

LABOR REGULATIONS AND THE COST OF CORRUPTION: EVIDENCE FROM THE INDIAN FIRM SIZE DISTRIBUTION

Amrit Amirapu and Michael Gechter*

Abstract—In this paper, we estimate the costs associated with an important suite of labor regulations in India by taking advantage of the fact that these regulations apply only to firms above a size threshold. Using distortions in the firm size distribution together with a structural model of firm size choice, we estimate that the regulations increase firms' unit labor costs by 35%. This estimate is robust to potential misreporting on the part of firms and enumerators. We also document a robust positive association between regulatory costs and exposure to corruption, which may explain why regulations appear to be so costly in developing countries.

I. Introduction

RESTRICTIVE labor regulations have been blamed for some of the most significant problems that developing countries face, including low labor force participation rates and low levels of employment in the formal sector (Besley & Burgess, 2004; Botero et al., 2004). It has even been suggested that regulations may distort the allocation of labor across firms, thus contributing to the substantially lower levels of aggregate productivity seen in developing countries (Hsieh & Klenow, 2009). What is not clear is why labor regulations should be so much costlier in a developing country setting, particularly since enforcement agencies there are typically characterized by severe resource constraints, low compliance, and widespread corruption (Svensson, 2005); (Chatterjee & Kanbur, 2013; Kanbur & Ronconi, 2015). Moreover, previous work on the subject in developing countries has focused on a small subset of labor regulations: namely, laws related to employment protection (e.g., firing restrictions) and minimum wages.¹ In actuality, labor regulations are multifaceted, encompassing many different types of employment-related laws, such as workplace safety requirements and the

provision of mandated benefits (including health insurance, social security legislation, and payment of gratuities).

In this paper, we address both of these gaps and make several further contributions to the growing literature on labor regulations in developing countries. In particular, we estimate the costs associated with a suite of labor regulations in India whose components include workplace safety regulations, social security taxes and business registration requirements.² What the regulations have in common is that they apply only to firms with ten or more workers, a feature we exploit to identify the magnitude of the costs they impose on firms. Because our methodology takes advantage of this objective feature of the laws, we do not need to rely for identification on inherently subjective assessments of differences in the text of the laws across regions, a criticism that has dogged some of the best-known work in the literature (see Besley & Burgess, 2004; Bhattacharjea, 2009, and Fagernas, 2010).

Instead, our methodology translates observed firm behavior in response to the ten-worker threshold into estimates of the increase in unit labor costs associated with these regulations. Because our estimates are derived from firm behavior in response to actual enforcement rather than from the text of the laws, we refer to our estimated labor cost increase as representing *de facto* regulatory costs in what follows. We find that these regulations effectively increase firms' unit labor costs by 35%, substantially distorting economic decisions relative to a counterfactual regime without these regulations. We also apply our method to India's most stringent, controversial piece of employment protection legislation, Chapter VB of the Industrial Disputes Act (IDA), which stipulates that any industrial establishment with more than 100 workers (in most states) must obtain prior permission from the state government before laying off workers or closing the establishment. In contrast to the substantial costs we uncover at the 10-worker threshold, we find only a small and statistically insignificant impact on unit labor costs from operating at or above the 100-worker threshold.

The next contribution of the paper is to provide suggestive evidence that the distortionary effect of regulations is associated with the quality of governance through the extent of corruption present in regulatory enforcement. We show

Received for publication February 10, 2017. Revision accepted for publication September 28, 2018. Editor: Rohini Pande.

*Amirapu: University of Kent; Gechter: Pennsylvania State University.

We thank Dilip Mookherjee, Arvind Subramanian, Kevin Lang, Hiroaki Kaido, Kehinde Ajayi, Sam Bazzi, and Marc Rysman for input, guidance and encouragement. We are grateful to the editor, Rohini Pande, and three anonymous referees for comments that substantially improved the paper. We also thank Sam Asher, Laveesh Bhandari, Aditya Bhattacharjea, Sharon Buteau, Urmila Chatterjee, Bibek Debroy, Andy Foster, François Gourio, Rana Hasan, Ravi Kanbur, Anirban Mitra, Devashish Mitra, Sandip Mitra, Paul Novosad, G. Sajeewan, Johannes Schmieder, Pronab Sen, Anushree Sinha, Sandip Sukhtankar, Shyam Sundar, Anil Suraj, John Van Reenen, and numerous seminar and conference participants for many useful comments. Representatives in the state governments of Bihar, Kerala, Maharashtra, Tamil Nadu, and West Bengal provided invaluable assistance in understanding the data collection process for the Economic Census. We gratefully acknowledge financial support from the Weiss Family Program Fund for Research in Development Economics. Liu Li, Stefan Winata, and Haiqing Zhao provided expert research assistance, supported by the BU MA-RA mentor program.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/rest_a_00837.

¹See Djankov and Ramalho (2009), Freeman (2010), and Nataraj et al. (2014) for excellent reviews of the literature, which reveal this focus, and Dougherty, Frisancho, and Krishna (2014) for a notable exception.

²Business registration requirements are generally considered separately from labor regulations. However, in our context, labor regulations intended to apply to all firms are much more likely to be enforced once enforcement agencies have records of a firm's existence obtained through registration. This view is consistent with recent research experimentally defraying the costs of registration (de Mel, McKenzie, & Woodruff, 2013; de Andrade, Bruhn, & McKenzie, 2014), which finds that informal firms behave as if registration imposes costs on them over and above the costs of registration alone.

that de facto costs are lower in states that reformed rules to constrain the power of inspectors and higher in states with greater levels of corruption. If extortionary corruption is a significant determinant of regulatory costs (as our results suggest), this may explain why regulations appear to be more costly in developing countries than in developed countries; it is not the regulations themselves that are particularly problematic but the way in which they are enforced.

We develop our argument as follows. We begin by exhibiting the Indian establishment size distribution using data from the Economic Census of India (EC). The EC aims to be a complete enumeration of all nonfarm establishments³ in India, and unlike all other Indian establishment-level data sets, it is not censored by size or restricted to include only the formal or informal sector. It is thus the only Indian data set that permits estimation of the complete establishment size distribution across all sizes and types of establishments. We find that a power law distribution fits the data well, except for a discontinuous and proportional decrease in the density of establishments with ten or more workers (see figures 1 and 2).

To understand and quantify the effect of the regulations on firm cost structure, we develop a simple model in which managers are endowed with heterogeneous productivities and must choose their optimal employment levels. Firms that report hiring more than a threshold number of workers face higher unit labor costs due to the presence of regulations and are thus smaller than they would be otherwise. Garicano, Lelarge, and Van Reenen (2016; henceforth, GLV) show that the magnitude of the increase in costs can be identified from characteristics of the distribution including, most importantly, the size of the downshift in the density above the threshold. Our model augments GLV to allow for the possibility of strategic misreporting. That is, managers may choose to deliberately misreport their employment levels at some cost, with the goal of avoiding some or all of the additional labor costs that apply to firms above the threshold size.⁴ Fitting the model's predicted size distribution to the one observed from the EC data, we generate an estimate of the additional labor costs that apply to firms above the ten-worker threshold that is robust to the possibility of strategic misreporting.

We show substantial heterogeneity in the magnitude of de facto regulatory costs along several dimensions, including state, industry and ownership type. In the first place, we find that privately owned establishments face the highest de facto regulatory costs, while government-owned establishments show no significant cost increase when employing ten or more workers. This supports our interpretation that the

downshift in the distribution starting at ten workers is indeed due to the regulations, since many regulations either do not apply to government-owned establishments or are less binding for such establishments in practice. We also document higher effective costs for businesses run by members of disadvantaged social groups (scheduled castes, scheduled tribes, and women), suggesting that regulatory enforcement is unequal and may be linked to the bargaining power of firm owners. Using variation across states and industries, we find a strong and robust positive correlation between our estimated regulatory costs and several different measures of corruption. The link between high regulatory costs and corruption may appear surprising if one thinks of corruption as collusive, "greasing the wheels" in a highly regulated economy by allowing firms to reduce their effective regulatory burden by bribing inspectors (Huntington, 1968). Our results are instead consistent with the concern that corrupt inspectors may overreport violations relative to honest inspectors in order to extract greater bribes, which we deem extortionary corruption.⁵

Our finding that the most contentious component of India's employment protection legislation, Chapter VB of the IDA, does not have a substantial effect on unit labor costs differs from much of the earlier academic work on the subject and belies the attention the IDA has received from academics (Besley & Burgess, 2004; Hasan, Mitra, & Ramaswamy, 2007; Aghion et al., 2008; Adhvaryu, Chari, & Sharma, 2013; Chaurey, 2015) and the business press (Bajaj, 2011; Ghosh, 2016) alike. We attribute this difference, first, to the fact that our methodology for estimating the impact of the legislation is very different from the strategies employed in previous papers. Most previous work identifies the effect of India's employment protection legislation based on differences in the growth of mean outcomes across states, which have been coded as initiating pro-worker or pro-employer reforms to the full IDA. The coding of states into these three groups (pro-employer, pro-worker, or neutral) has been the subject of controversy in the subsequent literature (Fagernas, 2010; Bhattacharjea, 2006, 2009), and an advantage of our identification strategy is that we can sidestep this controversy. Hsieh and Olken (2014), the only other paper in the literature to focus on establishment size distributions, report finding no visually striking change in the size distribution of Indian establishments at the 100-worker threshold, in accordance with our quantitative results. Another source of difference is our focus on Chapter VB rather than the full IDA. This is partially a question of the feasibility of applying our approach (Chapter VB is size based), but we also see focusing on Chapter VB as providing a proof of concept: if Chapter VB's complete restriction on firing is not very distortionary, it would be surprising if the law's more mild provisions are.

³The EC refers to these as "entrepreneurial units" and defines them as any unit "engaged in the production or distribution of goods or services other than for the sole purpose of own consumption." As is common in the literature, we occasionally refer to them as "firms" even though the unit of observation in the data is actually a factory or an establishment (only a minute proportion of establishments belong to multiestablishment firms).

⁴Note that "strategic misreporting" is distinct from the issue of corruption in the enforcement of labor regulations.

⁵See Banerjee (1994), Mookherjee (1997), Hindriks, Keen, and Muthoo (1999), Polinsky and Shavell (2001), and Mishra and Mookherjee (2013) for theoretical treatments. Empirically, Sequeira and Djankov (2014) and Asher and Novosad (2017) also provide evidence for the importance of extortionary corruption.

In addition to our contributions to the empirical literature on labor regulations, our extension of the GLV model to allow firms to strategically misreport their sizes should find applications in many other settings. Robustness to strategic misreporting in response to a size threshold is particularly crucial when working with survey data sets from developing countries because such data are typically self-reported. By contrast, in the high-income country administrative data used by studies such as GLV, it may be harder for establishment managers to misreport information when desired. We show that in the presence of strategic misreporting, a naive approach to estimating GLV's model can dramatically overestimate the increase in labor costs associated with a size-based regulation.

We identify firms' real responses under the theoretical restriction that the cost of misreporting be strictly convex in the degree of misreporting (in the online appendix, we show that the results are in fact robust to a range of other modeling assumptions). In explicitly modeling the decision to misreport, we show that while misreporting can be extensive near the threshold, the reported firm size distribution approaches the true distribution at large firm sizes, so one can minimize bias in the estimate of regulatory costs by focusing the estimation on large firm sizes and discarding the observations close to the threshold. In our case, if one fails to account for the possibility of misreporting, the estimated increase in per worker costs rises from 35% to 101%.⁶

Before closing this introduction, it is worth noting three significant limitations of our methodology. The first is that we cannot separately identify the costs of individual regulations. Our cost estimates refer to the costs associated with all of the regulations that become binding at the ten-worker threshold and are likely to also include effects of regulations at the twenty-worker threshold. The second limitation is that due to our misreporting framework, the only kinds of costs that we can capture are unit labor costs, because these are the only costs that result in a downshift of the log firm size distribution. Fixed costs affect the firm size distribution in a way that is similar to strategic misreporting: both lead to higher reported mass just below the threshold and lower mass just above the threshold. For this reason, their effects cannot be separately identified, so if the regulations we study have significant fixed cost components, our methodology will not detect these. Finally, because our methodology involves comparing firms below the size threshold with those far above the threshold, we lose the ability to compare firms of similar sizes (nine versus eleven), which are more likely to face similar production and demand conditions. For this reason and the fact that firm size is a choice variable, our methodology is quite unlike regression discontinuity studies and must in-

⁶The fact that misreporting is sufficiently large in magnitude to produce an almost threefold distortion in estimated costs speaks to low state capacity and serves as a cautionary tale for users of government statistics in such environments.

stead rely on assumptions regarding the distribution of firms and economic theory regarding firm behavior.

The rest of the paper is organized as follows. In section II, we provide an overview of the relevant institutional details regarding Indian labor and industrial regulations. Section III introduces the data. In section IV, we describe the theoretical model and empirical strategy. Section V provides the main results. In section VI, we interpret the findings and investigate the connection between our estimated costs and corruption. Section VII concludes.

II. Labor Regulations in India

Many labor regulations in India apply only to establishments that are larger than a certain threshold, where size is most often measured in terms of the number of workers in the establishment. There are several thresholds at which different labor regulations start to apply, but the two most prominent thresholds occur once an establishment employs at least 10 and at least 100 workers.⁷ In most states in India, establishments that employ more than 100 permanent workers (excluding contract workers) must abide by India's most controversial piece of employment protection legislation: Chapter VB of the IDA.⁸ Under this regulation, establishments over the threshold must be granted government permission before closing the establishment or laying off workers. It is the IDA, of which Chapter VB is a part, that has been the subject of most academic papers on labor regulations in India.

In contrast, the ten-worker threshold has received far less attention from academics, even though it is extremely important due to the large number of varied regulations that start to become binding at that threshold, as well as the fact that this threshold is most commonly associated with the formal/informal divide. The major regulations that start to apply once an establishment employs ten or more workers include the following: establishments must register with the government, meet various workplace safety requirements (under the Factories Act for manufacturing establishments that use power and the Building and Other Construction Workers' Act for construction-related establishments, for example), pay insurance and social security taxes (under the Employees' State Insurance Act), distribute gratuities (under the Payment of Gratuity Act), and must bear a greater administrative burden (under, for example, the Labor Laws Act).

In online appendix A, we provide a table that includes a comprehensive list of all central (i.e., federal) size-based

⁷There are other thresholds, such as at twenty workers (at which point establishments must contribute to the Employees' Provident Fund Organisation, which operates a pension program for formal sector workers) and at fifty workers (at which point, severance payment obligations increase under Chapter VA of the Industrial Disputes Act), but we do not separately analyze these thresholds because they are less contentious and do not appear to substantially distort the establishment size distribution.

⁸In 2005, the year to which our analysis applies, this threshold was 100 workers for all states except West Bengal, where the threshold was 50 workers.

labor regulations in India. For each law, we briefly describe the regulation as well as the nature of the size-based threshold. The table documents variation in the regulatory burden across industries and ownership type. Some regulations cover specific industries, while many others are explicitly universal in scope. We note that government establishments are explicitly included in some laws and explicitly excluded from others. Some important size-based laws (e.g., the Payment of Gratuity Act and the Payment of Bonus Act), which may apply to government establishments on paper, are not relevant in practice because gratuities and bonuses for government workers in establishments of all sizes are set by pay commissions and are far in excess of those required in these laws.

Other regulations are indirectly, though not explicitly, size based, because they reference laws with size-based aspects. For example, the Maternity Benefits Act applies only to establishments designated as “factories” under the Factories Act, which means it applies only to establishments with more than ten workers. Furthermore, there appears to be a salience effect associated with the ten-worker threshold: in interviews with small business owners in Chennai, several of them appeared to believe that certain regulations (such as the Provident Fund Act) apply once a business has ten workers, when in fact they did not.

In addition to, or in lieu of, the explicit costs associated with complying with the regulations, establishments with ten or more workers may be subject to implicit costs associated with increased interaction with labor inspectors, who have a large amount of discretion regarding the enforcement of administrative law (TeamLease Services, 2006) and may thus be able to extract bribes by tightening (or easing) the administrative burden firms face.

It has been argued that the ability to extract bribes is exacerbated by the antiquated or arbitrary nature of certain components of the laws (Debroy, 2013). TeamLease Services (2006) provides some telling examples: “Rules under the Factories Act, framed in 1948, provide for white washing of factories. Distemper won’t do. Earthen pots filled with water are required. Water coolers won’t suffice. Red-painted buckets filled with sand are required. Fire extinguishers won’t do.” The result of such rules is that almost all firms can be found guilty of some violation or another under the letter of the law even if they are in compliance with the spirit of the law. Firm owners who choose not to comply with such regulations face costs (fines and possible prison sentences) if discovered and convicted.

This kind of behavior has been referred to as “harassment bribery” (Basu, 2011). Anecdotal evidence of inspectors using the complexity, arbitrariness, and sheer amount of paperwork as a way to extract bribes is easy to come by. For example, we have included a selection of citizen reports from ipaidabribe.com in online appendix H, which demonstrates this kind of behavior.⁹ Interestingly, some of the reports sug-

gest that the size of the bribe paid is a direct linear function of the number of employees, which will be relevant for interpreting our results in section VI.

III. Data and the Size Distribution in India

A. Data

We use the Economic Census of India (EC) as our main data source to investigate the costs associated with the regulations described in the previous section. The EC is meant to be a complete enumeration of all formal and informal nonfarm business establishments in India at a given time. It contains a very large number of units: the 2005 wave, which we will principally use, has almost 42 million observations. It is the only Indian data set that represents the unconditional distribution of establishment size, which is essential for our analysis. Other data sets, such as the CMIE’s Prowess Database, the Annual Survey of Industries (ASI), and the National Sample Survey’s (NSS) Unorganized Manufacturing Surveys, cover only certain parts of the distribution and are thus unsuitable for our analysis. The ASI, for example, covers only establishments in the manufacturing sector that have registered with the government under the Factories Act. However, registration under this act is required only for establishments with ten or more workers if the unit uses power (twenty or more workers if the establishment uses no power). Therefore, the selection into the ASI varies discontinuously at precisely one of our points of interest.

The price to pay for uniform coverage and large sample size is that the EC does not contain very detailed information on each observation. For each establishment in the data, there is only information on a handful of variables, including the total number of workers usually working, the number of non-hired workers (such as family members working alongside the owner), registration status, four-digit NIC industry code, type of ownership (e.g., private, government), and source of funds for the establishment. There is no information on capital, output, or profits, and the data are cross-sectional.

We supplement our analysis with data from a variety of other sources. We get data on state- and industry-level corruption from Transparency International’s India Corruption Study 2005, the Reserve Bank of India (RBI), and the World Bank Enterprise Survey for India (2005). Data on state-level regulatory enforcement come from the Indian Labour Year Book.¹⁰ Other measures of state-level regulations come from Aghion et al. (2008) and Dougherty (2009).

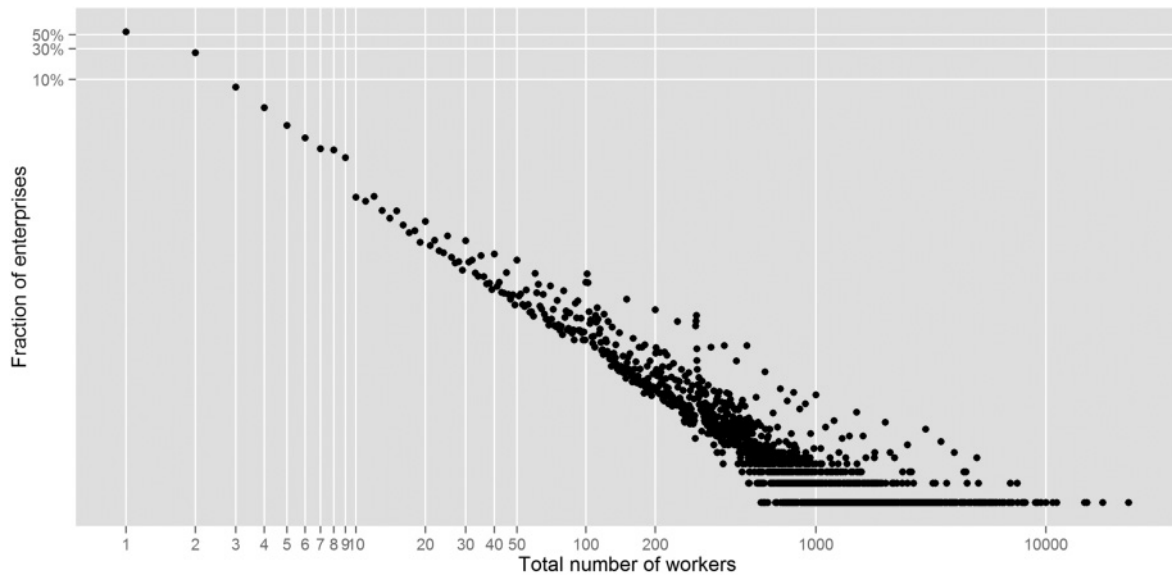
B. The Size Distribution of Establishments in India

Figure 1 shows the distribution of establishments by the number of total workers (hired plus nonhired) in 2005 on a log scale. Four things are striking about this figure. First,

¹⁰We thank Anushree Sinha and Avantika Prabhakar for their generous help in obtaining these data.

⁹We thank Andrew Foster for this suggestion.

FIGURE 1.—2005 LOG-LOG DISTRIBUTION OF ESTABLISHMENT SIZE



Both axes are on a log scale. The total number of workers is the number of workers usually working daily in an establishment.
Source: 2005 Economic Census of India.

the distribution is extraordinarily right-skewed. Indeed, about half of all establishments are single-person establishments. Second, the natural log of the density is a linear function of the natural log of the number of total workers. This implies that the unlogged distribution follows a power law in the number of total workers. This pattern will be important for the analysis that follows, but it is not very surprising in and of itself: power law distributions in firm sizes have been documented in many countries (Axtell, 2001, and Hernández-Pérez, Angulo-Brown, & Tun, 2006). Third, there appears to be a level shift downward in the log frequency for establishment sizes greater than or equal to 10. Finally, we do not see any discernible change in the distribution at 100 workers, the relevant threshold for employment protection legislation. We confirm this fact in our formal analysis.

Also apparent from the figures is a significant amount of nonclassical measurement error due to rounding of establishment sizes to multiples of five. The existence of rounding is not surprising given that the data are self-reported and that respondents are asked to give the “number of persons usually working [over the last year].” Our estimation procedure, described in the next section, accommodates this measurement error pattern.

IV. Model and Empirical Strategy

A. Modeling Size-Based Regulations with Strategic Misreporting

To interpret the downward shift from figure 1 in economic terms, we develop a model based on the framework from GLV but augmented to allow for the possibility that managers of

plants may strategically misreport their size to government officials—including labor inspectors and EC enumerators. For example, if plant managers are aware of the increased regulatory burden that is associated with employing ten or more workers and if they believe the EC enumerators will relay information to government regulatory bodies,¹¹ they may wish to hide the fact that their actual employment exceeds the threshold or more generally underreport their actual employment.

In the GLV framework, size-based regulations increase the unit labor costs of firms that exceed the size threshold, which results in a parallel downward shift in part of the theoretical logged firm size distribution. From the magnitude of the downshift observed in the empirical distribution, one can back out the additional labor costs imposed by the regulations. If firms are allowed to misreport their size, however, the reported firm size distribution may differ from the true distribution. In what follows, we show how a naive estimation procedure—which does not take misreporting into account—may result in biased estimates of the labor costs. We also present our solution, which minimizes the bias from misreporting.

The primitive object in our framework—following GLV as well as Lucas (1978), on which both our model and that of GLV are based—is the distribution of managerial ability ($\alpha \sim \phi : [\underline{\alpha}, \alpha_{\max}] \rightarrow \mathbb{R}$). Firms whose managers have higher ability (α) are more productive and can profitably employ more workers. Homogeneous workers are allocated to firms through a competitive labor market with a single

¹¹In fact, firms’ answers to EC enumerators have no impact on their regulatory burden, but it is quite possible that firms believe otherwise, and that is what is relevant.

market-clearing wage (w). As is common in the literature, we assume that the distribution of managerial ability follows a power law ($\phi(\alpha) = c_\alpha \alpha^{-\beta_\alpha}$), which then generates a power law in the theoretical firm size distribution. Our model differs from the basic GLV framework by allowing firms to choose not only their true employment (n) but also their reported employment (l).¹² Both are relevant when calculating expected costs due to the size-based regulations. In particular, a firm with productivity α faces the following profit-maximization problem,

$$\pi(\alpha) = \max_{n,l} \alpha f(n) - wn - \tau wl \times \mathbb{1}\{l > N\} - M(n, l), \tag{1}$$

where n is the number of workers a firm actually employs, l is the number of workers the firm reports to government officials (inspectors and enumerators alike), $f(n)$ is a production function (with $f'(n) > 0$ and $f''(n) < 0$), $\tau \geq 0$ is a proportional tax on labor that firms pay on their reported employment if their reported employment exceeds the regulatory threshold, and $M(n, l)$ is an expression that captures the expected costs of misreporting.¹³

The term capturing regulatory costs ($\tau wl \times \mathbb{1}\{l > N\}$) creates an incentive for firms to misreport their employment in a downward direction (i.e., to set $l < n$). Counteracting this incentive is that misreporting firms may be caught by the authorities and made to pay a fine. We think of the expected misreporting costs as being the product of three distinct terms: $M(n, l) = q(n) \times p(n, l) \times F(n, l)$. Firms are inspected with probability $q(n)$; conditional on being inspected, they are caught misreporting with probability $p(n, l)$, and if caught, they are made subject to a fine, $F(n, l)$. As written above, the probability of being inspected, the probability of being caught, and the magnitude of the fine may in general depend on n or l in an arbitrary way. Going forward, we will make the following assumptions regarding the general structure of $M(n, l)$, which enable us to identify τ in the presence of misreporting under minimal additional parametric assumptions. As we discuss further below, these assumptions are sufficient but not necessary to identify τ .

Assumption 1. Let $u \equiv n - l \geq 0$ denote the degree of misreporting and $M(u)$ denote the expected costs of misreporting. We assume that $M(0) = 0$ and that $M(u)$ is a continuous, increasing, and strictly convex function of u alone. Conditional on u , $M(u)$ is thus independent of firm size, n .

¹²A summary of the basic GLV framework (i.e., without misreporting) is provided in online appendix B.1. For a detailed derivation of the model, see Garicano et al. (2016).

¹³Note that labor demand will be lower in a regime with these regulations than without them. Therefore, the regulations will have a general equilibrium effect on employment and output through the wage, w . However, this will not affect our estimation of τ , our object of interest, because τ measures the increase in unit labor costs for larger firms as a proportion of the wage, which is common to all firms.

Under this assumption, we show that for large enough values of firm size, x , the difference between the log of the reported density, $\psi(x)$, and the log of the true density, $\chi(x)$, becomes vanishingly small.

Proposition 1. Suppose a firm’s profit maximization problem takes the form of equation (1) and assumption 1 holds. Then

$$\lim_{x \rightarrow \infty} \log \chi(x) - \log \psi(x) = 0.$$

Proof. See online appendix B.3.

Proposition 1 implies that an estimation based on large enough firm sizes will be minimally biased because the reported distribution becomes arbitrarily close to the true distribution at large sizes. We will demonstrate how this is accomplished with a specific example, but some discussion of the assumptions is necessary at this point.

The assumption that misreporting costs should be strictly convex in the degree of misreporting is both standard in the literature (Almunia & Lopez-Rodriguez, 2018; Kumler, Verhoogen, & Frias, 2015) and intuitive given our understanding of the context in which Indian businesses make such decisions.¹⁴ One important implication of assumption 1—that the extent of misreporting should be relatively lower for larger firms—finds empirical support in recent literature (Kouamé & Goyette, 2018). The other substantive assumption imposed above—that $M(u)$ is independent of firm size, n —is restrictive, but in fact neither it nor the convexity assumption is necessary; both are primarily useful in illustrating how identification proceeds without making parametric assumptions on the exact functional form of misreporting.

In online appendix B.4, we consider a range of alternative specifications for the functional form of misreporting, including several that depart from both of the assumptions above. What this exploration reveals is that most reasonable specifications either yield the same conclusions as proposition 1 or are incompatible with the observed data. The only specifications that are both consistent with the data and for which it is impossible to correctly identify τ are those that cause all firms to misreport a constant fraction of their true employment.

We now proceed by informally characterizing the solution to the firm’s problem, from equation (1), under assumption 1, for firms at every level of productivity. The lowest-productivity firms (those with α below some threshold, α_1) will be effectively unconstrained, in the sense that they choose to hire at most N workers ($n \leq N$) and thus do not fall under the purview of the size-based labor regulations. There

¹⁴The intuition is that hiding larger and larger numbers of employees from enumerators or inspectors gets increasingly difficult until at some point it is impossible. During our interviews with small firms, it was common to hear accounts of business owners ushering employees out the back door of the establishment whenever labor inspectors arrived, but this type of behavior is clearly possible for only relatively small numbers of employees. We thank Sharon Buteau and Balasekhar Sudalaimani from IFMR for helping to set up these interviews.

is no incentive for them to misreport, so they report truthfully ($l = n$). A second set of firms with higher productivity ($\alpha \in [\alpha_1, \alpha_2]$) find it optimal to exceed the regulatory threshold in practice (choosing $n > N$) but misreport their employment to avoid the higher regulatory costs (setting $l = N$). These firms only appear to be bunched up at N but in fact have higher employment.

The last category of firms are those with $\alpha > \alpha_2$, which are productive enough to warrant hiring workforces so large that they cannot completely avoid the regulation without being detected or fined with sufficiently high probability or severity, and thus report $l > N$. Even these firms, however, with both $n > N$ and $l > N$, do not find it profit maximizing to report truthfully. They can save on their unit labor costs by shading their reported employment and choose $l = n - M'^{-1}(\tau w)$. Note that the degree of misreporting is by a constant amount, which is a direct implication of assumption 1, as spelled out in online appendix B.3. Importantly, this last set of firms faces higher unit labor costs than in the absence of the regulations and therefore employ fewer workers by a constant proportion, resulting in a “downshift” in the logged firm size distribution. The fact that the degree of misreporting is by a constant amount implies that the difference between the true and reported distributions goes to 0 with size, and thus the downshift in the reported distribution will match that of the true distribution at large sizes.

To derive a closed-form solution for the true and reported firm size distributions, it is necessary to make some functional form assumptions. Doing so will clarify how we estimate τ and explore the implications of proposition 1 for our estimation procedure. The first parametric assumption we make, following GLV, is that firm output is a power function of labor: $f(n) = n^\theta$. The second is to impose a specific functional form for misreporting that satisfies assumption 1: $M(n, l) = \frac{F}{n_{\max}}(n - l)^2$.¹⁵ One way to generate this function is to suppose that the probability of inspection is proportional to firm size—for example, $q(n) = \frac{n}{n_{\max}}$; that the probability of being caught—conditional on being inspected—is proportional to the fraction of employees who are misreported ($p(n, l) = \frac{n-l}{n}$); and that the fine for those caught is proportional to the level of misreporting (i.e., $F(n, l) = F * (n - l)$). With these two substitutions, the firm’s profit maximization problem from equation (1) becomes:

$$\pi(\alpha) = \max_{n,l} \alpha n^\theta - wn - \tau wl * \mathbb{1}\{l > N\} - \frac{F}{n_{\max}}(n - l)^2.$$

As in our informal characterization the optimal choices of n and l will depend on the productivity of the firm. A full mapping between productivity α and the true firm size n , as well as between α and reported firm size l , is given by the

following equations:

$$n^*(\alpha) = \begin{cases} \left(\frac{\theta}{w}\right)^{\frac{1}{1-\theta}} (\alpha)^{\frac{1}{1-\theta}} \leq N & \text{if } \alpha \in [\underline{\alpha}, \alpha_1] \\ n_2^*(\alpha) & \text{if } \alpha \in (\alpha_1, \alpha_2], \\ \left(\frac{\theta}{w}\right)^{\frac{1}{1-\theta}} (1 + \tau)^{-\frac{1}{1-\theta}} (\alpha)^{\frac{1}{1-\theta}} > N & \text{if } \alpha > \alpha_2 \end{cases}$$

$$l^*(\alpha) = \begin{cases} \left(\frac{\theta}{w}\right)^{\frac{1}{1-\theta}} (\alpha)^{\frac{1}{1-\theta}} \leq N & \text{if } \alpha \in [\underline{\alpha}, \alpha_1] \\ N & \text{if } \alpha \in (\alpha_1, \alpha_2] \\ \left(\frac{\theta}{w}\right)^{\frac{1}{1-\theta}} (1 + \tau)^{-\frac{1}{1-\theta}} (\alpha)^{\frac{1}{1-\theta}} & \text{if } \alpha > \alpha_2. \\ -\frac{n_{\max}}{2F} w \tau > N \end{cases}$$

Because there is a strictly monotonic relationship between α and n , as well as α and l (except for the bunching), one can obtain expressions for the distributions of true and reported firm size, $\chi(n)$ and $\psi(l)$ as transformations of the distribution of managerial ability, $\phi(\alpha)$. Simplifying terms, one can write the log of the density of firms with true employment n as

$$\log \chi(n) = \begin{cases} \log A - \beta \log(n) & \text{if } n \in [n_{\min}, N) \\ \log[\xi(n)] & \text{if } n \in [N, n_m(\alpha_2)] \\ - & \text{if } n \in (n_m(\alpha_2), n_l(\alpha_2)) \\ \log A - \frac{\beta - 1}{1 - \theta} \log(1 + \tau) & \text{if } n \geq n_l(\alpha_2) \\ - \beta \log(n) & \end{cases} \quad (2)$$

and the log of the density of firms with reported employment l as

$$\log \psi(l) = \begin{cases} \log A - \beta \log(l) & \text{if } l \in [l_{\min}, N) \\ \log(\delta_l) & \text{if } l = N \\ - & \text{if } l \in (N, l_l(\alpha_2)) \\ \log A - \frac{\beta - 1}{1 - \theta} \log(1 + \tau) & \text{if } l \geq l_l(\alpha_2) \\ - \beta \log\left(l + \frac{n_{\max}}{2F} w \tau\right) & \end{cases}$$

where A is a function of constants and terms have been simplified and collected. Online appendix B.2 provides a derivation of this result, along with all missing steps.

Comparing the expressions for the reported and true size distributions above, there are several points worth noting.

¹⁵Again, we point out that this assumption is not necessary for identifying τ . The assumption that production is a power function, however, does have implications for identification, in the sense that our estimate of τ will depend on our estimate of θ .

First, For the range $l < N$, the true distribution coincides with the reported or observed distribution. Second, there appears to be bunching at N in the reported distribution, but some of these firms in fact have more than N workers. Third, compared to the distribution for $n < N$, both the true distribution and the reported distribution for $n \gg N$ are downshifted, and by exactly the same function of τ as in GLV's model without misreporting (the intercepts for both distributions are $\log A - \frac{\beta-1}{1-\theta} \log(1 + \tau)$ for larger firms versus $\log A$ for smaller firms).¹⁶ Fourth, as stated in proposition 1, the difference between the log of the reported distribution and the log of the true distribution converges to 0 for large firms. The intuition is straightforward: the only difference in the two expressions is the constant amount $\frac{n_{\max}}{2F} w\tau$, the contribution of which becomes negligible at large sizes.

Together, these observations allow us to back out an estimate of τ —the extra unit labor costs faced by firms above the size threshold. In particular, the first and fourth observations tell us that if we focus on firms below the size threshold and those well above it, the reported density will be arbitrarily close to the true density. The third observation—that a function of the tax enters additively in the log density for all firms above the threshold—tells us that τ can be determined from the size of the downshift observed in the log firm size distribution between large and small firm sizes. Given these observations, our identification strategy is quite simple. For very small and very large firm sizes ($n < N$ or $n \gg N$), one can express the log of the density according to the following equation:

$$\log(\chi(n)) = \log \left[\left(\frac{1-\theta}{\theta} \right)^{1-\beta} (\beta-1) \right] - \beta \log(n) + \log((1 + \tau)^{-\frac{\beta-1}{1-\theta}}) \mathbb{1}\{n > N\}, \quad (3)$$

where $\mathbb{1}\{\cdot\}$ is the indicator function. To see how τ is identified from $\chi(n)$, rewrite equation (3) as

$$\log(\chi(n)) = \alpha - \beta \log(n) + \delta \mathbb{1}\{n > N\}. \quad (4)$$

α , β , and δ can be identified by applying equation (4) to the observed size distribution. θ is a function of α and β and is thus also identified. τ is given by

$$\tau = \exp(\delta)^{-\frac{1-\theta}{\beta-1}} - 1,$$

which is identified as long as θ and β are identified.

In principle, by choosing some threshold n_L satisfying $n_L \gg N$, one should be able to produce a value for τ by using ordinary least squares to estimate the specification,

$$\log(\chi(n)) = \alpha - \beta \log(n) + \delta \mathbb{1}\{n > N\} + \epsilon(n), \quad (5)$$

¹⁶Online appendix B.1 includes the firm size distribution from GLV for comparison.

where $\epsilon(n)$ represents any deviation of the observed firm size distribution from the model coming from an idiosyncratic tendency for firms to cluster to or away from a particular size.

In practice, however, proposition 1 is problematic for estimating the parameters of the model from raw data using equation (5), because that equation must be estimated using data from relatively small establishments with sizes outside the range $[N, n_L]$. This is due to the fact that the empirical probability of observing an establishment of a given size is truncated at $\frac{1}{\# \text{ of observations}}$ (this is visually apparent in figure 1). Truncation makes the relationship between the log of the empirical probability and the log of the total number of workers nonlinear. To preserve the linear relationship, a researcher would have to omit establishment sizes large enough that truncation is not an issue.¹⁷ However, proposition 1 tells us that the misreported distribution is close to the true distribution only at large sizes and that the misreported distribution may be biased downward at establishment sizes close to a regulatory threshold. This leads to downward bias in δ (the downshift in the log density) and, consequently, upward bias in τ . Instead, we develop an empirical approach that deals with the truncation problem and allows us to focus on large firms, where the difference between the log of the reported distribution and the log of the true distribution is close to 0.

Since our approach involves estimating parameters of the firm's problem by fitting features of the theoretical density to the observed empirical density, it is worth noting the following discrepancy between the model and the data. The log density of reported employment that is generated by the model is undefined for $l \in (N, l_t(\alpha_2))$ because the density of reported employment contains a hole in this region (see equation [2]). Reporting $l \in (N, l_t(\alpha_2))$ is dominated by choosing either $l = N$ or $l \geq l_t(\alpha_2)$. However, figure 1 clearly shows that there are firms that report employing eleven workers.¹⁸ Based on interviews with firms and accountants, our understanding of the discrepancy is that small firms tend to be inattentive to the regulatory threshold, while large firms tend to be attentive. Attentive firms are aware of the regulations as well as the expected costs and benefits of misreporting, while inattentive firms are simply not aware of the relevant regulations—and hence do not bother to misreport their firm size. In online appendix B.6, we present a version of our model from the previous section that combines inattention and misreporting. In particular, we show there that if the fraction of inattentive firms is large at small firm sizes and small at large firm sizes,¹⁹

¹⁷The suggestion to focus on relatively smaller establishments appears in appendix B of Garicano et al. (2013).

¹⁸This type of discrepancy, in which many agents are observed to make strictly dominated choices, is common in analyses of behavior in response to “notches” or “kinks” (Kleven & Waseem, 2013).

¹⁹This assumption can be motivated theoretically if one imagines that managers must pay a fixed cost (which varies idiosyncratically across firms) in order to learn regulatory details—including the location of the thresholds. In practice, this would involve hiring an accountant or attorney who is knowledgeable about labor regulations. Under the plausible assumption

the model's predicted density will closely resemble the observed density, and τ remains identified using the method described here (in particular, by focusing primarily on firms with employment levels far above N).

Before proceeding, it is worth noting a second possible source of misreporting: EC enumerators themselves. EC enumerators were required to fill out an extra form containing the address of any establishment that reported ten or more workers. It is conceivable that enumerators might have found it preferable to underreport the number of workers for establishments with ten or more workers in order to avoid the extra burden of filling in the "Address Slip." However, as we show in online appendix B.5, this type of misreporting, like the previous one, generates bias only in the reported distribution for establishment sizes close to N . In particular, such enumerator-driven misreporting is likely to contribute to the "bunching" at ten and the "valley" just after ten, but it cannot lead to a downshift in the firm size distribution at large firm sizes, which is how we identify τ . Moreover, it is easy to show that any estimation technique that is robust to the possibility of manager-driven misreporting will also be robust to the possibility of enumerator-driven misreporting.

B. An Empirical Approach Robust to Strategic Misreporting

In this section, we develop a way of estimating equation (5) using establishments that are least affected by the misreporting. These include establishments that are below the bunching point, as well as those that are far above the size threshold. As we noted in section IVA, we cannot estimate equation (5) directly on large establishments because of truncation in the empirical probability of observing an establishment of a given size. Furthermore, as discussed in section III, the empirical size distribution is characterized by substantial rounding to multiples of five workers, especially at larger sizes. Setting aside the truncation problem, OLS estimation of equation (5) will produce downward bias in δ because sizes that are multiples of five are treated as single observations. Instead, their excess establishments should be distributed to nearby sizes.

To address both issues, we nonparametrically estimate the density associated with larger sizes using the method described in Markovitch and Krieger (2000; hereafter, MK). MK propose a nonparametric density estimator for heavy-tailed distributions that achieves L_1 consistency. L_1 consistency fails for any distribution with heavier tails than an exponential for the standard Parzen-Rosenblatt kernel density estimator,

$$\hat{f}(l) = \frac{1}{Eh} \sum_{i=1}^E K\left(\frac{l - L_i}{h}\right), \quad (6)$$

that the distribution of fixed costs does not vary with firm size, the fact that the benefits of adjusting employment in response to the threshold rise with size implies that all large firms will adjust, while only some small firms will.

where L_i , for our purposes, is establishment i 's total number of workers, l is a number of workers for which we would like to know the density, E is the total number of establishments in the 2005 EC, $K(\cdot)$ is a kernel function, and h a smoothing parameter or "bandwidth." L_1 consistency is known to hold for distributions with compact support, so MK suggest the simple approach of estimating the density of a transformation of L_i which has compact support, then inverting back for an estimate of the density of the original l .

Specifically, we first apply the transformation recommended by MK, $T(l) = \frac{2}{\pi} \arctan(l)$, to each establishment's number of workers. Our estimate of the density associated with a specific number of workers, l , is given by

$$\hat{\psi}(l) = \hat{f}(T(l))T'(l),$$

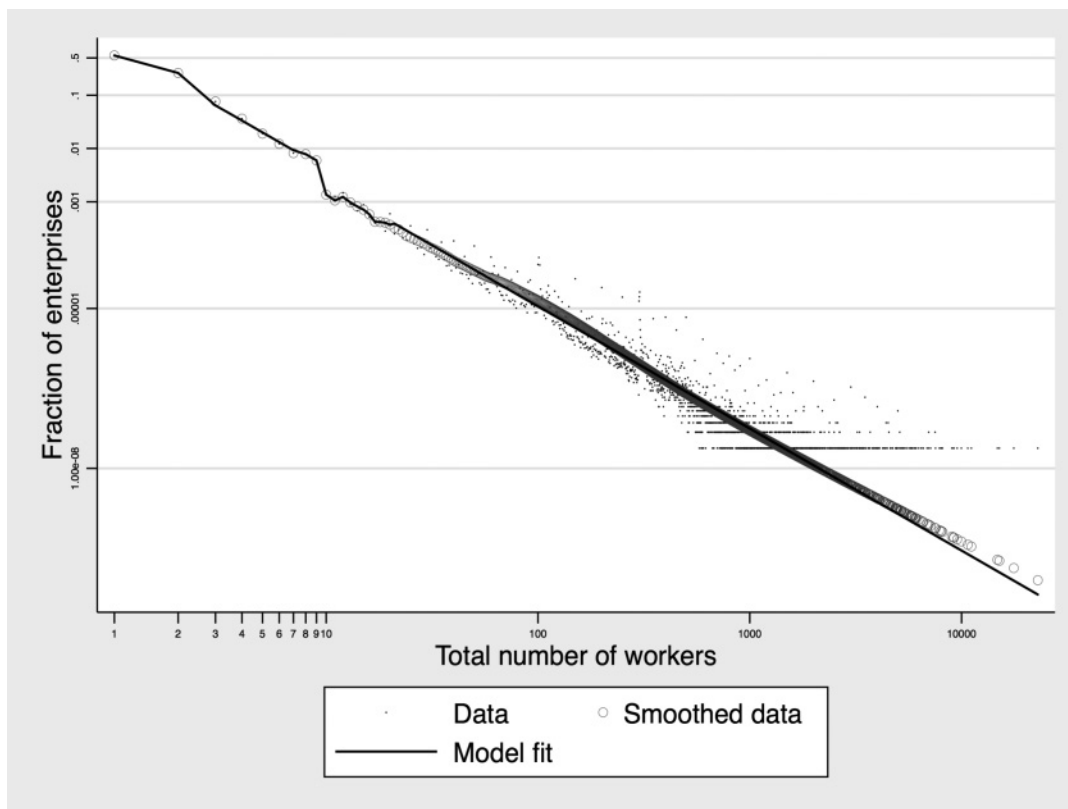
where $\hat{f}(T(l))$ applies equation (6) to the transformed data, $T(L_i)$, and evaluates at the transformed number of workers of interest $T(l)$. $T'(l)$ is the derivative of the transformation evaluated at l . We use the Epanechnikov kernel function. An advantage of this approach from our perspective is that a constant bandwidth applied to the transformed data expands asymmetrically with respect to the original data.²⁰ As we move to the right in the distribution, where data are more scarce, our kernel begins to put positive weight on observations further away. This accords with our observation that rounding in the reported distribution becomes more severe at larger sizes. We use the empirical probability for small sizes, where the establishment size distribution is better represented as a discrete variable.

We apply a modified version of equation (5) to the log of the estimated density $\hat{\psi}(l)$ for all observed sizes. For example, when analyzing the effect of regulations that apply to firms with ten or more workers, we remove the effect of misreporting close to the threshold by adding dummy variables for size 8 and 9 and for sizes 10 to 20. The choice of 20 as the largest size for which we include a dummy is unimportant and we show in table 5 in the online appendix that our results are robust to alternative choices of dummy variables. Since equation (5) treats each establishment size as one observation and since the range of establishment sizes in the 2005 EC runs from 1 to 22,901, the model is primarily estimated using data far from the 10-worker cutoff.²¹ Finally, we include dummies for having one or two workers because own account

²⁰Note that in this case, the bandwidth must be chosen. We cannot use cross-validation to choose the optimal bandwidth because it will recover the rounding pattern found in the data.

²¹We note here that our misreporting-robust strategy for estimating δ from equation (5) bears a resemblance to the estimator for the average treatment effect in "doughnut" regression discontinuity designs (Almond & Doyle, 2011; Barreca, Lindo, & Waddell, 2016; Eggers et al., 2018). However, the fact that our model estimates the log of the density associated with a firm size, rather than a regression function, generates a key difference in interpretation. Log density predictions generated by extrapolating the log density from the left of our doughnut to the right do not represent the counterfactual density with no regulation at the ten-worker threshold. Because of the requirement that the observed and counterfactual densities integrate to one, the density to the left of our doughnut would be reduced without regulations.

FIGURE 2.—MODEL FIT AND DATA



This figure shows the fit of the model described in section IVB (the black line) to the data (dark gray points). Model estimation involves nonparametric smoothing following Markovitch & Krieger (2000) with a bandwidth of 0.005 as a first step. The second step is to fit the black line to the light gray circles. Both axes in log scale. Source: 2005 Economic Census of India.

and two-worker establishments are likely to be household enterprises and may therefore differ fundamentally in character from their larger counterparts.²²

Figure 2 depicts the strategy. The dark gray dots show the raw data. The light gray circles represent the result of the first step: nonparametric density estimates associated with each establishment size. The line shows the fit of the model in equation (5), augmented by the dummy variables, to the nonparametric density estimates. Figure 2 also provides some evidence for the model described in section IVA. The observed establishment size distribution appears to converge back to a power law with the same slope as for establishments with fewer than ten workers but deviates slightly from that slope at sizes just above the ten-worker cutoff. In the next section, we report the results of the estimation.

V. Results

A. Regulations Applying to Firms Employing 10 or More Workers

Table 1 reports estimates for the increase in perworker costs associated with the increased regulatory burden of crossing

²²Table (5) shows that our results are robust excluding the two-worker dummy.

TABLE 1.—ESTIMATES OF τ AT THE TEN-WORKER THRESHOLD

Level	τ	SE
All-India	0.347	(0.081)
By state		
Bihar	0.693	(0.302)
Gujarat	0.165	(0.151)
Kerala	0.138	(0.196)
Karnataka	0.520	(0.156)
Uttar Pradesh	0.502	(0.254)
By industry		
Wholesale and retail trade	0.637	(0.094)
Manufacturing	0.268	(0.085)
Construction	0.478	(0.549)
Electricity, gas, and water	-0.367	(0.145)
By ownership type		
Government and PSU	-0.092	(0.128)
Unincorporated proprietary	0.430	(0.059)

This table presents estimates of regulatory costs faced by establishments with ten or more workers, using the methodology described in section IV with a bandwidth of 0.005. Standard errors were generated using a clustered bootstrap procedure with 200 replications. Clustering is done at the four-digit (NIC) industry level, following Garicano et al. (2016). Estimates are presented for a subset of states, industries, and ownership types. Results for all states, industries and ownership types are available in the online appendix C.

Source: 2005 Economic Census of India.

this threshold, τ , at the all-India level and for a selection of states, industries, and ownership types. Estimates for all states, industries, and ownership types are reported in online appendix C. Standard errors, displayed beside the point estimates in parentheses, are obtained from a clustered bootstrap procedure with 200 replications. Following GLV, we

cluster by industry at the four-digit NIC code level. This allows for the possibility that differences in production technology, which could affect the firm size distribution and therefore our estimates, may be correlated by industry.²³ The top panel of table 1 gives the all-India estimate of τ using our methodology. The point estimate is .35 and is significant at the < 1% level. This means that on average, establishments in India that employ more than nine workers act as though they must pay additional labor costs of 35% of the wage per additional worker.

By contrast, estimating the model without accounting for misreporting in any way yields much larger estimates. In particular, estimating equation (5) on the size distribution omitting sizes larger than 99 workers and including the same dummy variables as in our own specification would lead us to conclude that exceeding the 10-worker threshold increases per worker costs by 101%. This is due to a combination of rounding and the fact that the density associated with establishment sizes 21 to 99 converges only slowly back to the downshifted power law it follows at larger sizes, as predicted in our misreporting model. In other words, a naive estimation puts undue weight on firm sizes whose densities are biased downward by misreporting. In what follows, we focus our discussion on our misreporting-robust estimates of τ .

The lower panels of table 1 show substantial variation in the magnitude of our misreporting-robust estimates of the per worker tax by state, industry, and ownership type. For example, the point estimate on τ for the state of Kerala is .14 and is not statistically significant, while the estimate for Bihar is .70 and is statistically significant at the 5% level, implying that establishments in Bihar act as though they must pay a tax of 70% of the wage for each additional worker they employ past nine workers.

For industries, we see that de facto regulatory costs are high for establishments in manufacturing, construction, and retail and wholesale trade. Some industries have very noisy estimates, at times producing negative point estimates for τ . This is also true of some of the smaller states and ownership categories (as one can see in online appendix C) and is explained by the fact that the power law relationship can break down when there are a small number of observations in a category, as is the case for electricity, gas, and water. In a few cases, negative point estimates reflect the fact that the production and market characteristics of these industries can vary greatly from our model so that it provides a poor fit of the data.²⁴

When looking at the differences by ownership type, we find that the estimates for τ are highest for private firms (particularly unincorporated proprietorships, which form by

²³For robustness we have also tried alternative procedures, including a wild bootstrap and nonparametric bootstrap—both clustered at the firm size level.

²⁴Andhra Pradesh, the largest state to show a negative point estimate for τ , has a size distribution distorted in ways that are different from all other states which produces a poor fit. We have concluded that this is the result of errors in data collection or recording rather than deliberate misreporting.

TABLE 2.—ESTIMATES OF τ AT THE TEN-WORKER THRESHOLD BY OWNER'S SOCIAL GROUP OF OWNER

Level	τ	SE
By gender of owner		
Male	0.424	(0.060)
Female	0.525	(0.207)
By social group of owner		
Scheduled Tribe	1.016	(0.335)
Scheduled Caste	0.890	(0.233)
Other backward caste	0.425	(0.086)
Other	0.326	(0.059)

This table presents estimates of regulatory costs faced by establishments that employ ten or more workers, using the methodology described in section IV with a bandwidth of 0.005. Standard errors are calculated as in table 1.

Source: 2005 Economic Census of India.

far the largest category of private firms) and insignificant for government-owned firms. This is to be expected, since the regulatory burden does not vary as much across the ten-worker threshold for government establishments (see section II), and inspectors are less likely to engage in extortionary corruption with government establishments, which, we will argue in the following section, is a primary determinant of the high-effective regulatory costs.

For unincorporated proprietorships, we can observe information about the gender and social group of the owner. The results in table 2 show that the effective regulatory costs, τ , appear to be much higher for disadvantaged social groups (members of Scheduled Tribe and Scheduled Caste communities) that may lack bargaining power over government officials. The estimate of τ is also higher for female-owned establishments than for male-owned ones, although this difference is not statistically significant.²⁵ Since there is no difference in the substance of the law across gender or caste, the results imply that much of the variation in τ is driven by differences in how it is enforced. We explore this idea further in the next section.

The results above derive from the 2005 EC, but we have also used data from the 1998 EC to test whether there is intertemporal variation in regulatory costs. Using the same empirical methodology described in section IV, we estimate τ at the All India level to be equal to .48 (.12) in the earlier data. Although somewhat larger in magnitude, it lies within the confidence interval of our 2005 estimate.²⁶ Interestingly, the downshift in the 1998 firm size distribution is not as visually striking as that observed in the 2005 data, which may reflect the fact that incentives related to misreporting were different in the two time periods, for example, due to the address slip reporting requirement added in 2005. Let us reiterate, however, that while misreporting may be responsible for visible distortions around the threshold, it is not likely to

²⁵Note that while the estimated regulatory costs for establishments owned by members of disadvantaged communities are high, their contribution to the overall costs is relatively low, since there are few owners from these backgrounds (4.05% of proprietorships are members of Scheduled Tribes, 9.82% are members of Scheduled Castes, and 9.48% are women).

²⁶Relatedly, estimates of τ based on an alternative data set comprising the 2005/6 ASI and the 2005/6 NSSO Unorganized Manufacturing Enterprises Survey are also similar to our estimate of τ using the 2005 EC.

TABLE 3.—ESTIMATE OF τ AT 100 (HIRED)-WORKER THRESHOLD

	τ	SE
All-India	0.0107	(0.0287)

This table presents an estimate of regulatory costs faced by establishments that hire 100 or more workers, using the methodology described in section IV with a bandwidth of 0.005. Standard errors are calculated as in table 1. The estimate is presented for the All-India level (excluding West Bengal) using “hired workers” only.

Source: 2005 Economic Census of India.

affect our estimates of τ . This is especially true of enumerator misreporting (see online appendix B.5 for details).

B. Employment Protection Legislation

In this section, we report the results obtained by using our empirical strategy to test for an increase in per worker costs for establishments that hire more than 100 workers and thus fall under the ambit of Chapter VB of the IDA, the most stringent component of India’s employment protection legislation. As before, we run the test on the 2005 EC and report the standard error in parentheses. One difference in the estimation procedure is that we use the number of “hired workers” of the firm, as opposed to the “total workers” since the IDA excludes nonhired workers.²⁷ Another difference is that we now include dummy variables for firm sizes 1 to 20, so we are effectively comparing the distribution from 21 to 99 with that from 100 onward. We include the dummies from 1 to 9 because we do not want to conflate the effect of the 100-worker threshold with that of the 10-worker threshold, and we include the dummies from 10 to 20 because those values will be most contaminated by misreporting, as implied by our model. Finally, we exclude West Bengal in this analysis because its VB IDA threshold is different. The results, shown in table 3, largely conform to what the figures in section III informally suggest: there is little evidence of a downshift. The implied τ is only .01 and is not statistically significant.²⁸ Chapter VB of the IDA does not therefore appear to have an adverse effect on the unit labor costs of firms.²⁹

VI. Discussion and Investigation of Mechanisms

In the previous section, we documented considerable variation—across states, industries and ownership types—in our estimates of the costs of regulations (τ) applying to firms that employ ten or more workers. In this section, we explore the determinants of this variation and show that dif-

²⁷In fact, the threshold for Chapter VB of the IDA is meant to include only permanent workers, but the number of hired workers is the best proxy we have.

²⁸We obtain similar results when testing for a downshift at the fifty-worker threshold (Chapter VA of the IDA): τ is $-.069$ (.050). However, these results must be interpreted with some caution given that our model implies that firm sizes larger than ten can be affected by the regulations at ten (and twenty).

²⁹Note that our procedure is only capable of capturing distortions in the unit labor costs of firms, as those are the only ones that would show up as a downshift in the log firm size distribution. If the IDA imposes fixed costs, our procedure will not detect them. GLV identify fixed costs from bunching at N , but this is not possible for us because reported bunching may not reflect actual bunching, as discussed in section IVA.

ferences in regulatory enforcement across states (particularly inspector bargaining power and levels of corruption) help explain the variation in τ . Before getting to the results, we note that the analyses we run in this section are necessarily somewhat speculative, since we do not claim to have isolated as-good-as-random variation in regulatory enforcement. Note also that all relevant variables in the following analysis have been rescaled to have mean 0 and standard deviation 1, with the goal of allowing comparability between regression coefficients in different specifications.

A. τ versus Measures of Regulation and Corruption

We begin by regressing our state-level estimates of τ against other established measures of the regulatory environment.³⁰ These include the Besley-Burgess (BB) measure of labor regulations from Aghion et al. (2008), as well as several measures of regulatory reform from Dougherty (2009). The former is a measure of the number of amendments that a state government has made to the IDA in either a pro-worker or pro-employer direction, as interpreted by Aghion et al. (2008), who update the measure to include amendments up to 1997.³¹ Positive values of the BB measure indicate more pro-worker amendments, which are assumed to imply a more restrictive environment for firms operating in those states. Dougherty (2009) provides state-level reform indicators that reflect “the extent to which procedural or administrative changes have reduced transaction costs in relation to labor issues” by “limiting the scope of regulations, providing greater clarity in their application, or simplifying compliance procedures.”³² Higher values therefore indicate an improved environment for firms. Dougherty’s measures are unique in that they cover a wide range of labor-related issues—not just the IDA. In the analysis that follows, we focus on an overall measure of reforms from Dougherty (2009), as well as a measure of reforms regarding the role of inspectors, which aims to capture the extent to which states have reformed rules to constrain the influence of inspectors and includes such actions as limiting the number of inspector visits to one per year and requiring authorization for specific complaints.

Table 4 reports the results of regressing τ against the two measures from Dougherty (2009) and the BB measure. The

³⁰Using a dependent variable that is generated with error leads to standard errors that are biased upward. Weighted least squares is a standard approach for improving precision by weighting more heavily those observations that are estimated more precisely. We therefore weight observations using analytic weights inversely proportional to the variance of our estimate of τ in all regressions, which include τ as the regressand. Our conclusion does not depend on this procedure as we obtain qualitatively similar results when using unweighted regressions.

³¹Since there were no state-level amendments to the IDA between 1997 and 2005, this measure is appropriate for use with 2005 data.

³²These measures are the result of surveying “a labour expert designated by the AIOE [All-India Association of Employers] or Federation of Indian Chambers of Commerce and Industry (FICCI) affiliate in the state capital” of each state, and adjusting the answers “through discussions with local union leaders, independent labour experts, employers and state labour commissioners” (Dougherty, 2009).

TABLE 4.—TAU VS OTHER MEASURES OF REGULATIONS: ALL STATES AND UNION TERRITORIES

	(1) Tau	(2) Tau	(3) Tau	(4) Tau	(5) Tau	(6) Tau
Dougherty measure (all reforms)	−0.360 (0.169)	−0.394 (0.199)				
Dougherty measure (inspector reforms)			−0.480 (0.162)	−0.623 (0.148)		
Besley-Burgess measure (regulations)					0.223 (0.178)	0.235 (0.177)
Constant	0.131 (0.181)	2.900 (5.514)	0.209 (0.140)	−2.952 (5.401)	−0.00266 (0.280)	14.20 (7.402)
Observations	21	21	21	21	16	16
Controls	No	Yes	No	Yes	No	Yes

This table tests for correlations between our estimated regulatory costs (τ) and other established measures of the regulatory environment from the previous literature. Controls include the log of net state domestic product per capita in 2005 and the share of privately owned establishments. Robust SEs are reported in parentheses. Observations are weighted by the inverse variance of τ and include all Indian states and union territories for which data are available.

Sources: Dougherty (2009), Besley and Burgess (2004), and RBI.

TABLE 5.—TAU VERSUS STATE-LEVEL MEASURES OF CORRUPTION: ALL STATES AND UNION TERRITORIES

	(1) Tau	(2) Tau	(3) Tau	(4) Tau	(5) Tau	(6) Tau
TI corruption score	0.617 (0.286)	0.587 (0.321)	0.685 (0.127)			
Electricity losses				0.268 (0.303)	0.254 (0.226)	0.593 (0.153)
Dougherty measure (inspection reforms)			−0.594 (0.0910)			−0.494 (0.0917)
Electricity available (GWH)					0.139 (0.197)	
Constant	0.247 (0.233)	2.736 (4.438)	−2.864 (3.497)	−0.486 (0.251)	−4.347 (3.570)	−3.698 (3.937)
Observations	20	20	19	35	32	21
Controls	No	Yes	Yes	No	Yes	Yes

This table reports the results of our estimated regulatory costs (τ) regressed against two different measures of corruption. Controls include the log of net state domestic product per capita in 2005 and the share of privately owned establishments. Robust standard errors are reported in parentheses. Observations are weighted by the inverse variance of τ and include all Indian states and union territories for which data are available.

Sources: Transparency International (2005), RBI, and Dougherty (2009).

main finding is a robust correlation between the Dougherty (2009) measures and τ : states that saw more transaction-cost-reducing reforms—particularly if they constrained the power of inspectors—have significantly lower τ s.³³ This result is to be expected because Dougherty's measures include reforms that change how firms are affected by laws that vary across the ten-worker threshold. For example, reforms that affect the powers of inspectors certainly have a differential impact on firms above and below the threshold since firms above the threshold fall under the legal ambit of many more inspectors than firms below the threshold. By contrast, we find no strong correlation between τ and the BB measure. This is perhaps unsurprising, as the BB measure captures variation only due to state amendments to the IDA, which does not vary over the 10-person threshold. On the other hand, many studies use the BB measure to proxy for the general regulatory environment (Adhvaru et al., 2013), so we might expect it to correlate with our own measure of regulatory costs.

In table 6 of the online appendix, we report the results of regressing τ against other measures of the labor environment—in particular, per capita measures of strikes, worker-days lost

to strikes, lockouts, worker-days lost to lockouts, and the percentage of registered factories that have been inspected. The only measure that is significantly correlated with τ , echoing the results of table 4, is the percentage of registered factories inspected.

If imposing reforms that constrain the powers of inspectors is correlated with lower effective regulatory costs for firms, this might be because constraining inspectors allows firms to avoid the de jure costs associated with following the rules, or it might be because constraining inspectors makes it harder for them to extort firms for bribes. If the latter, we should expect a strong link between τ and the corruption level of the environment.

B. τ and Corruption

Indeed, the results of table 5 show a large and robust positive association between τ and two measures of corruption. The first three columns of table 5 report the results of regressing τ against state-level corruption as measured in a 2005 Transparency International (TI) Survey.³⁴ One might

³³ τ is not significantly correlated with the other subcomponent measures from Dougherty (2009), except for reforms related to the use of contract workers (not depicted here).

³⁴ The TI corruption measure is based on a survey of the perceptions and experiences regarding corruption in the public sector among 14,405 respondents in twenty Indian states.

be concerned, however, that the TI measure may be flawed as it is partly the result of individuals' perceptions. Therefore, columns 4 to 6 of table 5 report the results of τ regressed against the (normalized) percent of a state's available electricity that was lost in transmission and distribution in 2005. This variable has been used by other researchers as a proxy for corruption and poor state capacity and has the virtue of being an objective measure that does not depend on perceptions (Kochhar et al., 2006).³⁵

Although the state-level correlations between τ and corruption are robust, the regressions are subject to the concern that our measures of corruption may be correlated with omitted variables that also influence τ . To partially address this concern, we provide analysis in online appendix E.1 that corroborates our results using a conceptually different source of variation by taking advantage of within-state, industry-level heterogeneity in the exposure to corruption.

The implication that corruption may increase regulatory costs appears counterintuitive given that much of the literature on regulations and corruption (Khan, Khwaja, & Olken, 2016) has emphasized the role corruption may play in reducing regulatory burden. However, if one allows for the possibility that corrupt inspectors can extort firms by threatening to impose large fines for technical violations of the letter of the law while honest inspectors merely require firms to obey the spirit of the law (a more "reasonable" interpretation of the law that is less costly to abide by), the relationship between regulatory burden and corruption becomes theoretically ambiguous and can easily be positive. We sketch the basic points of such a framework in online appendix G, and in section II, we explain why it is likely that the Indian setting would provide a fertile ground for extortionary corruption.

VII. Conclusion

This paper makes several contributions to the literature on labor regulations in developing countries. We provide estimates of the unit labor costs associated with a suite of regulations whose components have hitherto received little attention. These regulations include mandatory benefits, workplace safety provisions, and reporting requirements where the literature has previously emphasized employment protection legislation and minimum wage laws. In the Indian context, we find that the costs associated with this suite of regulations are much larger than those associated with the most stringent portion of the country's employment protection legislation. Our results suggest that these types of regulations deserve more attention than they have received to this point.

Our results also suggest a mechanism that may explain why these regulations are so costly in a developing country

³⁵Because the samples vary significantly across specifications, we provide results in the online appendix (tables 7 and 8) that restrict the analysis to include only the eighteen largest states, for which data are most consistently available. We also provide partial residual plots associated with columns 3 and 6 of table 5 to demonstrate that the results are not driven by outliers (figures 1 and 2 in online appendix E).

context: high de facto regulatory costs appear to be driven by extortionary corruption on the part of inspectors. Specifically, we show that Indian states that have reformed their inspector-related regulations in a positive way face lower regulatory costs and states with the highest levels of corruption also have the highest levels of regulatory costs. This analysis points to the size of regulatory costs' having more to do with the way regulations are implemented than with the content of the specific laws themselves.

In addition, our paper makes a methodological contribution. We extend GLV's theoretical model to allow firms to strategically misreport their sizes and simultaneously develop an empirical strategy to estimate costs from a firm size distribution under the assumptions of our model. We show that ignoring the problem of misreporting can lead to vastly overestimating the actual costs of the regulations. We believe this contribution will find applications in other developing country settings, where the costs of strategic misreporting are typically low.

We close by noting that our analysis reveals the net costs of regulations borne by firms, but it does not speak directly to their possible benefits for workers. Our results do suggest that the current regulations make it easy for inspectors to penalize firms for technical violations rather than violations of grave consequence. To the extent that this is so, workers may not derive as much protective benefit from the regulations as they might otherwise. It is difficult to arrive at more concrete conclusions without data that would allow measuring how workers would benefit if their employers were made to follow the spirit rather than the letter of the law. However, the results hint at an intriguing possibility: by simplifying regulations identified as costly or by clarifying compliance and enforcement, it may be possible to reduce the costs borne by firms without diminishing effective protection for workers.

REFERENCES

- Adhvaryu, Achyuta, A. V. Chari, and Siddharth Sharma, "Firing Costs and Flexibility: Evidence from Firms' Labor Adjustments to Shocks in India," this REVIEW 95 (2013), 725–740.
- Aghion, Philippe, Robin Burgess, Stephen J. Redding, and Fabrizio Zilibotti, "The Unequal Effects of Liberalization: Evidence from Dismantling the License Raj in India," *American Economic Review* 98 (2008), 1397–1412.
- Almond, Douglas, and Joseph J. Doyle, "After Midnight: A Regression Discontinuity Design in Length of Postpartum Hospital Stays," *American Economic Journal: Economic Policy* 3:3 (2011), 1–34.
- Almunia, Miguel, and David Lopez-Rodriguez, "Under the Radar: The Effects of Monitoring Firms on Tax Compliance," *American Economic Journal: Economic Policy* 10:1 (2018), 1–38.
- Asher, Sam, and Paul Novosad, "Politics and Local Economic Growth: Evidence from India," *American Economic Journal: Applied Economics* 9 (2017), 229–273.
- Axtell, Robert, "Zipf Distribution of U.S. Firm Sizes," *Science* 293:5536 (2001), 1818–1820.
- Bajaj, Vikas, "Outsourcing Giant Finds It Must Be Client, Too," *New York Times*, November 30, 2001.
- Banerjee, Abhijit, "Eliminating Corruption," in M.G. Quibria ed., *Proceedings of the Third Annual Conference on Development Economics* (Mandaluyong, Philippines, Asian Development Bank, 1994).
- Barreca, Alan I., Jason M. Lindo, and Glen R. Waddell, "Heaping-Induced Bias in Regression-Discontinuity Designs," *Economic Inquiry* 54 (2016), 268–293.

- Basu, Kaushik, "Why, for a Class of Bribes, the Act of Giving a Bribe Should Be Treated as Legal," Ministry of Finance, Government of India working paper 1/2011 DEA (2011).
- Besley, Tim, and Robin Burgess, "Can Labor Regulation Hinder Economic Performance? Evidence from India," *Quarterly Journal of Economics* 119 (2004), 91–134.
- Bhattacharjya, Aditya, "Labour Market Regulation and Industrial Performance in India: A Critical Review of the Empirical Evidence," *Indian Journal of Labour Economics* 49:2 (2006), 211–232.
- , "The Effects of Employment Protection Legislation on Indian Manufacturing," *Economic and Political Weekly* 44:22 (2009), 55–62.
- Botero, Juan C, Simeon Djankov, Rafael La Porta, Florencio Lopez-De-Silanes, and Andrei Shleifer, "The Regulation of Labor," *Quarterly Journal of Economics* 119 (2004), 1339–1382.
- Chatterjee, Urmila, and Ravi Kanbur, "Regulation and Non-Compliance: Magnitudes and Patterns for India's Factories Act," (Washington, DC: World Bank, 2013).
- Chaurey, Ritam, "Labor Regulations and Contract Labor Use: Evidence from Indian Firms," *Journal of Development Economics* 114 (2015), 224–232.
- de Andrade, Gustavo Henrique, Miriam Bruhn, and David McKenzie, "A Helping Hand or the Long Arm of the Law? Experimental Evidence on What Governments Can Do to Formalize Firms," *World Bank Economic Review* 30:1 (2014), 24–54.
- de Mel, Suresh, David McKenzie, and Christopher Woodruff, "The Demand for, and Consequences of, Formalization among Informal Firms in Sri Lanka," *American Economic Journal: Applied Economics* 5 (2013), 122–150.
- Debroy, Bibek, "India's Segmented Labour Markets, Inter-State Differences, and the Scope for Labour Reforms," in Bibek Debroy, Laveesh Bhandari, Swaminathan Aiyar, and Ashok Gulati, eds., *Economic Freedom of the States of India* (New Delhi: Academic Foundation, 2013).
- Djankov, Simeon, and Rita Ramalho, "Employment Laws in Developing Countries," *Journal of Comparative Economics* 37:1 (2009), 3–13.
- Dougherty, S. M., "Labour Regulation and Employment Dynamics at the State Level in India," *Review of Market Integration* 1:3 (2009), 295–337.
- Dougherty, Sean, Veronica Frisancho, and Kala Krishna, "State-Level Labor Reform and Firm-Level Productivity in India," *India Policy Forum* 10 (2014), 1–56.
- Eggers, Andrew C., Ronny Freier, Veronica Grembi, and Tommaso Nannicini, "Regression Discontinuity Designs Based on Population Thresholds," *American Journal of Political Science* 62 (2018), 210–229.
- Fagernas, Sonja, "Labor Law, Judicial Efficiency, and Informal Employment in India," *Journal of Empirical Legal Studies* 7 (2010), 282–321.
- Freeman, Richard B., *Labor Regulations, Unions, and Social Protection in Developing Countries: Market Distortions or Efficient Institutions?* vol. 5 (Amsterdam: Elsevier, 2010).
- Garicano, Luis, Claire Lelarge, and John Van Reenen, "Firm Size Distortions and the Productivity Distribution: Evidence from France," NBER working paper 18841 (2013).
- , "Firm Size Distortions and the Productivity Distribution: Evidence from France," *American Economic Review* 106 (2016), 3439–3479.
- Ghosh, Abhik, "Why Make in India Is Stumbling over Our Labour Laws," *Business Standard*, February 28, 2016.
- Hasan, Rana, Devashish Mitra, and K. V. Ramaswamy, "Trade Reforms, Labor Regulations, and Labor-Demand Elasticities: Empirical Evidence from India," *this REVIEW* 89 (2007), 466–481.
- Hernández-Pérez, R., F. Angulo-Brown, and Dionisio Tun, "Company Size Distribution for Developing Countries," *Physica A: Statistical Mechanics and its Applications* 359 (2006), 607–618.
- Hindriks, Jean, Michael Keen, and Abhinay Muthoo, "Corruption, Extortion and Evasion," *Journal of Public Economics* 74 (1999), 395–430.
- Hsieh, Chang-Tai, and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics* 124 (2009), 1403–1448.
- Hsieh, Chang-Tai, and Benjamin A. Olken, "The Missing 'Missing Middle'," *Journal of Economic Perspectives* 28 (2014), 89–108.
- Huntington, Samuel P., *Political Order in Changing Societies* (New Haven, CT: Yale University Press, 1968).
- Kanbur, Ravi, and Lucas Ronconi, "Enforcement Matters: The Effective Regulation of Labor," CEPR discussion paper 11098 (2015).
- Khan, Adnan, Asim Khwaja, and Benjamin Olken, "Tax Farming Redux: Experimental Evidence on Performance Pay for Tax Collectors," *Quarterly Journal of Economics* 131 (2016), 219–271.
- Kleven, Henrik J., and Mazhar Waseem, "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan," *Quarterly Journal of Economics* 128 (2013), 669–723.
- Kochhar, Kalpana, Utsav Kumar, Raghuram Rajan, Arvind Subramanian, and Ioannis Tokatlidis, "India's Pattern of Development: What Happened, What Follows?," *Journal of Monetary Economics* 53 (2006), 981–1019.
- Kouamé, Wilfried A., and Jonathan Goyette, "Tax Evasion in Africa and Latin America: The Role of Distortionary Infrastructures and Policies," World Bank policy research working paper 8522 (2018).
- Kumler, Todd, Eric Verhoogen, and Judith Frias, "Enlisting Employees in Improving Payroll-Tax Compliance: Evidence from Mexico," NBER working paper 19385 (2015).
- Lucas, Robert E., "On the Size Distribution of Business Firms," *Bell Journal of Economics* 9:2 (1978), 508–523.
- Markovitch, Natalia, and Udo Krieger, "Nonparametric Estimation of Long-Tailed Density Functions and Its Application to the Analysis of World Wide Web Traffic," *Performance Evaluation* 42 (2000), 205–222.
- Mishra, Ajit, and Dilip Mookherjee, "Controlling Collusion and Extortion: The Twin Faces of Corruption," Boston University working paper (2013).
- Mookherjee, Dilip, "Incentive Reforms in Developing Country Bureaucracies," in *Annual World Bank Conference on Development Economics* (Washington, DC: World Bank, 1997).
- Nataraj, Shanthi, Francisco Perez-Arce, Krishna B. Kumar, and Sinduja V. Srinivasan, "The Impact of Labor Market Regulation on Employment in Low-Income Countries: A Meta-Analysis," *Journal of Economic Surveys* 28 (2014), 551–572.
- Polinsky, A. Mitchell, and Steven Shavell, "Corruption and Optimal Law Enforcement," *Journal of Public Economics* 81 (2001), 1–24.
- Sequeira, Sandra, and Simeon Djankov, "Corruption and Firm Behavior: Evidence from African Ports," *Journal of International Economics* 94 (2014), 277–294.
- Svensson, Jakob, "Eight Questions about Corruption," *Journal of Economic Perspectives* 19 (2005), 19–42.
- TeamLease Services, "India Labour Report 2006: A Ranking of Indian States by Their Labour Ecosystem," TeamLease technical report (2006).
- Transparency International, *India Corruption Study 2005* (New Delhi: Transparency International India, 2005). <http://transparencyindia.org/wp-content/uploads/2019/04/India-Corruption-Study-2005.pdf>