www.elsevier.com/locate/worlddev
World Development Vol. 64, pp. 169–180, 2014

http://dx.doi.org/10.1016/j.worlddev.2014.05.024

Does Affirmative Action Reduce Productivity? A Case Study of the Indian Railways

ASHWINI DESHPANDE [a] and THOMAS E. WEISSKOPF [b,*]
[a] *University of Delhi, India*
[b] *University of Michigan, Ann Arbor, USA*

Summary. — Our objective is to shed empirical light on a claim often made by critics of affirmative action policies: that increasing the representation of members of marginalized communities in jobs comes at the cost of reduced productive efficiency. We undertake a systematic empirical analysis of productivity in the Indian Railways—the world's largest employer subject to affirmative action—in order to assess whether higher proportions of affirmative action beneficiaries in employment have reduced efficiency in the railway system. We find no evidence for such an effect; indeed, some of our results suggest that the opposite is true.

Key words — affirmative action, labor markets, productivity, railways, Asia, India

1. INTRODUCTION

Affirmative action (AA) encompasses public policy measures designed to reduce the marginalization of members of groups that have historically suffered from discrimination, exclusion, or worse. Starting in India a century ago, and accelerating after World War II, a great variety of AA policies have been applied in many countries of the world. Such policies are most often highly controversial, and their efficacy is highly contested. In particular, it is often argued by critics that any possible gains in inclusivity are outweighed by significant costs in economic efficiency. Our objective in this study is to subject this argument to rigorous empirical testing, in the context of a particularly important case of AA that has implications for many similar AA policies around the world.

India has not only the longest history of AA policies but also the most comprehensive system of AA, reaching far more people than all such policies elsewhere. In India the most prominent form of AA takes the form of "reservations" or quotas for the "Scheduled Castes" (SCs), now called Dalits, and the "Scheduled Tribes" (STs), called Adivasis. Together 22.5% of all seats in central-government-supported higher educational institutions and public sector jobs are reserved for these groups, corresponding to their share of the overall population in the 1950s. [1]

Criticism of AA policies in India is much the same as in most other countries where AA policies have been implemented. It is argued that such policies conflict with considerations of merit because less qualified candidates are selected in place of more qualified candidates, so that poorer academic performance and poorer quality of work on the job is to be expected from AA beneficiaries. [2] But advocates of AA—in India as elsewhere—argue that hiring is otherwise often far from truly meritocratic, and that workforce diversity may actually generate efficiency gains. [3]

To shed empirical light on this debate we focus on the world's largest employer subject to AA—the Indian Railways (IR), with roughly a million and a half employees—in an effort to assess the effects on productive efficiency of its reservations on behalf of Scheduled Castes and Tribes (henceforth "SC/STs"). [4] In the United States, where AA in hiring has been practiced in many industries since the 1960s, a variety of studies of this kind have been carried out. [5] In developing countries, however, such studies are very few in number. The only studies assessing the impact of AA in India focus either on electoral representation (Besley, Pande, Rahman, & Rao, 2004; Munshi & Rosenzweig, 2009), or on higher education (Bertrand, Hanna, & Mullainathan, 2010; Robles & Krishna, 2012). To our knowledge there has not yet been any systematic quantitative study of the effect of AA in the labor market on enterprise efficiency.

For our study of the IR we first compiled data from various zonal annual reports on productive inputs and outputs, distinguishing SC/ST employees from non-SC/ST employees at different job levels, for eight regional railway zones from 1980 through 2002. [6] Using the employment data we then constructed variables representing the SC/ST percentage of IR employees (SCS/T%), first for all employees and then for high-level employees only. We consider the latter SC/ST% to be the better indicator of the effect of AA on IR operations, because almost all SC/ST employees in high-level positions are AA beneficiaries—i.e., they would not have been able to

* Financial support for this paper was provided by the Research Office of the Dean, College of Literature, Science & the Arts; by the Center for South Asian Studies; and by the Residential College (all at the University of Michigan, Ann Arbor); as well as by a research grant from Anthony Heath, one of the organizers of the British Academy conference on international experiences of affirmative action. Smriti Sharma provided sterling research assistance. We are especially grateful to K.L, Krishna, B.N. Goldar, J.V. Meenakshi, G. Alivelu, Charles Manski, Wiji Arulampalam, William R. Johnson, and Gary Solon for critical insights and suggestions. Comments and suggestions received at conferences at the British Academy, London; Institute for Development Studies, Jaipur; Delhi School of Economics; Indian Statistical Institute (New Delhi); the University of Johannesburg, and the University of Warwick, where an earlier draft of this paper was presented, have been very helpful. Staff of the Indian Railway Board library and offices were helpful during the data collection process. Needless to add, we are responsible for all errors and omissions. Final revision accepted: May 20, 2014.

reach such positions in the absence of India's reservation policies.

Our approach to analyzing the effect of reservations on productivity in the IR is as follows. First, we estimate total factor productivity (TFP) in each zone-year using a Cobb–Douglas production-function framework, accounting for zone-level fixed effects and employing the Levinson–Petrin correction for simultaneity (i.e., the possibility that input use could itself be influenced by anticipated productivity shocks). In some specifications we include a measure of SC/ST% as an independent variable and examine its significance. In other specifications we proceed to a second stage, in which we either regress the TFP estimate on an SC/ST% variable or correlate it with an SC/ST% variable, and then examine the significance of the result.

As an alternative to traditional production function analysis, we make use of the non-parametric Data Envelopment Analysis (DEA) technique, which requires no *a priori* assumptions about the functional form of production relations and which allows for more disaggregation of input and output variables than is possible in production function analysis. [7] We use DEA to generate estimates of annual rates of change of TFP (henceforth "ΔTFP"), and then we examine whether variation in ΔTFP is related to variation in any SC/ST% variable.

The key findings of or study may be summarized as follows. The production function and data-envelopment analyses provide no evidence in support of the claim that higher proportions of jobs filled by SC/STs are associated with lower total factor productivity or its annual rate of change. Furthermore, under some specifications, higher proportions of SC/ST employees in high-level positions—who are most likely to be AA beneficiaries—are positively associated with higher TFP or ΔTFP. These findings resonate very strongly with studies assessing the impact of workforce diversity on enterprise productivity in the US, which have found either a positive or null effect, but no evidence of a negative effect (Barrington & Troske, 2001).

Our interpretation of the results of this empirical analysis might be contested on the grounds that we have not actually identified the causal relationship at issue. If SC/ST% were itself influenced by a productivity variable, or if both these variables were influenced by other relevant variables omitted from our analysis, then our statistical results could not be interpreted as suggesting the presence or absence of an impact of SC/ST% on productivity. We therefore examined in some detail the processes by which IR jobs are filled, and we considered several specific ways in which SC/ST% in the IR might be thought to be a function of IR productivity or of omitted variables reflecting SC/ST education or ability. We also addressed the concern that our quantitative measures of IR output—and hence productivity—do not encompass potentially qualitative aspects of IR performance that might be especially sensitive to the competence of railway employees. Our analysis of these issues gives us greater confidence that we can interpret the statistical findings of this study as shedding light on the effect of affirmative action on productivity in the IR.

The rest of this paper is organized as follows. In Section 2, we briefly describe the Indian Railway system and discuss the way in which we have compiled the data available from the IR; we pay close attention to the relationship between reservation policies and our SC/ST% variables. In Sections 3 and 4, we explain our production-function and DEA analyses, respectively, and we present the results of these analyses. In Section 5 we address several possible alternative explanations of our findings; and in Section 6, we consider the concern that we have failed to capture key qualitative aspects of IR performance. The concluding Section 7 returns to the general debate about the impact of AA on productivity: we suggest some mechanisms that could explain our findings in the case of the Indian Railways, and we discuss the implications of our analysis for other countries in which AA policies have been or may yet be introduced.

2. AN OVERVIEW OF THE DATA

As noted above, the IR is divided for administrative convenience into regional zones. [8] From 1952 through 2002, there were nine zones in operation:: Central Railway (CR), Eastern Railway (ER), Northern Railway (NR), North-Eastern Railway (NER), North-East Frontier Railway (NFR), Southern Railway (SR), South Central Railway (SCR), South Eastern Railway (SER) and Western Railway (WR). Because separate data on SC/ST employment were not available for the NR, we had to drop that zone from our database; and because of insufficient data availability prior to 1980, our time horizon for analysis was limited to the period from 1980 to 2002.

The IR as a whole in recent years has been operating about 9,000 passenger trains, which transport 18 million passengers daily; its freight operations involve the transport of bulk goods such as coal, cement, foodgrains, and iron ore. The IR makes around 65% of its revenues, and most of its profits, from the freight service; a significant part of these freight profits are used to cross-subsidize passenger service, enabling it to charge lower fares to consumers. During the period from 1980 to 2002, IR gross receipts (earned from passenger and freight traffic) grew consistently from 26 to 411 trillion rupees at current prices; this represents a fourfold increase at constant prices.

While total track kilometers in the Indian Railway system increased modestly from 104,880 km in 1980 to 109,221 in 2002, the percentage of electrified routes increased more rapidly, from just 7% to more than 20%. Coal had long been the main source of fuel for the IR; but by 2002 almost all IR's operations were fueled by more efficient (and less polluting) diesel or electric power. Since the 1980s there have also been significant technological improvements in the form of track modernization, gauge conversion, and upgrading of signaling and telecommunications equipment. In the 1990s the IR switched from small freight consignments to larger container movement, which helped to speed up its freight operations.

In specifying the variables needed for our production-function and data-envelopment analyses, we sought as far as possible to make use of physical rather than value measures. We did so because the IR is not a profit-oriented enterprise. While it does seek to cover its costs, it has numerous politically-determined objectives—as reflected in the cross-subsidization of passenger by freight traffic—that make profitability a poor standard by which to evaluate IR performance, and that lead to pricing decisions that do not necessarily reflect the marginal cost or benefit of the commodity in question. In the following paragraphs, we describe in broad terms how we defined and measured the variables used in our analyses. [9]

(a) *Output variables*

The output produced by the Indian Railways consists of passenger service and freight service, measured physically in terms of passenger-kilometers (PK) and net ton-kilometers (NTK), respectively. For both passenger and freight service, the IR also provides data on revenues from each type of passenger rail service and each type of transported commodity. All of the data are available in annual time series for each zone as well as for all-India.

We first generated time series indices for total passenger output and total freight output from underlying time series for passenger and freight transport of different types. We then generated time series indices for total railway output by weighting the indices for total passenger output and total freight output according to their percentage of total railway revenue generated. The construction of the final zonal indices was done so as to reflect the scale differences between zones.

Although we believe that the above measures of railway output, based on physical measures, are superior to any value measures of railway output, we do recognize that industry outputs are most often measured in terms of gross revenue or value added. We therefore compiled data on railway revenue at current prices and deflated these data to obtain an alternative constant price time series for total railway revenue. We could not work with the value-added variable since we did not have data on non-fuel material inputs.

(b) *Labor variables*

The Indian Railways, like all departments of the government, hire workers in four different labor categories: categories A and B are the top two tiers of employees, comprising of administrative officers and professional workers; category C includes semi-skilled and clerical staff; and category D includes relatively unskilled attendants, peons and cleaning staff. The total employment figures provided by the IR serve as raw measures of the overall volume of labor input, [10] but they fail to reflect changes in the average quality of labor that result from changes in the category-composition of the labor force. We posit that the average quality of labor improves to the extent that the category-composition of jobs (the pattern of A + B, C, and D employment) shifts in the direction of a greater proportion of higher-skilled jobs. In order to take account of the effects of changes in the average quality of labor over time, we constructed time series indices for a new variable measuring the volume of "effective labor" input for each zone (and for all-India).

For the purposes of our analysis we need to be able to distinguish AA beneficiaries from other employees. The IR provide data on the number of employees in job categories A + B, C, and D, who declare themselves to belong to a Scheduled Caste or Tribe. Such a declaration is necessary for an SC or ST applicant to avail of reservations. This is in fact done by almost all SC/ST applicants for category A + B jobs, because they know that they are unlikely to score high enough on the required exams to gain access to a non-reserved job. In the event that an SC/ST applicant for a reserved job actually scores higher than the cut-off for a non-reserved job, he or she is not included in the count of SC/ST employees. [11] The available data on SC/ST employment in the high-skilled A + B category therefore measure fairly accurately the corresponding number of AA beneficiaries. However, many SC/ST applicants for the low-skilled D-category jobs do not avail of reservations, because they know they can meet the qualifications for a non-reserved job. [12] Thus the IR data on SC/ST employment in D jobs significantly over-estimate the number of AA beneficiaries, because many self-declared SC/ST employees would have been hired even in the absence of reservations.

The ultimate objective of our quantitative analysis is to examine the relationship between AA and productivity in the IR. Toward this end, in both the production-function and the data-envelopment analyses, we worked with two measures of the SC/ST percentage: the first is the ratio of all SC/ST employees to total employees, and the second is the ratio of SC/ST employees to total employees in labor categories A + B only. The second variable is considerably more accurate in measuring the percentage of AA beneficiaries than the first, because the latter is biased significantly upward by the over-estimation of the number of SC/ST employees in D-level jobs, which account for roughly half of all-India IR employment. [13]

There might appear to be reason for skepticism about our use of variation in SC/ST% to measure the impact of India's reservation policies on IR performance. One could argue that differences in SC/ST% across zones and over time might simply be due to the fact that, in certain zones and in later years, there was a greater number of SC/ST candidates who met the minimum qualifications for being hired for or promoted to an IR job than in other zones and in earlier years. Even in the absence of reservations one would expect to find higher SC/ST% values in some zones than others, and higher SC/ST% values in later than in earlier years, and these differences would be attributable to differences in the ability levels of individual SC/ST job candidates. Most important for our analysis, however, is the fact that in the absence of reservations the percentage of SC/ST employees in A + B positions would have been very low, because few SC/ST applicants would have had formal job qualifications as good as those of their non-SC/ST peers. [14] The actual SC/ST% in A + B positions therefore measures the extent to which apparently lower-qualified AA-beneficiary SC/ST employees have displaced apparently higher-qualified non-SC/ST would-be employees. It is this displacement that most worries critics of India's reservation policies. Since AA manifests itself precisely in this displacement, the SC/ST% in A + B positions measures very well the impact of AA for the purposes of our analysis.

Finally, when we examined graphs of each of the eight zonal time series for the two SC/ST% variables, we discovered that there were some distinctly outlying observations that appear to have been subject to measurement error. We therefore generated a second set of pooled time series in which data for roughly a dozen zone-years were dropped from the full set of 184 observations because of highly questionable values for one or both of the SC/ST% variables.

(c) *Capital variables*

The IR distinguishes between three types of capital stock—structural engineering, rolling stock, and machinery and equipment—and makes available annual current-price data on book value and gross investment for each type of capital, going back to 1966 for each zone and to 1952 for all-India. We chose to work with estimates of gross rather than net capital stock, because measures of net capital stock decline in value as the number of its productive future years decline, whereas measures of gross capital stock tend to be proportional to the capital value actually consumed during a given year. Book value data on capital stock are notoriously poor measures of the value of capital inputs, because they aggregate annual additions to capital stock that are valued at different prices every year; so we made use of the perpetual inventory method (Christensen & Jorgenson, 1969) to generate time series of constant-price gross capital stock of each type.

Constant-price gross capital stock measures do have one important shortcoming, in that they fail to reflect the extent to which embodied-in-capital technological progress increases the productive potential of a piece of constant-price capital stock from year to year. Just as we sought to adjust a raw measure of labor input (employment) to take account of changes in labor quality associated with the category-composition of labor in generating a better measure (volume of effective labor), so we found it desirable to adjust our raw measure of capital input (constant-price gross capital stock) to take account of changes in capital quality associated with the age structure of capital to generate a better measure that we call "effective capital input." We based these calculations on the assumption that a unit of constant price gross investment loses 1% of its productive value for each year elapsed since it was created. [15]

(d) *Material input variables*

The main material input used by a railway system is fuel. Using standard conversion factors to convert all the measures of fuel in physical terms into their equivalent in coal-tonnes, we compiled time series of total coal-tonnes of fuel input for each zone and for all-India from 1980 through 2002.

In the case of fuel input, as with capital and labor inputs, we saw reason to generate a second, more nuanced variable to take account of changes in fuel quality associated with changes in the proportions of different kinds of fuel utilized by the IR. In particular, diesel- and electricity-powered locomotion is significantly more efficient than locomotion powered by other fuels, because it enables greater acceleration, allows for easier maintenance, and generates less pollution. [16] We sought therefore to construct an index of "effective fuel" that would take account of the extent to which locomotion is powered by the more efficient fuels.

3. PRODUCTION FUNCTION ANALYSIS

Although 22.5% of jobs in the IR are reserved in principle for SC/ST members, the actual percentage of SC/ST employees varies a lot by type of job, by zone and by year, because of the failure to fulfill quotas (in the case of high-level jobs) and the higher percentage of SCs and STs in low-level jobs, as explained in the previous section. The variation in SC/ST employee proportions facilitates econometric estimation of the impact of SC/ST employees on productivity.

Using the variables described in Section 2, [17] we estimated log-linear Cobb–Douglas production functions of the following form. [18]

$$\ln(\text{output}) = \beta_0 + \beta_1 \ln(\text{capital stock}) + \beta_2 \ln(\text{labor}) + \beta_3 \ln(\text{fuel}) + \beta_4 \text{ time} + u$$

where time was introduced to capture the effect of technical change and u is the error term. Our panel is a balanced macro panel, with the number of zones $(N) = 8$ lower than the number of time periods $(T) = 23$. We used fixed effects (*FE*) estimation, in order to control for time-invariant zone-specific unobservable factors.

We carried out a variety of different regression runs, which varied in terms of the variables included in the specification of the production function and/or the ways in which those variable were measured. Our runs varied along the following dimensions:

1. Which dependent variable we include in the regression: a physical measure, total output (q), or a value measure, total revenue (r). We believe that "q" is the more reliable measure, because "r" is dependent on pricing decisions and the IR, as a quasi-monopoly, does not set prices competitively.
2. Which measure of the three inputs we use as independent variables in the regression: the adjusted measures of effective labor (el), effective capital stock (ek), and effective fuel (ef); or the raw measures of total employment (l), unadjusted capital stock (k), and unadjusted fuel (f)—which in our view are considerably less accurate, because they fail to take account of differences in quality between different subcategories of each input.
3. Whether we include or exclude a variable representing the percentage of SC/ST employees among all employees ("%SC_ALL"), or alternatively the percentage of SC/ST among category A + B employees ("%SC_AB"), as an independent variable in the regression equation
4. Whether to include all 184 zone-year observations that we compiled, or to exclude zone-years in which the figures we had for the variable representing an SC/ST percentage of employees were highly questionable We found 15 observations of "%SC_ALL" that were highly questionable, and 12 observations of "%SC_AB" that were highly questionable (mostly for different zone-years in the two cases). We believe that the regressions and correlations in which the questionable zone-years are excluded provide more reliable results.

Thus we estimated the following specifications:
Specification 1: $\ln q$ on $\ln ek$, $\ln el$, $\ln ef$, t
Specification 2: $\ln q$ on $\ln k$, $\ln l$, $\ln f$, t
Specification 3: $\ln r$ on $\ln ek$, $\ln el$, $\ln ef$ and t
Specification 4: $\ln r$ on $\ln k$, $\ln l$, $\ln f$ and t
Specification 5a: $\ln q$ on $\ln ek$, $\ln el$, $\ln ef$, %SC_ALL, and t
Specification 5b: $\ln q$ on $\ln ek$, $\ln el$, $\ln ef$, %SC_AB and t.

In the case of specifications 5a and 5b, we ran the regressions first with all zone-year observations included and then with the questionable zone-years excluded; the latter runs were labeled 5ax and 5bx. In the case of specifications 1–4, we can interpret the residuals in each zone-year as estimates of total factor productivity (TFP), and we then needed as a second step to examine the relationship between variation in TFP and variation in SC/ST%. This we did by correlating TFP with each of four SC/ST% variables: "%SC_ALL" and "%SC_AB" (including all zone-year observations) and "%SC_ALLx" and "%SC_ABx" (excluding the questionable zone-year observations). We also regressed TFP on each of the four SC/ST% variables.

For the reasons given under points #1, #2, and #4 above, we believe that our most reliable specifications are 1, 5ax, and 5bx; the results of these runs are given in Table 1. To assure readers of the robustness of our findings, however, we report the results of all specifications of our FE regression runs in Appendix Table 3. One can see that the coefficients on the independent variables fluctuate within a fairly narrow band when the dependent variable is $\ln q$, and they do the same when the dependent variable is $\ln r$. In the cases of specifications 1 through 4, neither the correlations nor the regressions of TFP with the four alternative measures of SC/ST% generated a significant coefficient. In the cases of specification 5a(x) and 5b(x), we found that the coefficient on the independent SC/ST% variable in the productivity regression was significantly positive for our best SC/ST% variable—"%SC_ABx."

(a) *The simultaneity problem*

In our production function analysis up to this point, we did not take account of the possible correlation between input levels and productivity due to the fact that firms generally respond to anticipated changes in productivity by changing their usage of factor inputs. Because input levels may be

Table 1. *Production function estimates*

	Lev–Pet (1)	Lev–Pet (5a)	Lev–Pet (5b)	FE (1)	FE (5ax)	FE (5bx)
Constant				−4.50	−4.55	−2.63
				(10.3)	(9.07)	(10.4)
ln *ek*	**0.88**	**0.93**	0.79	.075	0.119	0.004
	(0.36)	(0.35)	(0.44)	(0.50)	(0.44)	(0.47)
ln *el*	0.184	0.177	0.216	0.41	0.35	0.39
	(0.37)	(0.50)	(0.52)	(0.31)	(0.27)	(0.29)
ln *ef*	.000	.000	.000	0.022	0.017	−0.006
	(0.37)	(0.40)	(0.42)	(0.12)	(0.13)	(0.11)
time				**0.048**	**0.048**	**0.049**
				(0.021)	(0.019)	(0.20)
%SC_ALL		−0.19				
		(1.03)				
%SC_ALLx					−0.47	
					(0.55)	
%SC_AB			0.32			
			(0.84)			
%SC_ABx						**0.87**
						(0.33)
CORRELATIONS						
%SC_ALL	−0.02			−0.01		
%SC_ALLx	−0.05			−0.04		
%SC_AB	0.12			−0.07		
%SC_ABx	**0.19**			0		

Note: In the case of the Lev–Pet estimation method, the constant term and the time variable are incorporated into an earlier stage of the analysis and thus do not appear in the production function regression. Standard errors are in parentheses. Figures in bold are significant at 5%.

correlated with unobserved productivity shocks, our independent variables may be correlated with the error term, which would result in biased OLS estimates of the production function. This is the well-known simultaneity (or endogeneity) problem in production function estimation.

Many alternatives to OLS have been proposed to address this problem; for instance, Olley and Pakes (1996) have derived conditions under which use of an investment proxy variable eliminates variation in the error term that could be related to unobserved productivity shocks. Levinsohn and Petrin (2003) (henceforth Lev–Pet) modify this approach by using intermediate inputs rather than investment as a proxy, since this has some advantages over the Olley and Pakes method. In particular, it is less costly to adjust material inputs than investment in response to productivity shocks and, unlike investment which could often be zero, material inputs are always non-zero. Intermediate inputs are thus considerably more sensitive to productivity changes, and their use as a proxy is more likely to control for variation in the error term due to unobserved productivity shocks.

A detailed exposition of the Lev–Pet method can be found in their paper; here we summarize the essence of their method. Their estimation method takes account of unobserved productivity shocks by treating the error term of a standard production function regression equation as the sum of two components—a transmitted productivity component and an error term uncorrelated with input choices. They show that, under the reasonable assumption that the demand function for an intermediate input is monotonically increasing in the unobserved productivity component, that demand function can be inverted. This allows one to model the unobserved productivity component as a function of an intermediate input variable and a state variable such as capital stock.

We ran Lev–Pet regressions for all of the specifications of our productivity model for which we ran fixed-effect regressions. In the case of specifications 1–4, we again interpreted the residuals in each zone-year as estimates of TFP, and—where appropriate—we went on to correlate TFP with each SC/ST% variable and to regress TFP on each such variable.

Appendix Table 4 shows the complete set of results of our Lev–Pet estimations. The Lev–Pet literature indicates that any result in which an input coefficient turns out to be exactly 1 should be discarded as invalid, so we had to set aside specifications 2–4 and 5ax. This left us with three usable specifications—those numbered 1, 5a, and 5b. The results for specifications 5a and 5b did not generate significant coefficients on the independent SC/ST% variable. In the case of specification 1, we did find a statistically significant positive correlation, as well as a statistically significant positive regression coefficient, for "%SC_ABx"—our best SC/ST% variable. The usable Lev–Pet estimation results thus provide no evidence of a negative effect of SC/ST employment on productivity, and they provide some support for a positive effect on productivity of the SC/ST percentage of workers in A + B jobs.

In Table 1 we bring together the three usable results of our Lev–Pet estimations, which correct for simultaneity bias, as well as the results from the three most reliable specifications of our FE regressions, which account for zone-specific unobservable factors. This table illustrates our findings succinctly. All of our production function results, taken together, reject the hypothesis that higher proportions of SC/ST employees in A and B jobs contribute negatively to productivity levels in the Indian Railways. Indeed, they provide some evidence that higher proportions of SC/ST employees in A and B jobs—predominantly beneficiaries of affirmative action—may actually contribute positively to IR productivity.

4. DATA ENVELOPMENT ANALYSIS

As explained in Section 2, we also tried an alternative approach to investigate the impact of AA on productivity in the Indian Railways: a two-stage procedure in which the first stage was the use of the non-parametric method called Data Envelopment Analysis (DEA) of productivity changes, and the second stage was an econometric analysis of factors potentially influencing those productivity changes. DEA allows one to analyze productivity in the context of a pooled data set of time series data on inputs and outputs for multiple production units within a given industry. It does not require specification of any particular functional relationship between input and output variables; and it allows one to work with more than one output variable as well as multiple input variables. In essence it fits a frontier that represents maximum technical efficiency, enveloping the outermost data points; the distance of each zone-year observation from the frontier provides a measure of its technical efficiency. [19]

For the first stage of our alternative approach we tried two different variants. In Variant One, we initially used DEA to estimate annual changes in total factor productivity (ΔTFP) from 1980–81 to 2001–02 in each railway zone, taking into account two output variables (passenger transport and freight transport) and eight input variables (employment in each of the four labor categories A, B, C, and D; constant-price gross capital stock of each type—structural engineering, rolling stock, machinery & equipment; and total fuel input (in coal-tonnes). Then for the second stage we sought to explain our estimated ΔTFP values (for each zone and pair of years) in terms of several variables that appeared likely to influence annual total factor productivity change. The independent variables consisted of three that were designed to capture the quality of the three types of inputs (labor, capital, and fuel) and one to reflect the scale of production. For labor quality we used%SC_AB or %SC_ABx, since our primary focus is on the impact of SC/ST as opposed to other labor on productivity; [20] for capital we used the average vintage of gross capital stock (of all types); and for fuel we used fuel quality (the share of coal-tonnes of fuel accounted for by diesel oil and electricity). For the scale of production, we used our aggregated measure of total railway output. We regressed the estimated values of ΔTFP (from year "t" to year "$t + 1$") on the four independent variables (measured in year "t"), thus pooling 22 time series observations for each zone. We conducted tests for choosing between RE/FE and serial correlation; the tests indicated the use of FE estimation with no significant presence of serial correlation.

Subsequently we undertook a slightly different variant of our alternative approach. For the first stage of Variant Two we did a new DEA run in which we used the "effective" measures of the capital stock and fuel input variables instead of the unadjusted "raw" measures of the first variant. In other words, we incorporated the "quality" of the capital stock and fuel inputs into the first stage of the analysis, making it unnecessary to consider them in the second stage. For the second stage of this variant we simply correlated the estimated ΔTFP values (from year "t" to year "$t + 1$") from the first stage with the various SC/ST variables (measured in year "t"). [21]

For each of the variants of our two-stage DEA-based approach we undertook two separate analyses—one including observations for all eight zones, and the other including observations for seven zones, excluding the NFR zone. The reason for excluding this zone is that the figures for NFR constant-price gross rolling stock indicated a substantial and implausible decline throughout the period 1980–2002; in no other zone did we encounter such an implausible trend for any variable. [22] All of the first-stage and second-stage results we obtained for each DEA variant are available on request from the authors.

The key results of our DEA-based analyses are those that indicate the extent to which variation in ΔTFP is associated with variation in SC/ST%—the latter measured alternatively by %SC_ALL and %SC_AB, or (excluding highly questionable observations in the underlying data) %SC_ALLx and %SC_ABx. In the case of our first variant, using raw measures of capital stock and fuel inputs and undertaking a second-stage regression analysis of ΔTFP, the association is given by the estimated coefficient on the SC/ST variable. In the case of the second variant, using effective measures of capital stock and fuel inputs, the association is given by the correlation of ΔTFP with the SC/ST variable. The key results we obtained are given in Table 2. [23]

Table 2 indicates that, under Variant One, the 2nd-stage regression runs to explain ΔTFP yielded positive coefficients on the SC/ST variables in all but one of the eight cases, but no coefficient was even close to being significant at 5%. The correlations under Variant Two are all positive, though in five of the eight cases they were not significant at 5%. There is clearly no support here for the claim that higher proportions of SC/ST employees result in slower growth in total factor productivity.

Under Variant Two, significant positive correlations with ΔTFP were obtained for %SC_ABx in the 8-zone case and for both %SC_AB and %SC_ABx in the 7-zone case. In particular, the correlation in the case of %SC_ABx in the 7-zone case is 24% and the p-value just .4%, reflecting a remarkably high level of significance. This is especially noteworthy because we have every reason to believe that the results of 7-zone runs are more reliable than the results of 8-zone runs and that the %SC_ABx tests are more reliable than the %SC_AB tests. [24] Thus here we find evidence in support of the claim that higher proportions of SC/ST employees in A and B jobs contributes to more rapid total factor productivity growth, reinforcing our conclusion from the previous section that affirmative action has if anything improved productivity in the Indian Railways.

5. ALTERNATIVE EXPLANATIONS OF OUR FINDINGS

We have found that there is little statistical relationship between the SC/ST percentage of employees and total-factor productivity (or its growth) in the Indian Railways, after controlling for factor inputs and technological change, except for some cases of a positive association between the SC/ST percentage of upper-level (A + B category) employees and productivity. We believe that these findings imply that increasing the share of SC/ST employees via reservations in the IR does not impair productive efficiency—and may in some cases actually increase efficiency.

To sustain the latter interpretation, we must show that we have actually identified a causal relationship running from the percentage of SC/ST employees to productivity. This may be questioned, however, on the ground that our SC/ST% variables are not exogenous, but instead influenced by IR productivity and/or by omitted variables that are simultaneously

Table 2. *Association of ΔTFP with SC/ST variables*

	Variant #1		Variant #2	
	8 zones	7 zones	8 zones	7 zones
%SC_ALL	0.2	0.32	0.1	0.02
	(0.32)	(0.22)	(0.171)	(0.788)
%SC_ALLx	0.24	0.4	0.13	0.1
	(0.38)	(0.26)	(0.096)	(0.213)
%SC_AB	−0.03	0.12	0.13	**0.17**
	(0.3)	(0.18)	(0.079)	(0.028)
%SC_ABx	0.01	0.21	**0.15**	**0.24**
	(0.44)	(0.27)	(0.049)	(0.004)

Note: Figures in parentheses under variant #1 are standard errors; figures in parentheses under variant #2 are *p*-values.
Figures in bold are significant at 5%.

affecting both SC/ST% and productivity. If this were the case, it could be argued that findings of an association between SC/ST% and productivity—or findings of no such association—reflect forces other than the effect of changing percentages of SC/ST employment. [25] To address this concern one must first of all understand the processes that determine the variation in percentages of SC/ST employees across zones and over time in each zone.

In any given zone-year, the number of employees in any given job category (level A, B, C, or D) is the consequence of additions via new hiring, promotion (from a lower category), or transfer (from another zone) and subtractions via retirement, promotion (to a higher category), or transfer (to another zone). New hiring to fill vacancies is done on the basis of a "post-based roster", whereby a vacancy is designated either as an unreserved seat or a seat reserved for a particular AA-beneficiary group (e.g., SC or ST). [26] Candidates for placement in category-A positions have to take an all-India civil services exam. Allocation to IR zones depends primarily on the availability of seats in his/her group and his/her position in the exam score ranking; there are also transfers across the various high-level administrative services. Placement in category B positions is done entirely via promotion of category C employees, on the basis of seniority and/or competitive exams. Recruitment of group C and D employees are handled in a decentralized manner by regional recruitment boards (RRBs) and railway recruitment cells (RRCs), respectively; a candidate can apply to any of these but, if successful, will be placed in the zone to which the RRB or RRC is linked. In all cases where exams come into play, the cut-off scores in any given job category are significantly lower for applicants for reserved seats than for applicants for unreserved seats. [27] Retirements depend mainly on how many employees happen to reach the mandatory retirement age in a given year. Finally, the number of employees in any given zone-year seeking promotion to a higher-level job, or transfer to a different zone, varies on the basis of a variety of person-specific factors affecting decision-making by individual IR employees. It is therefore no surprise that there is considerable variation in the SC/ST% variables both between zones and from year to year—the latter at times upward and at times downward.

In view of all the considerations raised in the previous paragraph, it would certainly appear that variation across zones and over time in the percentage of SC/ST employees in any given job category is largely unrelated to any factor systematically related to productivity. It behooves us nonetheless to consider several specific concerns raised by critics about the possible endogeneity of our SC/ST% variables.

(a) *Reverse causality*

Is it not possible that some zones, or some years, are characterized by better management, which achieves higher total-factor productivity and also attracts higher-quality SC/ST employees (or is more willing to respond to AA-based pressures to hire more SC/ST employees)? For example, the central IR authorities might be allocating the highest-scoring SC/ST employees to the most productive zones in order to maximize overall IR output. Or the management of any given IR unit might be simultaneously balancing two goals: to hire or promote more SC/ST employees in order to satisfy AA requirements, and to produce output in an efficient manner, in which case managers' willingness to hire SC/ST employees would vary positively with the residual productivity of that IR unit. [28] As a concrete example, those units that are most confident in having non-SC/ST staff capable of covering for any SC/ST failures could be most disposed to hire SC/STs to high-level positions. Alternatively, it might be the case that units in zones with stronger economies face greater demand for their services, enabling them to operate more efficiently and also increase their proportion of SC/ST workers more rapidly. [29] In all of the above cases, the line of causality would run from productivity to the SC/ST percentage of employees, rather than the other way round.

These hypotheses are unlikely to stand up to scrutiny, however, in light of the actual processes of selection and allocation of labor within the IR. For one thing, IR central authorities have only limited influence over the selection of A-level employees, who must first pass general Indian civil service examinations and who then may or may not wind up in the service of the IR, depending on available vacancies, their own preferences for areas of service, and their exam scores. At all other levels, candidates may choose where to apply for their initial job, and after holding a job in a given zone they may be transferred to another zone—voluntarily or involuntarily. More importantly, the allocation of SC/ST employees across zones depends in considerable part on the distribution of vacancies in reserved SC/ST jobs, which in turn depends on prior patterns of promotion, transfer, or retirement. There is therefore very little leeway in the multi-tiered IR employee recruitment structure for either central IR authorities or zonal managers to control the allocation of SC/ST employees.

(b) *Omission of changes in SC/ST ability over time*

Is it not possible that any correlation over time between productivity and the percentage of SC/ST employees reflects the fact that both are increasing over time—the latter because of improvements over time in the unobserved average ability of SC/ST employees, rather than increases in their numbers due to IR reservations? Although the exam cut-off scores for SC/ST candidates are always considerably lower than for non-SC/ST candidates, these thresholds have been rising over time as the average level of educational achievement has been rising for both SC/ST and non-SC/ST Indians. Inclusion of measures of average ability in SC/ST and non-SC/ST employees might reduce positive correlations of SC/ST employee percentage and productivity to insignificance and might turn insignificant correlations negative.

This concern appears to be plausible, since it is certainly true that the average qualifications of SC/ST candidates have been trending upward with each passing year—thanks in part to India's policy of reservations in higher educational institutions. However, technological progress over time is controlled for in our regression equations, so there is no *a priori* reason to expect that total-factor productivity will also be rising steadily from year to year—and indeed our estimates show no such time trend in productivity.

(c) *Omission of differences in education across zones*

Is it not possible that higher proportions of SC/ST employees are hired or promoted in zones whose regions provide a greater quantity and/or quality of education? Such zones would then be characterized both by increased productivity and by greater numbers of SC/ST applicants who meet the minimum qualifications for being hired or promoted to high-level IR positions. Inclusion of measures of zonal education might therefore reduce positive correlations of SC/ST employee percentage and productivity to insignificance and might turn insignificant correlations negative.

The plausibility of this concern is heightened by the fact that there are significant differences in the quantity and quality of education as between different states in India. It is nonetheless doubtful that the omission of a zonal education variable could be a significant problem, for two reasons. First, zones contain within them inter-state heterogeneity in education. Second, the processes by which railway jobs are allocated at the different levels weaken the connection between zonal education and SC/ST%. Category A officers are recruited through national examinations and allocated to zones as explained above; there is no reason to believe that these employees end up working in the state where they did their schooling. Category C (some of whom are promoted to category B) and category D recruits are hired through the RRBs and RRCs, which have regional jurisdictions, but the link between the place of work and place of schooling is not clear-cut. Finally, within each category, part of the work force consists of inter-zonal transfers. For all employees who are transferred from other zones, there is no link between education in the regions under the jurisdiction of that zone and the productivity of employees in that zone.

Our rejection of alternative explanations of our findings in this section does not enable us to claim that the correlations we have found between SC/ST percentage of employees and total-factor productivity in the IR show definitively that higher proportions of SC/ST employment have not hurt—and may even have helped—productivity. But we do think that skeptics of our conclusion now bear a heavier burden of refutation.

6. QUALITATIVE ASPECTS OF EFFICIENCY

Considering our econometric analysis of railway productivity, one might be concerned that the measures of output we use are mostly quantitative in nature and that such measures fail to take account of possible changes in the *quality* of railway output. Elements of railway output quality include timeliness of arrivals at destinations, passenger comfort, and overall safety (i.e., freedom from accidents). Conceivably growth in the proportion of SC/ST labor could lead to a diminution in the quality of railway services provided, even as the quantity of services was not adversely affected.

Responding to this concern, we note first that one of our output measures—railway revenue—does reflect quality as well as quantity, insofar as higher quality is reflected in higher prices charged for railway services. Moreover, our quantitative measure of railway passenger output also reflects an element of quality because we weight the growth of passenger output according to class of service—thus giving more weight to the higher classes that provide more comfortable service. Nonetheless, it would be desirable to incorporate more fully into our work various indicators of timeliness, comfort, and accident-free service. We were not able to find systematic data on such indicators for the years and the zones of our statistical analysis; but we do recommend further research along these lines. [30]

Railway accidents, though rare, are obviously an important source of poor railway performance; they generate adverse consequences that go far beyond the loss of damaged equipment and the failure to complete a planned passenger or freight trip. Moreover, as noted in the introductory section, critics of reservation policies have suggested that higher proportions of SC/ST labor might well result in higher frequencies of railway accidents. We therefore thought it useful to see if trends in Indian Railway accident rates could be related in any way to trends in SC/ST labor percentages.

Correlating the all-India yearly railway accident rate (the total number of accidents per million train kilometers) over the period of our study (1980–2002) with the corresponding all-India figures for the percentage of SC/ST employees in total employment, we found correlation coefficients of $-.69$ for all employees and $-.93$ (both correlations significant at 1%) for employees in the upper-level A and B categories. [31] The second, higher correlation is the most relevant, both because IR employees serving in management and professional positions are especially responsible for guarding against accidents and because the data on SC/ST employees in the C and D categories fail to count many SC/ST employees who do not declare themselves as such.

Our finding of a highly significant negative correlation between the all-India accident rate and the SC/ST percentage of A + B-category employment results from the fact that the former has been declining and the latter rising (both fairly steadily) over the last few decades. This is strong evidence that higher SC/ST employment proportions are not resulting in higher accident rates—unless, of course, there are other likely determinants of the accident rate that have also shown steady trends (and the appropriate sign) over the same period. The most plausible alternative explanations for decreasing accident rates are increasing electrification of signals, improvement in track quality, and safer track crossings (including better-guarded level crossings and more bridges over tracks). There is indeed evidence of positive time trends in each of these alternative determinants (see GOI, Ministry of Railways, 2005–06 Yearbook, especially pp. 18–25). There are insufficiently detailed data, however, to include such variables in a multivariate regression analysis of accident rates. While such an analysis might well counter the notion that higher SC/ST employment proportions actually promote greater safety, it seems unlikely that it could undermine the conclusion that higher SC/ST employment proportions do no harm.

7. CONCLUDING OBSERVATIONS

Analyzing an extensive data set on the operations of the world's largest employer subject to affirmative action, the Indian Railways, we have found no evidence whatsoever to support the claim of critics of AA that increasing the proportion of AA beneficiaries adversely affects productivity or productivity growth. On the contrary, some of the results of our analysis suggest that the proportion of SC/ST employees in high-level positions (at the A and B job levels) is positively associated with IR productivity and productivity growth.

Our finding of such positive associations in the case of SC/ST employees in A and B jobs is especially relevant to debates about the effects of AA, for two reasons. First, the efficacy with which high-level managerial and decision-making jobs are carried out is likely to have a considerably bigger impact on overall productivity than the efficacy with which lower-level semi-skilled and unskilled jobs are fulfilled. Thus critics of AA are likely to be much more concerned about the potentially adverse effects of favoring SC/ST candidates for A and B jobs than for C and D jobs. Second, it is precisely in the A and B jobs—far more than in C and D jobs—that reservations have been indispensable for raising the proportion of SC/ST employees. Even without reservations, one would expect substantial numbers of SC/ST applicants to be hired into C and D jobs; but without reservations very few SC/ST applicants would have been able to attain jobs at the A and B level.

The results that we have obtained from our analysis of productivity in the Indian Railways are quite suggestive for other developing economies in which AA is practiced. These results are consistent with results obtained from productivity studies in the United States, in that there is no statistically significant evidence that AA in the labor market has an adverse effect on productivity. Our results are stronger, however, in that we do find some suggestive evidence that AA at the upper levels of the labor market actually has a favorable effect in contributing to greater productivity.[32]

It is beyond the scope of our paper to explain just how and why AA in the labor market might have such a favorable effect. We can, however, adduce some relevant evidence from research carried out by others. A number of studies in India and elsewhere indicate that hiring practices are often far from meritocratic in the absence of AA. For example, in a study of modern urban Indian highly-skilled labor markets for private sector jobs (often assumed to be among the most meritocratic), Deshpande and Newman (2007) show how caste and religious affiliations of job applicants shape employers' beliefs about their intrinsic merit. This confirms the findings of other studies that uncover labor market discrimination and point out how social identities impact hiring and wage offers (for instance, Bertrand & Mullainathan, 2004; Pager & Western, 2005; Siddique, 2008; Thorat & Attewell, 2007, and Jodhka & Newman, 2007).

Furthermore, there are numerous *a priori* reasons to expect that AA in hiring might improve economic performance—particularly in high-level jobs. Individuals from marginalized groups may well be especially highly motivated to perform well when they attain decision-making and managerial positions, because of the fact that they have reached these positions in the face of claims that they are not sufficiently capable, and they may consequently have a strong desire to prove their detractors wrong. Or such individuals may simply believe that they have to work doubly hard to prove that they are just as good as their peers, so they may actually work harder. To have greater numbers of managers and professionals from disadvantaged groups working in high-level positions might well increase productivity because their backgrounds make them more effective in supervising and motivating other workers from their own communities.[33] Finally, improvements in organizational productivity may well result from the greater diversity of talents and perspectives made possible by the integration of more members of marginalized groups into high-level decision-making teams.[34]

NOTES

1. India also reserves seats for SC/ST members in central and state legislative assemblies, according to the SC/ST proportion of the population. Moreover, there are additional reserved seats in public higher educational institutions and public enterprises for "Other Backward Classes" (OBCs)—castes and communities low in the socio-economic hierarchy, but not formerly considered "untouchable" like Dalits. Such OBC reservations were introduced at the central level in 1991 and at the state level in various years after independence. We focus on SC/ST reservations in this paper because that policy has been stable over the time period of our study.

2. See, for instance, Guha (1990a, 1990b), and Shah (1991). Some critics have even suggested that the failure to allocate key jobs on a strictly meritocratic basis has resulted in serious harm as well as gross inefficiency. For example, in "Job Reservation in Railways and Accidents," *Indian Express*, September 19, 1990 (cited by Kumar, 1992: 301), it is charged that the frequency of Indian Railway accidents would likely increase because reservation policies result in a larger proportion of less competent railway officials and lower overall staff morale.

3. See Deshpande (2011) and Weisskopf (2004) for details of India's reservation policies, as well as discussion of the debates surrounding these policies.

4. The IR is one of the most important industries of any kind in India: it is the dominant industry providing essential freight and passenger transport services to Indians throughout the country, and its 1.4 million workers are far more numerous than in any other Indian public sector enterprise.

5. Some of the US studies have estimated industry-level production or cost functions, augmented by information on the extent and/or way in which labor inputs were affected by affirmative action. Other studies have analyzed company-level financial data to determine whether and how stock prices have been affected by evidence of affirmative action. Yet others have compared supervisor performance ratings of individual employees in establishments that do and do not practice affirmative action. The most comprehensive survey of such studies in the United States concludes that "There is virtually no evidence of significantly weaker qualifications or performance among white women in establishments that practice affirmative action..." and that "There is some evidence of lower qualifications for minorities hired under affirmative action programs..." but "Evidence of lower performance among these minorities appears much less consistently or convincingly..." (Holzer & Neumark, 2000).

6. No previous quantitative study of productivity in the IR, as far as we are aware, has been based on data disaggregated by zone.

7. For a thorough explication of the DEA approach, see Ray (2004). We relied on Coelli (1996) as a guide to our use of the technique.

8. The information contained in the first three paragraphs of this section is based on Government of India, Ministry of Railways, Directorate of Economics and Statistics (2008), *Key Statistics (1950–51 to 2006–07)*, and Government of India, Ministry of Railways (annual), *Appropriation Accounts*, Annexure G.

9. The sources of all the underlying data, as well as detailed explanations of how we constructed variables such as capital stock and "effective" labor, capital and fuel, are provided in a Statistical Appendix available on request from the authors.

10. Unfortunately, the IR does not compile and publish data on actual labor hours worked, so we had to work with employment data for each job category.

11. We owe this observation to Professor S. Thorat.

12. Given the stigma attached to being identified as an SC, those SC candidates who do not wish to avail of reservations typically do not declare themselves as SCs. However, cleaners are already stigmatized because of the "unclean" nature of their work, so many of them have no reluctance in identifying themselves as SC, irrespective of whether they occupy reserved seats. For example, an overwhelming percentage of the IR cleaning staff belongs to a formerly untouchable caste whose traditional occupation has been cleaning. The IR provide figures separately for cleaners (part of category D), and the percentage of those who declare themselves to be SC/ST is far higher than 22.5%.

13. The share of D-level employees in all-India IR employment fell gradually from 54% in 1980 to 40% in 2002.

14. This is because educational attainment among SC/ST members is much lower than for higher-caste individuals; see Deshpande (2011, chap. 3).

15. We also repeated the estimation with a 2% capital obsolescence fraction, and that did not change our basic results.

16. Pollution from coal, coke, and wood used in steam engines has adverse effects on railway employees and passengers as well as on the countryside.

17. All of the data used in our empirical analyses—at the all-India as well as the zonal level—are available from the authors on request.

18. We also estimated a translog production function, and we carried out a weighted-least-squares estimation to correct for possible heteroskedasticity. In the first case the results indicated a poor fit: most of the coefficients were not significant, the coefficients associated with the inputs did not satisfy monotonicity, and the estimated equation displayed high multicollinearity (We estimated the function using the actual values of the variables and then the normalized values (around the mean); both specifications displayed similar problems.) In the second case, the heteroskedasticity correction did not significantly alter the findings. All of these estimations are available from the authors on request.

19. Because the frontier is generated from the data, it is not based on stochastic processes and therefore does not produce any measures of the statistical significance of the results obtained.

20. Since DEA analysis permits us to use four separate labor input variables for the four different labor categories, there is no need to adjust for the category-composition of labor—as we had to do in our production-function analysis.

21. Thus in the second variant we also dropped from consideration the possible effect of the scale of production, which had proven insignificant in the results for the first variant.

22. We chose not to exclude the NFR zone from our production-function analyses, because then we were using a single aggregate capital stock variable for capital input—and it showed substantial and plausible growth over the period from 1980 to 2002.

23. The results for variant #1 are based on 2nd-stage regressions excluding the scale of production variable, whose estimated coefficient value proved to be quite insignificant under most specifications.

24. We also examined 2nd-stage correlations of SC/ST variables with estimates of ΔTFP values generated from a 1st-stage DEA run in which raw measures—rather than effective measures—of capital stock and fuel inputs were used; this resulted in correlation results very similar to those shown under Variant #2 in Table 2.

25. We are grateful to several readers of earlier versions of this paper for raising the concerns that we address in this section.

26. Our source for the IR hiring procedures described in this paragraph is <http://indianrailwayemployee.com/content/recruitment-and-selection>.

27. For instance, for category C and D employees who were recruited through the 19 Railway Recruitment Boards across the country in 2003, the minimum qualification scores were 40%, 30%, and 25% for unreserved, OBCs, and SC/STs, respectively.

28. We are indebted to William R. Johnson for drawing our attention to this possibility.

29. The last two possibilities were suggested to us by an anonymous reviewer.

30. The Indian Railways have only recently started to maintain figures on punctuality of long-distance trains. From January 2009 the Railway Board is analyzing punctuality performance by means of an "Integrated Coaching Management System"—a computer-based on-line system for accurate reporting and analysis of the voluminous data of long-distance train operations. (See http://www.indianrailways.gov.in/deptts/yearbook/ANNUAL_REPORT_08 09/Railway_Annual_Report_English_08-09.pdf.)

31. Source: GOI, Ministry of Railways, Annual Statistical Statements.

32. Unfortunately, the data available from the IR did not allow us to test for the importance of identity match between managers (at the A and B level) and workers (at the C and D level), as recommended by an anonymous reviewer.

33. This recalls the arguments in favor of AA in US educational institutions made to the Supreme Court by US military officers, who want to avoid having just white men in charge of troops that are disproportionately of color (See Weisskopf, 2004, preface).

34. Page (2007) shows convincingly how groups that display a wide range of perspectives outperform groups of like-minded experts.

REFERENCES

Barrington, L., & Troske, K. (2001). *Workforce diversity and productivity: An analysis of employer-employee matched data*. Economics Program Working Paper Series 01-02. The Conference Board, <http://gatton.uky.edu/faculty/Troske/working_pap/barrington_troske.pdf>.

Bertrand, M., & Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal: A field experiment on labor market discrimination. *American Economic Review, 94*, 991–1013.

Bertrand, M., Hanna, R., & Mullainathan, S. (2010). Affirmative action in education: Evidence from engineering college admissions in India. *Journal of Public Economics, 94*(1–2), 16–29.

Besley, T., Pande, R., Rahman, L., & Rao, V. (2004). The politics of public good provision: Evidence from Indian local governments. *Journal of the European Economic Association, 2*(2–3), 416–426.

Christensen, L. R., & Jorgenson, D. W. (1969). The measurement of US real capital input, 1929–1957. *Review of Income and Wealth, 15*(4), 293–320.

Coelli, T. (1996). *A guide to DEAP version 2.1: A data envelopment analysis (computer) program*. Working Paper No. 96/08 of the Centre for Efficiency and Productivity Analysis, Department of Econometrics, University of New England, Armidale, New South Wales, Australia.

Deshpande, A. (2011). *Grammar of caste: Economic discrimination in contemporary India*. New Delhi: Oxford University Press.

Deshpande, A., & Newman, K. (2007). Where the path leads: The role of caste in post university employment expectations. *Economic and Political Weekly, 42*, 4133–4140.

Guha, A. (1990a). The Mandal mythology. *Seminar, 375*, 51–53.

Guha, A. (1990b). Reservations in myth and reality. *Economic and Political Weekly, 25*, 2716–2718.

Holzer, H. J., & Neumark, D. (2000). Assessing affirmative action. *Journal of Economic Literature, 38*, 483–568.

Jodhka, S., & Newman, K. (2007). In the name of globalization: meritocracy, productivity, and the hidden language of caste. *Economic and Political Weekly, 42*, 4125–4132.

Kumar, D. (1992). The affirmative action debate in India. *Asian Survey, 32*, 290–302.

Levinsohn, J., & Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies, 70*, 317–341.

Munshi, K., & Rosenzweig, M. (2009). *The efficacy of parochial politics: Caste, commitments and competence in Indian local governments*. National Bureau of Economic Research Working Paper No. 14335.

Olley, S., & Pakes, A. (1996). The dynamics of productivity in the telecommunications industry. *Econometrica, 64*(6), 1263–1298.

Page, S. (2007). *The difference: How the power of diversity creates better groups, firms, schools, and societies*. Princeton: Princeton University Press.

Pager, D., & Western, B. (2005). *Race at work: Realities of race and criminal record in the NYC job market*. New York: Schomburg Center for Research in Black Culture.

Ray, S. C. (2004). *Data envelopment analysis: Theory and techniques for economics and operations research*. Cambridge, UK: Cambridge University Press.

Robles, V.C.F., & Krishna, K. (2012). *Affirmative action in higher education in India: Targeting, catch up, and mismatch*. Working Paper 17727. National Bureau of Economic Research, <http://www.nber.org/papers/w17727>.

Shah, A. M. (1991). Job reservations and efficiency. *Economic and Political Weekly, 26*, 1732–1734.

Siddique, Z. (2008). *Caste based discrimination: Evidence and policy*. IZA-Institute for the Study of Labor discussion paper 3737.

Thorat, S., & Attewell, P. (2007). The legacy of social exclusion: A correspondence study of job discrimination in India. *Economic and Political Weekly, 42*, 4141–4145.

Weisskopf, T. E. (2004). *Affirmative action in the United States and India: A comparative perspective*. London: Routledge.

APPENDIX A

See Tables 3 and 4.

Table 3. *Fixed effects estimation results*

Specification	1	2	3	4	5a	5ax	5b	5bx
Constant	−4.50	−4.46	−6.33	−2.03	−5.25	−4.55	−4.13	−2.63
	(10.3)	(9.90)	(14.0)	(8.19)	(9.71)	(9.07)	(10.7)	(10.4)
ln *ek* (ln *k*)	0.075	*0.062*	−0.145	*−0.012*	0.100	0.119	0.077	0.004
	(.50)	*(.51)*	(.38)	*(.30)*	(.48)	(.44)	(.50)	(.47)
ln *el* (ln *l*)	0.41	*0.46*	0.65	***1.15***	0.44	0.35	0.37	0.39
	(.31)	*(.37)*	(.61)	*(.40)*	(.31)	(.27)	(.31)	(.29)
ln *ef* (ln *f*)	0.022	*−0.013*	−0.019	*−0.03*	0.20	0.017	0.20	−0.006
	(.12)	*(.076)*	(.091)	*(.076)*	(.13)	(.13)	(.12)	(.11)
Time	**0.048**	**0.052**	**0.061**	**0.060**	**0.047**	**0.048**	**0.048**	**0.049**
	(.021)	(.018)	(.017)	(.011)	(.021)	(.019)	(.021)	(.020)
CORRELATIONS					COEFF'S			
%SC_ALL	−0.01	−0.01	0	0.04	−0.20			
					(.29)			
%SC_ALLx	−0.04	0.03	−0.03	0.02		−0.47		
						(.55)		
%SC_AB	−0.07	−0.06	−0.07	0.026			0.15	
							(.37)	
%SC_ABx	0	0	0	0.11				**0.87**
								(.33)
COEFFICIENTS								
%SC_ALL	−0.19	−0.14	−0.30	−0.25				
	(.31)	(.32)	(.41)	(.39)				
%SC_ALLx	−0.35	−0.28	−.57	−0.46				
	(.48)	(.50)	(.60)	(.59)				
%SC_AB	0.09	0.15	−0.39	−0.34				
	(.41)	(.39)	(.38)	(.36)				
%SC_ABx	0.46	0.52	−0.19	−0.15				
	(.50)	(.49)	(.48)	(.47)				

Note: For specifications 3 and 4 the dependent variable is $\ln r$; for all other specifications it is $\ln q$.
Standard errors are in parentheses. Figures in bold are significant at 5%.

Table 4. *Lev–Pet regressions*

Specification	1	2	3	4	5a	5ax	5b	5bx
ln *ek* (ln *k*)	**0.88**	**1**	**1**	**1**	**0.93**	**1**	0.79	0.41
	(.36)	(.39)	(.40)	(.34)	(.35)	(.42)	(.44)	(.38)
ln *el* (ln *l*)	0.184	0.011	0.060	−0.15	0.177	0.079	0.216	0.344
	(.37)	(.47)	(.49)	(0.49)	(.50)	(.49)	(.52)	(.49)
ln *ef* (ln *f*)	.000	.000	.000	.000	.000	.000	.000	**1**
	(.37)	(0.40)	(0.32)	(0.41)	(0.40)	(0.44)	(0.42)	(0.43)
CORRELATIONS					COEFF'S			
%SC_ALL	−0.02				−0.19			
					(1.03)			
%SC_ALLx	−0.05					−1.02		
						(1.34)		
%SC_AB	0.12						0.32	
							(.84)	
%SC_ABx	**0.19**							1.37
								(1.24)
COEFFICIENTS								
%SC_ALL	.000							
	(.00)							
%SC_ALLx	.000							
	(.00)							
%SC_AB	.000							
	(.00)							
%SC_ABx	.000							
	(.00)							

Note: For specifications 3 and 4 the dependent variable is ln r; for all other specifications it is ln q.
Standard errors are in parentheses. Figures in bold are significant at 5%.