

NBER WORKING PAPER SERIES

BACK ON THE RAILS:
COMPETITION AND PRODUCTIVITY IN STATE-OWNED INDUSTRY

Sanghamitra Das
Kala Krishna
Sergey Lychagin
Rohini Somanathan

Working Paper 15976
<http://www.nber.org/papers/w15976>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2010

Sanghamitra Das passed away unexpectedly in December 2008 and is greatly missed. We are grateful to James A. Schmitz, Ana Cecilia Fieler and George Alessandria for comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 by Sanghamitra Das, Kala Krishna, Sergey Lychagin, and Rohini Somanathan. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Back on the Rails: Competition and Productivity in State-owned Industry
Sanghamitra Das, Kala Krishna, Sergey Lychagin, and Rohini Somanathan
NBER Working Paper No. 15976
May 2010
JEL No. J24,J48,L32,L61

ABSTRACT

We use a proprietary data set on the floor-level operations at the Bhilai Rail and Structural Mill in India to understand the determinants of changes in plant productivity in 2000-2003. During this period there was a 35 percent increase in output with minimal changes in factors of production, but sizable reductions in production delays. We model interruptions to the production process and find that a large part of these reductions are attributable to training. Our results suggest that specific knowledge-enhancing investments can have very high returns, and that the threat of competition provides powerful incentives to undertake such investments.

Sanghamitra Das†
Indian Statistical Institute
7, S. J. S. Sansanwal Marg
New Delhi - 110016
India
dasm@isid.ac.in

Sergey Lychagin
Department of Economics
The Pennsylvania State University
University Park, PA 16802
and
sul178@psu.edu

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

Rohini Somanathan
Delhi School of Economics
University of Delhi
Delhi 110007 INDIA
rohini@econdse.org

1 Introduction

Each year, the Indian prime minister announces labor awards to workers employed in government departments or public sector undertakings. In 2003, the most prestigious of these was awarded to a team of 5 from the Rail and Structural Mill of the Bhilai Steel Plant in recognition of their “outstanding contribution in the field of productivity”.¹ The Bhilai Steel Plant (BSP) is one of the five integrated steel plants of the Steel Authority of India Limited (SAIL), the company that has dominated the Indian steel sector since it was set up in the 1960s. SAIL is largely state-owned with 86 percent of equity and voting rights held by the Indian government.² Until the early nineties strict licensing rules restricted entry into Indian industry and SAIL, and many other manufacturing companies, survived with limited changes in technology and negative total factor productivity growth (Katja Schumacher & Jayant Sathaye 1998).

The industrial liberalization measures introduced in the early 1990s combined with a fall in world steel prices resulted in a series of operating losses for SAIL. Between 1992 and 2000 the share of the private sector in steel production went up from 45 to 68 per cent and the company faced the threat of its plants being labeled as *sick* industrial units (Ministry of Steel 1998-2009). Then began a remarkable revival. Between 1999-2003 SAIL production went

¹See <http://labour.nic.in/award/shram2003.htm> and “Bhilai team bags PM’s Shram Ratna for 2003”, *SAIL News, June-August 2003*, p.18

²www.sail.co.in

up by 12 percent even though the number of employees went down by 21 percent.³ This process has since continued with record profits in 2008, rising relative share prices and dividends of over 25 percent for several years.

In this paper we study this revival through the analysis of detailed data from one part of SAIL, namely the Bhilai Rail and Structural Mill (RSM) in the Bhilai Steel Plant over the period 1999-2003. 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 adequate 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 led to investigations which found sub-standard rails to be a major cause. To lower accident probabilities, the railways decided to procure longer rails and limit their hydrogen content. It was initially unclear whether the Bhilai RSM could provide these and for four months ending April 1999 purchases by the

³SAIL Performance Report, FY 2006 available at www.steel.co.in

railways were interrupted.⁴ Second, track replacements and an expansion in the network led to an accelerated demand for rails and it was suggested that private players and imports be allowed to supplement the capacity at Bhilai. Finally, the industrial liberalization measures had already brought private capital into mid-sized steel plants and these firms were keen to diversify into larger high-value products with stable demand. SAIL executives and workers understood that in the absence of significant quantity and quality improvements, the market share of the company was severely threatened.⁵

Of all the SAIL steel plants, labor productivity went up fastest at the

⁴The *Hindu Business Line* for June 8, 2000 (Calcutta) reports:

The drop in orders, particularly in 1998-99 and 1999-2000, was because of imports resorted to by the Railways on the plea that BSP was unable to supply rails as per its specifications. While the Railways has been a traditional buyer of BSP's rails with hydrogen content above three ppm (parts per million), it revised the specification in recent years and sought rails with hydrogen content of less than three ppm.

The situation forced SAIL to invest over Rs. 100 crores to equip BSP to meet the stringent requirement of the Railways for rails with less than three ppm hydrogen. Four Rail Quality Improvement Schemes were completed by BSP in 1999-2000.

⁵Newspaper 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)

Bhilai steel plant over the period 1999-2003. The production of crude steel per man year went from 121 metric tonnes in 1999-2000 to 129 in the next fiscal year and 153 in the year ending March 2003.⁶The output changes in the Rail and Structural Mill (RSM) far outstripped changes in the rest of the plant. In 1999-2000, the Indian railways bought a little over 300 thousand metric tonnes of rails from the plant. By 2002-2003, procurement was more than double this amount.⁷ Although production in the Rail and Structural Mill, like many other parts of the plant, is continuous, and largely automated, it relies more heavily on labor than in some of the other departments and is therefore well suited to a study of changes in productivity that rely on worker effort. The Bhilai steel plant is part of a township created by SAIL with subsidized schools and health facilities and is, in this sense, an island of skilled and highly paid labor in an otherwise poor state of central India. Both workers and management are unlikely to find comparable employment were the plant to close. It is therefore understandable that they had the right incentives to raise productivity when faced with competition. Our purpose in this paper is to explore the particular methods through which they achieved this productivity change.

The dataset that guides us in this effort contains detailed information on daily operations at the Bhilai RSM. The mill operates continuously with

⁶Each year, there is a special audit of major public sector undertakings of the Union government. These productivity figures are taken from this audit report for 2004 (Comptroller and Auditor General of India 2003-2004).

⁷Based on administrative data provided to us by the Ministry of Railways

3 production shifts per day. We obtained shift-wise data on the number of steel blooms rolled into rails in each shift, a list of all workers present during the shift, with their designations, and all delay episodes with their duration and a description of the cause of the delay. We combine these data on the production process with administrative data on worker demographics and all episodes of training. Even though the overall number of workers did not change much during the three-year period we consider, the combination of workers on the floor changes from one shift to another; we are able to control for this variation using worker fixed effects.

There is now considerable evidence of total factor productivity (TFP) differences across countries and firms. (Robert Hall & Charles Jones 1999) find that of the 35-fold difference in output per worker between the United States and Niger, TFP differences explain about twice as much as differences in physical and human capital. (Peter Klenow & Chang-Tai Hsieh 2009) use plant level data from India and China and show that the variance across firms within these countries is much larger than in the U.S. and the rationalization of production could raise output by as much as 50 percent. A range of institutional and policy variables could lie behind these TFP patterns, such as access to credit, physical and social infrastructure, technological spillovers and managerial practices.

A recent “bottom-up” approach in economics and management research tries to uncover particular sources of productivity differences by modeling the production process within particular industries. Our study is closely

related to this work and to a related literature that estimates productivity responses to greater market competition. (Casey Ichniowski, Kathryn Shaw & Giovanna Prennushi 1997) examine the productivity effects of human resource management practices using monthly data for 36 steel finishing lines across the United States. They find that workers in plants with traditional employment contracts and hierarchical supervisory structures are less productive than those in firms with innovative practices and that some of these productivity gains are realized through increased uptime. (Sanghamitra Das & Ramprasad Sengupta 2007) attribute the productivity increases of blast furnaces in Indian steel plants to improved coal quality and find that additional managers did not contribute to production unless they were also trained. (Nicholas Bloom & John Van Reenen 2007) combine surveys on management practices with TFP estimates from balance sheet data to examine the influence of such practices on firm productivity. This approach is in contrast to the traditional one that uses aggregate factors of production like labor and capital without specifying explicitly what happens inside the firm.

On the effects of competition, (Jose E. Galdon-Sanchez & James A. Schmitz, Jr. 2005) show that when the market for steel collapsed in the early 1980s, countries with iron-ore mines that were close to becoming non-competitive increased efficiency, while others did not. (James A. Schmitz, Jr. 2005) argues this efficiency increase resulted from less restrictive labor contracts which allowed more flexible allocation of labor time. Other

work on competition and productivity includes (Douglas W. Caves & Laurits R. Christensen 1980), (Jamie de Melo Tybout, James & Vittorio Corbo 1991), (James R. Tybout & M. Daniel Westbrook 1991), (Stephen J. Nickell 1996), (Francisco Rodriguez & Dani Rodrik 1999), (Daniel Treffer 2004) and (Shawn Klimek Dunne, Timothy & James A. Schmitz, Jr. 2009). (Chad Syverson 2010) provides a recent survey of this field. While we do not have data on plants not threatened by closure to compare to those that were, we do have much finer data on a single plant during events that led to increased competitive pressure. This allows us to focus on the shop floor for a finer investigation of apparent productivity improvements than is usually possible.

Our results attribute most of the observed productivity change to efficiency improvements resulting from fewer preventable delays and less production downtime. These changes in turn are explained by short and relatively inexpensive bouts of productivity training. Although several programs of managerial, motivation and technical training were conducted for the mill workers over this period, the only type of training that appears to have significant causal effects is training targeted at specifically improving rail quality. An interesting contrast with other studies is that we find increased uptime in the absence of any systematic changes in the numbers employed.

The rest of the paper is organized as follows. Section 2 provides background on the steel plant and rail mill at Bhilai. Section 3 describes our data set. Section 4 uses simple growth accounting identities to decompose output changes into its component parts, namely changes in rates of production, in

delays and in the fraction cobbled (the rails that were visually defective during the process of rolling before the final cooling process) and explores the patterns in each of these.⁸ Section 5 explicitly models the random processes resulting in production delays and Section 6 fits these to the data. Section 7 develops some counterfactual experiments to help identify the contribution of each of a number of factors to productivity and Section 8 summarizes the lessons learnt and concludes.

2 The Rail and Structural Mill in the Bhilai Steel Plant

The steel plant at Bhilai covers about 17 square km and currently employs about 34,000 workers.⁹ Until the plant was built in the mid 1950s, Bhilai was a small and remote village and the plant was located there as part of a planning strategy to bring jobs to remote areas. With the coming of the plant, Bhilai and 96 of its surrounding villages were transformed into a *company town* and the former owners of land were compensated in part by being given preference in employment. The jobs of regular workers are secure, with excellent fringe benefits including schooling, health care and housing, travel benefits and paid leave. These jobs have always been highly valued. (Jonathan P. Parry 1999) in his detailed and entertaining account

⁸Tests to identify defective rails are performed after cooling and so are not in our data.

⁹(Steel Authority of India 2008).

of labor conditions and performance at Bhilai refers to workers there as the *aristocracy of labor*. He notes that although tasks on the plant can be “extremely demanding, the amount of the working day spent on them is not.” There was, in the late nineties, a 15-20 percent surplus in manning levels, a strong correlation between seniority and pay and a very weak one between pay and performance.¹⁰

In spite of this, the Bhilai plant is widely regarded as the most successful plant in the public sector and perhaps surprisingly, the structure of incentives led to unequal distribution of work and absenteeism rather than low average effort. 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. Parry noted that although large public industrial units recruited and organized labor differently, underutilization was pervasive in many of these firms. Our findings on competition and productivity growth in the state-owned sector may therefore be applicable more broadly.

The Rail and Structural Mill (RSM) is an integral part of the plant. It was commissioned in 1960 with enough capacity to satisfy domestic demand at that time. Since then, it has been the sole supplier of rails for Indian Railways. In addition to producing rails, the mill produces a variety of different products (beams, slabs, channels, angles) that are collectively called

¹⁰(Parry 1999, pp. 17-19).

structurals and are either used directly in major infrastructural projects or as intermediate inputs into industries producing heavy machinery. Each shift at the plant is typically devoted to either rails or structurals, with a very few shifts that are mixed. These mixed shifts are dropped from the data. Table 1 shows the total number of shifts during which only rails and only structurals were produced for each of the four years of our study. Output during a structural shift is both sensitive to product type and the production process is typically more time consuming than that of rails. We therefore restrict our study productivity changes in the mill to those shifts that produced rails.¹¹

Table 1: RSM shifts worked by year and product type

Year	Rails	Structurals
April 1999-March 2000	743	279
April 2000-March 2001	754	258
April 2001-March 2002	839	182
April 2002-March 2003	958	79

Source: BSP Operational Statistics, 2003-2004, Table 9.18

Before describing our data in detail, it is useful to briefly outline the production technology. Figure 1 is a schematic representation of this process. In a nutshell, the the main input is a long rectangular block of steel call a *bloom*. These are stored in the *bloom yard* and pass through different sections in a sequential process which converts them into rail tracks. They first enter

¹¹From discussions with the management we gathered that the product mix is not primarily driven by price-cost margins. As the mill is the sole supplier for Indian Railways, they have a mandate to first meet orders from the Railways. Structurals are more profitable and their prices have been rising but their production depends on the capacity remaining after the demand for rails has been met.

one of four *furnaces* where they are heated. They then move through a series of work tables in the *mill area* where they are shaped and then pass to the *hot saw area* where they are cut to ordered lengths, 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 24 hours a day 7 days a week with very rare shutdowns for service and repairs. Production workers are rotated among three 8-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 are not. In the furnace area, the *services team* does the recording of blooms, *control men* move the blooms in and out of the furnace, while the *furnace maintenance* team looks after the furnace. In the mill area, *ground staff* are on the floor of the mill ensuring the smooth flow of production. The SCM team, a group of *senior control men* and *motor operators*, 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.

Some groups of workers are not restricted to particular areas of the RSM: for example, *crane operators* man cranes that transport blooms at various stages of production, *technicians* are responsible for fixing mechanical problems in the different machines, while the executives oversee the operation as a whole.

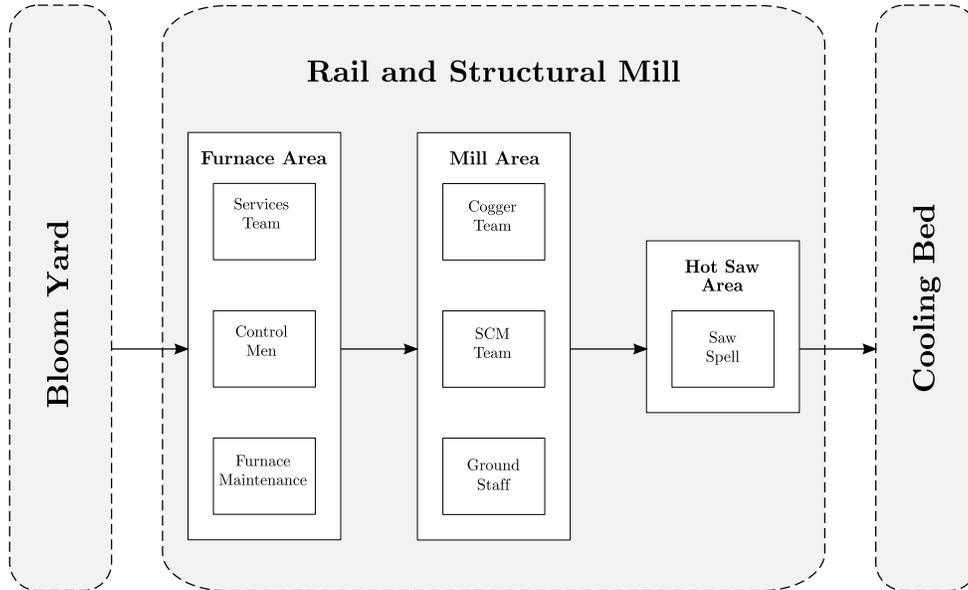


Figure 1: The process for rail production in the RSM

The model of production we estimate in Section 6 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 1, it is switched to work the afternoon shift for week 2 and and the night shift for week 3.

3 Our data

We have data on a total of 3558 shifts covering the period January 1, 2000 to March 31, 2003. There are two types of logs kept by the plant for each shift. The first of these is a *delay report* which records of the total input of steel, total output, the share of defective blooms and the length and cause of each interruption or delay in the production process. The second log is called the *daily presentee report* or the *dpr* and this records worker attendance. Each employee is assigned a unique identification number or *personal number* at the time they join the company. For each shift, and separately for the furnace and mill areas, the *dpr* lists the personal numbers of all workers on the floor during that shift. These two shift-level logs form the core of our data set and we describe them in some detail below.

During active shifts, the delay report allows us to classify all production delays into four classes. *Outside delays*, denoted by us as D_o , usually occur due to events outside the control of the managers and workers in the mill. These may be unanticipated, as in the case with gas shortages or electrical faults, or anticipated but unavoidable as in the case of some regular 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, there are planned delays, D_p , which are used for scheduled maintenance or adjustments of equipment. The fourth class 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. Avoidable delays are generated as the sum of mechanical, operational and electrical delays that are classified as avoidable. Although a description of the cause for each delay is available in the data, it is not possible to locate the source of all delays on the process chart in Figure 1 and we have no reference to the person or group at fault. We argue in the following sections that reductions in avoidable delays made possible notable productivity improvements at the RSM during the time period of interest.

The delay report sheet is filled in even if there was no production or no delay during that shift. For example, if there were inadequate orders, the delay report would record 480 minutes as the delay time and would list *no order* as the cause for downtime. There is therefore a delay report for every shift. We obtained access to paper copies of both delay and attendance logs, though for a very few shifts these logs were either missing, incomplete or illegible. In addition, data cleaning led to further attrition. Overall, the usable data covers about 94 percent of all the shifts in the period.

The attendance log (*dpr*) provides us with the composition of the workforce for each shift. The report records the brigade on the floor, lists the workers of the brigade that were present and also the reasons for the absence

Table 2: Categories of training received, Jan 2000-March 2003

Category	Total person days	Share in total, %
Motivational	511	27
Productivity	479	26
Environmental	184	10
Quality Control	165	9
Cost Reduction	135	7
Safety	131	7
Computer Skills (IT)	61	3
Job Instruction	57	3
Other	151	8
Total	1874	100

Source: Personnel Records, Rail and Structural Mill

of each worker in the brigade but not on the floor. We also have the designation of each worker by shift and can therefore track workers as they move across brigades, get hired, fired or promoted. We combine *dpr* data from administrative records on the social background of the personnel, including their caste affiliation and home state. Finally, as mentioned in Section 2, we have records of all episodes of training undertaken by employees of the mill. This includes a brief description of the training program, start and end dates and a list of employees trained.

There was an emphasis on training programs at the plant following the dismal performance of the company in the late nineties. Workers are trained both on the floor of the mill and in programs organized by the human resource department. Table 2 classifies these programs into nine categories based roughly on the types of skills that the program targeted.

Although there were a large number of programs, some of them lasted only

Table 3: 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

a couple days and involved very few employees. On average, the recipients of training were less experienced than their peers. This is especially noticeable for computer skills, cost reduction, safety and motivational training. Some training was conducted because it helped in obtaining International Organization for Standardization (ISO) certification. Most programs did not seem to target any particular designation; only few of them focused on a narrow workplace-specific skill. For the most part, training was administered to big groups of workers rather than to individuals. There were four large programs which together account for two thirds of total training time over our period. These are listed in Table 3. In Section 6 we examine the role of different types of training on productivity.

By combining *dpr*, personnel and training data we can generate shift-wise data on mean worker characteristics and estimate their effect on productivity. For example, the dates of training and the list of employees trained allows us

to generate training stocks for different types of training for each employee on each date and we aggregate this by shift to examine the role of variations in training on total output. Similarly, we have shift-wise compositions of worker designations and backgrounds. We know the number of executives, coppers, control men, ground staff, etc. on the floor for each shift, the share of migrant and local workers and their caste composition. (Parry 1999) observed some tension between the local population of Bhilai and migrants from other states. Potentially, communal conflicts like this may be strong enough to impair cooperation at the workplace and decrease productivity. We use both caste and home state data to account for this possibility. These data allows us to control for the composition of labor force at much greater detail than in possible in most studies of productivity. Although the overall composition of the workforce in the mill changed very little over this period, there is considerable variation in mean characteristics by shift and our strategy exploits this variation.

We have data available for 9 months of 1999 but ignore these because the delays logs during this period suggest that the mill was intentionally operated below full capacity. There were also several accounts in the press expressing concern by the railways about high hydrogen content of rails from Bhilai and reports of downtime in the mill frequently record “no order” as a cause. This problem was resolved in the following year with the installation of new equipment and orders from the railways went up again. As seen from the scatter plot in Figure 7, there was a jump in production at the very end

of 1999. Since our focus is on changes in productivity, we feared these might be overestimated with the inclusion of data from this 1999.¹²

4 Decomposition of Output Growth.

In this section we perform a simple decomposition of output growth over our period. Output is determined by the rate at which blooms are rolled and uptime. The latter is defined as the total shift time less the delay time in each of our categories. We then use the time pattern of output and delays to attribute output changes to changes in delay times for each of the delay classes. The main difference between this procedure and commonly observed decompositions is that we rely on the internal structure of the production process rather than on totals of raw inputs and the output.

Figure 2 summarizes the dynamics of output change. As seen there, average output during rail shifts expanded from 158 blooms in the first quarter of 2000 to 214 blooms in the first quarter of 2003, a 35 percent increase.

¹² The *Hindu Business Line* for June 8 , 2000 (Calcutta) reports:

The drop in orders, particularly in 1998-99 and 1999-2000, was because of imports resorted to by the Railways on the plea that BSP was unable to supply rails as per its specifications. While the Railways has been a traditional buyer of BSP's rails with hydrogen content above three ppm (parts per million), it revised the specification in recent years and sought rails with hydrogen content of less than three ppm.

The situation forced SAIL to invest over Rs. 100 crores to equip BSP to meet the stringent requirement of the Railways for rails with less than three ppm hydrogen. Four Rail Quality Improvement Schemes were completed by BSP in 1999-2000.

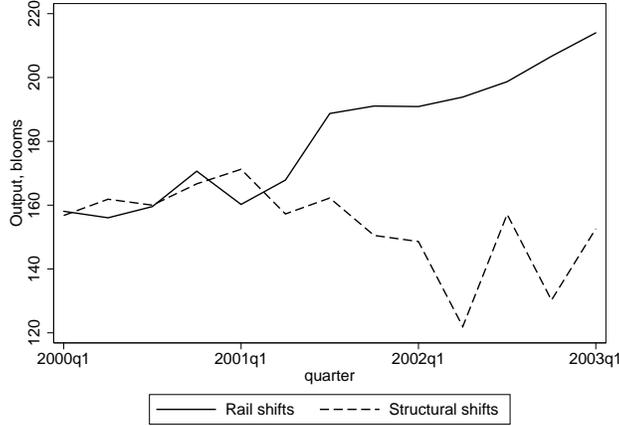


Figure 2: Blooms rolled for rails and structurals

Proceeding with our decomposition for rail shifts, let X_s denote the total number of steel blooms used by the brigade on duty during a rail shift s . Some proportion, p_s , of these blooms is successfully rolled into rails, while the remaining blooms are *cobbled* and removed from the line as defective. The final output is

$$Y_s = p_s X_s$$

The number of blooms the brigade is able to process is the product of uptime T_s and rolling rate R_s ¹³:

$$X_s = R_s T_s$$

Uptime is the total shift time of 480 minutes less time lost because of delays.

¹³(Das & Sengupta 2007) refer to R and T , as the *rate of output* the *rate of utilization* respectively.

We denote by D_{xs} the delay time resulting from a type x delay in shift s . Uptime is then given by

$$T_s = 480 - D_{os} - D_{ps} - D_{fs} - D_{as}$$

Given uptime and the input of blooms for each shift we infer the processing rate as X_s/T_s and combine the above equations to obtain:

$$Y_s = p_s R_s (480 - D_{os} - D_{ps} - D_{fs} - D_{as}) \quad (1)$$

Using Equation 1, we can attribute output growth in rails to the growth in 6 components: the fraction defective, p_s , the rolling rate R_s , and the duration of the four types delays D_{xs} .¹⁴ By definition, the percentage change in blooms rolled equals the sum of the percentage change in p_s , R_s , and uptime. The change in uptime can be further decomposed into its component parts:

$$\frac{dT_s}{dt} = -\frac{dD_{os}}{dt} - \frac{dD_{ps}}{dt} - \frac{dD_{fs}}{dt} - \frac{dD_{as}}{dt}$$

Dividing both sides by $\frac{dT_s}{dt}$, we find attribute changes in uptime to the different types of delays. Results of this decomposition are shown in the box contained in Table 4. For example, the contribution of outside delays to growth in total uptime is $\frac{(46.8-31.9)}{183-129} = 276 \approx .28$

¹⁴This is slightly richer than that in the literature. (Ichniowski, Kathryn Shaw & Giovanna Prennushi 1997) for example, focus only on the increase in uptime as the major source of productivity improvements.

Table 4: The decomposition of output growth into delay components.

	Q1 2000	Q1 2003	Growth
p	0.987	0.995	3%
R	0.54	0.614	42%
$(480 - D)$	297	351	55%
D_o	46.8	31.9	28%
D_p	90.4	72.5	33%
D_f	2.41	2.22	0.35%
D_a	43.6	22.7	39%
D	183	129	100%
Y	158	214	100%

According to Table 4, finishing downtime and fraction non-defective do not seem to be important contributors to output growth. The average fraction of defective blooms fell from 1.3 percent to 0.5 percent over the two-year period. While this change is significant the base year value is too low for it to have much effect on productivity. Outside delays and planned delays are important sources of output growth, but but are largely exogenous to a output in a particular shift.¹⁵ The remaining variables, the rolling rate R and avoidable delays, D_a , are determined on the mill floor and contribute substantially to the increase in output.

¹⁵The way in which workers operate machinery does influence the amount of time needed for its planned maintenance. These delays however, are not necessarily related to the composition of workers in the particular shift for which output is being measured.

4.1 Delays

A regular shift rarely runs without delays in production. Delays make a considerable part of a work day; during fiscal years 2001–2003 they accounted for 30 percent of an average shift time.

4.1.1 Avoidable delays

Most descriptions of avoidable delays contain one of these keywords: “not working”, “tripped”, “fallen”, “broken”, “jammed”, “grinding”, “adjustment”, “crane down”. The avoidable downtime decreased by almost a half, from 43 minutes per shift in the first quarter of 2000 to 22 minutes in the first quarter of 2003 (Figure 3). The second quarter of 2001 had an unusually long shutdown in production. According to our information, the time when the mill stood idle was used for training and equipment replacements. A training episode to raise productivity in rails termed the “acceptance of rails program” occurs at this time.¹⁶ The decline in delays that followed may thus have been caused by either training or equipment replacement. However, it is reasonable to expect, that the better equipment is likely to get broken less frequently in both rail and structural shifts, which is not observed in the data (for structurals, the avoidable downtime becomes even higher in Q4 2001). The training program explicitly focused on raising the output of rails which is consistent with the observed decrease downtimes during rail shifts, but not

¹⁶As we argue below, this training episode is the only one that looks like it actually worked.

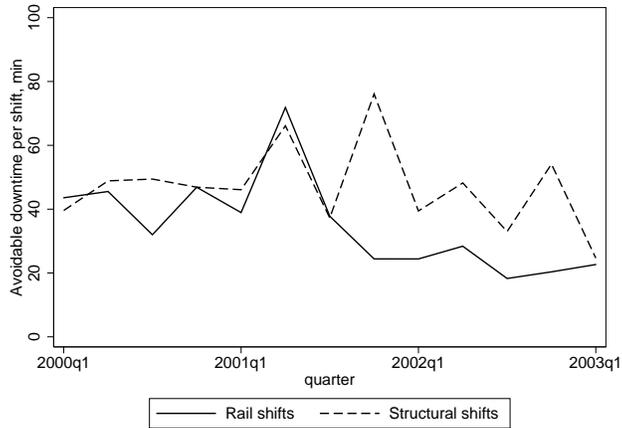


Figure 3: Avoidable delays for rails and structurals

structural ones.

4.1.2 Outside Delays

The patterns of outside delays were very different for rail and structural shifts (Figure 4). This is especially noticeable for year 2000, when RSM faced problems with the supply of low-hydrogen steel. Since hydrogen content is not as critical for the steel used in heavy structurals, as for rails, there was less outside downtime during structural shifts.

More than 60 percent of outside delays were associated with insufficient supply of inputs (keywords “shortage”, “voltage”, “restriction”) or their bad quality (keywords “lengthy”, “short”, “bad metal”, “asymmetry”).

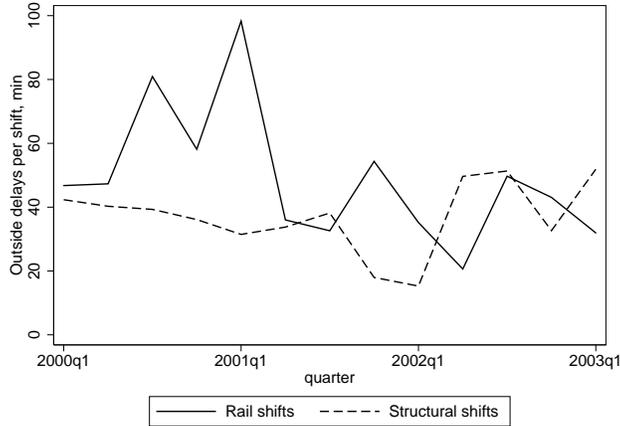


Figure 4: Outside delays

4.1.3 Planned Delays

After an initial drop in 2000, planned downtime has been slowly increasing until mid-2002 (Figure 5). We interpret this increase as a natural consequence of higher capital utilization. As output per shift grows over time, the equipment requires more frequent service. We did not observe any qualitative difference between rail and structural shifts which is consistent with this interpretation.

The descriptions of delay causes suggest that planned delays were primarily used for regular maintenance. More than 90 percent of planned delays were associated with “checking”, “adjustment” and “changing” of “section”, “stand” or “hot saw disc”.

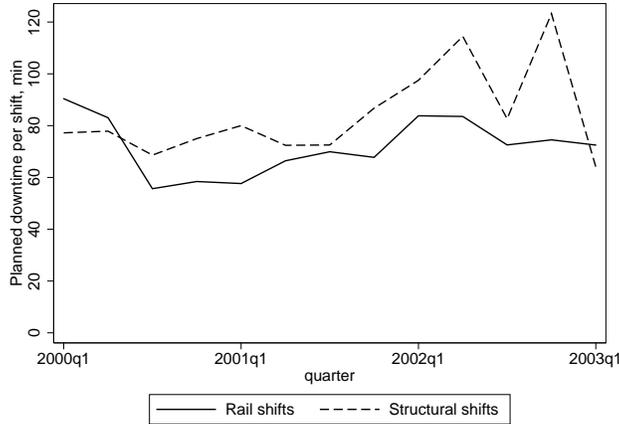


Figure 5: Planned delays

4.1.4 Finishing Delays

Finishing delays were only around 5–6 minutes per shift in 2000–2003 (Figure 6). Consequently, finishing downtime changes do not contribute much to output growth per se. They serve as a source of information about downstream bottlenecks that may restrict the productivity of the mill.

Finishing delays occur at the final phase of production – when the rails are coming from the Hot Saw section to the cooling bed. There is only one cause listed for all finishing delays: “cooling bed full”. If there is not sufficient space on the cooling bed, the operations at the Rail Mill are halted until the space becomes available.

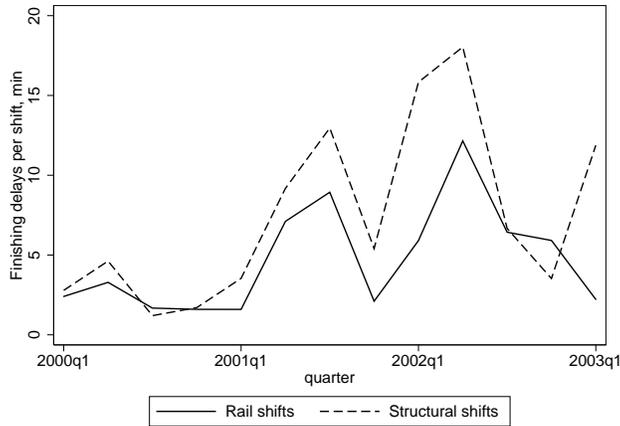


Figure 6: Finishing delays

4.2 Rolling Rates

Figure 7 plots the rolling rate over time. It includes both heavy structurals and rails.

It is evident from this figure that the rates seem to switch between discrete regimes. The switching clearly occurs at least three times: on September 15th 1999, November 7th 2000 and September 4th 2002. Within each regime the rates are dispersed around some average level that is stable over time which makes sense as dispersion will naturally arise in day to day operations.

Before September 1999, the mill had few orders as it was deemed incapable of producing the required quality. As a result, it was operating far below capacity. It is recorded that one furnace out of four was running between Sept. 1999 and November 1999, consistent with the low average rolling rate in this period.

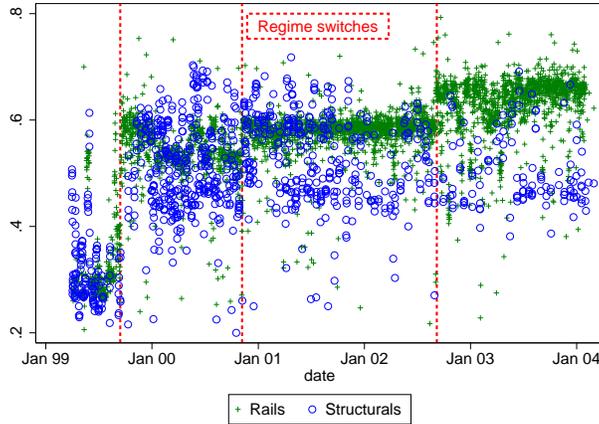


Figure 7: Rolling rates for rails and structurals, 1999-2003

After this first switch, between September 1999 and November 2000, there is a period where rolling rates fluctuated from one level to another. In this period the mill had limited access to low hydrogen steel from outside. There are two ways of reducing the hydrogen content in the rails. One is to use a degasser to make better steel. The other is to accept steel with a high hydrogen content but to cool the rail slowly allowing hydrogen to escape (see (Abhai Kumar Rai & Atul Agarwal 2007)). The Bhilai Plant installed a degasser in early 2000.¹⁷ It took six months or so to get consistent operation of this unit and until October of 2000, it was not fully effective. Note the high level of outside delays around the first switch (due to the lack of good steel) and high finishing delays afterwards (due to slow cooling) as depicted

¹⁷It is recorded that the degasser was put in for hot trials in March 2000 (Hindu Business Line Newspaper, June 9th 2000.) The degasser was effective October 1st 2000, as recorded in the controller general report 2003.

in Figure 8.¹⁸

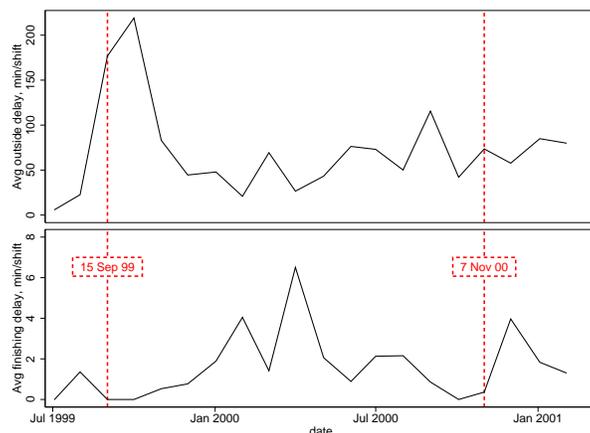


Figure 8: Outside and finishing delays, 1999-2001.

After November 2000, the degasser was running consistently and this is reflected in the higher more stable rate pattern.¹⁹ Finally, the regime switch that took place on September 4th 2002 is explained by the installation of some new equipment. We 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”. This delay occurred nine times in the first three days following the regime switching. Gradually, its frequency declined to five occurrences per quarter. Since the increase in the rolling

¹⁸For this reason, the RSM moved its output towards structurals in this period (where the hydrogen content was less of an issue) as far as possible. When it was forced to make rails, it did so, but could only use the slow cooling method as good steel was hard to come by. For this reason, before 2000, even when the share of rails was high, the output of rails was quite low.

¹⁹The degasser was installed in early 2000, started being tested in March, but did not function effectively until later in the year ((Hindu Business Line Newspaper, June 9th 2000, (Comptroller and Auditor General of India 2003-2004))

rate occurred simultaneously with the installation of the new equipment, we conclude that the former was likely to be caused by the latter.

Overall, it seems fair to say that the long run dynamics of the rolling rate seem to be determined by technological considerations and the outside constraints operating.

5 A Semi Structural Model of Production

We build a stylized model of the mill in which each bloom goes through a sequence of different stages before being finished. At each stage in the production process, 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 delays may 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. Finally, we discuss some alternative modeling approaches and justify the choice of the assumptions made in our model.

Each bloom goes through the following sequence of events and delays. We denote events and mistakes by M_x and delays by D_x , where $x \in \{o, p, a, f\}$ refers to the type of delay (outside, planned, avoidable and finishing).

1. A steel bloom is fed into the furnace area for reheating.
2. An *outside event* may occur at this point. If the event occurs, it triggers an outside delay of D_o minutes. D_o is drawn from the distribution

$F_o(D_o|\mathbf{Z})$ where \mathbf{Z} refers to the relevant characteristics of workers on the floor. We therefore assume that while the event is independent of the workers on the floor, the delay duration does depend on these characteristics because the events often require intervention by the mill personnel. In the absence of the event, there is no *outside delay*.²⁰

3. Next, workers may make an avoidable mistake causing an *avoidable delay* of D_a minutes before production is restored. D_a is sampled from $F_a(D_a|\mathbf{Z})$. In this case both the probability of the mistake and the delay time depends on \mathbf{Z} .
4. To roll the bloom into a final product, it takes time t where $t = 1/R$ and R is the rolling rate.
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 drawn from $F_f(D_f)$. This is unlikely to depend on \mathbf{Z} , since it is a downstream delay and mill personnel are not involved in clearing the cooling bed.
6. With probability p the final product is non-defective.
7. With some probability (which varies by quarter) the equipment requires maintenance and D_p minutes are spent in a planned delay, where D_p is drawn from $F_p(D_p|\mathbf{Z})$.

²⁰For 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.

8. The process is repeated starting from step 1.

For *outside, planned* and *finishing* delays, we assume that the *events causing the delays* are beyond the control of the mill workers. They occur with some *exogenous* probability that we allow to vary by the calendar quarter during which production takes place. *Avoidable delays* on the other hand are caused by worker mistakes which are allowed to depend on the composition of workers on the shift.

In each case, we use a logit model to approximate the process generating the event. For a shift with worker characteristics \mathbf{Z} , the probability of avoidable mistakes is given by

$$P_a(\mathbf{Z}, \theta_{\mathbf{a}}) = \Pr(M_a = 1|\mathbf{Z}) = \frac{1}{1 + e^{-\theta_{\mathbf{a}}'\mathbf{Z}}}$$

and for all other delay types the probability of events causing the delay in calendar quarter q is

$$\Pr\{M_x = 1|q\} = \frac{1}{1 + e^{-\theta_x(q)}}$$

If a delay occurs, we model the *duration* of delays that follow by gamma distributions. Each delay D_x is drawn from a gamma distribution $\Gamma(\alpha_x, \lambda)$ where $\alpha_x = \beta_{\mathbf{x}}\mathbf{Z}$ is the shape parameter and λ_x is the scale parameter (always independent of worker characteristics) for $x \in \{o, p, a\}$. The density function for delay durations of these three types is therefore

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

We assume that $F_f(D_f)$ takes a similar gamma form but with the additional restriction that the shape parameter is not dependent on \mathbf{Z} though it is allowed to vary by quarter. Thus

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

Recall that the mean of the gamma distribution is given by $\lambda\beta'\mathbf{Z}$, while the variance is $\lambda^2\beta'\mathbf{Z}$. Thus, our parametrization allows both the mean and the variance of avoidable, outside and planned delays to depend on who is on the floor. For finishing delays, it allows the mean and variance to vary by quarter only.

We chose the gamma distribution for its flexibility and as it fits the data quite well. In Figure 5 we compare our fitted gamma distributions with histograms based on the actual data and find the approximation to be very close. This parameterization also allows for a simple interpretation of the estimates. Assume the observed delay durations come from the sum of 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 an individual characteristic of worker i , by gamma-additivity, total delays are distributed as $\Gamma(\sum_i \alpha_i, \lambda)$. Our formulation is therefore consistent with, though not restricted by, a model in

which today's delays are the sum of delays caused by individual workers and individual delays in turn depend on the characteristics of the worker.²¹

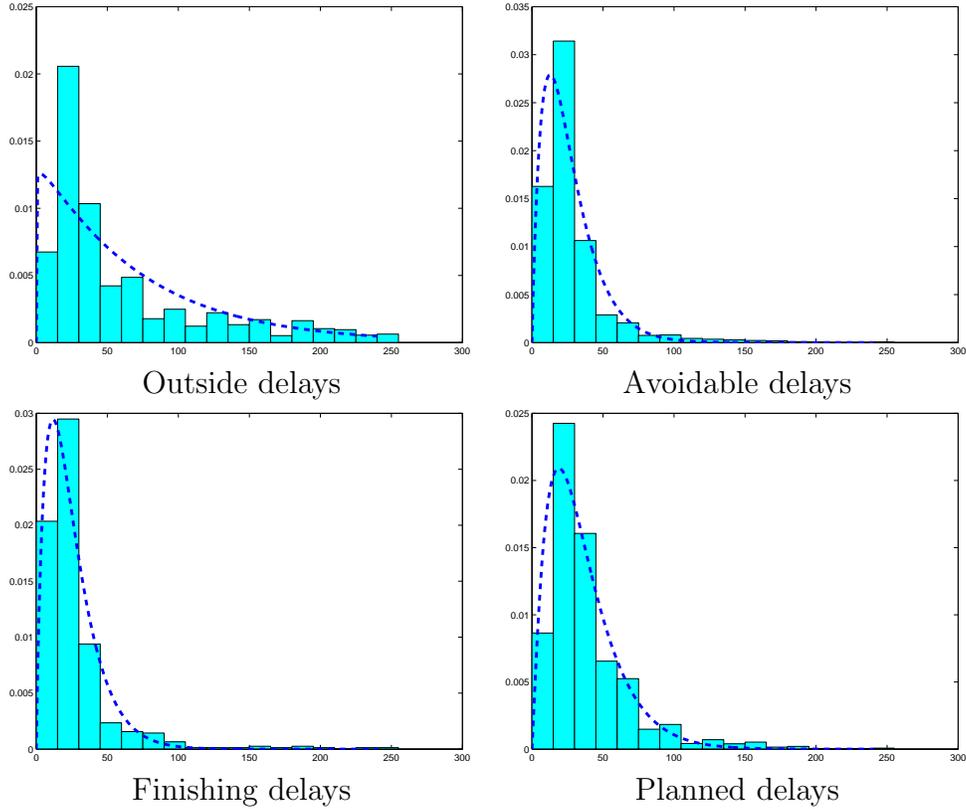


Figure 9: Gamma distribution approximations to observed delay durations

We estimate the probability of delays of various forms by applying the logit model to the sample of all rolled blooms. To keep our estimation

²¹For worker i we define α_i as follows: $\alpha_i = \gamma_1' \mathbf{Z}_i + \gamma_2' \frac{\mathbf{Y}}{N} + u_i$ where \mathbf{Z}_i contains worker-specific characteristics from \mathbf{Z} , \mathbf{Y} contains common characteristics, and u_i is worker i 's fixed effect. Then it is easily seen that $\sum \alpha_i = \sum (\gamma_1' \mathbf{Z}_i + \gamma_2' \frac{\mathbf{Y}}{N} + u_i) = \gamma_1 [\text{number of workers by designation, training by type}] + \gamma_2' \mathbf{Y} + \text{full set of fixed effects} = \beta' \mathbf{Z}$. For the model without worker fixed-effects, omit u_i from the above derivation. Tables 5 and 6 in the next section present estimates of the model with and without worker fixed effects.

tractable, we assume that all random processes in the model (M_x, D_x) are jointly independent conditional on \mathbf{Z} . The only permissible correlation between outside, avoidable and planned delays must therefore go through the brigade on the floor. This assumption allows avoidable, outside and planned delays to be estimated independently of each other. We estimate $\beta_{\mathbf{x}}$ and λ_x independently for avoidable, outside, planned and finishing delays by applying the method of maximum likelihood to the sub-sample of blooms with positive delay durations, D_x . It is well known that this likelihood function is concave and so has a unique maximum. (S. C. Choi & R. Wette 1969).

6 Estimation

We begin by describing the construction of variables that comprise the shift-wise characteristics of workers \mathbf{Z} in our model. These include the number of workers, disaggregated by their designation, diversity indices based on hometown and caste and training stocks by nine training categories.

The technological process is organized around ten groups of workers. As shown in Figure 1, there are seven teams of workers. In addition, there are three groups (Executives, Crane operators and Technicians) who may appear at any stage of the process. Since different groups perform different tasks, we treat them as separate types of labor and construct ten labor variables for the shift-wise numbers in each of these groups.

Worker diversity may affect cooperation among workers within a shift.

We use two indices to capture different dimensions of diversity, home state and caste affiliation. These are constructed as follows:

$$local_mix = \min(s_{local}, 1 - s_{local}),$$

$$caste_mix = \min(s_{scst}, 1 - s_{scst}),$$

where s_{local} is the share of shift workers originating from the local area around Bhilai and s_{scst} is the share of workers from the Scheduled Castes and Scheduled Tribes, the two groups that declared as disadvantaged by the Indian state and are entitled to affirmative action benefits.

To examine the effects of training, we construct training stocks for the nine categories in Table 2 for each worker on each date. So, for example, the total stock of safety training for worker w at any date equals the total number of days of such training administered to him by that date. We aggregate individual stocks for all workers on the attendance sheet for that shift to construct our nine stocks for every shift.

The model is estimated on the sample of shifts producing only rails during the period January 1, 2000-March 31, 2003. Table 5 presents our results. All estimates are presented as “average marginal effects”. For example, the effect of productivity training on the probability of an avoidable delay occurring has a marginal effect of -0.33. Recall that our estimates are scaled up by 10,000. This means that on average, an extra day of this training reduces the probability of a mistake on a single bloom by .000033 and with roughly

200 blooms a shift, by 0.66 percent a shift.

Column 1 of the table contains logit estimates based on the sample of all 418,819 blooms rolled over this period. 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.

Columns 2-4 contain maximum likelihood estimates based on the gamma distributions described above and the number of observations is therefore the number 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.

The estimates are consistent with the presence of over staffing. Note that in column 1 of Table 5 whenever the coefficient on the number of workers on the floor of a given type is significant, it is positive, indicating a higher probability of a mistake. The estimates on the determinants of delay durations in columns 2-4 also provide no clear evidence that the mistakes are fixed faster by larger brigades. Hence, a increase in the quantity of labor hired, all

else equal, is unlikely to reduce downtime and raise output per shift. These results are supported by the anecdotal evidence on over staffing at the BSP given in (Parry 1999).

Some types of training do seem to reduce avoidable mistakes as seen by the negative coefficients on environmental, motivational and productivity training. However we show below that of these, only productivity training seems to be robust to changes in model specification. We also find no evidence that caste diversity impairs productivity; if anything, workers in shifts that are heterogeneous in terms of caste seem to make fewer mistakes.

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 5, we augment \mathbf{Z} by individual worker dummies and reestimate these equations.

The estimates presented in Table 6 confirm our main result: productivity training significantly reduces downtime. In addition motivational training seems to reduce the average time spent in planned delays. The coefficient of job instruction training is now positive. This could be because job instruction occurs when new workers are hired or promoted. To check for this, we controlled for time in the job by putting in a novice dummy for those on the

job less than six months (these estimates are available on request), but this had no effect on estimates. Caste diversity is no longer significant. Thus, the only robust result is that productivity training helps.

To quantify and better illustrate the total effect of such training on output, we perform a set of counterfactual simulations in the following section where the stocks of such training and other possibly relevant explanatory variables are varied and the dynamics of simulated output are compared with observed output.

7 Counterfactual Experiments

In this section we use the estimates from the above section 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 model with individual fixed effects and use the estimates from Table 6.

We simulate production bloom by bloom, following the multi step procedure outlined in Section 5. In each of our counterfactual experiments, the set of brigade characteristics is split in two parts: $\mathbf{Z} = [\mathbf{Z}_1, \mathbf{Z}_2]$. The first part contains variables that we freeze at the level of quarter 1, 2000 in the simulation. The second consists of characteristics that 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{Z}_1 . By varying the composition of \mathbf{Z}_1 and \mathbf{Z}_2 from one simulation to the next, we sequentially examine the importance of different sets of explanatory variables.

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 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{Z}_2 contains all the covariates and \mathbf{Z}_1 is empty. We then shrink the list of variables in \mathbf{Z}_2 in stages and observe the response of simulated output. The results of these experiments are depicted in Figure 7.

Panel (a) shows that the full model fits the monthly output data very well. Recall from Section 5 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 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 under predict output starting early 2002. Since these predictions are not systematically outside the confidence bands, 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 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. Although the latter is the only covariate not frozen in time, the model still fits the data quite well suggesting that this other training was pretty useless. There is an 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, such as outside mistakes and variations in the processing rate. We now see a large discrepancy between prediction and the actual data. The last two panels suggest that productivity training was crucial in increasing output. Had management done nothing to train the employees, the growth in output would have been much more modest. The gap between simulated and actual output starts in the summer of 2001, which is precisely when the largest productivity training program took place (see Table 3).

8 Conclusions

We attempt to explain output growth in state-owned industry based on a proprietary dataset that documents floor-level operations at Bhilai Rail and Structural Mill, a unit of Steel Authority of India. During the three year period we consider, output increased by about a third in response to external pressures. Changes in the rolling account for 42 percent of this, while a fall in delay episodes and durations accounts for about 55 percent. Delays that are classified as *avoidable* by management in turn account for 39 percent of the time saved by fewer and shorter delays.

We then present and estimate a simple model of production that goes beyond the traditional production function approach and exploits the struc-

ture of the technological process. Our estimated model allows us to turn on and off various channels through which production could have increased. By conducting such counterfactual experiments, we show, for example, that most of the growth in production that came from reductions in avoidable delays occurred due to a single training episode during which workers were trained on raising rail quality.²²

We see the contribution of this paper as both methodological and empirical. The model that we propose is not specific to the steel industry. It may be applied to any production unit involved in a processing task of an arbitrary nature. Since many manufacturing firms are organized around tasks in established technological chains, our approach is most likely to be useful in the manufacturing sector.

By considering a firm in which aggregate labor adjustments were not possible, we highlight the role of other margins of productivity improvements, namely training and, to use Leibenstein's phrase, the existence of *X-efficiency* which made these changes possible.²³ The phase following industrial liberalization in India has seen a great diversity of experience with state-owned industry. While some industries, such as the state-owned airlines, have found it difficult to compete with private entrants, others like steel, heavy industry, telephone companies and state-owned banks have survived and, in some cases, increased their market share in response to competition.

²²Training episodes were by and large low cost operations as they were implemented at times when the mill was to be closed anyway.

²³(Harvey Leibenstein 1966)

The co-existence of state and private ownership within narrowly defined industries is intriguing given the differences in managerial practices, labor tenure systems and wage structures observed across the two forms of ownership. The changes we observed in Bhilai did not occur until workers perceived their jobs under threat. It may well be that the combination of job security and high wages that were associated with low productivity before industrial liberalization created the potential for the dramatic response that followed it.

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.

- Comptroller and Auditor General of India.** 2003-2004. "Report on the Union Government (Commercial): Public Sector Undertakings ." <http://www.cag.gov.in>.
- 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.
- Dunne, Timothy, Shawn Klimek, and James A. Schmitz, Jr.** 2009. "Does Foreign Competition Spur Productivity? Evidence From Post WWII U.S. Cement Manufacturing." Federal Reserve Bank of Minneapolis , Research Department Report.
- 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.
- 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.

- Klenow, Peter, and Chang-Tai Hsieh.** 2009. "Misallocation and Manufacturing TFP in China and India." *Quarterly Journal of Economics*, 124(4): 1404–1448.
- Leibenstein, Harvey.** 1966. "Allocative Efficiency and X-Efficiency." *The American Economic Review*, 56(3): 392–415.
- Ministry of Steel.** 1998-2009. "Annual Reports ." www.steel.gov.in.
- 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.
- Rai, Abhai Kumar, and Atul Agarwal.** 2007. "Rail Steel." Indian Railways Institute of Civil Engineering.
- 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.

- Schumacher, Katja, and Jayant Sathaye.** 1998. “Indias Iron and Steel Industry: Productivity, Energy Efficiency and Carbon Emissions .” *Ernest Orlando Lawrence Berkeley National Laboratory*, LBNL-41844.
- Steel Authority of India.** 2008. “Bhilai Steel Plant: Sustainability Report, 2007-2008 .”
- Syverson, Chad.** 2010. “What Determines Productivity.” NBER Working Paper 15712.
- 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.

Table 5: Estimates of downtime components

Dependent variable	M_a	D_a	D_o	D_p
Worker Teams (# workers)				
Control Men	1.91 [†]	-1.90	-9.65	4.77
Coggers	-0.86	-2.53	12.3	3.99
Crane Operators	-0.19	2.55	35.6*	-2.27
Executives	-3.63	-4.87	9.19	-12.8
Furnace Maintenance	-0.32	-1.31	-21	-1.63
Ground Staff	1.23 [†]	3.31 [†]	5.82	-0.31
Senior Control Men	1.76*	-4.31 [†]	-11.3	-1.12
Services	2.17 [†]	0.22	-10.1	-0.99
Saw Spell	0.45	-1.06	-9.16	0.21
Technicians	5.67**	-2.13	8.33	8.05
Training stocks (days)				
Cost Reduction	0.34	-0.04	-3.12	0.18
Environmental	-0.53**	0.06	-0.73	-0.72
IT	0.18	-1.20	3.03	0.29
Job Instruction	-0.27	1.30	0.73	1.75
Motivational	-0.12*	0.17	0.70	-0.12
Productivity	-0.33**	-0.13	-1.07 [†]	-0.08
Quality Control	-0.08	-0.15	-2.04	-0.17
Safety	-0.05	0.86	11.9**	0.89
Other	-0.33	-1.08	-0.14	0.79
Caste_mix	-146**	-108	-261	-185
Local_mix	23	-75	-367	156
Observations	418,819	3,366	1,788	5,219

Significance levels : †: 10 percent * : 5 percent ** : 1 percent

Column 1 reports average marginal effects on $\Pr\{M_a = 1\}$, scaled up by 10,000

Columns 2-4 report marginal effects on delay durations, scaled up by 10

Unlisted control variables: brigade dummies

Table 6: Estimates of downtime components (with worker dummies)

Dependent variable	M_a	D_a	D_o	D_p
Worker Teams (# workers)				
Control Men	29.2	-3.15	-83.3	12.9
Coggers	20.4	-3.59	17	12.5
Crane Operators	-21.7	52.2	58.1	12.1
Executives	-1.25	4.73	-11.5	-22.5
Furnace Maintenance	17.4	-23	-193	38.5
Ground Staff	15.5	18.4	-123	14.5
Senior Control Men	28	-1.80	-127	8.74
Services	26.5	27.1	-96	21.6
Saw Spell	14.1	25.2	-132	3.09
Technicians	-8.75	-15.9	201	-46.4
Training stocks (days)				
Cost Reduction	-0.67	2.11	4.18	-2.85
Environmental	-0.04	-0.75	-0.75	-0.95
IT	-0.32	0.26	21.5	2.62
Job Instruction	4.76**	4.63	-29.6	5.63
Motivational	0.07	-0.21	0.48	-0.60*
Productivity	-0.41**	-0.43	-1.88	-0.09
Quality Control	-0.22	-0.99	-1.64	-0.47
Safety	0.32	2.88	9.99	0.30
Other				
Caste_mix	-427	-926	-2638	34.1
Local_mix	11.2	-50	474	118

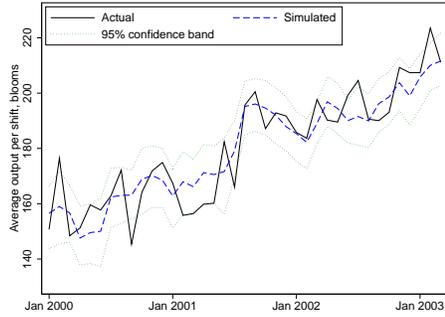
The number of observations are the same as in Table 5.

Significance levels : †: 10 percent * : 5 percent ** : 1 percent

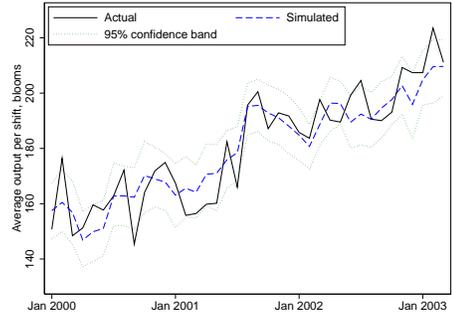
Column 1 reports average marginal effects on $\Pr\{M_a = 1\}$, scaled up by 10,000

Columns 2-4 report marginal effects on delay durations, scaled up by 10

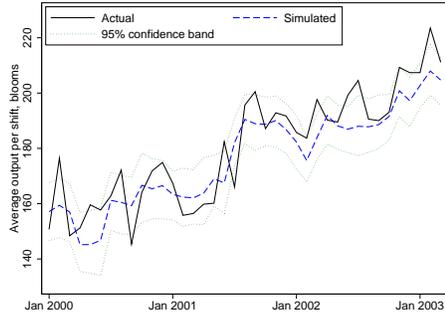
Unlisted control variables: individual worker dummies



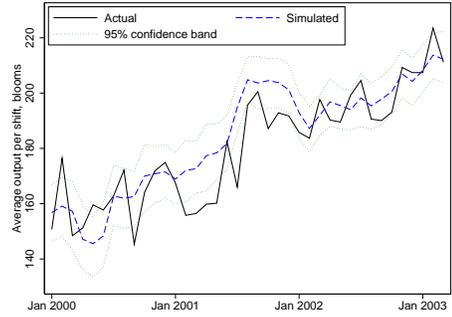
(a) Full model: $\mathbf{Z}_2 = \mathbf{Z}$



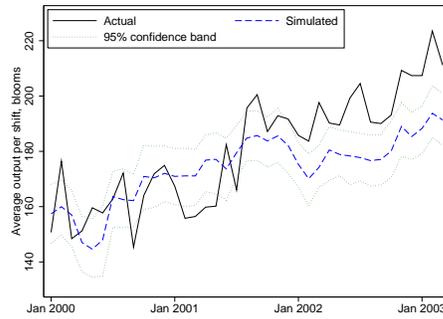
(b) Freeze diversity



(c) Freeze diversity, labor and worker dummies



(d) Freeze everything but productivity training



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

Figure 10: Diversity, labor and training, and their overall effects on output