

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Three Essays on Labor Markets in Telangana, India

Permalink

<https://escholarship.org/uc/item/4vg9r6pk>

Author

Brownstone, Steven

Publication Date

2025

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Three Essays on Labor Markets in Telangana, India

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Economics

by

Steven Brownstone

Committee in charge:

Professor Karthik Muralidharan, Chair
Professor Fabian Eckert
Professor Gaurav Khanna
Professor Craig McIntosh
Professor Gareth Nellis
Professor Paul Niehaus

2025

Copyright

Steven Brownstone, 2025

All rights reserved.

The Dissertation of Steven Brownstone is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2025

TABLE OF CONTENTS

Dissertation Approval Page	iii
Table of Contents	iv
List of Figures	vi
List of Tables	viii
Acknowledgements	xi
Vita	xiii
Abstract of the Dissertation	xiv
Chapter 1 The Labor Market Effects of Agricultural Mechanization: Experimental Evidence from India	1
1.1 Introduction	2
1.2 Background	8
1.3 Intervention	10
1.4 Data	12
1.5 Randomization	14
1.6 Empirical Strategy	15
1.7 Results	16
1.8 Model	21
1.8.1 Model Environment	21
1.9 Conclusion	28
1.10 Chapter Acknowledgments	30
1.11 Figures	30
1.12 Tables	37
1.13 Appendix	49
1.13.1 Additional Sample	49
1.13.2 Pooled Wage Specification	49
1.13.3 Utility Functional Forms	50
Bibliography	51
Chapter 2 The Long Shadow of Feudalism: Concentration of Land and Labor Market Power in India	54
2.1 Introduction	55
2.2 Institutional background	61
2.3 Data	64
2.4 Empirical strategy	65
2.5 Effects on land concentration	68

2.6	Effects on village labor markets	73
2.7	Alternate Explanations	82
2.8	Conclusion	86
2.9	Chapter Acknowledgments	87
2.10	Appendix	88
2.10.1	Model	88
2.10.2	Figures	92
2.10.3	Tables	98
	Bibliography	104
Chapter 3	Managing What You Can Measure? The Promises and Pitfalls of Measuring Bureaucrat Performance	107
3.1	Introduction	108
3.2	Institutional Setting	112
3.3	Treatment	116
3.3.1	Creating Scorecards	116
3.3.2	Sharing Scorecards	116
3.3.3	Validating Self-Reported Data Through MPO Inspections	117
3.4	Randomization and Timeline	119
3.5	Data	120
3.5.1	Validation with Citizen Surveys	120
3.5.2	“Truth” Scores	120
3.5.3	Score Inflation	121
3.5.4	Effort Metric	122
3.5.5	Balance	123
3.6	Main Results	125
3.7	Discussion	128
3.8	Conclusion	134
3.9	Chapter Acknowledgments	134
3.10	Appendix	135
3.10.1	Covariate Predictors of Individual Treatment Effects	135
3.10.2	Treatment Effect by Month	137
3.10.3	PS Score Rubric	139
3.10.4	Generic ML	141
3.11	Correlates of Supervisor Schore	143
3.12	Correlates of Inflation	143
	Bibliography	145

LIST OF FIGURES

Figure 1.1.	Reasons Farmers Stop Drum Seeding	30
Figure 1.2.	Distribution of Uptake Across Villages	31
Figure 1.3.	Main Samples	31
Figure 1.4.	Timeline of data collection	32
Figure 1.5.	Comparing Wage Distribution of Different Tasks	32
Figure 1.6.	Comparing Wage Distribution of Different Tasks	33
Figure 1.7.	Transplanting in Telangana	34
Figure 1.8.	Transplanting Through Time	35
Figure 1.9.	Project Villages	35
Figure 1.10.	CDF of Rural Household's Agricultural Land Owned by Women's Primary Occupation in DHS Survey in Study Districts	36
Figure 1.11.	Distribution of Household Income Effects	36
Figure 2.1.	Georeferencing of a historical map depicting different land tenure systems in princely Hyderabad	66
Figure 2.2.	Balance tests on RD specification	68
Figure 2.3.	Treatment effects on logged landholdings across distribution	71
Figure 2.4.	Regression discontinuity plots along landholding distribution	72
Figure 2.5.	Map of Hyderabad Princely State Showing Different Tenure Systems	92
Figure 2.6.	Phone survey sample	93
Figure 2.7.	Accuracy of Wages	94
Figure 2.8.	Jagirs within India	95
Figure 2.9.	Population dynamics	96
Figure 2.10.	Working population dynamics	96
Figure 2.11.	Wage HTEs by bank exposure	97

Figure 3.1.	Screenshots from PS App	117
Figure 3.2.	Scorecards Correct Misperception	129
Figure 3.3.	Generic ML heterogeneity	129
Figure 3.4.	Relationship between predicted treatment effects	132
Figure 3.5.	Predictors of Low vs High Individual Inflation Treatment Effect	135
Figure 3.6.	Predictors of Low vs High Supervisor Score Treatment Effect	136
Figure 3.7.	Effect of treatment on PS Score By month	137
Figure 3.8.	Effect of treatment on Inflation Score By Month	138

LIST OF TABLES

Table 1.1.	Balance Table	38
Table 1.2.	Sample NFHS/DHS Comparison	38
Table 1.3.	Sample NFHS/DHS Comparison Only Farmers	39
Table 1.4.	Sample NFHS/DHS Comparison Only Rural	39
Table 1.5.	Treatment Effect Village Engagement	40
Table 1.6.	Treatment Effect on Awareness	40
Table 1.7.	Farmer Drum Seeder Specific Knowledge	40
Table 1.8.	Treatment Effect on Uptake	41
Table 1.9.	Treatment Effect on Uptake	41
Table 1.10.	Who Adopts Drum Seeders in Wet Season	42
Table 1.11.	What Predicts Yield for Drum Seeding Farmers	43
Table 1.12.	Farmers Hiring Decisions Normalized by Area Farmed	43
Table 1.13.	Treatment Effect on Costs per Acre	44
Table 1.14.	Treatment Effect on Ln Revenue Per Acre	44
Table 1.15.	Treatment Effect on Paddy Yields and Acres	44
Table 1.16.	Treatment Effect on Profit per Acre in USD	45
Table 1.17.	Main Treatment Effect on Transplanting Wages	45
Table 1.18.	Treatment Pooled Wages	45
Table 1.19.	Labor Costs for Farmers who Transplant	46
Table 1.20.	Transplanters in Treatment Villages Have Greater Unmet Demand for NREGA	46
Table 1.21.	No Effect on Administrative NREGA Days Offered	47
Table 1.22.	Treatment Effect on Husband Allows to Work Outside	47

Table 1.23.	Treatment Effect on Husband Attitudes in Higher Predicted Uptake Strata .	48
Table 1.24.	Calibration Comparison	48
Table 1.25.	Model Simulation Results	48
Table 2.1.	Aggregate land concentration	69
Table 2.2.	Differences in the distribution of landholdings	70
Table 2.3.	Wages for agricultural labor	73
Table 2.4.	Labor demand	75
Table 2.5.	Productivity	75
Table 2.6.	Heterogeneity by Bank Access	77
Table 2.7.	NREGS implementation: registration and household-level access	79
Table 2.8.	NREGS implementation: log person days per month	80
Table 2.9.	Aggregate education levels	81
Table 2.10.	Health and education infrastructure	82
Table 2.11.	Investments in agriculture	84
Table 2.12.	Investments in agriculture	84
Table 2.13.	Agriculture infrastructure	85
Table 2.14.	Differences in violent events	98
Table 2.15.	Differences assembly election results	99
Table 2.16.	Linking large landowners and village leaders through caste	99
Table 2.17.	Balance table	100
Table 2.18.	Treatment effects on village composition	100
Table 2.19.	Wage agreement between workers and leaders in survey	100
Table 2.20.	Characteristics of survey	101
Table 2.21.	Wage survey with workers and officials	101

Table 2.22.	Heterogeneity by SEZ Placement	102
Table 2.23.	Heterogeneity by Road Access	102
Table 2.24.	Heterogeneity by low non-farm employment	103
Table 3.1.	GP performance/progress indicators - change this to a figure.....	116
Table 3.2.	Correlation Matrix: PS Performance Measures	120
Table 3.3.	Balance Table	124
Table 3.4.	Main Effects	126
Table 3.5.	Heterogeneity of main treatment effects by baseline misperception	130
Table 3.6.	Treatment Effects for JPS	131
Table 3.7.	Predicted treatment effects	133
Table 3.8.	Correlates of MPO Score	143
Table 3.9.	Effect of PS' Caste on score inflation percentage	144

ACKNOWLEDGEMENTS

I would like to thank Karthik and Craig for pushing me to pursue my research with a scale and ambition I did not think was possible. I thank Paul for providing thoughtful guidance across many difficult choices in my PhD. I thank Fabian and Gaurav for inspiring me to explore different styles of economics. I thank Gareth for helping me see my work through a different discipline's lens.

I am grateful to my co-authors on the panchayat secretary project: Advitha Arun, Siddhart George, and Naveen Kumar, who, by nature of the endless delays on the panchayat secretary project and their kindness and hospitality, became both unofficial advisors and friends during my PhD journey. I am grateful to my co-author, Kartik Srivastava, for his constant belief in me and our project together, even when it seemed our results were changing every week as we struggled to understand the data.

I thank the many faculty and fellow graduate students in the Department of Economics who, through countless conversations and presentations, improved my work and joy to my life. I want to especially thank Connor Goldstick from my cohort who lived with me for the first three years and helped get me through both the pandemic and my classes. I also want to thank Vivian Auloch who was a co-author on one of my favorite working papers. I also want to thank Erica Chuang who provided emotional and housing support during periods of lease tumult.

I am thankful for the research team at JPAL South Asia and CEGIS for helping me implement my projects and surveys. At JPAL, Smitha Cunnigal, Prerna Beebireddy, and Dimple Khattar got me through some of the most challenging fieldwork I have ever done. I also am thankful for their patience dealing with a first time PI. At CEGIS, Sriharini Narayanan, Sandesh Kotte, Amit Narendula, Javed Afroze, Soumya Jain, and Jhaswanth Naidu provided invaluable advice in dealing with the government of Telangana and providing a welcoming community of co-workers while I was in Hyderabad.

I was supported by many friends throughout this journey. In particular, I want to thank Ramya, Mansi, Luke, and Shilpa who provided a home away from home in India when I was

splitting my life between two continents. I also want to thank Karan, Sebastian, and Sid, who provided some continuity by also splitting their lives across continents alongside me.

Finally, I thank my parents, David Brownstone and Carole Uhlaner, for constantly supporting me in my journey as it ebbs and flows towards and away from the academic world they know best.

Chapter 1, in full, is currently being prepared for submission for publication of the material. Brownstone, Steven; The dissertation author was the primary investigator and sole author of this material.

Chapter 2, in full, is currently being prepared for submission for publication of the material. Kartik Srivastava and the dissertation author, Steven Brownstone, are co-authors of this chapter.

Chapter 3, in full, is currently being prepared for submission for publication of the material. Arun, Advitha; Brownstone, Steven; George, Siddhart; Kumar, Naveen; Muralidharan, Karthik are co-authors of this chapter. The dissertation author was one of the primary investigators and authors of this material.

VITA

2016	Bachelor of Arts, Brown University
2025	Doctor of Philosophy, University of California San Diego

ABSTRACT OF THE DISSERTATION

Three Essays on Labor Markets in Telangana, India

by

Steven Brownstone

Doctor of Philosophy in Economics

University of California San Diego, 2025

Professor Karthik Muralidharan, Chair

This dissertation comprises three essays exploring labor market dynamics in Telangana, India. The essays reveal the interplay of history, technology, and governance in shaping development outcomes.

The first essay examines the effects of agricultural mechanization on labor markets. I study a randomized intervention that helped farmers rent drum seeders to replace manual rice transplanting. While the technology boosts farmer profits by lowering labor costs without reducing yields, it significantly decreases labor demand for the displaced task, primarily performed by women, leading to lower wages. This essay illustrates the distributional tension inherent in technological adoption, emphasizing how labor market frictions shape the impact on workers

and farmers, and potentially slow further mechanization by reducing its relative profitability.

The second essay investigates the enduring legacy of pre-colonial feudal land tenure systems on contemporary labor markets. Utilizing a regression discontinuity design across historical boundaries within the former Hyderabad princely state, it demonstrates how villages with higher land concentration have more unequal labor markets. Notably, female agricultural laborers face significantly lower wages due to constrained mobility and outside options, and local elites strategically curtail public workfare programs during peak agricultural seasons, highlighting how historical power structures perpetuate inequality and undermine social safety nets.

The third essay examines the challenges of making the labor market for government jobs more dynamic. I study the impact of a performance monitoring system on local government bureaucrats, specifically Panchayat Secretaries, who are responsible for a wide range of village governance tasks. An experiment providing scorecards based on administrative data reveals a complex trade-off: while potentially aiding learning, such monitoring can incentivize data manipulation over tangible performance improvements, resulting in poorer service delivery outcomes. Notably, bureaucrats hired under a new scheme with job security contingent on performance were no more likely to respond positively or negatively to performance information than their colleagues with standard job security. Instead, difficult-to-observe types generated by complex interactions of covariates appear to be the primary driver of whether bureaucrats respond positively or negatively, highlighting the challenge of designing performance management for the bureaucracy.

Chapter 1

The Labor Market Effects of Agricultural Mechanization: Experimental Evidence from India

1.1 Introduction

No country has grown rich without also experiencing an increase in agricultural labor productivity. Developing countries' agriculture sectors are laggards in the quest for growth. The productivity gap between the wealthiest and the poorest 5% of countries is fifteen times greater in agriculture than in non-agricultural sectors (Restuccia et al., 2008). Mechanization significantly contributes to the agricultural labor productivity gap across countries, accounting for up to 37% of the difference (Caunedo and Keller, 2020). A key barrier to the adoption of mechanization is the small scale of farms in developing countries, which makes buying most machines unprofitable (Foster and Rosenzweig, 2022). One solution is rental markets, but these are often absent for new technologies.

While boosting productivity, mechanization also often has distributional consequences. Acemoglu et al. (2024) highlights how mechanization, defined as replacing labor with capital for a specific task, has a theoretically ambiguous effect on wages. If workers can compete for jobs across the economy, higher productivity per worker should translate to higher wages, but if workers' only option is the mechanized task itself, wages will fall in response to decreased labor demand. The degree to which workers are willing to continue to supply labor regardless of wage for the mechanized task, the labor supply elasticity, determines how far wages fall. The distributional impacts create a tension between improving productivity and protecting workers. The impact on workers feeds back into the equilibrium effect on productivity. Wages for the mechanized task determine the relative profitability of mechanization for later adopters (Hornbeck and Naidu, 2014; Manuelli and Seshadri, 2014). The wage for the mechanized task also partially reflects displaced workers' productivity in their new jobs, which determines how mechanization translates to changes in aggregate productivity driven by the new output displaced workers generate in their new jobs.

This paper studies mechanization and its consequences in the context of rice farming in the Indian state of Telangana. Drum seeders are low-cost devices that increase the labor

productivity of rice planting by eliminating transplanting, which comprises a fifth of the cost of rice cultivation. However, adoption requires access to both the devices and the knowledge to use the devices successfully. Although private rental markets exist for well-known technologies like tractors, they are largely absent for drum seeders. The need for skilled staff to train large groups of farmers clashes with the individualized nature of equipment rentals, making training a low-profit endeavor for rental entrepreneurs.¹ Governments provide agricultural training through a system of agricultural trainers and extensionists who offer skills but not capital. I worked with the Government of Telangana to help design a program to coordinate extension with village government drum seeder rentals facilitating drum seeder adoption. Beyond the interest in improving aggregate productivity, governments and carbon credit providers have additional interest in drum seeders due to the technology's potential to reduce water use and methane emissions from rice cultivation.²

Working with the government enabled deploying this program at scale across 205 randomly selected villages with over 70,000 farmers. Rice transplanting in Telangana is a setting where comparing villages closely approximates comparing separate rice transplanting labor markets. Transplanting is done by locally hired female labor. While men travel between villages to work, women do not, which keeps the labor market segmented at the village level. In addition to being segmented, village labor markets are observable. Unlike settings dominated by family labor, even small farmers in Telangana hire transplanting labor.

The study uses primary data collected from surveys of approximately 2,200 farmers in 342 villages sampled from the administrative register of rice farmers and phone surveys of 1,400 female workers. The rice farmer sample includes all types of rice farmers and reflects the demographics of the broader rural population in the study districts. The in-person survey

¹A related issue is that farmers might not trust entrepreneurs who are trying to profit from a new device to impartially explain its costs and benefits

²The Good Rice Alliance carbon project has enrolled enough farmers in India to target 100,000 tons equivalent in methane emission reductions from direct seeded rice and alternate wetting drying, which drum seeders facilitate (NDTV, 2024). Rice is responsible for 6-11% of global methane emissions (Smith et al., 2021). Crop burning is another potential environmental externality. However, in this study only 28 farmers self-reported rice stubble burning, one of which drum seeded.

encompasses data on all farming costs and revenues, with meticulous collection of task-specific wages. The survey also included detailed data on farmer training and extension. The worker phone survey included questions on work locations and government job scheme utilization during the rice transplanting period.

This study presents three main sets of empirical results: the impact on drum seeder adoption, the impact on farm productivity and profits, and finally the labor market consequences.

First, the program led to a nearly 300% increase in drum-seeded acreage in the primary rice season, an increase from 14 to 50 acres per village or 20% average utilization of the village's drum seeders.³ Mirroring a natural uptake process, more educated farmers with more land were more likely to adopt drum seeding. On the extensive margin, the fraction of farmers using drum seeders at all increased from 1% to 4% (about 12 additional farmers per village). Adoption was highly uneven across villages, with many villages with no adoption and others with up to 60%. In addition to using the rentals, farmers in treatment villages reported receiving more training on drum seeders, as designed. Further, the uptake effects appear durable, persisting into the secondary wet season.

Second, drum seeders improved farmers' labor productivity. Farmers in treatment villages employed 7% less hired female labor per acre. They did not use any more family labor or spend significantly more on other inputs, as compared to control farmers. Despite the reduced labor inputs, farmers in treatment villages experienced similar or even slightly higher yields than control farmers at the end of the season. With equivalent revenue and lower total costs, including the rental fee, farmers in treatment villages achieved profits per acre of rice that were 8% higher than in control villages. This large increase in profits relative to takeup is an intent-to-treat effect that includes impacts on both drum seeding and transplanting farmers in treatment villages.

Overall, farmers in treatment villages spent 10% less on hired labor. Even farmers who

³The adoption of mechanization is notoriously slow. The modest increases are expected as primary crops mechanize slowly over many seasons (Gulati and Juneja, 2020)) as farmers learn from their neighbors (Foster and Rosenzweig, 1995). Tractors took many years to spread in the US (Manuelli and Seshadri, 2014)

did not adopt drum seeders spent 7% less on hired labor.⁴ These savings came from transplanting farmers paying 6% lower piece-rate wages in treatment villages. While the identity of farmers who transplant rice in treatment villages is endogenous, the wage is an equilibrium object at the village level. Farmers competitively bid against each other within the village to attract transplanters on a given day. The wage decreases likely partially reflect women having few other job options outside these transplanting labor markets. There is suggestive evidence from the laborer survey that workers did not find other jobs in the short term. Unlike men, only 2.5% of women worked outside their village during the primary wet season, which is the focus of the results. Workers who were transplanters reported a greater unmet demand for days of paid manual work from a government job scheme, workfare, during the transplanting season. The laborer survey found no evidence that the transplanters immediately found jobs in other villages or sectors.

To explore the interplay between labor markets and mechanization, I developed a quantitative model examining how reducing barriers to mechanization affects productivity and wages under different labor market conditions. Through the model, I explore how, through wages, mechanization can sow the seeds of its own demise. In the model, the labor supply curve's slope determines how far wages fall in response to the mechanization labor demand shock, which in turn affects how many farmers stand to profit from mechanization. In the model, the equilibrium wage also serves as a proxy for the productivity of displaced workers in new jobs. Since agricultural output in the model is constrained by land, as is true in many developing country settings, productivity growth from mechanization is entirely driven by the output displaced workers achieve in their new sectors. Calibrating the model to the experiment allows me to quantify the role of wages in either amplifying or mitigating the effects of reduced mechanization fixed costs on equilibrium mechanization and productivity.

I explore how the experiment's impact on adoption and wages would change if the

⁴While this regression is not causal since who transplants in the treatment group is endogenous, the implied treatment on the treated effects, if there was no treatment effect on non-drum seeding farmers, are more than 100%.

reduction in fixed adoption costs remained the same, but the labor market changed. The first scenario simulates offering government jobs alongside mechanization subsidies. Absorbing half the jobs lost due to mechanization in a government jobs scheme does increase mechanization by 30% and reduce the wage decrease from 6.3% to 5.5%. However, this comes at a cost to productivity since the government jobs scheme is assumed to produce no output. The second scenario simulates improving the set of jobs women can access, thus reducing wage sensitivity to labor demand shocks.⁵ Making the labor supply 20% less elastic nearly triples mechanization in equilibrium from 3.42% to 9.13% and more than triples the aggregate productivity increases associated with reduced fixed cost of adoption from 5% to 19% since more workers move out of the agricultural sector and are more productive in their new jobs. The model illustrates the quantitative importance of labor market structure in shaping the impact of agricultural mechanization interventions in settings similar to the experiment.

The study's first contribution is demonstrating a way for governments to promote mechanization without large capital subsidies. The extensive literature on agricultural technology adoption focuses on technologies where the primary barrier is either skill (Conley and Udry, 2010) or capital (Caunedo and Kala, 2021) alone, but few papers address both barriers simultaneously. Prominent social enterprises, such as One Acre Fund, combine training and input access (Deutschmann et al., 2019), but they often need donor funding to scale. The rental aspect of this study is self-financing. The drum seeder's purchase cost was equivalent to 22 days of rentals at the study's price, at most two seasons. Other research emphasizes the importance of rental markets for improving access to capital goods (Bassi et al., 2022) but does not explore how these markets form. A private sector rental market had not emerged in the control arm of the study, suggesting government intervention contributed to kick-starting rentals.⁶ Integrating rental markets offers a promising avenue for governments to fully capitalize on existing agricultural

⁵This would correspond with government policy interventions to increase labor market opportunities for women, such as free access to inter-village buses.

⁶I briefly attempted to partner with an NGO that had identified village entrepreneurs interested in renting drum seeders, but the entrepreneurs were not willing to take a risk on buying the devices.

training investments.⁷

Second, the study contributes to the literature on the impact of mechanization on agricultural labor markets by using a market-level experiment to show how mechanization can affect wages and labor demand. National-level studies (Bustos et al., 2016) are limited to the long-run effects of labor displacement and cannot analyze mechanization's impact on non-adopting farmers. On the other hand, household-level randomization of mechanization subsidies (Caunedo and Kala, 2021) can only identify farm household labor supply decisions rather than changes in the labor market. The study bridges the national and household level literature to identify the impacts of mechanization by identifying spillovers from mechanization at the labor market level. Other evidence on technology adoption and labor markets demonstrates how labor shortages constrain the adoption of labor-complementing technologies like irrigation (Jones et al., 2022) and how higher rural wages can spur mechanization, in the long run (Hornbeck and Naidu, 2014; Manuelli and Seshadri, 2014). Given this evidence, the lower wages caused by mechanization in this paper will likely slow mechanization until workers adjust or until the next labor-complementing shock, such as an irrigation expansion, occurs.

Third, the study also contributes to the literature on rural labor markets. The results reinforce the finding that rural labor markets only experience seasonal slack (Breza et al., 2021). Further, in contrast to the overall downward rigidity in response to seasonal reductions in agricultural activity (Kaur, 2019), peak agricultural wages are very sensitive to changes in labor demand. The results in Muralidharan et al. (2023a) highlight how rural wages increase with more credible outside options. This study shows that when mechanization removes job options wages decrease.

Fourth, this study contributes to an emerging literature on low female labor force participation, particularly in India. This study provides micro-level experimental evidence consistent with a macro trend of declining female labor participation. The study reinforces the finding that

⁷A neighboring state, Andhra Pradesh, has combined agricultural equipment rentals with extension at integrated farmer services hubs called Rythu Bharosa Kendras (Reddy et al., 2023).

districts in India more agronomically suited to mechanization have lower female labor force participation but no reductions in male labor (Afridi et al., 2023). The study also provides another example of wages for Indian women decreasing in response to rural labor demand shocks (Afridi et al., 2022).

Finally, this study provides a clear illustration of the broader tension countries face between increasing productivity and protecting workers. This paper provides an empirical example of a mechanization event where an entire task is replaced, which is the focus of recent theoretical and empirical work highlighting the distributional consequences of mechanization (Acemoglu et al., 2024). These distributional consequences sometimes can create political push-back.⁸ The quantitative model suggests that social protection and labor mobility policies can address political concerns by making mechanization more equitable and, in equilibrium, more effective as well.

1.2 Background

Manual transplanting (Figure 1.7) remains the dominant method for growing rice in Telangana for two reasons. First, high transplanting wages and even the existence of wages are recent phenomena. The transplanting labor market is now tight. Second, switching rice cultivation to drum seeding is deceptively complex and requires substantial training.

This paper is not about the adoption of a newly invented complex technology. In fact, drum seeders were first promoted by the International Rice Research Institute (IRRI) in India over a decade ago. These simple plastic devices are now manufactured locally in India. They are simple enough that at least one farmer before the study created one from scratch inspired by a YouTube video. The reason farmers resorted to making drum seeders themselves is that they are only available for purchase in district headquarters and not available for rent in study districts' agricultural rental businesses. These rental businesses, custom hiring centers, provide

⁸From British Luddites to Gandhi, opposition to mechanization on behalf of workers has been a prominent political position. In India, politicians often conflate policies supporting inefficient firms with protecting workers (Muralidharan, 2024)

easy access to tractor and other device rentals. However, they do not provide any training.

Two big shifts in transplanting labor markets likely explain farmers' newfound interest in drum seeders. First, farmers shifted from exchanging transplanting labor to working in their neighbor's fields for wages. Second, these wages increased dramatically in the past 10 years, nearly 45% in real terms. During piloting, village leaders consistently cited a "labor shortage" as the most pressing concern facing farmers. Of course, what these farmers really meant was a labor shortage at the prevailing wage. I hope to explore the drivers of these changes in future work, but irrigation expansion leading to expanded rice cultivation seems to be an important driver of the wage increases. The irrigated area has increased by 117% and rice production by 342% in Telangana from 2015 to 2022 (Planning Department, 2023).

Wages also likely reflect a compensating differential for the unpleasantness of crouching in mud for hours at a time. Willingness to pay studies find Indian wives have a higher willingness to pay than their husbands for drum seeders (Khan et al., 2016) and mechanical transplanting machines (Gulati et al., 2024). There are also health risks associated with spending extended periods in flooded rice fields, including greater exposure to water-borne diseases, mosquitoes, and snakes (Vent et al., 2016). Thus, both men and women in farming households have substantial incentives to adopt drum seeders.

As has been found for other seemingly simple agricultural technologies (Aker and Jack, 2023), training is a major barrier to drum seeder adoption. Drum seeding simply involves rolling a drum seeder along a leveled and damp rice paddy to place the rice seeds in rows along the field. However, farmers must make other adjustments to ensure they drum seed successfully. Importantly, these adjustments take place at other points during the season before and after farmers use drum seeders. This makes a service provider model where service providers drum seed on behalf of farmers unable to fully address the training barrier to drum seeder adoption.

Farmers must adjust their land preparation, irrigation, and weed management when using drum seeders. The farmers must ensure the field is well leveled so that there are no deep puddles

of water that stop the rice from germinating.⁹ A major concern farmers have with drum seeding is weeds. While the relative weed susceptibility of drum-seed rice is debated by agronomists, farmers feel the method is more weed-prone. Weeds were the number one reason farmers who had tried drum seeding at some point in the past reported no longer using drum seeders (Figure 1.1). However, a field can become weed infested at any time so part of the complaints could be driven by farmers associating weed infestations with drum seeding. Very few farmers practice the recommended weed management regimen: re-flooding after planting, spraying a limited amount of herbicide before weeds emerge in the early season, followed by manual weeding in the late season. Most spray only after weeds emerge, and many skip manual weeding altogether. The strongest correlate of yields among drum seeding farmers was weed management knowledge, specifically to flood the field again after germination as a form of weed control (Table 1.11). Thus, access to training across the season and not just when farmers use the device is important.

Given the knowledge required to successfully drum seed, farmers clearly require both extension and access to make the best use of drum seeders. This was especially true in villages without any historical drum seeder use.

1.3 Intervention

The intervention's origin stems from the government of Telangana's interest in consolidating its investment in decentralization by strengthening village governments with its investment in agricultural training by hiring one college-educated extension officer for every 5,000 acres. Locally elected governments can mobilize citizens and wield fiscal autonomy but are short-staffed with one skilled bureaucrat saddled with a dizzying array of responsibilities (Brownstone et al. Forthcoming). On the other hand, extensionists have extensive agronomic knowledge but struggle to interest farmers in hypothetical technologies requiring upfront investment. I worked with the Telangana Department of Agriculture to develop a pilot program that would address

⁹A surprising cause of germination failure is birds eating the exposed seeds, which can be partially addressed by leaving excess seeds in piles near the fields for birds to eat

extensionists' and village governments' concerns. Local governments could leverage their fiscal autonomy to rent agricultural technology and leverage their popularity to mobilize farmers to receive training from extensionists on the same technology. Extensionists would be able to reach more farmers and provide training on a technology farmers could easily access.

Of the handful of priority technologies extensionists promoted, drum seeders interested village elected leaders, sarpanches, the most. During pilot field visits, village leaders consistently flagged transplanting labor shortages as farmers' top issue. Interestingly, while leaders recognized earnings from transplanting wages were important for many of their constituents, they felt promoting drum seeders would not hurt transplanters since there was a "labor shortage."¹⁰ Extension agents complained about an inability to mobilize farmers to attend meetings and the difficulty of promoting devices they could not directly provide.

The primary intervention was supplying village governments with drum seeders and allowing them to rent the drum seeders to their citizens. The devices were rented for 250 INR per day which was the median of rental prices in other districts with higher adoption. This price also ensured that after 22 days of rentals, approximately 2 seasons, the device's cost would be covered. Farmers were allowed to share the rental as one machine could cover up to 4 acres daily. The income from the rentals was deposited into the village government's bank account. Village governments were told the money could be spent on village agricultural development following the norms for other village spending. The actual government order formalizing the spending guidelines did not pass until after the experiment.

The role of the extension agent in the rentals was cross-randomized, with some extension agents directly managing the rentals on behalf of the village governments and other extension agents focusing solely on promotion activities, with the village government completely managing the rentals. The goal of this sub-experiment was to better understand the role village elected governments can play in extension, but in practice, local elected leaders involved themselves in

¹⁰Surprisingly, there was no correlation between the fraction of support elected leaders estimated came from laborers and the leaders' expectation of gaining electorally as a result of the program.

drum seeder promotion in both arms (Table 1.5) leading to similar levels of drum seeder adoption (Table 1.8). Thus, I will focus on the pooled treatment for the remainder of the paper.

1.4 Data

The project collected data from two main overlapping sampling frames. The farmers were surveyed from a database of farmers and cultivated land maintained by the agriculture department. Several aspects of this dataset strengthen its credibility. First, inclusion in this database is necessary for farmers to get a sizeable cash transfer called Rythu Bandhu or PM-Kisan. Thus, the department devotes considerable time and resources to ensuring these records are complete. The data is also linked to land records which further reduces the risk any farmed area is excluded. The sampling frame was restricted to farmers recorded as cultivating rice. I also drew two small supplemental samples of drum-seeding farmers discussed in Appendix 1.13.1. These supplemental samples were used for descriptive regressions looking at predictors of yield for drum seeding farmers.

The second sampling frame is female laborers from the workfare, National Rural Employment Guarantee Scheme (NREGS), rolls.¹¹ These are women from households that wanted the opportunity to take part in India's primary rural welfare scheme, NREGS. Eighty percent of households in the farming sample also had NREGS cards. While not all women who appear on the NREGS rolls are transplanters, nearly all transplanters take advantage of NREGS work during periods when other agricultural work is unavailable. See Figure 1.3 for a summary of the sampling frames.

While neither sampling frame perfectly represents the population of the study villages, it captures a wide swath of the relevant populations. Thus, the data can illustrate the heterogeneous impacts of easing access to drum seeders, which is the study's primary goal.

Due to the autonomy of the local governments, village elected leaders, sarpanches, had

¹¹I scraped job cards from the NREGS website and then asked village-level officials to track the women on the job cards. Note that with mandatory linking of biometrics, fake job cards have largely disappeared from Telangana.

to opt-in to participate in the study. Of those contacted, 33% of villages had leader was not interested in the program. Often, these were villages with less paddy cultivation, but they were also less likely to have any drum seeders and had slightly lower transplanting wages on average. Nonetheless, the selected villages broadly cover the selected districts geographically (Figure 1.9) and represent a wide range of previous drum seeder exposure and prevailing transplanting wages. Importantly, any scaled-up version of this program in India would have to operate on a similar opt-in basis. Some degree of self-targeting is likely optimal since local leaders are best informed about their community's agricultural needs.

For survey sampling, the randomization strata were grouped into four categories based on how the variables used for randomization predicted administrative drum seeder uptake. Based on predicted uptake, the sample size targets ranged from 4 to 12 to help ensure that drum-seeding farmers were included. In the lowest predicted uptake strata, only one of the two control villages was sampled. Thus, the primary results represent a hypothetical set of villages that were slightly better targeted. A scaled-up program could simply use the results of this study to screen villages better based on the data collected to form strata.

Notably, the final farmer sample matches the 2019 demographic and health survey (DHS) sample for the study districts relatively closely. The comparison between the sample and the DHS, designed to be representative at the district level, is complicated by the amorphous definition of rural in the Indian context. The DHS sample includes both urban and rural areas, but many areas classified as rural are, in fact, small towns with limited farming activity. Thus, I present a comparison to the full sample for the six districts (Table 1.2), the sample of individuals owning agricultural land (Table 1.3), and finally, individuals residing in areas classified as rural (Table 1.4). Notably, the three variables most important for drum seeder uptake are not significantly different between the study's sample and the DHS sample: acres owned, household head education, and irrigation pump ownership. Wealth proxied by televisions, pressure cookers, and livestock is similar across samples.

Telangana has two rice-growing seasons: a winter dry season (November to March) and

a summer wet season (June to October). While drum seeding during the wet season can make water management more challenging, there is also more rice cultivation during the wet season overall. The project began halfway through the planting period of dry season 2022-2023. I considered this a pilot season. The primary season discussed in this paper is the wet season of 2023. Before the start of this season, we surveyed village leaders and agricultural extensionists and identified workers from the NREGS rolls for the laborer survey. After planting but before harvest, we conducted the first farmer survey and asked detailed household and agronomic questions. Telangana then entered its state election period, which meant the state government was far less involved in monitoring the drum seeder program. Thus, the treatment effects in this second season can be seen as evidence of the durability of the uptake. However, increasing uptake driven by social learning will likely not occur until the subsequent dry and wet seasons since farmers consider the wet and dry seasons agronomically distinct. The primary survey with data on full wet season yields, costs, and wages took place after the state elections but before the national elections, from February to March of 2024. A follow-up phone survey of transplanter took place in March and April, and a final phone survey to capture dry season yields for farmers took place in May and June 2024. See figure 1.4 for a full timeline.

1.5 Randomization

This study was randomized at the village level, which is important for capturing labor market effects.

As mentioned above, villages were screened before randomization based on the village elected leader's willingness to participate. Part of this screening included a limited questionnaire for village elected leaders, which formed the basis for the randomization since the timeline did not accommodate a baseline. Before randomization, one village was selected per extension agent to ensure no agent served both a treatment and control village. Among the villages served by an extensionist willing to participate in the study, I included the village closest to

the extension agent’s office in the randomization. The randomization was stratified by initial transplanting wage and drum seeder uptake as well as the district. Randomizing only villages served by different extensionists led to greater spacing between treatment and control villages. The distance helped to limit the risk of spillovers driven by the 2.5% of female transplanters who report doing some work outside their village during the transplanting season. Since treatment and control villages were seldom neighbors (Figure 1.9), the likelihood that control villages pulled labor from treatment villages is very low.

The resulting randomization led to a sample of farmers that is balanced between treatment and control villages (Table 1.1). Unfortunately, the 2022 administrative rice cultivation data was unavailable for some farmers due to an inconsistency in IDs between different rounds of data shared by the government. Household characteristics are also missing for the 3% households that were only surveyed in the second round of data collection due to logistical issues in the first survey round.

1.6 Empirical Strategy

The main specification for the following results is the standard approach for cluster randomized controlled trials.

$$Y_{ivse} = \beta_1 T_v + \chi_e + \gamma_s + \varepsilon_{vi}$$

Y_{ivse} is the outcome variable. T_v is the treatment dummy for village v indicating whether drum seeders were rented. χ_e are enumerator fixed effects. γ_s are strata fixed effects. Finally, ε_{iv} errors clustered by village. In this specification, β_1 is the treatment effect. The limited control variables available only improve precision marginally,¹² so results are reported without control variables for simplicity.

Note that estimating the treatment on the treated regressions likely will not isolate the

¹²The most useful control variable is administrative data on rice cultivation prior to the experiment. There are also variables on historical wages, soil types, and educational attainment, which I assume are time-invariant and unaffected by the treatment.

effects of drum seeding in this study. This is because there are substantial spillovers from farmers drum seeding to those not drum seeding driven by the wage effects. This violates the exclusion restriction required to use village treatment assignment as an instrument for drum-seeding as in the standard treatment on the treated approach. The treatment affects non-compliers, e.g., farmers who do not drum seed, by lowering their labor costs. This means dividing the profit treatment effect by the fraction of drum-seeding farmers would overestimate the profitability of drum-seeding. Thus, in the results section, I focus on the overall intent to treat effects, which compares outcomes for all farmers sampled in treatment villages to all farmers sampled in control villages. The only identifying assumption required for these results is that the randomization is independent of any confounding variables.

1.7 Results

The project successfully increased farmers' awareness of drum seeders (Table 1.6). It also improved farmers knowledge of drum seeding practices correlated with yields (Table 1.7). It also successfully got local elected leaders to engage in promoting drum seeders (Table 1.5). However, as previously noted, there were no differences between the cross-randomized arms in leaders' involvement or uptake (Table 1.8).

The treatment did successfully lead to a large percentage increase in drum seeder cultivation in both the wet and dry seasons (Table 1.9). However, the actual number of farmers experimenting with drum seeders was relatively small. The treatment effects translate to approximately 10 acres in control villages as compared to 50 in treatment villages. However, 40 acres is still 400 days worth of transplanting work. There are several possible reasons for the low absolute uptake. First, farmers do not generally make dramatic changes to how they cultivate their primary crop quickly. Farmers wait to see the success of earlier adopters before making their own adoption decisions. Second, qualitative data showed farmers placed outsized signaling value on early adopters' yields. One negative or positive experience can shape an entire village's

attitude toward technology. For example, in some villages, farmers were convinced drum seeders led to higher yields because the early adopters had used drum seeding as an opportunity to try more expensive seeds. In a neighboring village, a farmer tried a different new seed along with the drum seeder, which the farmer blamed for the poor yields. However, the narrative in the village was focused on the drum seeder. The importance of farmers learning from others partially explains why uptake was not uniformly spread across study villages (Figure 1.2). The majority of treatment villages had no drum-seeding farmers in the sample, but in other villages, large numbers of farmers switched to drum-seeding. These high uptake villages tended to have some previous experience with drum seeders which likely primed farmers' interest in trying the technology.

The acreage converted to drum seeding matters more than the number of farmers adopting the technology for driving labor market effects. Importantly, larger farmers were much more likely to adopt drum seeding (Table 1.10) than small farmers.¹³ The smaller treatment effect on dry season acreage is likely driven by the overall reduction in paddy cultivation in that season. Only fields with a reliable water supply can be cultivated during the dry season. Thus, long-term adoption dynamics will require following these villages for many more seasons.¹⁴

Drum seeder adoption led to the expected decrease in hired female labor days (Table 1.12). This variable reflects women hired across all tasks, but since harvesting is mechanized, transplanting is the only major agricultural task requiring large numbers of women in this context. The intent to treat effect of .73 fewer hired female workers per acre roughly corresponds with the 10 fewer hired laborers per transplanted acre expected when the fraction of drum-seeding farmers is considered. Since these are intent to treat effects, any additional hiring by non-drum seeding farmers induced by the lower wages is taken into account in this estimate. To reconcile these

¹³In the administrative data, a handful of very large farmers adopted drum seeders. These farmers drum seeded 20 to 30 acres by themselves, but were not sampled in the survey. Thus, the acreage estimate is likely a conservative estimate of the overall effect.

¹⁴I am currently partnering with remote sensing experts in an effort to track long-run uptake at the village level using remotely sensed flooding patterns, which can provide a good proxy as to whether a field was transplanted or not.

results with the uptake results, there would have to be minimal additional hiring by non-drum seeding farmers. Since the transplanting operation takes a fixed amount of labor per land area, seeing no additional hiring makes sense. Thus, this result suggests that, aside from the choice of drum seeding, labor demand for female hired labor is inelastic.

This decrease in labor hired translated to decreases in labor expenditure in treatment villages (Table 1.13). There were no major compensating changes in family labor or other expenditures. The different numbers of observations in these regressions reflect both zeros as the variables are in logs and the exclusion of farmers who felt they couldn't accurately report expenditures despite surveyors' efforts to break up expenditures into sub-categories.

While it seems clear drum seeders saved farmers money, that would be meaningless if it came at the expense of revenues. However, there is some evidence that farmers in treatment villages were able to obtain higher revenues than those in control villages (Table 1.14). At the time of the survey, a number of farmers were not able to recall the price they sold their rice for the previous season. However, the distribution of prices is quite narrow, so it is possible to impute using district median prices. The modest increases in revenue could be driven by a few factors. One possibility is that by avoiding large transplanting labor expenses, farmers were able to make small additional non-labor expenditures that, while not statistically significant individually, mattered for revenue. For example, decisions by individual farmers to buy slightly more expensive seeds, fertilizers, additives, or herbicides could end up having an aggregate impact on yields, which also appear to rise modestly (Table 1.15). Another possibility is that farmers paid more attention to land leveling because they planned on using drum seeders. Other studies have shown that land leveling is important for rice yields overall since it ensures the even distribution of water and other inputs. (Ali et al., 2018) (Magnan et al., 2015).

Overall, the revenue and cost effects translate into a positive and significant treatment effect on winsorized and unwinsorized profits using the imputed revenues (Table 1.16). The effect is 9% for rice profits and 6% for total profits. Note that not all of these profit effects come from the drum-seeding farmers alone. Since wages for transplanting decrease in treatment villages,

even farmers who don't adopt drum seeding benefit from lower labor costs. As previously mentioned some of these lower labor costs could lead to small re-optimizations that improve yield and profits as well. Assuming labor demand for transplanting is perfectly inelastic, since it takes a fixed number of workers to do the task on a given farmer's land, the wage effects alone would drive a 5 percent decrease in labor costs as compared to the 10 percent observed.

Farmers pay transplanters in two ways: wages per acre transplanted and wages per day. I find wages decrease for both types of wages individually (Table 1.17) as well as when they are pooled (Table 1.18). The smaller sample sizes are because some farmers just reported total labor costs which can encompass multiple tasks. This data is included in the total labor cost per acre regression, which is significantly negative for transplanting farmers even though the total days hired is not (Table 1.19). The wage per acre, although the rarer type of wage, is the preferred result since it is a piece rate wage, so it should be unaffected by productivity differences between transplanters hired. This wage is 6% lower in treatment villages than control villages.

Wages are responsive because the transplanting labor market appears to be congested. Congestion in the transplanting market is evident when comparing the distribution of manual weeding, an optional time-flexible task done by the same women in the same fields, to transplanting wages. While both distributions have the same medians, there is a large mass of high transplanting wages that do not exist for weeding (Figure 1.5). Farmers report trying to stagger their transplanting to avoid these peak wages, but once the rice seedlings are matured in their nurseries farmers have a relatively short window to transplant them into the main field. Hydrology limits farmers' ability to stagger planting.¹⁵ The treatment villages have fewer peak wages associated with labor market congestion, as seen in this histogram comparing daily wages in treatment and control areas (Figure 1.6).

The fact that transplanting wages went down means that while the treatment appears to have been a success, focusing on the farmer sample, not everyone benefited. The exact

¹⁵Leveraging remote sensing to better understand the environmental determinants of observed transplanting wages is an area of ongoing analysis.

distribution of benefits depends on the joint distribution of labor supply and paddy cultivation. The population representative national health and family survey has data on both agricultural land ownership and women's primary occupation. Importantly, the survey distinguishes cultivation, which is household agricultural work, and agricultural labor. Comparing the cumulative distribution functions of land ownership for laborers and non-laborers shows that more than half of laborers come from land-owning households and a large fraction of laborers come from households with substantial amounts of land (Figure 1.10). Small landholders are marginally more likely to work as laborers, and households with very large land holdings are unlikely to work as laborers. Taking these insights and the experimental results, I estimate the overall change in household income driven by the experiment factoring in changing agricultural profits, fewer days of work, and lower wages (1.11). As expected, the impacts are asymmetric, with the bottom 20th percentile of households negatively affected and the largest benefits accruing to the largest landowning households.

There is evidence from the laborer survey that displaced workers were underemployed. Transplanters in the laborer sample reported 2 more days of unmet demand for workfare days in treatment villages in the preceding two weeks compared to 4 days of unmet demand in the control villages (Table 1.20). The workfare program known as the National Rural Employment Guarantee Scheme (NREGS) offers manual labor jobs for a daily wage. The outside option it provides plays an important role in rural labor markets (Muralidharan et al., 2023b). However, while jobs are meant to be offered year-round, the scheme is often unavailable during periods of peak agricultural activity. Although noisier, I find no evidence of more NREGS workfare days offered in treatment villages using the program's administrative data (Table 1.21).

In the longer run, the goal would be for the displaced transplanters to find jobs in other sectors. Increasing the participation of women outside the agricultural sector is a major policy goal in India. In the second laborer phone survey I asked whether workers' husbands would let them work outside their villages. Interestingly, the overall effect appears to be husbands getting more restrictive (40% approving in treatment compared to 50% in control), but husbands

of women who do transplanting work each season become 9 percentage points relatively less restrictive (Table 1.22). The pattern grows stronger when the sample is restricted to only strata with high predicted uptake (Table 1.23). The likely mechanism is that husbands who are benefiting from greater farming profits in treatment villages can afford to become more restrictive as has been the general trend in India (Afridi et al., 2023), but husbands who primarily experience the treatment as negative wage effects become more willing for their wives to search for work in other locations. There are no significant treatment effects on women working outside their home villages, but it would likely take women multiple seasons to identify new job opportunities. Further, there is still a substantial amount of transplanting work in treatment villages. Note that there is no treatment effect on the extensive margin of whether women did any transplanting work (Table 1.22).

1.8 Model

In this section, I develop a model of agricultural labor markets. The model's primary purpose is to illustrate how differences in labor market structure would have led to different wage, technology adoption, and productivity outcomes from a policy shock similar to the experiment. In particular, the experiment is conceptualized as a reduction in the fixed cost of adopting mechanization. The experiment is used to estimate the baseline calibration of the model. Then, the model is used to assess how the same reduction in the fixed cost of adoption implied by the experiment would impact the uptake of mechanization and wages under different labor market conditions. These results help illustrate how labor market and mechanization policies interact.

1.8.1 Model Environment

The core of the model is a technology and labor supply decision made by farmers endowed with land L_i . Farmers decide whether to adopt drum seeding $a_i = 1$ or continue to transplant, conditional on their landholdings and the equilibrium wage for transplanting w_T^* . Workers characterized by a representative household decide how much transplanting labor to

supply t_i . Equilibrium is reached when the labor market for transplanting clears. The final share of land drum seeded, and the final transplanting wages are endogenous objects in the model, which I target for calibration. The treatment is modeled as a decrease in the fixed cost of accessing a drum seeder. The model is static since I only observe the short-run equilibrium in the experiment.

Farmer Technology Adoption Problem

Each farmer i maximizes its agricultural profits, which are conditional on the equilibrium w_T^* and land endowment L_i . The household decision variable is a_i , which is a binary variable indicating whether or not they adopt drum seeding. I do not allow for partial adoption since a given plot can only be planted in one uniform way for agronomic reasons. I also assume that individual farmers are small relative to the labor market and take the equilibrium wage w_T^* as given when making their technology choice. The profit maximization function is:

$$\begin{aligned} \max_{a_i} \Pi(a_i, w_T^*, L_i) &= pf(a_i, L_i) - (1 - a_i)w_T^*h(L_i) - a_i\alpha \\ \text{s.t. } a_i &\in \{0, 1\} \end{aligned} \tag{1.1}$$

Farmer's revenues are a function of the land they cultivate and the exogenous price received for the output p . Due to widespread government procurement and centralized milling, rice output markets are relatively integrated, and there is no evidence of output price effects in the experiment's setting. If farmers do not adopt drum seeders they pay $w_T^*h(L_i)$ where w_T^* is the equilibrium wage and $h(L_i)$ is the labor required to transplant their land L_i . Note that the labor demand is modeled as purely a function of land since practically farmers plant the entirety of a rice plot. Other inputs, such as water and fertilizer, would be wasted if the entire plot is not planted. Thus, labor demand can be thought of as inelastic to wages conditional on the farmer's choice of planting technology. If farmers adopt drum seeders, they pay a fixed cost α to acquire the device but do not have to pay the variable cost of hiring transplanting workers.¹⁶

¹⁶In practice, there are small variable costs associated with drum seeding, but these tend to be done by male

For the remainder of this section, I assume linear production and labor demand functions. The labor demand is linear because the task is the same per unit area regardless of the amount of land planted. While yields vary with plot size, the exact nature of this variation is highly context-dependent, and assuming a constant yield is the most conservative assumption.

Proposition 1: With linear production $f(a_i, L_i) = \beta L_i$ and labor demand $h(L_i) = \gamma L_i$, profits per acre $\frac{\Pi_i}{L_i}$ are increasing in land for mechanized farmers $a_i = 1$ and unchanged for unmechanized farmers $a_i = 0$.

Proof:

Case 1: $a_i = 1$

$$\frac{\Pi_i}{L_i} = p\beta - \frac{\alpha}{L_i}$$

$$\frac{d\frac{\Pi_i}{L_i}}{dL_i} = \frac{\alpha}{L_i^2} > 0$$

Case 2: $a_i = 0$

$$\frac{\Pi_i}{L_i} = p\beta - w_T^* \gamma$$

$$\frac{d\frac{\Pi_i}{L_i}}{dL_i} = 0$$

Given these linear functional forms profits per acre are increasing in L if a farmer drum-seeds ($a_i = 1$), but not changing in L if farmers transplant $a_i = 0$. This matches the empirical result that farmers with large landholdings are more likely to adopt drum seeders than farmers with small landholdings.

Proposition 2: The lower wages decrease the acreage under mechanization.

Proof: The marginal farmer who adopts will be the farmer whose cost of cultivation equals the fixed cost of adoption e.g. $w_T \gamma L_i^* = \alpha$. The mechanized acreage L_m is:

workers or family members who would expend a similar amount of effort overseeing a hired transplanting team

$$\int_{\frac{\alpha}{w_T \gamma}}^{\bar{L}} L(i) di = L_m$$

Since $L(i)$ is always positive L_m is increasing in the interval of integration I

$$I = \bar{L} - \frac{\alpha}{w_T \gamma}$$

$$\frac{\partial I}{\partial w_T} = \frac{\alpha}{w_T^2 \gamma} > 0 \Rightarrow \frac{\partial L_m}{\partial w_T} > 0$$

Household Labor Problem

The labor supply decision is taken by a representative household. The assumption is that providing transplanting labor to a household's own farm is no different than providing it to another farmer in the village. An alternative motivation would be a scenario where labor is supplied by households who do not have any land and simply decide whether to transplant or take on other tasks based on wages. The household's labor supply is modeled as the standard labor leisure decision with days of transplanting work t_i . The time constraint T is the total number of working days available during the transplanting period.

$$\max_{t_i} U_i(w_t, t_i) = u(w_t t_i) + v(T - t_i) \quad (1.2)$$

See Appendix 1.13.3 for a derivation of the labor supply elasticity assuming specific functional forms for $u()$ and $v()$ to match the labor supply equation of the form $\log w_t = \frac{1}{\eta} \log t + \log \alpha$ where α is a constant.

Market Clearing

The demand for transplanting labor must equal the amount of transplanting labor supplied for the transplanting labor market to clear. This market-clearing condition pins down the

transplanting wage.

$$LD = \int (1 - a_i)h(L_i)di = t_i = LS$$

Solving

Since the experiment only estimates the labor supply elasticity, the quantitative model uses a constant elasticity labor supply curve. The labor supply elasticity is defined as the percent change in labor supplied over the percent change in wage. Mathematically:

$$\frac{d \ln LS}{d \ln w_T} = \eta_T$$

where η_T is the labor supply elasticity

This implies a constant inverse labor supply equation of the form

$$\alpha LS^{\frac{1}{\eta}} = w_t$$

where $\log \alpha$ is a constant corresponding to to the lowest log wage where any labor would be supplied. This can be seen by rearranging the above equation and taking logs

$$\frac{1}{\eta} \log LS = \log w_t - \log \alpha$$

To simplify the problem, I assume farmers predict the adoption decisions of other farmers and solve their adoption problem based on a correct expectation of the equilibrium wage. Thus, the farmers solve backward, correctly guessing adoption and thus labor demand and wage before solving their own technology adoption problem. Qualitatively, farmers do report anticipating fewer wage spikes due to drum seeder adoption, but they can also learn about adoption and wages from segments of the village that plant and transplant relatively early.

Simulation and Calibration

First, some values are exogenous. The yield of 2000 kg per acre ($f(L_i) = 2000L_i$), the labor demand of 10 days per acre $h(L_i) = 10L_i$, and the rice output price of 2000 INR per 100 kg (p) are median values for the setting. There are 700 farmers with land sizes ranging from 1 to 8 acres in steps of .01 acres. The linear land distribution ensures the effect of wages on mechanization is consistent across different wage levels. The elasticity η_T I take directly from the experiment dividing the wage effect by the labor demand effect. The assumption is that labor demand is inelastic conditional on the farmers' choice of technology so the experiment traces the labor supply curve. The other parameters, reservation wage R , initial α_0 , and final α_1 , are calibrated by matching moments so that the baseline wage and uptake and the changes match values observed in the experiment. To calibrate α , I use the fact that for the marginal farmer $\alpha = w_T 10L_m$ which implies $L_m = \frac{\alpha}{10w_T}$. If there are N total farmers then adoption is $\frac{(N-m)}{N}$. Conditional on the values of α , I get an estimated labor supply at w_t which I can plug back into the labor supply curve $w_t = LS^{\frac{1}{\eta_T}} + R$ to find R .

The model is solved algorithmically by guessing wages and then iterating until wages reach a steady state. First, each farmer's drum seeder adoption decision is calculated based on a wage guess. Second, based on drum seeder adoption, calculate the total labor demand. Third, update the wages based on the calculated labor demand. Finally, iterate until a steady state is reached.

To simulate the experiment, solve the model for a higher drum seeder adoption cost α_0 . Then, solve the model again with a lower α_1 using the equilibrium wages in the first step as the initial guess. When calibrated, the model matches the experiment closely (Table 1.24).

Policy Simulation Results

I simulate the same reduction in adoption fixed costs under two alternate labor market conditions in Table 1.25. The row agriculture productivity is the total output from agriculture divided by the labor in agriculture. The row aggregate productivity is the total output from

agriculture and the output from the displaced labor at the equilibrium wage, assuming the wage is equivalent to the marginal product of labor for workers' outside option, which could include work within the household. The first column in the table simply replicates the experiment in the model. The second simulation provides a benchmark for the counterfactual where labor supply is so elastic there is no wage effect.

The third simulation explores what would happen if the labor supply was such that the wages were less responsive to the labor demand shock. This is analogous to expanding the effective labor market for women by reducing search, transport, and cultural frictions. Alternatively, this simulation can be interpreted as the government introducing a rental market for a technology that replaces male labor, which other research suggests has much more muted responses to labor demand shocks in the Indian context (Afridi et al., 2022). In this world, labor supplied to agriculture would reduce much more rapidly in response to small wage changes. When the labor supply curve is flat, the same decrease in adoption costs leads to greater increases in adoption and limited wage effects. In equilibrium, mechanization is profitable for more farmers since wages remain high, more workers are displaced, and those displaced workers are more productive since their outside wages are higher.

The forth simulation explores if labor supply was even less elastic. This would result in further wage declines, less technology adoption, and lower productivity.

The simulations illustrate the importance of considering labor supply alongside mechanization policy and targeting labor markets with low labor supply elasticities. Importantly, for the government to achieve its goal of improving aggregate productivity, workers need to be able to move into other productive jobs. Otherwise, in equilibrium, farmers will produce the same output with less labor, but the freed-up labor will not produce anything new.

1.9 Conclusion

By leveraging the agricultural extension system of a major rice-growing state in India, the study tests intervention that reduces barriers to mechanization. The study leverages local governments to improve farmers' access to drum seeders, which allows farmers to avoid the laborious and costly process of transplanting. While only a few farmers adopt the technology in the initial seasons covered by the study, these farmers save on transplanting costs and earn higher profits. Even farmers who do not drum seed earn higher profits in treatment villages.

The reason non-adopters also profited is that wages were lower in treatment villages. The wage effect suggests a relatively steep labor supply curve for female transplanting labor. Because rice production is restricted by farmers' land and water access and labor requirements for transplanting are fixed, the shift in employment can be interpreted as a shift in labor demand. There is no evidence of women transitioning into other industries in the short run, but there is some evidence of women attempting to access, without success, more work from India's workfare scheme. This shows how the labor market, in the short-run, was not able to absorb the labor displaced by mechanization. The wage evidence, alongside workers' greater unmet demand for jobs, suggests that workers struggled to transition to other productive opportunities.

The period of this paper is too short to know whether the displaced women will eventually find potentially higher-paying jobs in other sectors, but in the short run, the rental scheme appears to have primarily translated into farmer profits rather than aggregate productivity growth. In future research, I aim to study potential interventions to facilitate women's finding new jobs in neighboring villages. Research in Italy studying the historical transition from manual transplanting in the 1950s suggests that in the long run, rice regions had higher female labor force participation and female political participation (Bagnato et al., 2023).

Political backlash to mechanization has been common throughout history. Mechanizing gendered agricultural tasks is a unique case where the winners and losers from mechanization are often in the same household. In my sample, only 22% of women in farming households

did not engage in agricultural work. While increasing female labor force participation is an important policy goal, the types of jobs policymakers envision are not rice transplanting. The experiment highlights the importance of identifying ways for women to stay in the labor force as they transition out of agriculture. Evidence from the experiment suggests husbands' opinions on where women can work shift along with mechanization and that women do look for new opportunities at the time of mechanization. Thus, pairing programs that target female labor force participation with programs targeting mechanization could be a promising solution. In the shorter run, governments can also ensure their workfare schemes continue to operate during peak agriculture seasons when they were historically less relevant.

In general, the study highlights how labor market frictions can sabotage productivity growth. When workers are trapped in a task, labor supply curves will be steeper, and the wage effects of mechanization will be more pronounced. These wage effects hurt both equity and efficiency, reducing worker incomes, the extent of mechanization, and the potential increase in aggregate productivity.

Social protection programs can blunt the wage and distributional effects of mechanization, but workers' job opportunities need to improve to fully realize the potential of agricultural mechanization to increase a country's aggregate productivity.

For developing countries trying to spur structural transformation, combining efforts to reduce labor market frictions and barriers to technology adoption might be a promising path forward. If displaced workers cannot find new jobs, the mechanization will become less profitable and slower. Sustained productivity growth through mechanization requires the labor market to absorb the displaced labor. This study highlights how mechanizing one of the least productive tasks in Indian agriculture, rice transplanting, displaced workers and reduced the profitability of further mechanization in the short run. Government intervention to give transplanters more employment options would not only benefit those vulnerable women but also accelerate productivity growth and structural transformation.

1.10 Chapter Acknowledgments

Chapter 1, in full, is currently being prepared for submission for publication of the material. Brownstone, Steven; The dissertation author was the primary investigator and author of this material.

1.11 Figures




Figure 1.1. Reasons Farmers Stop Drum Seeding

Farmers who reported having drum seeded in the past, but did not drum seed the season they were surveyed.




Figure 1.2. Distribution of Uptake Across Villages




Figure 1.3. Main Samples




Figure 1.4. Timeline of data collection




Figure 1.5. Comparing Wage Distribution of Different Tasks




Figure 1.6. Comparing Wage Distribution of Different Tasks



Figure 1.7. Transplanting in Telangana



Japan 1900



Italy 1948



India Now

Figure 1.8. Transplanting Through Time




Figure 1.9. Project Villages

Orange pins represent treatment and blue represent control




Figure 1.10. CDF of Rural Household's Agricultural Land Owned by Women's Primary Occupation in DHS Survey in Study Districts




Figure 1.11. Distribution of Household Income Effects

1.12 Tables

Table 1.1. Balance Table

Variable	(1)		(2)		(2)-(1)	
	N/Clusters	Mean/(SE)	N/Clusters	Mean/(SE)	Pairwise t-test	Mean difference
Total Admin 2022 Cult	796	1.751	1065	1.819	1861	0.068
	134	(0.065)	200	(0.060)	334	
Total Land Owned	877	4.977	1147	4.698	2024	-0.278
	138	(0.221)	204	(0.168)	342	
Household Size	850	3.922	1108	3.961	1958	0.039
	135	(0.072)	200	(0.063)	335	
Well Irrigation	877	0.800	1147	0.776	2024	-0.025
	138	(0.024)	204	(0.023)	342	
HH Head Secondary and Above	850	0.376	1108	0.366	1958	-0.011
	135	(0.018)	200	(0.016)	335	
F-test of joint significance (F-stat)						0.460
F-test, number of observations						1813
F-test, number of clusters						327

Fixed effect used in pairwise and f-test regressions: [strata_id]. Significance: ***=.01, **=.05, *=.1. Errors are clustered at variable: [panchayatid]. Note that with the exception of the administrative data on 2022 cultivation all variables come from the initial survey which took place after the start of some implementation

Table 1.2. Sample NFHS/DHS Comparison

Variable	(1)		(2)		(2)-(1)	
	N/Clusters	Mean/(SE)	N/Clusters	Mean/(SE)	Pairwise t-test	Mean difference
Acres Owned	2024	4.819	2219	4.733	4243	-0.086
	342	(0.135)	249	(0.441)	591	
HH Head Secondary and Above	1958	0.370	4114	0.431	6072	0.060***
	335	(0.012)	252	(0.014)	587	
Irrigation Pump	1958	0.477	4115	0.466	6073	-0.011
	335	(0.018)	252	(0.017)	587	
Pressure Cooker	1958	0.511	4115	0.561	6073	0.050**
	335	(0.016)	252	(0.013)	587	
Television	1958	0.877	4115	0.882	6073	0.005
	335	(0.008)	252	(0.006)	587	
Has Mobile	2024	0.954	4115	0.968	6139	0.014
	342	(0.012)	252	(0.004)	594	
Has Cow or Buffalo	2024	0.398	4115	0.212	6139	-0.186***
	342	(0.016)	252	(0.011)	594	
F-test of joint significance (F-stat)						3.199***
F-test, number of observations						4176
F-test, number of clusters						584

. Significance: ***=.01, **=.05, *=.1. Errors are clustered at variable: [panchayatid].

Table 1.3. Sample NFHS/DHS Comparison Only Farmers

Variable	(1) Farmer Sample		(2) DHS Ag Land Owners in Study Districts		(2)-(1) Pairwise t-test	
	N/Clusters	Mean/(SE)	N/Clusters	Mean/(SE)	N/Clusters	Mean difference
Acres Owned	2024	4.819	2219	4.733	4243	-0.086
	342	(0.135)	249	(0.441)	591	
HH Head Secondary and Above	1958	0.370	2218	0.349	4176	-0.022
	335	(0.012)	249	(0.013)	584	
Irrigation Pump	1958	0.477	2219	0.458	4177	-0.019
	335	(0.018)	249	(0.020)	584	
Pressure Cooker	1958	0.511	2219	0.531	4177	0.021
	335	(0.016)	249	(0.014)	584	
Television	1958	0.877	2219	0.890	4177	0.013
	335	(0.008)	249	(0.008)	584	
Has Mobile	2024	0.954	2219	0.972	4243	0.018
	342	(0.012)	249	(0.005)	591	
Has Cow or Buffalo	2024	0.398	2219	0.347	4243	-0.050**
	342	(0.016)	249	(0.014)	591	
F-test of joint significance (F-stat)						3.199***
F-test, number of observations						4176
F-test, number of clusters						584

Significance: ***=.01, **=.05, *=.1. Errors are clustered at variable: [panchayatid].

Table 1.4. Sample NFHS/DHS Comparison Only Rural

Variable	(1) Farmer Sample		(2) DHS Rural Sample in Study Districts		(2)-(1) Pairwise t-test	
	N/Clusters	Mean/(SE)	N/Clusters	Mean/(SE)	N/Clusters	Mean difference
Acres Owned	2024	4.819	1971	4.443	3995	-0.376
	342	(0.135)	194	(0.395)	536	
HH Head Secondary and Above	1958	0.370	1970	0.314	3928	-0.056***
	335	(0.012)	194	(0.012)	529	
Irrigation Pump	1958	0.477	1971	0.449	3929	-0.029
	335	(0.018)	194	(0.022)	529	
Pressure Cooker	1958	0.511	1971	0.510	3929	-0.001
	335	(0.016)	194	(0.014)	529	
Television	1958	0.877	1971	0.884	3929	0.006
	335	(0.008)	194	(0.008)	529	
Has Mobile	2024	0.954	1971	0.970	3995	0.016
	342	(0.012)	194	(0.006)	536	
Has Cow or Buffalo	2024	0.398	1971	0.377	3995	-0.021
	342	(0.016)	194	(0.014)	536	
F-test of joint significance (F-stat)						3.385***
F-test, number of observations						3928
F-test, number of clusters						529

Significance: ***=.01, **=.05, *=.1. Errors are clustered at variable: [panchayatid].

Table 1.5. Treatment Effect Village Engagement

	(1) Leader Gives Ag Advice	(2) Leader Refers to Extensionist	(3) Drum Seeder in Village Meeting
Treatment	0.05** (0.02)	0.07*** (0.02)	0.16*** (0.02)
Observations	1,879	1,939	1,890
Ctrl Average	0.38	0.49	0.19

Standard errors in parentheses. Strata and enumerator fixed effects. Standard errors clustered at the village level. Column (1) is farmer report village elected leader discussed agriculture. Column (2) is farmers reported village leaders telling them to contact extension agents. Column (3) is drum seeders mentioned at the official village meeting

Table 1.6. Treatment Effect on Awareness

	Ag Officer Info	Attend Demo	Heard Promotion
Extensionist Arm	0.06** (0.03)	0.10*** (0.03)	0.16*** (0.03)
Leader Arm	0.05 (0.03)	0.07** (0.03)	0.17*** (0.03)
Control Mean	0.43	0.49	0.32
R-squared	0.27	0.12	0.13
Observations	2205	1946	1893
Difference p-value	0.622	0.402	0.801

*** $p < .01$, ** $p < .05$, * $p < .1$ Standard errors clustered at the village level. Includes fixed effects for randomization strata. Ag Officer information means extensionist listed as source of agricultural advice in last month. Attend demo means the farmer ever attended a demonstration of drum seeders. Heard promotion means the farmer heard drum seeders promoted in their village. Don't know coded as missing.

Table 1.7. Farmer Drum Seeder Specific Knowledge

	(1) Irrig Info	(2) Weed Info
Treatment	0.08*** (0.02)	0.04*** (0.01)
Observations	2,023	2,023
Ctrl Average	0.19	0.06
Controls	N	N

Standard errors clustered at the community level. Enumerator and Strata FE. Column 1 is knows to delay irrigation post planting. Column 2 is knows to reflood to control weeds

Table 1.8. Treatment Effect on Uptake

	(1) Drum Wet	(2) Drum Dry	(3) Acres Drum Wet	(4) Acres Dry
Leader Arm	0.03*** (0.01)	0.02** (0.01)	0.10** (0.04)	0.05 (0.04)
Extensionist Arm	0.04*** (0.01)	0.03*** (0.01)	0.13*** (0.04)	0.07* (0.04)
Observations	2,023	2,023	2,023	2,023
Ctrl Average	0.01	0.01	0.04	0.05
Controls	N	N	N	N

Columns (1-2) are indicator variables for any drum seeding. Columns (3-4) are acres drum seeded. Standard errors in parentheses. Strata and enumerator fixed effects. Standard errors clustered at the village level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.9. Treatment Effect on Uptake

	(1) Drum Wet	(2) Drum Dry	(3) Acres Drum Wet	(4) Acres Dry
Treatment	0.04*** (0.01)	0.03*** (0.01)	0.11*** (0.03)	0.06** (0.03)
Observations	2,023	2,023	2,023	2,023
Ctrl Average	0.01	0.01	0.04	0.05
Controls	N	N	N	N

Columns (1-2) are indicator variables for any drum seeding. Columns (3-4) are acres drum seeded. Standard errors in parentheses. Strata and enumerator fixed effects. Standard errors clustered at the village level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.10. Who Adopts Drum Seeders in Wet Season

	(1)	(2)
	Acres Drum Wet	Acres Drum Wet
Total Cultivation	0.05*** (0.01)	0.06*** (0.01)
Saline Soil	-0.16 (0.11)	-0.24** (0.10)
HH Head Secondary	0.14** (0.06)	0.14** (0.06)
HH Head Higher Secondary	0.14* (0.07)	0.17** (0.07)
Well Irrigation	0.13** (0.06)	0.05 (0.06)
Observations	1,926	1,926
Ctrl Average	0.04	0.04
Fixed Effects	Village	Strata

Column (1) focuses on variation within village and Column (2) looks at variation across villages but within strata. HH Head means household head's total completed education and is coded as a binary variable. Total cultivation is administrative baseline data on rice cultivation. Standard errors in parentheses. Standard errors clustered at the village level. sym* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.11. What Predicts Yield for Drum Seeding Farmers

	Drum Seeder Yield
Knows to Delay Irrigation Post Planting	3.76 (2.56)
Knows to Reflood field to control weeds	3.56** (1.66)
Clay Soil	-5.50** (2.44)
Saline Soil	-1.45 (4.20)
total land	0.17 (0.26)
HH Head Secondary	2.75 (2.24)
HH Head Higher Secondary	0.82 (2.66)
Well Irrigation	-3.75 (2.53)
Observations	125
Ctrl Average	19.47

Standard errors in parentheses. Standard errors clustered at the village level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Data comes from only drum seeding farmers and includes a supplemental sample not used in the primary regressions.

Table 1.12. Farmers Hiring Decisions Normalized by Area Farmed

	(1) Workers Hired per Acre	(2) Workers Hired per Acre
Treatment	-0.72** (0.31)	-0.72** (0.32)
Observations	1,995	2,020
Ctrl Average	9.63	9.63
Controls	Y	N

Standard errors clustered at the community level. Enumerator and Strata FE. Total female workers hired per acre cultivated across all tasks. Excludes family work. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.13. Treatment Effect on Costs per Acre

	(1) log total labor expenditure	(2) total family days worked	(3) log total non-labor expenditure
Treatment	-0.10*** (0.03)	-0.02 (0.07)	0.03 (0.03)
Observations	1,967	2,023	1,789
Control Mean	8.9	1.5	9.4
Controls	N	N	N

Don't know coded as missing for expenditure variables

Table 1.14. Treatment Effect on Ln Revenue Per Acre

	(1) Paddy Rev.	(2) Total Rev.	(3) Impute Paddy Rev.	(4) Impute Total Rev.
Treatment	0.02 (0.02)	0.01 (0.02)	0.03* (0.02)	0.04** (0.02)
Observations	1,697	1,715	1,974	1,923
Ctrl Average	10.7	10.6	10.6	10.5
Controls	N	N	N	N

Actual revenue is reported sales (1-2) while imputed paddy revenue is production times district median price. The small increases are likely driven by the effects of small investments enabled by drum seeders, such as slightly better seeds. Strata and enumerator fixed effects. Standard errors clustered at the village level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.15. Treatment Effect on Paddy Yields and Acres

	(1) Paddy Yield	(2) Acres Paddy	(3) Acres Not Paddy	(4) Any Not Paddy
Treatment	0.59 (0.38)	0.12 (0.14)	-0.07 (0.06)	-0.03* (0.01)
Observations	2,023	2,023	2,023	2,023
Ctrl Average	18.3	3.3	0.6	0.2
Controls	N	N	N	N

Standard errors in parentheses. Standard errors clustered at the village level. Enumerator and Strata FE. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Yields and acres aggregate across all types of rice if farmer planted rice in multiple ways. Acres and any not paddy refer to any crop that is not rice

Table 1.16. Treatment Effect on Profit per Acre in USD

	(1)	(2)	(3)	(4)
	Paddy Prof	Total Prof	Winsor Paddy Prof	Winsor Total Prof
Treatment	20.28** (9.24)	15.15* (8.88)	17.70** (8.19)	13.25* (7.81)
Observations	1,844	1,895	1,844	1,895
Control Mean	216.11	246.98	215.87	245.76

Standard errors in parentheses. Standard errors clustered at the village level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Strata and enumerator fixed effects. Winsorized results winsorized at 5th and 95th Percentiles. Total aggregates across all crops and paddy is just rice. Imputed revenues used to calculate the profits.

Table 1.17. Main Treatment Effect on Transplanting Wages

	(1)	(2)
	Wage per Day	Wage per Acre
Treatment	-17.05** (8.54)	-335.84** (137.32)
Observations	895	418
Ctrl Average	557.43	5,776.75

Wage data from farmers hiring transplanting labor. Column (1) are wages based on time and column (2) are piece rate wages based on productivity. There is no treatment effect on wage type, but farmers with more land tend to pay piece rate wages. Standard errors in parentheses. Standard errors clustered at the village level. Enumerator and Strata FE. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.18. Treatment Pooled Wages

	(1)	(2)
	Reported Daily	Log Wage
Treatment	0.03 (0.02)	-0.03** (0.01)
Reported Wage Day		-2.18*** (0.03)
Observations	1,332	1,332
Ctrl Average	0.64	7.12

Standard errors in parentheses

Standard errors clustered at the village level. Enumerator and Strata FE.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.19. Labor Costs for Farmers who Transplant

	(1)	(2)
	Ln Labor Cost Trans Only	Days Hired Trans Only
Treatment	-0.07** (0.03)	-0.42 (0.33)
Observations	1,919	1,965
Ctrl Average	8.94	9.63

Standard errors in parentheses

Standard errors clustered at the village level. Enumerator and Strata FE.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.20. Transplanters in Treatment Villages Have Greater Unmet Demand for NREGA

	(1)	(2)
	Transplanted at Baseline	Unmet Demand
Treatment	0.01 (0.04)	-1.12 (0.85)
Transplanted X Treatment		1.93** (0.91)
Observations	553	553
Ctrl Average	0.7	4.6

Standard errors in parentheses. Data from phone survey of workers on NREGA rolls. Standard errors clustered at the community level. Column (1) is a dummy for whether the worker engages in transplanting every season at baseline since many workers only transplant a few days. Column (2) is the number of worker days the worker wanted minus the number of days they worked. Unmet demand is preferred since demand can appear to spike if the government exogenously starts a new project not considering local labor market conditions

Table 1.21. No Effect on Administrative NREGA Days Offered

	(1) Person Days July-October	(2) Person Days August
Treatment	116.09 (83.81)	-4.56 (26.61)
Lagged Dry Season NREGS	0.27*** (0.06)	0.06*** (0.02)
Lagged Wet Season	0.43*** (0.06)	0.08*** (0.02)
Double Lagged Wet Season	0.10** (0.04)	0.03** (0.01)
Observations	361	361
Control Mean	1867	466

Standard errors in parentheses. Standard errors clustered at the village level. Strata Fixed Effects.* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Data source is the NREGA program website

Table 1.22. Treatment Effect on Husband Allows to Work Outside

	(1) Trans at Baseline	(2) Work Outside OK
Treatment	0.00 (0.03)	-0.10*** (0.03)
Transplant at Baseline		0.06* (0.04)
Interaction		0.09* (0.05)
Observations	1,338	1,338
Ctrl Average	0.3	0.5

Standard errors in parentheses. Data from phone survey of workers on NREGA rolls. Standard errors clustered at the community level. Enumerator and Strata FE. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Column (1) is a dummy for whether the worker engages in transplanting every season at baseline since many workers only transplant a few days. Column (2) is whether women report that their husband allows them to work outside.

Table 1.23. Treatment Effect on Husband Attitudes in Higher Predicted Uptake Strata

	(1)	(2)
	Trans at Baseline	Work Outside OK
Treatment	-0.01 (0.03)	-0.14*** (0.04)
Transplant at Baseline		0.03 (0.05)
Interaction		0.19*** (0.07)
Observations	879	879
Ctrl Average	0.3	0.5
High Pred. Uptake Strata	Y	Y

Standard errors in parentheses

Standard errors clustered at the community level. Enumerator and Strata FE.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.24. Calibration Comparison

	Baseline Calibration	Experimental Data
Inverse Wage Elasticity η	0.86	0.86
Initial Adoption	1.14%	1.00%
Final Adoption	4.56%	4.00%
Initial Wage	557	557
Final Wage	525	523

Table 1.25. Model Simulation Results

	Baseline	No Wage Channel	Low Elas	High Elas
Change in Fixed Cost α	-8.59%	-8.59%	-8.59%	-8.59%
Elasticity η	0.86	0.07	0.43	1.30
Initial Adoption	1.14%	1.14%	1.14%	1.14%
Final Adoption	4.56%	9.56%	5.99%	3.71%
Wage Change	-5.77%	-1.28%	-4.42%	-6.41%
Ag Prod Change	6.06%	14.56%	8.53%	4.56%
Total Prod Delta	7.79%	19.62%	11.12%	5.83%

1.13 Appendix

1.13.1 Additional Sample

In addition to the primary sample, there are two supplemental samples. The first supplemental sample is 145 farmers listed on drum seeder rental records as a method for validating that administrative data. However, the comparison between implementation arms is not the focus of this paper.

The second supplemental sample is 186 farmers who were individually randomized to receive further cash incentives to encourage drum seeder adoption. Unfortunately, due to issues processing the payments this experiment did not have a large enough first stage to yield meaningful analysis.

However, both of these samples have more drum-seeding farmers than the main representative sample, so they are used in some auxiliary regressions exploring determinants of drum-seeding yield and profits.

1.13.2 Pooled Wage Specification

To pool wages across reporting types I use the following specification:

$$\ln(Wage_{ivse}) = \beta_1 T_g + \beta_2 D_i + \beta_3 L_i + \chi_e + \gamma_s + \varepsilon_v$$

$Wage_{ivse}$ wage as either cost per acre or cost per worker per day to transplant a field. Both forms of wage are standardized by their standard deviations to ensure both types of wages have a similar scale. T_v treatment dummy for village v . D_i indicator for how wage is reported for individual i . L_i acres of land owned by individual i . This is the best predictor of wage type with larger farmers paying by acre. χ_e enumerator fixed effects. γ_s strata fixed effects. ε_{vi} are errors clustered by village. In this specification, β_1 is the percent change in wage induced by the treatment.

1.13.3 Utility Functional Forms

Recall the basic labor supply problem where w_t is the transplanting wage and n is hours worked:

$$\max_n U_i(c, n) = u(c) - v(n) \quad (1.3)$$

Now we choose a functional form and rewrite this as:

$$U_i(c, n) = c - R \frac{n^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}}$$

subject to $c = w_t n$

The first-order conditions are

$$1 = \lambda$$

$$R n^{\frac{1}{\eta}} = \lambda w_t$$

The labor supply is:

$$R n^{\frac{1}{\eta}} = w_t$$

$$n^{\frac{1}{\eta}} = \frac{w_t}{R}$$

Taking logs, we get

$$\frac{1}{\eta} \log n = \log \frac{w_t}{R}$$

$$\frac{1}{\eta} \log n = \log w_t - \log R$$

$$\log n = \eta \log w_t - \eta \log R$$

Now we can directly find the elasticity:

$$\frac{d \ln n}{d \ln w_T} = \eta$$

Bibliography

Acemoglu, Daron, Fredric Kong, and Pascual Restrepo, “Tasks At Work: Comparative Advantage, Technology and Labor Demand,” Working Paper 32872, National Bureau of Economic Research August 2024.

Afridi, Farzana, Kanika Mahajan, and Nikita Sangwan, “The gendered effects of droughts: Production shocks and labor response in agriculture,” *Labour Economics*, 2022, 78, 102227.

—, **Monisankar Bishnu, and Kanika Mahajan**, “Gender and mechanization: Evidence from Indian agriculture,” *American Journal of Agricultural Economics*, 2023, 105 (1), 52–75.

Aker, Jenny C. and B. Kelsey Jack, “Harvesting the Rain: The Adoption of Environmental Technologies in the Sahel,” *The Review of Economics and Statistics*, 12 2023, pp. 1–52.

Ali, Akhter, Imtiaz Hussain, Dil Bahadur Rahut, and Olaf Erenstein, “Laser-land leveling adoption and its impact on water use, crop yields and household income: Empirical evidence from the rice-wheat system of Pakistan Punjab,” *Food Policy*, 2018, 77, 19–32.

Bagnato, Luca, Pan Yu, and Miriam Venturini, “Agricultural Practices, Organized Workers and Female Empowerment: Evidence from Italian Mondine,” Technical Report December 2023.

Bassi, Vittorio, Raffaella Muoio, Tommaso Porzio, Ritwika Sen, and Esau Tugume, “Achieving Scale Collectively,” *Econometrica*, 2022, 90 (6), 2937–2978.

Breza, Emily, Supreet Kaur, and Yogita Shamdasani, “Labor Rationing,” *American Economic Review*, October 2021, 111 (10), 3184–3224.

Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli, “Agricultural Productivity and Structural Transformation: Evidence from Brazil,” *American Economic Review*, June 2016, 106 (6), 1320–1365.

Caunedo, Julieta and Elisa Keller, “Capital Obsolescence and Agricultural Productivity*,” *The Quarterly Journal of Economics*, December 2020, 136 (1), 505–561.

- **and Namrata Kala**, “Mechanizing Agriculture,” Working Paper 29061, National Bureau of Economic Research July 2021.
- Conley, Timothy G. and Christopher R. Udry**, “Learning about a New Technology: Pineapple in Ghana,” *American Economic Review*, March 2010, 100 (1), 35–69.
- Deutschmann, Joshua W, Maya Duru, Kim Siegal, and Emilia Tjernström**, “Can Smallholder Extension Transform African Agriculture?,” Working Paper 26054, National Bureau of Economic Research July 2019.
- Foster, Andrew D. and Mark R. Rosenzweig**, “Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture,” *Journal of Political Economy*, 1995, 103 (6), 1176–1209.
- **and —**, “Are There Too Many Farms in the World? Labor Market Transaction Costs, Machine Capacities, and Optimal Farm Size,” *Journal of Political Economy*, 2022, 130 (3), 636–680.
- Gulati, Ashok and Ritika Juneja**, “Farm Mechanization in Indian Agriculture with Focus on Tractors,” *SSRN Electronic Journal*, 2020.
- Gulati, Kajal, Patrick S. Ward, Travis J. Lybbert, and David J. Spielman**, “Intrahousehold preference heterogeneity and demand for labor-saving agricultural technology,” *American Journal of Agricultural Economics*, 2024, 106 (2), 684–711.
- Hornbeck, Richard and Suresh Naidu**, “When the Levee Breaks: Black Migration and Economic Development in the American South,” *American Economic Review*, March 2014, 104 (3), 963–90.
- Jones, Maria, Florence Kondylis, John Loeser, and Jeremy Magruder**, “Factor Market Failures and the Adoption of Irrigation in Rwanda,” *American Economic Review*, July 2022, 112 (7), 2316–2352.
- Kaur, Supreet**, “Nominal Wage Rigidity in Village Labor Markets,” *The American Economic Review*, 2019, 109 (10), 3585–3616.
- Khan, Md. Tajuddin, Avinash Kishore, and P. Joshi**, “Gender dimensions on farmers’ preferences for direct-seeded rice with drum seeder in India,” Working Paper 01550, IFPRI 08 2016.
- Magnan, Nicholas, David J. Spielman, Travis J. Lybbert, and Kajal Gulati**, “Leveling with friends: Social networks and Indian farmers’ demand for a technology with heterogeneous benefits,” *Journal of Development Economics*, 2015, 116, 223–251.
- Manuelli, Rodolfo E. and Ananth Seshadri**, “Frictionless Technology Diffusion: The Case of

- Tractors,” *The American Economic Review*, 2014, 104 (4), 1368–1391.
- Muralidharan, Karthik**, *Accelerating India’s development: A state-led roadmap for effective governance*, London, England: Viking, April 2024.
- , **Paul Niehaus**, and **Sandip Sukhtankar**, “General Equilibrium Effects of (Improving) Public Employment Programs: Experimental Evidence From India,” *Econometrica*, 2023, 91 (4), 1261–1295.
- , —, and —, “General Equilibrium Effects of (Improving) Public Employment Programs: Experimental Evidence From India,” *Econometrica*, 2023, 91 (4), 1261–1295.
- NDTV**, “Bayer, GenZero, Shell, And Mitsubishi’s Good Rice Alliance Enrolls Over 10,000 Farmers In India,” *NDTV Profit*, October 2024.
- Planning Department**, “Telangana Socio Economic Outlook 2023,” Technical Report, Government of Telangana 2023.
- Reddy, Chiranjeevi, V Ratna Reddy, T Chiranjeevi, and M Srinivasa Reddy**, “Rythu Bharosa Kendras of Andhra Pradesh,” *Economic and Political Weekly*, November 2023, 58 (47).
- Restuccia, Diego, Dennis Tao Yang, and Xiaodong Zhu**, “Agriculture and aggregate productivity: A quantitative cross-country analysis,” *Journal of Monetary Economics*, 2008, 55 (2), 234–250.
- Smith, Pete, Dave Reay, and Jo Smith**, “Agricultural methane emissions and the potential formitigation,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 2021, 379 (2210), 20200451.
- Vent, Olivia, Sabarmatee, and Norman Uphoff**, “International - The System of Rice Intensification and Its Impacts on Women : Reducing Pain, Discomfort, and Labor in Rice Farming While Enhancing Households’ Food Security,” in A.J. Fletcher and W. Kubik, eds., *Women in Agriculture Worldwide*, Taylor & Francis, 2016, pp. 55–68.

Chapter 2

The Long Shadow of Feudalism: Concentration of Land and Labor Market Power in India

2.1 Introduction

Land ownership has historically been the main source of political power and control in agrarian economies. Despite seventy-five years of state-led efforts to dismantle feudal legacies and reduce inequality in India, progress has been uneven and incomplete. A vast literature finds that historical institutions have lasting and persistent impacts on a wide range of contemporary outcomes (Genicot and Ray, 2017; Dell, 2010; Kim and Wang, 2024) including land tenure in India (Banerjee and Iyer, 2005; Banerjee et al., 2002). A core driver of rural inequality is the labor market power of large landowners (Hornbeck and Naidu, 2014). Both the emergence of economic opportunities outside agriculture both through economic growth and social welfare policies (Muralidharan et al., 2023) should reduce inequality. We move beyond simply documenting persistent inequality to understanding what parts of the labor market and social programs are most distorted by land concentration.

In this paper, we use historic boundaries of feudal land institutions within the South Indian princely state of Hyderabad, active between 1724 and 1948, to compare present-day labor market and service delivery outcomes between nearby villages in the same modern districts with different concentrations of land ownership. We use granular, village-level variation to first illustrate where development gaps continue to persist and then highlight specific areas of the agricultural labor market and social welfare implementation where the disparities are concentrated.

Evidence from our setting suggests that while moderate increases in the number of small farms in non-feudal areas relative to feudal areas do not meaningfully affect yields or overall labor demand, agricultural wages for women are significantly lower in non-feudal areas. In this setting, gender norms restrict women’s labor market opportunities outside agriculture since gender norms restrict them from working outside their villages. Thus space continues to limit women’s economic opportunity compared to men whose mobility allows economic growth elsewhere to put upward pressure on even village agricultural wages.

We also find evidence that political power is more concentrated in villages with fewer small farms. We find political leaders in these villages reduce workfare job provision but only during periods of agriculture demand. These results are surprising under the priors of a perfectly competitive market, where even large farms do not have the ability to set wages. However, in this setting, collusion across larger landowners based on caste networks has long been documented (Anderson et al., 2015). We find that this plays an active role in our setting, where elected officials share deeply entrenched caste-based networks with the largest landowners in their villages.

This paper relies on variation in land tenure within the princely state of Hyderabad, which encompassed much of modern day south central India, covering more than 72 million people as of 2025. While some areas were directly taxed by the erstwhile ruler (the *Nizam*), other areas were feudal estates (*Jagirs*) established primarily in the early 19th century. These feudal areas, scattered across the state, never developed a land administration system and thus small farmers were unable to take advantage of land reforms in the 1950s which were made in response to a peasant revolution in 1947-48, focused on the exploitative nature of large landowners (Dhanagare, 1974; Sethi, 2023) ¹. The *Nizam* himself was reported to be the wealthiest person in the world in 1937 (TIME, 1937). A survey of comparable villages on either side of the former feudal borders in 1954 found that small farmers were much less likely to get tenure in the formerly feudal areas (Khusro, 1958). We georeferenced detailed maps from the colonial period to link regions of differential land tenure systems with modern-day administrative maps. We then compare present-day villages on either side of these borders in a spatial regression discontinuity framework, within narrow bandwidths (what we will refer to as “feudal” and “non-feudal”), exploiting variation *within* present-day administrative districts such that contemporary administrative quality is held constant and we are able to isolate village-level effects. We use publicly-available administrative data and data from our own primary survey from a subset of sampled villages to study a range of labor market and service delivery outcomes.

¹This movement is considered the antecedent to the Maoist movement in India which continued active violence well into the 2010s.

We first show that there are lasting and meaningful differences in the land structure across feudal and non-feudal areas, reflecting stickiness in land markets since the post-independence reforms in the 1950s (Bolhuis et al., 2021). As of 2022, land parcels are larger in formerly feudal areas. In particular, the 10th percentile of landholdings in feudal areas is 16% larger than the 10th percentile of landholdings in non-feudal areas just across the border. However, there is no difference in landholdings at the right tail of the distribution: the 90th percentile across feudal and non-feudal areas holds similar amounts of land. Both the total size of cultivated area, and the number of landowners, in feudal and non-feudal areas is similar. This indicates that the distribution of landholdings is shifted to the right in feudal areas (particularly at the bottom of the distribution), and is more compressed. This evidence is consistent historical surveys showing small farmers in feudal areas were unable to register their land and thus that land got divided among larger farmers (Khusro, 1958). We also find no differences in the stock or flow of land transactions across feudal and non-feudal areas, and find no differential movement of populations into or out of feudal regions over the previous three census rounds in aggregate and within population sub-groups. In order to test the validity of our spatial regression discontinuity design, we test it on a range of time invariant and geographic outcomes, and show that these vary smoothly at around the border.

The detailed nature of the spatial variation in our setting allows us to compare villages that have been part of the same district-level institutions and economies for the past seventy years. Unlike other settings where migration between across historical borders is a major threat to identification, the primary destination for economic migration is towns and cities far from the boundaries we study. To the extent these areas have economically integrated over time through trade and migration, that should attenuate our results towards zero. Villages on either side of the border have similar aggregate levels of overall development and poverty once we account for fixed effects at the higher administrative jurisdictions. Differences in some aspects of public goods provision persist, but the mechanism for these differences beyond persistence is difficult to pin down. On the other hand we can leverage more detailed detail on agricultural wages and a

government job scheme to better isolate the effect of land concentration on labor markets.

We select over 2,000 villages from among the universe of about 20,000 villages in our sample area within the spatial RD bandwidth where we conduct a primary survey. These telephonic surveys are conducted with elected representatives and workers, collecting data on outputs (yield), as well as inputs (labor demand and wages). Importantly, the latter are disaggregated by task type, with a particular focus on the disaggregation of tasks by gender.² Moreover, we are always able to account for differences in crop choice and seasonality. We supplement analysis of this data with a range of publicly-available datasets on service delivery commonly analyzed by other paper in this literature.

We find that feudal areas have lower wages for agricultural laborers in general, but the magnitudes of these difference are large enough to support statistical significance only among female agricultural laborers. Females working on relatively more “unskilled” tasks are paid 8% lower wages in feudal areas than non-feudal areas that are separated by short distances on either side of border, and the difference in the treatment effects across genders is significant at conventional levels.³ Importantly, we see these differences in the realized wages while seeing strong null effects on the labor demand and yield, indicating that the wage effects do not reflect differences in productivity or other aggregate differences in agricultural conditions across feudal and non-feudal areas. We also do not see any differences in aggregate labor demand that could explain these wage effects. We find that employers typically hire the same number of workers, in total and specifically on the gender-specific tasks on which we see wage effects. We also find that cultivation is equally productive in feudal and non-feudal areas. Notably, all our main results are always conditional on the crop and season, meaning the effects we find shut down any variation from the choice of crop and variation at the seasonal level.⁴ Another influential

²The key motivation for this granular focus is that agricultural labor markets are segmented by gender, among other cleavages (Brownstone, 2025).

³Our preferred bandwidth is 20 kilometers, which should be expected to take roughly one hour on a bicycle to traverse, and is thus conceptualized by us as an arbitrageable distance.

⁴In general, we do not find significant differences in extensive margin crop choice. Regardless of feudal status, crop choice very similar at the border segment level, i.e. across either side of 25 km segments of the border, which is another level on which we define fixed effects following Moscona et al. (2020). There are 25 unique crops grown

strand of the literature suggests that larger land sizes may allow cultivators to build economies of scale and increase productivity, stimulating local economic activity that translates to better developmental outcomes (Foster and Rosenzweig, 2022; de Janvry et al., 2015). However, these concentration effects likely only kick in at land sizes twenty five times greater than the median land holdings in our sample.

Our results are consistent with female workers having lower mobility and worse outside options than male counterparts, and thus being unable to exploit these wage differences over even very short distances. Women in this context face discontinuous costs around village borders, such that frictions in the labor market may be more binding for them (Cheema et al., 2024). While these constraints are intrinsically tied to social norms, we find somewhat similar effects for mobility constrained men as well. We test this by drawing on prior work on drivers of migration and arbitrage in rural labor markets in India (Jayachandran, 2006a). In particular, we use variation induced by the roll-out of local bank branches, with the implicit assumption that villages that got access to a local bank branch earlier would be associated with men in those villages being less credit constrained and thus less dependent on the agricultural labor market. We find that men belonging to villages that got a local bank branch after 2000 – the median year of the roll-out in our sample – have larger wage markdowns. More concentrated land ownership allows these landowners to exert their monopsony power more effectively on those workers who are intrinsically tied to the conditions of the local labor markets.

Given the median village has hundreds of landowning farmers, it may be surprising that landowners can exert monopsony power at all. There are two factors that contribute to this market power. First, as discussed at length in Anderson et al. (2015) dominant castes dominate large landholdings and political power. In a sub-sample of our villages, we found 71% of village leaders were the same sub-caste (*jati*) as the top landowner in the village and 91% were connected to one of the top twenty. Second, due to crops being planted at different dates and maturing at different speeds, the number of farmers hiring labor on a given day is surprisingly limited.

in our study region each season, but only 3.4 unique crops grown within a 25 km border segment.

Muralidharan et al. (2023) show that exogenous improvements in the implementation of the workfare programs increase both wages and employment in the private labor market, consistent with the absence of perfect competition in rural labor markets that overlap with those in our study.

The political dominance of large landowners allows them to affect the implementation of the workfare program, a key feature of the landscape of the rural economy and the most credible outside option for agricultural laborers. Feudal areas have 45% fewer person-days of NREGS work annually, with that difference increasing to 71% in peak agricultural months and disappearing entirely in lean agricultural months. A natural concern with our wage results is that feudal areas have systematically different labor markets. Another test of this is in our non-peak NREGS results: demand for the workfare program is no different between feudal and non-feudal areas during fallow periods when landowners are not competing with the state's workfare program, indicating that the labor market is not compositionally any different, and the treatment effects we see in peak agricultural months are likely driven by the supply side.

We also find that formerly feudal areas have worse provision of the kinds of public services that might help poorer residents economically diversify. We find that the share of population with primary education in feudal areas is almost 2 percentage points lower. Access to government seed centers and agricultural extension is also worse in formerly feudal areas, since larger landowners presumably are not dependent on the state's provision of these resources.

Our work relates to several literatures. We contribute to a growing literature on monopsony power in labor markets (Felix, 2024), with a particular focus on frictional land markets in agricultural settings (Deininger, 2003). Our work pushes on finding historical roots for present-day market power, and pays special attention to specific sub-populations of workers who are most exposed to employer market power. In our setting, these happen to be unskilled female agricultural workers who have the worst outside options. We also relate to a literature on tenancy and land reforms in developing countries, discussing the tradeoffs involved between agricultural productivity and the efficiency of labor markets (Besley et al., 2016; Banerjee et al., 2002). In

our setting, we do not see strong evidence of productivity or efficiency gains from larger average land parcels in feudal areas, but do see strong evidence of depressed wages for laborers.

Most directly, we contribute broadly to the literature on long-term effects of historical institutions (Smith, 2020), with a specific focus on land ownership in the Indian context (Banerjee and Iyer, 2005; Lee, 2019; Batra, 2024). Our first advance on this literature is better identification. We leverage regression discontinuities across many borders that no longer exist to identify the effect of historical institutions. Thus our estimates are less sensitive to geographic discontinuities across any specific border. Another advantage of the fine-level variation we leverage is that broad development outcomes such as overall poverty have largely converged, helping reveal the outcomes most effected by the persistent differences in land distribution.

We also contribute by focusing on detailed labor market outcomes not covered by the existing literature focused on present-day service delivery, and the performance of last-mile elected bodies and bureaucracies. Even papers with finer spatial variation (Ratnoo, 2024) focus broadly on welfare outcomes from administrative data, which tend to be universally worse, rather than isolating labor market effects. The primary contribution of this paper is highlighting the continued importance of land distribution for wages in rural India through large landowners control of social services and the wage setting process.

The rest of the paper proceeds as follows. In section 2, we discuss the institutional historical background. In section 3 and 4, we discuss the data and empirical strategy respectively. In section 5, we present the first-stage effects on land concentration, while in section 6, we discuss effects on the labor market. In section 7, we discuss effects on other service delivery outcomes. We conclude in section 8.

2.2 Institutional background

Within Hyderabad state a feudal land tenure and taxation system dating back to the 18th century was allowed to persist until 1949. While most of the land was directly administered and

taxed by the *Nizam*, the leader of the Hyderabad presidency, a significant amount of agricultural land was controlled by nobles who were given land grants in exchange for aiding the *Nizam's* ancestors in 18th century battles (Figure 2.5). Even though Hyderabad state stopped fighting wars by the early 19th century, these fiefs persisted. In general, the boundaries of these fiefs do not coincide with any modern administrative boundaries which allows for a sharper focus on village level variation in institutions. Even the Hyderabad Princely state itself was divided up between three modern Indian states: Maharashtra, Karnataka, and Telangana.

There were broadly four types of feudal estates in princely Hyderabad: *Jagir*, *Samasthans*, *Paigas*, and *Sarf-i-Khas*. Similar to the feudal estates still present in modern Britain, the boundaries of these territories is largely the product of innumerable historical clashes between great families. *Samasthans* were ancient Hindu kingdoms pre-dating the *Nizam's* rise to power in the 1700s. These Hindu rulers were given a degree of feudal autonomy over many aspects of governance including taxation in exchange for helping the first *Nizams* consolidate power. The second type of feudal land was a *Jagir*. A *Jagir* was simply a tract of land whose public revenue was assigned to an individual in exchange for rendering some service to the state. In most cases this had to do with raising troops, but it was possible to be granted a *Jagir* for other services to the *Nizam*. These grants were often hereditary although sometimes an additional tax needed to be paid to maintain the *Jagir* across generations. The *Paiga* lands were hereditary lands given to a particular *Mughal* general in the early 1700s who was helpful in the first *Nizam's* military campaigns and later divided among his descendants. The final category of land was the *Nizam's* personal estates for the maintenances of his own family, the *Sarf-i-Khas*. These areas typically faced higher land rents and greater land concentration. One account suggests some *Jagirs* had 1,000 of acres and charged tenants 10 times the rents of non-feudal areas. As of 1949, 36% of the area, 34% of the villages, and 29% of the population of the Hyderabad presidency belonged to one of the three types of feudal estates (Khusro, 1958): 2). For the purpose of our paper, we pool together *Jagirs*, *Samasthans*, *Sarf-i-Khas* and *Paigas*, since each of these is associated with a degree of feudal control awarded to local elites in the historical record. More importantly,

these areas were not part of British advised revenue modernization efforts which included the development of land records and a land administration bureaucracy ⁵.

When India gained independence from the British in 1947, the Hyderabad presidency refused to join the Indian union and was invaded by the Indian army. At the same time the *Nizam* was facing an armed communist rebellion from peasants opposing domination by the feudal landed gentry. Thus, when Hyderabad state was finally integrated into India, the abolition of the *Jagirdars* and land reform was an early priority. Notably, the peasant revolt became the roots of an armed maoist rebellion in rural Telangana that persisted until the 2000s. Although under-powered, we find no evidence that the formerly feudal areas have more political violence or are systematically more likely to elect left-wing politicians (Tables 2.14 and 2.15). There is some evidence that districts with more feudal land do have more maoist violence (Mukherjee, 2021), but this wider regional variation would be controlled for in our primary regression discontinuity specification.

As written, the laws placed ceilings on land ownership levels and established protections for tenant farmers, expecting land titles to be transferred to those who had cultivated these lands in feudal areas for a period of six years leading up to the legislation. In practice, this was not implemented in earnest. As we discuss below, large shares of tenant farmers were eased off of the land they were cultivating, with ownership reverting to local elites. Importantly, the formerly feudal areas including the *Sarf-i-Khas* lacked a modernized revenue system with local revenue bureaucrats. Survey evidence from 1953 suggests farmers in former feudal areas faced more barriers accessing the tenure they were entitled to by the land reforms.

”A significant degree of evasion is noticeable with respect to tenancy legislation and the law regarding ceilings on land. ... The tendency is for the smaller tenants to be more readily evicted than the larger ones and purchases of land have been undertaken more by the larger tenants than the smaller ones.” (Khusro, 1958): 167)

⁵An additional issue is that visually differentiating *Jagirs*, *Samasthans*, *Sarf-i-Khas* and *Paigas* holdings in the primary historical map we use is challenging

The same research surveys showed greater outmigration from the formerly feudal regions at the time of the reforms but also higher birthrates. Thus in the present day, these communities have roughly the same populations but have fewer households farming small landholdings.

Importantly, the boundaries of these feudal areas do not correspond to modern district or constituency boundaries, allowing us to isolate village level effects from other persistence channels that work through higher level institutions. This is a key difference in our setting relative to other work on colonial land tenure systems, wherein the variation resides primarily at the level of administrative districts, the borders of which are largely time-invariant. Instead, we look at fine village-level variation over very narrow geographic bandwidths of up to 20 kilometers on either side of the border.

2.3 Data

We procured land concentration data from a remote sensing company, who shared with us village-level indicators for our sample area across all three states.⁶ We use a range of publicly available datasets to estimate treatment effects on service delivery outcomes, including the population censuses of 1991, 2001 and 2011; the economic censuses of 2005 and 2013, the socio-economic census of 2012; and a village-level report of basic infrastructure (called the Mission Antyodaya data set) from 2020.⁷ In addition, we scraped data from the implementation monitoring portal of India's workfare program, the National Rural Employment Guarantee Scheme, detailing the take up and implementation quality in the villages in our sample area.

We also conducted a phone-based primary survey among 1641 village elected representatives and 459 residents. This survey was conducted July 2024 - September 2024 among respondents across study villages in the state of Telangana, focusing primarily on wages in

⁶We never received access to land parcel or owner-level data. Instead, they were able to run our analysis codes on their raw data to generate village-level aggregate indicators of land concentration, and percentiles of landholding which we use in our primary analysis of concentration.

⁷These data sets were procured from the SHRUG repository Asher et al. (2021). For a subset of villages, we also obtained individual land ownership data with names.

agricultural labor markets. We stratify our sample across border line segments of 25 kilometers, which are also the level at which we include fixed effects in our main empirical specification listed below. Each survey is conducted with a member of the village elected body, and they are asked to report typical wages and other indicators summarizing labor market dynamics in their respective villages. The sample of villages covered in these phone surveys is shown in the map in Figure 2.6 and the sample is detailed Table 2.20. Each elected leader was asked to provide phone numbers of agricultural labors hired. We sampled a subset of these numbers to verify the accuracy of leaders reported wages (Figure 2.7; Table 2.19)

A novel aspect of our survey is that we collect wages separately by task and gender for each of the two main agricultural seasons: rabi and kharif. In particular, we chose two common crop-agnostic agricultural tasks. For men, we asked about wages for fertilizer application. For women, we asked about wages for weeding. This disaggregated look at the agricultural labor market is rare in the vast literature on agriculture in India. This literature tends to consider agricultural laborers as a monolith, but in fact their reliance on these labor markets, what their outside options are, and what wages they are offered are all likely segmented by whether tasks are skilled and the gender of the workers.

2.4 Empirical strategy

We georeferenced maps from the National Archives of India, drawn by the Superintendent of the Revenue Survey and Settlement of the Hyderabad Presidency in 1854. This georeferencing allows us to delineate territories that belonged to *Jagir*, *Samasthans*, and *Paigas* vs. *Diwan-i-Khas*, i.e. feudal and non-feudal areas respectively, at the time of Indian independence. Since georeferencing is inherently laden with some degree of measurement error, we employ a “donut-hole” design in our main specification, omitting all units of observation 2 kilometers on either side of the border. We present our georeferenced study areas in figure 2.1, with the original map in figure 2.5 and the extent of our study area relative to the entire country in figure 2.8.

We follow Dell (2010); Moscona et al. (2020) in implementing a spatial regression discontinuity design. Our preferred specification uses modern 2022 district fixed effects since we use them in the outcome regressions to isolate the effects of village level institutions. The pre-independence 1931 districts correspond much more closely to the feudal areas and thus shut down the variation as expected. The preferred specification also includes a distance control mirroring the standard RD specification. In particular, this is the specification used in Moscona et al. (2020) who follow Gelman and Imbens (2019). We test a wide variety of different definitions of land concentration.

Our core specification is as follows:

$$Y_{id} = \beta_1 \text{Feudal}_{id} + \beta_2 \text{Distance}_{id} + \beta_3 \text{Distance}_{id} * \mathbb{1}_{\text{Feudal}_{id}=1} + \beta_4 X_{id} + \gamma_i + \epsilon_i$$

where Y_i is a land concentration or outcome variable of interest; Feudal_{id} is an indicator for whether a village i in district d was historically in a *jagir* or feudal area; Distance_{id} is the distance for village i in district d to the historical border separating feudal and non-feudal areas, positive for feudal areas and negative for non-feudal areas; $\text{Distance}_{id} * \mathbb{1}_{\text{Feudal}_{id}=1}$ is Distance_{id} censored at 0; X is a vector of controls, including the soil and ruggedness variables that we find are also (weakly) discontinuous at the boundary; and γ_i indicates fixed effects for the nearest line segment, where the border between feudal and non-feudal areas is split into segments of 25 kilometers each. We use robust standard errors.

For each unique data source (census, Mission Antyodaya, DHS, NREGS), we attempt to define geographic units separately and estimate distances from the border separating feudal and non-feudal areas, as well as identifying which line segment they are closest to and which present-day district they belong to. This process minimizes the need to create crosswalks across different data sources, as each strand of analysis is fully self-contained.

In figure 2.2, we present tests of continuity at the border for a range of variables to




Figure 2.2. Balance tests on RD specification

establish balance. This establishes broadly that most geographic features and time-invariant characteristics of feudal and non-feudal areas vary continuously across the threshold in our bandwidth of interest. Our results remain robust to the inclusion of these variables where appropriate.

2.5 Effects on land concentration

We focus first on the effects on land concentration. In table 2.1, we show the treatment effects on a range of land concentration measures, following our primary specification. In columns 1 and 2, we show that treated areas do not have significantly different normalized Herfindahl-Hirschmann Index or GINI indicators for landholdings. In columns 3 and 4, we report that the pareto shape parameters for the (80-20) and (90-10) splits are both 4% and 3% higher respectively in treated areas. In column 5, we report a novel indicator, the logged difference

between the 10th and 90th percentiles as a multiple of the 10th percentile. We find that this indicator is 7% smaller in treated areas, suggesting that the distribution of landholdings is more compressed in treated areas. Overall, we do not see consistent effects on *aggregate* measures of land concentration, though we have suggestive evidence of a more compressed distribution of landholdings in treated areas. Importantly, however, these measures mask differential impacts *across* the distribution of landholdings in both treatment and control areas.

Table 2.1. Aggregate land concentration

	Land concentration measures				
	Norm HHI	GINI	Pareto 80 : 20	Pareto 90 : 10	Class HHI
Feudal	-0.001 (0.002)	-0.010 (0.007)	0.019*** (0.005)	0.018*** (0.005)	0.028*** (0.009)
Non-feudal mean	0.016	0.168	0.450	0.598	0.466
F-stat	0.161	14.464	10.600	13.034	4.890
R ²	0.13	0.94	0.37	0.34	0.30
Observations	9836	9841	9841	9841	9841

This table shows results from our primary specification on a range of land concentration measures. In column (1), we show the effects on a normalized Herfindahl-Hirschmann index for land ownership. In column (2), we show the effects on the GINI index. In columns (3) and (4), we show the Pareto shape parameters for 90-10 and 80-20 splits. And in column (5) we show an intuitive measure of the spread between the 10th and 90th percentile as a share of the 10th percentile. All indicators are created using raw land parcel level data procured from a remote sensing company. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

In table 2.2, we show the treatment effect on key points of the (logged) landholdings distribution in our sample. The first percentile of landholdings in treatment villages is 23% higher than control villages. This effect dampens to being only about 5% for the 75th percentile, but attenuates fully at the very top of the distribution such that the 90th and 99th percentiles across treatment and control villages are not statistically significantly different in magnitude. We also show this dampening in the treatment effect over the distribution of (logged) landholdings in figure 2.3, and the RD plot on four points of the distribution in figure 2.4.

Feudal areas see larger land parcels at the left tail of their distribution: the smallest land

Table 2.2. Differences in the distribution of landholdings

	Percentiles of landholdings in hectares						
	1st	10th	25th	50th	75th	90th	99th
Feudal	319.648** (125.107)	1685.451*** (347.242)	3138.052*** (633.414)	3832.719*** (940.856)	3605.519*** (1169.158)	-540.583 (2454.059)	3897.611 (1.2e+04)
Non-feudal mean	1907.003	7631.200	16175.167	31626.314	52723.627	80647.307	199877.383
R ²	0.29	0.40	0.43	0.45	0.45	0.15	0.13
Observations	9749	9741	9743	9750	9760	9754	9758

This table shows results from our primary specification on percentiles across the distribution of logged landholdings in our study regions. All indicators are created using raw land parcel level data procured from a remote sensing company. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

parcel in feudal areas is approximately 15% larger than the smallest land parcel in non-feudal areas (figure 2.3). In figure 2.4, we validate our basic regression discontinuity specification. While there is a sharp discontinuity at the border when comparing the 1st percentile, the effect dampens somewhat for the 25th percentile, and diminishes markedly for the 75th percentile while disappearing at the right tail of the distribution. This pattern is consistent with the history of this period. The abolition of the *jagirdari* system was implemented imperfectly at best, and seemed to lead to the easing out of smaller tenants disproportionately in feudal areas.

Evidently, as has been documented elsewhere in India the small farmers that received land tenure in non-feudal areas retained their holdings leading to a persistent difference in the number of small farmers between feudal and non-feudal areas. In the next section, we test how these differences translate to key development outcomes, the implementation quality of arterial public programs, and local labor markets.




Figure 2.3. Treatment effects on logged landholdings across distribution

This figure plots the point estimates and 95% confidence intervals from our main regression specification, showing how the differences in landholdings across feudal and non-feudal areas are most pronounced at the left tail of the distribution and taper towards zero towards the right tail.




Figure 2.4. Regression discontinuity plots along landholding distribution

This figure shows RD plots from our key specification for four points along the landholding distribution.

2.6 Effects on village labor markets

We conducted detailed wage surveys with village elected leaders in our sample, carefully distinguishing between task and gender combinations. Focusing on these differences in wages is crucial, since agricultural labor markets are not uniform. Wage levels across villages differ by the skill level involved, and we use the task and gender to distinguish between these. In table 2.3, we report results from outcomes at the village - crop - season level. Each survey respondent is asked to report both peak and typical wages for both male and female tasks, for the primary and secondary crops across the two most recent agricultural seasons. This allows us to apply fixed effects at the crop and season level in our main regression discontinuity specification. We find that while wages are lower in feudal areas across the board, the magnitudes of the differences are much larger and statistically significant among females: women in feudal areas earn an average of 8% lower for the same task on the same crop in the same season in feudal areas, relative to non-feudal areas.

Table 2.3. Wages for agricultural labor

	Peak male fertilizer		Peak female weeding		Avg male fertilizer		Avg female weeding	
	Ln		Ln		Ln		Ln	
Feudal	-13.738	-0.021	-25.939***	-0.067***	-6.334	-0.008	-17.453**	-0.055**
	(10.707)	(0.022)	(9.879)	(0.025)	(9.847)	(0.021)	(7.432)	(0.024)
Non-feudal mean	524.630	6.235	371.358	5.854	475.762	6.128	294.371	5.631
Male vs. Female p-value			0.349	0.067			0.285	0.062
R ²	0.40	0.33	0.53	0.51	0.52	0.47	0.53	0.53
Observations	6935	6935	7313	7313	7603	7603	7792	7792

This table shows results from our main specification on wage data that we collected in our phone survey with a sample of 2000 respondents in our study area. Within each outcome, the first column reports the absolute wages in INR and the second column reports logged wages. Each regression is at the level of a village-crop-season combination, with crop and season fixed effects and standard errors clustered at the village level. The surveys ask both the peak and typical wages for each task-gender combination listed.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Women in these settings tend to have limited mobility. These mobility constraints may have discontinuous impacts along the village or neighborhood border Cheema et al. (2024)⁸. These mobility constraints faced by women may allow them to be disproportionately more exposed to local monopsony power in these labor markets. While the incidence of this monopsony may have reflected in lower wages even for male workers in contexts where male workers also faced constraints on mobility and migration, either due to travel costs or credit frictions Jayachandran (2006a)⁹.

Is this level of monopsony power reasonable? One approach following Muralidharan et al. (2023) is to compare the markdown to benchmarks in the literature. In Appendix 2.10.1 we place our results on wages and land concentration in a standard model of oligopolistic competition to calculate markdowns. This procedure highlights that to quantitatively reconcile the results workers in feudal areas need to have worse outside options which we provide evidence for below. We estimate workers receive 84% of their marginal product in feudal areas and 77% in non-feudal areas. To contextualize, a review by Sokolova and Sorensen (2021) estimating elasticities across many global labor markets finds an average elasticity of 7.1, implying that workers receive 88% of their marginal product, with a 95% confidence interval from 64% to 93%. Thus, the markdowns we estimate are well within the global range.

A natural concern following these results is that feudal and non-feudal areas have different labor demand levels and/or different productivity. In the survey, we asked representatives to detail the labor demand (in person-days) across all tasks for the primary and secondary crops in the most recent agricultural season for a typical 10 acre plot in their villages. We find no effects on the total labor demand, or the demand specifically for the tasks that we find the wage effects on. We similarly elicit the yields per acre for the primary and secondary crops across the two

⁸Cheema et al. (2024) show that rural women in Pakistan are four times as likely to visit a training center when they are based in their own village, with over half the penalty incurred upon crossing the village border. Field and Vyborny (2022) increased women's job search in Pakistan by providing a women-only transport service to and from work. Among rice transplanting women in Telangana, 47% report their husbands would not approve of them working outside their village including neighboring villages Brownstone (2025).

⁹We currently lack historical data on male wages to evaluate these effects.

most recent seasons, as well as the acreage in the entire village dedicated to planting the crops. Across all of these outcomes, we find no significant differences across feudal and non-feudal areas despite deploying the same specification as in the wage result.

Table 2.4. Labor demand

	Total	Task-specific	
		Male (fertilizers)	Female (weeding)
Feudal	-20.680 (24.848)	2.520 (4.274)	-2.211 (9.456)
Non-feudal mean	213.800	27.758	70.212
R ²	0.17	0.16	0.12
Observations	2393	2053	2053

Table 2.5. Productivity

	Yield		Planting acres	
		Ln		Ln
Feudal	-7.789 (6.255)	-0.079 (0.058)	-82.174 (136.365)	0.114 (0.110)
Non-feudal mean	107.767	4.426	773.851	5.883
R ²	0.56	0.50	0.25	0.46
Observations	4610	4557	7036	6975

We additionally collect information on a range of variables characterizing the aggregate labor markets in rural areas, including the total number of households, the number of employers, workers, and migrants. We find no systematic differences across feudal and non-feudal areas in

the size of thickness of the labor markets, as reported in table 2.18, reinforcing the explanation that any differences in wages are likely downstream of differential monopsony power. We also do not find differences in the number of land transactions in feudal areas, reflecting the stickiness in land markets in this setting.

A key determinant of rural labor market dynamics in this setting is the availability of outside options for workers. We see that the effects of land concentration on wages extend to men in labor markets where they have limited outside options. First, we leverage variation in the roll-out of bank branches which prior literature showed decreased labor supply elasticity dampening wage responses to agricultural productivity shocks (Jayachandran, 2006b). Credit access likely improves the range of economic opportunities available in a village over time. Notably, much of the work on bank branch expansion in India focuses on district level variation which we control for. In villages that got bank branches later opportunities we observe wage markdowns for men as well as women (Table 2.6). This result highlights the role women's lack of outside of opportunities likely play in our main results.

Table 2.6. Heterogeneity by Bank Access

	Ln Female weeding wages		Ln Male fertilizer wages	
	Peak	Avg	Peak	Avg
Feudal	-0.038 (0.032)	-0.059** (0.029)	0.019 (0.027)	0.038 (0.026)
Post 2000 bank	-0.017 (0.039)	-0.056* (0.033)	0.063* (0.034)	0.047 (0.031)
Feudal $\times$ Post 2000 bank	-0.030 (0.047)	0.020 (0.039)	-0.078** (0.039)	-0.097*** (0.037)
Non-feudal mean (logged)	5.857	5.629	6.241	6.132
Non-feudal mean (INR)	371.524	293.657	527.197	477.239
R ²	0.52	0.56	0.34	0.48
Observations	3499	3733	3301	3634

In table 2.7 we present a range of outcomes related to NREGS, the rural workfare program that provides a statutory guarantee of 100 days of paid work every financial year to at least one member of any household who seeks it. A vast literature on NREGS shows that demand for this work is high across rural areas and implementation varies substantially (Imbert and Papp, 2015). In our setting, NREGS implementation quality serves as a test of how healthy agricultural laborers' outside options are. Muralidharan et al. (2023) show that improving the quality of NREGS implementation drives up households' earnings by 14%, but that 86% of these gains come from increases in real wages *and* employment in the private agricultural labor market. This is reflected in increases in workers' reservation wages. In their setting, these effects on the private agricultural labor market are accentuated in villages where land ownership is more

concentrated, suggesting that these labor markets are operating under monopoly power in the status quo prior to the improvement in NREGS implementation quality.¹⁰ In column (2) of table 2.7, we show that feudal areas, where landownership is more concentrated, has 42% fewer active NREGS job cards, even though the number of registered job cards is (noisily) similar. In column (3), we show that the number of households employed in NREGS on the extensive margin is 58% lower, and the total person-days employed in feudal areas is 67% lower. These effects point to depressed NREGS implementation quality in feudal areas overall. These results could plausibly be explained by those with monopoly power in agricultural labor markets being able to exert influence on local elected bodies to constrain workers' outside options. The fact that the effects only appear during the peak agricultural months (Table 2.8) is evidence that the withholding of days may be strategic¹¹. This is also consistent with Anderson et al. (2015), who evaluate a workfare scheme in Maharashtra which was a precursor to the nationwide NREGS.

¹⁰Muralidharan et al. (2023) is based on a large-scale RCT in the undivided state of Andhra Pradesh, which neighbors (and minimally overlaps with) our study sample.

¹¹While there are many pathways to strategically decrease NREGA provision by slowing projects it is difficult for local officials to strategically increase NREGA provision through new projects (Muralidharan et al., 2016).

Table 2.7. NREGS implementation: registration and household-level access

	Job cards		Households employed	
	Ln(Registered)	Ln(Active)	Ln(Total)	I{Any for 100 days}
Feudal	-0.568*** (0.206)	-0.741*** (0.259)	-1.386** (0.554)	-0.508** (0.210)
Non-feudal mean (logged)	6.830	5.061	4.560	.
Non-feudal mean (raw)	1265.480	270.067	199.131	0.497
R ²	0.58	0.66	0.67	0.27
Observations	4890	4908	4818	4960
No. of villages	4739	4774	4694	4795

This table shows results from our main specification on NREGS implementation outcomes in 2023, using data scraped by us from the NREGS portal. Columns (1) and (2) depict the total number of registered and active NREGS job cards respectively. Columns (3) and (4) depict the number of households that received any work through NREGS and the number that received the statutory guarantee of 100 days of work, respectively. Column (5) depicts the total person-days of work provided. Regressions control for the total number of households as reported in the Mission Antyodaya 2020 data. Standard errors are robust.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2.8. NREGS implementation: log person days per month

	Annual		Peak ag months		Lean ag months	
Feudal	-0.719**	-0.711**	-0.799**	-0.817**	-0.180	-0.179
	(0.340)	(0.340)	(0.360)	(0.364)	(0.256)	(0.255)
Non-feudal mean (logged)	5.543	5.543	5.209	5.209	6.634	6.634
Non-feudal mean (raw)	669.463	669.463	465.784	465.784	2003.883	2003.883
Year FE		✓		✓		✓
R ²	0.58	0.59	0.53	0.55	0.57	0.59
Observations	49228	49228	48656	48656	43042	43042
No. of villages	6003	6003	6003	6003	5978	5978
No. of years	9	9	9	9	9	9

This table shows results from our main specification on (logged) NREGS person-days in the period 2016 - 2023, using data scraped by us from the NREGS portal. Columns (1) and (2) depict the total number of person-days at the monthly level. Columns (3) and (4) depict the number of person-days in peak agricultural months. Columns (5) and (6) depict the total person-days of work provided in lean agricultural months. Regressions have year fixed effects when indicated, and standard errors are clustered at the village level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Beyond NREGA, we find a number of systematic differences in village public goods provision across feudal and non-feudal areas. These even translate to notable differences in educational attainment despite these regions similar levels of per-capita consumption. In table 2.9, we report using census data that those in feudal areas are 3% less likely to have attained primary and middle school. Among both adults and children, feudal areas demonstrate long-run differences in health and education outcomes. These are also reflected in present-day access to health and education infrastructure. In table 2.10, we report that feudal areas are 88% less likely

to have middle schools, and are 91% farther away from primary and community health centers. Notably, education is important for facilitating workers outside labor market options so landed elites may have interest in suppressing its provision.

Table 2.9. Aggregate education levels

	Share of population with ... or above		
	Sec school	middle school	primary school
Feudal	-0.004	-0.009*	-0.016**
	(0.004)	(0.005)	(0.006)
Non-feudal mean	0.187	0.294	0.491
R ²	0.46	0.47	0.38
Observations	9153	9153	9153

This table shows results from our main specification on educational attainment outcomes, as reported in the 2011 population census. Standard errors are robust. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2.10. Health and education infrastructure

	School in village	Distance to ... (km)		
	Primary	Non-primary	PHC/CHC	Hospital
Feudal	-0.000	-0.549***	3.948**	2.324
	(0.025)	(0.203)	(1.969)	(2.884)
Non-feudal mean	0.973	0.579	4.299	6.821
R ²	0.18	0.23	0.22	0.26
Observations	5062	5062	5062	5062

This table shows results from our main specification on educational attainment outcomes, as reported in the 2020 Mission Antyodaya dataset. Column titles explain the construction of variables. SSC stands for senior secondary school. PHC and CHC stand for primary health center and community health center respectively. Standard errors are robust. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

2.7 Alternate Explanations

One concern is that despite survey reports of similar labor demand, differential agricultural investments lead to different agricultural productions functions in feudal and non-feudal areas. The lower wages could simply be a result of lower labor productivity. While we do find evidence of weakly lower investments in agricultural equipment and irrigation these differences are not large enough to explain the wage-gap and likely reflect differential provision of agricultural public goods. Using data from the 2012 SECC, we find that feudal areas have a 2 percentage points smaller share of households reporting that their agricultural fields are using any mechanized inputs, and a 2 percentage points smaller share of land irrigated overall (table 2.11). We also find that feudal areas are less likely to have seed centres and fertilizer shops, which are fundamental sources of support for farmers, aiding them with subsidized inputs as well

as information on agricultural best practices (table 2.13). Feudal areas are also directionally less likely to have soil testing programs and agricultural cooperative societies, though these effects are not significant by conventional standards of statistical significance. In our surveys, we elicit from GP heads information on how many households in their GPs own tractors and have made investments greater than INR 1 lakh (USD 1200) on their land. We do not see any systematic or statistically significant differences across feudal and non-feudal areas (table 2.12).

These differences in inputs and the quality of agricultural investments have modest impacts on yields themselves, as discussed previously in table 2.5. While we lack the ability to establish mechanisms for these results conclusively with individual level landholding-linked data on agricultural inputs or yields, we contend that these aggregate effects on agricultural outcomes mask substantive distributional differences that are downstream of land concentration. Larger landowners have the ability to drive higher mechanization in their own fields given economies of scale (Foster and Rosenzweig, 2022). At the same time, smaller landowners in these areas may be less able to invest in mechanization if they are closer to subsistence, and public goods that are meant to support cultivation may be weaker in these areas as a result of elite capture. The most important take-away from these results is that the feudal areas are not more productive than the non-feudal ones as a result of their land concentration nor are the productivity differences large enough to explain within season and within crop wage gaps.

Table 2.11. Investments in agriculture

	Agricultural equipment			Share of total land that is		
	Any	Mechanized	Irrigation	Irrigated	Irrigated for 2 crops	Unirrigated
Feudal	-0.024*	-0.013	-0.019	-0.020**	-0.140*	-0.513
	(0.013)	(0.018)	(0.017)	(0.010)	(0.085)	(0.498)
Non-feudal mean	0.938	0.854	0.884	0.409	1.468	2.795
R ²	0.06	0.18	0.08	0.23	0.06	0.03
Observations	9281	9281	9281	9079	9145	9145

This table shows results from our main specification on a range of outcomes denoting the use of agricultural inputs as reported in the SECC (2012). In columns (2) and (3) we use binaries for whether the village uses mechanized or irrigation equipment respectively, and column (1) equals 1 if they use either. Column (4) reports the share of land that is irrigated, while column (5) reports the share irrigated for 2 crops and column (5) denotes the share of unirrigated land. Standard errors are robust. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2.12. Investments in agriculture

	Households with tractors	Households with investments $\geq$ INR 1 lakh
Feudal	1.468	-2.722
	(3.152)	(3.232)
Non-feudal mean	25.033	18.841
R ²	0.10	0.15
Observations	2350	2334

This table shows results from our main specification on the number of households that own tractors and the number of households that have made investments of greater than INR 100,000. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2.13. Agriculture infrastructure

	Ag infrastructure			Water infrastructure		Other
	Seed centre	Fertilizer shop	Soil testing	Watershed	Rainwater harvesting	Ag coop society
Feudal	-0.300*	-0.408**	-0.119*	-0.005	-0.344	-0.197
	(0.159)	(0.202)	(0.070)	(0.178)	(0.309)	(0.318)
Non-feudal mean	0.127	0.276	0.057	0.302	0.507	0.387
R ²	0.22	0.26	0.23	0.26	0.32	0.31
Observations	5062	5062	5062	5062	5062	5062

This table shows results from our main specification on a range of outcomes denoting infrastructure supporting agricultural practices as reported in the 2011 population census. Columns (1), (2) and (3) report binaries for the presence of a center to buy subsidized seeds, a shop for subsidized fertilizers, and a soil testing center respectively. Columns (4) and (5) denote whether there is a watershed or rainwater harvesting infrastructure available. Column (6) denotes whether the village has an agricultural cooperative society. Standard errors are robust. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

One concern with our results is that changes in the composition of the population and differential migration may be linked to the quality of service delivery. If differential land concentration and agricultural labor market activity pushes certain types of households out of feudal areas, we might expect lower demand for key government programs, affecting the quantity and quality of their delivery. There is some evidence of differential migration in the early years of the reform:

“There has been a marked tendency for families to leave *ex-jagir* villages and such an exodus has accounted for a reduction in families by 8.6% in these villages” (Khusro, 1958, 170): 170

We address this concern in two ways. First, we normalize indicators by population (either in per capita terms or in terms of the total number of households) where appropriate. Second, we evaluate long-term population trajectories in our study areas, comparing inflows and outflows in feudal and non-feudal villages. We show trends for the overall population in figure 2.9, and trends specifically for working age population in figure 2.10. These results suggest that differences by population, both overall and working age, are relatively stable over the course of 1991 and 2011

as reported in three rounds of the population census. The historical literature suggests that the period immediately following the land reform legislation saw movement out of feudal areas, but there has not been meaningful *additional* movement in the last four decades. This is consistent with a broader literature that finds frictions in agricultural labor markets constraining worker mobility across space and sectors Baysan et al. (2024).

2.8 Conclusion

In this paper, we evaluate the long-run effects of historical differences in land tenure systems in princely Hyderabad, and the flawed application of land reform legislation in this region in the post-colonial period, on a range of rural service delivery outcomes as well as agricultural labor markets. We show that land concentration levels remain persistently elevated more than 100 years later in regions where land parcels were gifted to local elites. These areas also demonstrate poorer service delivery outcomes on a range of health and education variables.

Moreover, we find significant evidence of elite capture and monopsony power in rural labor markets. We find that these regions have worse implementation of the arterial workfare program, meant to provide an outside option for primarily landless agricultural laborers. Scuttling the implementation of this workfare program allows large landowners in feudal areas to sustain dramatically lower wages for agricultural workers in their own lands. The wage markdowns are significantly worse for workers with even poorer outside options owing to constraints on their mobility: in our setting, these are women working on unskilled tasks.

Our work extends recent advances in our understanding of monopsony as well as elite capture in rural labor markets, and highlights the importance of persistent land distributions in understanding the long run effects historical institutions on development.

2.9 Chapter Acknowledgments

Chapter 2, in full, is currently being prepared for submission for publication of the material. Brownstone, Steven; Srivastava, Kartik. The dissertation author was one of the primary investigators and authors of this material.

2.10 Appendix

2.10.1 Model

Empirical evidence indicates that although overall labor supply and demand in feudal and non-feudal areas are similar, the composition of workers' outside options differs. There are two channels we explore. First, in non-feudal the greater number of small farms makes collusion in wage setting more difficult, whereas in feudal areas the number of small farmers is lower, so farmers are more able to collude to restrict labor demand and drive up wages. We show in simple oligopolistic framework the difference in land concentration alone is unlikely to generate the wage difference we observe. However, allowing the labor supply elasticity to be lower in feudal villages, reflecting the reduced NREGA provision, rationalizes our results with other estimates of labor supply in this setting.

Oligopolistic wage setting is very natural in agricultural labor markets since farmers know the landholdings and crop choices of their neighbors and thus labor demand to a first approximation. Firms knowing other firms labor demand and optimizing accordingly is the core assumption for oligopolistic models of the labor market.

In this model firms are choosing employment to maximize revenues.

$$\max_{L_i} R(L_i) - w(L + L_i^*)L_i$$

The intuition in our case is farmers are choosing labor with an understanding of how their labor demand will effect wages because they know other farmers labor demand.

The first order condition of each firm leads to the markdown of:

$$\frac{MRP - w}{w}$$

The key intuition here is that larger farms MRP curves are shifted to the right as they are able to derive more revenue per given level employment. Note that it is still true the marginal worker is

more productive on a small firm than a large farm due to diminishing marginal returns. The idea is simply they diminish less steeply the more land a farmer has. Thus the more large farms are in a village the greater the wage markdown.

Note that the MRP translates directly to the employment share of each firm i conditional on the labor supply elasticity

$$\frac{MRP_i - w}{w} = \epsilon^{-1} \left[\sum_i^n \left(\frac{L_i}{L} \right)^2 \right]$$

Notably this is different from the raw HHI in land for two reasons. First, this should be the HHI in labor shares and since small farms use a much greater share of family labor the HHI in land doesn't capture this well. Secondly, larger fields do not operate as individual firms. There is evidence farmers of the same caste work together to restrict hiring to push down wages and optimize profits so these cast groups can be better conceptualized as one firm.

Rather than using individual farms we use land size categories, adapted from the agricultural census, as an approximation for caste groups. In particular, we use cutoffs c of 0.5, 1, 2, 3, and 4 hectares. Then the markdown becomes:

$$\frac{MRP - w}{w} = \epsilon^{-1} \left[\sum_c^5 \left(\frac{L_c}{L} \right)^2 \right] = \epsilon^{-1} H$$

Thus when the employment is more concentrated among the large farmers H increases and the markdown is bigger.

If we assume the labor supply elasticity is the same in feudal and non-feudal areas focusing on the effect of land concentration we get the following system of equations:

Non-Feudal:

$$\frac{MRP - w}{w} = \epsilon^{-1} H$$

Feudal

$$\frac{MRP - w'}{w'} = \varepsilon^{-1} \alpha^{-1} H'$$

We can then solve for ε by equating the MRP

$$\varepsilon = \frac{wHHI - w'HHI'}{w' - w}$$

We can take w, w', H, H' from our data. w, w' are simply the wages at the discontinuity for feudal and non-feudal areas. H, H' are the categorical versions of the HHI in feudal and non-feudal area. However, the result ε of -0.1082 is infeasible.

The second channel land concentration effects workers wages is through large landowners control of local governments. The large landowners are able to reduce workers access to the government workfare scheme which provides a credible outside option. This can be thought of as reducing workers labor supply elasticity ε by a multiplier $\alpha < 1$ further increasing the markdown on wages.

We now have a system of two equations describing wages in the feudal and non-feudal areas. Non-Feudal:

$$\frac{MRP - w}{w} = \varepsilon^{-1} H$$

Feudal

$$\frac{MRP - w'}{w'} = \varepsilon^{-1} \alpha^{-1} H'$$

Now rather than solving for ε we can take ε as the 3.07 estimated in Muralidharan et al. (2023) and solve for α that rationalizes are data.

$$\alpha = \frac{wHHI - w'HHI' \varepsilon^{-1}}{MRP - w'}$$

$$\alpha = \frac{wHHI - w'HHI' \varepsilon^{-1}}{wHHI \varepsilon^{-1} + w - w'}$$

The resulting value for α , .73, suggests that labor supply elasticity is 27% lower in-feudal villages which provide 70% fewer NREGS person days during the peak agricultural season. This value aligns with the limited literature exploring the sensitivity of labor supply elasticities to outside options. Caldwell and Oehlsen (2023) estimate an α of .53 implying the labor supply elasticity for ride-hailing firm drivers decreases from 1.5 to .8 when workers have no outside firm option.

2.10.2 Figures

Figure 2.5. Map of Hyderabad Princely State Showing Different Tenure Systems



Figure 2.6. Phone survey sample




Figure 2.7. Accuracy of Wages




Figure 2.8. Jagirs within India




Figure 2.9. Population dynamics




Figure 2.10. Working population dynamics





Figure 2.11. Wage HTEs by bank exposure



2.10.3 Tables

Table 2.14. Differences in violent events

	ACLED			UCDP		
	Total conflict	Land conflict	Fatal conflict	All conflicts	All deaths	a
Treatment	-0.005 (0.004)	-0.001 (0.001)	-0.001 (0.008)	-0.002 (0.008)	-0.219*** (0.031)	-0.123 (0.118)
Dep var mean	0.005	0.002	-4.591	-4.590	2.564	0.204
R ²	0.01	0.01	0.04	0.03	0.26	0.01
Observations	18583	18583	18583	18583	18566	18583

This table shows results from our primary specification on geo-coded violence data from ACLED. A major limitation is that many ACLED events appear to be coded to district centroids inducing measurement error.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2.15. Differences assembly election results

	Leftist party won assembly election = 1		
Treatment	-0.000	-0.001	0.003
	(0.007)	(0.007)	(0.007)
District FE	✓	✓	✓
Year FE		✓	✓
Constituency type			✓
Dep var mean	0.030	0.030	0.030
R ²	0.10	0.12	0.17
Observations	56525	56525	56525
Constituencies	123	123	123

This table shows results from our primary specification on state assembly elections. Unfortunately, there are a limited number of unique state assembly seats in our area of interest. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2.16. Linking large landowners and village leaders through caste

	Any match		Number of matches		Average match rank		Highest match rank	
	Last name	Caste	Last name	Caste	Last name	Caste	Last name	Caste
Treatment	0.022 (0.068)	-0.039 (0.045)	-1.045 (1.544)	-6.479 (4.991)	-0.001 (1.037)	0.917 (0.783)	-1.115 (1.183)	0.231 (0.833)
Non-feudal mean	0.689	0.948	6.704	45.556	10.792	4.061	9.452	2.815
R ²	0.16	0.12	0.16	0.32	0.20	0.20	0.20	0.15
Observations	238	238	238	238	238	238	238	238

For 238 villages compared the names of ward members to the names of the top 20 landowners. Matching done by local enumerators. Ambiguous cases are treated as not matches.

Table 2.17. Balance table

	Elevation		Temperature	Precipitation	Soil depth (sd)	Distance to	
	Mean	SD				river (ln km)	town (ln km)
Feudal	-1.687 (1.794)	-0.058 (0.252)	0.017 (0.011)	-0.053*** (0.006)	0.015 (0.023)	-0.287*** (0.057)	0.114*** (0.023)
Non-feudal mean	456.872	10.751	26.999	3.292	0.000	0.787	3.022
R ²	0.89	0.26	0.88	0.84	0.53	0.36	0.33
Observations	17916	17916	18561	18561	18561	16814	17664

Tests for balance across geographic characteristics. Note that some differences are statically significant, but very small in magnitude.

Table 2.18. Treatment effects on village composition

	# HHs			# workers	# land txns	# migrant HHs	
	Current	Recent exits	Employers			Cyclical	Permanent
Feudal	17.772 (91.131)	1.956 (4.108)	53.735 (38.439)	237.664 (344.169)	-6.220 (10.488)	-11.509 (18.527)	-13.761 (16.551)
Non-feudal mean	1010.216	22.623	142.883	1146.251	67.671	140.532	82.028
R ²	0.49	0.11	0.18	0.09	0.07	0.41	0.18
Observations	2366	2339	2330	2321	2209	2346	2356

This table shows results from our primary specification on the composition of village populations in feudal and non-feudal areas. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2.19. Wage agreement between workers and leaders in survey

Variable	Concordance correlation coefficient		95% CI	Mean difference		Observations
	Estimate	SE		INR	%	
Peak male wages	0.544	0.034	[0.477, 0.610]	14.822	2.34	426
Peak female wages	0.598	0.031	[0.538, 0.658]	33.675	6.88	425
Typical male wages	0.631	0.029	[0.574, 0.689]	6.641	1.29	425
Typical female wages	0.688	0.026	[0.638, 0.738]	2.851	0.91	425

Concordance correlation coefficient ranges from -1 to 1 and refers the degree of agreement between workers and leaders reports for the same villages.

Table 2.20. Characteristics of survey

Variable	Mean	SD	Min	Max	Observations
Panel A: Workers survey					
Number of surveys/village	1.475	0.768	1	5	459
Share of male respondents	0.961	0.182	0	1	459
Average respondent age	40.990	9.556	18	76	456
Share of married respondents	0.928	0.244	0	1	459
Share of respondents in cultivation last season	0.902	0.279	0	1	459
Share of upper caste respondents	0.209	0.387	0	1	459
Share of SC/ST respondents	0.196	0.376	0	1	459
Share of other caste respondents	0.400	0.463	0	1	459
Panel B: Officials survey					
Number of surveys/village	1.460	0.869	1	5	1641
Village head surveyed	0.593	0.457	0	1	1641
Village vice-head surveyed	0.057	0.212	0	1	1641
Village elected member surveyed	0.350	0.443	0	1	1641

Characteristics of respondents for wage survey. Not all officials were asked if they were knowledgeable about agricultural issues before being surveyed.

Table 2.21. Wage survey with workers and officials

	Peak male fertilizer		Peak female weeding		Avg male fertilizer		Avg female weeding	
	Ln		Ln		Ln		Ln	
Feudal	-8.763 (12.988)	-0.006 (0.024)	-27.704* (14.254)	-0.058* (0.030)	-5.504 (10.571)	-0.000 (0.022)	-18.538** (8.512)	-0.058** (0.027)
Non-feudal mean	567.073	6.297	405.867	5.912	490.715	6.151	294.894	5.632
Male vs. Female p-value			0.434	0.066			0.702	0.125
R ²	0.46	0.41	0.54	0.55	0.51	0.46	0.48	0.49
Observations	9318	9318	7809	7809	9619	9619	7882	7882

Note there are very few female worker respondents so those estimates do not change substantially.

Table 2.22. Heterogeneity by SEZ Placement

	Ln Female weeding wages		Ln Male fertilizer wages	
	Peak	Avg	Peak	Avg
Feudal	-0.065** (0.026)	-0.047** (0.021)	-0.043* (0.022)	-0.055*** (0.021)
Low SEZ exposure	0.034 (0.027)	0.022 (0.024)	-0.016 (0.023)	-0.052** (0.022)
Feudal $\times$ Low SEZ exposure	-0.004 (0.032)	-0.014 (0.028)	0.036 (0.027)	0.083*** (0.026)
Non-feudal mean (logged)	5.854	5.631	6.235	6.128
Non-feudal mean (INR)	371.209	294.302	524.602	475.710
R ²	0.51	0.53	0.33	0.47
Observations	7317	7796	6939	7607

Low SEZ exposure is defined as having more than the median number of SEZs within a 200km radius. SEZ locations from (Gallé et al. 2024)

Table 2.23. Heterogeneity by Road Access

	Ln Female weeding wages		Ln Male fertilizer wages	
	Peak	Avg	Peak	Avg
Feudal	-0.081*** (0.025)	-0.124*** (0.023)	-0.028 (0.023)	-0.038* (0.021)
No major road	-0.014 (0.028)	-0.094*** (0.026)	0.038 (0.025)	0.036 (0.024)
Feudal $\times$ No major road	-0.005 (0.034)	0.097*** (0.031)	-0.021 (0.029)	0.025 (0.028)
Non-feudal mean (logged)	5.869	5.645	6.240	6.136
Non-feudal mean (INR)	377.868	299.102	528.046	480.216
R ²	0.52	0.54	0.34	0.48
Observations	5962	6389	5641	6219

Village with access to a district road in the village census.

Table 2.24. Heterogeneity by low non-farm employment

	Ln Female weeding wages		Ln Male fertilizer wages	
	Peak	Avg	Peak	Avg
Feudal	-0.039 (0.025)	-0.018 (0.023)	0.009 (0.022)	0.015 (0.021)
Low non-farm employment	-0.023 (0.027)	0.031 (0.026)	0.069*** (0.024)	0.048** (0.024)
Feudal $\times$ Low non-farm employment	-0.009 (0.034)	-0.037 (0.032)	-0.057* (0.029)	-0.046 (0.029)
Non-feudal mean (logged)	5.787	5.561	6.204	6.077
Non-feudal mean (INR)	342.290	272.446	507.086	449.805
R ²	0.46	0.47	0.28	0.40
Observations	5215	5433	5053	5339

Data from the 2013 Economic Census, which records the total amount of non-agricultural employment in the country. We test for heterogeneous treatment effects for village with below-median non-agricultural enterprises

Bibliography

Anderson, Siwan, Patrick Francois, and Ashok Kotwal, “Clientelism in Indian Villages,” *American Economic Review*, 2015, 105 (6), 1780–1816.

Asher, Sam, Tobias Lunt, Ryu Matsuura, and Paul Novosad, “Development research at high geographic resolution: an analysis of night-lights, firms, and poverty in India using the SHRUG open data platform,” *The World Bank Economic Review*, 2021, 35 (4).

Banerjee and Iyer, “History, Institutions, and Economic Performance: The Legacy of Colonial Land Tenure Systems in India,” *American Economic Review*, 2005, 95 (4), 1190–1213.

—, **Paul Gertler, and Maitreesh Ghatak**, “Empowerment and Efficiency: Tenancy Reform in West Bengal,” *Journal of Political Economy*, 2002.

Batra, Kartikeya, “Long-run effects of land redistribution: evidence from India,” *Working Paper*, 2024.

Baysan, Ceren, Manzoor H. Dar, Kyle Emerick, Zhimin Li, and Elisabeth Sadoulet, “The agricultural wage gap within rural villages,” *Journal of Development Economics*, 2024, 168, 103270.

Besley, Timothy, Jessica Leight, Rohini Pande, and Vijayendra Rao, “Long-run impacts of land regulation: Evidence from tenancy reform in India,” *Journal of Development Economics*, 2016, 118, 72–87.

Bolhuis, Marijn A., Swapnika R. Rachapalli, and Diego Restuccia, “Misallocation in Indian Agriculture,” Technical Report 29363, National Bureau of Economic Research October 2021. Revised March 2025, JEL No. O11, O13, O4.

Brownstone, Steven, “The Labor Market Effects of Agriculture Mechanization: Experimental Evidence from India,” *Working Paper*, 2025.

Caldwell, Sydnee and Emily Oehlsen, “Gender, Outside Options, and Labor Supply: Experimental Evidence from the Gig Economy,” *Working Paper, University of California Berkeley*, 2023.

- Cheema, Ali, Asim Khwaja, M Farooq Naseer, and Jacob Shapiro**, “Glass Walls: Experimental Evidence on Constraints faced by Women in Accessing Valuable Skilling Opportunities,” *Working Paper*, 2024.
- de Janvry, Alain, Kyle Emerick, Marco Gonzalez-Navarro, and Elisabeth Sadoulet**, “Delinking Land Rights from Land Use: Certification and Migration in Mexico,” *American Economic Review*, October 2015, *105* (10), 3125–49.
- Deininger, K.**, “Land Markets in Developing and Transition Economies: Impact of Liberalization and Implications for Future Reform,” *American Journal of Agricultural Economics*, 2003.
- Dell, Melissa**, “The persistent effects of Peru’s mining mita,” *Econometrica*, 2010, *78*, 1863–1903.
- Dhanagare, D.N.**, “Social origins of the peasant insurrection in Telangana (1946-51),” *Contributions to Indian Sociology*, 1974, *8* (1), 109–134.
- Felix, M.**, “Trade, labor market concentration, and wages,” *Working Paper*, 2024.
- Field, Erica and Kate Vyborny**, “Women’s Mobility, and Labor Supply: Experimental Evidence from Pakistan,” *Asian Development Bank Economics Working Paper Series*, 2022, 655.
- Foster, Andrew D. and Mark R. Rosenzweig**, “Are There Too Many Farms in the World? Labor Market Transaction Costs, Machine Capacities, and Optimal Farm Size,” *Journal of Political Economy*, 2022, *130* (3), 636–680.
- Gelman, Andrew and Guido Imbens**, “Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs,” *Journal of Business & Economic Statistics*, 2019, *37* (3), 447–456.
- Genicot, Garance and Debraj Ray**, “Aspirations and Inequality,” *Econometrica*, 2017, *85* (2), 489–519.
- Hornbeck, Richard and Suresh Naidu**, “When the Levee Breaks: Black Migration and Economic Development in the American South,” *American Economic Review*, 2014, *104* (3), 963–990.
- Imbert, Clément and John Papp**, “Labor Market Effects of Social Programs: Evidence from India’s Employment Guarantee,” *American Economic Journal: Applied Economics*, April 2015, *7* (2), 233–63.
- Jayachandran, Seema**, “Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries,” *Journal of Political Economy*, 2006, *114* (3), 538–575.

- , “Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries,” *Journal of Political Economy*, 2006, 114 (3), 538–575.
- Khusro, A M**, *Economic and social effects of Jagirdari abolition and land reforms in Hyderabad*, Department of Publication & University Press, Osmania University, Hyderabad, 1958.
- Kim, Oliver and Jen-Kuan Wang**, “Land Reform in Taiwan, 1950-1961: Effects on Agriculture and Structural Change,” *Working Paper*, 2024.
- Lee, Alexander**, “Land, State Capacity and Colonialism: Evidence from India,” *Comparative Political Studies*, 2019, 52, 412–444.
- Moscona, Jacob, Nathan Nunn, and James A. Robinson**, “Segmentary Lineage Organization and Conflict in Sub-Saharan Africa,” *Econometrica*, 2020, 88 (5), 1999–2036.
- Mukherjee, Shivaji**, *Colonial Institutions and Civil War: Indirect Rule and Maoist Insurgency in India* Cambridge Studies in Contentious Politics, Cambridge University Press, 2021.
- Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar**, “Building State Capacity: Evidence from Biometric Smartcards in India,” *American Economic Review*, October 2016, 106 (10), 2895–2929.
- , —, and —, “General Equilibrium Effects of (Improving) Public Employment Programs: Experimental Evidence From India,” *Econometrica*, 2023, 91 (4), 1261–1295.
- Ratnoo, Vigyan D.**, “Persistent effects of colonial land tenure institutions: Village-level evidence from India,” *Journal of Development Economics*, 2024, 168, 103247.
- Sethi, Pankaj**, “Operation Polo 1948: How Hyderabad’s last Nizam agreed to give up his land for reform,” *The Siasat Daily*, 2023. Accessed: 2025-02-10.
- Smith, C**, “Land Concentration and Long-Run Development in the Frontier United States,” *Working Paper*, 2020.
- TIME**, “HYDERABAD: Silver Jubilee Durbar,” *Time Magazine*, 1937.

Chapter 3

Managing What You Can Measure? The Promises and Pitfalls of Measuring Bureaucrat Performance

3.1 Introduction

Most development programs are ultimately implemented at scale by bureaucrats. The effectiveness of development schemes thus depends critically on the performance of public officials like city managers, school teachers, primary health workers, tax collectors, and agricultural extension officers. Public sector organisations can improve bureaucrat performance and service delivery by adopting effective personnel policies and management practices (Rasul and Rogger, 2017). One common strategy, particularly in contexts where performance-based financial incentives are infeasible, is to monitor bureaucrat performance with dashboards not explicitly tied to high-powered incentives (Finan et al., 2017).

The impact of dashboards is unclear. They can help bureaucrats learn how to improve at their jobs and enable the use of non-pecuniary incentives like postings and awards (Khan et al., 2019). However, dashboards can also cause bureaucrats to focus on metrics over performance. In many cases, bureaucrats have some ability to manipulate the data used to monitor them, and even without manipulation, focusing only on tasks captured in monitoring dashboards can lead bureaucrats to ignore other services valued by citizens.

Dashboards without high-powered incentives and some degree of self-reporting are the most common form of performance management intervention in bureaucracies. Prior research focuses on cases where high-quality third-party data is available to form the basis of often higher-powered performance management systems (Finan et al., 2017; Besley et al., 2022). In this paper, we examine the impact of monitoring the performance of Gram Panchayat Secretaries in the Indian state of Telangana. Gram Panchayats (GPs) are the lowest tier of government in rural India, responsible for delivering local public services and running government schemes.

Each GP is managed by a PS, who is in charge of day-to-day operations and administration. Panchayat Secretaries hire local staff, manage public infrastructure (e.g. roads, drains, streetlights, plantations), oversee public services (e.g. sanitation), interact with citizens (e.g. registering births and deaths), and implement key development programs (e.g. the National

Rural Employment Guarantee Scheme). India has approximately 250,000 GPs serving more than 900 million citizens, each administered by a PS. Despite their prevalence and importance, which has been highlighted in qualitative research (Veeraraghavan, 2021; Rao et al., 2017) and even popular culture (e.g. this popular comedy-drama), Panchayat Secretaries remain virtually unstudied by economists.

We find evidence that on average the dashboards modestly decreased service delivery, as measured from citizen surveys, and increased misreporting, measured by comparing bureaucrat and supervisor reports. These average effects mask substantial individual heterogeneity. Using modern ML methods to construct individual treatment effects, we estimate that the dashboards had a positive impact on 34% of the bureaucrats studied. Although the positive effect is weakly related to the degree of learning, the individual treatment effects themselves are not significantly correlated with any specific observable covariate.

Measuring the performance of a PS is challenging, given the breadth of the job. We worked with the state to design and introduce a mobile application for Panchayat Secretaries to capture or report data across all of their responsibilities¹. Using the self-reported data from the performance management app we determine a performance score for each PS. The score is a weighted average of performance across more than ten responsibility areas. We experimentally evaluate the impact of giving performance scorecards generated using self-reported data. 75% of all Panchayat Secretaries were randomly selected to receive monthly scorecards.

With self-reported data, governments face a trade-off between using the data for performance management and trusting the integrity of data (Singh, 2020)². We evaluate the extent of this trade-off between learning and misreporting by collecting the same self-reported data from alternative sources. First, the mandal supervisors (MPOs) periodically visit villages under their supervision and collect very similar data to the Panchayat Secretaries daily reports. We

¹Not all performance indicators can be captured by devices, but instead will have to be self-reported by the bureaucrats. Like many apps for frontline workers, photos and capturing GPS locations tags were the primary methods to ensure data credibility.

²Collection of outcome data can be delegated to third parties to ensure reliability but may limit external validity as bureaucrats are usually tasked with collecting administrative data on the programs that they administer.

deliberately aligned these reports so that we could construct a measure of “truthful” reporting for the Panchayat Secretaries on the days where their supervisor came for an inspection. Even more rarely, state level visitors (SLVs) visit villages to collect data on a subset of the outcomes covered by the PS and their supervisor. Finally, the department has a call center which periodically conducts phone surveys of citizen satisfaction. The SLV and citizen data allow us to validate the supervisor data used in the main analysis.

The overall result that the scorecards backfired and led to, if anything, reduced performance masks important heterogeneity. First, we show that scorecards helped bureaucrats learn about the important aspects of their jobs and their own weak areas, and those that learned the most did improve their performance. Second, we show that scorecards had similar effects on the performance of tenured and untenured bureaucrats, suggesting that the availability and salience of performance data did not heighten the impact of career concerns in our setting. Third, we show using machine learning methods to estimate individual treatment effects that the performance and inflation results are highly correlated with large groups of bureaucrats reacting to the scorecards by improving performance and not inflating alongside the bureaucrats who drive the average treatment effects who did inflate while performing worse.

Overall, our results highlight conditions under which performance monitoring is most likely to be beneficial and improve bureaucrat performance.

Contribution to Literature.

Our paper contributes to two strands of literature. First, we contribute to a growing literature on the personnel economics of performance measurement and monitoring (See (Finan et al., 2017; Besley et al., 2022) for a review). Examples of measurement include tax collectors (Khan et al., 2016, 2019), teachers (Duflo et al., 2012), agricultural extension workers (Muralidharan et al., 2021) and officials tasked with road construction (Olken, 2007). We contribute to this literature by evaluating the impacts of performance monitoring for jobs that approximate the core civil service. In these multi-faceted jobs, monitoring may have more positive impacts in terms of

fostering learning, but also create more opportunities for multi-tasking and manipulation. We show that monitoring has less sanguine effects in our setting, given the opportunity to misreport data and inflate performance (Jacob and Levitt, 2003; Duflo et al., 2013; Dhaliwal and Hanna, 2017; Mattsson, 2021).

Second, we contribute to the literature on multi-tasking. Prior work has shown that high-powered incentives in government can backfire (Fisman and Wang, 2017; Chen et al., 2018; Acemoglu et al., 2020; Giné et al., 2022; de Janvry et al., 2023). We provide evidence of two types of multi-tasking. We show that Panchayat Secretaries respond to scorecards by misreporting data and inflating their performance. We also find that Panchayat Secretaries improve performance on high-weight indicators while worsening performance on low-weight indicators, including those that citizens value.

The first part of the paper describes the institutional setting and treatments. Then, we discuss the various data sources used and constructed for the study. We offer some initial correlations suggesting that the administrative scores do correlate with citizen opinions. We then show treatment effects on the different performance measures. Finally, we look at heterogeneous treatment effects both traditional and using machine learning. Finally, we discuss the mechanisms suggested by the results and directions for future research.

3.2 Institutional Setting

In this section, we describe the policy that led to the hiring of 9,355 Panchayat Secretaries to implement all public works related to the village assigned. We further describe the roles and responsibilities of Panchayat Secretaries, the performance measurement system, and inspections by Panchayat Secretaries supervisors. In particular, we explain the nature of the inspections and practical challenges that come with PS's job.

Hiring of Panchayat Secretaries

A state is administered through districts and each district is broken down into mandals (blocks). Each Mandal is then made up of several gram panchayats. Gram panchayat is the local level of administration in India. It provides a democratic structure for villages through elected members. Each village with a population of 5,000 or a group of villages with that population form a gram panchayat (GP).

Telangana region in the unified state of Andhra Pradesh was given statehood in 2014. Government of Telangana implemented the "Panchayat Raj" Act in 2018. To monitor and support a gram panchayat, the act allows the state to appoint panchayat secretaries to all gram panchayats. The act also created a number of new Gram Panchayats to bring local administration closer to citizens. Telangana has a population of 35 million and with 33 districts, 594 mandals, and 12,769 GPs.³

To ensure that each GP has a PS, the GoT allowed Commissioner Panchayat Raj & Rural Employment to fill 9,355 vacant posts of Junior Panchayat Secretaries (JPS) in 2018. The Junior Panchayat Secretaries are hired on a remuneration basis for 3 years and they were to be regularized as permanent government employees based on satisfactory performance. The junior panchayat secretaries join existing panchayat secretaries who are already working as frontline bureaucrats. When the study started in April 2021, the state of Telangana had 12,101

³At the time of formation, Telangana had 10 districts, 459 mandals and 8,368 gram panchayats.

panchayat secretaries for 12,769 GPs.⁴ Of the 12,769 GPS, junior secretaries manage 9,081 GPs and permanent secretaries manage 2,839 GPs.⁵

Roles and Responsibilities of Panchayat Secretaries

According to the Act, Panchayat Secretaries are responsible for all public works related to the GP assigned. The panchayat secretaries manage small teams of laborers to clean roads and drains, maintain streetlights, plant greenery, produce compost, and collect garbage. The secretaries are also responsible for managing village finances, collecting property taxes, and providing various certificates and approvals such as birth and death records. The act also defines Key Result Areas (KRA) that are a priority to the government.

Prioritizing key result areas largely takes place at the state level. At the time of the app release, sanitation is by far the most important result area to the administration followed by plantations. The Act specifies specific targets such as 90% working streetlights. Panchayat Secretaries have several general responsibilities in governance, administration, and tax collection, implementation of development schemes. The administration values general administrative result roughly the same as more measurable result areas.

Introduction of the Performance Measurement System

Unlike tax collectors, but similar to other multifaceted government jobs, capturing performance metrics is difficult. The panchayat secretaries may themselves not understand if they are performing well or poorly. We worked with the government of Telangana to design and introduce an Android mobile application for panchayat secretaries to collect and report data across all of their areas of responsibility.

The mobile application captured daily information on the status of randomly selected roads and drains within the GP. On a monthly basis, the application captured information on

⁴668 GPs were given as "full additional charge" or "incharge".

⁵The remaining 849 GPs are either managed by a contract PS or "Outsourcing" PS which have even less job security than junior panchayat secretaries.

GP governance and other aspects of GP performance. Like many apps for frontline workers, photos and capturing GPS location tags were the primary methods to ensure data credibility. Starting December 2020, the commissioner of panchayat raj instructed all panchayat secretaries of the state to use the mobile application for the reporting their daily, weekly and monthly responsibilities.

GP Inspections

All panchayat secretaries within a mandal are supervised by the Mandal Panchayat Officer (MPO). Recall that the Panchayat Raj Act 2018 states a JPS will be hired on probationary basis for 3 years and would be regularized based on satisfactory performance. The Panchayat Raj Act 2018 also instructs the MPOs to conduct GP inspections. This was the administration's plan to gather the performance indicators needed for the JPS regularization decision. Importantly, there is substantial heterogeneity in the number of GPs per mandal. MPOs can manage anywhere between 3 and 32 Panchayat Secretaries generating substantial variation in the intensity of supervision and frequency of inspection.

All MPOs within a district in turn report to the District Panchayat Officer (DPOs). The commissioner of Panchayat Raj, through DPOs, implements and supervises the administration's efforts to support and improve various aspects of GPs. The administration had also rolled out a system of state level visitors (SLVs) who periodically inspected GP to assess them. These visitors are generally retired officials which were specifically recruited for the purpose of auditing local projects. Their data serves as the external performance indicators of GPs and came to be particularly valued by senior state leadership for tracking progress. However, this external data cannot be used in the JPS regularization decision as the Act requires that only the Panchayat Raj Department's data (such as MPO inspections) shall be considered for regularization decision. Also, with only 103 state level visitors and other inspection responsibilities these state level visitors (SLVs) could only cover approximately 2,000 GPs per month which means the data couldn't be used to track GP performance at high frequency.

Practical Challenges that come with Panchayat Secretaries Job

Similar to how a state has a chief minister elected by the people and a bureaucrat (IAS) appointed by the state, a GP has a sarpanch (elected) and a PS (appointed). Sarpanchs and Panchayat Secretaries work very closely together for all GP related activities. However, their mixture of cooperation and conflict is described by some officials as akin to a wife and mother-in-law in a household. Aside from controlling the budget, the sarpanchs influence contracting and hiring for basic public works. There are a key village resources that sarpanch's can co-opt for example access to a tractor which is necessary for successful plantations. Similarly, having a computer in the GP is a key piece of enabling infrastructure.

There are a variety of socio-economic factors that affect PSs ability to achieve the key results enumerated in the Panchayat Raj act. GPs with large SC/ST populations can be more difficult to administer. For example, village sanitation personnel (even if they are SC/ST) are often unwilling to clean the SC/ST hamlet because it is beneath them. The physical size of a GP is important. The funding doesn't really take into account the physical spread of a GP and it can be physically impossible to clean larger GPs with the allotted funds. Finally, junior panchayat secretaries' own lack of training makes achieving the results envisioned in the act difficult. The government has not carried out a systematic training program for the JPSs.

The incentive for a JPS to perform well at job is clear - permanent employment. For panchayat secretaries on permanent contracts transfers, remain a powerful to reward or punish panchayat secretaries. Since the JPSs are younger and have less training than can be more easily bossed around by the Sarpanchs. The more permanent PSs are sometimes more powerful than the sarpanchs since they have more education. The permanent Panchayat Secretaries primary incentive is transfers to larger, more resourced, or more welcoming GPs. As mentioned above there is substantial heterogeneity in the difficulty of the PS job across GPs. Some GPs have substantial tax revenue which makes the Panchayat Secretaries job easier since they can hire more staff. Panchayat Secretaries cannot be directly promoted to MPOs. They must take another

competitive examination to be eligible for the MPO job.

3.3 Treatment

3.3.1 Creating Scorecards

Using the self-reported data from the app, we determine a performance score for each GP and assign that score to the PS. The score, which is out of 100, is a weighted average of performance across more than ten responsibility areas. Recall that the PS is responsible for all GP related tasks. Table x lists the composition of the overall score. The 100 points is distributed across y key result areas as follows: 40 points for sanitation, 8 points for streetlights, 2 points for water tank, 20 points for greenery, 20 points for GP administration and 10 points for NREGS. See appendix 3.10.3 for the exact rubric.

Table 3.1. GP performance/progress indicators - change this to a figure

Key Result Areas (PR Act)		Notes	Points
1	Sanitation	Road, Drain, and Institutional Cleanliness	40
2	Street lights	Installation, maintenance, performance	8
3	Water Tanks	Water tank cleanliness	2
4	Plantations	Nursery, plantations, survival	20
5	Administration	Meetings, Certificates, and Registers	12
6	NREGA	Days worked per active job card	18

3.3.2 Sharing Scorecards

(J)PSs of GPs in the treatment group received additional screens in their mobile applications which displayed the GP's overall score and rank based on the information Panchayat Secretaries had entered over the course of the month Figure 3.1 ⁶. This setting allows us to experimentally evaluate the effect of providing Panchayat Secretaries performance-based feedback

⁶While the value of each indicator was clearly displayed within the application itself, the actual method of score computation was a separate PDF document that panchayat secretaries had to access. Consequently, not all Panchayat secretaries were immediately aware that their own data was being used to generate the scores.




Figure 3.1. Screenshots from PS App

via app-based scorecards that show the PS her absolute level of performance, highlight areas for improvement, and convey relative performance within the district.

We hypothesize that the scorecards might improve true performance through providing information that helps Panchayat Secretaries learn how to do better, as well as by providing incentives to perform well and stand out as a top performer (or at least not stand out as a bottom performer). The app scorecards also displayed the relative rank band of panchayat secretaries within their district and the names of the top 20 panchayat secretaries in the district.

3.3.3 Validating Self-Reported Data Through MPO Inspections

We examine whether performance-based feedback actually improves true performance or merely leads to misreporting. Misreporting occurs when secretaries fudge self-reported performance data to score better on a metric that now has higher stakes. We leverage on MPO inspections to capture misreporting. In a mobile app designed for inspections, we also ask MPOs to enter data on all PS indicators that can be compared/verified. Additionally, we also modified the state-level visitors inspection app to include the same indicators to use as external audit.

However, since the external visitor app was built by a different technology contractor many of the crucial sanitation questions which relied on road level answers couldn't be used for direct comparison.

3.4 Randomization and Timeline

The mobile application for the PS to self-report on duties was launched in December 2020. It took some time to fix bugs and drive usage from January to March 2021. By April 2021 we were able to generate scores from the app. However, the department did not share a scorecard for April since that corresponded with India's intense but brief Delta COVID wave. The first scorecard was released for May on June 27th. Since some indicators are entered at the end of the month we still consider June part of the treatment period, but July was the first full month where treatment Panchayat Secretaries were able to react to scorecards. The July, August, and September scorecards were released closer to the end of their respective months. In October, there was a change in leadership of the Panchayat Raj department. The new leadership wanted a large number of changes in the application which resulted in data collection temporarily pausing. The app, with some significant modifications focusing on PS attendance, relaunched in November 2021. The department was about to relaunch the scorecards but then ran into contracting issues with the contractor for the mobile application. The mobile application stopped working again in May 2022 and was recently restarted under a new tech contractor in September 2022.

We conducted a stratified randomization at the mandal level using district and number of GPs per mandal as strata. 9,584 GPs in 405 mandals received scorecards and 3,186 GPs in 134 mandals did not receive scorecards. We did not stratify on baseline scores since the department required that we randomize prior to baseline scores becoming available. Further, since randomization took place at the mandal level averages of GP level characteristics were less useful. Mandal level randomization was necessary since Panchayat Secretaries within a mandal meet frequently and providing scorecards to some Panchayat Secretaries and not others would sow confusion and distort supervisor effort. The mandal size strata was chosen since we believed the intensity of supervisory visit would strongly influence PS's response to scorecards. We chose a split of 75% and 25% control based on plans to layer on further cross-randomized treatments.

However, these plans were disrupted by the October leadership change discussed above.

3.5 Data

3.5.1 Validation with Citizen Surveys

To validate the performance scores we conducted citizen phone surveys with numbers sampled from government databases. In total we sampled 26,775 citizens across 3,825 GPs. In these phone surveys we asked citizens to rate how their GP was performing on different result areas themselves using a simple five point likert scale. Table 3.2 shows that the PS scores, pooled across all months, correlated significantly with the citizen feedback. The scores awarded by inspectors (SLVs as well as MPOs) also correlates significantly with the citizen satisfaction score. This suggests that the inspections and the scores awarded by inspectors are a signal of PS performance on the ground.

Table 3.2. Correlation Matrix: PS Performance Measures

(1)				
	PS Score	MPO Score	SLV Score	Citizen Satisfaction Score
PS Score	1			
MPO Score	0.183***	1		
SLV Score	0.0310***	0.0305***	1	
Citizen Satisfaction Score	0.0495***	0.0462***	0.0125**	1

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

3.5.2 “Truth” Scores

To measure the level of “true” PS performance, we start by mapping the KPIs under the PS self-reported measures to the indicators under MPO inspections. It is important to note

that some KPIs cannot be verified for veracity accurately due to differing frequency of data collection. For example, while the Panchayat Secretaries self-report the cleaning frequency of roads and drains in their panchayat, the MPOs cannot ascertain the cleaning frequency during their inspection visit. The MPOs can only ascertain the state of cleanliness of roads and drains during their visit.

This mapping exercise gives us an estimate of the “truthable” PS score; the proportion of the PS score can be accurately verified as true PS performance by the MPOs. The remaining portion of the self-reported PS score can be attributed as inflation on the part of Panchayat Secretaries, to misreport their performance. Hence, we define score inflation as the difference between self-reported PS score and the score awarded to the same PS by the MPO, based on their inspections.

Missing data on KPIs, either from the PS or from the MPO, poses a challenge in our ability to validate PS performance. To account for this, we define an “adjusted PS score” on the basis of whether both the PS and MPO have reported a score for each KPI. A KPI is said to be “non-truthable” if the MPO has not entered a score for it, during the inspection. So the respective portion of the PS score for that indicator is subtracted from the sum total of the PS score, to obtain the adjusted PS score. Out of a total PS score worth 100 points we determined 52 points were theoretically “truthable”. Missing data further reduced the average “truthable” score to 34. For the SLV data the form was even more limited so there were only 18 points theoretically “truthable.” Missing data further reduced the average truthable score is 13.97.

3.5.3 Score Inflation

Since the self-reported performance score is prone to inflation, we define and use three complementary measures of score inflation to monitor actual PS performance on the ground.

Total inflation is the percentage-point difference between the PS score percentage (PS points as a share of the total possible points on every KPI the PS reports in that period) and the MPO score percentage (MPO points as a share of the total possible points on every KPI rated by

the MPO in that period).

$$\text{TotalInfl}(\%) = \left(\frac{\text{PS Score}}{\text{Max points on all PS KPIs}} - \frac{\text{MPO Score}}{\text{Max points on all MPO KPIs}} \right) \times 100$$

Verified inflation is the difference between the adjusted PS score and the adjusted MPO score computed only on KPIs observed in both sources during the same period, and is scaled by the total “truthable” points, where the truthable points are calculated as the sum total of all KPIs for which there are non-missing observations in the PS and MPO data.

$$\text{VerifiedInfl}(\%) = \left(\frac{\text{Adjusted PS score} - \text{Adjusted MPO score}}{\text{Total truthable points}} \right) \times 100$$

Finally, the *unverified inflation* proxies the inflation in the PS score that comes from KPIs without a counterpart in the MPO score. It equals the PS score on these unmatched KPIs expressed as a percentage of the total points available for those KPIs, minus the overall MPO score expressed on the same percentage-point scale.

$$\text{UnverifiedInfl}(\%) = \left(\frac{\text{PS points on PS-only KPIs}}{\text{Max points on PS-only KPIs}} \times 100 \right) - \text{MPO score } \%$$

Due to the carefully constructed administrative data sources discussed above we are able to measure score inflation directly. We find that there is substantial inflation in our setting. In GPs that didn’t receive the scorecard, the median total inflation was 15 percentage points.

3.5.4 Effort Metric

A unique feature of the mobile application is the requirement that panchayat secretaries submit a daily sanitation report which the app would only allow to submit if the phone’s GPS was within the village. While, in practice, some Panchayat Secretaries were able to find subordinates or other technical workarounds, reporting is a good proxy for PS effort. Daily physical presence

at the GP is a crucial and costly part of the job since nearly all Panchayat Secretaries do not reside in the villages in which they serve.

3.5.5 Balance

The treatment and control groups are generally balanced on observables. Importantly, the experiment is balanced with respect to the outcome variables in the April “baseline” period. Note that due to government hiring rules there are a substantial amount of Panchayat Secretaries that are woman or come from an SC/ST background. These Panchayat Secretaries from marginalized backgrounds may face further challenges dealing with village elected leaders.

Table 3.3. Balance Table

	Treatment		Control		Diff. (SE)
	Obs. (Cluster)	Mean (SD)	Obs. (Cluster)	Mean (SD)	
Panel A. Baseline Scores					
Supervisor Assessment (%)	9,482 (405)	67.64 (11.75)	3,126 (134)	67.07 (11.98)	0.569 (0.86)
Self-reported Assessment (%)	9,584 (405)	68.06 (7.56)	3,186 (135)	68.43 (7.16)	-0.377 (0.42)
Citizen Assessment (SD)	719 (325)	-0.20 (0.97)	249 (112)	-0.37 (1.03)	0.170** (0.07)
Auditor Assessment (%)	3,862 (404)	71.17 (12.81)	1,317 (133)	71.25 (12.17)	-0.076 (0.94)
Reporting Rate (%)	9,584 (405)	69.49 (8.03)	3,186 (135)	69.39 (8.02)	0.097 (0.48)
Total Score Inflation (%)	9,482 (405)	0.84 (13.64)	3,126 (134)	1.89 (14.19)	-1.054 (0.97)
Score Inflation (%)	9,482 (405)	9.19 (15.14)	3,126 (134)	9.92 (15.07)	-0.729 (0.88)
Unverified Score Inflation (%)	8,441 (405)	9.44 (27.18)	2,773 (134)	9.32 (28.23)	0.120 (1.67)
Panel B. PS Demographics					
Junior Panchayat Secretary (%)	9,584 (405)	71.01 (45.37)	3,186 (135)	71.41 (45.19)	-0.392 (1.26)
Female (%)	9,154 (405)	28.69 (45.23)	3,024 (135)	30.82 (46.18)	-2.133* (1.28)
Scheduled Caste (%)	9,154 (405)	18.39 (38.74)	3,024 (135)	17.56 (38.05)	0.826 (0.97)
Age (Years)	8,079 (405)	33.59 (6.71)	2,713 (135)	33.65 (6.72)	-0.066 (0.20)
Years in Service	8,070 (405)	4.44 (5.60)	2,709 (135)	4.30 (5.27)	0.141 (0.16)
Panel C. Village Characteristics					
Total Population (2011)	9,579 (405)	1,579.10 (1,345.04)	3,181 (135)	1,659.48 (1,608.09)	-80.38* (43.85)
Rural Monthly Consumption per capita (2011-12)	9,264 (405)	22,276.28 (4,303.33)	3,069 (134)	22,617.29 (4,308.57)	-341.01 (289.68)
Mean Night-time Light Index (2020)	9,379 (405)	0.78 (0.89)	3,109 (135)	0.81 (1.00)	-0.03 (0.05)
Wealth Index (Facebook Predicted) 2021	9,375 (405)	-0.15 (0.25)	3,107 (135)	-0.14 (0.26)	-0.02 (0.02)
Share Completed Secondary Education (2012)	9,287 (405)	0.18 (0.08)	3,077 (135)	0.19 (0.08)	-0.01 (0.01)
Education Index	9,273 (405)	-0.00 (0.38)	3,075 (134)	0.01 (0.37)	-0.02 (0.02)
Health Index	9,273 (405)	-0.00 (0.39)	3,075 (134)	0.00 (0.35)	-0.00 (0.02)
Sanitation Index	9,273 (405)	-0.03 (0.39)	3,075 (134)	0.01 (0.38)	-0.04 (0.02)
Connectivity Index	9,273 (405)	-0.00 (0.41)	3,075 (134)	0.01 (0.40)	-0.01 (0.02)
No. of Villages in Mandal	405 (405)	23.66 (8.70)	135 (135)	23.60 (9.18)	0.06 (0.90)

3.6 Main Results

Our primary specification is a basic difference-in-difference style estimator with standard errors clustered at the mandal level. We control for various covariates: PS gender, years of work experience as PS, PS caste, posting type of PS, MPO gender, years of experience as MPO, and posting type of MPO.

The difference in difference style estimator pools all of the treatment months together. For robustness we also run the event study style estimator. We find there are no significant month by month differences in the treatment effects on scores (Table 3.7). The inflation effects are less stable, but exhibit no clear trend (Table ??).

We observe a GP i in month t in strata s . The treatment months June, July, August, and September are represented by T . The preferred specification is:

$$Y_{it} = \alpha + \beta_1 Treat * 1_{t \in T} + \beta_2 1_{t \in T} + \beta_3 X + \mu_t + \mu_m$$

where μ_t are month fixed effects and μ_m are mandal fixed effects.

We find significant but small increases in panchayat secretary scores. Note that while on a scale out of 100 an increase of less than 1 may seem insignificant the scores are highly clustered. The inter-quartile range is just 13 from 67 to 80 and the standard deviation is 10. However, the scorecards effect seems to be contained to the Panchayat Secretaries themselves. The other bureaucrats inspecting GPs, MPOs and SLVs, do not report any differences in score between treatment and control GPs. In fact, the coefficients are, if anything, slightly negative. The negative effect on official scores matches the scorecards negative impact on citizen satisfaction. The effect translates to a 4% decrease relative to the control mean.

The logical explanation for the divergence in treatment effects between what the Panchayat Secretaries report and the citizen assessment is that the treatment induces Panchayat Secretaries to inflate their scores. We do in fact see a statistically significant increase in total

Table 3.4. Main Effects

Panel A:	(1) Self-Reported Assessment	(2) Supervisor Assessment	(3) Auditor Assessment	(4) Citizen Assessment
Score-cards x Post	1.019*** (0.392)	-0.377 (0.756)	-0.539 (0.840)	-0.145** (0.074)
Observations	255,300	111,253	17,363	11,138
Clusters	540	540	539	539
Mandal FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Control Mean	79	62.7	78.3	.02
Panel B:	(1) Total Inflation	(2) Verified Inflation	(3) Unverified Inflation	(4) Effort Metric
Score-cards x Post	1.707** (0.707)	1.387 (0.875)	1.328 (1.425)	-0.206 (0.439)
Observations	89,330	89,330	73,622	1,059,910
Clusters	540	540	540	540
Mandal FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Control Mean	16.1	23.99	19.61	65

Column (1) are the PS scores from the app which theoretically range from 0 to 100, but practically have an inter-quartile range of 66 to 80. Supervisor assessments are the scores from the MPO app on a similar scale. Auditor assessments are scores from SLV's. Citizen assessments are scores from citizen phone surveys that simply asked for satisfaction on a 1-5 scale. The inflation measures are discussed in detail in section 4.3. The effort metric is a measure of whether the Panchayat Secretaries attended their GPs every day.

inflation and the coefficients on the other inflation measures are also positive.

3.7 Discussion

There is evidence that the scorecard intervention did help Panchayat Secretaries learn more about their relative performance. In the April PS survey, Panchayat Secretaries self-perceived performance bore no relationship to the PS scores. After the scorecards, Panchayat Secretaries surveyed in July in the scorecard groups had scores that more closely matched their perceptions. Thus the scorecards did help some Panchayat Secretaries better internalize their relative performance in the eyes of the department. We see that precisely the panchayat secretaries that had the worst most misinformed about their performance at baseline that were the most likely to respond to the socrecards by performing better as measured by the MPO scores and the effort metric (Table 3.5).

score inflation in general does not appear to be strongly correlated with observables such as gender, caste, and type of posting. Regressions examining the correlates of inflation can be found in Appendix 3.12.

Differences in career incentives surprisingly do not drive the treatment effects. Junior panchayat secretaries are just as likely to reduce actual performance and inflate the self-reported scores as panchayat secretaries with more job security (Table 3.6). Both JPS and PS's broadly benefit from higher scores and the gradations of that benefit may matter less than differences between GPs and secretaries in terms of opportunities to inflate scores or change performance.

While demographic variables might not strongly predict heterogeneous effects on inflation individually, that does not mean there is not significant heterogeneity in treatment effects across Panchayat Secretaries. In praticular, there may be complicated combinations of PS and MPO characteristics and baseline scores that are highly predictive. To test for this we followed the generic machine learning method proposed in (Chernozhukov et al., 2017). The details of the method are discussed in appendix 3.10.4. The basic finding was that the ML algorithm can sort Panchayat Secretaries into groups that both significantly increased and decreased their scores in response to the scorecards treatment.




Figure 3.2. Scorecards Correct Misperception




Figure 3.3. Generic ML heterogeneity

Table 3.5. Heterogeneity of main treatment effects by baseline misperception

Panel A:	(1) Self-Reported Assessment	(2) Supervisor Assessment	(3) Auditor Assessment	(4) Citizen Assessment
Score-cards x Post x Misperception	1.194** (0.576)	2.447** (1.021)	-0.479 (1.925)	0.355 (0.255)
Post x Misperception	2.724*** (0.445)	-1.542* (0.871)	0.813 (1.534)	-0.293 (0.210)
Post x Score-cards	0.344 (0.430)	-1.556 (1.008)	0.344 (1.556)	-0.340** (0.170)
Observations	55,000	26,476	3,353	3,752
Clusters	538	538	430	482
Panchayat FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Panel B:	(1) Total Inflation	(2) Verified Inflation	(3) Unverified Inflation	(4) Effort Metric
Score-cards x Post x Misperception	-1.339 (1.127)	-0.431 (1.203)	-3.270 (2.778)	1.254* (0.680)
Post x Misperception	4.237*** (0.937)	3.255*** (0.896)	6.431*** (2.275)	1.019** (0.515)
Post x Score-cards	2.315** (0.957)	0.984 (1.104)	2.741 (2.177)	-0.811 (0.507)
Observations	21,200	21,200	17,334	228,250
Clusters	538	538	537	538
Panchayat FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes

Table 3.6. Treatment Effects for JPS

Panel A:	(1) Self-Reported Assessment	(2) Supervisor Assessment	(3) Auditor Assessment	(4) Citizen Assessment
Score-cards x Post x JPS	0.038 (0.397)	-0.603 (0.599)	0.014 (1.015)	0.140 (0.159)
Post x JPS	0.345 (0.342)	0.245 (0.535)	0.002 (0.873)	-0.027 (0.137)
Post x Score-cards	0.993** (0.427)	0.048 (0.893)	-0.553 (1.239)	-0.247* (0.136)
Observations	255,300	111,253	17,363	11,138
Clusters	540	540	539	539
Mandal FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Panel B:	(1) Total Inflation	(2) Verified Inflation	(3) Unverified Inflation	(4) Effort Metric
Score-cards x Post x JPS	1.071 (0.669)	1.167 (0.842)	1.096 (1.539)	-0.513 (0.446)
Post x JPS	-0.415 (0.588)	-1.421* (0.744)	2.284* (1.341)	0.284 (0.390)
Post x Score-cards	0.952 (0.876)	0.566 (1.074)	0.549 (1.876)	0.159 (0.555)
Observations	89,330	89,330	73,622	1,059,910
Clusters	540	540	540	540
Mandal FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes




Figure 3.4. Relationship between predicted treatment effects

Table 3.7. Predicted treatment effects

Category	Performance	Data Accuracy	Share
Clearly positive	+	+	18%
Weakly positive	+	0	7%
Weakly positive	0	+	9%
Positive impact			34%
Tradeoff	+	-	4%
Tradeoff	-	+	2%
Weakly negative	-	0	13%
Weakly negative	0	-	12%
Clearly negative	-	-	15%
Negative impact			40%
No impact	0	0	19%

Notably, none of the variables used to generate individual treatment effects are good predictors of being in a high or low treatment group by themselves when comparing the standardized difference between the highest and lowest treatment effect groups for each of the predictors used (3.10.1).

The correlation of individual treatment effects across outcomes, but limited direct relationship with observables, highlights the challenge of performance management in a bureaucracy. Employees can differ dramatically in earnestness and intrinsic motivation, resulting in heterogeneous responses to performance management, but it is difficult for policymakers to classify employees into motivation types from observables alone.

3.8 Conclusion

Do bureaucrats use self-reported scores as an opportunity to make themselves look better in the eyes of their departments or as a tool to learn about their own performance? Our study suggests bureaucrats use scoring systems in both ways. While some bureaucrats were content to simply learn about their performance, others, especially those who learn their performance is relatively low, took the next step of “correcting” their scores in the following month. The fact that not all bureaucrats inflated their scores to the same degree makes it more difficult to confidently target resources and track progress using the self-reported data. One promising solution is to incentivize accuracy in addition to performance. Rather than simply rewarding bureaucrats for having the highest scores, governments should also reward bureaucrats whose data stands up to the scrutiny when inspected. If bureaucrats can inflate their scores in response to simple information about relative performance, then they can also report honestly in response to simple information about their relative accuracy. Policy solutions to encourage accurate self-reporting are important. Governments will struggle to deliver services at scale until they can trust and use data generated from the bureaucrats implementing their programs.

3.9 Chapter Acknowledgments

Chapter 3, in full, is currently being prepared for submission for publication of the material. Arun, Advitha; Brownstone, Steven; George, Siddhart; Kumar, Naveen; Muralidharan, Karthik are co-authors of this chapter. The dissertation author was one of the primary investigators and authors of this material.

3.10 Appendix

3.10.1 Covariate Predictors of Individual Treatment Effects




Figure 3.5. Predictors of Low vs High Individual Inflation Treatment Effect

The graph plots the standardized difference between Panchayat Secretaries in the highest and lowest quintile of predicted treatment effects for each baseline characteristic used in the random forest. Each dot plots one split for a given covariate. Large circles with white centers plot the median. Colored covariates have at least 90% or more splits above (or below) zero; gray covariates do not




Figure 3.6. Predictors of Low vs High Supervisor Score Treatment Effect

The graph plots the standardized difference between Panchayat Secretaries in the highest and lowest quintile of predicted treatment effects for each baseline characteristic used in the random forest. Each dot plots one split for a given covariate. Large circles with white centers plot the median. Colored covariates have at least 90% or more splits above (or below) zero; gray covariates do not

3.10.2 Treatment Effect by Month




Figure 3.7. Effect of treatment on PS Score By month




Figure 3.8. Effect of treatment on Inflation Score By Month

3.10.3 PS Score Rubric

KRA 1: Sanitation (40 points)

1. Road cleaning frequency (4 pts): Proportional to the number of roads cleaned each day.
The score will be calculated only taking into consideration the 18 best days in the month.
2. Road cleaning (7 pts): Proportional to the percentage of randomly selected (DSR) roads which are clean. The score will be calculated only taking into consideration the 18 best days in the month.
3. Drain cleaning frequency (4 pts): Proportional to the number of drains cleaned each day.
The score will be calculated only taking into consideration the 18 best days in the month.
4. Drain cleaning (7 pts): Proportional to the percentage of randomly selected (DSR) drains which are clean. The score will be calculated only taking into consideration the 18 best days in the month.
5. Institution cleanliness (6 pts): Proportional to the percentage of randomly selected (DSR) institutions which are clean. The score will be calculated only taking into consideration the 18 best days in the month.
6. Garbage transport (4 pts): Proportional to the days in the month when garbage is transported to the segregation shed/ dumpyard. The score will be calculated only taking into consideration the 18 best days in the month.
7. Household waste segregation (4 pts): Proportional to the number of days in the month when segregated waste is collected from HHs (measured as days when segregated waste was collected from at least one HH)
8. Compost preparation (4 pts): 2 points if compost facilities present; 2 points if compost is prepared (from MAS)

KRA 2: Streetlights (8 points)

1. Streetlights (8 pts): Streetlights scores are based on the share of streetlights which are functional on an average day in the month (in the best 18 days). If this average is lower than 0.6, the score is 0, and increases linearly beyond that.

KRA 4: Greenery (20 points)

1. Nursery infrastructure (6 pts): 2 points each if the village nursery has watering facility, cattle trap, fencing.
2. Plantation watering (5 pts): Equal points (5/3) given if each of the 3 plantation types are watered regularly.
3. Plantation fencing (2 pts): Equal points (2/3) given if each of the 3 plantation types have fencing
4. Plantation survival (7 pts): Proportional to plantation survival rates, equally divided for the 3 plantation types.

KRA 5: Administration (12 points)

1. GP/GS meetings (3 pts): 2 points awarded if GP meeting held in the past month, 1 point awarded if GS meeting held in the past 2 months.
2. Register maintenance (2 pts): Proportional to the number of registers which are up to date.
3. Birth and death registration and certification (3 pts):

0.75 points each are proportional to the number of birth and death registered within timelines

0.75 points each are proportional to the number of birth and death certificates issued within timelines

4. E-panchayat computerization (2 pts): Proportional to the number of days in the month that the DSR has been filled.
5. CC charges paid on time (2 pts): Proportional to the share of bill amount for which cheque has been generated.

KRA 6: NREGA (18 points)

1. NREGA (18 pts): Proportional to the days of work per reported wage seeker. Graded relative to the 90th percentile GP in the mandal (GPs at and above the 90th percentile get full marks). GPs with less than or equal to 5 reported wage seekers are not eligible for NREGA score, and their overall PS score will be graded out of 82 instead of 100.

3.10.4 Generic ML

We estimate heterogeneous treatment effects using generic machine learning, a method proposed in (Chernozhukov et al., 2017). This is a tool-independent procedure to facilitate inference of individual-specific heterogeneous treatment effects in RCT settings. In this section, we briefly describe the generic ML procedure and highlight the results that emerge from applying this method in our setting.

The main stages of GenericML are: (i) estimate HTE via machine learning models with train/test data split, (ii) compute linear homogeneous TE proxies (BLP/CATE) from the machine learning models, (iii) rank observations by the predicted effect size (GATES); (ic) test difference in baseline characteristics (CLAN) between groups with the highest and lowest effect sizes.

GenericML makes the HTE estimation possible by predicting the counterfactual outcomes with ML models. It does so by repeating the following steps with a set of candidate ML models:

1. randomly 50/50 split the full data set into a main set M and an auxiliary set A . A and M both contain treated and untreated subject

2. train ML models with only samples from M to formulate the relationship between baseline characteristics and the outcome. In particular, we want learn such relationship respectively in the control and the treatment group.
3. Having formulated the relationships, we then can ‘see’ the counterfactual outcome for each individual in A (who was not involved in training models). Simply, for each individual a in A , we can predict its: treated outcome (by plugging its baseline characteristics in f_t , denote the predicted outcome as $f_t(a)$), control group outcome (by plugging in f_c , denote as $f_c(a)$) and predicted treatment effect (by taking the difference, $HTE = f_t(a) - f_c(a)$).

After repeating the above steps with multiple random A/M splitting each time, we have the statistics of HTE derived by apply different ML model, and move to the next step - horse racing between candidate models.

4. We obtain HTE estimates for each observation by pretending it is either in control and treatment group. However, each observation is actually either in the treatment or the control group, and its observed outcome is known. We can therefore assess prediction quality of each model by comparing the observed to the predicted outcome. Averaging such comparison across different individuals lead to a performance indicator for a model. We then proceed with the model of the highest performance indicator.
5. GATES: having fixed the best model, we predict HTE with that model and sort the individuals according to predicted size of treatment effect. We then split the individuals into 5 groups ranked by effect size; the average treatment effect is then estimated for each group.

GenericML has two main advantages over simpler methods of estimating heterogeneous treatment effects, such as including interactions between treatment dummies and particular covariates. First, GenericML captures more complex heterogeneity that sequential addition of interaction terms

By horse-racing multiple prediction models and including non-parametric models, GenericML enables data-driven model choice. Instead of discretionary choice of model, GenericML provides a data-driven approach to automated model selection. Moreover, GenericML pipeline makes the horse-racing of models feasible and, more importantly, credible by enforcing a shared performance indicator, based on which we decide which model is the best. With inclusion of non-parametric ML models, GenericML features better robustness in HTE estimation against high-dimensional data where the true data generating process is indeed high-dimensional. Consider the old-school way of detecting heterogeneity: by introducing interaction terms in regression analysis. When there are many regressors, fully saturated regressions are likely to be overfit and absorb too much noise.

3.11 Correlates of Supervisor Schore

Table 3.8. Correlates of MPO Score

	Supervisor Assessment			
	(1)	(2)	(3)	(4)
PS Scheduled Caste	-0.003 (0.003)			-0.003 (0.003)
Junior PS		0.004 (0.003)		0.002 (0.003)
Female PS			0.003 (0.003)	0.003 (0.003)
Observations	105893	111292	105893	105893
R^2	.14	.14	.14	.14
Month FE	Yes	Yes	Yes	Yes

3.12 Correlates of Inflation

Table 3.9. Effect of PS' Caste on score inflation percentage

	Total Inflation (1)	Verified Inflation (2)	Unverified Inflation (3)
PS Scheduled Caste	0.000 (0.003)	-0.007** (0.003)	0.002 (0.004)
Junior PS	0.030*** (0.003)	0.009*** (0.003)	0.028*** (0.004)
Female PS	-0.010*** (0.003)	-0.017*** (0.003)	-0.005 (0.004)
Observations	85014	85014	70048
R^2	.27	.21	.084
Month FE	Yes	Yes	Yes

Bibliography

Acemoglu, Daron, Leopoldo Fergusson, James Robinson, Dario Romero, and Juan F Vargas, “The perils of high-powered incentives: evidence from Colombia’s false positives,” *American Economic Journal: Economic Policy*, 2020, 12 (3), 1–43.

Besley, Timothy, Robin Burgess, Adnan Khan, and Guo Xu, “Bureaucracy and development,” *Annual Review of Economics*, 2022, 14, 397–424.

Chen, Yvonne Jie, Pei Li, and Yi Lu, “Career concerns and multitasking local bureaucrats: Evidence of a target-based performance evaluation system in China,” *Journal of Development Economics*, 2018, 133, 84–101.

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Iván Fernández-Val, “Generic Machine Learning Inference on Heterogenous Treatment Effects in Randomized Experiments,” 2017.

de Janvry, Alain, Guojun He, Elisabeth Sadoulet, Shaoda Wang, and Qiong Zhang, “Subjective Performance Evaluation, Influence Activities, and Bureaucratic Work Behavior: Evidence from China,” *American Economic Review*, March 2023, 113 (3), 766–99.

Dhaliwal, Iqbal and Rema Hanna, “The devil is in the details: The successes and limitations of bureaucratic reform in India,” *Journal of Development Economics*, 2017, 124, 1–21.

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan, “Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from India,” *The Quarterly Journal of Economics*, 2013, 128 (4), 1499–1545.

—, **Rema Hanna, and Stephen P Ryan**, “Incentives work: Getting teachers to come to school,” *American Economic Review*, 2012, 102 (4), 1241–78.

Finan, Frederico, Benjamin A Olken, and Rohini Pande, “The personnel economics of the developing state,” *Handbook of economic field experiments*, 2017, 2, 467–514.

Fisman, Raymond and Yongxiang Wang, “The distortionary effects of incentives in government: Evidence from China’s” death ceiling” program,” *American Economic Journal: Applied*

Economics, 2017, 9 (2), 202–18.

Giné, Xavier, Ghazala Mansuri, and Slesh A Shrestha, “Mission and the bottom line: Performance incentives in a multigoal organization,” *Review of Economics and Statistics*, 2022, 104 (4), 748–763.

Jacob, Brian A and Steven D Levitt, “Rotten apples: An investigation of the prevalence and predictors of teacher cheating,” *The Quarterly Journal of Economics*, 2003, 118 (3), 843–877.

Khan, Adnan Q, Asim I Khwaja, and Benjamin A Olken, “Tax farming redux: Experimental evidence on performance pay for tax collectors,” *The Quarterly Journal of Economics*, 2016, 131 (1), 219–271.

—, **Asim Ijaz Khwaja, and Benjamin A Olken**, “Making moves matter: Experimental evidence on incentivizing bureaucrats through performance-based postings,” *American Economic Review*, 2019, 109 (1), 237–70.

Mattsson, Martin, “Information systems, service delivery, and corruption: Evidence from the bangladesh civil service,” *Service Delivery, and Corruption: Evidence from the Bangladesh Civil Service (December 16, 2021)*, 2021.

Muralidharan, Karthik, Paul Niehaus, Sandip Sukhtankar, and Jeffrey Weaver, “Improving last-mile service delivery using phone-based monitoring,” *American Economic Journal: Applied Economics*, 2021, 13 (2), 52–82.

Olken, Benjamin A, “Monitoring corruption: evidence from a field experiment in Indonesia,” *Journal of political Economy*, 2007, 115 (2), 200–249.

Rao, Vijayendra, Kripa Ananthpur, and Kabir Malik, “The anatomy of failure: An ethnography of a randomized trial to deepen democracy in rural India,” *World development*, 2017, 99, 481–497.

Rasul, Imran and Daniel Rogger, “Management of Bureaucrats and Public Service Delivery: Evidence from the Nigerian Civil Service,” *The Economic Journal*, 04 2017, 128 (608), 413–446.

Singh, Abhijeet, “Myths of official measurement: Auditing and improving administrative data in developing countries,” *Research on Improving Systems of Education (RISE) Working Paper*, 2020, 42.

Veeraraghavan, Rajesh, *Patching development: Information politics and social change in India*, Oxford University Press, 2021.