Econometrica, Vol. 93, No. 4 (July, 2025), 1411–1448

CONTRACT LABOR AND ESTABLISHMENT GROWTH IN INDIA

MARIANNE BERTRAND
University of Chicago Booth School of Business, University of Chicago, NBER, CEPR, and IZA

CHANG-TAI HSIEH
University of Chicago Booth School of Business, University of Chicago, NBER, and CEPR

NICK TSIVANIDIS
Haas Business School, University of California, Berkeley

India's Industrial Disputes Act (IDA) requires large manufacturing plants to pay substantial costs if they wish to shrink their workforce. Since the early 2000s, these large plants have dramatically increased their use of contract workers who are not subject to these regulatory constraints. Between 2000 and 2015, the contract labor share in non-managerial employment nearly doubled at establishments with more than 100 workers (from 21 to 40 percentage points), while it only increased from 14 to 17 percentage points at establishments with less than 50 workers. Over the same period, the thickness of the right tail of the establishment size distribution in formal Indian manufacturing plants increased, the average product of labor at large plants declined, the job creation rate for large plants increased, and the probability that large plants introduced new products rose. We argue that these changes were caused by the increased adoption of contract labor. In a model of establishment growth subject to firing costs, we show that easing access to contract labor increased TFP in Indian manufacturing by 7.3% since the early 2000s, occurring all through a one-time reduction in misallocation between large and small plants with negligible change in the long-run growth rate.

KEYWORDS: Contract labor, growth, India.

1. INTRODUCTION

MANY OBSERVERS have pointed to the Industrial Disputes Act (IDA) of 1947 as an important constraint on growth in India. In particular, Chapter VB of the IDA requires manufacturing plants with more than 100 non-managerial workers that wish to shrink their employment to provide severance pay, mandatory notice, and obtain governmental retrenchment authorization.[1] The IDA thus potentially constrains growth in two ways. First, the most productive Indian plants are likely to be sub-optimally small. Consistent with this, the Indian manufacturing sector is characterized by a large number of informal plants, a small number of large plants, and a high marginal product of labor at large plants. Second, the higher costs faced by large plants in retrenching workers may dissuade

Marianne Bertrand: marianne.bertrand@chicagobooth.edu
Chang-Tai Hsieh: chang-tai.hsieh@chicagobooth.edu
Nick Tsivanidis: ntsivanidis@berkeley.edu

We thank the University of Chicago's India Trust for financial support and Manish Sabharwal for helping us understand the labor contracting landscape in India. Amedeus Dsouza, Adarsh Kumar, Kazuatsu Shimizu, and William Wang provided superb research assistance.

[1] A 1976 amendment to the IDA made layoff, retrenchment, and closure illegal for all plants with more than 300 non-managerial workers. The threshold was lowered to 100 in 1982, with some states further lowering it to 50. From here on, we use the terms workers and employment interchangeably to refer to non-managerial workers.

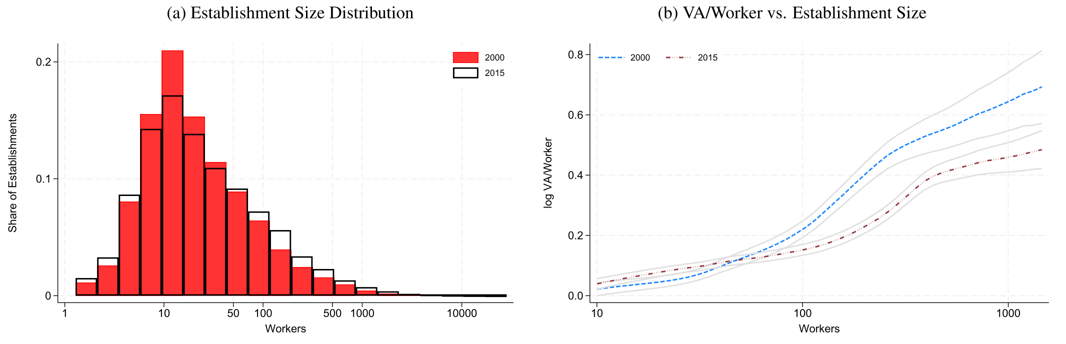

FIGURE 1.—Establishment size distribution and VA/Worker by establishment size, 2000 versus 2015. *Note*: Left panel shows distribution of plants by employment. Right panel shows coefficients and 95% confidence intervals from non-parametric regressions of log VA/Worker on log employment using Epanechnikov kernel with a bandwidth of 0.6. Log VA/Worker is residualized by industry-year fixed effects.

them from undertaking risky investments, which may be one of the forces behind the low life-cycle growth of Indian manufacturing plants.[2]

This paper argues that the constraints on large plants have diminished since the early 2000s, despite no change to the IDA.[3] Consider the evidence in Figure 1. The left panel shows that the thickness of the right tail of formal Indian manufacturing increased between 2000 and 2015. The right panel shows that average value-added/worker is increasing in establishment employment in 2000 and 2015, but this relationship is more attenuated in 2015 compared to 2000, particularly for larger plants. If the marginal product of labor is proportional to its average product, and employers equate the marginal product of labor to the cost, the effective cost of labor has diminished for large Indian plants compared to smaller ones since the early 2000s.[4]

We further argue that the main force behind the decline in labor constraints faced by large Indian plants since the early 2000s is the increased reliance of these plants on contract workers hired via staffing companies. The IDA rules regarding severance pay, mandatory notice, and governmental authorization for retrenchment only apply to an establishment's permanent workforce. Hence, plants have the flexibility to return contract workers to staffing companies without being in violation of the IDA.

While a legal framework for the deployment of contract labor has been in place in India since the early 1970s to limit its use to tasks that are non-perennial and not regularly performed by permanent workers, the staffing model started booming in the early 2000s. Figure 2 shows the fraction of establishment workers hired through contractors as a function of total establishment non-managerial employment. While there has been no sizeable change in the share of contract workers at smaller plants, there has been a dramatic increase among larger plants, especially those with more than 100 workers. By 2015, contract workers account for 40% of non-managerial employment at plants with more than 100 workers, compared to 21% in 2000. In contrast, the equivalent number at

[2]See Hsieh and Olken (2014) on plant-size distribution in India and Hsieh and Klenow (2014) on low growth over the life-cycle in Indian manufacturing.

[3]The Industrial Relations Code of 2020 consolidates and updates the Industrial Disputes Act of 1947, the Trade Unions Act of 1926, and the Industrial Employment (Standing Orders) Act of 1946. It has yet to come into force.

[4]Figures 1 and 2 are from India's Annual Survey of Industries described in Section 3.

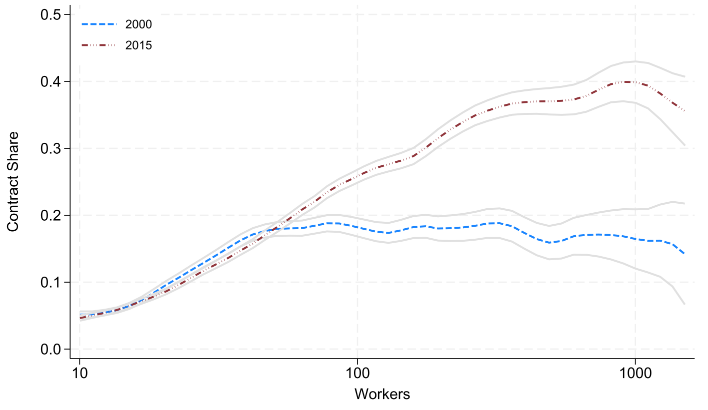

FIGURE 2.—Contract labor use and establishment size: 2000 versus 2015. *Note*: Plot shows point estimates and 95% confidence intervals from non-parametric regression of the share of workers hired through contractors on (log) employment.

plants with less than 50 workers is only 17% in 2015, which is only a 3 percentage points increase relative to 2000.

We argue that a decision by the Indian Supreme Court in 2001 played an important role in the increased adoption of contract labor in India, particularly at larger plants.[5] Prior to this decision, it was unclear whether plants that were "caught" improperly using contract workers in "core" activities would have to absorb these workers into their permanent workforce. This plausibly made large plants reticent to rely too much on such outsourcing. The 2001 Supreme Court decision clarified that such absorption was not required, therefore decreasing the cost of tapping into this alternative source of labor. We show that there was a marked change in the use of contract workers by large plants, in the employment share of large plants, and in the gap in labor productivity between large and small plants after 2001. In addition, these changes were more pronounced for plants with closer access to staffing centers prior to the Supreme Court decision.

There are two main channels via which a greater reliance on contract workers may have led to the expansion of large plants and the decline in the value-added per worker at these plants. First, the IDA places size-dependent restrictions on the ease of firing workers. Because these firing costs are lower for contract workers, employment at large plants that rely more on contract workers should be more responsive to productivity shocks. Consistent with this channel, the time series evidence reveals both an increased likelihood of large (more than 10%) employment change, as well as an increase in the standard deviation of employment growth at large plants starting in the early 2000s. Also, using Bartik-style labor demand shocks as well as rainfall shocks at the district level, we show that plants in districts where contract workers are more readily available are more responsive to such demand shocks.

Second, the availability of contract labor may have reduced the extent to which large plants face a higher marginal cost of labor because of greater unionization and other labor cost pressures disproportionately imposed on these plants by the regulatory environment. Consistent with this channel, we find that while there is a positive and quite stable elasticity of the average cost of labor to establishment size prior to 2000 of about 0.14, this

[5]The Supreme Court Case is "Steel Authority of India Ltd. vs. National Union Water Front Workers.".

elasticity starts declining in the early 2000s, dropping to 0.08 by 2015. This decline comes from two forces. First, the relative cost of contract labor compared to permanent labor is lower at larger plants, and hence the average cost of labor goes down for larger plants as they tap more into the contract labor pool. Second, the rise in contract labor exerts downward pressures on the wages of permanent workers at larger plants: the elasticity of permanent labor cost to establishment size is positive but started trending down in the early 2000s, especially in districts closer to staffing centers.

We corroborate all of these findings in an establishment-year panel that controls for establishment fixed effects as well as industry-year specific shocks. We show that an establishment's increased reliance on contract labor is associated with an increase in its size, a decrease in the average product of labor, an increase in employment variability, and a decrease in the average cost of labor. We also report evidence suggesting that reliance on contract labor makes plants more dynamic and more likely to change their product mix.

We use a model of creative destruction with heterogeneous establishments to quantify the effect of contract labor in the presence of the IDA. The model features two types of establishments. High-type establishments innovate more and have more products, and low-type establishments innovate less and have less products in steady state. We model the IDA as an adjustment cost faced by high-type plants whenever they fire workers. The anticipation of future retrenchment causes high-type plants to hire a sub-optimally small number of workers (increasing the average product of labor) and to invest less in innovation (reducing the likelihood they grow by adding new products). We model the use of contract workers by large plants by assuming that plants subject to the IDA can circumvent firing costs by hiring contract workers after paying a fixed cost in order to tap into this alternative employment pool.

We estimate the reduction in this fixed cost that matches the rise in the share of contract labor within large plants over the period under study. By simulating a counterfactual in which only this fixed cost changes while holding all other aspects of the model constant (including the retrenchment costs due to the IDA), we estimate the effect of the growing use of contract workers by large plants on the gap in value-added per worker between large and small plants, aggregate TFP, and the innovation rate.

We find that the increased adoption of contract labor explains all of the decline in the gap in value-added per worker between large and small plants seen in Figure 1. This decline represents less static misallocation of labor and accounts for 7.3% of the overall increase in manufacturing TFP over this period. However, our results suggest the aggregate growth rate did not change due to the proliferation of contract labor. On the one hand, the plants subject to the IDA innovate more as its bite is reduced when they use more contract labor. On the other hand, entrants (who are more likely to be of the low type) respond to this increased competition by innovating even less. Consistent with these predictions of the structural model, we find reduced form evidence of an increased innovation gap between large and small plants and a reduced employment share of entrants—both borne out in the data despite not being targeted by the model. On net, these two effects cancel each other out with no net impact on aggregate innovation (and hence aggregate growth).

This paper makes four main contributions. First, it contributes to the literature on the rise of non-standard work arrangements.[6] In the case of the rise of contract work in India, most relevant to our work is Chaurey (2015) who shows that establishments located

[6]Katz and Krueger (2019), Goldschmidt and Schmieder (2017), Drenik, Jager, Plotkin, and Schoefer (2020), and Felix and Wong (2021) document the rise of temporary work arrangements in the United States, Germany, Argentina, and Brazil, respectively.

in states with more stringent labor regulations hire relatively more contract workers in response to weather-induced transitory local demand shocks. The evidence in Chaurey (2015) is consistent with the hypothesis that establishments facing the most stringent labor regulation might be hiring contract workers to get around the strict labor laws. Saha, Sen, and Maiti (2013) find a positive relationship between import penetration and the share of contract workers; they also find that pro-worker legislation and greater bargaining power of permanent workers increase the share of contract workers. Kapoor and Krishnapriya (2017) study the rise in the use of contract labor in the Indian manufacturing sector between 2000 and 2012 across states and industries. They document greater use of contract workers in establishments where the gap in average wage between permanent and contract workers is smaller, which suggests that contract workers are used to suppress the bargaining power of regular workers. We are not aware of any research directly using the SAIL judgment as a source of temporal variation in the size of the contract workforce to explore its implications for misallocation and productivity growth in the formal Indian manufacturing sector.

Second, we develop a model of creative destruction with heterogeneous plants and size-dependent firing costs to quantify their effect on aggregate TFP. A strand of existing theoretical work seeks to quantify the effects of such policies.[7] However, in these models, establishment productivity is exogenous and thus cannot account for how such policies might affect productivity growth. In this paper, we extend Klette and Kortum's (2004) model of endogenous innovation to incorporate size-dependent firing costs.[8] This allows us to quantify the effect of these frictions on TFP, through impacts on both static misallocation and long-run productivity growth.

Third, we provide new evidence on the source of misallocation in developing countries. A large literature has documented misallocation as a potential source of low TFP in developing countries, and much work since has sought to isolate particular causes of misallocation. This paper provides evidence for Indian manufacturing that it is large rather than small plants that face higher frictions in the labor markets, and shows that this gap is driven in part by size-dependent firing costs using a quasi-experimental policy change that reduced these frictions.[9]

Lastly, our paper contributes to the literature on the economic impact of labor regulation in India.[10] We provide new evidence on how large plants subject to the IDA were able to circumvent this law by employing contract workers and trace the impacts of this de facto reduction in labor regulation on the Indian formal manufacturing sector.

[7]See, for example, Hopenhayn and Rogerson (1993) and Guner, Ventura, and Xu (2008).

[8]Other papers that extend Klette and Kortum (2004) include Acemoglu, Akcigit, Alp, Bloom, and Kerr (2018), Akcigit and Kerr (2018) and Aghion, Bergeaud, and Van Reenen (2021). Our model is closest to that in this last paper, which analyzes the effect of size-dependent restrictions in France.

[9]Gourio and Roys (2014) and Garicano, Lelarge, and Van Reenen (2016) study the effects of size-dependent labor regulations in France.

[10]Fallon and Lucas (1993) show that the 1976 amendment of the IDA, which mandated plants employing 300 or more workers to request permission from the government prior to retrenchment, lowered formal employment by 17.5%. Dutta Roy (2004) also finds that plants subject to the IDA face substantial adjustment costs, but that the 1982 amendment to the IDA, which extended the prohibition to retrench workers without government authorization to plants that employed 100 or more workers, did not change these costs. Besley and Burgess (2004) exploit the state-level variation induced by the state-level amendments to the IDA and find that states which amended the IDA in a pro-worker direction experienced lowered output, employment, investment, and productivity in formal manufacturing. Hasan, Mitra, and Ramaswamy (2007) and Aghion, Burgess, Redding, and Zilibotti (2008) show that pro-worker states are less responsive to trade reform and industrial licensing reform, respectively.

The paper is organized as follows. Section 2 lays out the institutional background. Section 3 describes data sources. Section 4 provides a simple theory to illustrate the possible forces behind the observed increased use of contract labor, highlighting differences between markers in the data of supply versus demand shocks. Section 5 provides empirical evidence suggestive of a causal relationship between the rise in the supply of contract labor and the increase in establishment size and decline in average product labor at larger plants. In Section 6, we investigate two mechanisms via which contract labor may have freed up establishment growth: reduction in labor adjustment costs and reduction in the cost of labor. Section 7 develops and estimates our structural model. Section 8 discusses the implications of the rise of the staffing model for workers and sketches a model to quantify its distributional effect across groups of more and less educated workers. We conclude in Section 9.

2. CONTRACT LABOR IN INDIA

A central piece of labor legislation in India is the Industrial Disputes Act (IDA, 1947), which lays out the conditions for hiring and retrenching workers, as well as for the closure of plants. In particular, a 1976 amendment to the IDA (Chapter VB) stipulates that all plants with more than 300 workers need to get government authorization for any layoff, retrenchment, or closure. This coverage was extended in 1982 to all plants with more than 100 workers, with some states further reducing this threshold to 50 workers.[11] Unapproved separations carry a potential punishment of both a substantial fine and a prison sentence for the employer.[12, 13]

However, a loophole exists in the IDA that can theoretically enable large plants to skirt some of its requirements. The application of severance pay, mandatory notice, or governmental retrenchment authorization only applies to permanent workers. Hence, by resorting to contract labor, a large employer could theoretically bypass some of the most restrictive regulations of the IDA.

Indian legislators began to address this loophole in 1970 when they passed the Contract Labour (Regulation and Abolition) Act (CLA, 1970). This Act was enacted "to regulate the employment of contract labour in certain plants and to provide for its abolition in certain circumstances."[14] The CLA requires that all plants with 20 or more contract workers obtain a registration for employing such labor and that all staffing agencies with

[11] The IDA defines a worker, which it refers to as a "workman" as: "any person (including an apprentice) employed in any industry to do any manual, unskilled, skilled, technical, operational, clerical or supervisory work for hire, whether the terms of employment be expressed or implied..." The definition further explicitly excludes individuals working in managerial and administrative tasks. It does not differentiate between part-time and permanent work. Establishment size for purpose of IDA coverage is based on the number of "workmen... employed on an average per working day for the preceding twelve months.".

[12] Actual compensation for retrenchment if granted is quite low by international standards: any worker (as defined by the IDA) with more than 240 days of service is entitled to one month's notice and 15 days of compensation for every year of service at 50 percent of basic wages plus dearness allowance.

[13] Other aspects of the IDA and related laws impose additional costs on plants with a large number of workers. For example, the Industrial Employment (Standing Orders) Act requires establishments of more than 100 employees (and in some states 50) to specify to workers the terms and conditions of their employment, while the IDA requires employers to provide Notice of Change (Section 9-A), meaning that no employer can effectuate any change in the conditions of service of any worker without giving 21 days of notice. The IDA also sets conciliation, arbitration and adjudication procedures to be followed in the case of an industrial dispute and empowers national or state governments to constitute Labour Courts, Tribunals, National Tribunals, Courts of Inquiry, and Boards of Conciliation. See Ahsan and Pages (2007).

[14] https://clc.gov.in/clc/acts-rules/contract-labour-regulation-abolition-act-1970.

20 or more employees be government-licensed.[15] Contract workers covered under the CLA have rights related to working hours, safety and health, social security (under the Employees' State Insurance Act of 1948), and retirement benefits (under the Employees' Provident Funds and Miscellaneous Provisions Act of 1952).

Most importantly, Section 10 of the CLA limits employers' ability to deploy contract workers as a way to get around the IDA's requirements. Under Section 10, contract workers are de jure not supposed to be in charge of tasks within an establishment that are perennial in nature and typically completed by permanent workers in that industry. The Act gives government the authority to prohibit or "abolish" contract labor at any establishment that uses this labor for its "core" operations.

The CLA, however, left vague what would happen to the contract workers at an establishment subsequent to the government issuing a notification under Section 10 banning the establishment from using this labor. In particular, there was uncertainty as to whether, subsequent to an abolition notification, the employer would be required to automatically absorb the contract workers into its permanent workforce. While such absorption would seem to be in the spirit of Section 10 (e.g., not using contract labor as a loophole around the IDA) and might have been implicitly assumed, the Act was not explicit.

The liberalization of the Indian economy in 1991 gave Indian employers a stronger impetus to get around the IDA and find ways to bring in more contract workers in their workforce (Gopalakrishnan and Mirer, 2013). In response to industry pressures, some state governments eased up the licensing procedures for labor contractors and started making amendments to the legislation on contract labor (Saha and Sen, 2014). Employers also started lobbying for a reform of Indian labor laws, including Chapter VB of the IDA, as well as for the scrapping of Section 10 of the CLA (Gopalakrishnan and Mirer, 2013). These legislative lobbying efforts went nowhere, likely because of strong opposition from the trade unions.

At the same time, employers also started arguing before the Courts that the CLA did not require the absorption of contract workers into the permanent workforce. Following a series of earlier judicial decisions (some pro-employers; some pro-workers), a 2001 ruling by the Supreme Court of India, which overturned a prior 1997 ruling, lifted the uncertainty about absorption requirements in employers' favor. In its *Steel Authority of India Limited v. National Union Water Front Workers* judgment (the "SAIL" judgment), the Supreme Court ruled that there is no requirement for automatic absorption of contract workers in the permanent workforce subsequent to an abolition notification.

The SAIL judgment has been deemed by various observers as critical in the rise of contract labor in India. Gonsalves (2011) writes: "A legal right, to permanent employment of the contract workers where the contract labour system has been abolished, goes a long way to reducing the prevalence of the contract labour system throughout India. The stand that no such right exists on abolition will achieve quite the opposite." Similarly, Landau, Mahy, and Mitchell (2015) note: "The implications of this shift have proven significant and contentious, with unions abandoning their strategic use of s. 10(1) of the Act as a means of securing permanency for contract workers…". Several authors (e.g., Sankaran, 2012, Cox, 2012; Sundar, 2012) describe the SAIL decision as "de facto" deregulation without any changes to the labor laws.

[15]According to a report by *Staffing Industry Analysts*, the three largest staffing companies in India by 2012 were Adecco, Teamlease, and Randstad and these three companies accounted for about 15 percent of the total market; the market share of the top ten staffing companies was about 26 percent. See https://www2.staffingindustry.com/row/Editorial/Daily-News/India-Adecco-is-largest-staffing-firm-in-a-USD-5-Billion-market-28900.

TABLE I

TOP 10 OCCUPATIONS FOR CONTRACT AND PERMANENT WORKERS IN FORMAL MANUFACTURING.

	Contract Workers		Permanent Workers	
Rank	Occupation	Share	Occupation	Share
1	Industrial and machine workers	51.45%	Industrial and machine workers	42.16%
2	Plant and machine, industrial machine operators	8.21%	Supervisors, Shift-in-charge, Workshop Managers	9.04%
3	Supervisors, Shift-in-charge, Workshop Managers	7.78%	Plant and machine, industrial machine operators	7.30%
4	Tailors, Dressmakers, Dress designers	3.56%	Metal Moulders, Welders	6.24%
5	Metal Moulders, Welders	2.86%	Office, bank clerks, court clerks, office assistants	4.63%
6	Peons, cleaners and helpers	2.25%	Tailors, Dressmakers, Dress designers	3.74%
7	Liftmen, watchmen, security guards	1.96%	Plant and Machinery Mechanics and Repairers	2.49%
8	Plant and Machinery Mechanics and Repairers	1.62%	Engineers	2.30%
9	Engineering and Industrial Designers	1.52%	Traditional hand embroiders, cloth block printers	1.97%
10	Engineers	1.52%	Machine technicians, Mechanical engineering technicians	1.93%

Note: Sample includes temporary and permanent workers employed in non-managerial occupations in the formal manufacturing sector. Table reports the mean across three waves of the CPHS between Jan 2018 and Dec 2018, where each observation within a wave is weighted with the CPHS sampling weight.

With the absorption requirement gone, employers may have become more willing to operate in a legal "gray zone" and rely on contract labor for core operations. In a survey of about 100 Haryana-based manufacturing plants conducted in 2015, Singh, Kusum Das, Choudhury, and Kukreja (2016) found that the large majority of surveyed plants that use contract workers report having contract and permanent workers work side by side. Singh et al. (2016) write: "We can thus broadly make the inference that the survey supports the hypothesis that contract workers are not confined to peripheral activities but rather substitute for regular workers in the core tasks of plants."

In support of this claim, Table I shows the top 10 occupations of permanent and contract workers in formal manufacturing in 2018 based on worker-level data from India's Consumer Pyramids Household Survey (CPHS).[16] In particular, both groups of workers share the same three most common occupational categories: industrial and machine workers; plant and machine, industrial machine operators; supervisors, shift-in-charge, workshop managers. These three occupations account for nearly 70 percent of contract workers' employment and nearly 60 percent of permanent workers' employment. Seven out of the 10 most common occupations for permanent workers are also represented in the top 10 for contract workers. This is in contrast to developed country settings where contract workers tend to work in occupations peripheral to the plant, such as security, cleaning, logistics, and catering.[17]

Table II further shows that there are no marked differences in educational attainment between both types of workers. On the other hand, permanent workers are substantially

[16]We defer to the next section (and Supplemental Appendix Section C in Bertrand et al. (2025)) for a fuller description of the CPHS data. 2018 is the earliest year for which we can perform this tabulation.

[17]See Goldschmidt and Schmieder (2017).

TABLE II

CHARACTERISTICS OF FORMAL MANUFACTURING WORKERS: PERMANENT VERSUS CONTRACT.

	Age	> 10 Years	> 12 Years	Female	Upper Caste	Scheduled Caste	> 5 Years
Permanent	3.770	0.004	−0.004	−0.034	0.024	−0.082	−0.023
Worker	(0.450)	(0.018)	(0.022)	(0.010)	(0.016)	(0.019)	(0.048)
Constant	34.989	0.805	0.541	0.073	0.180	0.261	0.631
	(0.320)	(0.013)	(0.016)	(0.007)	(0.012)	(0.015)	(0.033)
R^2	0.13	0.07	0.11	0.11	0.15	0.06	0.13

Note: Table shows coefficients from a regression of the worker's characteristic (shown in each column) on a permanent status dummy variable with industry and state fixed effects. Sample in columns 1–6 includes temporary and permanent workers employed in non-managerial occupations in formal manufacturing (Nobs=3860). Column 7 only includes workers with < 10 years of schooling (Nobs=727). Analysis is based on CPHS May–Aug 2017 wave using weights for population aged 15 or higher. Robust standard errors reported.

older (by nearly 4 years) than contract workers. There is also some evidence that permanent workers belong to more advantaged groups in society (more likely to be males, less likely to belong to the schedule castes), which may reflect discriminatory barriers in accessing the rare permanent positions in the Indian formal economy. The last column restricts the sample to workers with less than 10 years of schooling, and shows that, among this less educated group, there are no differences in the fraction with more than 5 years of education between permanent and contract workers.

3. DATA

Our primary source of data is the Annual Survey of Industries (ASI) conducted by India's Statistical Office (NSSO). The ASI collects data between April of a given year until the end of March the following year. When we refer to the year of the ASI, we refer to a survey that began in April of that year. The ASI is a census of "large" formal Indian manufacturing plants and a random sample of "smaller" formal plants.[18] For most years, plants with more than 100 workers are in the census sector, although the size threshold for inclusion in the census sector changes over time.[19] Our main analysis is based on the ASI from 1980 to 2015.

The key variables from the ASI are establishment ID (available between 1993 and 2015), district identifiers (available until 2009), value-added, employment, labor compensation, electricity usage, book value of capital, and main industry of the establishment at either the 4- or 5-digit level. The ASI provides information on the number of workers directly employed by the establishment and workers hired through contractors (here-

[18]The 1948 Factories Act requires that plants with more than 20 workers be formally registered (the threshold is 10 workers if the plant uses electricity).

[19]Up until 1996, the census sector consisted of plants with more than 100 workers, and plants not in the census sector were sampled by state and 3-digit sector, with roughly one-third probability. Between 1997 and 2003, only plants with 200 or more workers were included in the census sector, and smaller plants were sampled by state and 3-digit sector roughly with one-seventh probability. The census sector reverted to all plants over 100 workers between 2004 and 2014, and plants outside the census sector were sampled by state and 4-digit industry with roughly one-fifth probability. Starting in 2015, the size threshold for inclusion in the census sector varied by state.

after referred to as "permanent" and "contract" workers).[20] Wages, bonuses, and benefits for all workers are reported in all years, and a breakdown between permanent and contract workers is provided in a subset of years.[21] The ASI also provides information on the number of "managerial" and "non-managerial" workers, as well as wages, bonuses, and benefits for these two types of workers.

Our secondary data set is the Center of Monitoring of the Indian Economy (CMIE)'s India-wide representative Consumer Pyramids Household Survey (CPHS). It is a panel survey of nearly 160,000 households across India. CPHS surveys are carried out in a "wave" of 4 months, where each household (and its members) are surveyed three times a year.

The information we use from the CPHS are income, sex, occupation, educational attainment, age, industry (at the 2-digit level), labor market arrangement (self-employed, permanent, contract, daily-wage, not-employed), caste, access to a provident fund, and the sampling weight that makes each wave nationally representative. We define a worker as "formal" if they report having access to a provident fund, and informal if they do not. Supplemental Appendix Section C (Bertrand, Hsieh, and Tsivanidis (2025)) shows that the total number of workers and the share of contract workers in the CPHS sample of formal workers in manufacturing are comparable to those in the ASI (which only surveys formal manufacturing plants). Therefore, unless otherwise indicated, we restrict the CPHS to workers in the formal manufacturing sector to make the CPHS sample comparable to the ASI.[22]

4. CONTRACT LABOR: SUPPLY VERSUS DEMAND

We sketch a model of supply and demand for contract workers to illustrate the forces behind the increased use of contract labor observed in the data. The goal is to show that the increased use of contract labor by large plants can be driven either by an increase in the supply of contract labor or by an increase in the demand for contract labor, but that the declining gap in the average product of labor between large and small plants can only be due to higher supply of contract labor. In Section 7, we use this model, after endogenizing the innovation rates, to estimate the effect of an increase in the supply of contract labor on aggregate TFP.

Aggregate output is $Y = (\int_0^1 (q_j y_j)^{\frac{\sigma-1}{\sigma}} dj)^{\frac{\sigma}{\sigma-1}}$, where y_j denotes quantity and q_j quality of variety j. Output of a variety is given by $y_j = \ell_j$, where ℓ_j denotes the number of workers used to produce variety j. A worker can be permanent (employed directly by the establishment) or contract (employed via a staffing company). The two types of workers are perfect substitutes in production and are paid the same wage w.[23]

[20]Workers in the ASI "include all persons employed directly or through any agency whether for wages or not, and engaged in any manufacturing process,…, the repair and maintenance or production of fixed assets or for generating electricity or producing coal, gas etc.".

[21]Wages for permanent and contract workers are separately provided between 1998 and 2015; the same is true for bonuses and benefits between 1998 and 2007.

[22]See Supplemental Appendix Section C (Bertrand, Hsieh, and Tsivanidis (2025)) for more details on how we use the CPHS to construct the relevant samples of workers. For some of the results, we also use the Economic Censuses, rainfall data from Matsuura and Willmott (2012), and measures of industry-level reforms occurring between 1985 and 1997 from Aghion et al. (2008).

[23]This formulation is isomorphic to one where the two types of workers differ in quality but are perfect substitutes when adjusted for quality. In this case, l_j is the number of workers in quality-adjusted units, and the observed wage gap reflects the quality gap between the two types of workers. Supplemental Appendix

An establishment is a collection of varieties so differences in establishment size reflect differences in the number of products they own and the average quality of these products. There are two types of plants, a "high" type and a "low" type, that differ in two ways. First, high-type plants, on average, own a larger number of products compared to low-type plants.[24] Second, to capture the effect of the IDA, we assume that high-type establishments face firing costs for their permanent workers while low-type establishments do not.

We assume contract workers can be fired at zero cost (by all plants) but the employer needs to pay a fixed cost F for each product line they want to staff with such workers. Given this cost and the assumption that permanent and contract workers are perfect substitutes, a low-type establishment will always employ permanent workers. Revenue per worker of such plants is given by $(\frac{\sigma}{\sigma-1})w$ and the same for all varieties owned by low-type plants.

A high-type establishment may choose to employ permanent workers on some product lines and contract workers on other product lines. The critical variable is the probability that the establishment will be forced to retrench when another establishment innovates on its products, which occurs with probability x. If the high-type establishment chooses to employ permanent workers on a product line, it faces an additional labor cost $x\kappa w\ell$. Conditional on employing permanent workers, profit-maximizing labor productivity is given by $(\frac{\sigma}{\sigma-1})w(1+x\kappa)$, which is higher than the labor productivity of low-type plants due to the firing cost.

A high-type establishment can avoid the retrenchment cost by paying a fixed cost F to contract with a staffing company. In this case, profit-maximizing labor productivity is given by $(\frac{\sigma}{\sigma-1})w$. It will choose to do this when the flexibility gains from employing contract relative to permanent workers exceeds the fixed cost F.[25] Average labor productivity of high-type plants is thus a weighted average of labor productivity of products that use permanent workers and products that use contract workers. The gap in average labor productivity between high- and low-type plants thus depends on the share of products for which the high-type establishments employ permanent workers, where this share depends on the fixed cost F of hiring contract workers.

This model captures two key facts about contract labor in India. First, larger plants are more likely to be high-type plants because such plants have a larger number of products, and thus are more likely to employ contract labor. This is consistent with the evidence in Figure 2 that larger plants are more likely to hire contract labor. Second, larger plants pay on average higher labor costs because they are more likely to be high-type plants that face a higher cost for the product lines on which they choose to only employ permanent workers. This captures the fact in Figure 1 that the average product of labor is higher in larger plants compared to smaller plants.

We interpret a reduction in the fixed cost F as a metaphor for forces such as an expansion in the number of staffing companies that make it easier for companies subject to

Section F.5 (Bertrand, Hsieh, and Tsivanidis (2025)) considers a model where full-time and contract workers are imperfect substitutes even after adjusting for quality. Bertrand, Hsieh, and Tsivanidis (2025) also presents a model where a plant adds management layers as a function of the demand for its products.

[24] In Section 7, we follow Klette and Kortum (2004) and endogenize the distribution of products across plants as the result of an innovation process where high-type plants innovate more frequently compared to low-type plants. Therefore, high-type plants have on average more products compared to low-type plants in a steady state.

[25] A high-type establishment will employ contract labor on product lines where the quality exceeds the threshold quality $q^* \equiv \frac{\sigma}{\sigma-1}w[1-(1+x\kappa)^{1-\sigma}]^{\frac{1}{1-\sigma}}(\sigma F)^{\frac{1}{\sigma-1}}$.

the IDA to hire contract workers. Specifically, a reduction in the fixed cost of contracting with staffing companies makes high-type firms choose to employ contract labor for more of their products. This increases the share of contract labor in the employment of high-type firms. It also decreases the average product of labor for such firms because they now employ costly permanent workers on a smaller number of their products.

Now suppose instead that there is no change in the fixed cost of using contract workers but instead there is an increase in the share of establishments that are high-type. Think of this as an increase in the demand for contract workers, which also increases the share of firms that use contract workers. However, more demand for contract labor does not lower the average cost of labor faced by high-type plants, and thus cannot explain a decrease in the average product of labor of large establishments relative to that of small establishments seen in the data.

5. DID THE RISE IN CONTRACT LABOR FREE UP ESTABLISHMENT GROWTH?

In this section, we examine the effect of specific supply shifters on contract labor use and the average product of labor at plants that increased their use of contract workers. First, we conduct an event study analysis around the SAIL judgment in 2001, which we argued plausibly lifted the constraints on the use of contract labor by large Indian plants, that is, decreasing the cost of tapping into this alternative pool of workers. Second, we examine the heterogeneous effect of the SAIL judgment in districts with greater versus lesser proximity to staffing centers prior to SAIL. Third, we conduct within-establishment analysis to study changes in establishment outcomes associated with the hiring of contract workers.

5.1. *Effect of SAIL Event in the Time Series*

The SAIL judgment in 2001, by freeing up the use of contract workers, may have weakened the additional constraints large plants faced compared to smaller plants because of the IDA. Under this hypothesis, we expect the year 2001 to mark a break in trend for the motivating patterns documented in the Introduction.

Figure 3 shows the use of contract workers across plants in different size categories over time. Panel (a) regresses the share of contract labor in non-managerial employment on year dummies interacted with establishment size indicators: a dummy for whether the plant has 20–49, 50–99, 100–499, and more than 500 workers (relative to the omitted category of plants with less than 20 workers).[26] We then plot the estimated establishment size coefficients for each year, as well as the 95% confidence intervals.

The figure reveals some divergence starting around the SAIL decision between larger and smaller plants. In particular, the relative representation of contract workers at plants with less than 20 workers and plants with between 20 and 49 workers has remained roughly stable throughout the time period under study. In contrast, while there is substantial overlap in the contract labor share between plants with between 20 and 49 workers and those with more than 50 workers over the 1990s, a statistically and economically significant gap emerges post-SAIL. In particular, plants with 100 workers or more, but

[26]Also included in the regression are industry-year fixed effects, and standard errors are clustered by industry. All figures include establishments with more than 10 workers (the cutoff for inclusion in the ASI given that all plants in the ASI use electricity) and less than the 99th percentile of workers. We also winsorize the 1% tails of continuous, unbounded variables.

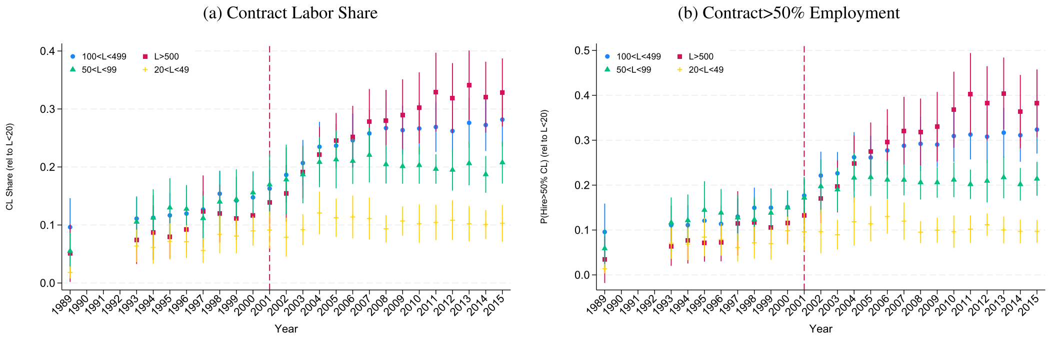

FIGURE 3.—SAIL and contract labor use by establishment size. *Note*: Plot shows coefficients and 95% confidence intervals on regression of outcome on year-industry dummies and year dummies interacted with each size category for employment (with less than 20 workers the omitted category). Standard errors are clustered at industry level.

especially those with 500 workers or more, experience a continuous relative increase in their contract labor share until the early 2010s.

The figure also indicates some rise in the contract labor share at establishments above the 50 workers threshold in the 1990s, which is likely a reflection of some of the easing on contract labor usage in some Indian states following economic liberalization measures in 1991 (see Section 2). The figure also makes it clear that a break in that pre-trend emerges around the SAIL decision, with a remarkable acceleration in the use of contract labor at larger establishments.[27]

Panel (b) estimates the same regression as Panel (a) but instead uses a dummy for whether contract labor represents at least 50% of an establishment's workers as a dependent variable. This is a relevant alternative dependent variable as the Supreme Court's ruling that no absorption of the contract workforce is required may have made plants more willing to take the risk of relying on contract workers for a large share of their operations. The patterns in Panel (b) are consistent with those in Panel (a). While only 14 percent of plants with 20 to 49 workers relied on contract labor for at least half of their workers in 2015, 34 (40) percent of plants with 100–499 (more than 500) workers did.

We next examine whether the timing of the changes in the establishment size distribution also coincides with the SAIL case. Figure 4 plots the 50th, 75th, 90th, and 95th establishment size percentiles for manufacturing and services using the Economic Census rounds from 1990, 1998, 2005, and 2013. The left panel shows the sustained growth in the upper percentiles of manufacturing establishments around SAIL. The 90th and 95th percentile establishments have around 55 percent larger employment in 2013 compared to 1998, with slightly less growth at the 75th percentile.[28]

[27]The contract labor share also increases in plants with between 50 and 99 workers until about 2005, when it stabilizes until the end of the sample period. It is unsurprising that at least some of these plants may have opted to increase their reliance on contract labor, as they might be on the margin of exposure to Chapter VB through future employment growth. It is also possible that the use of contract workers lowers the bargaining power of permanent workers that increases with establishment size (we will later show evidence consistent with this). If so, some establishments with less than 100 workers may also employ contract labor for this purpose when secured that they will not need to absorb this workforce upon "abolishment.".

[28]Supplemental Appendix Figure B.1 (Bertrand, Hsieh, and Tsivanidis (2025)) repeats this figure using the annual manufacturing data from the ASI, and shows this growth in the right tail of the size distribution starts right after the SAIL decision in 2001. The Economic Census data consider the universe of formal and informal

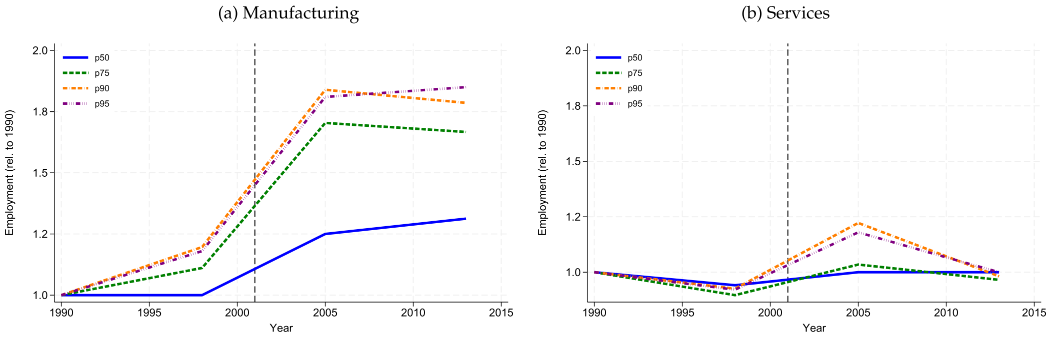

FIGURE 4.—Manufacturing versus services plant size distributions over time. *Note*: Plot shows percentiles of plant employment in manufacturing and service sector establishments from the Economic Census. Data come from 1990, 1998, 2005, and 2013. See Figure B.1 for finer plots for manufacturing data using annual data from the ASI.

The right panel in Figure 4 shows the growth in the upper percentiles of service sector establishments. While the IDA as a whole applies to all sectors, Section VB, which covers the majority of restrictions on retrenchments, applies only to manufacturing establishments, mines, and plantations with more than 10 workers. Since the firing restrictions of the IDA did not apply to services, there was no comparable change in the size distribution of establishments in services over the period when contract labor grew.

Figure 5 examines the timing of the change in the elasticity of value-added per worker with respect to establishment size. If the marginal product of labor is proportional to the average product and plants equate marginal products with factor costs, this elasticity measures the extent to which large plants face higher effective costs of labor than small plants. We regress log value-added per worker on log employment interacted with year dummies (and a full set of industry by year fixed effects) from 1980 to 2015, and plot the coefficient on log employment in each year.[29] The elasticity shows some increase between the late 1980s and the early 2000s, possibly because the reforms that began in 1991 removed most licensing restrictions and reservations for small plants, which may have made the labor constraints of the IDA more binding.[30] More importantly, Figure 5 shows that the elasticity of the average product of labor to establishment size fell after the early 2000s, possibly due to the SAIL event.[31]

establishments, while the ASI reports only formal employment, and therefore the changes in percentiles differ between the ASI and the Economic Census.

[29]Unlike Figure 3, this plot relies only on total employment, not broken down by permanent or contract, and therefore can be provided for every year from 1980 to 2015. We trim the 1% tails of VA/Worker to reduce the influence of outliers, as we do with other continuous and unbounded variables used in the analysis. Supplemental Appendix Section E.3 (Bertrand, Hsieh, and Tsivanidis (2025)) recreates this plot using the wage-bill to measure labor inputs.

[30]An earlier version of the paper showed that industries for which restrictions on FDI were lifted experienced large increases in this elasticity during the 1990s, with a smaller and imprecise increase in industries which delicensed.

[31]Supplemental Appendix Section E.10 (Bertrand, Hsieh, and Tsivanidis (2025)) shows that the break in trend in these elasticities around SAIL is significant. The Supplemental Appendix also reports three robustness checks. First, Figure E.5 in the Supplemental Appendix adjusts for possible differences in effective labor supplied by permanent and contract workers. Second, while the ASI does not provide firm identifiers so we cannot group plants into firms, it does provide information in certain years on the number of establishments operated by the firm which operates the plant. Supplemental Appendix Section E.1 reproduces the results on a

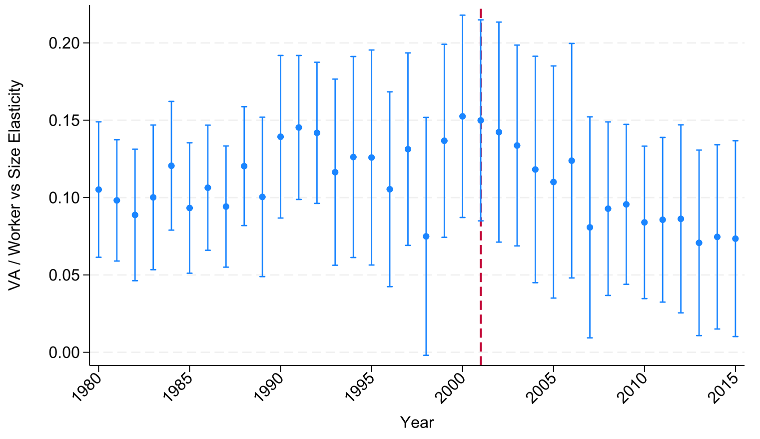

FIGURE 5.—Elasticity of VA/Worker to establishment size. *Note*: Plot shows coefficients and 95% confidence intervals from regressions of log VA/Worker on log plant employment interacted with year fixed effects. Regressions also include full set of industry-year fixed effects. Standard errors clustered at the industry level.

5.2. *Heterogeneity of SAIL Event*

We examine heterogeneity across Indian districts based on the initial supply of staffing companies in the district. The Contract Labor Act requires that plants access contract workers through government-licensed contractors or staffing companies. It is therefore likely that the SAIL shock we identified in the time series was larger for plants that were geographically closer to such staffing centers.

To isolate the supply of staffing plants uncorrelated with demand-side forces that may have spurred growth of the sector after the SAIL decision, we measure a district's proximity to staffing plants in the 1990 Economic Census. We use a distance-weighted proximity measure rather than the staffing employment within a district to capture that, although most districts in 1990 did not have plants providing staffing services, those close by still had access to these plants and the staffing industry as a whole radiated outwards from these initial clusters over time.[32]

We measure district d's proximity to staffing employment in 1990 as $\sum_{k \neq d} e^{-\kappa \text{dist}_{kd}} L_{k,1990}^{\text{Staffing}}$, where dist_{kd} is the number of kilometers between the centroids of districts d and k, $L_{k,1990}^{\text{Staffing}}$ is the number of workers employed by staffing plants in district k in 1990, and κ controls the rate at which the weight on surrounding staffing employment decays with distance.[33]

sample of plants owned by single-plant firms, and shows that the results are virtually identical. Third, the SAIL judgment may have changed the incentives for (large) plants to misreport employment of contract workers, so that the changes we document could be driven by changes in reported rather than actual employment. Supplemental Appendix Section E.2 argues against this misreporting concern by analyzing how both electricity use and sales respond when plants hire permanent versus. contract workers.

[32]See Supplemental Appendix Figure B.2 (Bertrand, Hsieh, and Tsivanidis (2025)) for this evidence. Using the distance-weighted proximity will be valid so long as the location of staffing plants in 1990 was unrelated to future trends in unobservables that affect manufacturing labor demand. Supplemental Appendix Section E.4 shows the characteristics of districts where staffing plants (i) initially located in 1990 and (ii) grew between 1990 and 2013. The evidence suggests that location choices of staffing plants in the early 1990s were unrelated to labor demand from manufacturing and instead were correlated with demand from the service sector.

[33]We exclude a district's own staffing employment since this may be endogenous to future outcome growth. We use a decay parameter of $\kappa = 0.0075$ in the main specifications, and vary this parameter in Supplemental

TABLE III

1990 STAFFING AND GROWTH OF CONTRACT LABOR USE IN MANUFACTURING PLANTS.

	CL Share	CL Share	P(Hire>50%)	P(Hire>50%)
ln Staffing × Post	0.012	0.011	0.015	0.014
	(0.005)	(0.007)	(0.006)	(0.008)
N Obs	593,338	562,468	593,338	562,468
N Clusters	437	367	437	367
R^2	0.20	0.20	0.17	0.17
State × Year FE	X	X	X	X
Industry × Year FE	X	X	X	X
District Controls × Year FE		X		X

Note: Observation is an establishment-year. Dependent variables are the establishment's share of workers that are contract workers (CL Share) and a dummy for whether the plant hires more than 50 percent of workers through contract labor (P(Hire>50%)). Post is a dummy for after 2001. Staffing is the weighted staffing employment in 1990 in all other districts with a decay rate of 0.0075. Controls include log district total employment, average formal manufacturing plant size, the 90th percentile of log formal manufacturing plant size, average log value-added per worker, the difference in log value-added per worker between large and small plants, log proximity to manufacturing employment in 1990 constructed in the same way as the staffing measure, and log proximity to staffing employment constructed in the same way, all measured in 1990. Standard errors clustered at district level.

TABLE IV

HETEROGENEOUS OUTCOME GROWTH POST-SAIL BY 1990 STAFFING.

	(1)	(2)	(3)	(4)	(5)	(6)
log Employment	0.044	0.061	0.056	0.057	0.046	0.055
	(0.017)	(0.021)	(0.016)	(0.016)	(0.017)	(0.032)
N Obs.	589,576	589,576	577,649	547,668	547,668	547,668
N Clusters	437	437	437	367	367	367
log VA/Worker	−0.022	−0.024	−0.024	−0.016	−0.020	−0.053
	(0.011)	(0.015)	(0.013)	(0.014)	(0.017)	(0.033)
N Obs.	469,598	469,598	459,745	436,595	436,595	436,595
N Clusters	436	436	436	367	367	367
State × Year FE	X	X	X	X	X	X
Industry × Year FE	X	X	X	X	X	X
Wght Man Emp × Year FE		X	X	X	X	X
Establishment Controls × Year FE			X	X	X	X
District Controls × Year FE				X	X	X
Wght Serv Emp × Year FE					X	X
District Emp Share Wghts						X

Note: Observation is an establishment-year. Each entry corresponds to the coefficient from a regression of the outcome in each row on the log staffing measure interacted with a post-SAIL dummy. Each column corresponds to a specification. Wght Man Emp refers to log proximity to manufacturing employment in 1990 constructed in the same way as the staffing measure. Establishment controls include dummies for plant ownership and organization type as well as a polynomial in establishment age. District controls include log district total employment, average formal manufacturing plant size, the 90th percentile of log formal manufacturing plant size, average log value-added per worker, and the difference in log value-added per worker between large and small plants. Wght Serv Emp refers to an additional control that is log proximity to service employment measure (computed similarly to the log staffing measure). The last column weights by each establishment's share of district employment (all other columns weight by sampling weights). Standard errors clustered at the district level.

Appendix Table A.1 (Bertrand, Hsieh, and Tsivanidis (2025)). The weight falls by $\kappa \times$ dist_{ij} percent for districts dist_{ij} km away from each other. For example, the average distance between all districts in Maharashtra is

Table III assesses whether the increase in the use of contract labor after SAIL was larger among plants located in districts that were closer to staffing employment. We regress a plant's contract labor share (columns 1 and 2) and the probability a plant hires more than 50% of its workforce through contractors (columns 3 and 4) on a post-SAIL dummy interacted with 1990 district-level staffing. Each regression includes industry-year and state-year fixed effects. Table III shows that the 2001 SAIL shock is more pronounced in districts that are closer to staffing centers. The point estimates are essentially unchanged when we include interactions between year dummies and a vector of 1990 district-level controls (even columns).[34]

Table IV then analyzes how employment and output per worker of the average plant in the district changed after SAIL in districts with greater access to staffing plants in 1990. We measure the change in these two outcomes for the average plant in a district, which includes the change within individual plants and the effect of plant entry and exit. Each entry corresponds to a different regression and reports the estimated coefficient on the interaction term between the post-SAIL dummy and the 1990 district-level staffing exposure measure. The table begins by including state-year and industry-year fixed effects, and then sequentially adds in establishment and district controls.

Table IV shows that average plant employment grew by more (row 1) and value-added per worker declines by more (row 2) after 2001 in districts with greater exposure to staffing, although the the differential decline in the average product of labor is noisy.[35] The results in columns 1–5 are an unweighted average across plants within a district. Since the variation in staffing is at the district level, column 6 shows the results are robust to using a plant's share of district employment instead.

5.3. *Establishment-Year Panel Analysis*

Since the ASI includes plant identifiers from 1993 to 2015, we can exploit the panel structure to estimate within-establishment changes associated with the use of contract labor in regressions with establishment fixed effects. We report this analysis in Table V. Each entry in the table corresponds to a different regression. Reported in the cell is the coefficient on the contract labor use variable. We consider both a dummy variable for any contract labor use and a dummy variable for contract labor accounting for at least 50% of employment. All regressions control for state-year fixed effects and industry-year fixed effects. We restrict the sample to the set of plants that use contract labor at any point in time, whether or not they use contract labor in a particular year. The results show that size increases (column 1) and value-added per worker falls (column 2) when plants begin to use contract labor.[36]

---

358 km with a minimum of 32 km and maximum of 891 km. With $\kappa = 0.0075$, this implies a weight of 0.06 on the average district and 0.001 on the furthest district in the state.

[34]These 1990 district-level controls, listed in the table notes, are the seven variables that are significantly associated with the staffing measure in Supplemental Appendix Table E.8 (Bertrand, Hsieh, and Tsivanidis (2025)).

[35]Supplemental Appendix Table E.6 (Bertrand, Hsieh, and Tsivanidis (2025)) shows the effect becomes sharper when measuring labor inputs via the wage-bill.

[36]Supplemental Appendix Section E.6 (Bertrand, Hsieh, and Tsivanidis (2025)) presents event study analyses that show the evolution of these key outcomes in the years that precede and follow the first hiring of contract labor.

TABLE V

CORRELATES OF CONTRACT LABOR HIRING WITHIN PLANTS.

	Employment	VA/Worker	Inaction	Job Creation	Add Product
Contract	0.369	−0.210	−0.028	0.115	0.008
	(0.013)	(0.010)	(0.005)	(0.006)	(0.004)
Contract > 50%	0.365	−0.236	−0.026	0.136	0.009
	(0.016)	(0.017)	(0.007)	(0.006)	(0.005)

Note: Sample are plants that use contract labor at any point in time. Entries report the coefficient from a regression of the outcome on a dummy for the years in which the establishment hires contract workers (any contract worker in row 1 and contract workers for more than 50% of its employment in row 2). Employment and VA/Worker are in logs. Inaction is defined as a dummy for whether a plant's employment growth rate is less than 10% in absolute value. Job creation rate is a plant's employment growth rate $g_{it} = \frac{L_{it}-L_{it-1}}{0.5\times(L_{it}+L_{it-1})}$ for expanding plants and zero otherwise. Add Output Product is a dummy for whether a plant adds a new 5-digit product to its output lineup relative to the previous year. All regressions include state-year, industry-year, and establishment fixed effects. Standard errors clustered at the industry level.

6. HOW DID THE RISE OF CONTRACT LABOR FREE UP ESTABLISHMENT GROWTH?

In this section, we look for evidence for two channels through which the use of contract workers may have benefited large plants in India. First, the more widespread availability of contract workers may have prompted large Indian plants to employ more workers and undertake more risky investments because they are no longer subject to firing costs. Second, contract workers may also have increased the bargaining power of large employers with respect to their permanent workers.

6.1. *Reductions in Labor Adjustment Costs*

When an establishment receives a positive labor demand shock that may be reversed in the future, the firing cost can make it reluctant to expand—large plants subject to a moderate positive shock today will not hire additional workers with the knowledge that they will most likely have to fire them in the future. The firing cost could also discourage plants from undertaking risky investments. The use of contract workers, by reducing firing costs, could reduce the inaction band in employment and prompt plants to undertake risky investments. In this subsection, we look for evidence consistent with these mechanisms.

Consider first the time series around SAIL in Figure 6. For Panel (a), we first define in the establishment-year data a variable called "inaction" to which we assign a value of 1 if the establishment did not change its employment by more than 10% (in absolute value) from one year to the next. We regress this inaction dummy on log employment interacted with year dummies, as well as a full set of industry-year and state-year dummies. Figure 6 shows the coefficients on log employment for each year. Throughout the sample period, the likelihood of inaction increases with establishment size. Most relevant to us, and consistent with a decrease in relative adjustment costs at large plants post-SAIL, is the decline in the strength of this inaction to establishment size elasticity after 2001.[37] Panel (b) displays the gross job creation rates by establishment size over time (relative to 1994).[38] Here again, we observe an uptick in job creation by larger plants starting in the early 2000s relative to plants with fewer than 100 workers.

[37]Figure B.3 in Supplemental Appendix (Bertrand, Hsieh, and Tsivanidis (2025)) repeats Panel (a) for two alternative ways of measuring plant "inaction" or employment dynamism.

[38]We calculate the job creation rate from expanding plants in each size bin. From this sample, the job creation rate for each size is the ratio of the sum of employment change across expanding plants in each size

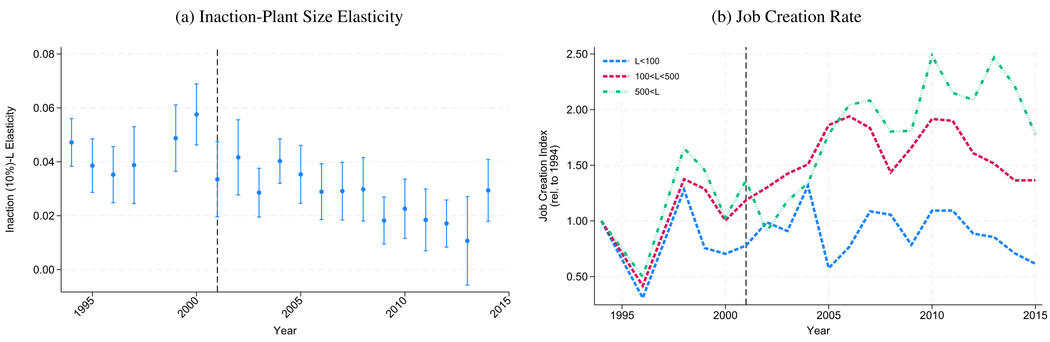

FIGURE 6.—SAIL and employment dynamics. *Note*: Panel (a) shows coefficients and 95% confidence intervals from regressions of a dummy for whether a plant's annual employment growth rate exceeds 0.1 in absolute value on log plant employment interacted with year fixed effects. Regression also includes full set of industry-year and state-year fixed effects, a fourth-order polynomial in plant age, and dummies for the organization and ownership type of the establishment, all interacted with year fixed effects. Standard errors clustered at industry level. Panel (b) shows job creation rates by size bin over time (relative to 1994), defined as the positive employment change in each size bin divided by the average aggregate employment across both start and end years. 1995 and 1997 are omitted due to large spikes in those years (one positive and negative, so no substantive impact on the trend in the pre-period).

Another way to examine whether labor started to appear like a more flexible input post-SAIL is to compare it to another flexible input such as electricity. We do this by comparing the dispersion of the average revenue products of labor and electricity (within industry-year cells) over time in Figure 7, where the dispersion of each input is normalized to 1 in 2001.[39] Panel (a) shows that while the dispersion of the average revenue product of electricity is stable before and after SAIL, the dispersion of the average product of labor falls by about 25% beginning right after the SAIL decision. Panel (b) shows this is driven mostly by large plants with more than 100 workers.

The establishment-year panel analysis in Table V also provides evidence of such greater dynamism when a given establishment uses contract labor. The last three columns in Table V show that an establishment is less likely to be in the inaction range, the job creation rate is higher, and is more likely to add new products to its output portfolio when it has contract workers on its rolls.[40]

We next examine the effect of contract labor to the sensitivity of plant employment to economic shocks. We present two approaches. We first consider how districts differentially respond to local shocks based on their usage of contract labor. We construct Bartik-style instruments for growth in manufacturing employment and run regressions of the form

$$g_d = \beta_0 + \beta_1 \widehat{g}_d + \beta_2 \text{Contract Init}_d + \beta_3 \widehat{g}_d \cdot \text{Contract Init}_d + \gamma_s.$$

bin divided by the average of total establishment employment in each size bin at the beginning and end of each period. A period is one year.

[39]We define dispersion of a variable as its standard deviation. The figure plots the dispersion of the residuals from a regression of the log average product of labor and the log average product of electricity on industry-year fixed effects.

[40]Supplemental Appendix E.5 (Bertrand, Hsieh, and Tsivanidis (2025)) examines whether plants produce riskier products when they use contract labor. We find that the added products are neither more nor less risky than the products they made previously.

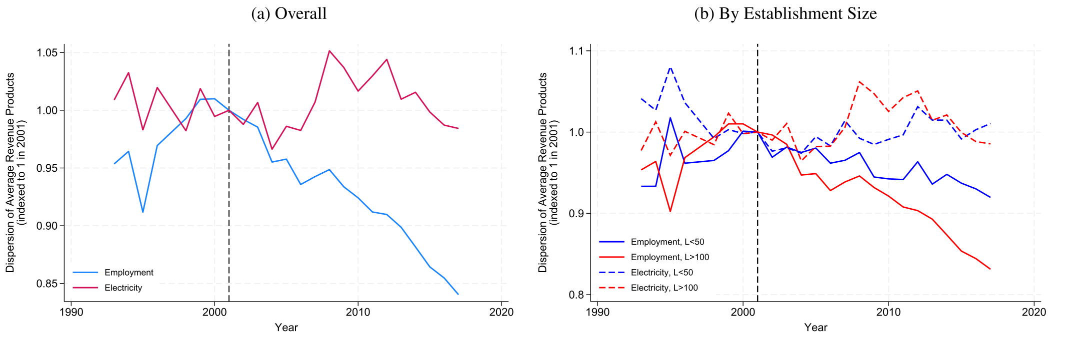

FIGURE 7.—SAIL and the dispersion of average revenue products of labor and electricity. *Note*: Both figures begin by regressing the average revenue products of labor (VA/Worker) and electricity (VA/KWH) on industry-year fixed effects. Panel (a) then computes the dispersion of the residuals from this regression for each input, normalizing each series to 1 in 2001. Panel (b) does the same but breaking down by plants with less than 50 workers or more than 100 workers. Dispersion defined as the standard deviation.

Here $g_d \equiv (L_{d,t+k} - L_{dt})/L_{dt}$ is the growth rate of manufacturing employment in district d between dates t (1997–1999) and $t+k$ (2007–2009), Contract Init_d is the share of manufacturing plants using contract workers in the initial period (1997–1999), $\widehat{g}_d \equiv (\widehat{L}_{d,t+k} - L_{dt})/L_{dt}$ is the predicted growth rate in employment in the district, and γ_s are state fixed effects.[41] To measure predicted employment growth in a district, we start by computing growth rates of employment at the industry level between dates t and $t+k$, and then take the weighted average of these industry-specific growth rates, using initial district industry employment share as weights. We then define predicted employment in a district $\hat{L}_{d,t+k}$ by multiplying initial district employment by the predicted growth rate.[42]

Column 1 of Table VI presents the first stage, which shows that the instrument has good predictive power for district-level employment changes. The slope is 0.789, and the F-stat is 36.41. Column 2 reports an alternative first stage that additionally controls for the vector of district-level conditions in 1990 from Table III. Again, the instrument has good predictive power.

We then examine how the initial contract share of a district (computed between 1997 and 1999) affects the responsiveness of actual employment growth to predicted employment growth during the 2000s. The results suggest that contract labor has a significant effect on responsiveness to shocks: increasing the contract share by one standard deviation raises the elasticity by about.43 (columns 3 and 4, where the latter controls for district-level conditions in 1990). Columns 5 and 6 show that this result strengthens as we allow for differential responsiveness to such economic shocks across Indian states, while column 7 shows it is robust to allowing for differential responsiveness by district characteristics (by adding interactions between predicted employment growth and district controls). Columns 8 and 9 replicate columns 5 and 6 but use the district's access to staffing employment in 1990 (as computed in Section 5.2) as an alternative measure of access to staffing employment.

[41]The ASI only provides district identifiers until 2009. We pool years into a pre- and post-period to increase precision.

[42]We exclude own-district employment when computing national industry growth rates, and standardize the contract share to have zero mean and unit standard deviation.

TABLE VI

CONTRACT LABOR AND RESPONSIVENESS OF DISTRICT EMPLOYMENT TO LOCAL BARTIK LABOR DEMAND SHOCKS.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Predicted Emp	0.789	0.761	0.978	0.892					
Growth	(0.130)	(0.147)	(0.126)	(0.174)					
Initial Contract Measure			−0.065	−0.064	−0.063	−0.027	−0.022	−0.140	−0.088
			(0.082)	(0.080)	(0.090)	(0.077)	(0.080)	(0.091)	(0.109)
Predicted Emp Growth ×			0.449	0.421	0.616	0.515	0.491	0.794	0.561
Initial Contract Measure			(0.222)	(0.220)	(0.263)	(0.212)	(0.234)	(0.193)	(0.202)
R^2	0.48	0.47	0.50	0.49	0.54	0.60	0.62	0.57	0.59
N Obs	390	335	390	335	385	335	335	385	335
F-Stat	36.41	27.73							
State FE	X	X	X	X	X	X	X	X	X
District Controls		X		X		X			X
State FE × Pred. Emp Growth					X	X	X	X	X
District Cont. × Pred. Emp Growth							X		
Staffing Measure								X	X

Note: Observations at district level. Outcome is the growth in district ASI employment between 1997–1999 and 2007–2009. Predicted Emp Growth is predicted employment growth rate according to the Bartik measure using the aggregate rate of employment growth across industries in all other districts. Initial contract measure is share of plants using contract labor in the district between 1997 and 1999, standardized to have unit standard deviation (except in columns (8) and (9) where it is the log staffing measure, also standardized). Regressions weighted by the district's average number of observations across both pre- and post-periods. Controls are same as in Table III. Standard errors clustered by district.

Table VII moves the analysis back to the establishment-year panel. Here we follow Chaurey (2015) and use annual rainfall in a district as an alternative economic shock. The variable "shock" in Table VII takes the value of 1 if rainfall in the establishment's district in that year is below the 20th percentile in that district's average annual rainfall distribution between 1990 and 2010, −1 if rainfall in the establishment's district in that year is above the 80th percentile in the district's distribution, and 0 otherwise.[43] The dependent variable in all regressions is log employment.[44]

The patterns in Table VII are consistent with the view that the use of contract labor helped reduce labor adjustment costs after the SAIL decision. Column 1 shows that there was no differential responsiveness of employment to shocks overall in the post-SAIL period. However, column 2 shows that plants employing a higher fraction of their workforce through contractors fired more workers in response to negative rainfall shocks than those employing less contract labor after SAIL. Column 3 shows that the employment responses to rainfall shocks become more pronounced amongst large plants relative to small ones. Finally, column 4 shows greater employment responses to rainfall shocks post-SAIL in plants located in districts that are closer to staffing employment in 1990.

6.2. *Reductions in the Cost of Labor*

A more widespread reliance on contract labor may reduce the cost of labor, especially so for larger plants. The difference in cost between contract and permanent workers might

[43] Adhvaryu, Chari, and Sharma (2013) show that rainfall shocks are associated with drops in agricultural production, wages, and district per capita expenditure.

[44] All regressions include establishment fixed effects, state-year fixed effects, industry-year fixed effects, and interact district-level conditions in 1990 with year dummies.

TABLE VII

CONTRACT LABOR AND RESPONSIVENESS OF ESTABLISHMENT EMPLOYMENT TO RAINFALL SHOCK.

	(1)	(2)	(3)	(4)
Shock	0.007	0.002	0.004	−0.032
	(0.006)	(0.006)	(0.006)	(0.018)
Shock × Post	−0.005	0.007	0.000	0.057
	(0.008)	(0.009)	(0.008)	(0.024)
Contract		0.609		
		(0.022)		
Contract × Post		0.266		
		(0.024)		
Shock × Contract		0.042		
		(0.017)		
Shock × Contract × Post		−0.089		
		(0.022)		
Large			1.127	
			(0.020)	
Large × Post			0.139	
			(0.013)	
Shock × Large			0.015	
			(0.010)	
Shock × Large × Post			−0.039	
			(0.011)	
Staffing × Post				0.030
				(0.014)
Shock × Staffing				0.013
				(0.005)
Shock × Staffing × Post				−0.019
				(0.007)

Note: Observations at the establishment-year level. Outcome is log employment. Post is a dummy for after 2001. (Negative) Shock is defined at the district level and defined by relative rainfall in a year relative to the average. It takes a value of 1 when rainfall is below the 20th percentile of a district's distribution, −1 when above the 80th percentile, and 0 otherwise. Contract Share is the establishment's share of workers hired through contractors. Large is a dummy for whether the establishment has more than 100 workers. Staffing is the log weighted staffing employment in 1990 in all other districts with a decay rate of 0.0075. All regressions include establishment fixed effects, state-year fixed effects, industry-year fixed effects, and district controls-year fixed effects. District controls are same as in Table III. Standard errors clustered at the district level.

be particularly large for plants that are more regulated, as these plants are more likely to be unionized, more subject to strikes and other forms of "labor militancy," all of which may drive up the wage of their permanent workers. Furthermore, an increased reliance on contract labor may lower the wage of permanent workers by lowering their bargaining power.

Does Contract Labor Cost Less? Panel (a) in Figure 8 plots the elasticity of average wages with respect to total employment over time. The average wage is defined as the total wage bill for non-managerial workers divided by the number of workers. As before, we first regress log average wage on log employment interacted with year dummies, as well as a full set of industry-year dummies, and report in the figure coefficients on log employment for each year. There is a positive elasticity of the average wage to employment throughout the sample period. While this positive elasticity is quite stable at about .14 from 1980 to 2001, there is a break in trend in the early 2000s when the elasticity starts sharply declining, dropping to about.075 by 2013–2015. This time series evidence therefore shows that the SAIL event also coincided with a sharp decline in the gap in average

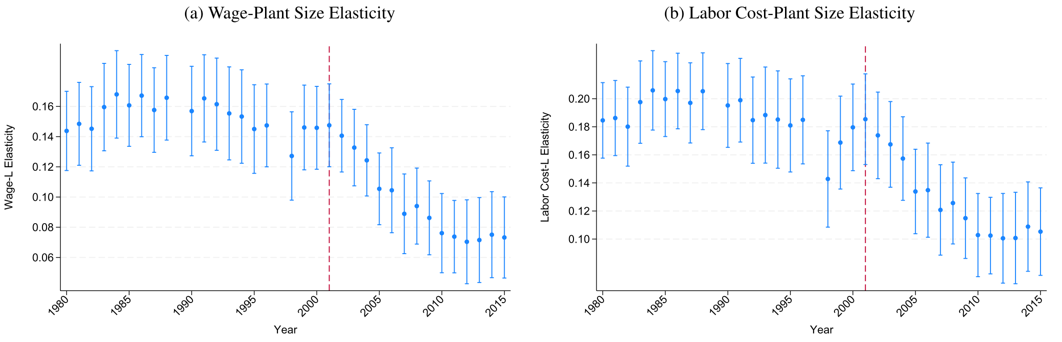

FIGURE 8.—Labor cost and plant size over time. *Note*: Figures plot coefficients and 95% confidence intervals from regression of log wage per worker (Panel (a)) and log labor cost per worker (defined as wages, bonuses, and benefits in Panel (b)) on log employment interacted with year fixed effects. Regressions also include full set of year-industry fixed effects. Standard errors clustered at the industry level.

wage between large and small plants. Panel (b) replicates Panel (a) but focuses on average daily labor cost. Labor cost sums wages, bonuses, as well as various benefit payments (such as contributions to provident and other funds and other welfare expenses). Again, we see a positive and rather stable elasticity of daily labor cost to establishment size (except for two outlier years) from 1980 to 2001 of about.18, and a break in trend in the early 2000s, with the elasticity reaching.1 by the end of the sample period.

While suggestive of a reduction in the average cost of labor induced by the rise of contract labor, it is possible that they reflect changes in the composition and quality of workers.

To address this possibility, we measure the wage gap between permanent and contract workers holding constant the composition of employment within the establishment in a flexible fashion. For each type of labor $\ell \in \{\text{Contract}, \text{Permanent}\}$, we run the following specification for establishment i in year t:

$$\ln W_{\ell itb} = \gamma_{ktb} + \beta_t \mathbb{I}\{\ell = \text{Contract}\},$$

where $W_{\ell it}$ is the average daily wage of type-ℓ workers, γ_{ktb} are industry-year-bin fixed effects, and b is a group indicator for the share of contract workers at establishment i. We bin the contract labor share into five groups depending on whether the establishment employs no contract workers, between 0 and 24%, 25–49%, 50–74%, or 75–100% of workers through contracting. By controlling for the composition (e.g., share contract vs. permanent workers) of employment by industry-year cell, we hope to capture differences in the type of permanent and contract workers employed by plants with different shares of work contracted out. Although our evidence in Table I suggests that the tasks performed by contract and permanent workers are quite comparable, these controls allow for the effect of worker composition to vary within each industry-year cell.

Figure 9 plots the estimates of β_t, which identifies the average wage difference between contract and permanent workers. We also repeat the analysis using total labor costs (wages, bonuses, and benefits) as the outcome variable.[45] Contract workers are about 25% cheaper than permanent workers in terms of wages, and about 30% cheaper in terms of

[45]Recall wages for permanent and contract workers are separately provided between 1998 and 2015; while the same is true for bonuses and benefits between 1998 and 2007.

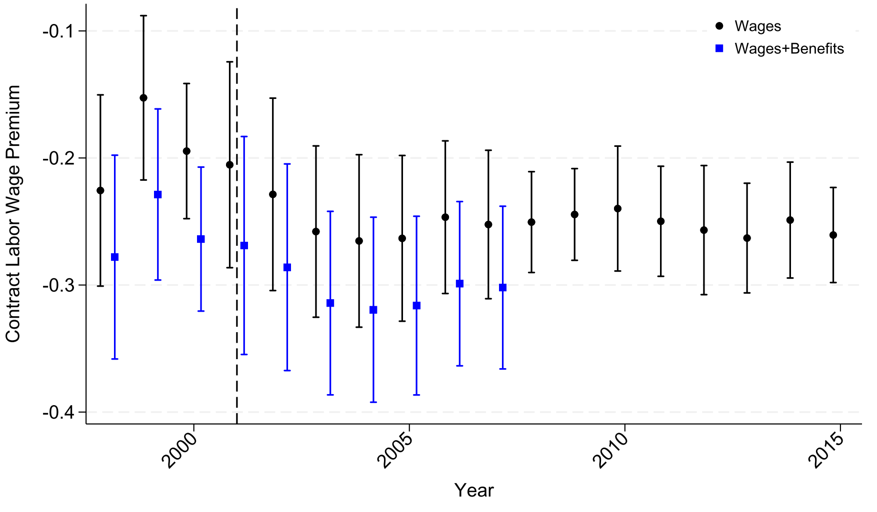

FIGURE 9.—Relative cost of contract labor. *Note*: Figures plot coefficients and 95% confidence intervals from a regression of log wage per worker on a dummy for whether the worker category is contract (relative to the omitted category of permanent workers) interacted with year fixed effects, as well as a full set of industry-year-contract labor share bin fixed effects, where contract labor share bins are dummies for whether the plant hires no contract workers, between 0 and 24%, 25–49%, 50–74%, or 75–100% of workers through contracting. Wages+Benefits cover wages, bonuses, and benefits. Wages and benefits are only provided separately by type of worker from 1998 to 2007. Standard errors clustered at the industry level.

overall payments. After some swings in the late 1990s and early 2000s, these wage and labor cost differences appear quite stable across the 2000s, even though confidence intervals are large.

In Supplemental Appendix Section E.8 (Bertrand, Hsieh, and Tsivanidis (2025)), we use CPHS data to measure the wage gap between permanent and contract workers after controlling for worker characteristics. We estimate that contract workers earn about 25% less than observationally similar permanent workers, which lines up with the wage gap estimated in the ASI.

Does Contract Labor Bring Down Costs Disproportionately for Large Plants? More relevant for our purpose is assessing whether the relative price of contract workers (compared to permanent workers) differs by establishment size. In Panel (a) of Figure 10, we plot the raw non-parametric relationship between the relative wage of contract workers and total plant employment in 2000 and 2015. Consistent with the view that the difference in the cost between permanent and contract workers is greater for large plants, we observe a downward sloping relationship. Consider the relationship in 2000: in plants with 10 workers, contractors are paid about 10% less than permanent staff, but this difference increases to 25% (65%) in plants with 100 (1000) workers. In Panel (b), we observe the same qualitative pattern when we residualize the relative wage by industry-year-contract labor share bin fixed effects. The downward slope is particularly steep in 2000 for plants above the 100 permanent workers mark. This downward slope is consistent with the additional bargaining power we hypothesize permanent workers to have in larger plants, and one of the reasons why the rise in contract labor may have reduced the gap in labor cost between larger and smaller plants.

Panel (b) also shows that the relationship between the relative wage of contract versus permanent workers and plant size flattens in 2015 compared to 2000 for plants with more

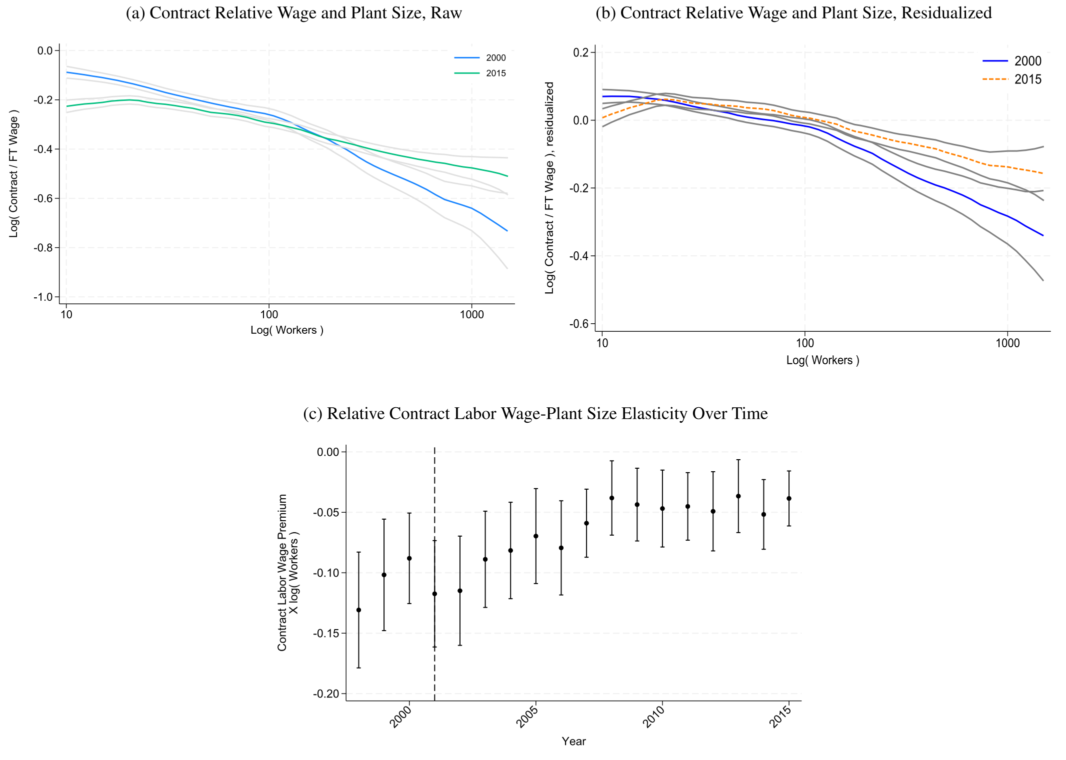

FIGURE 10.—Contract and permanent relative wages and plant size over time. *Note*: Sample are plants which hire contract workers. Panel (a) plots the non-parametric relationship between plant size and the log relative average wage per worker between contract and permanent workers. In Panel (b), we repeat the exercise but first regress the log relative average wage per worker on a set of industry-year-contract labor share bin fixed effects. We then plot the residualized relative wage against plant employment. In Panel (c) we run the same specification as the regression in Panel (b) and add interactions between the contract × year dummies with log employment. We then plot the coefficients and 95% confidence intervals on the contract × year × log employment.

than 100 permanent workers, yet is almost identical in the two years for smaller plants. Panel (c) explores the timing of this change by estimating the following regression:

$$\ln W_{\ell itb} = \gamma_{ktb} + \beta_{1t}\mathbb{I}\{\ell = \text{Contract}\} + \beta_{2t}\ln L_{it} + \beta_t^{Size}\mathbb{I}\{\ell = \text{Contract}\} \times \ln L_{it},$$

where L_{it} is the number of workers. Panel (c) of Figure 10 plots β_t^{Size} from this regression which capture the extent to which the wage differential between contract and permanent workers varies with the number of workers employed at the establishment. While our data only allow us to examine this relationship from 1998 onwards, it appears that the fall in the wage premia of permanent workers within large plants began around or just after the SAIL adjudication.

Did Permanent Labor Become Cheaper for Large Plants? Figure 10 suggests that the cost of permanent workers relative to contractors fell for larger plants during the 2000s. In Figure 11, we diagnose whether this was driven by an increase in contract wages or a fall in permanent wages at larger plants during the 2000s. Panel (a) plots the elasticity over time of the average wage per contract worker to the number of workers (constructed in the same way as the previous elasticity plots). There is a positive elasticity of around 0.05

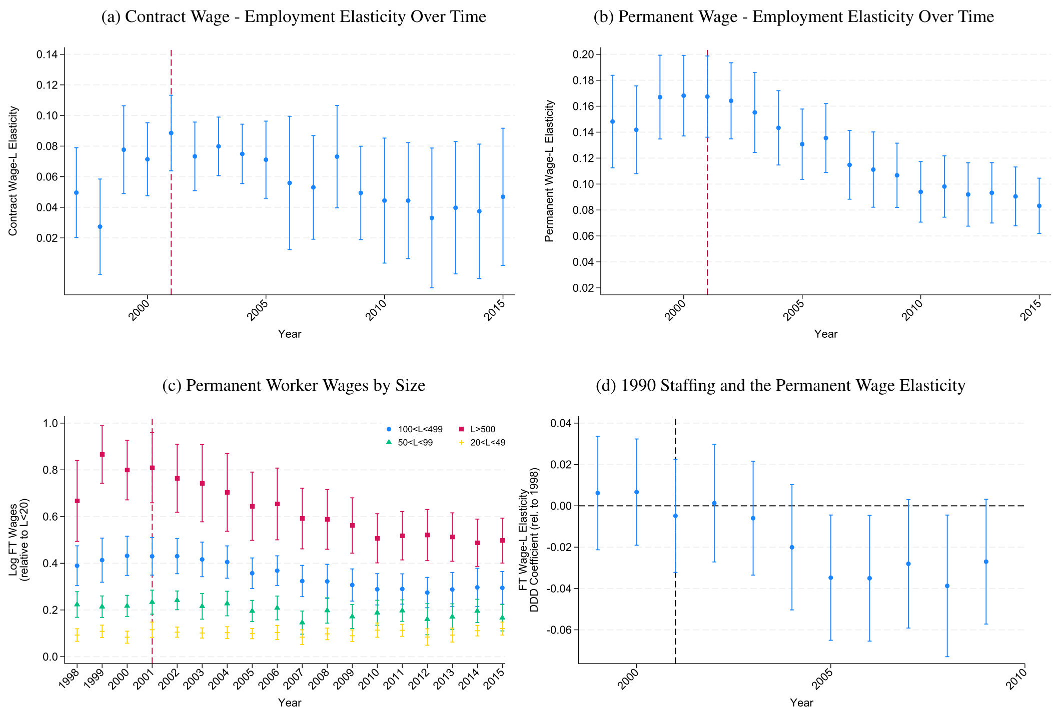

FIGURE 11.—Contract and permanent wages and plant size over time. *Note*: In Panel (a), we regress the log average contract wage on log plant number of workers interacted with year dummies (as well as a set of industry-year-contract labor share bin fixed effects) and plot the employment-year coefficients along with the 95% confidence intervals. In Panel (b), we do the same for the wages of permanent workers. Panel (c) examines the relationship less parametrically by looking at average wages paid to permanent workers by plants with different numbers of employees. Panel (d) regresses log average wage of permanent workers on a full interaction of log employment, log 1990 staffing and year dummies, as well as a set of district, state-year, industry-year-contract labor bin fixed effects, and 1990 district characteristics interacted with year fixed effects. The triple interaction coefficients (and the corresponding 95% confidence intervals from standard errors clustered at the district-year level) are plotted, and are interpreted as the change in the permanent wage plant size elasticity in a given year for a 1% increase in the 1990 staffing measure, relative to the omitted category of 1998. District identifiers are not provided after 2009, so Panel (d) ends then. In Panels (a)–(c), standard errors are clustered by industry, while in Panel (d), they are clustered by district (since the 1990 staffing variation is at the district level).

over the period, suggesting that larger plants faced higher wages to hire contract workers. This elasticity rises and then falls slightly around the SAIL event, but the magnitude of the change is relatively small.

Panel (b) repeats this analysis but focuses on changes over time in the elasticity of the average wage of permanent workers to the number of workers. Here, we observe a pronounced drop post-SAIL. Permanent workers became disproportionately cheaper for large plants starting in 2001. While our wage data by worker category only begin in 1998, the lack of a pre-trend in the wage elasticity in Figure 9 suggests the permanent wage elasticity was likely constant prior to 1998 given the dominance of permanent vis-a-vis contract workers during those early years. Panel (c) runs a less parametric regression to examine how average wages paid to permanent workers by plants with different numbers of employees changed over time. Relative to small plants with less than 49 workers, we see some convergence in permanent wages after the SAIL decision.

Overall, Panels (a) to (c) of Figure 11 suggest that the rise in the relative cost of contract workers amongst large plants documented in Figure 10 was driven by a fall in the cost of permanent workers rather than a rise in the cost of contract workers. Panels (b) and (c) suggest that this change lines up fairly closely with the SAIL decision.

Panel (d) in Figure 11 examines how the relationship between the elasticity of permanent wages to employment and a district's level of staffing in 1990 evolved over time. If the SAIL decision was the principal factor driving the downward trend in this elasticity during the 2000s, then we expect that districts with more staffing available (in 1990) should experience a larger decline after 2001. To test this, we regress log wage of permanent workers on a full interaction of log employment, log 1990 staffing and year dummies, as well as a set of district, industry-year-contract labor share bin and state-year fixed effects and 1990 district characteristics interacted with year fixed effects. The estimated triple interaction terms in this difference-in-difference-in-difference regression, interpreted as the change in the permanent wage plant size elasticity in a given year for a 1% increase in the 1990 staffing measure, relative to the omitted category of 1998, are reported in Panel (d). They are mostly negative post-SAIL, consistent with the permanent worker wage-plant size elasticity falling more in districts with greater exposure to staffing in 1990.[46]

7. A MODEL OF ESTABLISHMENT GROWTH, INNOVATION, AND FIRING COSTS

The empirical evidence above suggests that the increased use of contract labor is likely due to a decrease in the cost of tapping into this alternative employment source rather than higher demand for such workers. Furthermore, it suggests that plants that choose to employ more contract labor do so because it reduces future labor adjustment costs. Lastly, it also suggests that easing access to contract labor allowed large plants to grow.

In this section, we use a model to quantify the effect of the IDA, and of the expansion of contract labor in the presence of the IDA, on aggregate TFP. If all we were interested in is the static output loss from the misallocation of labor due to the IDA, then the model in Section 4 where the probability a firm's product is destroyed is an exogenous parameter would suffice. However, we also want to know the effect of the expansion of contract labor on the growth rate by changing the incentives to innovate. Therefore, we extend the model in Section 4 by endogenizing the innovation rate. We assume innovation takes the form of improving the quality of another firm's product. Therefore, the probability a firm's product is destroyed is endogenous to the innovation rate.[47]

Finally, to capture the fact that higher use of contract labor is associated with a decline in the average product of labor among the plants that increase their use of contract labor, we assume that the expansion of contract labor is driven by a reduction in the fixed cost of using contract labor F. We then use the model to assess the effect of a decline in F on static misallocation and on the aggregate innovation rate by changing the innovation incentives of high-type establishments.

[46]In Supplemental Appendix Section E.9 (Bertrand, Hsieh, and Tsivanidis (2025)), we use CPHS data to assess whether the patterns in Figure 11 could be explained by increased negative selection of permanent workers when the contract labor share increases. We find that selection looks positive, with permanent workers being somewhat older and slightly more educated as the contract labor share in their industry increases.

[47]It is possible that innovation also takes the form of new products or quality improvements in the firm's own products. However, these alternative models of innovation cannot generate the negative shocks that are essential for the IDA to have empirical bite. The reason is that innovation in the form of new products or own quality improvements raises the wage but otherwise has no negative implications for incumbents that do not innovate.

7.1. *A Model of Endogenous Innovation*

When an establishment successfully innovates, it improves upon the quality of a randomly chosen product with step-size λ, where the step-size follows a Pareto distribution with unit scale and shape parameter θ. The cost of innovation (in units of the final good) *per product* is $c_H(x_H) = (\frac{x_H}{\xi_H})^{\frac{1}{1-\beta}} Y$, where x_H is the flow rate of innovation per product owned by a high-type establishment and ξ_H is the productivity of the high-type establishment in R&D. The cost of innovation for the low-type establishment is given by a similar expression, with x_H and ξ_H replaced by x_L and ξ_L.

The marginal private benefit of resources spent on innovation is the product of the marginal increase in innovation from additional R&D and the expected value of a variety obtained through innovation. Equating the marginal cost with the marginal benefit of the innovation, the optimal innovation rate of the high-type establishment x_H is

$$x_H = \tilde{\beta} \xi_H^{1/\beta} \mathbb{E}\big[v_H(\lambda \hat{q}_j)\big]^{\frac{1-\beta}{\beta}},$$

where v_H is the normalized value of a variety of high-type plants and $\hat{q}_j$ is the normalized quality of a product q_j.[48] The optimal innovation rate for a low-type establishment x_L is given by the same expression with v_H and ξ_H replaced by v_L and ξ_L. Holding the expected value of a variety constant, $\xi_H > \xi_L$ implies that the innovation rate of high-type plants is higher than that of low-type plants.

Turning to entrants, their cost of innovation is $(x_E/\xi_E)^{\frac{1}{1-\beta}} Y$, where x_E denotes the innovation intensity of an entrant. The type of an entrant (high or low type) is realized after they invest in R&D, where $\alpha \leq 1$ denotes the ex ante probability an entrant is a high-type establishment. The expected return to innovation for an entrant is a weighted average of the value of a product for a high- and a low-type establishment. Equating the cost to the benefit, the optimal innovation rate for an entrant is

$$x_E = \tilde{\beta} \xi_E^{1/\beta} \big\{\alpha \mathbb{E}\big[v_H(\lambda \hat{q})\big] + (1-\alpha)\mathbb{E}\big[v_L(\lambda \hat{q})\big]\big\}^{\frac{1-\beta}{\beta}}.$$

The innovation rate of entrants is increasing in the productivity of entrants in R&D and in the weighted average of the value of a variety for a high- and low-type establishment.

The key endogenous variables in the expressions for the optimal innovation rates are the expected value of a variety for high- and low-type plants. For a low-type establishment, the expected value of a variety is given by a standard arbitrage equation:

$$\mathbb{E}\big[v_L(\hat{q})\big] = \frac{\sigma^{-1}\mathbb{E}\big[\hat{q}^{\sigma-1}\big]}{\rho + x + (\sigma-1)g} + \frac{\beta\tilde{\beta}\xi_L}{\rho + x}\mathbb{E}\big[v_L(\lambda \hat{q})\big]^{1/\beta},$$

where g and x denote the growth rate and the innovation rate. The first term is the expected value of the flow of profits from owning a variety. The second term is the expected value from innovating and possibly grabbing another variety.

[48]Both are normalized by aggregate output per worker. Also $\tilde{\beta} \equiv (1-\beta)^{\frac{1-\beta}{\beta}}$.

For a high-type establishment, the expected value of a variety v_H is given by a similar arbitrage condition:

$$\begin{aligned}\mathbb{E}\big[v_H(\hat{q})\big] = {} & \big(\mathbb{P}(\hat{q} < \hat{q}^*)(1+x\kappa)^{1-\sigma}\mathbb{E}\big[\hat{q}^{\sigma-1}|\hat{q} < \hat{q}^*\big] + \mathbb{P}(\hat{q} > \hat{q}^*)\big(\mathbb{E}\big[\hat{q}^{\sigma-1}|\hat{q} > \hat{q}^*\big] \\ & + \mathbb{E}\big[e^{-(r-g+x)\tilde{t}(\hat{q})}|\hat{q} > \hat{q}^*\big]\varepsilon\big)\big) \\ & /\big(\sigma\big(\rho + x + (\sigma-1)g\big)\big) \\ & - \mathbb{P}(\hat{q} > \hat{q}^*)\frac{F}{\rho + x} + \frac{\beta\tilde{\beta}\xi_H}{\rho + x}\mathbb{E}\big[v_H(\lambda\hat{q})\big]^{1/\beta},\end{aligned}$$

where $\mathbb{P}(\hat{q} > \hat{q}^*)$ denotes the probability that the normalized quality of the innovated variety exceeds the normalized threshold quality, and $\tilde{t}(\hat{q}) = g^{-1}[\ln(\hat{q}) - \ln(\hat{q}^*)]$ denotes the duration for which the normalized quality of the innovated variety remains above the normalized threshold quality.[49]

The expected value of a variety for a high-type establishment can be interpreted as follows. The first line is the expected flow of profits from owning a variety. Since $(1 + x\kappa)^{1-\sigma} < 1$, the value of a product of a given quality is lower for a high-type establishment because of the possibility it will be forced to hire too few permanent workers due to the firing cost. The value is adjusted for the length of time $\tilde{t}$ the establishment expects to hire contract labor conditional on drawing a productivity initial above the cutoff $\hat{q}^*$. The second term is the fixed cost paid if the product is of high enough quality to employ contract workers. The third term is the expected gain from innovation.

The share of products owned by the two types of plants and the aggregate rate of innovation are then pinned down by the rates of innovation. Specifically, the share of products owned by high-type plants ϕ in a steady state is

$$\phi = \frac{(x_H - x_L - x_E) + \sqrt{(x_H - x_L - x_E)^2 + 4(x_H - x_L)\alpha}}{2(x_H - x_L)}.$$

The steady-state share of high-type plants is increasing in x_H and decreasing in x_L. The aggregate rate of innovation x is then given by $x = \phi x_H + (1-\phi)x_L + x_E$. To get a stationary quality distribution, we add a reflecting barrier where the bottom ψ percent of products draw new qualities from $j \in [\psi, 1]$ and the quality of the top ψ percent of products is not upgraded. The expected growth rate of aggregate output Y is then given by

$$g = \frac{1}{1-\psi}\left(x \cdot \frac{1}{\theta - (\sigma - 1)}\right),$$

which is the product of the innovation rate x and the average step size $\frac{1}{\theta-(\sigma-1)}$ adjusted by $\frac{1}{1-\psi}$.

7.2. *Parameter Estimates*

The model is characterized by 10 parameters: $\{\rho, \sigma, \beta, \alpha, \kappa, \xi_H, \xi_L, \xi_E, \theta, F\}$. We impose values for ρ, β, and σ and estimate the remaining parameters in three steps.[50]

[49]And $\varepsilon \equiv \frac{1}{\sigma}\frac{(1+x\kappa)^{1-\sigma}-1}{[\rho+x+(\sigma-1)g]}\hat{q}^* + \frac{F}{\rho+x}$.

[50]We pick $\beta = 0.5$ to match the elasticity of successful innovation with respect to R&D. In addition, we assume the elasticity of substitution across products $\sigma = 2$, a discount rate $\rho = 0.05$, and a reflecting barrier

First, we treat the innovation rates as exogenous and estimate $\{x_H, x_L, x_E, \alpha, \kappa, F, \theta\}$ to match the moments between 1999 and 2001 (the "pre-period") shown in Table VIII. The innovation rates for high- and low-type plants x_H and x_L are jointly identified by the job creation rate by incumbents and the 75th percentile of the establishment age distribution. The incumbent job creation rate identifies total innovation by incumbents $x_H + x_L$ since greater innovation by all incumbents increases the rate they add products and, in turn, job creation by incumbent plants. The 75th percentile of the establishment age distribution pins down the ratio of x_H to x_L since a smaller gap between x_L and x_H implies that the two types of plants are more similar and the dispersion in age is smaller. Innovation by entrants x_E is identified by the share of total employment by entrants. The more innovative are entrants, the more product lines they hold and the higher their share of overall employment.

To identify the firing costs and contract labor adoption cost, we target the elasticity of value-added per wage-bill with respect to the establishment's wage-bill and the share of large plants using contract labor intensively. We use the wage-bill instead of employment to account for differences in labor quality across plants, including differences in labor quality between contract and permanent workers. Since the ratio of average revenue products of labor between plants employing and not employing contract labor is $1 + x\kappa$ in the model, the former identifies κ given a value of x pinned down by the innovation rates x_H, x_L, and x_E. The latter identifies the fixed cost of hiring contract workers since this parameter determines the extensive margin of how many plants adopt contract labor. The share of high-type plants among entrants α is pinned down by the difference in exit rates between young and old plants conditional on establishment size. Due to selection, the pool of older plants will contain more high-type plants, and so a higher α manifests itself through more high-type plants amongst old plants and a larger gap in exit rates. Lastly, the shape parameter of the distribution of innovation draws θ is set to match the growth rate conditional on x.[51]

The top panel in Table IX shows the values of $\{x_H, x_L, x_E, \alpha, \kappa, F, \theta\}$ that most closely match these moments. High-type plants are much more innovative than low-type plants (which conduct almost no innovation) and about 6.3 times as innovative as entrants. There are many more low-type plants than high-type plants in the economy, with only around 2.6% of entrants likely to be high-type. Firing costs κ are estimated to be around 6.9 times the wage rate, while the cost of adopting contract labor F is around 1.2 units of the final good. A shape parameter θ for the distribution of the quality draws of around 3 is needed to match the aggregate growth rate.

In the second step, we invert the expressions for optimal innovation decisions to recover the productivity parameters in the R&D sector ξ_H, ξ_L, and ξ_E from the innovation rates x_H, x_L, and x_E shown in the top panel of Table IX. The second panel in Table IX shows the resulting productivity parameters in R&D. The productivity of high-type plants in R&D is about 2.6 times higher than the productivity of entrants and 93 times larger than the productivity of low-type plants.

The third step is to estimate the change in fixed cost F necessary to explain the change in the share of large plants that use contract labor intensively between the pre- and the post-period (between 1999–2001 and 2013–2015). Specifically, in the post-period we assume

parameter $\psi = 0.02$. These numbers are standard; see, for example, Acemoglu et al. (2018) and the references therein.

[51]We minimize the weighted sum of squared percentage deviations of the model-generated moments from the data moments, where the weights on the TFP growth and intensive contract labor usage are five times the weights on other moments.

TABLE VIII

MOMENTS IN MODEL AND DATA.

Moments	Parameter Identified	Period	Data	Model
Rate of Job Creation by Incumbents	$x_H + x_L$	Pre	0.0434	0.0526
Share of Employments in Entrants	x_E	Pre	0.0456	0.0365
Difference in Exit Rates, Young vs. Old	α	Pre	0.0132	0.0134
VA/Wage-Bill vs. Wage-Bill Elasticity	κ	Pre	0.0444	0.0408
Pre-Period, % Large Plants With Intensive Contract Labor Use	F^{Pre}	Pre	0.218	0.217
Post-Period, % Large Plants With Intensive Contract Labor Use	F^{Post}	Post	0.358	0.360
75th Percentile of Establishment Age	x_H/x_L	Pre	21.00	15.95
TFP Growth	θ	Pre	1.0740	1.0695

Note: Table shows averages over 1999–2001 (pre-period) and 2013–2015 (post-period), unless otherwise noted, and includes plants with more than five workers. Job creation by incumbents is sum over incumbent plants with increasing employment in each one-year period divided by average of employment in the initial and final year. Share of employment of entrants in year t is employment of plants in year t that did not exist in year $t-1$ as a share of total employment in year t. Young plants are those with age < 10 and old plants are those with age > 10. Exit rate for young versus old plants is computed by regressing an indicator variable for exit over one-year dummies for young and old, establishment employment, and state-year and industry-year fixed effects. The VA/Wage-Bill versus Wage-Bill elasticity is computed by regressing log VA/Wage-Bill on log Wage-Bill. Intensive contract labor use defined as hiring more than 50% of workers through contractors. Large defined as more than 100 workers. TFP growth for manufacturing plants over 1993–2007, taken from Bollard, Klenow, and Sharma (2013).

the innovation parameters ξ_H, ξ_L, ξ_E remain fixed and choose the change in F that most closely matches the share of large plants that use contract labor intensively in 2013–2015.[52] The third panel in Table IX shows that to "explain" an increase from 22% to 36% in the share of labor plants that use contract labor intensively, the cost of adopting contract labor must have fallen from 1.21 to 0.65 (in units of the final good).

TABLE IX

ESTIMATES OF MODEL PARAMETERS.

Parameters	Description	Estimate
Pre-Period (Step 1)		
x_H	Innovation rate for high-type incumbents	0.1203
x_L	Innovation rate for low-type incumbents	0.0000
x_E	Innovation rate for entrants	0.0191
α	Proportion of high-type plants among entrants	0.0255
κ	Firing cost of permanent labor	6.9007
F	Fixed cost of adopting contract labor, pre-period	1.2119
θ	Shape parameter of Pareto distribution	2.9953
Pre-Period (Step 2)		
ξ_H	Innovation parameter for high-type incumbents	0.2516
ξ_L	Innovation parameter for low-type incumbents	0.0027
ξ_E	Innovation parameter for entrants	0.0954
Post-Period (Step 3)		
F	Fixed cost of adopting contract labor, post-period	0.6489

[52]See Supplemental Appendix Section G (Bertrand, Hsieh, and Tsivanidis (2025)) for the simulation algorithm. We also keep fixed α, κ, and θ.

TABLE X

SIMULATING THE IMPACT OF THE FALL IN CONTRACT LABOR ADOPTION COSTS.

	Model	Data
Δ ln Aggregate TFP	0.073	
Δ ln VA/Wage-Bill vs. Wage-Bill Elasticity	−0.817	−0.693
Δ Growth Rate Aggregate TFP	−0.052	
Δ ln Job Creation Rate, Large vs. Small	0.177	0.369
Δ ln TFP, Large vs. Small	0.165	0.212
Δ ln Prob Add or Drop Products, Large vs. Small	0.028	0.043
Δ ln Share of Employment in Entrants	−0.187	−0.118

Note: Column 1 shows the change in the moment predicted by the model in response to the change in the fixed cost of using contract labor between the pre-period (1999–2001) and the post-period (2013–2015). Column 2 shows the same moment in the data over the same time period. The top panel shows the change in log aggregate TFP and the elasticity of the average product of labor with the total wage-bill across firms. The bottom panel shows the change in the growth rate of aggregate TFP, three measures of innovation by large versus small firms (job creation rate, firm TFP, the fraction of firms that add or drop at least one product), and the employment share of entrants. Large (small) firms are defined by the top (bottom) 25% of firm employment in the model and data.

7.3. *Quantifying the Impacts of Contract Labor Growth*

How did the proliferation of contract labor reshape Indian manufacturing? The goal of the third step in Table IX is to answer this question. By holding all model parameters fixed at their values estimated to match moments in the pre-period (1999–2001), and varying only the fixed cost of contract labor adoption to match its increased use amongst large plants by 2013–2015, this counterfactual allows us to quantify what Indian manufacturing would look like had only this change occurred over the 15-year period.

We begin by showing the static gain in aggregate TFP. The first row in Table X shows that aggregate TFP rises by 0.073 log points in response to the decreased cost of contract labor adoption. As more large, high-type plants hire contract labor, the model predicts that the gap in the marginal product of labor between large and small firms becomes smaller. Specifically, the elasticity of the average product of labor with respect to establishment size falls by 0.82 log points in the model. This is close to the reduction of 0.69 log points observed in the data during the 2000s, as shown in the second column of Table X. So both in the model and in the data, the gap in the marginal product of labor between large and small firms fell, and this decline drives the gain in aggregate TFP.

We then turn to the dynamic effect of the reduction in the fixed cost of using contract labor. The model predicts that, in contrast to the effect on the level of aggregate TFP, the *growth rate* of aggregate TFP does *not* change. This is because of two offsetting effects on innovation when the fixed cost of contract labor adoption falls. On the one hand, high-type firms innovate more when it is easier to circumvent the IDA by using contract labor. On the other hand, the model predicts that innovation by entrants *declines* when the high-type firms use more contract labor. Specifically, when high-type plants innovate more, the value of low-type plants falls as their products are more likely to be stolen by high-type plants. As a consequence, innovation by low-type firms falls.

While we do not directly observe high- versus low-type firms in the data, viewed through the lens of our model, the majority of large incumbent firms are likely to be high-type firms, and the majority of small incumbent firms are likely to be low-type firms. Similarly, the majority of entrants are likely to be low-type firms. Also, while we cannot directly measure innovation in the data, we can measure proxies of innovation. Innovation by incumbent firms shows up in the job creation rate, firm TFP, and the probability of adding or dropping an additional product. For entrants, changes in their innovation rates will show up as changes in their employment share. We can therefore verify several of this model's predictions in the data.

Table X shows the predicted change in the model between the pre- and the post-periods in the gap in job creation rate, TFP, and the probability of adding or dropping products between large and small firms.[53] The second column reports these same changes in the data: all three proxies of innovation do in fact increase for large versus small firms. We note that none of these measures of innovation are targeted in the estimation of the model.[54] So both in the data and in the model, innovation by large firms appears to increase relative to small firms.

Finally, the last entry in Table X reports on the change in the employment share of entrants between the pre- and the post-periods as a proxy for the change in innovation by entrants. This moment is also not targeted in the model estimation. In the data, the employment share of entrants falls by 0.12 log points, compared to a 0.19 log points decline in the model. This evidence suggests that innovation by entrants, most of which are low-type firms, does in fact fall when large incumbent firms increasingly adopt contract labor.

8. IMPLICATIONS OF THE RISE OF THE STAFFING INDUSTRY FOR WORKERS

The model in the last section suggests that the rise in contract labor may have increased aggregate TFP, and thus, the average wage. In this section, we explore the effect of the increase in contract labor on the *distribution* of wages between workers.

We begin by showing wages associated with contract work compared to other employment arrangements. We use the CPHS sample, but this time expand the sample to include workers in informal plants as well as daily-wage workers and self-employed.[55] Supplemental Appendix Table E.15 (Bertrand, Hsieh, and Tsivanidis (2025)) shows regressions of log earnings on employment type in this sample in 2017, separately by education group (< 10, 10–12, and > 12 years of schooling). Across all education groups, all employment arrangements vastly dominate daily-wage/casual employment (the omitted category). Informal contract work is associated with the second-worst outcomes for workers.

Table E.16 in the Supplemental Appendix (Bertrand, Hsieh, and Tsivanidis (2025)) shows the share of workers in each employment status. The two worst employment arrangements in terms of worker's earnings, daily-wage employment, and informal sector contracting, are by far the most common in Indian manufacturing, with 23% of manufacturing workers in India being daily-wage workers and another 26% being temporary workers in the informal sector. Less-educated workers are particularly over-represented in these lower-paying employment arrangements, with 37% of workers with less than 10 years of schooling in daily-wage employment. Hence, the growth of the contracting model in formal manufacturing has the potential to add better opportunities to workers stuck in worse employment arrangements.

Table E.17 in the Supplemental Appendix (Bertrand, Hsieh, and Tsivanidis (2025)) exploits the panel characteristics of the CPHS to calculate year-to-year transitions between

---

[53]TFP is measured by Value-added$^{\sigma/(\sigma-1)}$/(wage-bill$^{(1-\alpha)}$capital$^{\alpha}$), where σ is the elasticity of substitution and α is the capital share.

[54]Large (small) firms are defined by the largest (smallest) 25% of firms measured by employment in the model and in the data.

[55]See Supplemental Appendix C (Bertrand, Hsieh, and Tsivanidis (2025)) for details.

employment categories. We broaden the CPHS sample to include workers outside of manufacturing as well as the non-employed.[56] The high transition rate into permanent employment for workers that are employed as formal contract workers is notable. Remember that these are the workers likely to be supplied by the staffing companies that expanded after the late 1990s. Conditional on formal contract work, the transition rate into permanent employment is 32%, which is the highest transition rate into permanent employment among all employment arrangements. This evidence suggests that an important "effect" of accessing formal contract employment is that it increases the probability of permanent work.

Motivated by this evidence, we sketch a model to estimate the equilibrium effect of the increased availability of contract labor on workers with different skill sets. We assume that a worker is characterized by their group g and labor market status s. Empirically, we measure labor market status s as permanent, daily-wage, self-employed, temporary with provident fund, and temporary without provident fund, and worker type g as years of schooling (< 10, 10–12, and > 12). In addition, to match the fact that workers switch between labor market status, we assume that labor market status is a choice variable. Intuitively, the lifetime utility from a specific labor market status is not simply the wage from being in that status but also its "effect" on the probability of transitioning into other labor market statuses.

Specifically, the utility of a worker of type g and labor market status s in the current period from switching to labor market r in a future period is $\frac{w_{rg}\epsilon_r}{d_{sr}}$, where w_{rg} is the wage, d_{sr} is the cost of switching from s to r, and ϵ_r is the worker's idiosyncratic preference for sector r.

We then "explain" transitions in labor market status by assuming that a worker gets a new draw of ϵ (for each of the labor market states) in each period. Assuming that ϵ follows a Fréchet distribution with shape parameter θ, the probability that a worker of type g in labor market status s switches to labor market status r is given by

$$P_{rsg} = \frac{\left(\dfrac{w_{rg}}{dsr}\right)^{\theta}}{\displaystyle\sum_k \left(\dfrac{w_{kg}}{d_{sk}}\right)^{\theta}}. \tag{1}$$

The steady-state share of workers of group g in each labor market status l_{rg} is then given by

$$\{l_{rg}\} = \{P_{rsg}\}'\{l_{rg}\}. \tag{2}$$

Average lifetime utility of a worker of group g is then given by[57]

$$V_g = \sum_s l_{sg}\left(\sum_k \left(\frac{w_{kg}}{d_{sk}}\right)^{\theta}\right)^{1/\theta}. \tag{3}$$

Average lifetime utility is a weighted average of the utility associated with each labor market state $(\sum_k (\frac{w_{kg}}{d_{sk}})^{\theta})^{1/\theta}$, where the weights are the steady-state share of workers in each state l_{rg}.

[56]We calculate transition rates between employment arrangement over a year in three waves of the CPHS data (May–Aug 17, Sep–Dec 17, Jan–Apr 18) and take the average transition rate across the three waves.

[57]We assume no discounting of future income for simplicity.

TABLE XI

LIFETIME UTILITY AND SELF-EMPLOYMENT SHARE BY SCHOOLING GROUPS.

	Lifetime Utility		% Self-Employed	
Schooling Group	1999	2017	1999	2017
< 10 Years	1	1.08	14.8%	10.0%
10–12 Years	1.21	1.32	14.0%	9.8%
> 12 Years	1.45	1.71	20.1%	16.0%

Note: We assume the only change in 1999 relative to 2017 is the wage from contract work. Columns 1 and 2 show the PDV of lifetime utility in 1999 and 2017 computed from equations (1), (2), and (3). Lifetime utility is normalized by utility in 1999 for workers with < 10 years of schooling. Columns 3 and 4 show the share of workers in each schooling group that are self-employed in 1999 and 2017 computed from equations (1) and (2).

We then estimate the effect of the expansion of contract labor between 1999 and 2017 on the lifetime utility of each group in three steps.

First, we use CPHS data to calculate the transition matrices P_{rsg} and the wage w_{rg} in 2017. These two data moments are calculated for the three schooling groups and for all individuals over age 15 in all sectors. Then, assuming the cost of staying in the same labor market status is zero ($d_{rr} = 1$), we use equation (1) to infer the transition cost d_{sr} from the ratio of the transition probability from r to s to the transition probability of staying in r.[58]

Second, we assume the wage of contract workers changes between 1999 and 2017 uniformly for all worker groups g. For this exercise, we assume that wages for all the other labor market states other than contract labor and the transition costs d_{rg} are unchanged at their 2017 levels. That is, $\Delta w_{rg} = \gamma_r$, where r are contract workers and $\gamma_s = 1$ for $s \neq r$. We then choose γ_r such that the change in the steady-state share of contract workers in the model is equal to the change observed in the data for formal manufacturing workers between 1999–2000 and 2017.[59]

Third, we impute the steady-state labor shares $\tilde{l}_{rg}$ implied by γ_g we calculated in the previous step using equations (1) and (2). We then calculate expected lifetime welfare in 1999 from equation (3), where l_{rg} is replaced by $\tilde{l}_{rg}$, the wage of contract workers is replaced by w_{rg}/γ_r, and d_{rg} and the wage in other sectors are held fixed.

Table XI shows the results of this calculation. The first column shows average lifetime utility of each schooling group in 1999 relative to workers with < 10 years of schooling. Not surprisingly, average utility is increasing with schooling. The second column shows average utility in 2017 after the expansion of contract labor. The largest increase in utility is for the most educated group (> 12 years), the next largest is for the least educated group (< 10 years), and the smallest increase is workers with 10–12 years of schooling.

We next show the model's estimate of the effect of the expansion of contract labor on the share of workers that choose self-employment (in columns 3 and 4). The self-employment share declines for all three schooling groups, roughly equally for all educational groups. Since it is likely that the vast majority of the self-employed are the owners of what the model calls "low-type" firms, this calculation suggests that the expansion of

[58]We also use CPHS data on the wage of each worker type by labor market status w_{rg}. We assume $\theta = 3$.

[59]The share of formal manufacturing workers increased by 13.8 percentage points between 1999–2000 and 2017, as calculated in the sample of formal workers in the ASI in 1990–2000 and the CPHS in 2017. The total share of workers with labor market status r in steady state is given by $l_r = \sum_g l_{rg}$, where l_{rg} is given by equations (1) and (2).

contract labor also resulted in a decline in the number of such firms. This prediction is consistent with the fact shown earlier (in Table X) that the employment share of entrants declined over the period when contract labor expanded.

9. CONCLUSION

We provide evidence that the employment restrictions on large Indian plants appear to have diminished since the early 2000s. We argue that this is driven by the expansion of formal staffing companies that provide contract workers primarily to large plants. The use of contract labor allows large Indian plants to respond to shocks to profitability, expand employment, and invest in new products. This shows up as an increase in the thickness of the right tail of the establishment size distribution in India and a decrease in the average product of labor of large Indian plants. Our quantitative exercise suggests that the increased use of contract labor can "explain" the declining gap in the average product of labor, which accounts for 7.3% of the increase in aggregate TFP over this period.

The quantitative exercise also suggests that entrants will *lower* innovation because of the increased competition from large incumbent plants. In the data, this is consistent with a decline in the employment share of entrants. The decline in innovation among entrants entirely offsets the increased innovation by incumbents, so there is no net increase in innovation.

Despite the decline in the average product of labor seen in the data since the early 2000s, it is still the case that the average product of labor in large Indian plants is substantially higher than that of smaller plants, and that Indian manufacturing is still dominated by a large number of small informal plants. This suggests that a greater reliance on contract labor is only a partial answer to the constraints faced by formal manufacturing in India. It is also possible that the reliance on contract labor can lead to longer-term problems for Indian manufacturing.

Finally, we made an attempt to quantify the distributional effects of contract employment, but our answer should be viewed as preliminary. For example, our estimation takes as given the transition matrices between the different employment arrangements, but it is possible that they are endogenous to the aggregate magnitude of contract employment.

REFERENCES

ACEMOGLU, DARON, UFUK AKCIGIT, HARUN ALP, NICHOLAS BLOOM, AND WILLIAM KERR (2018): "Innovation, Reallocation and Growth," *American Economic Review*, 108 (11), 3450–3491. [1415,1440]

ADHVARYU, ACHYUTA, A. V. CHARI, AND SIDDHARTH SHARMA (2013): "Firing Costs and Flexibility: Evidence From Firms' Employment Responses to Shocks in India," *Review of Economics and Statistics*, 95 (3), 725–740. [1431]

AGHION, PHILIPPE, ANTONIN BERGEAUD, AND JOHN VAN REENEN (2021): "The Impact of Regulation on Innovation," NBER Working Paper, No. 28381. [1415]

AGHION, PHILIPPE, ROBIN BURGESS, STEPHEN REDDING, AND FABRIZIO ZILIBOTTI (2008): "The Unequal Effects of Liberalization: Evidence From Dismantling the License Raj in India," *American Economic Review*, 98 (4), 1397–1412. [1415,1420]

AHSAN, AHMAD, AND CARMEN PAGES (2007): "Are All Labor Regulations Equal? Assessing the Effects of Job Security, Labor Dispute, and Contract Labor Laws in India," Policy Research Working Paper 4259. World Bank, Washington, DC. [1416]

AKCIGIT, UFUK, AND WILLIAM KERR (2018): "Growth Through Heterogeneous Innovations," *Journal of Political Economy*, 126 (4), 1374–1443. [1415]

BERTRAND, MARIANNE, CHANG-TAI HSIEH, AND NICK TSIVANIDIS (2025): "Supplement to 'Contract Labor and Establishment Growth in India'," *Econometrica Supplemental Material*, 93, https://doi.org/10.3982/ECTA20046. [1418,1420,1421,1423-1429,1434,1437,1441,1443]

BESLEY, TIMOTHY, AND ROBIN BURGESS (2004): "Can Labor Regulation Hinder Economic Performance? Evidence From India," *Quarterly Journal of Economics*, 119 (1), 91–134. [1415]

BOLLARD, ALBERT, PETER J. KLENOW, AND GUNJAN SHARMA (2013): "India's Mysterious Manufacturing Miracle," *Review of Economic Dynamics*, 16 (1), 59–85. [1441]

CHAUREY, RITAM (2015): "Labor Regulations and Contract Labor Use: Evidence From Indian Plants," *Journal of Development Economics*, 114 (C), 224–232. [1414,1415,1431]

COX, J. (2012): "With Eyes Wide Open: Recent Trends in Supreme Court Labour and Industrial Judgments," in *Contract Labour in India: Issues and Perspectives*, ed. by K. R. Shyam Sundar. New Delhi: Daanish Books. [1417]

DRENIK, ANDRES, SIMON JAGER, PASCUEL PLOTKIN, AND BENJAMIN SCHOEFER (2020): "Paying Outsourced Labor: Direct Evidence From Linked Temp Agency-Worker-Client Data," Working Paper. [1414]

DUTTA ROY, SUDIPTA (2004): "Employment Dynamics in Indian Industry: Adjustment Lags and the Impact of Job Security Regulations," *Journal of Development Economics*, 73 (1), 233–256. [1415]

FALLON, PETER R., AND ROBERT E. B. LUCAS (1993): "Job Security Regulations and the Dynamic Demand for Industrial Labour in India and Zimbabwe," *Journal of Development Economics*, 40, 241–275. [1415]

FELIX, MAYARA, AND MICHAEL WONG (2021): "Labor Market Consequences of Domestic Outsourcing: Evidence From Legalization in Brazil," Working paper. [1414]

GARICANO, LUIS, CLAIRE LELARGE, AND JOHN VAN REENEN (2016): "Firm Size Distortions and the Productivity Distribution: Evidence From France," *American Economic Review*, 106 (11), 3439–3479. [1415]

GOLDSCHMIDT, DEBORAH, AND JOHANNES SCHMIEDER (2017): "The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure," *Quarterly Journal of Economics*, 132 (3), 1165–1217. [1414,1418]

GONSALVES, COLIN (2011): *Kaliug: The Decline of Human Rights Law in the Period of Globalisation*. New Delhi, India: Human Rights Law Network. [1417]

GOPALAKRISHNAN, RAMAPRIYA, AND JEANNE MIRER (2013): "The Regulation of Non-Standard Forms of Employment in India, Indonesia and Vietnam," ILO's Conditions of Work and Employment Series, n. 63. [1417]

GOURIO, FRANCOIS, AND NICOLAS ROYS (2014): "Size Dependent Regulations, Firm Size Distribution, and Reallocation," *Quantitative Economics*, 5, 377–416. [1415]

GUNER, NEZIH, GUSTAVO VENTURA, AND YI XU (2008): "Macroeconomic Implications of Size-Dependent Policies," *Review of Economic Dynamics*, 11 (4), 721–744. [1415]

HASAN, RANA, DEVASHISH MITRA, AND K. V. RAMASWAMY (2007): "Technology Trade Reforms, Labor Regulations, and Labor-Demand Elasticities: Empirical Evidence From India," *Review of Economics and Statistics*, 89 (3), 466–481. [1415]

HOPENHAYN, HUGO, AND RICHARD ROGERSON (1993): "Job Turnover and Policy Evaluation: A General Equilibrium Analysis," *Journal of Political Economy*, 101 (5), 915–938. [1415]

HSIEH, CHANG-TAI, AND PETER J. KLENOW (2014): "The Life-Cycle of Manufacturing Plants in India and Mexico," *Quarterly Journal of Economics*, 129 (3), 89–108. [1412]

HSIEH, CHANG-TAI, AND BENJAMIN OLKEN (2014): "The Missing Missing Middle," *Journal of Economic Perspectives*, 28 (3), 1403–1448. [1412]

KAPOOR, RADHICKA, AND P. P. KRISHNAPRIYA (2017): "Informality in the Formal Sector: Evidence From Indian Manufacturing," IGC Working Paper. [1415]

KATZ, LAWRENCE F., AND ALAN B. KRUEGER (2019): "The Rise and Nature of Alternative Work Arrangements in the United States, 1995-2015," *ILR Review*, 72 (2), 382–416. [1414]

KLETTE, TOR, AND SAMUEL KORTUM (2004): "Innovating Plants and Aggregate Innovation," *Journal of Political Economy*, 112 (5), 986–1018. [1415,1421]

LANDAU, INGRID, PETRA MAHY, AND RICHARD MITCHELL (2015): "The Regulation of Non-Standard Forms of Employment in India, Indonesia and Viet Nam," Conditions of Work and Employment Series No. 63. International Labour Office. [1417]

MATSUURA, KENJI, AND CORT J. WILLMOTT (2012): "Terrestrial Precipitation: 1900-2010 Gridded Monthly Time Series," University of Delaware. [1420]

SAHA, BIBHAS, AND KUNAL SEN (2014): "What Explains the Increasing Use of Contract Workers in Indian Manufacturing?" https://www.ideasforindia.in/topics/governance/what-explains-the-increasing-use-of-contract-workers-in-indian-manufacturing.html. [1417]

SAHA, BIBHAS, KUNAL SEN, AND DIBYENDU MAITI (2013): "Trade Openness, Labour Institutions and Flexibilisation: Theory and Evidence From India," *Labour Economics*, 24, 180–195. [1415]

SANKARAN, KAMALA (2012): "Flexibility and Informalisation of Employment Relationships," in *Challenging the Legal Boundaries of Work Regulation*, ed. by J. Fudge, S. McCrystal, and K. Sankaran. London: Hart Publishing. [1417]

SINGH, JAIVIR, DEB KUSUM DAS, HOMAGNI CHOUDHURY, AND PRATEEK KUKREJA (2016): "Law, Skills and the Creation of Jobs as "Contract" Work in India: Exploring Survey Data to Make Inferences for Labour Law Reform," ICRIER Working Paper 330. [1418]

Co-editor Oriana Bandiera handled this manuscript.

Manuscript received 6 August, 2021; final version accepted 22 March, 2025; available online 27 March, 2025.

The replication package for this paper is available at https://doi.org/10.5281/zenodo.15021961. The Journal checked the data and codes included in the package for their ability to reproduce the results in the paper and approved online appendices.