when they join the plant. In addition, each worker is assigned to a *brigade* or work team.

As mentioned above, the mill has historically been the sole supplier of rails to the Indian Railways and there exists an informal understanding that this will continue unless the plant at Bhilai fails to provide needed quantities of rails of appropriate quality. Over the period we consider, the plant's orders were threatened for several reasons. First, a series of train accidents, culminating in a major train wreck in 1998 that killed 210 people, led to investigations that found sub-standard rails to be a major cause. Consequently, the railways decided to procure longer rails (reducing the number of joins) and limit their hydrogen content (since excess hydrogen makes rails brittle and prone to fracture). Second, track replacements and an expansion in the network led to an accelerated demand for rails and it was suggested that procurement from private players and imports be allowed to supplement the capacity at the Bhilai RSM. Finally, industrial liberalization measures of the early nineties had already brought private capital into mid-sized steel plants and these firms were keen to diversify into larger high-value products, like rails, with a stable demand. Executives and workers understood that in the absence of significant quantity and quality improvements, the market share of the mill was severely threatened.² In this setting, the RSM faced not just competition, but a threat to its very existence. A productivity surge ensued. How was this surge obtained? What lessons can be drawn from their experience? These are our main questions.

2.2 Our Data

There are a total of 3,558 shifts covering the period of January 1, 2000 to March 31, 2003. The core of our data set comes from two types of logs kept by the plant for each shift. The first of these is a *delay report* which records the total input of steel (number of blooms), total output (non-defective blooms rolled), the share of defective blooms and the length and cause of each interruption or delay in the production process. It also lists the brigade of workers on the floor during the shift. The second log is called the *daily presentee report (DPR)* and this records worker attendance by listing the personal numbers of all workers on the floor during a shift. It also lists reasons for the absence of each worker in the brigade but not on

²Newspaper reports at the time frequently discussed the breaking of SAIL's monopoly on rails. The Indian Express (June 9, 2000) reports:

Purchases from SAIL were stopped for a brief period of four months following the accident in Khanna, when the quality of rail was questioned by the Railway Safety Committee. However, purchases were later resumed in April 1999...

Jindal Steel and Power plans to break SAIL's hold over the huge orders by manufacturing rail for the domestic market from the next year. The company will manufacture 78 metre long rail, by acquiring and relocating a rail and structural mill in South Africa, near Raigarh in MP.

("Railways to procure Rs 400 cr worth rails from Bhilai Steel" by Jyoti Mukul)

the floor.

For active shifts, the delay report records and classifies all delays into four types. *Outside delays*, denoted by us as D_o , occur due to events outside the control of the managers and workers in the mill. These may be unanticipated, such as those caused by gas shortages or electrical faults or anticipated but unavoidable as in the case of electricity rationing or an inadequate supply of rail steel. *Finishing delays*, D_f , result mostly from the cooling bed for finished rails being full and unable to accept more rails. This is a downstream constraint that can shut down or slow down production in the mill. Third are *planned delays*, D_p , which are used for scheduled maintenance or adjustments of equipment. The fourth type of delay is the most important one for our analysis; it consists of unplanned and *avoidable delays*, D_a that result from workers making mistakes. We argue in the following sections that it was reductions in avoidable delays made possible notable productivity improvements at the RSM. Although a description of the cause for each delay episode is available in the data, we have no reference to the person or group at fault. We therefore model these delays as a function of the characteristics of the entire team present during the shift.

We supplement the shift-level logs described above with administrative data on worker designations and demographics and all programs of formal training undertaken by employees of the mill. In terms of designations, the technological process is organized around ten groups of workers. Seven of them man the stations along the production line. The other three groups (Executives, Crane operators and Technicians) participate at various stages of the process. Since different groups perform different tasks, we treat them as separate types of labor and construct ten labor variables from the shift-wise numbers in each of these groups.

It has been argued Parry (1999) that worker diversity by caste and residence may affect cooperation among workers within a shift. We use two indices to capture different dimensions of diversity:

$$I_{local} = \min(S_{local}, 1 - S_{local}),$$
$$I_{caste} = \min(S_{caste}, 1 - S_{caste}),$$

S_{local} is the share of shift workers originating from the local area around Bhilai and S_{caste} is the share of workers from the Scheduled Castes and Scheduled Tribes. These are two aggregate categories that comprise all the castes and tribes that have been declared as disadvantaged by the Indian state and thereby eligible for affirmative action. Historically, there was limited contact between these groups and the rest of the population. If there is tension between these groups, mixed brigades may be less productive.

There was an emphasis on training programs at the plant following the dismal performance of the company in the late nineties. Some training was also conducted because it helped in obtaining International Organization for Standardization (ISO) certification. Although there were a large number of programs, many of them lasted only a couple days and involved

few employees. Training records include a brief description of each program, start and end dates and a list of employees trained. We break up the total training time into the nine categories listed in Table 1, based roughly on the types of skills being targeted. There were four large programs which together account for two thirds of total training time over our period. These are listed in Table 2.

By combining the dates of training and the employees trained with attendance data from the DPR, we are able to generate training stocks for each shift in our data. So, for example, the total stock of safety training for a worker at any date equals the total number of days of such training administered to him by that date. We aggregate individual stocks of all workers on the attendance sheet for that shift to construct our nine stocks for every shift.

The delay report from which we obtain output and delays data was only available to us in the form of paper slips, one for each shift. Some of these were illegible and a few were missing. In addition, some shifts were dropped because there were devoted entirely to maintenance. These factors and some data cleaning (described in Appendix B) leave us with 3,121 shifts or 88% of the 3,558 possible shifts in our period.

The summary statistics for our main variables are reported in Table 1. Average output per shift increased by 46 blooms between the first quarter of 2000 and 2003. The number of defective blooms was halved and the number of interruptions caused by avoidable delays went from 1.4 per shift to 0.8. Conditional on the occurrence of a delay, the minutes of downtime did not change much. Employment over the period was fairly stable with approximately 66 workers on the floor in an average shift.³ The number of shifts devoted to rail production went from 133 in the first quarter of 2000 to 243 in the first quarter of 2003. There was also an overall increase in the number of shifts worked with fewer shutdowns. The bulk of the training given was in the four categories listed in Table 2. The table also lists the name of the most important training program in each category.

3 Reduced Form Estimates

We begin with reduced form estimates which reveal correlations between inputs and output without forcing a particular production model on the data. These estimates are qualitatively similar but larger than those in our preferred model in Section 4.⁴ They are useful in

³This is obtained by summing across the different designations. In total about 330 different workers were employed by the RSM over this period. Turnover and days off account for the difference in this number and the number on the floor across three shifts.

⁴However, the reduced form says little about how productivity increases might have occurred. By modeling production more carefully, we can better isolate the factors behind such improvements and reduce the chances that we are picking up a spurious correlation in the reduced form. For this reason we estimate a

Table 1: Summary statistics

	Mean	Std.dev.	2000 Q1	2003 Q1
Output, blooms	182	53	165	211
Defective output, blooms	1.7	1.7	2.6	1.3
Rolling rate, blooms/min	0.57	0.07	0.51	0.6
Delays per shift (#)				
Avoidable	1.3	1.4	1.4	0.8
Finishing	0.22	0.43	0.07	0.1
Outside	0.76	0.92	0.86	0.51
Planned	1.85	1.2	1.8	2.1
Downtime per delay (min.)				
Avoidable	30	39	30	29
Finishing	25	26	21	26
Outside	63	76	50	61
Planned	36	32	36	35
Workers (#)				
Control Men	5.4	1.6	5.2	6
Coggers	4.8	1.2	4.3	4.4
Crane Operators	4.9	1.1	5.2	5.5
Executives	1.9	0.33	1.6	2.1
Furnace Maintenance	6.2	1.3	6.4	6
Ground Staff	19	2.2	19	20
SCM Team	12	2	12	12
Services	5.8	1.3	5.7	6
Saw Spell	4.7	1.3	5	4.4
Technicians	1.2	0.73	1.2	1
Training Stocks (days)				
Productivity	63	54	0	110
Cost Reduction	24	11	19	33
Environmental	14	20	0	44
IT	6.5	4.8	1.7	11
Job Instruction	3.2	2.7	0	5.9
Motivational	21	34	3.6	119
Quality Control	19	19	4.2	42
Safety	20	9	7.5	29
Other	21	12	14	35
Diversity Indices				
Backward vs. Other Castes	0.34	0.03	0.34	0.35
Locals vs Migrants	0.41	0.05	0.41	0.42
Shifts producing rails	2,609		133	243
Shifts producing structurals	512		92	16

Table 2: The four biggest training programs of RSM employees, Jan 2000-March 2003

Name of program	Dates	Category	% Training time
Acceptance of rails program	June 2001 – July 2001	productivity	22
ISO-9000 workshop	May 2001, March 2002	quality control	9
ISO-14001 workshop	Jan 2002, July 2002	environmental	10
Success through empowerment of people	Oct 2002 – Jan 2003	motivational	24

Source: Personnel Records, Rail and Structural Mill

establishing that our overall conclusions do not rely on our structural assumptions.

We estimate a shift-level output equation:

$$Y_s = \beta \mathbf{X}_s + \varepsilon_s$$

where Y_s refers to the blooms rolled during shift s and \mathbf{X}_s is a vector of explanatory variables for the shift.

The variables that comprise the shift-wise characteristics \mathbf{X}_s in our model include the number of workers disaggregated by their designation, diversity indices based on hometown and caste, and training stocks by nine training categories. In every specification of our reduced form model, we control for type of output being produced as heavier products are typically slower to roll. We also control for the time of day, as morning shifts are usually less productive, and for the brigade-specific fixed unobservables.

Our results are reported in Table 3. The estimates of the baseline specification are shown in the first column. There are two training categories that are statistically significant and have positive coefficients: productivity and motivational training. Safety training is marginally significant but with a negative sign. Surprisingly, the numbers of non-executive workers on the floor do not seem to be systematically correlated with output; they are at most marginally significant.⁵ Brigades where employees from different castes are mixed together do not seem to do worse in terms of productivity than more homogeneous brigades. Brigades where locals work together with migrants appear to be slightly less productive, consistent with the anecdotal stories from Parry (1999), but the relationship is not statistically significant.

semi structural model of production in the next section.

⁵This number may be endogenous and we address this issue using an instrumental variable strategy described towards the end of this section.

In Columns 2 and 3 we attempt to isolate the effects of unobservables that could create an upward trend in output. By construction, training stocks are weakly increasing, and for this reason may be spuriously correlated with the observed growth of output. Most training programs are administered to groups of workers so the stocks of training in work teams change sharply over time. If the trend in unobservables varies smoothly, its spurious effect on the training coefficients should be greatly reduced with the introduction of quarter fixed effects in Column 2 and quarter-brigade fixed effects in Column 3. However, the estimated effects of both productivity and motivational training stocks remain the same, although the latter estimate becomes less precise consistent with motivational training spuriously entering the regression. The estimates also suggest that, among the different employee designations, only the number of executive workers on the floor is consistently associated with a higher output per shift.

In Table 3 we also report the p-values of the F-test that labor (of all designations) collectively makes no contribution to output. We find that even with the inclusion of executive workers, this test statistic is only marginally significant. Similarly we test for whether all training other than productivity training matters: in our preferred specification (Column 3) it does. We suspect that motivational training drives this result and we revisit this question at the end of this section.

To examine the effect of productivity training more closely, we take a difference-in-differences approach. We use the fact that productivity training was focused primarily on rails, as opposed to structurals; in fact, the training episode that accounted for a bulk of productivity training was termed the “acceptance of rails program”. Any worker-level unobservables such as the level of motivation or cooperation on the floor are likely to affect both rail and structural shifts⁶. Rails and structurals follow the same technological path in the production process, hence any unobserved improvements in the quality of physical capital are likely to improve output of both types of products as well. For these reasons, if one wishes to measure the effect of productivity training, structural shifts would be a good control for rail shifts. We therefore add the interaction of a rail shift dummy and stock of productivity training to our regression, thereby allowing productivity training to have a different effect on rails and structurals.

The results of this exercise are reported in Column 4 and confirm that the effect of such training is coming from rails and not structurals and that the size of the effect is approximately the same as in specifications (1)-(3). Figure 1 depicts this result in a more informal way by plotting the average output of rail and structural shifts and indicating the time period when most of the productivity training took place. The increase in rail output and the absence of such an increase in structurals is evident.⁷

⁶Rails and structurals are produced by the same brigades of workers; there are no “specialist” brigades.

⁷In Figure 1 we control for the type of structurals produced in order to eliminate any effects due to changes in the composition of structurals as these are very diverse. This is done by first estimating an OLS regression of blooms rolled on indicator variables for each product (whether rails or structurals). Recall that

Table 3: Output equation

	(1)	(2)	(3)	(4)	(5)	(6)
Workers (#)						
Control men	-0.95	-0.89	-0.95	-0.97	-0.83	0.24
Coggers	-1.14	-0.87	-2.08*	-2.07*	-1.67	-0.60
Crane Operators	-1.86*	-1.60	-1.30	-1.26	-0.94	-1.96
Executives	6.94*	7.47*	6.75*	6.72*	6.74*	6.52*
Furnace Maintenance	1.39	1.44	1.73*	1.77*	1.25	2.70*
Ground Staff	-0.41	-0.18	-0.44	-0.39	-0.86	-0.50
SCM Team	-0.24	0.03	0.27	0.26	0.43	1.53
Services	0.16	0.02	0.28	0.11	0.28	-0.88
Saw Spell	-0.99	-1.17	-0.89	-0.85	-0.58	-0.22
Technicians	0.17	0.59	0.59	0.38	0.29	-0.43
Training Stocks (days)						
Productivity training	0.27**	0.28**	0.34**	0.09	0.31*	0.28*
Productivity X Rail Shift				0.29**		
Cost Reduction	0.08	0.08	0.60	0.57	0.44	0.50
Environmental	-0.07	-0.25	-0.63*	-0.71**	-0.60*	-0.72*
IT	-0.44	-0.41	-0.23	-0.26	-0.25	-0.12
Job Instruction	0.19	-0.06	-0.52	-0.40	-0.46	-0.46
Motivational	0.21**	0.23*	0.27*	0.26*	0.18	0.28*
Quality Control	0.20	0.18	0.39*	0.38*	0.26	0.36*
Safety	-0.45*	-0.45	-0.68	-0.57	-0.38	-0.71
Other	0.07	0.10	-0.19	-0.21	-0.21	-0.19
Diversity Indices						
Backward vs. Other Castes	43	40	46	48	67*	51
Locals vs Migrants	-52	-52	-59	-52	-49	-54
Joint Tests, p-values						
Labor=0	0.04	0.09	0.03	0.03	0.05	0.08
Non-prod. training=0	0.00	0.15	0.01	0.01	0.04	0.01
Quarter fixed effect	No	Yes	Yes	Yes	Yes	Yes
Brigade X Quarter FE	No	No	Yes	Yes	Yes	Yes
Observations	3,121	3,121	3,121	3,121	3,121	3,121

Significance levels * 5% ** 1%

Unreported controls: Brigade, product, time of day fixed effects

Columns 1-4: OLS; dependent variable – blooms rolled (per shift)

Column 5: OLS; dependent variable – (blooms rolled/rolling rate) × (average rolling rate)

Column 6: IV; dependent variable – blooms rolled; endogenous variables: non-executive labor; instruments: extraordinary absence, interactions of all fixed effects

Under this specification, an extra day’s productivity training for the brigade raises the output by 0.29 blooms per shift. During the episode in which productivity training was concentrated (June–July 2001), an average worker received 1.6 days of training. As there are roughly 66 workers on the floor in a rail shift, a total stock of training grew by 106 days for an average working brigade. This translates into an output increase of 31 blooms per shift. Comparing this number to the actual increase in output of rails that took place during the summer of 2001 (see Figure 1), one can easily see that the productivity training played a major role in the rapid growth of the rail output.

As it happened, this training episode took place when the plant was shut down for maintenance and machinery replacement thereby effectively reducing the opportunity cost of the work time lost to training to zero. However, even stopping the mill for the duration of this training program would result in negligible costs relative to the productivity gains. The amount of training administered could be completed in two days (i.e. six shifts), leading to approximately $160 \times 6 = 960$ blooms in forfeited output. Given the productivity improvement of 31 blooms per shift, it would take only 31 shifts (i.e. less than two weeks) for this investment of time to pay off. We do not have figures for any other costs associated with the training but conversations with management suggest these were not large.

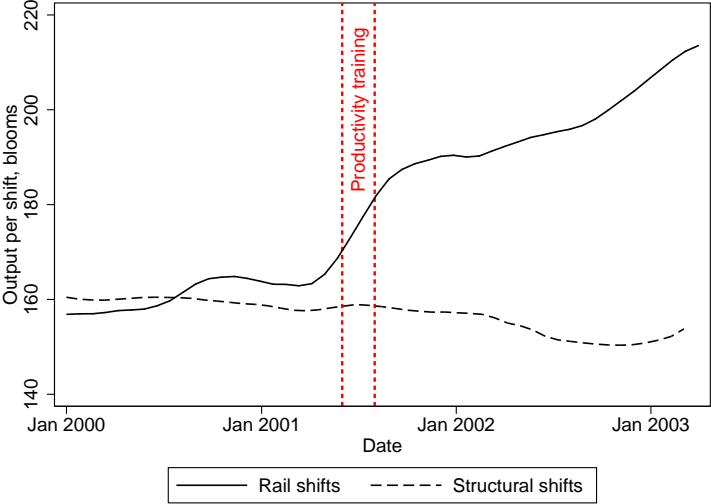


Figure 1: Effect of productivity training on output per shift (adjusted for product type)

While the effect of productivity training was large and systematic in all our specifications, we did not find similar evidence on the role of motivational training. Motivational training programs commenced on October 3rd, 2002, and were administered very gradually. For this

we dropped all shifts with mixed products (different types of rails or different types of structurals or mixed rails and structurals). For each shift we then adjust the output mix by subtracting the fixed effect of the product actually produced in the shift and adding the fixed effect of the base product. The adjusted output is then plotted in figure 1.

reason, we could not rely on a sharp increase in output to identify the effect of motivational training, as we did for productivity training.

Moreover, on September 4th, 2002, one month before the first person was trained, there was a systematic increase in the rate of production defined by the rolling rate, as depicted in Figure 2. This increase in the blooms rolled per minute of uptime coincided with the installation of new equipment which appears to have increased rolling rates.⁸ To ensure that our estimates are not confounded by the higher rolling rates associated with the new equipment, we estimate a model in which rolling rates are forced to be equal across shifts. This focuses on the output effects of production uptime. We do this by dividing output per shift by the rolling rate for the shift and multiplying by the average rolling rate in the entire data. We then re-estimate the same equation as in Column 3 and report the results in Column 5 of Table 3. Motivational training is no longer statistically significant and its size falls by about a third compared to that in Column 3. In contrast, the size of the coefficient on productivity training remains roughly the same. This also suggests that productivity training worked mainly through increases in uptime rather than by increasing the number of blooms rolled per minute of uptime.

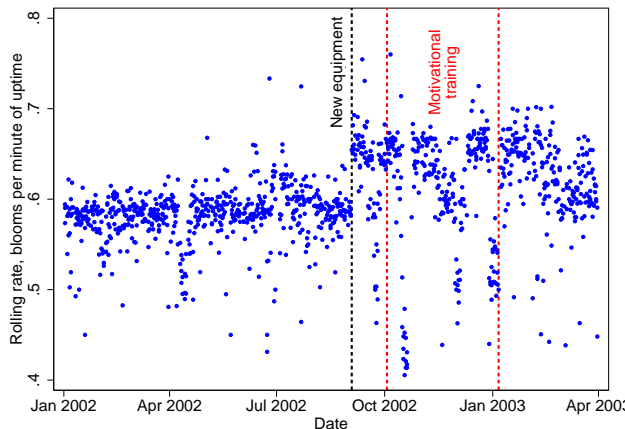


Figure 2: Coincidence of motivational training and investment into new equipment

Our final specification addresses the issue of labor endogeneity. If managers expect breakdowns, they may ask employees to work overtime or postpone earned leave. As a result, the estimated effect of labor would be biased downwards and may even become negative. We explore this possibility by using worker absence as an instrument for labor. Every time

⁸We identify this using delay cause descriptions. On exactly September 4th a new delay cause started appearing in the data; it is listed as “jamming at new descaling unit”. Since the increase in the rolling rate occurred simultaneously with the installation of the new equipment, we conclude that the former was likely to be caused by the latter. Our treatment of the rolling rate as rigidly determined by the technical parameters of the mill is in line with Ichniowski, Shaw and Prennushi (1997) who study a very similar production process.

an employee does not show up on the floor with his brigade, the reason for his absence is recorded. Usually, the absence is predictable and is controlled by the management to some degree: a worker may have a regular day-off, be on a holiday, or on an earned leave. However, there are instances when workers skip shifts due to a medical emergency, an injury at work, or some other extraordinary event. The total number of employees of each kind who do not show up at work at a given shift for those unusual reasons is clearly a good instrument for labor. Thus, we construct the number of workers predicted to be on the floor for that brigade, in that quarter and time of day of the shift, given the number of workers taking extraordinary leave. We use this instrument to estimate the same equation as in (3) and report our results in column 6. The two sets of results are similar suggesting that such endogeneity is not a major issue.⁹

Before proceeding to a more structural model, it is worth summarizing the consistent patterns that emerge from our reduced form approach. Table 3 suggests that there is at best a weak link between labor and output. With the exception of executive workers, the number of workers on the floor is not consistently associated with higher output. One possible explanation of this result is that the RSM is overstaffed, which is in line with popular beliefs about public sector enterprises in India. We do not find any support for the number of workers being endogenous. Diversity indices seem to have little systematic effect on output in our data. Training seems to matter: in particular, productivity training through a program designed to increase the output of acceptable rails has large and durable effects.

4 A Semi Structural Model of Production

The reduced form specification in the previous section suggests that productivity training might be responsible for most of the improvement in output per shift. However, it says little about how such an improvement might have occurred. If the improvement occurred through the operation of factors potentially subject to control like avoidable delays, then the evidence is consistent with our story. However, if, for example, outside delays fell in the period of productivity training, we could mistakenly identify productivity training as being behind the increases in output per shift. For this reason we estimate a semi structural model of production that allows us to rule this out.

⁹We do not have data on the reason for absences for executives and are therefore unable to construct a separate instrument for them. Their numbers are treated as exogenous.

4.1 The Production Process

We model a stylized version of the production process in which each bloom goes through a sequence of different stages before being finished. At each stage, either events outside the control of the brigade may occur or workers on the floor may make mistakes. These events and mistakes result in delays. The duration of most delays are allowed to depend on the worker characteristics during the shift. We choose what we believe are reasonable distributions for delay times caused by these events and then estimate the parameters of these distributions. We denote events and mistakes by M_x , their probability of occurrence by P_x and delay times by D_x , where $x \in \{o, p, a, f\}$ refers to the type of delay (outside, planned, avoidable and finishing).

The stages of production are as follows:

1. A steel bloom is fed into the furnace area for reheating.
2. An *outside event* may occur at this point triggering an outside delay of D_o minutes.
3. Next, workers may make an avoidable mistake causing an *avoidable delay* of D_a minutes before production is restored.
4. To roll the bloom into a final product, it takes time t where $t = 1/R$ and R is the rolling rate which may vary by quarter.
5. When the bloom is rolled, there is some chance of the cooling bed being full, resulting in a *finishing delay* of D_f minutes.
6. With probability p the final product is non-defective.
7. With some probability the equipment requires maintenance and D_p minutes are spent in a *planned delay*.
8. The process is repeated starting from step 1.

In the case of outside, planned and finishing delays, the events causing these delays are usually beyond the control of the mill workers. These include poor quality steel, flooding, electricity shortages, voltage fluctuations or a full cooling bed. We assume these delays occur with some exogenous probability that we allow to vary by the calendar quarter during which production takes place. We use a logit model to approximate the probability that each of these delays will occur. For calendar quarter q

$$P_x(\theta_q) = \Pr\{M_x = 1|q\} = \frac{1}{1 + e^{-\theta_{xq}}}$$

where $x \in (o, p, f)$.

Avoidable delays, in contrast, are caused by worker negligence and their occurrence depends on worker characteristics. For a shift with worker characteristics \mathbf{X} , the probability of avoidable mistakes is assumed to be given by

$$P_a(\mathbf{X}, \theta_a) = \Pr(M_a = 1 | \mathbf{X}) = \frac{1}{1 + e^{-\theta_a \mathbf{X}}}$$

Once a delay occurs, worker characteristics often influence how quickly the underlying problem is resolved. For example, flooding in the rainy season requires drainage of the affected area before production can be resumed and fluctuations in electrical voltage or broken equipment may have to be reported. When rails are defective or machinery is jammed, workers are actively involved in clearing the work area and equipment. This is true for all but finishing delays which occur downstream and do not involve mill personnel. We therefore allow delay durations for all but finishing delays to depend on shift characteristics.

If a delay of type x occurs, its duration, D_x , is drawn from a gamma distribution $\Gamma(\gamma_x \mathbf{X}, \lambda_x)$, where the shape parameter is dependent on shift characteristics \mathbf{X} and the scale parameter, λ_x , is constant for a given delay type. For $x \in \{o, p, a\}$, the density function for delay durations is therefore

$$f(D_x | M_x = 1, \mathbf{X}) = D_x^{\gamma_x \mathbf{X} - 1} \frac{e^{-D_x / \lambda_x}}{\lambda_x^{\gamma_x \mathbf{X}} \Gamma(\gamma_x \mathbf{X})}, \quad x = a, o, p.$$

For finishing delays $F_f(D_f)$ takes a similar gamma form with the shape parameter independent of \mathbf{X} but varying by quarter:

$$f(D_f | M_f = 1, q) = D_f^{\gamma_{fq} - 1} \frac{e^{-D_f / \lambda_f}}{\lambda_f^{\gamma_{fq}} \Gamma(\gamma_{fq})}.$$

We chose the gamma distribution because of its flexibility and find that it fits the data quite well. It also has a convenient invariance property allowing for a simple interpretation of the estimates. Assume the observed delay durations come from the sum of some baseline delays caused by each individual on the floor and that these individually generated delays are independently generated. Then if the delays of a single worker come from the gamma distribution $\Gamma(\alpha_i, \lambda)$ where α_i is worker i 's propensity to generate delays, by gamma-additivity, total delays are distributed as $\Gamma(\sum_i \alpha_i, \lambda)$. Our formulation is therefore consistent with, though not restricted to, a model in which delays are the sum of delays caused by individual workers and individual delays in turn depend on the characteristics of the worker, the product and the brigade.

To keep our estimation simple, we assume all random processes (M_x, D_x) are jointly independent conditional on covariates \mathbf{X} . The only permissible correlation between outside, avoidable and planned delays must therefore go through the observed characteristics of the brigade on the floor. This allows us to estimate avoidable, outside and planned delays independently of each other. We estimate γ_x and λ_x by applying the method of maximum likelihood to the sub-sample of blooms with positive delay durations, D_x ¹⁰.

It is worth emphasizing that if the probability of outside delays coincidentally fell just when productivity training was given, this would be picked up by the quarterly dummies in the model and would not spuriously be attributed to productivity training.

4.2 Results

The model is estimated on the sample of 2,609 shifts producing only rails during the period of January 1, 2000 to March 31, 2003.¹¹ Table 4 presents our estimates of θ_a , γ_a , γ_o and γ_p . To facilitate the interpretation of the results, we rescaled the estimates so that they represent the marginal effects of the components of \mathbf{X} on output. Thus the coefficient on productivity training of .06 in Column 1 means that an extra day of such training reduces the probability of avoidable delays in a way that results in .06 additional blooms rolled per shift. Appendix C describes the derivation of these marginal effects.

Column 1 of the table contains logit estimates of the determinants of delay episodes. If an avoidable delay occurs during a shift, we set $M_a = 1$ for the first bloom in that shift. For each subsequent avoidable delay on that shift, we assign that delay to subsequent blooms in that same shift. If, say, four avoidable delay episodes occur during that shift, $M_a = 1$ for the first four blooms of that shift. Given that the worker characteristics are the same for all blooms rolled in the shift and delays are assumed to be independent, all assignments of delays to blooms have equal probability and the way in which delays are assigned to particular blooms does not affect our estimates. The number of observations is the total number of blooms rolled into rails during the period.

The next three columns contain maximum likelihood estimates based on the gamma distributions described above and the number of observations is therefore the number of all delay episodes during this period for each type of delay. The dependent variable is the amount of delay time, in minutes. There are therefore as many observations per shift as delay episodes of the category that is being explained.

¹⁰This likelihood function has been shown to be concave and therefore has a unique maximum Choi and Wette (1969).

¹¹We exclude structurals to ensure maximum product homogeneity. The number of rail and structural shifts is given in Table 1

Total labor variables alone might not be able to capture all the relevant dynamics in workforce composition. When using aggregate numbers for different designations and training stocks, we implicitly assume that replacing one worker with another will not change the outcome as long as the workers' training stocks, etc., are the same. If this is not the case, our estimates may be subject to the omitted variable bias. To check the robustness of the estimates in Table 5 we augment \mathbf{X} by individual worker dummies and reestimate these equations.

The estimates presented in Tables 4 and 5 confirm the main result from the previous section: productivity training significantly increases output mainly through decreases in the occurrence of avoidable delays. The total effect of an additional day of productivity training for a work shift can be obtained by summing the coefficients of productivity training across the columns in Table 5). This gives us 0.14 blooms per day of training which is less than half the size of the reduced form coefficient of 0.29 blooms reported in Table 3. This translates into an output increase of 14.84 blooms per shift as opposed to the 31 blooms obtained from reduced form estimates. This comparison suggests that while the reduced form estimates do pick up some spurious correlation between output and the training received by the workers, the direct effect of productivity training on the probability of avoidable delays remains significant in a more controlled analysis.¹²

Our estimates in this section do not support the hypothesis of a positive effect of motivational training. The total effect of motivational training here is an order of magnitude lower than the corresponding estimates in the reduced form section. In addition, after the brigade composition is fully controlled by the inclusion of worker dummies, the effect of motivational training on any component of downtime becomes statistically insignificant. This supports our interpretation from the previous section: the positive estimates of the motivational training effect in Table 3 seem to pick up a coincidental shock to productivity caused by outside factors.

As in the previous section, the link between the number of workers on the floor and the level of output is insignificant. There is no evidence that larger teams make fewer mistakes or fix these mistakes faster. This lack of statistical significance persists even after controlling for the identities of workers and further supports the anecdotal evidence on overstaffing at the Bhilai Steel Plant given in Parry (1999). We also find no evidence that diversity impairs productivity; if anything, workers in shifts that are heterogeneous in terms of caste seem to make fewer mistakes.

Productivity training seems to be the only explanatory variable that is positively associated with higher output irrespective of the model's specification. To quantify and better illustrate the total effect of such training on output, we perform a set of counterfactual simulations below where the stocks of such training and other possibly relevant explanatory variables

¹²This may be because the reduced form does not fully pick up shift-wise factors outside the workers' control such as power outages or the quality of steel.

Table 4: Estimates of downtime components

	M_a	D_a	D_o	D_p
Workers (#)				
Control Men	-0.25	0.11	0.15	-0.29
Coggers	0.38	-0.04	-0.19	-0.31
Crane operators	0.26	-0.15	-0.88*	0.24
Executives	0.97	-0.01	-0.59	0.62
Furnace Maintenance	0.05	0.04	0.55	0.17
Ground Staff	-0.25	-0.11	-0.10	0.05
SCM Team	-0.33	0.17	0.25	0.12
Services	-0.19	-0.05	0.24	0.11
Saw Spell	0.02	-0.00	0.32	0.09
Technicians	-0.94*	-0.04	-0.43	-0.44
Training Stocks (days)				
Productivity	0.06**	0.00	0.02	0.01
Cost Reduction	-0.15**	-0.02	0.03	0.04
Environmental	0.11**	-0.02	0.03	0.05
IT	-0.02	0.02	-0.07	-0.03
Job Instruction	-0.36*	0.06	0.37	-0.24
Motivational	0.03*	-0.00	-0.01	0.01
Quality Control	0.02	0.02	0.05	0.01
Safety	0.07	-0.02	-0.31**	-0.12
Other	0.11*	0.01	-0.03	-0.04
Diversity Indices				
Backward vs. Other Castes	21**	1.2	2.7	13
Locals vs. Migrants	-7.0	0.57	9.8	-6.9
Observations:	481,003	3,329	1,800	5,209

Significance levels : * 5 percent ** 1 percent

Unreported controls: Brigade, product, time of day fixed effects

Only shifts producing rails are included.

Column 1: marginal effects on output via the probability of avoidable delays

Columns 2–4: marginal effects via the respective delay durations

are varied and the dynamics of simulated output are compared with observed output.

4.3 Counterfactual Experiments

We now use our estimates from the above model to study the impact of counterfactual changes in labor, diversity and the stock of training on the overall output. To avoid omitted variable bias, we restrict ourselves to the specification with individual fixed effects and use the estimates from Table 5.

We simulate production bloom by bloom, following the multistep procedure outlined in the beginning of this section. In each of our counterfactual experiments, the set of brigade characteristics is split in two parts: $\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2]$. In our simulations we freeze the components of \mathbf{X}_1 at their level in the first quarter of 2000. The components of \mathbf{X}_2 are allowed to change over time as observed in the data. This way, we predict the time path of output that would occur, had the management chosen not to adjust the variables in \mathbf{X}_1 . By changing the composition of \mathbf{X}_1 and \mathbf{X}_2 from one simulation to the next, we sequentially examine the importance of different sets of explanatory variables.

If R is the quarterly rolling rate, each bloom that enters the mill takes time $1/R$ to be processed if no mistakes or delays occur. If the model generates the event that a draw from a delay distribution is warranted, then the delay drawn is added to this time. Many delays may occur and these are additively incorporated. There is a probability, which varies by calendar quarter, that the bloom may be defective or cobbled, in which case the simulation will throw this bloom out. This continues until the 480 minutes of the shift are over. At the end of each shift, the total blooms rolled are generated. We take the monthly output generated by the simulation and label this to be the simulated output.

We start with simulating a full model in which \mathbf{X}_2 contains all the covariates and \mathbf{X}_1 is empty. We then shrink the list of variables in \mathbf{X}_2 in stages and observe the response of simulated output. The results of these experiments are depicted in Figure 3. In this figure, we report an average output across one thousand simulations, as well as the confidence bounds that account for randomness in the process that generates downtime.

Panel (a) shows that the full model fits the monthly output data very well. Recall that so as not to over parametrize the model, we only allow for quarterly changes in the probabilities of outside, planned and finishing delays. As a result, if outside delays, for example, are frequent in a particular week or month, the model will not take this into account and will tend to overestimate output for that month. This is why the model does not track the data spike by spike, but it does track it well on average.

In panel (b), we assume that the diversity indices are kept at their average level in the first

Table 5: Estimates of downtime components (with worker dummies)

Dependent variable	M_a	D_a	D_o	D_p
Workers (#)				
Control Men	-1.24	2.28	1.02	-2.55
Coggers	-0.00	2.56	-1.80	-1.92
Crane operators	2.51	0.89	-1.28	-2.90
Executives	0.39	-0.99	3.38	0.36
Furnace Maintenance	-1.94	2.30	7.14	-0.71
Ground Staff	-2.67	-0.04	4.63	1.31
SCM Team	-1.09	2.54	1.70	-1.84
Services	-5.37	-1.48	3.19	0.77
Saw Spell	-2.27	-0.70	4.93	2.31
Technicians	-1.18	1.28	-1.29	2.80
Training Stocks (days)				
Productivity	0.08**	0.01	0.03	0.02
Cost Reduction	0.11	-0.11	-0.10	0.26
Environmental	-0.05	-0.03	0.05	0.07
IT	0.10	0.02	-0.73	-0.22
Job Instruction	-1.20**	-0.23	1.28*	-0.53
Motivational	0.00	0.02	-0.00	0.04
Quality Control	0.02	0.05	0.06	0.03
Safety	-0.04	-0.10	-0.27	-0.30
Other	0.09	-0.01	-0.01	-0.04
Diversity Indices				
Backward vs. Other Castes	64	22	111	-38
Locals vs. Migrants	-13	-0.43	1.5	-5.9
Observations:	481,003	3,329	1,800	5,209

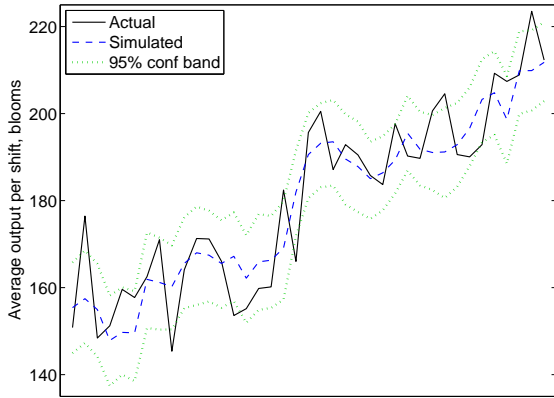
Significance levels : * 5 percent ** 1 percent

Unreported controls: Worker, product, time of day fixed effects

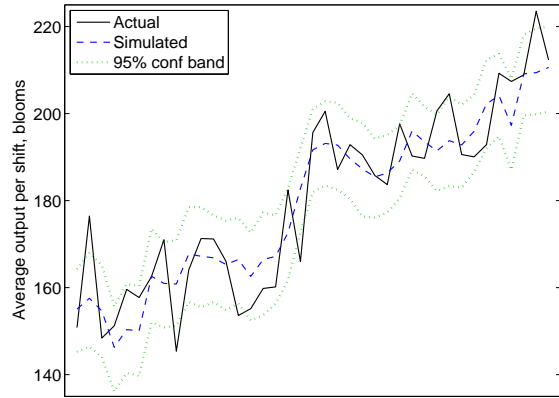
Only shifts producing rails are included.

Column 1: marginal effects on output via the probability of avoidable delays

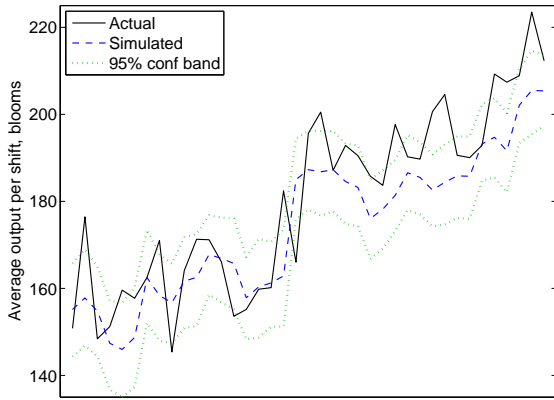
Columns 2-4: marginal effects via the respective delay durations



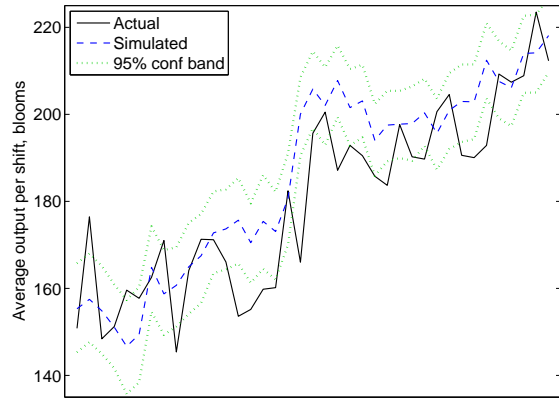
(a) Full model: $\mathbf{X}_2 = \mathbf{X}$



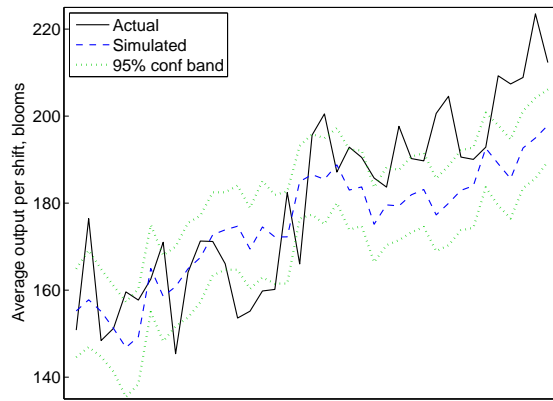
(b) Freeze diversity



(c) Freeze diversity, labor and worker dummies



(d) Freeze everything but productivity training



(e) Freeze all covariates: $\mathbf{X}_2 = \emptyset$

Figure 3: Diversity, labor and training, and their overall effects on output per shift

quarter of 2000. By comparing simulated output here with that in panel (a), we see that this restriction does very little to estimated output: therefore we can say that the effect of changing diversity is very small.

In the next panel, we impose restrictions on an additional set of variables: the total labor in each team and all worker dummies. This does not allow the management to control the composition of the workforce at all. The fluctuations in output are now driven only by training stocks and the quarterly dummies. This causes a model to slightly underpredict output starting early 2002. Since the data is not systematically outside the confidence bands of the predictions, we can say that changes in these labor related variables were not the primary determinants of output growth. In the RSM, total labor used did not change by much in this period and none of these coefficients is significant so it is not surprising that the change that did occur has little impact.

Panel (d) shows what output would be produced if no changes in diversity or labor composition were allowed and only the productivity training was administered to the workers. This way, we shut down the effects of all training that does not belong to the productivity category. There is some over-prediction in the first and second quarter of 2001 in panel (d), but the fit is good in later periods.

Finally, in panel (e) we fix all covariates at their level of their average value in the first quarter of 2001. The predicted time path of output is driven by the outside factors only, i.e, variations captured by the quarterly dummies in the probability and duration of delays and the quarterly variations in the processing rate. We see a large discrepancy between the predictions. Eyeballing the last two panels we see that productivity training was crucial in increasing output, with the total effect of the productivity training program equal to approximately 20 blooms of additional output per shift.¹³ Had management done nothing to train the employees, the growth in output would have been much more modest.

5 Conclusions

We see the contribution of this paper as both empirical and methodological. We attempt to explain output growth at Bhilai Rail and Structural Mill, a unit of the state-owned Steel Authority of India. In line with recent work in industrial organization we find worker training to be a low cost approach to improving productivity. Our analysis is based on a proprietary dataset that documents floor-level operations in greater detail than available in most studies within the productivity literature. This allows us to control for unobservable worker characteristics, shift-wise changes in the composition of work teams and the effects

¹³Since output is a non linear function of training, the marginal effects need not be exactly the same as those from the counterfactuals.

of periodic maintenance of equipment or availability of inputs. The access to data on all training programs allows us to isolate the types of training that have high returns. We attribute output growth to training that is specifically targeted to improving the output of acceptable rails by reducing the likelihood of mistakes on the factory floor. Training that is designed to achieve formal quality certification or improve safety and worker motivation appears to have no systematic effects in our case. In line with Galdon-Sanchez and Schmitz (2005) and Schmitz (2005), we suggest that competitive pressure was primary driver of productivity improvements.

On the methodological side, we explicitly model the process of production. The model we propose is not specific to the steel industry and could be applied to any manufacturing process that is organized around tasks in an established technological chain. This structural approach also allows us to perform series of counterfactual experiments on the role of different potential influences on output.

6 Appendix

6.1 Production Process

Figure 4 is a schematic representation of the production process. The main input is a long rectangular block of steel called a *bloom*. These are stored in the *bloom yard* and passed through different sections in a sequential process which converts them into rail tracks. They first enter one of four *furnaces* where they are reheated. They then move through a series of work tables in the *mill area* where they are shaped. The shaped blooms are cut to ordered lengths in the *hot saw area* and then stamped and moved to a *cooling bed*. Defective or misshapen blooms are referred to as *cobbed* and are set aside, the rest are classified as *rolled*. The mill runs twenty four hours a day, seven days a week with rare shutdowns for service and repairs. Production workers are rotated among three eight-hour production shifts.

Each worker, at any point in time, has a designation based on their job description and their seniority. Designations can be usefully divided into a few groups. Some workers are restricted to a particular location in the production process while others move around and ensure its smooth operation. In the furnace area, the *services team* keeps records of blooms, *control men* move the blooms in and out of the furnace and the *furnace maintenance* team, as the name suggests, ensures the furnace is in working order. In the mill area, *ground staff* are on the floor of the mill to ensure production flows smoothly. The *SCM team* is a group of *senior control men* and *motor operators* who along with the *coggers* sit in pulpits and direct the actual rolling of the rails in this part of the plant. In the hot saw area we have the *saw spell* team. *Crane operators* man cranes that transport blooms at various stages

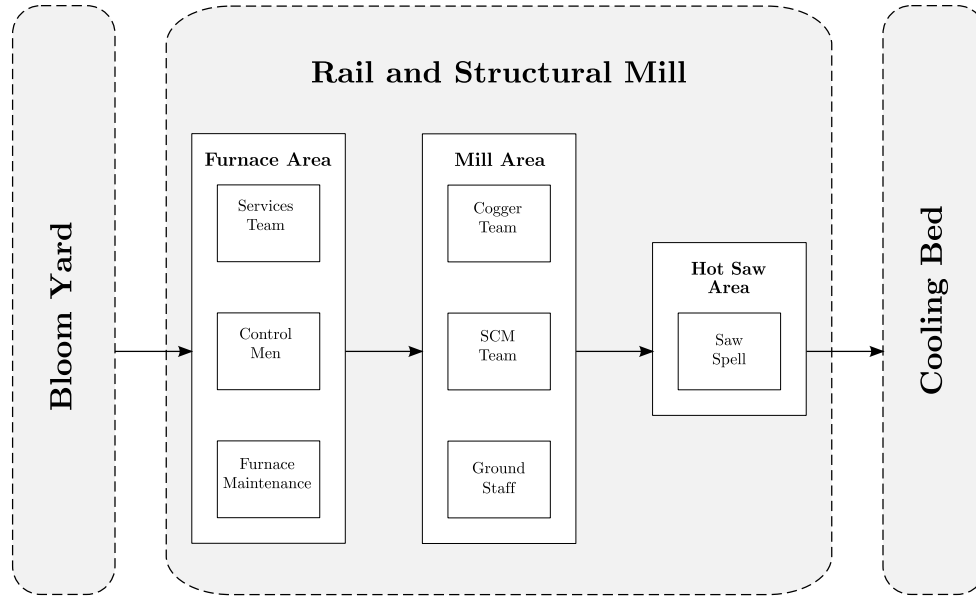


Figure 4: The process for rail production in the RSM

of production, *technicians* are responsible for fixing mechanical problems in the machines, while the *executives* oversee the operation as a whole. The model of production we estimate in Section 4 uses shift-wise data on the numbers of workers in each of these categories. There are shift-wise variations in these numbers generated by the number and types of workers on leave during any particular shift.

Shifts are operated by groups of workers called *brigades* that remain relatively stable over time. Each worker, at the time of joining the mill is assigned to one of these brigades. Brigade membership can be changed based on worker preferences and decisions of the supervisory and executive staff but these movements are infrequent. There are more people in a brigade than typically work in a shift, allowing for weekly days off and other types of leave. Brigades are rotated weekly across shifts: if a brigade works the morning shift in week one, it is switched to work the afternoon shift for week two and the night shift for week three.

6.2 Data Construction

There are a maximum of 3,558 shifts that cover the period of January 1, 2000 to March 31, 2003. Our data cleaning procedure drops 76 shifts with missing delay data and 126 shifts with missing or zero output; the majority of these shifts occur during regular maintenance days, when the entire mill is shut down. We also exclude 23 shifts with evident inconsistencies in the data such as those with zero uptime and non-zero output, or abnormally high or low

Table 6: Shifts in the main sample, by product type

Product	# shifts	Product	# shifts
Rails:		Structurals: Beams	
R-52	1,428	B-250	89
R-60	1,149	B-300	82
BS-90	32	B-350	17
Structurals: Channels		B-400	27
CH-250	54	B-450	26
CH-300	19	B-500	50
CH-400	40	B-600	108

rolling rates.

For each shift in the remaining sample we have a description of the type of output being produced. All products fall into two broad categories: rails and heavy structurals. We eliminate 212 shifts producing exotic products (such as crossing sleepers) or more than one type of output. The remaining 3,121 shifts constitute our main sample and are listed in Table 6 by product type.

6.3 Marginal Effects on Output per Shift

To facilitate the interpretation of the results obtained in the semi structural model of Section 4, it is useful to express all the estimates in terms of implied output gains. Let the expected output per shift be $\bar{Y} = E[Y]$. Suppose that one is interested in the effect of productivity training (denoted as X_{pt}). To derive the expression for this effect, one has to relate \bar{Y} to the model's parameters, θ and γ . This is done as follows.

Let T be the time needed to roll one bloom, taking into account all delays that occur along the way. This is a random variable; its distribution is determined by probabilities of delays of each type and the respective distributions of downtime. Specifically, the distribution of T is fully determined by the model's parameters, θ and γ , brigade characteristics, \mathbf{X} , and quarter, q .

Recall that P_x is the probability of a delay of type x for a single bloom and D_x is the delay time in minutes if the delay occurs. If R is the rolling rate for the quarter in question, then the expected time required to role a bloom is the inverse of the rolling rate plus all expected

delays for a single bloom. Denoting the expected value of T by \bar{T} , we have

$$\begin{aligned}\bar{T}(\mathbf{X}) = & (1/R) + P_a(\mathbf{X}, \theta_a) E[D_a | M_a = 1, \mathbf{X}, \gamma_a] + \\ & P_f(q, \theta_f) E[D_f | M_f = 1, q, \gamma_f] + \\ & P_o(q, \theta_o) E[D_o | M_o = 1, \mathbf{X}, \gamma_o] + \\ & P_p(q, \theta_p) E[D_p | M_p = 1, \mathbf{X}, \gamma_p].\end{aligned}$$

The expected number of blooms rolled in one minute is then $1/\bar{T}$ and expected output per shift is

$$\bar{Y} = p \frac{480}{\bar{T}}$$

where p is the probability of a non defective bloom. After substituting the above expression for \bar{T} into the expression for \bar{Y} , and noting that p and R vary only by quarter and so are unaffected by \mathbf{X} , one can derive the marginal effect of any component of \mathbf{X} . In particular, the effect of productivity training equals

$$\frac{d\bar{Y}(\mathbf{X})}{dX_{pt}} = \underbrace{-\frac{480p}{\bar{T}^2} \frac{dP_a}{dX_{pt}} \bar{D}_a}_{\text{column 1}} - \underbrace{\frac{480p}{\bar{T}^2} P_a \frac{d\bar{D}_a}{dX_{pt}}}_{\text{column 2}} - \underbrace{\frac{480p}{\bar{T}^2} P_o \frac{d\bar{D}_o}{dX_{pt}}}_{\text{column 3}} - \underbrace{\frac{480p}{\bar{T}^2} P_p \frac{d\bar{D}_p}{dX_{pt}}}_{\text{column 4}}$$

where $\bar{D}_j = E[D_j | M_j = 1]$ where $j \in (a, f, o, p)$. The labels in the above equation refer to columns in Tables 4 and 5. Note that as $\frac{d\bar{Y}(\mathbf{X})}{dX_{pt}}$ is a function of \mathbf{X} , we need to specify where it is evaluated: the numbers reported in Tables 4 and 5 are the values of the components of the function evaluated at the sample mean of \mathbf{X} .

References

- Bloom, Nicholas, and John Van Reenen.** 2007. "Measuring and Explaining Management Practices Across Firms and Countries." *Quarterly Journal of Economics*, 122(4): 1351–1408.
- Caves, Douglas W., and Laurits R. Christensen.** 1980. "The Relative Efficiency of Public and Private Firms in a Competitive Environment: The Case of Canadian Railroads." *Journal of Political Economy*, 88(5): 958–976.
- Choi, S. C., and R. Wette.** 1969. "Maximum Likelihood Estimation of the Parameters of the Gamma Distribution and Their Bias." *Technometrics*, 11(2): 683–690.
- Das, Sanghamitra, and Ramprasad Sengupta.** 2007. "Projection Pursuit Regression and Disaggregate Productivity Effects: The Case of the Indian Blast Furnaces." *Journal of Applied Econometrics*, 19(3): 397–418.

- Galdon-Sanchez, Jose E., and James A. Schmitz, Jr.** 2005. “Competitive Pressure and Labor Productivity: World Iron-Ore Markets in the 1980’s.” *The American Economic Review*, 92(4): 1222–1235.
- Hall, Robert, and Charles Jones.** 1999. “Why Do Some Countries Produce So Much More Output Per Worker Than Others?” *Quarterly Journal of Economics*, 114(1): 83–116.
- Holmes, T.J., D.K. Levine, and J.A. Schmitz.** 2008. “Monopoly and the incentive to innovate when adoption involves switchover disruptions.” National Bureau of Economic Research.
- Hsieh, Chang-Tai, and Peter Klenow.** 2009. “Misallocation and Manufacturing TFP in China and India.” *Quarterly Journal of Economics*, 124(4): 1404–1448.
- Ichniowski, Casey, Kathryn Shaw, and Giovanna Prennushi.** 1997. “The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines.” *The American Economic Review*, 87(3): 291–313.
- Nickell, Stephen J.** 1996. “Competition and Corporate Performance.” *Journal of Political Economy*, 104(4): 724–746.
- Parry, Jonathan P.** 1999. “Lords of Labour: Working and Shirking in Bhilai.” *Contributions to Indian Sociology*, 33(1-2): 107–140.
- Rodriguez, Francisco, and Dani Rodrik.** 1999. “Trade Policy and Economic Growth: A Skeptic’s Guide to the Cross-National Evidence.” NBER Working Paper 7081.
- Schmitz, Jr., James A.** 2005. “What Determines Productivity? Lessons From the Dramatic Recovery of the U.S. and Canadian Iron Ore Industries Following Their Early 1980s Crisis.” *Journal of Political Economy*, 113(3): 582–625.
- Shaw, K.** 2009. “Insider econometrics: A roadmap with stops along the way.” *Labour Economics*, 16(6): 607–617.
- Syverson, C.** 2004. “Product substitutability and productivity dispersion.” *Review of Economics and Statistics*, 86(2): 534–550.
- Trefler, Daniel.** 2004. “The Long and Short of the Canada–U.S. Free Trade Agreement.” *The American Economic Review*, 94(4): 870–895.
- Tybout, James, Jamie de Melo, and Vittorio Corbo.** 1991. “The Effects of Trade Reforms on Scale and Technical Efficiency: New Evidence From Chile.” *Journal of International Economics*, 31(3-4): 231–250.
- Tybout, James R., and M. Daniel Westbrook.** 1991. “Trade Liberalization and the Dimensions of Efficiency Change in Mexican Manufacturing Industries.” *Journal of International Economics*, 39(1-2): 53–78.