

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Essays in Development Economics

Permalink

<https://escholarship.org/uc/item/5wn5n2hc>

ISBN

9798189275814

Author

Sarkar, Shreya

Publication Date

2026-05-01

Peer reviewed|Thesis/dissertation

Essays in Development Economics

by

Shreya Sarkar

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

Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Jeremy Magruder, Chair

Professor Supreet Kaur

Professor Aprajit Mahajan

Professor Gautam Rao

Spring 2026

Essays in Development Economics

Copyright © 2026

by

Shreya Sarkar

Abstract

Essays in Development Economics

by

Shreya Sarkar

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Jeremy Magruder, Chair

This dissertation comprises three essays on why workers and households in low-income settings underinvest in welfare-improving decisions, despite potentially high returns. A unifying answer emerges: frictions that concentrate costs at the moment of action impose large and regressive burdens on those least equipped to absorb them. Chapter 1 shows that liquidity constraints at the point of purchase suppress investment in worker-owned capital, the portable productive assets workers finance themselves but deploy on the job. A randomized controlled trial with a large Indian home-services platform firm finds that subsidizing the scooter down-payment increases adoption by 32 percentage points, raises worker weekly earnings by 13 percent, and generates a 153 percent firm return within the first year. A discrete-choice experiment confirms that adoption is 2.5 times more sensitive to the down-payment than to total price, isolating cash at hand as the binding constraint. Chapter 2 identifies the same purchase-time liquidity mechanism in India's LPG clean cooking fuel subsidy program, where a cash-back design inadvertently depresses adoption among the poor and worsens child health outcomes. Chapter 3 shows that extreme rainfall imposes sharply unequal income losses on urban gig workers, with peripheral, low-infrastructure neighborhoods bearing losses 3.5 times larger than central ones, losses that platform flexibility cannot offset. Together, the three essays show that modest reforms to the timing and delivery of costs and transfers can substantially improve welfare for the poor, and that the expansion of flexible labor market arrangements, while increasing workers' ability to smooth income, does not eliminate the structural inequalities that leave low-income workers disproportionately exposed to economic and environmental shocks.

Contents

Acknowledgments	ix
General Introduction	1
1 Worker-Owned Capital and Productivity: Experimental Evidence from a Large Platform Firm in India	4
1.1 Introduction	5
1.2 Study Setting	13
1.2.1 The Firm	13
1.2.2 Mobility as a Production Input for Jobs	14
1.3 Field Experiment: Subsidy for Worker Owned Capital	15
1.3.1 Intervention Details	15
1.3.2 Data Sources	16
1.3.3 Summary Statistics and Baseline Balance	18
1.3.4 Empirical Strategy	19
1.4 Results: Impact of Subsidy for Worker Owned Capital	19
1.4.1 Take-up of the Subsidy	19
1.4.2 Effects on Worker Outcomes	20
1.4.3 Effects on Firm Outcomes	23
1.5 Return to Investment	24
1.5.1 Returns to Subsidy for the Firm	24
1.5.2 Returns to Scooter Purchase for the Worker	25
1.6 Mechanism Experiment: Liquidity Constraint vs. Price Response	27
1.6.1 Model of Investment Choice by Workers	27
1.6.2 Testing Model Predictions: Discrete Choice Experiment	30
1.6.3 Results: Discrete Choice Experiment	31
1.7 Conclusion	32
Appendix 1.A Appendix A: Additional Results	46
Appendix 1.B Appendix B: Return to the Firm and Worker: Details	66
1.B.1 Return on Investment (ROI) for the Firm	66
1.B.2 Worker ROI: New Adopters	68
1.B.3 Worker Internal Rate of Return (IRR): New Adopters	69
Appendix 1.C Appendix C: Investment Modeled as Discrete Choice	72
1.C.1 Alternate Modeling Choice: Investment Decision as Discrete	72

2	Timing the Transfer: Liquidity Constraints and the Transition to Clean Fuels	76
2.1	Introduction	77
2.2	Institutional background	82
2.3	Theory	83
2.4	Data	85
2.4.1	Administrative sales data	85
2.4.2	Household-level data	86
2.4.3	Other datasets	87
2.4.4	Summary statistics	87
2.5	Empirical strategy	89
2.6	Results	90
2.6.1	Impact of LPG refill price and subsidy on refill purchase	90
2.6.2	Impact of LPG refill price on solid fuel expenditure	92
2.6.3	Robustness	93
2.7	Impact on health	93
2.8	Mechanism - Liquidity constraints	95
2.8.1	Alternative measures and sources of liquidity	95
2.8.2	Alternative explanations	97
2.9	Conclusion	99
	Appendix 2.A Online Appendix: Additional Results	113
	Appendix 2.B Supplementary Results: Reproduction of analyses in the main paper with LPG Refills recoded as missing until consumer obtains an LPG account during our study period (Nov 2017 - Dec 2019)	138
3	Climate Shocks and the Limits of Labor Market Adaptation	143
3.1	Introduction	144
3.2	Context	146
3.2.1	Comparisons to the broader urban Indian labor market	147
3.2.2	Delhi-National Capital Region (Delhi-NCR)	148
3.3	Data	148
3.3.1	Worker Data	148
3.3.2	Platform Demand Data	149
3.3.3	Rainfall Data	149
3.3.4	Infrastructure Data	150
3.3.5	Geographic and Topographic Data	150
3.3.6	Wealth Index Data	150
3.3.7	Weather Forecast Data	150
3.4	Descriptive Statistics	151
3.4.1	Labor Market	151
3.4.2	Rainfall Shocks	151
3.4.3	Economic Geography of the City: Distance, Infrastructure, and Inequality	152
3.5	Empirical Strategy	153
3.6	Results	154

3.6.1	Heterogeneity by Time Sensitivity of Occupation Category	155
3.6.2	Intertemporal Substitution	156
3.7	Spatial Heterogeneity in the Impacts of Rainfall Shocks	158
3.7.1	Grid Level Impacts	158
3.7.2	Role of Geography: Distance from City Center	159
3.7.3	Do Rainfall Shocks Exacerbate Economic Inequality Within Cities? .	160
3.7.4	What kind of infrastructure might mitigate income losses?	160
3.7.5	Policy Implications	161
3.8	Conclusion	161
Appendix 3.A	Appendix	180

List of Figures

1.1	Gender Differences in Actual and Preferred Mode of Commute (Pre-Experiment Survey)	34
1.2	Correlation Between Baseline Mode of Commute, Productivity and Travel Expenditure	35
1.3	Correlation Between Job Value and Job Distance and Timing	36
1.4	Correlation Between Job Completion Rate and Job Start Time	37
1.5	Impact on Job Outcomes Over Time for New Adopters	38
1.6	Discrete Choice Incentive Compatible Experiment: Contract Take-Up Rates	39
A1	Example of Intervention Scooters	46
A2	Correlation Between Wealth and Scooter Use and Productivity at Baseline	47
A3	Descriptive Facts: Daily Calendar Availability and Job Completion Rate for Workers	48
A4	Worker's Credit Score (CIBIL) Distribution	49
A5	Take-up Rate Over Time	50
A6	Impact on Job Outcomes Over Time for Incumbent-Users	51
A7	Impact on Jobs of Control Workers Over Time	52
A8	Baseline Beliefs About Jobs and Revenue Increase Conditional on Scooter Adoption	53
A9	Heterogeneity in Treatment Effects by Beliefs	54
A10	Heterogeneity in Treatment Effects by Wealth	55
2.1	Over-the-counter Price, Effective Price and Subsidy on LPG (November 2017 - December 2019)	101
2.2	Correlation between International LPG Price and Over-the-counter Domestic LPG Price (November 2017 - December 2019)	102
2.3	Households with wealth higher than $\widehat{W}(q)$ purchase fuel, while those with lower wealth do not. The wealth cutoff would be the dash-dotted line absent any liquidity constraints.	103
A1	Delay in LPG Subsidy Transfer	113
A2	Average Cash on Hand (Rural Households)	114
A3	Correlation of Over-the-counter LPG Price Across Districts (November 2017 - December 2019)	115
A4	Market-level Over-the-counter and Effective LPG Prices (November 2017 - December 2019)	116
3.1	Distribution of Providers Across Occupation	163

3.2	Proportion of Grids Shocks	164
3.3	Spatial distribution of workers and customers by tercile	165
3.4	Intertemporal Substitution of Earnings and Labor Supply	166
3.5	Warnings or Forecasts Before Heavy Rain	167
3.6	Forecasted vs Unforecasted shocks	168
3.7	Spatial heterogeneity in shock effects across grids	169
3.8	Daily Income Loss Due to Rainfall Shocks, by Distance from Center	170
3.9	Dynamic Adaptation, by Distance from Center	171
3.10	Correlation between Relative Wealth and Income Loss in the Grid	172
A1	Grid Level Estimates with P-Value less than 0.1	180

List of Tables

1.1	Impact on Take-Up of Subsidy and Scooter Purchase	40
1.2	Impact on Jobs, Earnings and Productivity	41
1.3	Impact on Labor Supply Reallocation Across Space, Time, and Job Quality on New Adopters	42
1.4	Impact on Firm Revenue, Profits and Worker Attrition	43
1.5	Return on Investment (ROI) for the Firm and the Workers	44
1.6	Discrete Choice Incentive Compatible Mechanism Experiment	45
A1	Baseline Balance by CIBIL (Credit) Score Among (All) Applicants	56
A2	Baseline Balance	57
A3	Take Up Characteristics	58
A4	Robustness: Impact on Jobs and Earnings of New Adopters	59
A5	Robustness: Impact on Jobs and Earnings of Incumbent Users	60
A6	Robustness: Impact on Jobs and Earnings of Pooled Sample	61
A7	Impact on Travel Expenditure, Commute Time or Workers, and Spillover on Family Members' Earnings and Jobs	62
A8	Impact on Distance Traveled and Job Timing: Incumbent Users	63
A9	Treatment Effects on Job Quality Indicators: Incumbent Users	64
A10	Contract Menus and Payment Burden (Amounts in thousands of INR)	65
A11	ROI Sensitivity to the Discount Rate (r) – 12 months	67
2.1	LPG Consumer Characteristics (Administrative Data)	104
2.2	Household Characteristics (Survey Data)	105
2.3	Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills (Administrative Data)	106
2.4	Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills (Survey Data)	107
2.5	Impact of Over-the-counter Price and Cash-back Subsidy on Solid Fuel Expenditure (Survey Data)	108
2.6	Impact of Over-the-counter Price and Cash-back Subsidy on Birth and Post- birth Child Health Outcomes	109
2.7	Impact of Over-the-counter Price on Monthly Refills: Heterogeneity by Household Attributes (Survey Data)	110
2.8	Impact of Over-the-counter Price and Cash on Hand (COH) on Monthly Refills (Survey Data)	111
2.9	Effect of High Rainfall in Sowing Season on Harvest Season Refills	112
A1	Average Cash-back Subsidy, Over-the-counter Price, and Effective Price	117

A2	Characteristics of Non-PMUY and PMUY Households (Survey Data)	118
A3	Cash on Hand (COH)	119
A4	Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills (Administrative Data)	120
A5	Effect of LPG Subsidy Delay on Refill Demand, by PMUY status (Administrative Data)	121
A6	IV Estimation: Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills (Administrative Data)	122
A7	IV Estimation: Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills (Survey Data)	123
A8	Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills (Administrative Data with Bootstrapped Standard Errors)	124
A9	Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills (Survey Data with Bootstrapped Standard Errors)	125
A10	Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills (Administrative Data) - Poisson Pseudo-Maximum Likelihood Estimation . .	126
A11	Birth and Post-Birth Child Health Outcomes	127
A12	Impact of Over-the-counter Price and Cash-back Subsidy on Birth and Post-birth Child Health Outcomes (with Sampling Weights)	128
A13	Effect of Over-the-counter Price and Cash-back Subsidy on Post-birth Child Health Outcomes	129
A14	Effect of Over-the-counter Price and Cash-back Subsidy on Child Anthropometric Outcomes	130
A15	Impact of Over-the-counter Price on Monthly Refills: Heterogeneity by Asset Ownership (Survey Data)	131
A16	Correlation of Consumer Price Index with Nominal Over-the-counter Price of LPG	132
A17	Impact of Over-the-counter Price on Monthly Refills: Heterogeneity by Knowledge About Subsidy (Survey Data)	133
A18	Impact of Over-the-counter Price on Monthly Refills: Heterogeneity by Trust in ASHA Worker (Survey Data)	134
A19	Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills (Administrative Data) (Sample Restricted to experienced LPG Users)	135
A20	Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills (Administrative Data) including Unemployment as a Control	136
A21	Impact of Over-the-counter Price and Cash-back Subsidy on Solid Fuel Expenditure (Survey Data) (Adding Year FE)	137
B1	LPG Consumer Characteristics (Administrative Data)	138
B2	Household Characteristics (Survey Data)	139
B3	Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills (Administrative Data)	140
B4	Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills (Survey Data)	141
B5	Impact of Over-the-counter Price on Monthly Refills: Heterogeneity by Household Attributes (Survey data)	142

3.1	Labor Market Summary Statistics	173
3.2	Statistics for the Median Worker, by Distance from Center Terciles	174
3.3	Wealth and Infrastructure of Neighborhoods, by Distance from Center Terciles	175
3.4	Rainfall Shock Impact	176
3.5	Labor Demand vs Supply	177
3.6	Heterogeneity: Urgent vs Non-Urgent	178
3.7	Infrastructure and Rain Impacts	179

Acknowledgments

This dissertation would not exist without the generosity, patience, and support of many people, and I am glad to have the opportunity to acknowledge them here.

My deepest gratitude goes to my committee. Jeremy Magruder has been everything one could hope for in an advisor: intellectually generous, unfailingly honest, and consistently present at every stage of this work. From the early days of encouraging me to launch the randomized controlled trial with the platform firm, always believing in the importance of the question it was answering, to brainstorming modeling ideas in his office, to attending my job market paper seminars and meeting with me weekly to work through the arguments and results, his investment in my development as a researcher has been extraordinary. His ability to see the strongest version of an argument, and to push me toward it even when I could not see it myself, has shaped not only this dissertation but the way I think about research.

Supreet Kaur and Aprajit Mahajan brought me onto a collaborative project on opioid experimentation among youth in Punjab, a question of real urgency in a region where drug dependence among the rural poor is severe, and one that eventually scaled into a government-wide program. Working closely with both of them on a project of that consequence raised my standards for what careful, policy-relevant research can achieve, and their generosity in engaging with my other work throughout the PhD improved my thinking in ways I continue to draw on. They remained consistently generous with their time across all of my papers and projects, engaging carefully with the details, critiquing arguments, and pushing me to sharpen my thinking in ways that improved the work substantially. Gautam Rao's encouragement came at exactly the moments I needed it most. His support, and his push to trust my core ideas and stay with the hard questions, gave me a confidence in my own judgment that I carried forward through the rest of the PhD. I am grateful to all four for the seriousness with which they took my work and for the confidence they showed in me.

I also owe a great deal to my co-authors. Prabhat Barnwal has been not just a co-author but a trusted colleague and advisor, consistently generous with his time and thinking across every stage of my PhD. Farzana Afridi and Christian Ahlin brought rigor and depth to Chapter 2, and I have learned enormously from working with both of them. Aakash Bhalotia has been a wonderful collaborator on Chapter 3, and working with him has been one of the genuine pleasures of the last several years.

This research would not have been possible without the partner firms and institutions that opened their data and operations to outside scrutiny. I am grateful to the platform firm at the center of Chapter 1 for their willingness to run an experiment that was genuinely uncertain, and for scaling the program afterward, which remains one of the most satisfying

outcomes of this project. I am also grateful to Mukesh Gulati and the Foundation for MSME Clusters for their generous support and partnership in the design and execution of a firm-level randomized controlled trial on barriers to technology adoption in India's garment manufacturing sector. Their field knowledge and logistical generosity made that work possible in ways that institutional support alone could not have. I gratefully acknowledge financial support from the IZA Gender, Growth and Labor Markets in Low-Income Countries Programme, the Center for Effective Global Action, the Weiss Fund, and IWWAGE-IFMR. I also received the Berkeley Fellowship for Graduate Study, for which I am deeply thankful.

The Berkeley ARE community has been an exceptional intellectual home. I am grateful to the faculty, staff, and fellow students who attended seminars, asked hard questions, and made Giannini Hall a place worth coming to every day. To my fellow graduate students: the friendships formed here have been one of the unexpected gifts of this process. You made the hard stretches bearable and the good stretches genuinely fun, and I will carry those years with me long after I have forgotten the details of any particular regression.

My parents gave me the foundation without which none of this would have been possible. My father, Dr. Tamal Sarkar, is himself a development economist with a lifelong passion for research and policy on industrial cluster development in India. Growing up watching him think carefully about firms, workers, and the constraints they face planted seeds that I did not fully recognize until I found myself asking similar questions in my own work. His generosity extended well beyond intellectual inspiration: it was through his connections and encouragement that I was introduced to the Foundation for MSME Clusters, which opened the door to a firm-level randomized controlled trial on barriers to technology adoption, a project whose data collection is now complete and whose analysis and writing are underway. I could not have designed or executed that work without him. My mother, Amita Sarkar, raised me to believe that intellectual ambition was worth pursuing, and that hard work was its own form of integrity. On the days when the work felt impossible, she would remind me of something simple: that the word *research* means, at its root, to search again. To fail, to lose the thread, and then to go looking once more. That, she would say, is not a sign that something has gone wrong. That is the job. I returned to that reminder more times than I can count.

Ronak, there is no adequate way to thank you in a paragraph, but I will try. You moved cities, adjusted plans, and reorganized your life around my PhD. You sat with me through the nights when nothing was working, celebrated the small wins with the same enthusiasm as the large ones, and never wavered across all the years this took. The PhD was harder than I expected, and the hardest parts of it were made manageable entirely because of you. What has come after has already exceeded everything I let myself hope for, and the best of it is the life we have built together.

To my in-laws, Rakesh Jain, Mamta Jain, and Rohan Jain: you welcomed me into your family with a generosity and openness that I am deeply grateful for. You supported this path with grace, asked after the work with genuine curiosity, and gave Ronak and me the space and encouragement to build the life we wanted. I am in awe of your kindness and grateful for the warmth with which you have always treated this PhD as a shared endeavor.

To my best friends Tanvi Jaipurkar and Lavangi Naithani: you have been constants across every phase of this journey, showing up with the kind of warmth and steadiness that makes even the most difficult periods feel less isolating. Through every stretch of this PhD, you

were a source of joy that had nothing to do with how the work was going, and a reminder that peace and love were always in abundance. That reminder gave me the strength to keep going more times than you know. I am deeply lucky to have you both.

My aunt Arti Sarkar and my cousins Ipshita Sarkar and Uddalak Sengupta have been a source of warmth and encouragement that I have drawn on throughout these years. Their belief in me has been quiet, consistent, and entirely unconditional, and there is a particular kind of strength that comes from knowing that certain people are simply always in your corner. I am fortunate to have them in mine.

General Introduction

A central question in development economics is why workers and households in low-income settings persistently underinvest in welfare-improving decisions, despite potentially high returns. This dissertation argues that an important part of the answer lies in frictions that distort the timing of costs and benefits. The first two essays study purchase-time liquidity constraints: situations in which the costs of a welfare-improving decision must be borne upfront while returns accrue gradually, deterring investment even when net returns are substantial. The third essay examines whether the flexibility of platform-mediated labor markets insulates workers from climate shocks, finding that income losses from extreme rainfall are large, persistent, and fall disproportionately on workers in poorer neighborhoods with weaker infrastructure. The three essays examine these mechanisms using high-frequency administrative data from platform firms and government records, large-scale field experiments, and quasi-experimental research designs, all in the context of India's expanding gig-economy platform sector and the government's flagship clean cooking fuel program. The first two chapters share a common mechanism, namely purchase-time liquidity constraints that suppress welfare-improving investment, and study it in two institutionally distinct settings. The third chapter takes a different approach, examining how environmental shocks interact with spatial inequality in urban infrastructure to produce regressive income losses among gig workers.

Chapter 1 introduces the concept of *worker-owned capital*: portable productive assets that workers finance and own themselves but deploy on the job, such as corrective eyewear for a tailor, non-slip footwear for a healthcare worker, or the scooter that enables a home-services worker to reach customers efficiently across a sprawling city. Unlike firm-owned capital, worker-owned capital retains its value if the worker changes employers, which means neither the firm nor the worker captures the full social return. If workers face liquidity constraints and firms can appropriate only part of the productivity gain, equilibrium investment falls below the joint-surplus benchmark, suppressing both earnings and firm profits. To identify this mechanism and measure its consequences, the chapter reports the results of a randomized controlled trial conducted with a large home-services platform firm in India and its female workforce. The experiment subsidizes the down payment on a standard consumer scooter loan, relaxing the upfront financing constraint without altering the total asset cost. The down-payment subsidy increases scooter adoption by 32 percentage points, a 188 percent increase over the control group, indicating that demand was constrained by purchase-time liquidity rather than long-run affordability. New adopters increase weekly jobs completed by 9 percent and weekly earnings by 13 percent, a 10 percent gain in labor productivity achieved with no increase in hours worked. Treated workers expand the extensive margin of their labor

supply, accepting more geographically distant jobs and more late-evening assignments that would otherwise go unmatched. The firm recovers the full subsidy cost within 18 weeks and earns a 153 percent return within the first year, while workers recover the full asset cost within thirteen months even without the subsidy. An incentive-compatible discrete-choice experiment confirms the mechanism: adoption is approximately 2.5 times more sensitive to a reduction in the required down payment than to an equivalent reduction in total price, ruling out long-run affordability and isolating the cash-at-hand constraint.

Chapter 2 identifies the same purchase-time liquidity mechanism in a different institutional context: India’s Pradhan Mantri Ujjwala Yojana (PMUY) LPG subsidy program. Under this program, households pay the full over-the-counter market price at purchase and reclaim a cash-back subsidy afterward. Because the net-of-subsidy price remains nearly constant, standard price theory predicts that fluctuations in the over-the-counter price should not affect purchasing behavior. Using rich administrative data on LPG refill purchases and quasi-experimental variation in the over-the-counter price, the chapter finds that a 1 percent increase in the upfront price reduces purchases by approximately 1.5 percent of the mean for low-asset households, with no corresponding effect for wealthier ones. This sensitivity to cash at hand, rather than to the effective price, is consistent with a binding liquidity constraint. Constrained households substitute toward solid fuels such as firewood and dung, and household survey data document deterioration in child health: reduced birth weight among neonates and increased incidence of coughing and rapid breathing among children under five, both scaling with the upfront price increase. A program designed to promote cleaner indoor air thus inadvertently undermines that goal through its payment architecture. The chapter shows that redesigning the timing of the transfer, by minimizing reimbursement delays, enabling purchase-point credit, or issuing electronic vouchers that apply the subsidy at point of sale, can substantially reverse these effects, and cautions that delayed cash transfers impose unintended regressive costs on the poorest beneficiaries.

Chapter 3 examines a distinct but complementary question: whether the flexibility of platform-mediated labor markets enables workers to adapt to environmental shocks, and how the consequences are distributed across the urban income distribution. Using a dataset of more than four million daily job records from a home-services platform in Delhi, linked to worker residential locations, gridded rainfall data, and neighborhood infrastructure measures, the chapter estimates the income consequences of extreme rain events. Earnings fall by approximately 4 percent on days when rainfall exceeds 10mm, driven by a reduction in labor supply rather than by falling demand. Workers engage in intertemporal substitution, partially smoothing income by shifting effort to adjacent days, but show no spatial reallocation toward less-affected areas, even when substantial cross-neighborhood variation in shock severity exists on a given day. Physical barriers to mobility, including waterlogged roads, absent metro connections, and poor drainage, constrain workers’ responses. The distributional consequences are unequal: workers in the most peripheral tercile of neighborhoods experience income losses approximately 3.5 times larger than those in the innermost tercile, a gradient mediated by drainage quality and metro access. These results suggest that rapid urban expansion without commensurate infrastructure investment concentrates climate costs among low-income workers in the urban periphery, and that platform flexibility provides limited insulation against physical

mobility constraints.

A common thread across all three essays is the use of high-frequency administrative data from platform firms, combined with credible identification strategies that exploit exogenous variation in treatment status or policy-relevant prices. Chapter 1 uses a pre-registered randomized controlled trial and an incentive-compatible discrete-choice experiment to provide experimental estimates of adoption, productivity, and firm returns. Chapter 2 exploits quasi-random variation in India’s over-the-counter LPG price, set at the national level and tracking international oil markets, to identify causal effects of purchase-time costs on household behavior and health outcomes. Chapter 3 uses within-worker, within-grid variation in daily rainfall exposure to identify the causal income consequences of extreme weather, with rich spatial heterogeneity analysis enabled by the granularity of the platform data. This combination of large-scale administrative data with experimental and quasi-experimental identification allows the essays to trace the mechanisms through which frictions operate at a level of precision rarely achievable with conventional data sources.

Chapter 1

Worker-Owned Capital and Productivity: Experimental Evidence from a Large Platform Firm in India

Abstract. Worker productivity often depends on privately owned, portable inputs that workers finance themselves. When liquidity constraints bind, investment in these inputs may be inefficiently low, depressing both worker earnings and firm profits. I study whether firm subsidies that enable acquisition of worker-owned capital can raise productivity while remaining privately profitable for the firm. I partner with a large home-services platform firm in India where workers travel to customer homes to complete service jobs. In this setting, efficient last-mile connectivity is a key input for worker productivity, making a scooter essential worker-owned capital. I implement a randomized experiment in which workers receive subsidies toward scooter loan down payments. The subsidy increases scooter adoption by 32 percentage points. Administrative data show large increases in jobs completed and earnings. Treated workers travel farther, accept more lucrative early-morning and late-evening jobs, and deliver quality improvements. The firm recovers the subsidy within 18–29 weeks through increased revenue share and reduced attrition, subsequently scaling the program to additional cities. A complementary discrete-choice experiment shows demand for this worker-owned capital is more responsive to reducing upfront payment requirements than to lowering overall prices, consistent with liquidity constraints as the primary barrier. The results demonstrate that enabling investment in worker-owned capital can be both productivity-enhancing and privately optimal for firms while benefiting workers.

Keywords: *productive assets, physical capital, investment, labor market outcomes, firm productivity, female workers, liquidity constraints*

1.1 Introduction

Worker productivity often hinges on worker-financed, portable capital – assets employees purchase and use on the job but that firms typically do not provide. Think of corrective lenses that let a tailor stitch without rework; cushioned, non-slip footwear that allows a security guard or a healthcare worker to stand longer and miss fewer shifts; basic tools owned by construction workers; or a vehicle, smartphone, or reliable internet connection that expands a service or delivery worker’s feasible set of tasks and coordination. While *capital* is typically analyzed in the context of entrepreneurs or firms, these individually-owned, job-complementary assets might also shape worker productivity (De Mel et al., 2008; Bandiera et al., 2017). Such worker-owned capital has three salient features: (i) it raises output in the worker’s current tasks (faster completion, fewer errors, less idle time); (ii) it is general, in the sense that the same asset retains value across employers and matches; and (iii) because the asset is owned by the worker, it can deliver private gains — higher earnings, safer conditions, more stable schedules — even when the employer does not directly invest.

Investment in worker-owned capital may remain inefficiently low for two reasons.¹ The first reason is surplus-sharing and appropriability. When workers bear the up-front cost but the marginal product is shared with the firm, they rationally underinvest; this is the classic rent-sharing (hold-up) problem observed in many settings (e.g., sharecropping and tenancy), where the investor captures only a fraction of the marginal return (Laffont and Matoussi, 1995; Shaban, 1987; Jacoby and Mansuri, 2008). Symmetrically, firms have weak incentives to finance such general (portable) assets because returns are only partly appropriable as the same asset raises the worker’s productivity with other employers, so the firm cannot capture the full marginal product (Becker, 1964). The second reason is liquidity constraints. Even when private returns are high, lumpy up-front payments can block adoption, especially in developing-country settings. In equilibrium, these two forces depress privately optimal adoption below the joint-surplus benchmark.

If workers are liquidity constrained and the firm can appropriate only part of the resulting productivity gains, an employer subsidy for such capital could raise worker welfare and firm earnings simultaneously. Motivated by this, and in collaboration with a large Indian platform and its female workforce, I examine three questions: (i) to what extent does a firm-provided subsidy induce adoption of worker-owned capital; (ii) how does induced adoption affect labor supply, on-the-job productivity, and the firm’s profits; and (iii) which frictions does the subsidy relax — purchase-time liquidity versus surplus-sharing concerns.

To answer these questions, I study an employer program that subsidizes the down payment on a standard consumer loan used to purchase a scooter on credit, thereby relaxing the worker’s upfront financing constraint (while simultaneously reducing the scooter price). The platform matches workers to in-home service jobs, for example at-home salon and spa services including haircuts, facials, and manicures. As these jobs are completed at the customer’s homes, reliable last-mile mobility is a first-order production input. A scooter is

¹By “inefficiently low” I mean underinvestment relative to a joint-surplus benchmark: a planner maximizing the sum of firm profits and worker utility, net of the full economic cost of the asset and financing, taking wages and output prices as given and abstracting from externalities and general-equilibrium effects. In the absence of (i) the firm’s appropriation problem for general capital and (ii) workers’ purchase-time liquidity/present-bias frictions, adoption would be higher.

a worker-owned capital good that could directly expand the feasible set of jobs and reduce per-job costs. In the study cities — typical of many low and middle income countries’ urban environments — public transit does not reliably provide last-mile connectivity.² As a result, workers commute either by personally owned vehicles (typically scooters) or by private taxis (e.g. ride-hailing). In a baseline survey of 1,531 workers, only about 31% of women reported commuting by their own scooter, yet 79% stated a preference to do so, suggesting a usage-to-preference gap of about 50 percentage points.³ Respondents also rated scooter commuting as faster, more reliable, cheaper, and safer than taxis, consistent with higher on-the-job effectiveness in this setting. Given female workers’ low current scooter use despite strong stated demand, and the centrality of last-mile mobility for platform work, the firm’s female workforce is the natural population in which to test whether easing access to worker-financed, portable capital (here, a scooter) raises productivity and firm profits.

The experiment enrolls 350 full time women workers in Delhi, Bengaluru, and Pune. Participants applied through an in app invitation and met a pre specified credit access screen.⁴ Control workers were eligible for standard two wheeler credit to purchase a USD 1,200 scooter with a 15 percent down payment of USD 180. Treated workers received a USD 120 subsidy applied at the time of down payment. This policy delivers a 10% price rebate and, crucially, covers two thirds of the down payment, reducing the out-of-pocket payment at purchase from USD 180 to USD 60. I observe an eight-month panel at daily job level for each worker, yielding high-frequency measures of jobs, hours, and earnings linked to verified scooter purchase and subsequent use. This worker–platform panel - combining granular, job-level outcomes with a direct input–output mapping — is unusual in developing-country settings and rare for female workers in particular. Because treatment status is randomized and scooter adoption is directly observed, I can cleanly link worker use of own capital to productivity and earnings at the worker level and to profits at the firm level, isolating the adoption margin rather than confounding it with selection or access to credit.

Ex ante, adoption of a credit-financed scooter is uncertain given the asset’s size relative to income: a USD 1,200 scooter is about 20 weeks of earnings, and the loan entails fixed monthly installments over about 26 months. In the experiment, 49% of treated workers purchased a scooter upon offer of the subsidy. Because the subsidy is disbursed only at purchase and paid directly to the lender at the down payment, this purchase implies take-up of credit rather than cash financing. The randomized offer generates a 32 percentage-point increase in purchase relative to control (control mean 17%), a 188% relative rise, providing a strong first stage on the adoption margin. Adoption among workers in the control group is concentrated

²For instance see: *India—Making Public Transport and Urban Spaces Safer and Inclusive for Women: Policy Brief* (2022); *World Bank* (2023); *Urban Transport: Policy Note Series* (2021); *Shah et al.* (2017); *Gender Equality in Transport: Toolkit* (2021).

³Among men in the same survey, approximately 77% reported commuting by own scooter and 92% stated a preference to do so, implying a usage-to-preference gap of about 10 percentage points, much smaller than for women (see Figure 1.1).

⁴From August to October 2024 the firm sent in app invitations to 3,000 randomly selected female salon workers across Delhi, Pune, and Bengaluru. About 1,200 workers responded which is a 40 percent rate with 50 percent from Delhi, 25 percent from Bengaluru, and 25 percent from Pune. The research team completed light verification of government ID and address to screen out low commitment inquiries. Enumerators then administered a baseline survey and collected documentary proof of each applicant’s credit score. Eligibility required a baseline CIBIL score of at least 700. Operationally 74 percent of firm workers met this threshold.

4 to 6 months after the launch of the experiment, consistent with gradual self-financed accumulation. In contrast, treatment induces earlier and substantially higher take-up, i.e., both an extensive-margin increase and an acceleration in the timing of adoption.⁵

Subsidy treatment assignment was randomized within strata defined by baseline commute mode. Nearly 70% of the workers in the sample were not using a scooter at baseline and constitute the main analysis sample (“*new adopters*,” defined by non-use of scooter at baseline); the remaining 30% were already using a scooter – but wanted to buy a new one – and serve as a placebo comparison (“*incumbent users*”).⁶ If the subsidy operates by inducing a travel mode shift — from taxis and ride-hailing to own-scooter commuting — no productivity effect is expected among incumbent users, for whom commuting mode is unchanged. Accordingly, all treatment effects are reported separately for new adopters and incumbents.

For new adopters, the intent-to-treat effect on the number of jobs is an increase of 0.54 jobs per week ($p < 0.05$) relative to a control mean of 6.12, which corresponds to a 9% increase weekly over an 8-month period. Weekly earnings increase by INR 602 ($p < 0.01$) over a control mean of INR 4,758, a 13% increase. Labor supply – the sum of hours marked available to work, waiting for work, time at the customer’s location completing work, and travel time – does not change, remaining near 53 hours per week.⁷ Since workers earn more with unchanged hours, weekly productivity (earnings per total hours) increases by INR 9.82 ($p < 0.01$) from a control mean of INR 94, a 10% gain.

A potential threat to interpretation is endogenous job assignment: if the platform began routing more (or more lucrative) jobs to treated workers in response to their higher completion rates, the observed gains could reflect reallocation by the matching algorithm rather than improved mobility. I find no evidence for such a mechanism. First, event-study estimates show an immediate jump in completed jobs at the time of treatment, and this higher level is sustained over the panel; there is no delayed jump. This timing is inconsistent with a platform-reassignment story, which would require a period of improved post-treatment performance before routing rules are updated. Instead, the pattern points to a direct, contemporaneous effect of the subsidy-induced mobility change on output. Second, treated and control workers receive statistically indistinguishable numbers of job offers over the study horizon. The offer mix is likewise balanced across treatment arms — by price, distance, and time of day — with no shift toward higher-priced or otherwise favorable jobs for the treated. Finally, institutional details from the firm indicate that the assignment algorithm does not use completion rates to preferentially route demand

⁵We restrict to women who had already expressed interest in buying a scooter. This is appropriate for our question—how contract *design* affects take-up among inframarginal (already credit-eligible, interested) workers—rather than whether expanding access brings in new borrowers. The approach aligns with evidence that capital interventions yield returns mainly for people who have previously sought credit (Beaman et al., 2023).

⁶“New adopters” or “Incumbent users” are defined by *self-use of a scooter for commuting* at baseline, not by household ownership. Some respondents reported a scooter in the household (e.g., owned by a spouse or relative) that they did not use themselves. Stratification therefore hinges on baseline *use* because the mechanism of interest is a modal shift in the worker’s own commute; household ownership without self-use does not change the worker’s commuting technology. Baseline status is recorded in the pre-treatment survey.

⁷The lack of change in hours is consistent with the absence of a gift-exchange or other behavioral mechanism that would operate through increased effort.

for equity reasons; persistent non-completion is handled via retraining and counseling rather than through advantaged or disadvantaged job allocation. Together, these facts argue against assignment-driven reallocation and support increased productivity operating through improved mobility.

For the incumbent users (the placebo stratum), coefficients are small and statistically indistinguishable from zero, as expected given that these women already commuted by scooter at baseline. In the pooled sample, including both new adopters and incumbent users, earnings and productivity rise (significant at the 10 percent level), with effects driven entirely by new adopters. Taken together, these results indicate that the program’s benefits stem from new (first-time) adopters switching to a more efficient mode of transportation, rather than from behavioral responses among workers who already owned scooters.

The richness of the firm’s job-level data allows me to further unpack the productivity gains. Each transaction links a worker identifier to timestamps, locations (and thus home-to-job distances), and job attributes, allowing me to estimate treatment effects on job selection and quality rather than only on aggregates. Among new adopters, treatment expands the spatial (and temporal) reach of completed jobs: weekly distance traveled rises by 6.6 km — an 11% increase — with the entire effect concentrated in 10–25 km trips. There is no detectable change for <10 km jobs, indicating that the margin of adjustment is additional medium-to-long trips rather than substitution away from nearby work.⁸ Further, I find a reallocation of work toward peak periods. Among new adopters, the number of jobs completed between 8–11 a.m. and after 6 p.m. increases by 0.37 per week (control mean 3.10). In the firm’s administrative data, these windows traditionally show female-supply shortfalls and the highest unfilled-job rates; shifting work into them therefore expands output precisely where capacity is most constrained.

An important question is whether the increase in jobs completed comes at the expense of service quality. To investigate this quantity-quality tradeoff, I focus on two outcomes: (i) customer ratings and (ii) late arrival at customer’s location. Average ratings among treated workers increase by 0.03 stars ($p < 0.05$) relative to a control mean of 4.83 on a 5-star scale. This closes the gap to a perfect rating by approximately 18%. Given that ratings are heavily skewed above 4, such movements are nontrivial. Moreover, the improvement pushes the average past the firm’s internal quality threshold of 4.85, which the platform considers a “good” rating, making this shift operationally meaningful. By contrast, I find no statistically significant change in late arrivals to jobs. This null result is consistent with post-investment selection into farther or peak-hour assignments, which raises travel-time variance even as underlying service quality improves. Follow-up survey evidence supports the quality channel. Treated workers report less back pain and lower stress, plausibly contributing to higher on-the-job service quality as reflected in ratings. Among incumbent users, I find no treatment-induced reallocation – no detectable changes in distance traveled, time-of-day patterns, or job quality – consistent with their placebo role and the fact that their commuting technology was unchanged at baseline.

The firm earns a commission on each completed service, so increases in the number and value of jobs mechanically raise firm revenue. Among new adopters, total revenue generated per worker rises by INR 1,180, a 12% increase ($p < 0.01$). Because this gain coincides with

⁸Distances are measured door-to-door and include both outbound and return segments.

higher worker earnings and more completed jobs, it suggests an expansion of realized output or joint surplus, not a pure transfer. A natural concern is that the rise in total revenue reflects business stealing rather than higher output. I rule this out in two ways. First, an event-study restricted to control workers shows no contemporaneous decline in their outcomes around the treatment date: completed jobs and earnings for control workers remains flat around exposure and then drift up modestly in later months, mirroring platform growth. Second, the treatment-induced gains are concentrated in far-distance and late-evening jobs — the very segments with the highest pre-treatment non-fill rates — so added supply is most plausibly converting previously unserved demand into completed transactions rather than displacing other workers. Taken together, these facts indicate that the revenue increase reflects higher aggregate output rather than reallocation across workers.

Turning from revenue to surplus, I estimate profit per worker-week. Among new adopters, profit rises by INR 265 per worker-week, a 13% increase relative to the control mean of INR 2,108. The sign and magnitude are consistent with the platform’s 20–30 percent commission on job value and the documented shift toward higher-value jobs. A natural concern is that higher productivity might accelerate exit if outside options are attractive (three rival platforms operate in these cities and offline salon jobs are available). I therefore examine retention directly. I find no evidence of treatment-induced attrition. In the control group, one-year attrition is about 10%; treatment lowers this by about 4 percentage points ($p=0.15$). The estimate is imprecise but economically meaningful and indicates that profit gains are not offset by worker exit. Among the incumbent users, the control group attrition rate is approximately 2% and the treatment effect is near zero, consistent with a retention ceiling and with the premise that mobility was not binding for this group. Taken together, the profit increase for new adopters is not undone by higher attrition despite substantial outside opportunities.

To evaluate the firm’s return on the subsidy, I focus on the treatment effect on the treated (TOT). Take-up of the subsidy was around 50% among those offered, and no workers in the control group were offered the subsidy; with this one-sided non-compliance, I estimate TOT by instrumenting take-up of the subsidy with the randomized offer of the subsidy. This estimand is the relevant one for the firm’s accounting, since the firm incurs the cost only when a worker accepts the subsidy. The IV estimate implies that profit per treated worker increases by about INR 558 per week, from which the firm recovers the subsidy in 18 weeks and the implied 12-month return on investment is about 153%.

The high measured return naturally raises the question of why the firm had not already implemented such subsidies. While the experiment cannot isolate the precise source of ex-ante underinvestment, discussions with the firm point to two first-order frictions. First, uncertainty about return on investment: prior to the experiment, the firm had weak priors about the magnitude and durability of productivity gains and about the share of those gains it would capture through its commission structure, rendering the investment’s net present value highly uncertain. Second, experimentation to resolve this uncertainty is costly, and the firm faced constrained administrative capacity and competing operational priorities that delayed both the experiment and subsequent scaling decisions. This pattern is consistent with growing evidence that firms may underexperiment with potentially profit-improving practices (Bloom et al., 2013; Hanna et al., 2014). Consistent with the experimental results demonstrating program profitability, the firm scaled the subsidy program to three additional

cities, extending the offer to over 3,000 additional female workers. This ex-post expansion provides revealed-preference evidence that, once belief frictions were attenuated, the subsidy is viewed internally as a profitable lever for growth.⁹

To quantify returns to workers, I instrument scooter purchase with randomized subsidy assignment. Treatment assignment increased scooter purchase by 32 percentage points. With two-sided noncompliance, the IV estimates identify a local average treatment effect (LATE) for compliers. The estimated (IV) effect of purchasing a scooter is an increase of INR 2,018 in weekly earnings. The discounted payback period of the scooter for the new adopters is thus 54 weeks.¹⁰ Thirteen months after rollout, 90% of treated workers remain active, with attrition below that of the control group. Extrapolating the observed effects and survival rate for an additional 12 months implies a 24-month ROI of 73%.¹¹

The magnitude of these returns sharpens the central puzzle – why did workers not invest pre-subsidy? I address this next by examining the frictions the subsidy relaxes for the workers. I explore two reasons for underinvestment by workers: (i) incomplete internalization of surplus (workers bear the full cost but capture only a share of output) and (ii) a binding liquidity constraint at adoption (difficulty assembling the down payment). Guided by a simple model of the worker’s investment decision under these frictions, I implement an incentive-compatible discrete-choice exercise with a new set of workers. Each worker evaluates a within-person, randomized menu of four contracts: (1) the regular loan — 15% down payment with the remaining 85% financed through loan (similar to the field experiment “control” contract); (2) a proportional price rebate that reduces the total purchase price and scales both the down payment and financed principal proportionally — primarily lowering lifetime cost while doing little to relax upfront cash; (3) a low down-payment contract that cuts cash due today by shifting it into financed principal (raising lifetime cost); and (4) the subsidy contract — a grant that directly reduces the upfront payment (akin to the field experiment “treatment” contract).¹²

Overall demand varies sharply by contract features: 16% under the regular loan, 31% with a pure price rebate, 66% with a low down-payment contract, and 95% with a subsidy. This already indicates that workers are much more responsive to low down-payment contracts than to a price rebate despite their high overall expense. The mechanism is clearest in the heterogeneity of contract demand by liquidity. Among liquidity-constrained workers (those who report they are unable to assemble INR 15,000 at purchase), demand is largely unresponsive to the regular loan and to a price rebate, but rises to 70% with a low down payment and to 98% with a down-payment subsidy. By contrast, among workers not liquidity-constrained, all alternative contracts raise adoption relative to the benchmark. The two price-reducing offers (rebate and subsidy) perform similarly and yield higher take-up than the low-down-payment option. Taken together, I find that liquidity at the time of purchase

⁹Ongoing work analyzes the performance of the scaled program and whether the high returns generalize outside the original experimental sites.

¹⁰Payback is computed using incremental LATE weekly earnings (INR 2,018) against the present value of the financed purchase cost, discounted. We do not net out the weekly EMI in the earnings flow—doing so would double count financing costs. Details in Section 1.5 and Appendix 1.B.

¹¹Even without taking the subsidy into account, the implied 24-month ROI for the worker is 57%.

¹²One stated “yes” is implemented by lottery, ensuring truthful revelation and yielding policy-relevant elasticities.

(down payment) is the pivotal margin: relief targeted at the upfront cash requirement generates three-to-four times larger adoption responses than an equal-value reduction in lifetime price. Further, any incomplete internalization of the firm’s surplus (a pure price effect) is, at most, second-order relative to the binding liquidity constraint at adoption.

Contribution. My first contribution is to provide experimental evidence on returns to worker-owned capital within a firm, and to quantify the surplus it creates on both sides of the employment relationship. A large theoretical literature predicts underinvestment by firms in workers’ general (portable) human capital because employers cannot fully appropriate the resulting productivity gains (Becker, 1964).¹³ When labor markets are perfectly competitive, even high-return general training is privately unattractive to firms. By contrast, with empirically relevant frictions — informational asymmetries, slow employer learning, search and mobility costs — wages can lie below marginal products, creating rents that make firm investment in even general capital privately profitable (Acemoglu and Pischke, 1999; Autor, 2001; Katz and Ziderman, 1990).

I extend this framework from general *human* capital to general *physical* capital owned by workers. This margin is understudied yet first-order: ownership of simple, portable assets can differentially raise on-the-job productivity among co-workers who have access versus those who do not¹⁴. In low- and middle-income settings, such assets are often lumpy and expensive relative to cash flow, making adoption sensitive to purchase-time liquidity. This paper shows that a firm-side contract that lowers the upfront cost of acquisition enables workers to make this investment; enabling worker-owned capital raises measured productivity and generates positive private returns to the firm. The results demonstrate that even in a modern platform-firm environment with comparatively low switching costs, labor-market frictions are sufficient for employers to profitably induce investment in worker-owned, general capital and experience no treatment-induced attrition.

My second contribution lies in studying liquidity for workers. This contribution connects two strands of literature. A long line of field experiments studies capital placed with independent producers (entrepreneurs or households), finding high returns to in-kind assets or cash and more mixed effects for standard microcredit (De Mel et al., 2008; Fafchamps et al., 2014; Banerjee et al., 2015b; Meager, 2019; Crépon et al., 2024; Bandiera et al., 2017). In contrast, I bring the capital-provision question *inside a production organization*, relevant directly for workers. There is emerging work showing consequence of liquidity for job choice: a modest, well-timed cash infusion moves Teach For America applicants into teaching (Coffman et al., 2019).

In my experiment, by focusing on workers who already qualify for standard loans yet are unable to access credit, I shift the focus from expanding credit access for marginal borrowers to identifying constraints for a more understudied population, the inframarginal borrowers (Banerjee et al., 2015a).¹⁵ By randomizing employer-mediated contracts that decouple (i)

¹³Recent empirical work has shown evidence of this in agricultural setting with planting technology and in garment factory setting with soft skills (Adhvaryu et al., 2023; Cefala et al., 2024).

¹⁴Recent work studying randomized provision of eyeglasses that correct presbyopia generated sizable productivity gains for manual workers at minimal cost (Reddy et al., 2018).

¹⁵The workers in my experiment meet the lender’s underwriting criteria but still cannot execute the loan (in my setting - due to the down-payment requirement).

relief of the upfront down payment from (ii) a reduction in total lifetime price, I find evidence that upfront liquidity — not a pure price cut — moves investment decisions among workers.¹⁶ Advancing a large literature — mostly in agriculture — that shows that advancing cash at the moment inputs must be purchased materially raises adoption of high-return inputs or insurance contracts, consistent with binding short-run liquidity and present-bias (Duflo et al., 2008, 2011; Casaburi and Willis, 2018; Jack et al., 2025).

The third contribution of my paper is to the literature on mobility and productivity, especially for female workers already in the labor force. Prior work shows that lowering the generalized cost of travel raises work and earnings: subsidizing seasonal bus travel in Bangladesh increased migration and income, and randomized access to motorcycle financing in Brazil boosted formal employment and wages (Bryan et al., 2014; Van Doornik et al., 2024). I build on these insights within a firm–worker context and for women specifically.

A large, complementary literature documents that mobility constraints suppress women’s labor supply in both rich and poor settings (Le Barbanchon et al., 2021; Abou Daher et al., 2023; Borker et al., 2021; Garlick et al., 2025). I extend the evidence from labor-force entry to what improved mobility does for productivity conditional on employment. By enabling private mobility through worker-owned scooters, I show that employed women complete more jobs, shift toward farther and late-evening tasks where returns are higher, earn more, and deliver at least as high service quality; in parallel, savings improve and attrition falls. These results establish that mobility is not only an entry constraint but a first-order production input for working women, with measurable gains in output, job mix, and retention inside the firm.

Finally, these results hold direct policy relevance in India. Tamil Nadu’s Amma Two-Wheeler Scheme, launched in 2018, offers a 50% subsidy (up to INR 25,000) explicitly to help low-income working women to purchase a scooter.¹⁷ Uttar Pradesh has earmarked INR 400 crore (USD 46 million) for the new Rani Laxmi Bai Scooty Scheme to provide scooters to meritorious Class-12 girls entering college, and Rajasthan continues statewide distribution under the Kalibai Bheel Medhavi Chhatra Scooty program through its Higher Education Department — both indicating substantial, recurring budget commitments.¹⁸ My evidence points to a complementary, private-sector path: a within-firm contract that eases the upfront cost of a scooter for working women, financed and implemented by employers who can appropriate part of the productivity gains.

The rest of the paper is organized as follows. Section 1.2 describes the setting, the platform firm’s operations, and the role of mobility in the platform job. Section 1.3 presents the field experiment, including the intervention, data sources, summary statistics, and empirical specification. Section 1.4 reports the main results, with subsections on worker outcomes and firm outcomes. Section 1.5 quantifies returns on investment for the firm from the subsidy and for workers from scooter purchase. Section 1.6 develops the mechanism test, guided by a simple model and an incentive-compatible discrete-choice experiment, it

¹⁶The measurement strategy follows the contract-choice tradition in consumer credit—e.g., subprime auto studies that elicit preferences over down-payment and price menus (Adams et al., 2009) — but extends it in two important ways: the choices are incentive compatible with real stakes, and the setting is an employment relationship where the firm can appropriate part of the surplus generated by worker-owned capital.

¹⁷[News Article - Launch of Tamil Nadu Scheme. Further details of Tamil Nadu Scooter Scheme.](#)

¹⁸[News Article - Launch of Uttar Pradesh Scheme](#)

separates liquidity constraints from price rebates to explain why the subsidy raised adoption.

1.2 Study Setting

1.2.1 The Firm

For the study I partner with a large app-based platform-firm in India that intermediates home services. The firm matches customers with electricians, plumbers, beauticians, cleaners, and other specialists through a mobile application. As of early 2025, it operates in over 30 cities, engages roughly 50,000 workers, and has served over 10 million customers. This organizational form exemplifies the formalization of service work through digital intermediation in the Global South (NITI-Aayog, 2022; Datta et al., 2023). Conceptually, the setting is closest to U.S. platforms such as Thumbtack or Taskrabbit, but it differs along several institutional margins. Unlike the U.S. service platforms, where customers select an independent worker based on observables such as hourly fee and experience, in this firm customers book a service and the firm assigns the worker via an internal assignment algorithm; services are standardized and accompanied by firm-led onboarding and periodic training; worker performance is routinely monitored against operational thresholds; and worker onboarding includes rigorous background checks and identity verification.

The experiment focuses on the firm’s beauty services category, in which technicians deliver at-home salon services—hair care, waxing, facials, pedicures, and manicures. Tasks are short to medium in duration (typically one to three hours), drawn from standardized menus with standardized inputs supplied by the worker, and demand exhibits pronounced diurnal peaks in early-morning and late-evening slots. Both men and women work in this category; service is gender-matched (female technicians serve female clients; male technicians serve male clients). Customers select bundles of services in the app, which determines the booking value. Because services are performed at the customer’s residence, efficient travel is a first-order input to effective labor supply and an important determinant of margins.

Workers face two types of salient outside options: (i) employment with similar app-based intermediaries - there are about two to three national competitors in the at-home beauty segment - and (ii) employment at brick-and-mortar salons (fixed location, manager-set schedules, lower travel but less autonomy).¹⁹ Multi-homing is feasible - the platform does not require exclusivity, so workers could allocate time across competing apps and salon jobs. In practice, most female workers on this platform neither multi-home nor hold other jobs; in my experimental sample, 99% report this platform as their sole source of labor income.

At the start of each week, workers declare availability windows in the app (days and time blocks) and can also renege by toggling off in real time. In this setting, I measure *available hours* as the total time a worker is online and eligible to receive offers —explicitly inclusive of time waiting for jobs, performing jobs, and traveling between locations — which is the operative measure of labor supply. Customer requests are routed to available workers via a centralized algorithm that *sequentially* dispatches offers; workers observe job details (expected earnings, distance, duration, time) and choose to accept or decline. Service prices

¹⁹A third option—independent customer acquisition via neighborhood networks or digital classifieds—exists but is less common in our sample.

are standardized and fixed at the category level; there is no intra-day dynamic pricing to clear the market. Pre-set surge multipliers apply only to early-morning and late-evening hours. The firm takes a 20–30% commission on completed jobs.

The platform continuously tracks punctuality, acceptance, cancellations, and customer ratings, enforcing minimum thresholds. Frequent cancellations or sustained inactivity trigger short blocks that temporarily remove the worker from the pool until re-training. All workers complete a formal on-boarding prior to activation and periodic firm-led skill upgrades.

1.2.2 Mobility as a Production Input for Jobs

As service delivery occurs at customers’ residences, commute time, reliability, and routing flexibility are first-order inputs to production. In these cities — broadly representative of urban areas in developing economies — last-mile public transport is patchy; own two-wheeler scooters and private taxis (including ride-hailing) are the relevant modes.²⁰ In this context, a scooter is *worker-owned capital* that augments effective labor supply by potentially lowering per-job travel time and variance, expanding feasible spatial-temporal coverage, and reducing variable transport costs.

Prior to the experiment, we conducted a phone survey of 1531 randomly selected platform workers which showed a large gender gap in actual commute modes for the same salon-service job on the platform. About 77% of men report commuting by their own two-wheeler scooters, relative to 31% of women (Figure 1.1, Panel A). However, preferences are similarly high across genders: 79% women and 92% men state that they would prefer commuting by own two-wheeler scooter rather than private bike-taxis or three-wheelers (Figure 1.1, Panel B). This usage-preference gap motivates focusing the intervention on women, for whom adoption is far from saturation. Baseline survey highlights key issues with relying on ride-hailing rather than using one’s own scooter. The main ones include difficulty obtaining rides at peak times; higher and more volatile fares in early/late hours; refusals by bike-taxi drivers to carry large service kits; concerns about harassment and unsafe driving. Workers also emphasize that search and waiting time displaces productive hours and rest between appointments.

Administrative data show suggestive productivity differences at baseline. Relative to workers who depend on private taxis for commute, women who already commute by scooter show higher daily earnings and spend less on transport, while the number of services completed per day is similar across groups (Figure 1.2). These patterns are consistent with scooters enabling acceptance of better-paying jobs at lower effective cost.²¹

These institutional facts — (i) mobility as a key input, (ii) large gender gaps in usage despite similar preferences, and (iii) baseline productivity gradients correlated with scooter use — motivate an experimental subsidy targeted to women. In the analysis, I separate *new adopters* (women not using a scooter at baseline) from *incumbent users* (already using a scooter at baseline). If the subsidy operates by inducing a *change in mobility modality*

²⁰Car ownership and use are rare in this worker demographic.

²¹Wealth covaries strongly with both productivity and scooter use (Figure A2); the probability of commuting by own scooter rises monotonically across wealth quintiles (Figure A2). Household scooter ownership is also monotone in an asset-based wealth index: 37% of households in the poorest quintile possess at least one scooter, 63% in the second, and over 90% from the third through the richest quintiles.

(from taxis/three-wheelers to own-scooter commuting), we should not expect productivity effects on the platform among incumbent users, for whom commuting modality is unchanged. Accordingly, all estimates are reported separately for new adopters and incumbent users.

1.3 Field Experiment: Subsidy for Worker Owned Capital

1.3.1 Intervention Details

The field experiment randomly offered a one-time purchase subsidy for a two-wheeler (“scooter”) — the worker-owned capital in this setting — to female workers who had ex-ante indicated interest in buying a scooter. Figure A1 shows an example scooter in this setting.²²

In the study period, the average invoice price of a scooter was INR 106,000 (approximately USD 1,200). Prices varied across makes and models but only modestly: the interquartile range spanned INR 90,000–123,000 (USD 1,094–1,440). At prevailing earnings, the purchase price corresponds to about twenty weeks of income for this population, implying that acquisition is typically credit-financed. Consistent with standard vehicle-finance contracts in this market, the loan requires a down payment equal to 15% of the invoice price, due at purchase; the remaining 85% is financed at a fixed nominal annual interest rate of 18%. Borrowers select the tenure within a customary range of 12–36 months; the realized mean in the experiment is 26.625 months (s.d. = 7). The loan is collateralized by the asset, so persistent delinquency can trigger repossession of the scooter.

Eligibility required a baseline CIBIL credit score ≥ 700 . Operationally, 74% of firm workers satisfied this threshold. Two considerations motivated this criterion. First, it guarantees that any worker randomly offered the subsidy could, if she chose, transact and purchase the scooter on credit. Second, it permits a clean focus on *post-credit-access* frictions: the study concentrates on infra-marginal borrowers to whom the market was already willing to lend, yet who nevertheless did not purchase at the prevailing terms.

CIBIL is one of the main credit bureaus in India. Because its scoring model differs from U.S. FICO, direct level-by-level comparisons are not meaningful; heuristically, a CIBIL score of 700 is often mapped to 600–650 FICO, but this correspondence is only approximate.²³

²²Throughout this paper the term scooter (or the colloquial “scooty”) denotes the lightweight, step-through, gear-less two-wheelers, typically 90-125 cc petrol or electric vehicles such as the Honda Activa and TVS Scooty, that account for 6.85 million wholesale sales in FY 2025, a 17% year-on-year increase and a quarter of India’s total two-wheeler market (Autocar Pro). Unlike the stand-up, dockless “kick-scooters” popularised by Lime or Bird in North America, these Indian scooters are ridden seated, licensed for road traffic, and designed for trips well beyond first- and last-mile range. Their step-through frame, continuously-variable (gearless) transmission and low weight reduce physical exertion and make them especially attractive to novice and female riders. Manufacturers increasingly market scooters to women because females constitute only about 10% of current two-wheeler owners, leaving substantial head-room for growth (The Economic Times 2025). At the national level nearly half of all Indian households already own some form of two-wheeler (The Times of India 2023).

²³There is no authoritative CIBIL to U.S. (FICO) conversion. Some providers (e.g., Nova Credit) offer proprietary, odds-based translations that remap a foreign bureau score into a U.S.-equivalent score range;

I document two diagnostic facts about the eligibility screen. First, the empirical CDF of baseline credit scores is smooth in a neighborhood of 700, with no visible bunching, consistent with the absence of manipulation; moreover, a majority of applicants lie above 700 (Figure A4). Second, baseline covariates are well balanced just below versus just above 700: education, marital status, recent earnings, jobs completed, travel expenditure, pre-experiment scooter-loan activity, and household composition show no economically meaningful discontinuities (some individual p -values are marginal, but the joint pattern indicates tight balance). Note that a lower CIBIL score at a point in time may also reflect a recently originated loan. These checks support the interpretation that conditioning on $\text{CIBIL} \geq 700$ secures credit availability without inducing material composition shifts.

Treated workers receive a flat, one-time purchase subsidy of INR 10,000 (USD 120), paid directly to the dealership at the time of purchase rather than to the worker. The transfer reduces the effective purchase price by about 10% relative to the sample mean scooter price. Because it is applied at the point of sale, the subsidy is maximally salient at the down-payment stage: it lowers the scooter’s net price exactly when the down payment is due, while leaving the financed principal unchanged. Control workers receive no subsidy; they are, however, connected to the lender and may purchase on the standard market contract. Later, in Section 1.6, I present a simple model that separates liquidity from price channels to interpret the mechanisms.

From August to October 2024, an in-app invitation to apply for the scooter-subsidy study was sent to 3,000 randomly selected female salon workers across Delhi, Pune, and Bengaluru. Approximately 1,200 workers (a 40% response rate) contacted the research team: 50% from Delhi, 25% from Bengaluru, and 25% from Pune. A light-touch verification (government ID and address check) was completed by the research team to screen out low-commitment inquiries. Enumerators then administered a baseline survey and obtained documentary proof of each applicant’s credit score. Workers meeting the credit-access criterion ($\text{CIBIL} \geq 700$) entered rolling, city-specific cohorts. Within each cohort, I stratified by baseline commute mode (already using a scooter vs. not using a scooter) and randomized assignment, offering the subsidy to 60% of eligible workers. The resulting experimental sample comprises 350 women.²⁴

1.3.2 Data Sources

The empirical analysis combines three complementary sources: (i) high-frequency administrative records on worker productivity and firm revenue, (ii) a pre-intervention baseline survey, and (iii) three waves of post-intervention follow-up surveys.

Nova’s documentation and sample reports describe this as a mapping based on underlying default-odds tables rather than a deterministic crosswalk. For intuition only (not a numeric crosswalk), note that Indian bureaus use a 300–900 scale and describe $\approx 700+$ as “good” while FICO scores typically span 300–850 with “good” defined as 670–739. (Source).

²⁴Sample-selection funnel: From 1,200 initial inquiries, 697 were screened out at the light ID and address verification and baseline scheduling stage, leaving 503 candidates (41.9% of initial). The documentary credit-score screen ($\text{CIBIL} \geq 700$) excluded a further 153, yielding 350 eligible workers (29.2% of initial; 69.6% of post-ID stage). Randomization was conducted within rolling city cohorts, stratified by baseline scooter use, with the subsidy offered to 60% of eligible workers.

The administrative data spans March 2024–June 2025 and comes from two complementary files maintained by the platform. The *calendar dataset* provides a day-by-day snapshot of intended and realized availability. In particular, workers declare the hours they are willing to work (up to a week in advance) and may revise these declarations at any time; the file therefore records, for each worker–day, (i) whether the worker reported herself available or absent, (ii) whether availability covered the full day or specific hourly blocks, and (iii) whether the worker was temporarily suspended for performance shortfalls. Because availability is the gateway to matching, these records define the *active matching pool* by day. I then merge the calendar data with the *job-level master data* at the worker–day level, and later aggregate to worker–weeks to smooth transitory shocks.²⁵

The *job-level master dataset* records *every offer* sent to a worker on a day she has marked herself available—irrespective of whether the offer is accepted—thereby revealing the full *choice set* conditional on availability. For each offer the file records: the worker’s location identifier and the client location (used to compute home–to–client distance), the scheduled start time and expected duration, the combination of service-tasks involved, the worker payout (income) that would be earned if accepted, and the platform’s commission, revenue per job. Conditional on acceptance, the dataset adds operational and quality outcomes: completion status, arrival-time stamps (punctuality), cancellations or reschedules, the client’s rating, and any complaint flags. This structure allows me to (i) trace *where* and *when* work occurs (distance bands and time-of-day windows), (ii) measure reliability and quality using the same indicators the firm uses operationally, and (iii) observe *both sides of the transaction* — worker earnings and the firm’s realized profits — at high frequency.

The baseline survey, administered immediately before the intervention, records a rich set of demographic and socioeconomic characteristics such as household composition, primary commuting mode, whether the platform job is the respondent’s sole employment, prior scooter-loan uptake, and related covariates. To measure pre-specified, mechanism-relevant outcomes that are not observable in administrative records, I fielded three follow-up surveys at approximately equal intervals. Each wave collects a job-by-job commuting log for the preceding day — transport mode and round-trip expenditure for every completed job. Waves 1 and 3 additionally elicit the employment status and earnings of all adult household members to measure any spillovers. To assess women’s autonomy, the instruments include modules on self-confidence, decision-making authority, and freedom of movement. I construct a self-confidence index by averaging standardized responses to four Likert-scale items covering household decision-making, task self-efficacy, pursuit of personal interests, and independent mobility. Leisure is summarized by an index based on four binary indicators for the prior week: visiting friends, visiting relatives, undertaking an outside leisure outing (cinema, mall, or park), and engaging in self-care. Survey attrition remains below ten percent in every wave and is balanced across treatment arms.

²⁵Results using daily data are qualitatively identical and sufficiently powered. I report worker–week aggregates to smooth idiosyncratic day shocks (e.g., weather, illness, platform outages), absorb strong day-of-week seasonality; the weekly horizon also aligns with the platform’s operational accounting cycle.

1.3.3 Summary Statistics and Baseline Balance

Table A2 presents summary statistics and baseline balance. Looking at the full sample, a typical worker is about 34 years old. About 75% are married and nearly 88% report having at least one kid. About 30% of the sample already use their own scooter for work, and among women who don't use their own mode to commute, about 80% hail two or three wheeler taxis (remaining rely on a family member to drop them, while use of public transport like metros or buses is negligible - due to the last mile connectivity requirement for these jobs). 90% of the sample report knowing how to drive a scooter at baseline. On average, respondents report two earning members in a household, with about 8% reporting having an unemployed husband. For 98% workers, platform work is their sole and regular source of earnings.

The wealth index is constructed as the first principal component from a principal component analysis of household economic indicators, including housing ownership (owned or rental) and ownership of consumer durables such as a refrigerator, television, laptop, personal vehicle, air conditioner, and washing machine.²⁶ At baseline, scooter use is strongly increasing in wealth: the share who drive their own scooter rises from 13.5% in the poorest quintile to 21.2%, 30.8%, 40.3%, and 45.8% in successively richer quintiles (Figure A2).

At baseline, workers had about three years of platform tenure, completed about one service per day, and earned around 790 per day and 5,200 per week; the mean app rating was 4.84 out of 5. Baseline means are similar across treatment and control in the pooled sample and within the incumbent-user sub group and new-adopter sub-groups; differences are small and rarely statistically significant.

Figure 1.3 documents two empirical regularities in the job-level administrative data that motivate the subsequent analysis of scooter adoption. Panel (a) of Figure 1.3 shows that job-value rise with client distance from the worker's home location; Panel (b) of the same figure shows that earnings per minute are higher for jobs booked in the early-morning (7–11h) and late-evening (17–19h) windows than for mid-day slots. These unconditional patterns indicate that the return to mobility is not homogeneous: the marginal value of an additional kilometer of reach and of availability at off-peak hours is higher. Because scooters primarily relax travel-time and travel-variance constraints, they are plausibly most productive exactly on these margins.

Similarly, Figure 1.4 shows that, on average, 84% of jobs for a given day are fulfilled (completed) in the system. However, completion rates are substantially lower at the early and late time windows relative to mid-day, and decline with client distance, consistent with higher travel frictions and tighter punctuality constraints in those segments. Put differently, the platform compensates workers more for far-away and early or late jobs, but those jobs also carry a higher probability of non-completion. Taken together, these facts rationalize why a scooter *could* raise earnings not merely by doing more of the same jobs or by lowering per-job travel costs. Adoption may expand the feasible set toward (i) farther jobs and (ii) early or late bookings that pay more, while (iii) improving reliability on those riskier segments by reducing travel time and its variance. The figures thus provide descriptive context and pre-analysis guidance for heterogeneous treatment effects by distance and booking hour.

²⁶This is the standard method used to create wealth indices in development economics literature such as in De Mel et al. (2008).

1.3.4 Empirical Strategy

To estimate the impact of the subsidy offered, I use the following empirical specification:

$$Y_{it} = \beta_1 \text{Subsidy}_i + \beta_0 Y_{0iDPL} + \delta_C \times \lambda_S + \gamma_t + \varepsilon_{it} \quad (1.1)$$

where Y_{it} is the outcome variable of interest (e.g. worker’s weekly jobs or earnings or firm’s profits), Subsidy_i is the indicator for whether the worker is offered the subsidy (i.e. treatment), Y_{0iDPL} is the baseline value of the dependent variable, δ_C are randomization cohort $\times$ λ_S randomization strata fixed effects, γ_t represents week fixed effects for all administrative data analysis - as data is aggregated to the weekly level. Standard errors are clustered at the worker level. The analysis uses 8 months of post-randomization data collapsed at the weekly level. $\hat{\beta}_1$ thus provides the intent-to-treat causal estimate of being offered a subsidy.

For outcomes that are observed only at discrete survey waves (e.g., travel expenditure or household income), I re-specify equation (1.1) with survey-round fixed effects (instead of calendar week fixed effects) so that treated and control workers interviewed in the same wave are compared after netting out any round-specific shocks.

1.4 Results: Impact of Subsidy for Worker Owned Capital

1.4.1 Take-up of the Subsidy

I first study take-up of the *subsidy offer* and, separately, take-up of the *scooter purchase*. By design, only treated workers could redeem the subsidy: 49% of the treatment group used the INR 10,000 offer to purchase a scooter on credit, whereas 0% in control did so because no subsidy was offered. Turning to adoption, over the period from rollout to the final survey round (approximately ten months), the *treatment effect on purchasing a scooter* is 32 percentage points: 49% in treatment versus 17% in control, with the control rate plausibly reflecting baseline purchasing absent the program (as shown in Table 1.1). Figure A5 plots cumulative take-up in the treatment arm. Take-up rises sharply from 15% at randomization to just above 35% after one month, reaching nearly 45% by month 2. Thus, about 90% of the eventual treatment response materializes within the first two months, consistent with a prompt demand response to the cash incentive. Thereafter, the curves largely plateau, over the subsequent eight months, additional uptake is modest.

Table A3 shows that take-up rates are statistically indistinguishable between two baseline groups: *incumbent users* (those already commuting by scooter) and *new adopters* (those not using a private scooter at baseline). Two mechanisms can rationalize comparably high take-up among incumbents. First, replacement and quality upgrading. Many incumbents relied on aging or shared vehicles with high maintenance costs (the mean scooter age among incumbents is seven years), so a discounted new scooter is privately attractive under the subsidy. Second, intra-household reallocation. An additional vehicle may be valuable to other household members, which raises the marginal benefit of purchase even for incumbents.

The estimated treatment effect on purchase is highly stable across observable characteristics. Interacting treatment with more than twenty baseline covariates yields effects concentrated in the 0.30–0.36 range, indicating that the subsidy increases purchases broadly rather than only for specific types. Only a small subset of interactions reaches conventional significance. Among new adopters, take-up is increasing in age and in the number of household members who can ride a scooter.

1.4.2 Effects on Worker Outcomes

Jobs, Earnings and Productivity

Using an eight-month post-randomization panel aggregated at the worker-week level, I estimate intent-to-treat (ITT) effects from administrative data via Equation 1.1. The main sample of interest comprises *new adopters* — workers who did not use a private scooter at baseline (just over 70% of the experimental sample).²⁷ Panel A Column 1 of Table 1.2 shows that treated new adopters complete 0.54 additional jobs per week (s.e. 0.23; $p < 0.05$) relative to a control mean of 6.12. This is a 9% increase, sizable at the weekly horizon and persistent over the eight-month window.

Weekly earnings rise by INR 602.40 (s.e. 204.41; $p < 0.01$) over a control mean of INR 4,757.98, a 13% gain. The earnings effect exceeds the output effect, indicating a composition margin toward higher-value jobs rather than a pure increase in volume (Table 1.2 Panel A Column 2). To measure the response on the extensive margin of labor supply, I use “hours online” from the administrative records, defined as all time windows during which the worker is in the active matching pool—waiting for matches, traveling to jobs, or performing jobs. Using this measure, I find no statistically significant change in time supplied relative to control. The null effect therefore indicates no meaningful expansion in the total time the worker allocates to the platform ecosystem; the earnings gains documented above come from higher productivity within given time supplied rather than from supplying more time.

Because hours do not rise, the earnings gain translates into a higher return per hour online. Earnings-per-online-hour increase by INR 9.82 (s.e. 3.59; $p < 0.01$) relative to a control mean of INR 94.05, an improvement of about 10%. Taken together — +9% jobs, +13% earnings, and no detectable change in hours — these estimates imply that the subsidy allows new adopters to reallocate into higher-value jobs and to earn more per unit of time, rather than simply working longer. The magnitudes are large in weekly terms and remain stable when I vary controls, sample windows, or aggregation.

Incumbent users — workers already commuting by scooter at baseline — serve as a natural placebo group for productivity effects: purchases here are predominantly replacement (or upgrading an older scooter) or intra-household reallocation of the new scooter rather than first-time expansions of mobility. Consistent with that interpretation, Panel B shows no detectable effects on any outcome. The ITT estimates are small and imprecise relative to the corresponding control means, suggesting that the subsidy does not raise measured productivity for incumbents (Table 1.2 Panel B).

In the pooled sample (Table 1.2 Panel C), the subsidy raises weekly earnings by INR 322 (s.e. 170; $p < 0.10$) from a control mean of INR 5,098, and increases earnings per online hour

²⁷Recall that randomization was stratified at baseline by new adopters vs. incumbent users.

by INR 5.37 (s.e. 3.21; $p < 0.10$); effects on jobs and hours online are small and imprecise. These averages mask sharp heterogeneity: the gains are concentrated among *new adopters* (baseline non-users), who earn more without supplying additional online time—consistent with reallocation toward higher-value tasks—whereas *incumbent users* (baseline users) exhibit no detectable effects on any margin, consistent with replacement or intra-household reallocation rather than capacity expansion.²⁸

Commuting impacts. Using three follow-up surveys per respondent, I collect detailed, job-level travel data that allow precise measurement of per-job commuting expenditures and time. As shown in Table A7 column 1, the subsidy offer reduces travel expenditure for new adopters (Panel A) and incumbent users (Panel B). For new adopters, per-job travel expenditure falls by about 16%. For incumbent users — who already rode scooters prior to treatment — the decline is about 20 percent. The larger percentage drop for incumbents may seem at odds with the bigger overall gains for new adopters, but the comparison reflects composition: treated new adopters reallocate toward farther jobs (as discussed in next sub-section and Table 1.3), which mechanically dampens the per-job saving even as their total earnings and job counts rise. In fact, it might appear surprising that travel expenditure should fall for incumbent users. The fall however is mechanically consistent with upgrading an aging asset: replacing the median seven-year-old scooter raises kilometers-per-liter by 25 percent, which, holding fuel prices and routing fixed, maps one-for-one into lower cost per kilometer and therefore lower cost per job. Time outcomes display the expected asymmetry: travel time per job declines for incumbents, while first-time users show no precise change—plausibly reflecting learning and parking frictions during the transition to own-vehicle commuting. Overall, these patterns are exactly what a worker-owned-capital view predicts: the offer compresses the shadow price of mobility for all workers, and when the asset was already in use, quality upgrading (higher mileage) and smoother operations translate into both monetary savings and time efficiencies.

Spillover effects on other household members. As shown in Table A7 column 3, the subsidy raises family-member earnings, and this gain is concentrated in households of incumbent users. There is no movement in the husband’s unemployment rate (Table A7 column 4), so the adjustment operates on the intensive margin — more hours or higher-pay tasks — rather than through new employment. This pattern is consistent with the worker-owned-capital mechanism operating on the *other* household worker: incumbent user’s were able to allocate an extra scooter to the second worker in the household, who was previously constrained in accessing productive work. In contrast, for new adopters the estimated change in family-member income is not statistically distinguishable from zero, consistent with the idea that the primary margin for the households of new adopters is the treated worker’s own transition to scooter use rather than an immediate reallocation to another member. Taken together, the evidence points to worker-owned capital relaxing binding constraints within

²⁸Daily level estimates are qualitatively identical and adequately powered; I present worker-week aggregates to smooth day shocks. Tables A4, A5 and A6 show robustness for main results for new adopters, incumbent users and pooled sample. The results are similar to the main specification presented and discussed in Table 1.2.

the household and lifting earnings where capacity to use the asset pre-existed.

Reallocation by Distance and Time Window

I next examine how the subsidy changes the composition of completed jobs for new adopters. Table 1.3 shows a sizable expansion in spatial reach: treated workers travel an additional 6.67 kilometers per week (s.e. 2.55, $p < 0.01$) relative to a control mean of 58.40 km, an 11.4% increase. Given that the new adopters are doing more jobs, the fact that they travel more is not surprising.

However, interestingly, the composition of trips shifts toward longer hauls (Table 1.3, Columns 2-3). There is no detectable change in *near jobs*, defined as jobs with round-trip distance ≤ 10 km, the estimate statistically indistinguishable from zero relative to the control mean of 3.49. By contrast, *far jobs* — assignments with round-trip distance > 10 km and ≤ 25 km (i.e., 5–12.5 km each way) — increase by 0.54 per week (s.e. 0.23, $p < 0.05$) over a control mean of 2.68, a 20% rise. A Wald test rejects the null that the near- and far-job coefficients are equal ($p = 0.005$), indicating that the distance margin is driven by additional medium-to-long trips rather than changes in nearby work.

Consistent with the descriptive facts that early-morning (before 11h) and late-evening (after 17h) bookings carry both higher pay and higher non-completion risk, while mid-day (12–16h) jobs are easier to fulfill, scooter access reallocates labor supply toward the riskier, better-paid windows (Table 1.3 Column 5). Mid-day activity is essentially unchanged (Table 1.3 Column 4). By contrast, early and late jobs rise by 0.37 per week (s.e. 0.21, $p < 0.10$) from a control mean of 3.10 — about 12% increase. A formal test cannot reject equality of the early/late and mid-day effects ($p = 0.195$), but paired with the distance results, the pattern is economically coherent: adoption expands workers’ feasible set exactly where travel-time and travel-variance constraints bite most (farther, off-peak work).

This reallocation toward segments with both higher posted returns and higher baseline failure risk translates into substantial gains in contingent pay: surge earnings increase by INR 18.67 per week (s.e. 6.66, $p < 0.01$) over a control mean of INR 43.97 — a 40% rise as seen in Table 1.3 Column 6.

Job Quality

Client ratings rise by 0.03 points (s.e. 0.01, $p < 0.10$) from a baseline mean of 4.83 on a 5-point scale. Because ratings are tightly bunched above 4 (top-coding is pervasive), this is a meaningful shift: it closes roughly 17% of the residual gap to a perfect score and pushes the average past 4.85, a threshold the firm informally treats as “excellent” (Table 1.3 Column 7).

Despite the reallocation toward farther and off-peak jobs, the “late to job” indicator does not improve, as seen in Table 1.3 Column 8: +0.18 (s.e. 0.16) on a control mean of 2.49 incidents per worker, statistically indistinguishable from zero. This pattern is consistent with offsetting forces: scooters reduce travel time and its variance (mitigating delays), but learning and parking or search frictions for new adopters, along with longer average trip distances, limit mechanical gains in punctuality. The net quality effect therefore appears to operate primarily through the interaction margin — workers arrive less stressed and with

lower physical strain (e.g., less back pain from ride-hailing and waiting), leading to modest but economically relevant improvements in client ratings without a detectable change in tardiness.

1.4.3 Effects on Firm Outcomes

Firm Revenue and Profits

For new adopters, i.e. workers who were not using a scooter pre-experiment, the subsidy translates into sizable, internally consistent gains for the platform. Table 1.4 shows impact on the firm. Panel A Column 1 shows that revenue rises by INR 1,178.95 per worker-week (s.e. 404.54; $p < 0.01$) on a control mean of INR 9,369.21, an increase of 13%. After paying workers, the platform’s revenue net of payouts increases by INR 575.72 (s.e. 202.63; $p < 0.01$) on a control mean of INR 4,611.23 (Panel A Column 2). Profits move one-for-one with these margins (in percentage terms): +INR 265.26 (s.e. 91.02; $p < 0.01$) relative to a mean of INR 2,108.07 (12–13%) (Table 1.4 Column 3). The pattern is consistent with the mechanism that by relaxing mobility constraints, adoption shifts labor supply toward farther and off-peak jobs with higher marginal revenue, raising both gross receipts and the platform’s commission base.

Among incumbent users, point estimates are small and imprecise across gross revenue, net revenue, and profits (Table 1.4 Panel B). This null is predicted by the mechanism. The subsidy does not relax a binding mobility constraint for incumbents, who had already selected into using a scooter and into the high-return distance or time margins. Control means in fact show that incumbent users generate more revenue than new adopters at baseline — gross revenue INR 11,219 vs. 9,369, net revenue 5,420 vs. 4,611, and profits 2,524 vs. 2,108 — consistent with prior selection into scooter use and high-return segments. The subsidy lifts new adopters’ gross revenue by INR 1,179, net revenue by INR 576, and profits by INR 265 per week, closing about two-thirds of the incumbent–new baseline gap on each margin.

When pooling the full sample, the ITT estimates are positive and significant (Table 1.4 Panel C). Gross revenue rises by INR 681.59 per worker-week, revenue net of payouts by INR 336.30, and profits by INR 153.36 (all $p < 0.05$). These pooled estimates are a weighted average of large, precisely estimated gains for first-time adopters and near-zero effects for incumbents, so they primarily reflect the share of new adopters in the sample and the magnitude of their treatment effects. In aggregation, pronounced heterogeneity at the subgroup level appears as moderate but statistically meaningful improvements at the firm level.

Attrition

Table 1.4 Column 4 examines whether improved mobility induces exits from the platform over the subsequent year. In the control group, inactivity at one year is 0.109 for new adopters and 0.022 for incumbent users, which indicates that incumbents are much less likely to leave in control group. The ITT effect on attrition for new adopters is -0.040 with s.e. 0.045 (Table 1.4 Panel A Column 4), which implies a reduction in attrition of about four percentage points from a 10.9 percent baseline although the estimate is imprecise. For

incumbents the ITT effect is -0.015 with s.e. 0.037 on a 2.2 percent baseline. Pooled across all workers the ITT effect is -0.036 with s.e. 0.031 on a 8.5 percent baseline.

Taken together, these estimates show higher productivity for treated workers with no detectable increase in exits. This pattern matches the [Acemoglu and Pischke \(1999\)](#) view that firm investment is sustainable when labor market frictions or wage compression allow the firm to retain part of the surplus generated by the investment. Here the scooter subsidy is akin to firm-sponsored training. Workers become more productive, yet separations do not rise.

Several mechanisms could plausibly explain this behavior on the ground. The platform aggregates time-sensitive demand and reduces search costs, which are difficult to reproduce off platform. Reputation and ratings are partly platform specific, which limits portability. Switching costs are present because new adopters learn routes and customer routines over time. Outside opportunities may offer less schedule flexibility and weaker payment reliability, features that are especially valuable for women who face tighter household time constraints. Safety and the option value of surge access further raise the value of staying. These frictions help the platform retain a share of the marginal product created by the subsidy, which rationalizes firm investment in worker-owned capital without inducing higher attrition.

1.5 Return to Investment

1.5.1 Returns to Subsidy for the Firm

To compute returns to the firm, the relevant causal estimate is the treatment effect on the treated. The firm pays the subsidy only when a worker takes it up. Since the subsidy was offered only to the treatment group and never to the control group, the experiment generates one-sided noncompliance. Take-up was 49 percent in treatment and zero in control, which delivers a strong first stage. I therefore estimate the effect of subsidy by instrumenting take-up of subsidy with the offer of subsidy. Under the standard exclusion and monotonicity assumptions, the IV coefficient identifies the average causal effect for adopters in the treatment arm. This is the treatment effect on the treated that is required for the return calculation.

Before moving to the returns to the firm, I first establish that the measured gains do not come from work displacement of other workers. A controls-only event study that indexes weeks by each control worker's work week shows no decline in weekly jobs after assignment as seen in [Figure A7](#). Second, the intensive-margin gains we saw above for treated workers are concentrated in early and late hours and at farther distances where the platform records high unfilled demand and higher surge pay, which points to previously unused capacity rather than crowd-out. Third, a representative city has roughly three thousand active workers and about thirty compliers in our experiment. The estimated effect is about two extra jobs per complier per week. Even if every additional job performed by a complier were taken from someone else, the implied loss would be about thirty times two divided by three thousand which is approximately two hundredths of a job per worker per week. This magnitude is economically small. Taken together, these facts indicate that the revenue and profit gains reflect added production rather than reallocation.

With job displacement unlikely to be first order, I proceed to quantify the payback period and the firm’s return on investment. Table 1.5 Panel A shows that the TOT weekly profits rise by INR 557.87 for first-time users and by INR 345.95 in the pooled sample. To compute recovery time, I compare the subsidy cost to the weekly profit lift. The implied payback period is about eighteen weeks for the new adopter group and about twenty-nine weeks for the pooled sample (that is, including new adopters and incumbent users).

For the twelve-month return, I cumulate incremental weekly profits over the first year and account for observed survival. Months one through eight are realized in the data. For months nine through twelve I scale the weekly effect by the observed twelve-month survival rate. I then subtract the subsidy cost and divide by the cost to obtain an annualized return per subsidy delivered. Using this procedure, the return per treated worker is about 153% for new adopters. The corresponding figure for the pooled sample is about 60%. These estimates use conservative assumptions and match the implied payback windows reported above.

Two conclusions follow. First, for workers who newly adopt a scooter, the firm recovers its outlay within the first half-year and earns a large positive return by month twelve. Second, even after averaging in incumbents who were unconstrained at baseline, the investment remains profitable over the first year. Combined with the flat attrition results, the pattern is consistent with the firm recouping part of the productivity gain in a labor market with meaningful frictions, and the subsidy therefore expands joint surplus rather than re-shuffling output across workers.

Given large returns to subsidy for the firm, a natural question that arises is why the firm had not already financed scooter adoption. Three considerations emerge consistently in interviews and internal discussions with firm managers. First, contractibility is imperfect. Scooters are a general-purpose asset that workers can redeploy outside the platform, which makes full financing risky when separation risk is non-trivial and collateral is weak. Second, ex ante uncertainty about payback was material. Prior to the experiment, managers could not verify whether incremental profits would be sufficient to recover the subsidy within a reasonable horizon, especially given concerns about displacement and survival. Third, organizational frictions matter. Administrative capacity was limited, competing projects drew on the same teams, and piloting a new reimbursement workflow had non-trivial fixed costs.

The experiment reduces each of these barriers. It produces clean evidence on take-up, payback, and survival, and it shows no displacement of untreated workers. In response, the firm has scaled the program to additional cities. With frictions that allow the platform to retain part of the surplus, the firm now finds it optimal to invest in worker-owned capital.

1.5.2 Returns to Scooter Purchase for the Worker

For workers the object of interest is the payoff to purchasing a scooter on credit. In the treatment group, 49% purchase; in control, 17% also purchase. Because some treated workers do not buy and some control workers do, adoption exhibits two-sided noncompliance. I therefore estimate the effect of purchasing a scooter by instrumenting actual purchase with randomized offer of the subsidy. Under standard exclusion and monotonicity assumptions the IV coefficient identifies the local average treatment effect (LATE) for compliers. This is

the effect relevant for workers who adopt when offered.

A potential threat to the exclusion restriction is a direct income effect for always-takers in the treatment group. About 17% would have purchased a scooter irrespective of assignment. In treatment they also receive the INR 10,000 subsidy, which could raise earnings or hours through a wealth shock rather than through adoption itself. Three considerations make this unlikely to matter for our estimates. First, the transfer is small relative to baseline earnings. INR 10,000 corresponds to 2.3 weeks of income, and labor supply is not very responsive to such modest wealth shocks at weekly horizons. Second, even if there is some response, the elasticity of labor supply to small wealth changes is plausibly very low, so any bias would be limited in magnitude. Third, compliers constitute the majority of treated adopters, which mechanically attenuates any influence from always-takers. Taken together, these points make a direct-income violation of exclusion unlikely. I therefore interpret the IV coefficient as the causal effect of purchasing a scooter for compliers.

Table 1.5 Panel B shows that for new adopters the IV estimate of weekly earnings is INR 2,018 with a standard error of INR 768. The weekly loan payment in the standard credit contract is approximately INR 1,040. Incremental earnings thus exceed the installment from the first week, so the investment is cash-flow positive immediately. These estimates imply an expected payback period of approximately 54 weeks (see Appendix 1.B for details).

I next quantify a two-year return on investment for adopters. A two-year horizon is natural for workers in this setting. In my experimental sample, women’s average tenure on the platform is close to three years, and conditional on surviving the first year many remain attached for six years. I therefore report a 24-month return that aligns with the asset’s useful life and observed firm attachment. The calculation cumulates incremental weekly earnings over two years, discounts future flows at the worker’s borrowing rate, and applies survival adjustment for months nine through twenty-four using the observed hazard rate. I then compare the present value of incremental earnings with the financed purchase cost.²⁹ The implied two-year ROI for the new adopters is about 73%. While this computation assumes the worker’s cost net of subsidy, even without accounting for the subsidy, the ROI over 24 months is 57%. For completeness I also compute the annual internal rate of return for the worker’s asset which is over 100%.³⁰

Why, then, did many workers not invest prior to the experiment? Mis-calibrated beliefs do not appear to be the primary explanation. Baseline surveys show that the average worker expected larger gains than those realized ex post. Mean expected monthly revenue gains are around INR 27,400 while realized gains are closer to INR 13,600, and expected job increases average 19 per month versus a realized increase closer to 8 (as seen in Figure A8). Belief heterogeneity does predict heterogeneity in treatment effects, but on average expectations are optimistic rather than pessimistic (see Figure A9).

These facts point away from under-expectation and toward direct frictions that the subsidy relaxes, such as liquidity, risk, and learning costs associated with first-time adoption. I investigate these mechanisms in the next section using a targeted mechanism experiment.

²⁹See Appendix 1.B for details.

³⁰The IRR is the discount rate that equates the present value of incremental earnings to the purchase cost, whereas the ROI compares the present value of incremental earnings minus cost to cost over a fixed horizon. Using the same inputs as above, the implied annual IRR is over 100 percent. Full details are in Appendix 1.B.

1.6 Mechanism Experiment: Liquidity Constraint vs. Price Response

1.6.1 Model of Investment Choice by Workers

In the field intervention the subsidy lowers the required down payment at purchase. Subsidy bundles two channels that can raise adoption. It reduces the lifetime user cost of the asset and it relaxes the up-front liquidity constraint. The resulting take-up response does not identify which margin is pivotal, and that distinction is central for identifying the operative constraint on worker investment. Informal interviews with workers pointed to two potential distinct constraints. Some workers report not investing because they would bear the full acquisition cost while capturing only part of the surplus, which weakens private incentives. Others report not investing because the down payment is binding given limited cash on hand at the time of purchase.

To isolate these channels, consider two counterfactual contracts that hold everything else fixed. The first is a pure price contract that reduces the purchase price and leaves the down payment nearly unchanged. The second is a financing contract that smooths the initial cash requirement by allowing the worker to pay only a small amount at purchase and to roll the remaining up-front amount into the financed principal. These two thought experiments motivate the model. The first maps to a return-appropriation friction that depresses the private value of investment relative to its social value. The second maps to a liquidity constraint at the down payment margin. The model below organizes predictions for adoption under each contract, which I test in the mechanism experiment that follows.

Model Set Up.

Consistent with the economic environment in this setting, assume workers are loan-eligible conditional on bringing the required down payment. The lender requires a fixed share $\delta \in (0, 1)$ of the list price $S > 0$ of scooter at time of purchase; the remainder ($\phi \equiv 1 - \delta$) is financed on credit. The financed principal accrues interest at a per-rupee present value interest $\iota \geq 0$. The worker chooses the level of investment k , given liquid cash L , retained share $\alpha \in (0, 1]$, and technology f with $f' > 0 > f''$. In Appendix 1.C, I re-write the model using a k as a discrete lumpy investment decision instead of continuous (as shown here) - and the predictions are the same. Presenting continuous k version here for simplicity.

The two levers that the firm can implement to impact the investment decision of the worker are, first, provide a **price rebate** contract, where the rebate is $\theta \in [0, 1)$ so the list price is $(1 - \theta)S$. Or second, a **low down payment (no discount) contract** (henceforth referred to as **low-DP** contract) $g \geq 0$ (reduces cash due today by g and adds g to financed principal).

Cost to the Worker. Hence, the down payment required is $D(\theta, g) = \delta(1 - \theta)S - g$. The financed principal becomes $P(\theta, g) = (1 - \delta)(1 - \theta)S + g$. And the present value interest charge becomes:

$$\iota \times P(\theta, g) = \iota[\phi(1 - \theta)S + g]$$

Since $D + P = (1 - \theta)S$, the total present value cost per unit becomes
 $c(\theta, g) = (1 - \theta)S + \iota[\phi(1 - \theta)S + g] = (1 - \theta)\bar{P} + \iota g$, where, $\bar{P} \equiv S[1 + \iota(1 - \delta)]$.

Worker solves. The worker's maximization problem is given by:

$$\max_{k \geq 0} \alpha f(k) - c(\theta, g)k \quad \text{s.t.} \quad D(\theta, g)k \leq L.$$

where, the down-payment she can offer is constrained by the liquidity L . In order to solve this, let's consider two cases. One where liquidity does not bind, and the other where liquidity does bind.

Case 1. Liquidity does not bind (interior solution). If the constraint is slack, the multiplier is $\lambda = 0$ and the interior solution satisfies:

$$\boxed{\alpha f'(k^*) = (1 - \theta)\bar{P} + \iota g, \quad \alpha f''(k^*) < 0}.$$

Prediction 1A Since $\partial c / \partial \theta = -\bar{P}$, the comparative static gives:

$$\frac{\partial k^*}{\partial \theta} = \frac{\bar{P}}{\alpha(-f''(k^*))} > 0.$$

This shows, that if the worker receives a price rebate contract, it should strictly increase investment.

Prediction 1B Because $\partial c / \partial g = \iota > 0$,

$$\frac{\partial k^*}{\partial g} = \frac{\iota}{\alpha f''(k^*)} < 0.$$

Low-DP contract should in fact reduce investment when liquidity does not bind (as it raises the present value marginal cost through ιg).

Corollary 1: Price rebates act like higher worker shares When liquidity is not binding (the down-payment constraint is slack), a price rebate θ is *equivalent* to raising the worker's retained claim α . Intuitively, a proportional discount scales down every rupee the worker must repay by the same factor $(1 - \theta)$, so the problem is identical to one with a larger claim on returns by worker. Formally, along the price menu (g fixed, in particular $g = 0$),

$$\boxed{k^*(\alpha, \theta; g = 0) = k^*\left(\frac{\alpha}{1 - \theta}, 0; g = 0\right), \quad \alpha^{\text{eff}} = \frac{\alpha}{1 - \theta}.$$

Case 2. When liquidity binds, $D(\theta, g)k = L$ binds, the choice is pinned by the cash cap:

$$\boxed{k^*(\theta, g) = \frac{L}{D(\theta, g)} = \frac{L}{\delta(1 - \theta)S - g}}.$$

Prediction 2A Since $D_\theta = -\delta S$,

$$\frac{\partial k^*}{\partial \theta} = -\frac{L}{D(\theta, g)^2} D_\theta = \frac{L \delta S}{D(\theta, g)^2} > 0.$$

A price rebate raises purchase because it reduces the cash due today by $\delta S \theta$.

Prediction 2B Since $D_g = -1$,

$$\frac{\partial k^*}{\partial g} = -\frac{L}{D(\theta, g)^2} D_g = \frac{L}{D(\theta, g)^2} > 0.$$

A low-DP contract raises purchase because it reduces the cash due today one-for-one. This is in contrast with the Case 1 when liquidity did not bind.

Corollary 2: Magnitude of investment response between two levers when liquidity binds. While the constraint binds,

$$\frac{\partial k^* / \partial g}{\partial k^* / \partial \theta} = \frac{1}{\delta S}.$$

A price rebate reduces the required down payment only by the *fraction* δ of the price cut (change in D equals $\delta S \Delta \theta$), whereas a low-DP contract g reduces the down payment *one-for-one* (change in D equals Δg). Therefore, while liquidity binds:

A change in g should increase k^* much more than an equal-sized change in θ

because g relaxes the cash hurdle directly, whereas θ only relaxes it through the down-payment share δ .³¹³²

Summary of the Model Predictions.

Worker's decision to invest under a price rebate contract vs. a low-DP contract is critically dependent on her liquidity for down payment.

1. If liquidity does not bind:

- (a) **Price rebate (θ) increases investment.** Along any interior solution, $\theta \uparrow$ lowers the lifetime unit cost and *raises* k^* .
- (b) **Down payment relief (g) reduces investment.** Holding liquidity slack, $g \uparrow$ increases the lifetime unit cost and *lowers* k^* .

2. When liquidity does bind:

- (a) **Price rebate increases purchase.** Because $D_\theta = -\delta S$, a rebate reduces the down payment by $\delta S \theta$ and hence *raises* k^* .

³¹For example, 10% rebate with $\delta = 0.15$ lowers the down payment by $0.15 \times S \times 0.10$, while the low-DP contract lowers it by g one-for-one; the latter is mechanically much larger.

³²That is, a one rupee reduction in downpayment via g vs. a one rupee reduction via a price rebate θ .

- (b) **Repayable low-DP increases purchase.** Because $D_g = -1$, g reduces the down payment one-for-one and thus *raises* k^* .

1.6.2 Testing Model Predictions: Discrete Choice Experiment

Guided by the model, I test take-up using an incentive compatible discrete choice experiment. Each respondent evaluated four financing contracts that correspond directly to the model’s two levers. The contracts were (also summarized in Table A10):

1. **Benchmark:** List price INR 100,000, down payment 15 percent (INR 15,000), the remaining INR 85,000 financed at the market rate over 24 months.
2. **Price rebate (θ):** The list price is reduced by INR 10,000 to INR 90,000, the down payment remains 15 percent of price which implies INR 13,500 due today, and INR 76,500 is financed. This primarily lowers lifetime user cost.
3. **Low-DP, no rebate (g):** The list price is held at INR 100,000, the down payment is reduced to INR 5,000, and INR 95,000 is financed. This relaxes the day-of-purchase liquidity requirement without changing lifetime price.
4. **Subsidy** The worker brings the standard down payment of INR 15,000, immediately receives INR 10,000 at purchase, and finances the usual 85 percent. This reproduces the field intervention and combines a price rebate with low-DP relief at the decision moment.

For the discrete choice experiment women were recruited from Lucknow who were active in the platform’s salon vertical. Lucknow was selected as a non-study city to avoid interference with the field experiment while keeping technology and demand comparable. Recruitment proceeded via in-app invitations sent to all eligible female workers on the Lucknow roster, following the same protocol as in the main study. The invitation was issued twice over a three-week window and asked workers who were considering a scooter purchase to contact the research team.

Workers were first invited to contact the research team if they were considering purchasing a scooter in the near future. The invitation explained that the firm sought to learn about preferences over alternative financing contracts. Enrolled respondents were then asked to evaluate four contracts. For each contract, enumerators explained the cash required today, the monthly installment, total interest, and the implied effective annual rate; respondents recorded a yes or no purchase decision for that contract. Contract order was randomized within respondent across the 24 possible permutations. To anchor choices in incentives, a post-survey lottery selected ten respondents; for each selected respondent one contract was drawn at random from the set just evaluated and implemented if the respondent had said yes to it. This procedure yields Becker–DeGroot–Marschak incentive compatibility and one hundred decisions per contract, for a total of 400 worker–contract observations.

Comparing acceptance across contracts identifies the elasticity to benchmark “control” contract, contract with reduced lifetime price holding liquidity nearly fixed, the elasticity to liquidity holding lifetime price fixed, and the response to the bundled policy that mirrors

the field intervention. The design therefore takes the model’s two levers, price rebate θ and low-DP (g), to the data in a way that isolates the operative constraint behind adoption.

1.6.3 Results: Discrete Choice Experiment

Table 1.6 Column 1 shows that the take-up under the benchmark (control) contract is about 16 percent, which mirrors the control-group purchase rate in the field experiment and supports external comparability. The price rebate contract lifts take-up to about 37 percent, a gain of 21 percentage points. Low DP (no discount) raises take-up to about 69 percent, a gain of 53 percentage points. The increase under low DP is about two and a half times the increase under the price rebate (53 pp vs. 21 pp), so the elasticity of adoption to relaxing the down-payment requirement is much larger than the elasticity to lowering lifetime price by a comparable amount. The subsidy arm that bundles both margins reaches about 95 percent (79 pp increase over 16% baseline) and is weakly preferred to either single lever. This ordering is exactly what the model predicts in the liquidity-binding case, where shifting cash due today generates a much larger movement in adoption than reducing the asset’s lifetime cost.

I next examine heterogeneity that speaks directly to the liquidity mechanism. Before the survey, each respondent stated whether she could mobilize the standard down payment of INR 15,000 within a week. This split yields two pre-specified groups that align with the model’s cases.

Workers who cannot mobilize INR 15,000 (liquidity constrained). Table 1.6 Column 2 shows that take-up under the benchmark is low, around 13 percent. Moving lifetime price with a price rebate produces only a small change that is not distinguishable from the benchmark contract. By contrast, relaxing the cash hurdle with low DP (no discount) raises take-up to roughly 70–75 percent, a large and precisely estimated jump. The subsidy that combines both margins pushes take-up to the top of the scale, near full adoption. Within this group, the effect of low DP is statistically and economically larger than the effect of the price rebate. This is exactly the pattern predicted by the model when the down-payment constraint binds.

Workers who can mobilize INR 15,000 (not liquidity constrained). Table 1.6 Column 3 shows that take-up in the control (benchmark) contract itself is higher - 23% - for this group relative to the liquidity constrained. The Low-DP (no discount) contract increases take-up by 0.39 over this benchmark, even though it raises lifetime cost by rolling the down payment into principal. This sign conflicts with the interior-case prediction (Prediction 1B), which, under slack liquidity, implies lower adoption than the benchmark. At the same time, Low-DP remains significantly below both price-reducing contracts (Price Rebate and Subsidy), consistent with the model’s ranking on lifetime price. The positive Low-DP effect within the non-liquidity-constrained group is consistent with a shadow value of liquidity even when the INR 15,000 constraint does not strictly bind. In particular, if agents value cash-on-hand for buffer-stock or risk-management reasons (e.g., income volatility early in tenure) or face non-pecuniary costs of assembling the upfront cash (time, hassle, and screening frictions), then the marginal utility of liquidity enters the first-order condition and makes a reduction

in attractive despite a higher lifetime price.³³

Among take up rates in the non-benchmark contracts, take-up is lower under low DP (no discount) than under the price rebate, consistent with the model’s prediction when liquidity does not bind and rolling cash into principal raises lifetime cost. The subsidy also yields high take-up but is statistically indistinguishable from the price rebate in this subgroup, while both exceed low DP. This pattern aligns with the interior-case prediction of the model - which should be the take-up patten among the non-liquidity constrained group.

Thus, the discrete choice experiment separates the price and liquidity margins in a way that maps directly to the model. Where liquidity binds, moving cash due today dominates a proportional cut in lifetime price. Where liquidity does not bind, the two levers perform similarly and the subsidy does not add much beyond either lever on its own. The evidence therefore points to the down-payment constraint as the operative margin for a large share of workers, while also indicating that for workers with adequate cash on hand either contract is sufficient to induce adoption.

1.7 Conclusion

In this paper, I examine the contribution of worker-owned capital to worker productivity and firm profitability. In a platform-firm setting in India, where salon services are delivered at customers’ homes and efficient mobility is a key production input, I evaluate an innovative subsidy program that incentivizes investment in “scooters,” a form of worker-owned capital. Half (50%) of treated workers take up the subsidy to purchase a scooter, priced at 20 weeks of their income, on credit. New adopters — workers who switch from unreliable private taxis or ride-hailing — experience large productivity gains, whereas incumbent users already operating a scooter for work exhibit no detectable effects. Using daily job-level data over more than eight months, intent-to-treat estimates show that new adopters increase jobs completed by 9% and weekly earnings by 13%, with no increase in labor supply, implying a 10% rise in weekly labor productivity. These workers disproportionately complete jobs that would otherwise be at risk of going unmatched in the system — specifically, farther-distance assignments and late-evening jobs — and they record higher work quality as measured by improved ratings. On the firm side, revenue and profit per worker-week rise in ways not explained by business stealing, indicating a clear increase in total surplus. Finally, I find no treatment-induced attrition, despite the platform context, eliminating the principal channel through which the firm would have been unable to recoup its subsidy.

Consequently, both the firm and workers realize large returns on investment. The firm recoups the subsidy within 18 weeks and, within the first year, earns a 153% return on investment per new adopter. On the worker side, the cost of capital is recovered within thirteen months (even absent the subsidy), and under minimal assumptions, the two-year return exceeds 73%. These magnitudes naturally raise the question of why the firm did not previously incentivize such investment. While the paper cannot adjudicate this

³³Not liquidity constrained’ is based on a coarse threshold (ability to mobilize INR 15,000). Given that all subjects are low-income platform workers, this partition is unlikely to carve out a group with zero shadow value of liquidity. Measurement error in self-reported liquidity and heterogeneity in cash-flow risk might blur the strict interior-case prediction.

experimentally, conversations with firm leadership indicate that the principal deterrent was the absence of credible evidence on the subsidy’s effectiveness — specifically, its returns for workers and for the firm, as well as implications for attrition. Following the intervention, the firm scaled the subsidy program to additional cities, providing a revealed-preference measure of the firm’s assessment of the program’s value.


On the worker side, underinvestment is likewise puzzling given sizable returns even without the subsidy. To probe mechanisms, I conduct an incentive-compatible discrete-choice experiment with a new sample of workers that orthogonally varies lifetime price and day-of-purchase liquidity via the required down payment, thereby disentangling a liquidity mechanism from a pure price reduction (which could operate through surplus sharing). I find that many workers are sharply liquidity constrained at the down-payment margin: take-up is substantially higher under a contract with a low down payment and a higher total price than under an alternative with a higher down payment but a lower total price. This pattern indicates that purchase-time liquidity can be a binding constraint for workers.

These findings matter for three reasons. First, the evidence indicates that liquidity constraints play a first-order role in labor market functioning — not only for firms and entrepreneurs, where most of the literature has focused, but for workers as well. In environments where workers are liquidity constrained, firm-financed or subsidized investment in productivity-enhancing, worker-owned capital can shift both workers and the firm onto a higher earnings path. Second, even in a platform setting with ostensibly low switching costs and ready outside options, the firm recovers its subsidy because attrition remains low; workers do not exit, allowing the firm to appropriate part of the surplus created by the investment. Third, the results open a clear agenda for future work: to identify the specific frictions that sustain retention in such settings and to assess external validity across contexts and populations. In particular, it is important to test whether similar retention and surplus sharing would arise in male-dominated segments of the platform economy, or whether the predominantly female workforce studied here plays a role in the observed low attrition.


Finally, analogous policies are proliferating in developed-country settings in which firms help workers overcome liquidity or credit constraints. In the United States, for example, ride-hailing platforms have partnered with vehicle providers (e.g., Uber–Hertz) to facilitate access to high-quality cars for work. These initiatives are explicitly aimed at improving retention and productivity. It is therefore natural to ask how such models perform in developing-country contexts, where liquidity constraints may bind more tightly, and to study which program designs most effectively translate employer support into persistent gains for both workers and firms.

Figure 1.1: Gender Differences in Actual and Preferred Mode of Commute
(Pre-Experiment Survey)

(a) Actual Mode of Commute




(b) Preferred Mode of Commute




Notes: Figure reports evidence from a pre-experiment phone survey of N=1,531 randomly selected platform workers performing the same salon-service job. Panel A shows actual primary commute mode; Panel B shows stated preference. “Own 2-wheeler” denotes using a personally owned scooter; “3-/2-wheeler taxis” denotes private bike-taxi or three-wheelers (auto-rickshaws). In Panel A, 77% of men versus 31% of women report commuting by own scooter (the remainder use taxis: 23% and 69%, respectively). In Panel B, 92% of men and 79% of women state they would prefer to commute by own scooter rather than taxis. The sizeable usage–preference gap for women contrasted with near-saturation among men motivates focusing the intervention on women and is consistent with frictions of ride-hailing highlighted in the baseline (limited availability at peak times, higher/volatile off-hour fares, refusals to carry equipment, safety concerns, and search/wait time).

Figure 1.2: Correlation Between Baseline Mode of Commute, Productivity and Travel Expenditure

(a) Correlation: Baseline Mode of Commute and Productivity




(b) Correlation: Baseline Mode of Commute and Travel Expenditure for Work




Notes: This figure compares outcomes by baseline commuting mode using pre-experiment administrative data for women drawn from a phone survey of N=1,531 randomly selected platform workers (the survey was 50% women). “Scooter user” denotes workers who primarily commute on a personally owned two-wheeler. “2/3-wheeler taxi user” denotes reliance on private motorcycle taxis (two-wheelers) or auto-rickshaws (three-wheelers). Each panel plots histograms of standardized variables (mean = 0, s.d. = 1). Panel A displays productivity, measured by daily earnings. Panel B displays daily travel expenditure for work. At baseline, scooter users earn more per day and spend less on transport than taxi users. These are descriptive correlations (unadjusted) and should not be interpreted causally.

Figure 1.3: Correlation Between Job Value and Job Distance and Timing

(a) Correlation: Job Value and Job Distance




(b) Correlation: Job Value and Job Timing



Notes: Pre-experiment administrative job-level data from the study platform. Panel (a) plots mean booking value (INR) against straight-line distance from the worker's registered home location; booking value increases with distance. Panel (b) plots mean value per minute—booking value divided by service-time minutes—by booking hour; earnings per minute are higher for early-morning (7–11 h) and late-evening (17–20 h) bookings than for mid-day (12–16 h). These are unconditional correlations (non-causal). They motivate analyzing scooter adoption as a technology that relaxes travel-time and travel-variance constraints precisely where the returns to mobility — longer distances and off-peak hours — are higher.


Figure 1.4: Correlation Between Job Completion Rate and Job Start Time




Notes: This figure uses pre-experiment administrative, job-level data from the study platform to show how fulfillment varies over the day. The outcome is the job completion rate, defined as the share of booked jobs that are marked completed in the system on the scheduled day. The x-axis is the booking (scheduled start) hour; the line plots the mean completion rate by hour. The overall mean completion rate in the baseline data is 0.84. Completion is highest around mid-day and is noticeably lower at the edges of the working day. In particular, completion rates in the early window (7-11 h) and the late window (17-19 h) are below those in the mid-day window (12-16 h). These are unconditional, descriptive patterns (non-causal). They are consistent with greater travel and punctuality frictions for early/late and farther jobs, and they motivate analyzing heterogeneous treatment effects by booking hour and distance when studying scooter adoption.

Figure 1.5: Impact on Job Outcomes Over Time for New Adopters

(a) Impact on Jobs Over Time for New Adopters




(b) Impact on Earnings Over Time for New Adopters




Notes: This figure plots the event-time difference-in-differences estimates for new adopters in the experimental sample. The horizontal axis reports months relative to scooter subsidy offer, with 0 denoting the month of offer. Estimates are obtained from weekly panel regressions that include cohort fixed effects and calendar-week fixed effects. Event-time indicators are interacted with the treatment group; for controls these indicators are zero in all weeks. The indicator for month -1 (the last pre-adoption month) is omitted, so each coefficient measures the average outcome at event time k relative to month -1 , net of cohort and week effects. Weekly coefficients are binned to monthly event times for display. Squares plot pre-trend coefficients for $k \in (-6, \dots, -2)$; circles plot post-adoption coefficients for $k \in (0, \dots, +8)$. Vertical bars show 95 percent confidence intervals with standard errors clustered at the worker level. Panel (a) reports weekly jobs; Panel (b) reports weekly earnings (INR). Pre-adoption estimates are near zero, indicating no detectable pre-trends relative to month -1 . Beginning in month 0 — when exposure is partial by construction — and persisting thereafter, both weekly jobs and weekly earnings are higher than in the last untreated month.

Figure 1.6: Discrete Choice Incentive Compatible Experiment: Contract Take-Up Rates

(a) Full Sample



(b) Sample Split by Liquidity Constraint for Down Payment



Notes: This figure shows take-up in an incentive-compatible discrete-choice experiment where workers chose among four scooter-finance contracts that differentially relax the day-of-purchase cash requirement or reduce lifetime price. “Control” is the benchmark loan: list price INR 100,000 with a 15% down payment (INR 15,000) and the remaining INR 85,000 financed over 24 months at the market rate. “Price Rebate” lowers the list price to INR 90,000 while keeping the 15% down payment (INR 13,500) and financing INR 76,500; this primarily reduces lifetime user cost without easing upfront liquidity. “Low-Down Payment (DP), no rebate” keeps the list price at INR 100,000 but reduces the down payment to INR 5,000, financing INR 95,000; this relaxes the day-of-purchase liquidity requirement while holding lifetime price fixed. “Subsidy” replicates the field intervention: the worker brings the standard INR 15,000 down payment, receives INR 10,000 at purchase, and finances the usual 85%; this combines a net price reduction with immediate relief of the upfront cash requirement. Bars report the percent choosing to adopt under each contract; whiskers are 95% confidence intervals. Panel (a) shows the full sample. Panel (b) splits the sample by whether the worker can assemble INR 15,000 at purchase (a measure of liquidity constraint).

Table 1.1: Impact on Take-Up of Subsidy and Scooter Purchase

	(1)	(2)	(3)	(4)
	Subsidy	Scooter	Subsidy	Scooter
Treatment	0.50*** (0.04)	0.32*** (0.05)	0.47*** (0.04)	0.30*** (0.06)
Treatment $\times$ Incumbent Users			0.09 (0.08)	0.05 (0.11)
Control Mean	0	0.17	0	0.17
N: Workers	350	350	350	350

Notes: This table reports the first-stage effects of the randomized offer on two binary outcomes: (i) **Subsidy**: whether the worker accepts the financing subsidy contract; and (ii) **Scooter**: whether the worker purchased a scooter during the study window. “Treatment” is an indicator for being randomly offered the subsidy. “Incumbent User” is an indicator for using a scooter at baseline for completing jobs on the platform. Columns (1)–(2) estimate intent-to-treat effects in the pooled sample. Columns (3)–(4) add an interaction with “Incumbent User” to allow the effect of the offer to differ for incumbents versus new adopters; in those columns, the coefficient on *Treatment* is the effect for baseline non-users, while *Treatment* $\times$ *Incumbent User* is the difference relative to baseline users of scooters. Outcomes are indicators. “Control Mean” reports the mean of the outcome in the control group (no offer). “N: Workers” is the number of workers in the specification. All regressions include randomization-cohort (version) fixed effects; standard errors are clustered at the worker-level. Significance: * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

Table 1.2: Impact on Jobs, Earnings and Productivity

	(1) Jobs (weekly)	(2) Earnings (weekly)	(3) Hours-online (weekly)	(4) Earnings- by-hours-online
Panel A: New Adopters				
Treatment Effect (ITT)	0.54** (0.23)	602.40*** (204.41)	1.29 (1.28)	9.82*** (3.59)
Control Mean	6.12	4,757.98	53.08	94.05
N: Worker-Weeks	8,167	8,167	8,167	8,167
Panel B: Incumbent-Users				
Treatment Effect (ITT)	-0.34 (0.34)	-89.14 (289.32)	-2.51 (1.99)	-0.92 (5.50)
Control Mean	7.24	5,799.22	55.97	106.64
N: Worker-Weeks	3,570	3,570	3,570	3,570
Panel C: Pooled Sample				
Treatment Effect (ITT)	0.18 (0.19)	322.23* (170.48)	0.13 (1.01)	5.37* (3.21)
Control Mean	6.49	5,097.52	54.03	98.19
N: Worker-Weeks	11,737	11,737	11,737	11,737

Notes: This table presents intent-to-treat (ITT) estimates of the effect of being offered a scooter subsidy on weekly outcomes constructed from platform administrative records. The unit of observation is the worker-week; daily records are aggregated to the weekly level over the eight months following randomization. Columns report four outcomes: (1) Jobs - the count of services completed in the week; (2) Earnings - total platform earnings in Indian rupees (INR) in the week; (3) Hours-online (Labor Supply) - the total number of hours the worker is “online,” defined as all time windows in which the worker is in the active matching pool (waiting for jobs, traveling to clients after being matched, or performing jobs); and (4) Earnings-by-hours-online - weekly earnings divided by hours-online, a measure of productivity. Regression follows estimation strategy described in Equation 1.1 that includes cohort (randomization batch) fixed effects and calendar-week fixed effects. Standard errors, in parentheses, are clustered at the worker level. “Control Mean” gives the mean of the outcome for control observations in the corresponding panel. “N: Worker-Weeks” is the total number of worker-week observations used in that specification. Currency is INR. Stars denote p-values (* p<0.10, ** p<0.05, *** p<0.01). Panel A restricts the sample to new adopters (70% of the sample), defined as workers who did not commute by their own scooter at baseline (the main analysis for the intervention). Panel B restricts to Incumbent-Users, defined as workers who were already commuting by their own scooter at baseline; this stratum serves as a placebo group for productivity effects. Panel C pools both strata to report intent-to-treat effects in the combined experimental sample. Randomization was stratified by these baseline adoption strata.

Table 1.3: Impact on Labor Supply Reallocation Across Space, Time, and Job Quality on New Adopters

	Job Selection Effects						Job Quality Metrics	
	(1) Total KMs	(2) Near (0–10 km)	(3) Far (10–25 km)	(4) Mid-day (12–16)	(5) Early/Late (<11 or >17)	(6) Surge Earnings	(7) Job Rating	(8) Late to Job
Treatment effect	6.67*** (2.55)	-0.29 (0.24)	0.54** (0.23)	0.21 (0.19)	0.37* (0.21)	18.67*** (6.66)	0.03* (0.01)	0.18 (0.16)
<i>p</i> -value (Near = Far)	[0.005]							
<i>p</i> -value (Late = Afternoon)					[0.195]			
Control mean	58.40	3.49	2.68	3.07	3.10	43.97	4.83	2.49
N	8,167	8,167	8,167	8,167	8,167	8,167	5,695	8,061

Notes: This table defines and estimates treatment effects on margins of job selection (space and time) and on job quality for the ****new-adopter**** stratum only. The unit of observation is the worker–week; outcomes are built from administrative logs and aggregated to weekly frequency.

- Job Selection Effects (Cols. 1–6). “Total KMs” is the total kilometers associated with all completed jobs in the week, computed using the worker’s registered home location and the client location for each job. “Near (0–10 km)” and “Far (10–25 km)” are counts of completed jobs whose distance from home falls in those bins. The residual >25 km bin increases modestly, not shown here. “Mid-day (12–16)” and “Early/Late (<11 or >17)” are counts of completed jobs whose scheduled start falls in those booking-hour windows. “Surge Earnings” is the extra rupee amount of weekly earnings attributed to surge pricing premia recorded by the platform.
- Job Quality Metrics (Cols. 7–8). “Job Rating” is the client rating of the service (platform 1–5 scale), averaged across the worker’s completed jobs in the week. “Late to Job” is the incidents of completed jobs in the week for which the platform marks the worker as late to the scheduled start time.
- Rows labeled “*p*-value (Near = Far)” and “*p*-value (Late = Afternoon)” report Wald tests of equality between the corresponding treatment effects in columns (2) vs. (3) and (5) vs. (4), respectively.
- Regression follows the estimation strategy in Equation 1.1 and includes cohort (randomization batch) fixed effects and calendar-week fixed effects. Standard errors, in parentheses, are clustered at the worker level. “Control Mean” is the mean outcome for controls in this stratum. “N: Worker–Weeks” is the number of worker–week observations. Currency is INR. Stars denote *p*-values (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$).

Table 1.4: Impact on Firm Revenue, Profits and Worker Attrition

	Firm Earnings			Worker Attrition
	(1) Revenue Total	(2) Revenue Net-of-Worker-Payout	(3) Profits	(4) Inactive Post-1-Year
Panel A: New Adopters				
Treatment Effect (ITT)	1,178.95*** (404.54)	575.72*** (202.63)	265.26*** (91.02)	-0.040 (0.045)
Control Mean	9,369.21	4,611.23	2,108.07	0.109
N:Worker-Weeks	8,167	8,167	8,167	245
Panel B: Incumbent-Users				
Treatment Effect (ITT)	-154.35 (572.98)	-77.38 (287.75)	-34.73 (128.92)	-0.015 (0.037)
Control Mean	11,219.38	5,420.16	2,524.36	0.022
N:Worker-Weeks	3,570	3,570	3,570	105
Panel C: Pooled Sample				
Treatment Effect (ITT)	681.59** (330.25)	336.30** (165.15)	153.36** (74.31)	-0.036 (0.031)
Control Mean	9,972.53	4,875.01	2,243.82	0.085
N:Worker-Weeks	11,737	11,737	11,737	350

Notes: This table reports intent-to-treat (ITT) effects of being offered the scooter subsidy on firm-side outcomes constructed from platform administrative records. Panels A–C mirror the experimental strata: Panel A restricts to new adopters (baseline non-users of a private scooter), Panel B to incumbent users (baseline users), and Panel C pools both strata. The unit of observation in columns (1)–(3) is the worker–week over the eight months after randomization. Column (4) is measured at the worker level one year after randomization; the sample size in that column is the number of workers observed at that horizon. Outcomes: Column (1) “Revenue—Total” is the platform’s gross booking value attributed to the worker’s completed jobs in the week (INR). Column (2) “Revenue—Net-of-Worker-Payout” deducts the worker’s payout for those jobs from gross revenue, i.e., the platform’s net revenue before other costs. Column (3) “Profits” is the platform’s accounting profit for those jobs as recorded in the administrative data (net revenue after subtracting recorded variable costs such as refunds/adjustments and platform incentives). Column (4) “Inactive Post-1-Year” is an indicator equal to 1 if the worker is inactive on the platform at the one-year mark and 0 otherwise. Regression follows the estimation strategy described in Equation 1.1: cohort (randomization-batch) fixed effects and calendar-week fixed effects are included. Standard errors, in parentheses, are clustered at the worker level. “Control Mean” is the mean of the outcome for control observations in the corresponding panel. “N: Worker–Weeks” is the number of worker–week observations in columns (1)–(3); column (4) reports the number of workers. Currency is INR. Stars denote p-values (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$).

Table 1.5: Return on Investment (ROI) for the Firm and the Workers

	Treatment Effect IV	Recovery Period	ROI 12 months	ROI 24 months
Panel A: Firm’s return to subsidy				
New Adopters	557.87*** (197.29)	18 weeks	153%	345%
Pooled Sample	345.95*** (151.17)	29 weeks	60%	180%
Panel B: Workers’ return to scooter purchase				
New Adopters	2,018*** (768.03)	54 weeks	-2.8%	73%

Notes: This table reports return-on-investment (ROI) measures for the platform (Panel A) and for workers (Panel B) using administrative outcomes. Causal effects are estimated by instrumental variables, with the randomized subsidy offer as the instrument. All payoff calculations are *discounted present values* computed at a weekly rate.

- Horizon and survival: For months 1–12 I cumulate discounted weekly increments from weeks $t = 1$ to 52. Weeks 1–35 are realized in the data; weeks 36–52 are multiplied by the 12-month survival probability s . (See Section 2.B for formulas and parameter values.)
- Panel A (firm).
 - Column 1, “Treatment Effect (IV)”: IV estimate of the weekly increment in the platform’s profit per worker.
 - Column 2, “Recovery Period”: *discounted payback time*—the smallest week T such that $\sum_{t=1}^T \beta^t \times (\text{IV weekly profit}) \geq C$, where C is the subsidy cost per adopting worker.
 - Columns 3–4, “ROI 12 months / 24 months”: percentage return over the stated horizon with PV computed under the horizon and survival conventions above.
- Panel B (workers).
 - Column 1, “Treatment Effect (IV)”: IV estimate of the weekly increment in worker earnings, instrumenting actual purchase with the subsidy offer.
 - Column 2, “Recovery Period”: *discounted payback time*—the smallest week T such that $\sum_{t=1}^T \beta^t \times (\text{IV weekly earnings}) \geq \text{PV}(\text{financed purchase cost})$. The financed cost is the *present value of the loan payment stream* (principal and interest, net of subsidy), as detailed in Appendix 1.B.
 - Columns 3–4, “ROI 12 months / 24 months”: percentage returns computed as $(\text{PV of incremental weekly earnings} - \text{PV}(\text{financed cost})) / \text{PV}(\text{financed cost})$ over the stated horizons, under the same discounting and survival conventions.
 - The negative 12-month ROI for workers reflects a discounted payback of roughly 54 weeks; returns turn positive only once the present value of incremental earnings exceeds the present value of loan obligations.
- Currency is INR; time is in calendar weeks. Sample restricted to **new adopters** because impacts for incumbent users are indistinguishable from zero on platform jobs (Table 1.2), so pooling would attenuate policy-relevant returns to adoption. Stars denote p-values (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$). All formulas, parameter choices (r, s, δ, C), and robustness alternatives are provided in Appendix 1.B.

Table 1.6: Discrete Choice Incentive Compatible Mechanism Experiment

	Full Sample	Liquidity Constrained	Not Liquidity Constrained
	(1)	(2)	(3)
Price Rebate	0.210*** (0.059)	0.058 (0.061)	0.548*** (0.106)
Low-DP Contract	0.530*** (0.060)	0.594*** (0.069)	0.387*** (0.124)
Subsidy	0.790*** (0.044)	0.841*** (0.045)	0.677*** (0.105)
Control Mean	0.16	0.13	0.23
p(Price Rebate = Low-DP)	0.00	0.00	0.14
p(Price Rebate = Subsidy)	0.00	0.00	0.13
p(Low-DP = Subsidy)	0.00	0.00	0.01
N: Worker-Contract	400	276	124

Notes: This table reports results from the discrete-choice experiment. The dependent variable equals one if the worker states she would accept the displayed contract. Regressors are indicators for contract menus. “Control” is the benchmark loan: list price INR 100,000 with a 15% down payment (INR 15,000) and the remaining INR 85,000 financed over 24 months at the market rate. “Price Rebate” lowers the list price to INR 90,000 while keeping the 15% down payment (INR 13,500) and financing INR 76,500; this primarily reduces lifetime user cost without easing upfront liquidity. “Low-Down Payment (DP), no rebate” keeps the list price at INR 100,000 but reduces the down payment to INR 5,000, financing INR 95,000; this relaxes the day-of-purchase liquidity requirement while holding lifetime price fixed. “Subsidy” replicates the field intervention: the worker brings the standard INR 15,000 down payment, receives INR 10,000 at purchase, and finances the usual 85%; this combines a net price reduction with immediate relief of the upfront cash requirement. Each worker evaluated the four contracts, yielding $100 \times 4 = 400$ worker-contract observations. Column 1 reports the results for full sample. Column 2 restricts the sample to workers who report being unable to assemble INR 15,000 at the time of purchase, which we use as an ex-ante indicator of a binding purchase-time liquidity constraint (INR 15,000 corresponds to the benchmark down payment). Column 3 restricts the sample to workers who report being able to assemble INR 15,000 at the time of purchase, i.e., workers without a binding purchase-time liquidity constraint under the benchmark down payment. All specifications include version (order) fixed effects. SEs are clustered at the worker level. Significant at * 10%, ** 5% and ***1% level.

Appendix 1.A Appendix A: Additional Results


Figure A1: Example of Intervention Scooters




Notes: This figure illustrates the type of lightweight, step-through, gear-less two-wheeler scooters, often referred to as “scooty”, that are prevalent in India. **Unlike the stand-up, dockless “kick-scooters”** popularized by companies like Lime or Bird in North America, these Indian scooters are ridden seated, licensed for road traffic, and designed for trips well beyond first- and last-mile range (more like the popular Italian brand scooters ‘Vespas’). These vehicles, typically 90-125 cc petrol or electric models such as the Honda Activa and TVS Scooty, feature a step-through frame, continuously-variable transmission, and low kerb weight (e.g., the 93 kg TVS Scooty Pep Plus), making them particularly appealing to novice and female riders due to reduced physical exertion.

Figure A2: Correlation Between Wealth and Scooter Use and Productivity at Baseline

(a) Correlation Between Wealth and Baseline Scooter Use




(b) Correlation Between Wealth and Baseline Productivity




Notes: This figure documents baseline correlations between wealth and mobility/productivity measures. Wealth is measured at baseline and used to partition workers into wealth groups. Panel (a) plots the share of workers who report commuting by their own scooter at baseline by wealth quintile (poorest to richest). Percent labels indicate the fraction using an own scooter within each quintile. Panel (b) presents box-and-whisker plots comparing outcomes for workers above vs. below the median of the baseline wealth distribution. Outcomes are: daily earnings (INR/day), monthly travel expense for work (INR/month), daily labor supply (hours online per day), and daily services completed (jobs/day). For each outcome, the box shows the interquartile range with the median line; whiskers extend to $1.5 \times$ the interquartile range (points beyond are plotted as outliers, if any). These are descriptive correlations at baseline and should not be interpreted causally.

Figure A3: Descriptive Facts: Daily Calendar Availability and Job Completion Rate for Workers




Notes: This figure uses administrative logs to describe scheduling marks on the worker calendar and average daily activity; all statistics are descriptive. Panel (a) shows, for each status label, the **average number of days per week per worker that a given day is marked with that label in the platform calendar: Working (any day), Working >2h (platform), Working full day (>8h), Working partial day (<8h), Leave, Temp blocked, Unmarked, Working <2h (platform), and Missing. Points are sample means; whiskers are 95% confidence intervals. Panel (b) shows, on an average worker-day, the number of jobs the platform sends (offers) versus the number of jobs completed, separately for weekdays (Mon–Wed, dashed) and weekend/high-demand days (Thu–Sun, solid). “Sent per available day (UB)” counts offers on days the worker is marked available; “Done per available day” counts completed jobs on those days. “Sent per high-demand day (UB)” and “Done per high-demand day” restrict to high-demand days. Whiskers indicate the interquartile range (p25–p75).

Figure A4: Worker's Credit Score (CIBIL) Distribution



Notes: This figure shows the empirical cumulative distribution function (CDF) of workers' baseline credit-bureau scores (CIBIL). The x-axis reports the credit score on the CIBIL scale; the y-axis shows the fraction of workers with a score less than or equal to each value. Scores are measured prior to the experiment. The vertical dashed line marks a score of 700 for reference.


Figure A5: Take-up Rate Over Time




Notes: This figure plots the cumulative share of workers who purchase a scooter over event time. Event time is measured in months since each worker's cohort randomization; month 0 is the randomization month. Curves are shown separately by treatment status and by baseline riding status (driver = reported commuting by own scooter at baseline; non-driver otherwise). Solid lines denote non-drivers; dashed lines denote drivers.

Figure A6: Impact on Job Outcomes Over Time for Incumbent-Users

(a) Impact on Jobs Over Time for Incumbent Users




(b) Impact on Earnings Over Time for Incumbent Users



Notes: This figure plots the event-time difference-in-differences estimates for incumbent users in the experimental sample. The horizontal axis reports months relative to scooter subsidy offer, with 0 denoting the month of offer. Estimates are obtained from weekly panel regressions that include cohort fixed effects and calendar-week fixed effects. Event-time indicators are interacted with the treatment group; for controls these indicators are zero in all weeks. The indicator for month -1 (the last pre-adoption month) is omitted, so each coefficient measures the average outcome at event time k relative to month -1 , net of cohort and week effects. Weekly coefficients are binned to monthly event times for display. Squares plot pre-trend coefficients for $k \in (-6, \dots, -2)$; circles plot post-adoption coefficients for $k \in (0, \dots, +8)$. Vertical bars show 95 percent confidence intervals with standard errors clustered at the worker level. Panel (a) reports weekly jobs; Panel (b) reports weekly earnings (INR). Pre-adoption estimates are near zero, indicating no detectable pre-trends relative to month 1. Beginning in month 0—when exposure is partial by construction—and persisting thereafter, both weekly jobs and weekly earnings are higher than in the last untreated month.


Figure A7: Impact on Jobs of Control Workers Over Time




Notes: This figure reports a placebo event-study using only *control* workers. Outcome: weekly jobs. Event time k is measured in weeks relative to each control worker's cohort week 0, defined as the calendar week in which that cohort was assigned treatment in the experiment (month 0 in the main design). The omitted category is $k = -1$ (the last pre-period). Regressions include cohort (randomization-batch) fixed effects and calendar-week fixed effects. Points plot the event-time coefficients; vertical bars are 95% confidence intervals (standard errors clustered at the worker level).

Figure A8: Baseline Beliefs About Jobs and Revenue Increase Conditional on Scooter Adoption

(a) Baseline Beliefs: Jobs




(b) Baseline Beliefs: Earnings




Notes: This figure summarizes baseline beliefs about gains from adopting a scooter, elicited prior to the experiment. Panel (a) plots the distribution of workers' expected monthly increase in jobs, defined as believed total jobs with a scooter minus self-reported baseline jobs (count/month). Bars show the empirical histogram (y-axis = fraction of respondents). The vertical dashed lines mark the sample mean of believed gains (blue) and the realized experimental gain used for reference (red). Panel (b) plots the distribution of workers' expected monthly increase in revenue, defined analogously as believed revenue with a scooter minus self-reported baseline revenue (INR/month). Bars show the histogram (bin width = INR 500) with a kernel density overlaid; the vertical dashed lines again mark the mean believed gain (blue) and the realized experimental gain used for reference (red).

Figure A9: Heterogeneity in Treatment Effects by Beliefs

(a) Heterogeneity in Impact on Jobs by Beliefs




(b) Heterogeneity in Impact on Earnings by Beliefs




Notes: This figure shows heterogeneity in treatment effects by baseline beliefs about gains from adopting a scooter. Beliefs are elicited pre-experiment and standardized to a z-score (mean = 0, s.d. = 1); higher values indicate larger expected gains. For each outcome, treatment effects are estimated using Equation 1.1 with cohort (randomization-batch) fixed effects and calendar-week fixed effects. Panel (a) reports weekly jobs; Panel (b) reports weekly earnings (INR).

Figure A10: Heterogeneity in Treatment Effects by Wealth

(a) Heterogeneity in Impact on Jobs by Wealth



(b) Heterogeneity in Impact on Earnings by Wealth



Notes: This figure examines heterogeneity in treatment effects by baseline wealth. Wealth is measured at baseline and partitioned into quartiles (1 = poorest, 4 = richest). For each quartile, treatment effects are estimated using Equation 1.1 with cohort (randomization-batch) fixed effects and calendar-week fixed effects. Points plot the estimated treatment effect within each wealth quartile (from interactions of the treatment indicator with quartile dummies). Vertical bars are 95 percent confidence intervals with standard errors clustered at the worker level. Panel (a) reports weekly jobs; Panel (b) reports weekly earnings (INR).

Table A1: Baseline Balance by CIBIL (Credit) Score Among (All) Applicants

Outcome	Coef	p-val
Applied / tried for scooter loan	-0.056	0.208
Children in HH	-0.071	0.056
Earnings (last month)	-1.4e+03	0.632
Education (years)	-0.030	0.831
Got scooter loan (baseline)	-0.092	0.226
Jobs (last month)	-0.304	0.859
Married	-0.047	0.568
Travel expenditure (weekly)	-558.566	0.179

Notes: Notes: This table shows baseline balance in observables among workers who applied to the subsidy program. The comparison is by credit-score group defined at the CIBIL threshold of 700: ($N < 700$) (“lower score”) versus ($N \geq 700$) (“higher score”), where (N) denotes the CIBIL score. For each outcome, the “Coef” column reports the difference in means from an OLS regression of the baseline variable on an indicator for ($N < 700$). Thus, a negative coefficient indicates a lower average for the ($N < 700$) group. The “(p)-val” column reports the two-sided (p)-value for the null of equal means across the two groups. Standard errors are clustered at worker level. All variables are measured prior to treatment using baseline data. Variable definitions: “Applied / tried for scooter loan” indicates any prior application attempt; “Children in HH” is the number of children residing in the household; “Earnings (last month)” is total labor income in the month before baseline; “Education (years)” is completed schooling; “Got scooter loan (baseline)” indicates pre-existing vehicle finance; “Jobs (last month)” counts platform jobs completed in the month before baseline; “Married” is marital status at baseline; and “Travel expenditure (weekly)” is total transport spending during the baseline week. Monetary values are INR. Overall, the two score groups appear similar on observables: estimated differences are small in magnitude and generally not statistically significant.

Table A2: Baseline Balance

	Pooled Sample			New Adopters			Incumbent Users		
	Control	Treat	P-val (Diff)	Control	Treat	P-val (Diff)	Control	Treat	P-val (Diff)
Panel A: Baseline Survey Data									
Age	34.97	34.90	0.961	34.40	34.50	0.831	36.17	35.95	0.739
Married	0.75	0.75	0.501	0.75	0.72	0.949	0.73	0.85	0.173
Have Kids	0.90	0.89	0.944	0.84	0.88	0.379	1.00	0.93	0.088
Use 3-Wheelers/Bike-Taxis	0.55	0.57	1.000	0.81	0.78	.	0.00	0.00	.
Drives own scooter	0.32	0.27	0.236	0.00	0.00	.	1.00	1.00	.
Knows scooty driving	0.89	0.86	0.339	0.85	0.81	0.378	0.98	1.00	0.348
Total HH mems	3.83	4.00	0.418	3.72	4.01	0.145	4.04	3.98	0.694
Num HH know driving	2.26	2.25	0.605	2.11	2.17	0.856	2.58	2.47	0.380
Total earning members	1.96	2.08	0.235	1.97	2.08	0.460	1.94	2.08	0.446
Wealth index	-0.00	0.06	0.885	-0.36	-0.17	0.498	0.76	0.66	0.718
Husband unemployed	0.08	0.07	0.991	0.10	0.07	0.806	0.04	0.05	0.790
Service subcategory	4.40	4.26	0.211	4.42	4.26	0.080	4.34	4.25	0.781
Travel expenditure	5,359	5,673	0.574	6,223	6,578	0.334	3,542	3,350	0.372
Panel B: Baseline Admin Data									
Platform Experience	3.36	3.15	0.294	3.02	3.15	0.815	4.06	3.17	0.037
Daily working hours	7.92	7.93	0.675	7.69	7.90	0.318	8.40	7.99	0.529
Weekly working hours	51.14	50.93	0.991	49.26	50.62	0.567	54.99	51.72	0.395
Daily services delivered	1.01	1.01	0.790	1.00	1.01	0.997	1.05	1.03	0.614
Weekly services delivered	6.67	6.69	0.818	6.53	6.64	0.969	6.95	6.80	0.584
Daily app earnings	792	774	0.672	763	766	0.611	852	794	0.206
Weekly app earnings	5,212	5,108	0.688	5,009	5,058	0.587	5,628	5,235	0.185
Daily rating	4.84	4.84	0.766	4.84	4.83	0.939	4.83	4.86	0.170

Notes: This table shows balance at baseline. Columns (1)–(3) report statistics pooled sample; columns (4)–(6) new adopters; columns (7)–(9) for the incumbent users. Within each triplet, the first two columns report Control and Treatment means, respectively; the third column reports the two-sided p -value from a test of equality of means between Treatment and Control in that subgroup. Panel A variables come from the baseline survey. Panel B variables come from platform administrative records measured at baseline; daily measures are averaged to the worker level and weekly measures are totals over the baseline week. “Daily rating” is on the platform’s 1–5 scale. Among incumbent users, baseline “Platform Experience” is modestly imbalanced (control 4.06 years vs. treatment 3.17 years, $p=0.037$); all incumbent placebo results are unchanged when we control for baseline experience. Monetary variables are INR. The wealth index is constructed from baseline assets (first principal component).

Table A3: Take Up Characteristics

	New Adopters		Incumbent-Users		Pooled Sample	
	(1)	(2)	(3)	(4)	(5)	(6)
	Treatment	Treatment $\times$ Cov	Treatment	Treatment $\times$ Cov	Treatment	Treatment $\times$ Cov
Panel A: Survey Covariates						
Treatment (no covariate)	0.317***		0.355***		0.319***	
Drives own scooter					0.305***	0.038
Age	0.350***	0.138**	0.341***	0.153	0.338***	0.137**
Married	0.316***	0.039	0.360***	-0.013	0.319***	0.018
Have Kids	0.320***	0.069	0.371***	0.000	0.326***	0.083
Use 3-Wheelers/Bike-Taxis	0.357***	-0.072	0.355***	0.000	0.319***	-0.029
Knows scooty driving	0.313***	-0.004	0.352***	0.000	0.318***	-0.005
Total HH mems	0.313***	0.161***	0.362***	-0.131	0.317***	0.103**
Num HH know driving	0.338***	0.124*	0.340***	0.070	0.324***	0.093*
Total earning members	0.313***	0.078	0.334***	-0.203**	0.313***	-0.009
Wealth index	0.340***	0.099	0.340***	0.044	0.322***	0.064
Husband unemployed	0.318***	0.034	0.336***	-0.007	0.318***	0.055
Service subcategory	0.322***	-0.039	0.358***	-0.009	0.321***	-0.014
Travel expenditure	0.302***	0.033	0.411***	0.099	0.321***	0.036
Panel B: Admin Covariates						
Platform Experience	0.322***	0.087	0.334***	-0.051	0.311***	0.044
Daily working hours	0.309***	-0.090	0.367***	-0.141*	0.315***	-0.083
Weekly working hours	0.309***	-0.077	0.375***	-0.170**	0.314***	-0.080
Daily services delivered	0.314***	-0.061	0.372***	-0.233**	0.316***	-0.073
Weekly services delivered	0.315***	-0.053	0.373***	-0.237**	0.316***	-0.068
Daily app earnings	0.312***	-0.054	0.365***	-0.068	0.318***	-0.015
Weekly app earnings	0.313***	-0.047	0.365***	-0.073	0.318***	-0.011
Daily rating	0.322***	0.036	0.333***	-0.015	0.317***	0.044

Notes: The dependent variable is an indicator for scooter take-up during the study window (intent-to-treat; equals 1 if the worker purchases a scooter on credit by the end of follow-up). Each row reports coefficients from a model of the form where X_i is the baseline covariate named in the row. The coefficient reported under the “Treatment” column is the impact of Treatment on take-up, the coefficient under “Treatment $\times$ Cov” is heterogeneity in the treatment effect with respect to various covariates. The row “Treatment (no covariate)” reports the take-up effect with no interaction. Columns (1)–(2) restrict to new adopters (baseline non-drivers), Columns (3)–(4) to incumbent-users (baseline drivers), and Columns (5)–(6) to the pooled sample. Panel A covariates come from the baseline survey; Panel B covariates come from baseline administrative records (daily measures averaged, weekly measures summed, at the worker level). All specifications include cohort-stratum fixed effects (randomization strata). SEs are clustered at worker level. Significant at * 10%, ** 5% and ***1% level.

Table A4: Robustness: Impact on Jobs and Earnings of New Adopters

	New Adopters					
	(1) Jobs	(2) Earnings	(3) Jobs	(4) Earnings	(5) Jobs	(6) Earnings
Treatment effect (ITT)	0.54** (0.23)	602.40*** (204.41)	0.56** (0.23)	633.75*** (207.36)	0.54** (0.25)	603.99*** (224.51)
Treatment effect (LATE)	1.79** (0.79)	2,018.46*** (768.03)	1.86** (0.85)	1,898.57** (767.17)	1.79** (0.87)	2,021.89** (842.71)
Control mean	6.12	4,757.98	6.12	4,757.98	6.08	4,651.74
<i>N</i> (worker-weeks)	8,167	8,167	8,167	8,167	7,787	7,787
Strata FE	✓	✓	✓	✓		
Week FE	✓	✓	✓	✓		
Baseline Y	✓	✓	✓	✓	✓	✓
Lin Adjustment			✓	✓		
Strata × Week FE					✓	✓

Notes: This table presents robustness checks for the main effects on weekly jobs and weekly earnings among new adopters using administrative worker-week data over the eight months after randomization. Each column is a separate regression; odd columns report jobs and even columns report earnings. “Treatment effect (ITT)” is the coefficient on the randomized offer indicator. “Treatment effect (LATE)” is an instrumental-variables estimate that instruments actual subsidy receipt with the randomized offer and is interpreted for compliers (see Section 1.5). Standard errors, in parentheses, are clustered at the worker level. “Control mean” is the mean for controls in this stratum; (*N*) is the number of worker-week observations. Specification features indicated by check marks: “Strata FE” = randomization-cohort fixed effects; “Week FE” = calendar-week fixed effects; “Baseline Y” = the corresponding baseline outcome included as a covariate; “Lin Adjustment” = the Lin (2013) regression-adjusted specification—baseline covariates (including the baseline outcome) are demeaned within strata and fully interacted with treatment to improve precision under randomization; “Strata × Week FE” = cohort-by-week fixed effects allowing fully flexible common shocks by cohort.

Table A5: Robustness: Impact on Jobs and Earnings of Incumbent Users

	Incumbent Users					
Treatment effect (ITT)	−0.34 (0.34)	−89.14 (289.32)	−0.33 (0.33)	−83.45 (291.94)	−0.36 (0.37)	−105.22 (318.16)
Treatment effect (LATE)	−0.89 (0.88)	−236.72 (764.43)	−1.58 (1.44)	−273.25 (726.09)	−0.95 (0.97)	−280.07 (841.57)
Control mean	7.24	5,799.22	7.24	5,799.22	7.41	5,932.15
<i>N</i> (worker-weeks)	3,570	3,570	3,570	3,570	3,415	3,415
Strata FE	✓	✓	✓	✓	✓	✓
Week FE	✓	✓	✓	✓	✓	✓
Lin Adjustment			✓	✓		
Strata × Week FE					✓	✓

Notes: This table presents robustness checks for the main effects on weekly jobs and weekly earnings among incumbent users using administrative worker-week data over the eight months after randomization. Each column is a separate regression; odd columns report jobs and even columns report earnings. “Treatment effect (ITT)” is the coefficient on the randomized offer indicator. “Treatment effect (LATE)” is an instrumental-variables estimate that instruments actual subsidy receipt with the randomized offer and is interpreted for compliers (see Section 1.5). Standard errors, in parentheses, are clustered at the worker level. “Control mean” is the mean for controls in this stratum; (*N*) is the number of worker-week observations. Specification features indicated by check marks: “Strata FE” = randomization-cohort fixed effects; “Week FE” = calendar-week fixed effects; “Baseline Y” = the corresponding baseline outcome included as a covariate; “Lin Adjustment” = the Lin (2013) regression-adjusted specification—baseline covariates (including the baseline outcome) are demeaned within strata and fully interacted with treatment to improve precision under randomization; “Strata × Week FE” = cohort-by-week fixed effects allowing fully flexible common shocks by cohort.

Table A6: Robustness: Impact on Jobs and Earnings of Pooled Sample

	Pooled Sample					
Treatment effect (ITT)	0.18 (0.19)	322.23* (170.48)	0.17 (0.20)	330.89* (173.06)	0.17 (0.21)	318.98* (184.41)
Treatment effect (LATE)	0.65 (0.62)	1,103.50* (568.02)	0.65 (0.64)	1,028.48* (553.30)	0.64 (0.67)	1,093.73* (613.90)
Control mean	6.49	5,097.52	6.49	5,097.52	6.46	5,041.09
<i>N</i> (worker-weeks)	11,737	11,737	11,737	11,737	11,512	11,512
Strata FE	✓	✓	✓	✓	✓	✓
Week FE	✓	✓	✓	✓	✓	✓
Lin Adjustment			✓	✓		
Strata × Week FE					✓	✓

Notes: This table presents robustness checks for the main effects on weekly jobs and weekly earnings among pooled sample using administrative worker-week data over the eight months after randomization. Each column is a separate regression; odd columns report jobs and even columns report earnings. “Treatment effect (ITT)” is the coefficient on the randomized offer indicator. “Treatment effect (LATE)” is an instrumental-variables estimate that instruments actual subsidy receipt with the randomized offer and is interpreted for compliers (see Section 1.5). Standard errors, in parentheses, are clustered at the worker level. “Control mean” is the mean for controls in this stratum; (*N*) is the number of worker-week observations. Specification features indicated by check marks: “Strata FE” = randomization-cohort fixed effects; “Week FE” = calendar-week fixed effects; “Baseline Y” = the corresponding baseline outcome included as a covariate; “Lin Adjustment” = the Lin (2013) regression-adjusted specification—baseline covariates (including the baseline outcome) are demeaned within strata and fully interacted with treatment to improve precision under randomization; “Strata × Week FE” = cohort-by-week fixed effects allowing fully flexible common shocks by cohort.

Table A7: Impact on Travel Expenditure, Commute Time of Workers, and Spillover on Family Members' Earnings and Jobs

	Commuting outcomes		Spillovers on family	
	Travel Cost (INR)	Travel Time (Minutes)	Household Income (INR) (Non-Worker)	Prob. Husband Unemployed
	(1)	(2)	(3)	(4)
Panel A: New Adopters				
Treatment effect	-18.16* (9.60)	-3.35 (4.41)	2,407.86 (2,612.37)	0.02 (0.04)
Control mean	120.47	48.51	23,141.25	0.11
N: Worker-Survey Round	971	1,029	325	456
Panel B: Incumbent Users				
Treatment effect	-13.56* (7.61)	-8.15** (3.96)	7,742.88** (3,251.54)	-0.03 (0.04)
Control mean	65.36	45.47	18,038.76	0.11
N: Worker-Survey Round	448	467	147	206
Panel C: Pooled Sample				
Treatment effect	-15.77** (6.53)	-4.63 (3.28)	4,889.24** (1,990.01)	0.00 (0.03)
Control mean	100.13	47.42	21,311.01	0.11
N: Worker-Survey Round	1,426	1,500	472	662

Notes: This table reports intent-to-treat estimates on commuting outcomes for workers and spillovers on family members, using follow-up survey data. Commuting outcomes (Cols. 1–2): Workers were contacted in three follow up survey rounds after randomization. In each call, they were asked to list every job performed on the day prior to the call and, for each job, to report (i) the out-of-pocket travel cost (INR) and (ii) the minutes spent traveling for that job. The analysis is conducted at the job level (one observation per job reported across the three waves); coefficients therefore capture per-job effects on travel cost and travel time. Household spillovers (Cols. 3–4): Household income (excluding the platform worker) and the spouse's employment status were collected in two follow-up survey rounds. These outcomes are analyzed at the **worker-survey-round** level. Panels A–C correspond to new adopters, incumbent users, and the pooled sample, respectively. Regressions follow Equation 1.1 and include cohort (randomization-batch) fixed effects and **survey-wave** fixed effects; standard errors, in parentheses, are clustered at the worker level. “Control mean” is the mean for controls in the corresponding panel; (N) reports the number of observations for that column (jobs for Cols. 1–2; worker-survey rounds for Cols. 3–4). Currency is INR. Stars denote p-values (* (p<0.10), ** (p<0.05), *** (p<0.01)).

Table A8: Impact on Distance Traveled and Job Timing: Incumbent Users

Panel A: Distance						
	Total KMs	0–6 KMs	6–12 KMs	12–18 KMs	18–24 KMs	24 KMs+
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment Effect (ITT)	-2.62 (5.13)	-0.19 (0.17)	-0.17 (0.18)	-0.09 (0.12)	-0.17** (0.08)	-0.04 (0.12)
Treatment Effect (TOT)	-6.76 (13.14)	-0.51 (0.49)	-0.45 (0.47)	-0.25 (0.33)	-0.46* (0.24)	-0.09 (0.30)
Control Mean	93.69	1.93	2.40	1.29	0.70	0.94
N: worker-weeks	3,467	3,467	3,467	3,467	3,467	3,467
Panel B: Timing						
	Pre 10 AM	Pre 11 AM	12–4 PM	Post 5 PM	Post 6 PM	Surge
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment Effect (ITT)	-0.18 (0.14)	-0.29 (0.20)	-0.21 (0.26)	-0.11 (0.14)	-0.06 (0.06)	-6.08 (9.51)
Treatment Effect (TOT)	-0.53 (0.46)	-0.74 (0.65)	-0.62 (0.77)	-0.33 (0.41)	-0.17 (0.18)	-26.75 (29.24)
Control Mean	1.65	2.46	3.59	1.28	0.46	69.49
N: worker-weeks	3,499	3,499	3,499	3,499	3,499	3,499

Notes: This table shows the impact of treatment on distance (Panel A) and job-timing (Panel B) for the incumbent-users sub-group. Panel A column 1 shows the impact on total distance traveled in a week for jobs, columns 2-6 show the impact on total number of jobs completed in a week within 0-6 KMs (2), 6-12 KMs (3), 12-18 KMs (4), 18-24 KMs (5), and 24+ KMs (6); round-trip from worker's home to customer's home location. Panel B columns 1-5 show the total number of jobs completed in a week that had a start time before 10 AM (1), before 11 AM (2), between 12 noon to 4 PM (3), post 5 PM (4) and post 6 PM (5). Panel B column 6 shows total amount earned in terms of 'surge' value in a given week. First row of each panel shows the intent-to-treat (ITT) estimates and the second row of each panel shows the treatment-on-the-treated (TOT) effects. Specifications include baseline value of dependent variable, randomization-cohort $\times$ randomization-strata fixed effects, and week fixed effects. Observations are at the worker-week level for 8 months post randomization period for each worker from administrative firm data. SEs are clustered at the worker-level. Significant at * 10%, ** 5% and ***1% level.

Table A9: Treatment Effects on Job Quality Indicators: Incumbent Users

	Job Rating	High Rating	Bad Rating	Absenteeism	Temp Block
	(1)	(2)	(3)	(4)	(5)
Treatment Effect (ITT)	0.00 (0.02)	0.01 (0.07)	0.01 (0.01)	0.15 (0.16)	0.07 (0.20)
Treatment (TOT)	0.01 (0.05)	0.03 (0.19)	0.01 (0.03)	0.30 (0.29)	0.19 (0.54)
Control Mean	4.85	1.22	0.08	1.25	0.51
N: Worker-Weeks	2,693	3,570	3,570	3,570	3,570

Notes: This table shows the impact of treatment on job quality indicators for the incumbent users sub-group. Column 1-3 show the quality improvement on the customer side, i.e., improvement in worker's work quality. Column 1 shows the impact on ratings at the job rating, column 2 counts the total number of jobs that received a full five-star rating, column 3 counts the total number of jobs that received a bad rating of 0 or 1 in a given week. Columns 4-5 show the quality improvement from the firm's perspective. Column 4 shows the impact on total days in a week when the worker was fully unavailable to work. Column 5 shows the total number of days in any given week when the worker was temporarily blocked and asked to undergo additional training due to poor quality performance. First row shows the intent-to-treat (ITT) estimates and the second row of each panel shows the treatment-on-the-treated (TOT) effects. Specifications include baseline value of dependent variable, randomization-cohort $\times$ randomization-strata fixed effects, and week fixed effects. Observations are at the worker-week level for 8 months post randomization period for each worker from administrative firm data. SEs are clustered at the worker-level. Significant at * 10%, ** 5% and ***1% level.

Table A10: Contract Menus and Payment Burden (Amounts in thousands of INR)

Contract	Down Payment	Principal	Interest ^a	Total PV of Total	^b
Benchmark	15.0	85.0	33.2	133.2	112.2
Price discount	13.5	76.5	29.8	119.8	101.0
Down payment relief	5.0	95.0	37.1	137.1	113.7
Down payment discount	5.0 (net) ^c	85.0	33.2	123.2	102.2

^a “Interest” is computed at a *flat* 18% per annum over a 26-month tenure on the original financed balance: $\text{Interest} = \text{Principal} \times 0.18 \times (26/12) = 0.39 \times \text{Principal}$.

^b $\text{PV of Payments} = \text{Down Payment} + \sum_{t=1}^{26} \frac{\text{EMI}}{(1+0.18/12)^t}$, where the monthly installment is $\text{EMI} = (\text{Principal} + \text{Interest})/26$. Equivalently, $\text{PV} = \text{Down Payment} + \text{EMI} \cdot \frac{1-(1+0.18/12)^{-26}}{0.18/12}$. Values are rounded to one decimal (thousands of INR).

^c “Down payment discount” keeps the contractual down payment at 15% of list price (INR 15,000), but the firm contributes INR 10,000 at purchase; the worker pays INR 5,000 up front. The financed principal and loan terms are identical to the Benchmark.

Appendix 1.B Appendix B: Return to the Firm and Worker: Details

1.B.1 Return on Investment (ROI) for the Firm

This appendix details the computation of the firm's return on its scooter-subsidy program. The object is the *present value* (PV) of incremental profits generated by a treated worker over twelve months, net of the subsidy cost. Cash flows are discounted at an annual nominal rate of $r = 22\%$, consistent with a risk-adjusted hurdle rate for short-cycle working-capital investments in this setting for the firm.

Setup and timing

Let

- δ denote incremental *weekly* profit per treated worker (INR/week);
- s denote the 12-month survival rate (probability the worker remains active at 12 months) (estimated in Table 1.4 Column 4);
- C denote the per-worker subsidy cost (INR);
- $\beta \equiv (1 + r)^{-1/52}$ denote the weekly discount factor.

I adopt a 52-week horizon with weeks $t = 1, \dots, 35$ realized and weeks $t = 36, \dots, 52$ projected using the observed survival s . Under no differential attrition by treatment, the survival profile is common, so I scale only the projected weeks by s .

Present value and ROI

The PV of incremental gains is

$$PV_G(r) = \sum_{t=1}^{35} \beta^t \delta + \sum_{t=36}^{52} \beta^t s \delta. \quad (2)$$

The corresponding ROI is

$$ROI(r) = \frac{PV_G(r) - C}{C}. \quad (3)$$

Parameterization and estimate

For new adopters (compliers), I use

$$\delta = 557.87 \text{ INR/week (from Table 1.5),} \quad s = 0.891, \quad C = 10,000 \text{ INR.}$$

Evaluating (2)–(3) at $r = 22\%$ (so $\beta = (1.22)^{-1/52}$):

$$PV_G(0.22) = 25,382.96 \text{ INR}, \quad ROI(0.22) = \frac{25,382.96 - 10,000}{10,000} = 1.538 \text{ (or 153.8\%)}.$$

For comparison, the *undiscounted* calculation ($r = 0$) yields

$$PV_G(0) = 35\delta + 17s\delta = 27,975.51 \text{ INR}, \quad ROI(0) = 1.798 \text{ (or 179.8\%)}.$$

Sensitivity of ROI: Based on different discount rate of the firm, the ROI varies only slightly as presented in table below. (Calculations in the table are for the **12-month** horizon.)

Table A11: ROI Sensitivity to the Discount Rate (r) – 12 months

Annual discount rate r	20%	22%	24%
$PV_G(r)$ (INR)	25,585.84	25,382.96	25,185.49
$ROI(r)$	1.559 (155.9%)	1.538 (153.8%)	1.519 (151.9%)

Notes: Computed from (2)–(3) with $\delta = 557.87$, $s = 0.891$, $C = 10,000$. A one–percentage–point change in r around 22% moves ROI by roughly 0.7–1.0 percentage points.

Twenty-four months (New Adopters/compliers). Extending (2) to $T = 104$ weeks and scaling weeks 36–104 by the observed 12-month survival s , the present value is

$$PV_G^{24m}(0.22) = \sum_{t=1}^{35} \beta^t \delta + \sum_{t=36}^{104} \beta^t s \delta = 44,559.00 \text{ INR}.$$

With $C = 10,000$, the **24-month ROI** is

$$ROI^{24m}(0.22) = \frac{44,559.00 - 10,000}{10,000} = 3.456 \text{ (or 345.6\%)}.$$

Pooled Sample

For the pooled sample (compliers), the incremental weekly profit is $\delta = 345.95$ INR (see Table 1.5 Column 1 Pooled Sample), the 12-month survival rate is $s = 0.915$, and the per-worker subsidy cost is $C = 10,000$ INR. Using the same timing convention and discount factor $\beta = (1 + r)^{-1/52}$ with $r = 22\%$, the present value of incremental gains from (2) is

$$PV_G^{\text{pooled}}(0.22) = \sum_{t=1}^{35} \beta^t \delta + \sum_{t=36}^{52} \beta^t s \delta = 15,859.96 \text{ INR}.$$

The implied **12-month ROI** from (3) is therefore

$$ROI_{12m}^{\text{pooled}}(0.22) = \frac{15,859.96 - 10,000}{10,000} = 0.586 \text{ (or 58.6\%)}.$$

Twenty-four months (pooled sample). Extending the horizon to 104 weeks (scaling weeks 36–104 by $s = 0.915$) gives

$$\text{PV}_G^{\text{pooled}, 24m}(0.22) = \sum_{t=1}^{35} \beta^t \delta + \sum_{t=36}^{104} \beta^t s \delta = 28,071.84 \text{ INR},$$

and hence the **24-month ROI** is

$$\text{ROI}_{24m}^{\text{pooled}}(0.22) = \frac{28,071.84 - 10,000}{10,000} = 1.807 \text{ (or 180.7\%)}. \quad (4)$$

1.B.2 Worker ROI: New Adopters

This section computes the worker’s discounted return on the scooter purchase. Let δ_w denote the incremental *weekly* earnings attributable to adoption (INR/week), C_w the worker’s *investment cost*, and r_w the annual discount rate relevant for the worker (APR). We take $r_w = 18\%$ as in the main text; the weekly discount factor is $\beta_w = (1 + r_w)^{-1/52}$. Survival is incorporated analogously to the firm calculation: weeks $t = 1, \dots, 35$ are realized; weeks $t = 36, \dots, 52$ are multiplied by the observed 12-month survival $S_{52} = 0.891$. For the 24-month horizon, we scale weeks 36, $\dots$, 104 by the same survival (constant-hazard approximation).

Parameters (new adopters).

$$\delta_w = 2,018 \text{ INR/week}, \quad r_w = 18\% \Rightarrow \beta_w = (1.18)^{-1/52} \approx 0.99682, \quad S_{52} = 0.891.$$

We report returns using the worker’s *out-of-pocket* cost that nets out the firm subsidy: $C_w^{\text{net}} = 106,000 - 10,000 = 96,000 \text{ INR}$.³⁴

Definition

Let $S_t = 1$ for $t \leq 35$ and $S_t = S_{52}$ for $t \geq 36$. The present value (PV) of incremental earnings over horizon T weeks is

$$G_{1:T}(r_w) = \sum_{t=1}^T \beta_w^t \delta_w S_t. \quad (5)$$

Worker ROI over horizon T is

$$\text{ROI}_T^{\text{worker}}(r_w) = \frac{G_{1:T}(r_w) - C_w}{C_w}. \quad (6)$$

³⁴For transparency, we also report the “gross” figures using $C_w^{\text{gross}} = 106,000 \text{ INR}$.

Twelve months (52 weeks)

Evaluating (4) at $T = 52$ gives

$$G_{1:52}(0.18) = \sum_{t=1}^{35} \beta_w^t \delta_w + \sum_{t=36}^{52} \beta_w^t S_{52} \delta_w = 93,306.82 \text{ INR.}$$

Using the net cost $C_w^{\text{net}} = 96,000$,

$$\text{ROI}_{12m}^{\text{worker}}(0.18; \text{net}) = \frac{93,306.82 - 96,000}{96,000} = -0.028 \text{ (or } -2.8\%).$$

For reference, with the gross cost $C_w^{\text{gross}} = 106,000$,

$$\text{ROI}_{12m}^{\text{worker}}(0.18; \text{gross}) = \frac{93,306.82 - 106,000}{106,000} = -0.120 \text{ (or } -12.0\%).$$

The *discounted* payback time τ solves $\sum_{t=1}^{\tau} \beta_w^t \delta_w S_t \geq C_w$. With the net cost, $\tau \approx 54$ weeks; with the gross cost, $\tau \approx 61$ weeks.

Twenty-four months (104 weeks)

Under the same survival adjustment for weeks 36–104,

$$G_{1:104}(0.18) = \sum_{t=1}^{35} \beta_w^t \delta_w + \sum_{t=36}^{104} \beta_w^t S_{52} \delta_w = 166,216.21 \text{ INR.}$$

The corresponding ROIs are

$$\text{ROI}_{24m}^{\text{worker}}(0.18; \text{net}) = \frac{166,216.21 - 96,000}{96,000} = 0.731 \text{ (or } 73.1\%)$$

$$\text{ROI}_{24m}^{\text{worker}}(0.18; \text{gross}) = \frac{166,216.21 - 106,000}{106,000} = 0.568 \text{ (or } 56.8\%).$$

Interpretation. At a worker APR of 18% with survival-adjusted projections, the investment is close to break-even at twelve months on a discounted basis when the firm subsidy is netted from the worker's cost, with discounted payback arriving around week 54. Over twenty-four months, the worker's discounted ROI is large and positive (about 73% net of the subsidy, 57% on a gross-cost basis).

1.B.3 Worker Internal Rate of Return (IRR): New Adopters

Object and rationale. Following Bandiera et al. and Doornik et al., we compute the worker's internal rate of return (IRR) as the constant weekly rate r^* that equates the present value of expected incremental earnings from adoption to the upfront investment

cost. Relative to ROI at a fixed discount rate, IRR is a discount-rate-free summary of the project’s profitability from the worker’s perspective.

Cash-flow mapping. Let δ_w denote incremental *weekly* earnings attributable to adoption (INR/week), C_w the upfront investment cost (INR), and S_t the probability the worker remains active at week t (survival). The weekly IRR r^* solves

$$C_w = \sum_{t=1}^T \frac{\delta_w S_t}{(1 + r^*)^t}, \quad T = 104 \text{ weeks (24 months)}. \quad (6)$$

We use the survival profile from the firm/worker ROI sections: $S_t = 1$ for $t = 1, \dots, 35$ (realized) and $S_t = S_{52} = 0.891$ for $t = 36, \dots, 104$ (projection).

Parameterization (new adopters).

$$\delta_w = 2,018 \text{ INR/week}, \quad C_w = 106,000 \text{ INR}, \quad S_{52} = 0.891.$$

Result. Solving (6) for r^* yields a weekly IRR of

$$r^* = 0.013565 \Rightarrow \text{IRR}_{\text{annual}} = (1 + r^*)^{52} - 1 = 1.01509 \approx \mathbf{101.5\% \text{ per year}}.$$

Hence, the worker’s asset purchase clears a very high internal rate of return on a 24-month, survival-adjusted cash-flow basis.

Comparison to Bandiera et al. (2017) and Van Doornik et al. (2024). [Bandiera et al. \(2017\)](#) estimate the internal rate of return to BRAC’s asset-transfer program in rural Bangladesh, finding IRR between **16%** and **22%** annually, depending on the assumed opportunity cost of beneficiaries’ time and on the persistence of benefits beyond the evaluation window. The horizon there spans multiple years (two-year program with follow-ups up to 5–7 years), and benefits are measured as sustained increases in earnings/consumption and asset accumulation.

[Van Doornik et al. \(2024\)](#) study access to motorcycle credit in Brazil and summarize the investment’s profitability as an *annual real rate of return* of roughly **17%** over a **ten-year** horizon (earlier working-paper versions report $\sim 12\%$); their return consolidates persistent gains in formal employment and salaries relative to the financed asset’s cost.

Methodologically, our IRR is the same object—solve $C = \sum_t \delta_w S_t / (1 + r^*)^t$ —but two features of this platform context push our estimate much higher: *(i) timing and intensity of cash flows*, with sizeable weekly earnings gains beginning immediately (no buildup/harvest cycle), and *(ii) horizon/compounding*, as we compute a weekly IRR over 24 months with survival-adjusted flows. With $\delta_w \approx \text{INR } 2,018/\text{week}$ against $C_w = 106,000$, net benefits arrive quickly, so the break-even rate r^* annualizes to $\approx 101\%$ —an outcome consistent with very front-loaded payoffs. By contrast, Bandiera et al. and Van Doornik et al. report *annual* returns over much longer asset/program horizons with smaller per-period gains relative to cost, yielding IRRs near 16–22% and annual returns near 12–17%, respectively. In

short, the first order importance of scooters in this work setting (directly complementing production process), differences in horizon and cash-flow timing (immediate and continual vs. seasonal/persistent), account for the gap; the underlying definition of IRR is the same across studies.

Appendix 1.C Appendix C: Investment Modeled as Discrete Choice

1.C.1 Alternate Modeling Choice: Investment Decision as Discrete

Motivation. I present a discrete adoption model as a robustness check because the scooter purchase is lumpy decision. This formulation captures nonconvexities (eligibility thresholds, fixed fees) that plausibly matter at the decision margin. Importantly, for the *predictions we need for the incentive/choice experiment*—comparative statics of take-up when we vary lifetime price versus the day-of-purchase liquidity requirement—both the discrete and the continuous- k specifications imply the same qualitative mapping (with the discrete model imposing a threshold rule and the continuous model providing a smooth approximation to that threshold). We therefore keep the continuous- k benchmark as our main specification because it is more transparent for interpretation (closed-form comparative statics, units in INR per week and INR of down payment), and yields the same policy-relevant predictions for contract design. Given these equivalences and the empirical fit, it is reasonable to proceed with the continuous formulation in main text, while reporting the discrete model to demonstrate robustness.

Binary (Lumpy) Adoption

Environment and present value. Time is discrete (e.g., weekly). Adoption is binary, $k \in \{0, 1\}$, and generates a constant incremental flow surplus $f(1) > 0$ each period relative to non-adoption $f(0) = 0$. The worker receives share $\alpha \in (0, 1]$, discounts between adjacent periods at factor $\beta \in (0, 1)$, and faces a per-period separation (exit) hazard $s \in [0, 1)$. Hence the expected one-step continuation factor is $\beta(1 - s)$. For notational economy define

$$\rho \equiv 1 - \beta(1 - s) \in (0, 1],$$

so that the infinite-horizon present value of the constant flow is

$$F \equiv \sum_{t=0}^{\infty} [\beta(1 - s)]^t f(1) = \frac{f(1)}{1 - \beta(1 - s)} = \frac{f(1)}{\rho}.$$

The worker's present-value gain from adoption is therefore $\alpha F = \alpha f(1)/\rho$.

Price, financing, and liquidity. The list price is $S > 0$. A proportional price rebate $\theta \in [0, 1)$ scales the price to $(1 - \theta)S$. The lender requires a fixed down-payment share $\delta \in (0, 1)$; the remaining $(1 - \delta)$ is financed. A *repayable low-down-payment relief* $g \geq 0$ reduces cash due today one-for-one but is rolled into principal. Let $\iota \geq 0$ denote the per-rupee *present-value* finance charge on principal.

Define the PV cost and the cash due today:

$$\underbrace{C(\theta, g)}_{\text{PV cost}} = (1 - \theta) \left[S + \iota(1 - \delta)S \right] + \iota g \equiv (1 - \theta) \bar{C} + \iota g, \quad \underbrace{D(\theta, g)}_{\text{cash today}} = \delta(1 - \theta)S - g,$$

where $\bar{C} \equiv S + \iota(1 - \delta)S$. Let $L \geq 0$ denote liquid wealth at $t = 0$.

Worker's problem (lumpy choice).

$$\max_{k \in \{0,1\}} \alpha F k - C(\theta, g) k \quad \text{s.t.} \quad D(\theta, g) k \leq L.$$

Because k is binary, solve by comparing the two feasible alternatives.

Case 1: Liquidity does not bind

Assume $D(\theta, g) \leq L$ so $k = 0$ and $k = 1$ are both feasible.

Objective values: $V(0) = 0$ and $V(1) = \alpha F - C(\theta, g)$.

Result (adoption condition).

$$k^* = 1 \iff \alpha F \geq C(\theta, g) \iff \alpha \geq \alpha^*(\theta, g) \equiv \frac{C(\theta, g)}{F} = \frac{(1 - \theta)\bar{C} + \iota g}{F}.$$

Comparative statics (non-binding liquidity). Differentiate the *threshold* α^* :

$$\frac{\partial \alpha^*}{\partial \theta} = -\frac{\bar{C}}{F} < 0, \quad \frac{\partial \alpha^*}{\partial g} = \frac{\iota}{F} > 0.$$

Thus, with liquidity slack: (i) a larger rebate θ lowers PV cost and *raises* adoption; (ii) larger repayable relief g *raises* PV cost via interest (ιg) and *reduces* adoption.

Price–share equivalence (interior). Set $g = 0$. Suppose liquidity is slack for all θ under consideration, e.g. $L \geq \delta S$ so $D(\theta, 0) \leq L$. The condition for adoption is

$$\alpha F \geq (1 - \theta)\bar{C} \iff \frac{\alpha}{1 - \theta} F \geq \bar{C}.$$

Define the transformed share $\alpha' \equiv \alpha/(1 - \theta)$. Then

$$\mathbf{1}\{\alpha F \geq (1 - \theta)\bar{C}\} = \mathbf{1}\{\alpha' F \geq \bar{C}\}.$$

Proposition. With $g = 0$ and non-binding liquidity, a proportional rebate θ is behaviorally equivalent to raising the worker's effective share from α to $\alpha' = \alpha/(1 - \theta)$. $\square$

Case 2: Liquidity binds

Assume $D(\theta, g) > L$. Then $k = 1$ is infeasible and

$$k^* = 0 \quad (\text{regardless of NPV}).$$

Restoring feasibility. Feasibility requires $D(\theta, g) \leq L$, i.e.

$$\delta(1 - \theta)S - g \leq L.$$

Two equivalent ways to view the policy instruments:

(i) *Minimal rebate to clear liquidity (given g):*

$$\theta_{\text{liq}}(L, g) \equiv 1 - \frac{L + g}{\delta S}, \quad \text{so } \theta \geq \theta_{\text{liq}}(L, g) \iff D(\theta, g) \leq L.$$

(ii) *Minimal repayable relief to clear liquidity (given θ):*

$$g_{\text{liq}}(L, \theta) \equiv \delta(1 - \theta)S - L, \quad \text{so } g \geq g_{\text{liq}}(L, \theta) \iff D(\theta, g) \leq L.$$

Once feasibility is restored, adoption also requires NPV non-negativity:

$$\alpha F \geq C(\theta, g) = (1 - \theta)\bar{C} + \iota g.$$

Comparative statics (binding region). On the *liquidity* margin,

$$\frac{\partial D}{\partial \theta} = -\delta S < 0, \quad \frac{\partial D}{\partial g} = -1 < 0.$$

Hence both θ and g relax the cash constraint; locally, g relaxes it one-for-one, while a rebate relaxes it by $\delta S d\theta$.

On the *NPV* margin (once feasible), the same signs as in Case 1 apply:

$$\frac{\partial C}{\partial \theta} = -\bar{C} < 0, \quad \frac{\partial C}{\partial g} = \iota > 0.$$

Thus a higher θ helps both on liquidity and NPV; a higher g helps liquidity but hurts NPV.

Magnitude on the liquidity margin. Consider a small policy *budget* of Δ rupees.

Equal-budget rebate: a price discount of Δ implies $d\theta = \Delta/S$, reducing the down payment by

$$dD_{\text{rebate}} = \frac{\partial D}{\partial \theta} d\theta = (-\delta S) \frac{\Delta}{S} = -\delta \Delta.$$

Repayable relief: setting $dg = \Delta$ reduces the down payment by

$$dD_{\text{relief}} = \frac{\partial D}{\partial g} dg = -\Delta.$$

Therefore, per rupee of policy outlay, repayable relief moves the cash constraint by a factor of about $1/\delta$ more:

$$\left| \frac{dD_{\text{relief}}}{dD_{\text{rebate}}} \right| = \frac{1}{\delta} > 1.$$

If adoption at the extensive margin is governed by the density of liquidity $f_L(\cdot)$ at the

threshold $L^* \equiv D(\theta, g)$, the induced change in the adoption rate satisfies, to first order,

$$dA_{\text{rebate}} \approx f_L(L^*) \delta \Delta, \quad dA_{\text{relief}} \approx f_L(L^*) \Delta,$$

so $dA_{\text{relief}}/dA_{\text{rebate}} \approx 1/\delta$.

Chapter 2

Timing the Transfer: Liquidity Constraints and the Transition to Clean Fuels

Abstract. Cash-back subsidy programs require households to pay full price for the subsidized good and be reimbursed later, thus creating liquidity hurdles. We exploit India's LPG subsidy program, where the net-of-subsidy price is fixed while the over-the-counter price varies quasi-exogenously, to identify the causal effect of liquidity constraints. For low-asset households, a 1% increase in the over-the-counter price (with net price unchanged) reduces LPG purchases by 1.5% of the mean. Survey evidence shows substitution towards polluting solid fuels and worse child health. Transfer programs that neglect seemingly minor liquidity frictions may reduce essential consumption and under-deliver on welfare goals.

JEL Codes: D12, I38, Q41 **Keywords:** *subsidy, liquidity constraints, energy transition*

2.1 Introduction

We pursue two related objectives in this paper. First, we aim to estimate the effect of liquidity constraints on household consumption, using large-scale administrative data and a distinctive feature of a universal subsidy program in India. The program requires households to overcome an exogenously varying liquidity hurdle in order to purchase the subsidized good. If liquidity constraints are slack, such hurdles should be inconsequential;¹ if binding they should depress demand. This setting therefore provides a direct test of the role of liquidity constraints and enables quantification of the impact of liquidity hurdles on consumption of the subsidized good and its substitutes.

Second, we attempt to provide evidence helpful for effective subsidy delivery in a broad range of programs. Effective delivery of subsidies can be a critical component of public policy, particularly those that aim to encourage purchase of targeted goods and services.² But the mode of subsidy delivery can be critical. For example, consider the issue of timing. A simple approach is to credit the subsidy at the point of sale, but this has several drawbacks and many subsidy programs include an ex-post subsidy design, where households pay the unsubsidized price upfront and then receive the subsidy later.³ The ex-post approach has its own potential drawbacks, this paper argues. In particular, even if the ex-post subsidized price would be affordable, requiring the full unsubsidized payment upfront creates an extra liquidity hurdle that may stretch household liquidity too far – especially for poorer households. Thus, we may encounter (partial) policy failure: intended recipients, even some of the most targeted, may fail to take up the subsidy.

To pursue these objectives, we study an at-scale, cash-back energy subsidy program of the Indian government designed to encourage household use of LPG (“liquefied petroleum gas”) for clean cooking by keeping its price stable and affordable.⁴ The program holds the net-of-subsidy price of the product, 14.2-kg LPG refill cylinders, virtually unchanged over

¹Liquidity constraints refer here to limitations on the availability of household assets to use as payment for a good or service today. Availability would come through efficient credit markets that enable households to tap into their assets (including human capital, i.e., future earnings), or could come partly through access to and ability to make optimal use of safe, liquid savings vehicles. A “liquidity hurdle” here is a feature of the environment that raises the household’s need for access to its assets.

²Governments across the world provide subsidies or conditional cash transfers to correct for externalities, lack of information or other market failures, e.g., for preventive health, clean energy, schooling fees and opportunity costs, among others.

³For example, the ex post direct debit approach for nationwide LPG subsidies in India has been shown to decrease leakage and corruption (Barnwal, 2024). Other examples of ex-post subsidy design include electric vehicle or other green technology subsidies provided in the form of tax credit or rebate checks in the USA, Norway, Germany, China, and other countries (Bibra et al., 2021), the health expense tax credit in Canada (Canada Revenue Agency, 2025), and fertilizer subsidization through reimbursement in the Philippines (Department of Agriculture, Republic of the Philippines, 2020). Reasons for preferring the ex post subsidy design over crediting the subsidy at point of sale include lowering administrative cost, allowing more time to verify conditionalities, discouragement of over-consumption and waste, and opportunity to curb theft of public money, leakage, and black marketing.

⁴Substitutes for solid fuels – gas and electricity – are costly and require subsidy in low-income countries to sustain usage and provide insurance against price risk. Since the current electricity grid is carbon intensive; full electrification may be more CO₂-intensive than cooking with solid-fuels (Floess et al., 2023). The average PM_{2.5} exposure for exclusive LPG use to be roughly $35 \mu\text{gm}^{-3}$, compared to $250 \mu\text{gm}^{-3}$ for exclusive biomass use (Gould et al., 2024).

time by offering a time-varying subsidy; see [Figure 2.1](#), which depicts the co-movement in the unsubsidized, over-the-counter LPG refill price and the household LPG refill subsidy, alongside the resulting constancy in the subsidized price of LPG refills.

The subsidy is delivered *ex post*: to purchase a refill, consumers must pay the current over-the-counter price, and then receive the subsidy in their bank account after one week, on average. The key to our empirical strategy is that the over-the-counter price co-moves virtually one-for-one with the plausibly exogenous international price of LPG. Indeed, [Figure 2.2](#) shows a nearly perfect correlation between the (domestic) over-the-counter LPG price and the international price of LPG lagged by a month during our study period. It follows that when the international price of LPG is high, consumers must pay a high price but get a high subsidy rebated; when it is low, the upfront price and the subsidy are also low. This creates a liquidity hurdle that varies exogenously while the effective (subsidized) price is held fixed, enabling us to identify the quantitative impact of monetary liquidity requirements on access to the subsidized good.⁵

We thus empirically investigate how variation in the over-the-counter price of LPG – which measures the size of the liquidity hurdle – impacts purchases of this clean fuel by both non-poor and poor households. We also estimate the impact of this over-the-counter LPG price on substitution towards inferior and more polluting solid fuels, and we trace the impact further toward a health burden on children likely resulting from the inferior fuels’ indoor pollution. We further explore evidence pinpointing the liquidity effect and ruling out other explanations for the results.

Our analysis draws on four distinct data sources. First, we use transaction-level administrative records covering roughly one million LPG-using households in Indore, a major district in central India. These data document about 24 million LPG refill transactions over a 26-month period (November 2017–December 2019). Second, we link these administrative records to a primary survey of about 3,000 households from the same sample, conducted in two rounds between October 2018 and December 2019, which provides detailed information on biomass fuel consumption, demographic characteristics, and asset ownership. Finally, we incorporate two nationally representative secondary datasets: the Demographic and Health Survey (DHS) from India (also known as the National Family Health Survey or NFHS), which provides information on births and child-level health outcomes, and the Periodic Labor Force Survey (PLFS), which offers information on employment and households’ cash flow.

Our results indicate a large and significant impact of the cash-back subsidy program’s liquidity hurdle on LPG refill purchases by households. In particular, low-asset households decrease their refill purchases by 1.5% of their average when the over-the-counter refill price increases by 1%. The effect is similarly negative but much smaller and not always statistically significant for middle-to high-income households. There is also suggestive evidence that the effect is stronger the longer is the delay in subsidy receipt.

This core finding is bolstered by a second, that households respond to higher over-the-counter prices for LPG by substituting toward polluting solid fuels. Specifically, we find

⁵We think of this natural experiment as approximating an experiment in which households can pay a fixed, subsidized price for an LPG refill on the condition of offering the government a short-term interest-free loan of random size (equal to the subsidy).

that monthly explicit expenditure on solid fuels by low-asset households increases by 1.2% of the mean in response to a 1% increase in the over-the-counter LPG price; the figure is 1.6% for implicit expenditure (including imputed time cost of gathering/making solid fuel). Substitution toward more-polluting solid fuels can impact health adversely, especially of children. Results from the DHS data suggest an adverse impact of LPG price increases on child health due to a significant increase in the incidence of rapid breathing.

We provide further analysis aimed at pinpointing liquidity sensitivity as the key force behind the results. First, we show that the results are primarily driven by low-asset households, and that this conclusion is robust to using several alternative proxies for wealth. Second, we use the PLFS to construct a direct proxy for liquidity – specifically, cash on hand – and show that the results are driven by low cash-on-hand household-months. Third, we restrict the analysis to harvest months and check for heterogeneous impacts by an exogenous liquidity shifter: crop growing-season rainfall. The harvest-time effects of over-the-counter price are significantly attenuated after good rainfall, again suggesting that the amount of available household liquidity is a key moderator of the estimated effect.

Finally, our findings are robust to alternative explanations. First, one might be concerned that LPG price increases reflect general inflation, or acting as a negative supply shock to the macroeconomy. If so, the effects we find could be interpreted as wealth or income effects rather than liquidity effects. But we find no significant correlation between the overall Consumer Price Index and the over-the-counter price of LPG refills. Additionally, the results are robust to controlling for monthly unemployment, a proxy for macroeconomic conditions. Second, lack of ‘subsidy salience’, i.e. lack of awareness or understanding of the subsidy scheme, may be an issue for households, causing them to interpret changes in the over-the-counter price as real price changes. However, we find no heterogeneity in our results by measures of household awareness and understanding of the subsidy scheme. Related, we test for the possibility that imperfect trust in the government is driving results. If households doubt that the government would credit the subsidy, they could respond to the over-the-counter price as reflecting the true price with some probability. But we find no difference in estimated effects by whether respondents exhibit trust in government functionaries. We also find that the results change little when we restrict the sample to experienced LPG user households, who have more exposure to faithful delivery of the subsidy. In general, these tests engender confidence that the observed response of poorer households to the price of over-the-counter LPG is due to liquidity constraints.

Our paper makes several contributions. First, we identify and quantify pure liquidity effects on purchases of a key household staple, and find that for poor households they can be substantial. This is all the more remarkable since the liquidity needs in our setting are moderate and quite temporary; reimbursement typically comes within a week. Thus, using a plausible identification strategy exploiting a natural experiment setting embedded in a developing country program at scale, we add to the literature documenting that liquidity constraints can be significant, especially for poorer households, and hence that credit markets operate imperfectly in developing countries. Our approach further allows us to approximate a novel parameter: the elasticity of consumption to liquidity requirements.

Previous literature on the relevance of liquidity constraints is extensive (for a seminal contribution, see [Zeldes \(1989\)](#)). However, causal effects of liquidity constraints are difficult to estimate, given the challenge of varying only liquidity availability or requirements while

keeping wealth, prices, and preferences fixed. Much of the literature seeking to identify liquidity constraints examines consumption responses to variation in liquidity available to the household, whether through anticipated or unanticipated income infusions (see, e.g., [Bazzi et al. \(2015\)](#), [Aguila et al. \(2017\)](#), [Gelman et al. \(2020\)](#), and [Gross et al. \(2022\)](#), and the references therein), exogenous credit availability (e.g. [Tarozi et al. \(2014\)](#), [Berkouwer and Dean \(2022\)](#)) or savings opportunities ([Brune et al. \(2021\)](#), [Kansikas et al. \(2023\)](#)). But there are challenges involved in isolating liquidity effects using these methods. Timing purchases with anticipated income receipt may be more about cash management or optimal lumpy consumption (e.g. the occasional party) than truly distorted consumption, while credit can be viewed as a change in net worth (not merely liquidity) depending on what the household believes about the consequences for default.

A second approach to identifying liquidity constraints is much rarer: varying liquidity requirements rather than liquidity availability. Using RCT methodology, [Casaburi and Willis \(2018\)](#) and [Levine et al. \(2018\)](#) vary the timing of payment required to purchase two significant goods: crop insurance and cookstoves, respectively. Both find large effects of allowing households more time (at perhaps a better time) to come up with payment.⁶ We follow this general approach with one key difference: we vary liquidity requirements not via timing of payment but via the size of the required payment itself, holding price fixed. This enables us to cleanly identify a unique parameter: the response of purchases to the monetary amount of liquidity required.

Second, we shed further light on a potentially critical issue in the design of programs featuring subsidies or conditional cash transfers, relevant to both equity and targeting: how should the transfers be timed? While we are unable to estimate the full tradeoffs involved, this paper documents a potentially significant cost to delay in providing the subsidy. The implications are relevant more broadly than for classic subsidy programs. Consider conditional cash transfer (CCT) programs, e.g. Progresa and similar CCT programs aimed at increasing investment in human capital. To meet these programs' conditions, households had to attend school for some period of time before receiving the transfer; thus they needed to pay any fees and the opportunity cost of foregone labor upfront. The logic of this paper is that households were thus facing program-induced liquidity hurdles that may have reduced program impact.⁷ Even minor delay creates liquidity hurdles and can limit access to the subsidized good or conditional transfer, this paper argues, and especially for poorer households.

The point extends beyond the timing of subsidies to the level of subsidies. Any required positive co-payment represents not only a marginal price, but also a liquidity hurdle. Thus, the optimal co-payment in the presence of liquidity constraints will typically be smaller than the one in the absence of such constraints. This applies to subsidized food, fuel, and even medicine (as [Gross et al. \(2022\)](#) argue).

Taken together with recent research on administrative reforms in welfare delivery in developing countries ([Barnwal, 2024](#); [Muralidharan et al., 2020](#)) that aim at improving efficiency of public programs, the results highlight a possible equity-efficiency trade-off in

⁶As with credit, one might be concerned that the opportunity to default when delaying paying for the good affects the implicit price, not liquidity requirements alone.

⁷This does not contradict positive impact from these programs, but suggests the impact was not as large as it could have been without liquidity issues.

overhauling public welfare programs, as outlined by [Kleven and Kopczuk \(2011\)](#). They suggest that while delay may eliminate a significant amount of mistargeted benefits, it may also exclude key targeted beneficiaries; and conversely for point-of-sale reimbursement. Put differently, reaching some of the poorest may be costly in terms of leakage. Ways to design subsidy policy that may soften or eliminate this tradeoff, however, are discussed in the Conclusion.

Third, we add to the existing evidence for policies seeking to reduce pollution and health externalities by encouraging transition to clean energy in developing countries.⁸ Subsidy policies can be useful in such transitions, but how those subsidies are delivered can determine their effectiveness, and particularly for poorer households, which may be at the greatest risk of using more damaging technologies. Several countries run household subsidy programs to accelerate the transition to cleaner energy sources. For environmental justice reasons, energy subsidy programs often include provisions that target low-income households. But the effect of these programs' administrative design on efficiency and equity in the take-up of clean energy remains understudied. Policy-makers usually pay less attention to how the subsidy should be provided, focusing instead on what the subsidy level should be. The findings of our paper suggest that the design of these programs, in terms of their timing, can be critical for clean energy access.

Recent research shows that households' financial constraints play a critical role in adoption of clean energy in low-income contexts. Using LPG sales records, [Afridi et al. \(2021\)](#) underscore the salience of financial constraints in switching to clean fuels in spite of the LPG subsidy program in rural India. Close to our setting, [Berkouwer and Dean \(2022\)](#) find that households' willingness to pay for energy-efficient cooking stoves increases when credit constraints are loosened via a loan offer. Our study provides evidence that liquidity constraints limit households' ability to sustain the transition to clean energy through regular usage of such fuels, and especially poorer households – even when they have already acquired more efficient cooking technology.

The remainder of the paper is organized as follows. We explain the institutional background in Section 2.2 before developing basic theoretical predictions for liquidity-constrained consumption choice in Section 2.3. Section 2.4 describes the data and presents summary statistics, while Section 2.5 outlines our empirical strategy. Baseline results on purchases of LPG and alternative fuels are in Section 2.6, while health results are in Section 2.7. Section 2.8 presents results buttressing the liquidity interpretation, and Section 2.9 concludes.

⁸There are several rationales for providing LPG subsidies, including to reduce indoor pollution and the resulting health burden ([Gould et al., 2024](#)), and to save time spent on using, collecting or making solid fuels ([Afridi et al., 2023](#)). The impact of indoor air pollution on health and other related outcomes is extensively discussed in the literature. A smaller number of articles provide causal evidence on how the transition from biomass-based fuel to clean energy for cooking helps improve health outcomes (e.g., [Cesur et al. \(2017\)](#); [Imelda \(2020\)](#)), although this may not always be the case ([Berkouwer and Dean, 2023](#); [Thivillon, 2023](#)). LPG subsidies have been particularly highlighted to have a high potential to reduce the USD 1.6 trillion global health burden of biomass-based fuel ([Gould et al., 2024](#)).

2.2 Institutional background

The first step a household has to take in order to use LPG is to obtain a “connection”, or “account” – i.e. register with one of the three state-owned Oil Marketing Companies (OMCs), India’s exclusive LPG suppliers, through their network of local “dealers”. This is an upfront cost of about INR 3200 (45 USD), about two weeks’ worth of household income in rural areas.⁹ To expand access to LPG, the Government of India launched the *Pradhan Mantri Ujjwala Yojana* (PMUY) in April 2016. The PMUY is the largest program for access to clean fuel in the world, reaching 72 million poor families between April 2016 and June 2019. The program mandates that women in socioeconomically disadvantaged households can obtain an LPG account at no upfront cost. The government bears the beneficiary’s security deposit and administrative charges for an LPG account (USD 25). The beneficiary also receives an interest-free loan (USD 20) to purchase the stove and the gas cylinder.¹⁰

An LPG account entitled the consumer to a universal, nationally set subsidy on refills of the LPG cylinder between 2014 and mid-2020. All residential LPG consumers in India, irrespective of PMUY status, received a subsidy or so-called “direct benefit transfer” (DBT), for up to 12 LPG cylinder refills a year.¹¹ The DBT for LPG subsidy functioned in the following way during this period. The consumer paid the entire market price (over-the-counter) to her local LPG dealer at the time of refill cylinder purchase. The dealer recorded the refill purchase against the consumer’s ID in a centralized database. The subsidy amount was then directly deposited into the consumer’s linked bank account, on average within 7 days (see [Figure A1](#) in [section 2.A](#)). As a consequence of biometric identification of consumers, the possibility of resale of LPG refills is minimal, if not absent.

The subsidy program initially aimed to maintain a stable net nominal price (over-the-counter price minus subsidy) of around INR 500 per cylinder during our study period (equivalent to INR 350 in the base year 2012), which each consumer was entitled to for up to 12 cylinder refills per year. Districts are typically broken down into multiple LPG markets by distance from the LPG cylinder bottling plant. LPG refill prices vary by a small amount (INR 2 - 5) between markets, primarily depending on the transportation cost of moving refill cylinders from the bottling plant. Hence the variation in refill prices is mainly due to monthly changes in market price and minimally due to spatial variation. Beginning in April 2016, the government modified the pricing by implementing pre-announced small monthly increases (INR 2-4 per refill); this resulted in a roughly constant real price during

⁹“Connection” is the official term that refers to registration for obtaining a gas pressure regulator, a consumer booklet with unique ID along with the first LPG cylinder refill. To register for a connection, a consumer has to provide proof of identity, residential address, bank account information, and submit a security deposit of 25 USD with her local LPG dealer. The consumer has to pay the market price separately for the gas in the cylinder (10 USD) and a stove (10 USD). Note that the average rural (urban) household income was approximately INR 7,215 (INR 10,000) per month in 2011, the latest year for which reliable estimates are available ([Desai et al., 2011](#)). 1 USD = INR 72 during our study period.

¹⁰The loan from the OMC to the PMUY consumer was to be recovered from her LPG refill subsidy entitlement, discussed next, but this recovery was not effectuated during the period of our study.

¹¹Throughout, we refer to the standard-sized LPG cylinder in the Indian market, which holds 14.2 kg of LPG. All registered consumers are assigned a unique consumer ID and a booklet that records, among other details, the date of LPG account creation, the LPG dealer holding the account, and every LPG refill purchase. Consumers can purchase refills from any of the three approved dealers. All dealers offer LPG accounts and cylinder refills at the same market price within their market.

our sample period, as Figure 2.1 shows. The Figure also shows that the subsidy delivered via DBT varied monthly with the LPG market price, which itself moves in tandem with the international market price of LPG. Specifically, the Indian government regulates the over-the-counter price every month based on the Saudi Aramco LPG price in the preceding month; see Figure 2.2. In short, consumers pay an exogenously varying over-the-counter price, and receive a subsidy debited to their bank account shortly thereafter, which is designed to keep the net real price constant. For instance, if the market price of an LPG cylinder happened to be INR 820, given a target net price of 500, a subsidy amount of INR 320 would be directly deposited into the consumer’s bank account within 7 days of purchase on average.

Nationwide, an estimated 79% of all households had LPG accounts in 2018 (the period of our study) but with a significant rural-urban gap of almost 40 percentage points (PPAC Report, 2018). The PMUY program has been successful in significantly improving rural households’ access to LPG, but has yet to ensure widespread consistent LPG usage (Pandey, 2019). LPG use continues to be much lower in rural areas, where mean annual consumption is about four LPG cylinders – roughly half that in urban areas (Mishra, 2019).

2.3 Theory

We next present a model to illustrate how the liquidity hurdle imposed by the LPG subsidy program may negatively impact consumption of LPG. The model clearly separates the liquidity effect from price and wealth effects, and makes predictions about which kinds of households are most impacted by the liquidity requirements.

Time is divided into a short, current period – corresponding to the time a household must wait for its subsidy rebate, often a week – and the future. A household maximizes utility by choosing current consumption of a general good C , at price normalized to 1, and whether to make an indivisible fuel purchase $X \in \{0, 1\}$. The program fixes the subsidized price of X at p , but also requires the household to pay the market price $p + q$ for good X upfront, and only in the future get reimbursed q , the subsidy amount. Our theoretical focus is on how consumption of X varies with the liquidity hurdle q .

The household is endowed with net worth or “wealth” W (interpreted broadly to include future earnings, e.g.), so that $W - C - pX$ is available for future consumption. The liquidity constraint is modeled as a fraction $\lambda \in [0, 1)$ of total wealth $W + qX$ being unavailable for purchases in the current period.¹² If credit markets are perfect, the household has immediate access to all its wealth and $\lambda = 0$. On the other hand, imperfect credit markets imply that the household is limited in ability to borrow from its wealth, whether future earnings, illiquid tangible assets, or the promised subsidy, leaving $\lambda > 0$.¹³

¹²Counting total net worth for current-period borrowing purposes as $W + qX$ rather than W is equivalent to assuming that the promised future subsidy can be borrowed against, conditional on X being purchased, as well (or poorly) as other sources of net worth. This symmetric treatment allows the liquidity constraint to be captured by a single parameter (λ), but could be generalized without changing results.

¹³Savings constraints, whether structural or behavioral, are unmodeled but may be thought of as incorporated in λ in a reduced-form way; if saving is difficult, the fraction of net worth that can be accessed in the short run $(1 - \lambda)$ may be lower.

Assuming a log utility specification, for $a \in (0, 1)$ and $b > 0$, the household's problem is:

$$\begin{aligned} \max_{C, X \in \{0, 1\}} \quad & (1 - a) \ln(C) + bX + a \ln(W - C - pX) \\ \text{s.t.} \quad & C + (p + q)X \leq (1 - \lambda)(W + qX) \end{aligned} \quad (2.1)$$

Consider first the unconstrained maximization problem, where $\lambda = 0$. If the household chooses $X = 0$, it is optimal to set $C = (1 - a)W$, leaving aW for future consumption. If the household chooses $X = 1$, it is optimal to set $C = (1 - a)(W - p)$, leaving $a(W - p)$ for future consumption. It is straightforward to show that buying fuel gives greater utility than abstaining from fuel in this unconstrained case iff $W \geq p/(1 - e^{-b})$. Households of sufficiently low net-worth decide that buying a unit of fuel would leave the rest of their consumption too low.

Thus, key determinants of gas purchase in the absence of liquidity constraints are the subsidized price p and net worth W . The liquidity hurdle q has no impact, however. This provides a testable implication of the absence of liquidity constraints: if subsidized price and net worth are held fixed, varying liquidity requirements should have no impact on consumption of good X .

Consider next the constrained problem. Note that the liquidity constraint is not binding at the optimum when $X = 0$ if $(1 - a)W < (1 - \lambda)W$, i.e. $\lambda < a$; for brevity, we restrict attention to this case. The liquidity constraint does not bind at the optimum when $X = 1$ iff $(1 - a)(W - p) + p + q \leq (1 - \lambda)(W + q)$, i.e. $W \geq (ap + \lambda q)/(a - \lambda)$.

Putting these facts together, it is clear that $X = 1$ is optimal if

$$W \geq \max \left\{ \frac{p}{1 - e^{-b}}, \frac{ap + \lambda q}{a - \lambda} \right\},$$

since the unconstrained allocation with $X = 1$ is both feasible and preferable to the unconstrained allocation with $X = 0$. It is also clear that if $W < p/(1 - e^{-b})$, $X = 0$ is optimal because the unconstrained allocation with $X = 0$ is both feasible and preferable to the unconstrained allocation with $X = 1$.

The remaining possibility involves

$$W \in \left(\frac{p}{1 - e^{-b}}, \frac{ap + \lambda q}{a - \lambda} \right).$$

This interval is non-empty iff $q > p(a - \lambda e^b)/[\lambda(e^b - 1)]$, which the rest of the paragraph assumes. In this interval, purchasing X would be preferable if there were no liquidity constraints, but the unconstrained C is not affordable if the household purchases X . So, the choice is between purchasing X by underconsuming C in the current period, or abstaining from X . One can show that there is a $\widehat{W}(q)$ in the interior of this interval such that the household purchases X iff $W > \widehat{W}(q)$.¹⁴ One can also show that $\widehat{W}'(q) > 0$. This is

¹⁴If W is near the top of this interval, the unconstrained allocation with $X = 1$ is nearly affordable, and strictly preferable; so the household does better to purchase X . If W is near the bottom of this interval, the household is almost indifferent between the unconstrained allocations with $X = 0$ and $X = 1$, but is not close to affording the unconstrained allocation with $X = 1$; so abstaining is better.

because as q rises, the sacrifice in C required in order to purchase X gets larger, so marginal households – for whom W and C are lowest and the marginal utility from C highest – abstain from X .

The optimal purchase decision is pictured in (q, W) -space in Figure 2.3.¹⁵ It is clear that as q rises, so does the wealth cutoff for purchasing fuel, $\widehat{W}(q)$. Thus, we should see fewer fuel purchases when q is higher – fewer households find it optimal to purchase when they must pay a higher upfront cost, since in the presence of binding liquidity constraints this necessitates greater sacrifice in current consumption. Further, the model predicts that the drop in fuel purchases resulting from a higher q is concentrated among households with relatively low wealth. The lowest wealth households, e.g. those with $W < p/(1 - e^{-b})$, are so poor they would never purchase fuel, even without liquidity constraints. But of those who would ever purchase fuel, it is the lowest-wealth households that are sensitive to q ; see Figure 2.3. We should thus see the effects most strongly among the poorest of those households ever purchasing fuel. The focus of this paper is on empirically testing this paragraph’s predictions.

The model is stylized and does not fully capture the dynamics of the decision-making involved in purchasing LPG refills. For example, it does not model the dependence of demand for fuel on the accumulated stock of fuel already purchased. While a fully dynamic model is beyond the scope of this paper, the basic impact of liquidity constraints in distorting purchases should be robust to that setting.

2.4 Data

Our main data come from Madhya Pradesh (MP), the second-largest Indian state by area and the fifth largest by population with over 75 million residents. Over 60% of MP households (rural and urban) had LPG access in January 2018 (PPAC Report, 2018). We focus on the richest district of MP, Indore.

We use two main sources of data: (1) administrative, consumer-level data from the LPG sales records of all three OMCs for all household-level LPG consumers in Indore district, and (2) household-level survey data from approximately 3000 randomly sampled rural households in Indore district.¹⁶

2.4.1 Administrative sales data

The OMC administrative sales data cover all LPG refill transactions of all consumers for all three OMCs in Indore for 26 months, November 2017 – December 2019. Besides unique consumer IDs, the data contain information on the date of LPG account registration, PMUY status, LPG dealer, residential address of each consumer, and date (day, month and year) of every LPG refill transaction of each consumer. Using a machine learning algorithm, we

¹⁵The Figure is for parameter values satisfying $\lambda \in (a/e^b, a)$, which guarantees that the above interval is non-empty for any $q \geq 0$. The case where $\lambda \in (0, a/e^b)$ is similar for $q \geq p(a - \lambda e^b)/[\lambda(e^b - 1)]$ but flat for $q \in [0, p(a - \lambda e^b)/[\lambda(e^b - 1)]]$.

¹⁶We use the term consumer and household interchangeably because each household typically has only one LPG account with their local LPG dealer.

generated the gender variable of the consumer (from PMUY status and the consumer’s name) and their rural-urban location (using Google API and address pin-code), allowing us to present results on Rural and Urban subsamples.

Monthly market prices from the 26 markets of Indore district are available from all three OMCs’ (i.e. IOCL, BPCL and HPCL) administrative data. Note that refill prices vary slightly between markets but not between dealers within a market. The administrative data link all LPG dealers to their corresponding market and thus to the corresponding monthly LPG price. Given that each consumer’s dealer is identified in the data, we have information on the monthly market-level LPG refill price data for all consumers for the entire study period.

In addition to monthly price data for refill cylinders, we also have the LPG subsidy per refill in each month during the study period. After merging, we have a database of all customers’ LPG refill purchase quantities, refill subsidy in each month, and refill market price for 676 market-months (26 Indore markets x 26 months). We drop outliers, consumers who purchased an abnormally high number of LPG refills (0.024% of the administrative sample, who consumed more than 40 refills in a year).

The administrative data are unbalanced – LPG refill consumption is missing in months prior to registration for consumers who registered for an LPG account during our sample period. We balance the sample by recoding these missing values to 0. For instance, if the consumer’s account was registered in January 2018, her LPG refill purchase is entered as 0 in November and December 2017. Section 2.B reproduces all main results with the unbalanced panel that results from coding LPG consumption as missing prior to account initiation.

Data from one OMC, BPCL, also contain the date the subsidy is credited to the consumer for a majority of transactions. This allows us to measure delay, defined as the number of calendar days between the delivery date of a refill and the associated subsidy bank deposit. These are aggregated to the mean or median level within-month to create a proxy for the typical delay prevailing in the month a refill is purchased.

2.4.2 Household-level data

We utilize data from a cluster randomized control trial (RCT) implemented in rural areas of Indore district by Afridi et al. (2021) between October 2018 and December 2019.

150 villages were randomly sampled from four census blocks in Indore district excluding the city of Indore. In the sampled villages, a household was deemed eligible for the study if it had a currently residing member younger than 10 or older than 55. Twenty eligible households were randomly sampled in each of these villages by systematic random sampling during a baseline survey in November and December 2018. Following the completion of a cluster-randomized information campaign on the health benefits of LPG and information on the universal LPG subsidy between January and September 2019, the endline survey was conducted between October and December 2019. Thus, the households surveyed in the baseline were revisited during the same season approximately a year later, with negligible attrition (< 2%).

Households in the sample were asked whether they currently had an LPG account. If they did, account details, including their unique consumer ID and number of refills in the past year, were recorded from their consumer booklets and photographed. Detailed information

on household composition, assets,¹⁷ socio-economic characteristics, use and collection of fuel in the previous month, PMUY status, and health awareness was gathered for all households irrespective of whether they had an LPG account or not. The same household level data were gathered at baseline and endline.

Households with LPG accounts were linked to the OMCs’ sales data (described above) using the unique consumer ID, enabling us to track these households’ LPG consumption in all 26 months of the administrative data. This paper’s analysis is restricted to the subsample of surveyed households who either already had an LPG account or obtained one during the sample period of the administrative data, November 2017 – December 2019.¹⁸

2.4.3 Other datasets

Since we did not collect information on income or expenditure in our household Survey, we use individual and household level data from multiple rounds of the nationally representative Periodic Labor Force Survey (PLFS rounds 2017-18, 2018-19 and 2019-20) that overlap with our study period of November 2017 – December 2019. Data from the rural sample is used to estimate total monthly household income (earned, unearned, self-employment, and other sources of income) and total monthly household expenditure. The difference between the two is our estimate of a household’s “cash-on-hand” in a month.¹⁹ Data on caste and occupation is then used to average cash-on-hand by month-caste-occupation cells, using 12 months, 4 main occupation categories (salaried, self-employed on-farm, self-employed off-farm and casual labor), and 4 caste groups (SC, ST, OBC and General or higher-caste categories) to create 192 cells.

We use the DHS 2019-21 India data to construct household-level health outcomes: (i) at the infant level, birth weight for births within a year prior to the survey, and (ii) at the young child level, indicators for cough and rapid breathing in the two weeks before interview for living children under five. To cover the birth recall window, we extend the monthly LPG price series backward to June 2018 and forward to May 2021, using BPCL–Indore price data and the CPI to build a continuous real, month-level series. Pregnancy timing for birth outcomes is recovered from the DHS reproductive calendar and linked to the mother’s location-month. For birth weight, we average the real LPG price over the gestation months; for cough and rapid breathing, we use the interview-month price.

2.4.4 Summary statistics

The domestic, over-the-counter market price of an LPG cylinder is determined by the government, based on the international market price for LPG. Specifically, the domestic price of LPG is expected to be reset by the government every month to reflect the world price. This can be seen in [Figure 2.2](#) – the domestic (over-the-counter) price of LPG is

¹⁷We use the first component of a principal component analysis over several indicators measuring the economic status of a household to generate the asset index. These indicators include ownership of land and farm animals, *pucca* house, and a list of consumer durables.

¹⁸The earliest (latest) LPG account date is July 1986 (September 2019) in our household survey data.

¹⁹More extensive data that could better capture stock rather than flow of cash-on-hand, e.g. using techniques of Samphantharak and Townsend (2010), is unfortunately lacking.

nearly perfectly correlated with the one-month lag of the international LPG price in INR (correlation coefficient of 0.99, p -value=0.0004).

The administrative sales data are summarized in [Table 2.1](#), which shows LPG refill consumption from administrative sales records of all three OMCs for the entire district of Indore, Nov 2017 - Dec 2019. The data are at the consumer level and represent almost a million customers. 7.4% of the consumers are PMUY account holders, while 32% are women and 57.2% are rural, as shown in Panel A. Panel B provides statistics on LPG refill transactions, overall and by rural-urban residence. The average number of LPG refills per month is 0.578, but just 0.252 for PMUY consumers. This PMUY–non-PMUY consumption gap holds in both rural and urban areas.

The average LPG over-the-counter price and subsidy amount per kg of gas for each of the 26 months in our data are shown in [Figure 2.1](#). Note that the LPG refill subsidy tracks the over-the-counter price as it varies each month, leaving the real effective price (over-the-counter price minus subsidy) essentially constant during our study period at approximately INR 25 per kg.²⁰ Note that the subsidy amount ranges from less than 20% to almost 100% of the subsidized price during this time period.

[Table 2.2](#) summarizes the characteristics of the rural households with an LPG account in our Survey data. Of the 3003 households in the sample, almost 70% had an LPG account during the study period (i.e. 2091 households), of which 39.1% were PMUY consumers. Self-employment or casual labor are the dominant occupations of household heads, as shown in Panel A. Panel B shows that the explicit, average monthly expenditure on solid fuels purchased by the households is INR 56.8 and INR 151.2 for firewood and dung cakes, respectively. However, monetizing the reported time spent by the household in solid fuel collection/making in the survey data by multiplying the time spent in fuel collection/making with the minimum unskilled wage in rural MP, we find that the implicit monthly household expenditure on either solid fuel is more than INR 500 – higher than the average effective price of one LPG refill (see [Appendix Table A1](#)). Households spent almost 4 hours in a typical week of the previous month collecting or making solid fuels, according to our survey data. Panel C shows that average LPG refills purchased per month by households with LPG accounts was 0.41 (using OMC sales data from Nov 2017 - Dec 2019), lower than from the full administrative data ([Table 2.1](#)). PMUY households’ annual refill consumption was just over half of non-PMUY households’ consumption.

In [Appendix Table A2](#) we compare the characteristics of the PMUY and non-PMUY consumers using the household survey data. PMUY households have a lower asset index, are less likely to belong to the general caste group (relative to socio-economically disadvantaged Scheduled Caste (SC), Scheduled Tribe (ST) or Other Backward Castes (OBC)), are less likely to own or lease land, have lower education and are more likely to have casual labor as their main occupation. These data, therefore, underline the fact that PMUY households are significantly more impoverished, and thus more likely to be liquidity constrained, than non-PMUY households.

[Appendix Figure A2](#) shows the average per month cash-on-hand (COH) of rural

²⁰[Appendix Table A1](#) summarizes these prices for a standard 14.2 KG refill cylinder for the 676 market-months in our data. As mentioned earlier, the over-the-counter price of LPG may slightly vary across LPG distributorship markets in a given month, due to small differences in transport costs and local taxes.

households in India, calculated by taking the difference in monthly total household income (from all sources) and the total household expenditure in the PLFS data. The figure indicates that households have less cash-on-hand during the off-agricultural season (April-September, prior to the winter cropping season), suggesting cyclicity of liquidity. In Appendix [Table A3](#), we summarize the cash-on-hand data that we obtain for the 416 (=4x4x26) occupation-caste-month groups, which we use later to analyze heterogeneous effects of over-the-counter price on LPG refill take-up in our Survey sample of households in rural Indore.²¹

2.5 Empirical strategy

Our empirical strategy is to estimate responses of LPG purchases to variation over time in the over-the-counter LPG price, and with it the subsidy amount, particularly among low-wealth households. Since the effective (net) LPG price – over-the-counter price net of subsidy – remains essentially constant over time, while the amount that must be paid upfront varies significantly and exogenously, our setting provides a unique opportunity to quantify the impact of liquidity requirements.

[Equation 2.2](#) shows our main specification to estimate the impact of changes in the over-the-counter price of LPG (or cash-back subsidy) on outcomes pertaining to LPG and solid fuel usage:

$$Y_{it} = \alpha + \beta Price_t + \gamma Price_t \times PMUY_i + \mu_i + \delta_{m(t)} + \tau_{y(t)} + \epsilon_{it} . \quad (2.2)$$

Y_{it} denotes two main outcomes of interest – (1) the number of 14.2kg LPG refills purchased and (2) the expenditure on solid fuels, by household i at time t . $Price_t$ denotes our main variable of interest: the log monthly over-the-counter (real) price per kg of LPG (alternatively, the cash-back subsidy per kg), which varies mainly over time.

We are especially interested in estimating the effect of changes in over-the-counter price (or cash-back subsidy) on low-wealth households; according to theory ([Section 2.3](#)), this is where liquidity constraints should have their greatest impact. The binary variable $PMUY_i$ equals 1 if the household has its gas connection under PMUY, a strong indicator of lower socio-economic status, and 0 otherwise. Thus, β captures the change in the monthly number of refills purchased by a regular (non-PMUY) household due to a one percent increase in the over-the-counter price of LPG (or cash-back subsidy), while γ captures the differential effect on the number of refills purchased by a PMUY household.

Our preferred specification ([Equation 2.2](#)) includes household fixed effects (μ_i) to control for time-invariant household-specific factors. This fixed effect subsumes the PMUY indicator and any persistent market-specific variation. Controlling for arbitrary time effects (ν_t) is not possible since our empirical strategy relies primarily on monthly variation in the LPG price p_t ; more generally, there is a tradeoff between absorbing generic time variation and retaining sufficient variation in the LPG price. Bearing in mind the tradeoff, we take a middle ground by including month and year effects in some specifications, to control flexibly

²¹The PLFS data are not representative at the district level, and the rural MP sample is relatively small, especially within occupation-caste-month cells. Hence we compute cash-on-hand at the all-India level.

for seasonality and time trends: $\delta_{m(t)}$ and $\tau_{y(t)}$, where $m(t)$ and $y(t)$ refer to the month and year, respectively, of month x year t . The error term is captured by ϵ_{it} , and we allow for clustering at the market-time level.

For causal interpretation of our β and γ estimates, the key assumption is that the over-the-counter price varies exogenously, independently of local supply and demand factors.²² That local factors do not affect the price is virtually certain, since the over-the-counter price of LPG follows the international market price for LPG. Even national-level factors generally have a non-negligible effect; India can be considered a price-taker for petroleum products in the world market. To interpret the price impacts as liquidity effects, the further necessary assumption is that the LPG price only impacts LPG purchases through the liquidity channel. This also is likely, first and foremost since the price channel is ruled out by the publicly observable constancy of the net price. An income or wealth effect on LPG purchases is also possible, if LPG price changes impact or reflect general macroeconomic activity, and thus real wealth or incomes. This too appears implausible, as the LPG price does not go to extremes over the sample period (see [Figure 2.1](#)), and anyway is not seen as a significant driver of the Indian business cycle or consumer price index. Nonetheless, we variously control for time effects and macroeconomic variables like unemployment (in [Section 2.8.2](#)) to bolster this assumption.

2.6 Results

2.6.1 Impact of LPG refill price and subsidy on refill purchase

We first analyze how changes in the over-the-counter price (or cash-back subsidy) affect LPG refill purchases. Recall that since the net price (post-subsidy) is not changing, this is interpretable as the effect of liquidity requirements on participation in the program. [Table 2.3](#) shows our main results. Negative effects of the price and cash-back subsidy on non-PMUY household purchases are estimated, though the magnitude and statistical significance depend on whether year FE are included (even vs odd columns). But for PMUY (i.e., low-income) households, differential and total purchases are highly sensitive to changes in over-the-counter prices. In the specification controlling for year FE (Column 2, Panel A), the full-sample coefficients suggest that a 1% increase in the over-the-counter price leads PMUY households to purchase 0.0039 fewer LPG refills per month. Given a mean of 0.25 refills per month purchased by PMUY households, this translates to a roughly 15% decrease in LPG purchases from a 10% increase in the over-the-counter price.²³ A very similar pattern emerges in the rural and urban samples of our administrative data (Panels B and C in [Table 2.3](#)). If anything, the effects are more pronounced in rural areas, perhaps because liquidity constraints are generally tighter in rural areas, or because PMUY households in rural areas are poorer. Columns 3 and 4 show consistent results when we estimate the impact of changes in cash-back subsidy instead of the over-the-counter price. The magnitudes are smaller as expected, since gross price and subsidy co-move essentially one-for-one in levels,

²²This guarantees the same for the cash-back subsidy, since it is nearly perfectly correlated with the price.

²³The standard deviation of over-the-counter price is about 12% of its mean; see [Appendix Table A1](#).

and any absolute change is a larger percentage change in the subsidy.²⁴

Next, Table 2.4 shows that results are similar when we use the much smaller sample derived from the rural household survey matched to administrative data. A key difference, compared to the most relevant results of Table 2.3 in Panel B (rural sample), is that the effect for non-PMUY households is larger and more statistically significant. One plausible explanation is that non-PMUY households in the survey sample are relatively poorer than the non-PMUY rural households in the administrative data; the gap in the average monthly LPG refill purchase between the two samples, as shown in Tables 2.1 and 2.2, lends some credence to this possibility. Effects on our focal PMUY households are quite consistent with Table 2.3 results, if slightly smaller.²⁵

In summary, the estimated elasticity of LPG demand to liquidity requirements – captured by the over-the-counter price or cash-back subsidy *holding subsidized price constant* – is substantial for poorer households. Hence, the usage of LPG for cooking is likely to fall in these households with an increase in the over-the-counter LPG price, in spite of the subsidy program designed to keep the effective price basically constant for consumers. To contextualize these results, note the following. First, an increase in the over-the-counter price raises the liquidity hurdle regardless of whether the subsidy responds; so changes in upfront price in general could lead to simultaneous changes in both liquidity requirements and net price. What is unique about this setting is that the subsidy moves in tandem with the over-the-counter price, neutralizing the price effect and isolating the liquidity effect. Second, overall household income levels are important. Our study area is located in a state that has low rural household income. In contrast, using data from one of the richest states in India, Jeuland et al. (2023) estimate the price elasticity of LPG demand to be 1 among PMUY households. Third, the biomass alternatives to LPG are readily accessible in rural and semi-urban areas; hence the cost that households face to switch away from LPG temporarily may not be high.

Do the liquidity constraints driving this behavior operate mostly on the extensive or the intensive margin? That is, does it matter how long households have to part with their money before reimbursement, or does simply having to pay upfront drive the result, however fast or slow is the reimbursement? To understand this, we explore the length of the reimbursement delay as an intermediating factor. We measure delay – day of reimbursement minus day of purchase – as the median across delays for consumer transactions within a given month-year from BPCL (the only OMC with available delay data); see Appendix Figure A1. As shown in Appendix Table A5, purchases are generally more responsive to the over-the-counter price the *longer* the delay in reimbursement is. For example, based on Column 2 estimates both non-PMUY and PMUY households respond more negatively to a one percent increase in price by about -0.085 refills per monthly median day of delay.²⁶ Given a standard deviation

²⁴Thus, their effect in levels rather than logs should be essentially identical; we verify this in Appendix Table A4. Specifically, our estimates suggest a 0.01 reduction in LPG refills purchased by PMUY households from a one INR per kg increase in over-the-counter price or cash-back subsidy (Columns 2 and 4 of Panel A, Appendix Table A4). Thus a one standard deviation increase in the per kg price, i.e., INR 4.4 (or about INR 62 for a 14.2 kg cylinder, see Appendix Table A1), lowers PMUY refill purchases by about 0.044, about 17% of the PMUY mean LPG refill of 0.25.

²⁵The difference could be due in part to the fact that PMUY households in the survey sample appear slightly richer than those in the larger sample.

²⁶Specifically, the estimated response in refills purchase to a one percent price increase is $-0.023 - 0.084 \times$

of median delay days across months of about 2/3, a standard deviation increase in delay increases the magnitude of the price impact on PMUY households by about 19% of the impact at the mean ($(2/3) \times (0.086/0.296)$).

While delay is not necessarily exogenously varying, the results are suggestive that the intensive margin matters, i.e. that households are sensitive to the length of time they have to part with the expected subsidy.²⁷ Parting with cash for longer, even a few days, may indeed be harder.

2.6.2 Impact of LPG refill price on solid fuel expenditure

Solid fuels, particularly firewood and dung, are regularly used for cooking in rural India. A major part of the motivation behind the LPG subsidy policy and the PMUY program is precisely to discourage solid fuel usage because of its negative health and pollution externalities (Barua and Agarwalla, 2018; Kar et al., 2019) and to encourage energy transition to a clean fuel. In our survey data, although 70% of households have an LPG account, 75% report using either firewood or dung for cooking in the previous month. Since solid fuels and LPG are substitutes, an increase in the LPG over-the-counter price that lowers LPG purchases is likely to affect solid fuel usage.²⁸

We use two measures of household expenditure on solid fuels - explicit and implicit. Rural households often collect (e.g. firewood) or make (e.g. cow dung cakes) solid fuels themselves, in addition to buying them on the market. Explicit firewood and dung expenditure per month by the household is reported in the survey, as is the time spent on solid fuel collection. “Implicit” expenditure is constructed by multiplying the total number of hours the household spent collecting/making the solid fuel in a typical week (in the previous month) by four and by the minimum hourly wage in MP for unskilled labor (i.e., INR 35 per hour $\times$ hours spent on fuel collection/making per week $\times$ 4) to get a monthly estimate of the implicit expense on solid fuels.

Table 2.5 shows the effect of an increase in over-the-counter price on explicit (Columns 1-2), implicit (Columns 3-4) and total (explicit + implicit, Columns 5-6) monthly solid fuel expenditure by the household. The results document substitution toward solid fuels when the LPG over-the-counter price or subsidy rises. This is true for both explicit and implicit expenditure for all households, but especially for PMUY (i.e. the poorer) households. A ten percent increase in the over-the-counter price leads to INR 25 more explicit expenditure (12% of the unconditional mean) and INR 163 more implicit expenditure (15% of the unconditional mean) for PMUY households. Including year fixed effects, which is not preferred due to having only five months, leads to noisier but qualitatively similar results (see Appendix Table A21). Thus, upfront price increases with the net-of-subsidy price held fixed impact substantially not only LPG purchases but also usage of solid fuels that the program is seeking

Delay for non-PMUY and $-0.296 - 0.086 * \text{Delay}$ for PMUY, where *Delay* is demeaned. Thus the effect at mean delay is close to zero for non-PMUY and about -0.3 for PMUY.

²⁷Drawing conclusions about the extensive margin from these results would require extrapolation well outside the range of the data.

²⁸We asked the household primary cook to list all the fuels used in preparing the last meal. More than half of the households reported using solid fuels exclusively, even though more than two-thirds had LPG connections. Only 29% of households report using LPG exclusively in preparing the last meal.

to discourage.

2.6.3 Robustness

Our results are robust to a number of checks. First, to address concerns regarding the potential endogeneity of the domestic LPG price, we run 2SLS estimation using 1-month lags of the international LPG price as an instrument for the over-the-counter price (and cash-back subsidy). Our results are virtually unchanged for both the administrative data (Appendix Table A6) and the household survey data (Appendix Table A7) analyses. Second, we alternatively cluster-bootstrap the standard errors at the LPG market level (26 markets), instead of market-month as in our standard analysis, and show that our findings do not change in Appendix Table A8 (administrative data) and Appendix Table A9 (household survey data).

Finally, our results are robust to an alternative model of Poisson pseudo-maximum likelihood estimation (Appendix Table A10). We also recode consumer refills as missing (instead of 0 as in our main analysis) in months where the consumer had not yet obtained an LPG account, and rerun our analysis with an unbalanced panel in Section 2.B. The results are qualitatively unchanged, though magnitudes decline somewhat.

To summarize, our results so far indicate that poor households are extremely sensitive to increases in the market price of LPG despite a public subsidy program that holds the subsidized price essentially constant. We hypothesize that the take-up of this clean cooking fuel is adversely affected due to liquidity constraints faced by low-income households.

2.7 Impact on health

In this section, we focus on the health rationale for the LPG subsidy program. While an overarching objective of the LPG subsidy program is to protect consumers from increases in the price of LPG, our analysis shows that the program falls short of achieving this goal. Poor consumers are less likely to purchase LPG at higher prices and switch instead to biomass-based solid fuels. Burning solid fuels, in turn, increases indoor air pollution and thus leads to adverse health effects, especially for young children and pregnant women. What are the health consequences of households' sensitivity to LPG price fluctuations despite the active subsidy program?

Relevant data are available from the Demographic and Health Survey or the National Family Health Survey (2019-21) of India, summarized in Appendix Table A11 for the entire country. Combining this with our administrative data on LPG refill prices, we estimate the effects of over-the-counter prices of LPG on infant and child health, using Equation 2.3:

$$Y_{it}^{ds} = \alpha + \beta Price_{it} + \mu Child_i + \gamma_d + \delta_{m(t)} + \tau_{y(t)} + \phi^s * t + \epsilon_{it}^{ds} . \quad (2.3)$$

For our infant health analysis, Y_{it}^{ds} denotes the birth weight (in kilograms) of infant i , born in district d , state s , and conceived in month (m)-year (y), t . The sample includes pregnancies conceived within one year prior to the DHS interview date. $Price_{it}$ in the infant health analysis is the log of the average over-the-counter LPG refill price during the mother's

gestation period (including less than full-term pregnancies) for birth i .²⁹ For the child (under 5 years of age) health analysis, Y_{it}^{ds} denotes the other two main outcomes of interest – an indicator variable that equals 1 if there is an incidence of (1) cough and (2) rapid breathing during the previous two weeks in child i in district d , state s , interviewed in month-year t . $Price_{it}$ for child health outcomes denotes the log of over-the-counter LPG price at the time of the survey interview. We include infant- or child-specific characteristics along with mother and household controls in $Child_i$ (e.g. gender, birth order, whether child is part of a multiple birth, quadratics in age of child (for child health outcomes) and age of mother, mother’s education, religion, and caste), as well as district-level unobservables (γ_d), month ($\delta_{t(m)}$), year ($\tau_{t(y)}$), and state-specific linear time trends. The error term is captured by ϵ_{it}^{ds} and is clustered by time (of conception for infant weight outcomes, of interview for child outcomes). We report standard p -values of the estimates and p -values after correcting for multiple hypothesis testing for all the outcomes and samples. Identification of health effects relies on quasi-random variation in LPG price over different gestational periods in the infant health case, and across different household interview dates in the child health case.

Columns 1 and 2 (Panel A) in Table 2.6 show the effect on birth weight of the over-the-counter refill price (or subsidy), in log terms after averaging over the gestation period. Though not quite significant at conventional levels, a 10% increase in the over-the-counter LPG refill price during gestation is associated with a 44 gram reduction in infant birth weight, 1.6% of the mean. The more immediate effects of LPG price increases on child health are seen in Columns 3-6 (Panel A) of Table 2.6. A 10% increase in the over-the-counter LPG refill price is associated with a 0.85 percentage point increase in the likelihood of reported coughing during the previous two weeks (7% of the mean) and a 1.01 percentage point increase in the likelihood of reported rapid breathing (16% of the mean) among children under five (Columns 3 and 5, respectively). Table 2.6 further suggests that these results are stronger for rural households (Panel B), especially for the rapid breathing outcome, consistent with our main results.

Results are weaker but similar when we run household weighted regression analysis, as shown in Appendix Table A12. They are also mostly robust to cluster-bootstrapping the standard errors, as shown in Appendix Table A13, which may be preferred given the small number of month-year clusters for ‘cough’ and ‘rapid breathing’ outcomes.³⁰ Furthermore, we run a falsification exercise to rule out spurious findings by analyzing the effect of LPG refill prices on child height-for-age and weight-for-age Z-scores, outcomes that are unlikely to be significantly impacted by contemporaneous exposure to smoke inhalation. We find no detectable effects, as reported in Appendix Table A14.

Overall, our estimated health effects are substantial, and likely a lower bound for the impacts on LPG-using low-income households.³¹ These findings are unsurprising given our

²⁹We do not have refill price data for all districts of India. Instead, we impute our refill prices from the district of Indore (rural and urban) to all districts in our India-wide health analysis here. There are negligible differences in LPG price levels across districts, and marginal changes in prices occur more or less in tandem across India (Appendix Figure A3). The alternative strategy of restricting the DHS sample to only MP, Indore’s state, would result in a very small sample.

³⁰We have only 18 clusters for month-year of interview in the DHS sample.

³¹The DHS does not provide information on households’ LPG usage, PMUY status, or income/expenditure. Thus, we cannot condition the DHS sample on whether the household uses LPG.

earlier results that higher upfront LPG prices, with the subsidized price held fixed, decrease LPG purchases and push households to more household-polluting cooking fuel. Further, they demonstrate that these liquidity effects on consumption can have real health consequences.

We turn next to further evidence that this responsiveness to LPG gross price with net price fixed is due to liquidity constraints.

2.8 Mechanism - Liquidity constraints

In this Section, we bolster the liquidity interpretation of our results by showing robustness to alternative measures and sources of liquidity, including direct evidence from exogenous weather shocks, and by ruling out alternative non-liquidity based explanations. This increases confidence that the sensitivity of poor households to fluctuations in LPG prices despite the subsidy program is attributable to liquidity constraints.

2.8.1 Alternative measures and sources of liquidity

First, given that these constraints are likely more frequently binding for poor households, we use our Survey data to test for heterogeneous impacts by alternative household proxies for wealth. [Table 2.7](#) shows that the observed responsiveness to over-the-counter LPG price depends significantly on household wealth, measured along multiple dimensions: an asset index (Columns 1 and 2), a dummy for belonging to General or upper caste category (0 if SC/ST or OBC) (Columns 3 and 4), and a dummy for household head being a casual laborer (Columns 5 and 6).

The negative impact of higher over-the-counter LPG prices on refill purchases is significantly larger for households with lower asset holdings, as shown in Column 1 of [Table 2.7](#). A two standard-deviation increase in the asset index (about 1.5, see [Table 2.2](#)) lowers the impact magnitude by 60% of the mean effect. This pattern suggests that poorer households are more price-sensitive, likely due to greater liquidity constraints. [Appendix Table A15](#) breaks down the analysis by asset quartiles. We find consistent results, i.e. the observed effect is significantly higher for households with below median assets and in the lower asset quartiles.

Furthermore, LPG refill purchases by higher caste (i.e., general caste) households, who are economically better-off on average, decline less when the over-the-counter price rises than do purchases of lower caste households (i.e. SC, ST or OBC), as shown by the interaction term in Columns 3 and 4 in [Table 2.7](#) (p -value<0.10). Households that are engaged in more precarious sources of livelihood, e.g. daily wage earners from casual labor, are also likely to be poorer and experience more frequently binding liquidity constraints. The estimated coefficients in Column 5 therefore support the liquidity interpretation, as the negative interaction coefficient (p -value<0.10) implies that casual laborers' LPG refill purchases are more responsive to the over-the-counter price.

As [Appendix Table A2](#) shows, the PMUY program with its one-time start-up subsidy to obtain an LPG account effectively targeted less well-off households, by the measures used here ([Table 2.7](#)). Our results thus indicate that lower-wealth households, proxied by

a number of measures, differentially reduce take-up of LPG in response to increases in the over-the-counter price, even though the LPG net price is held fixed.

Second, we construct a direct measure of household liquidity – cash-on-hand in a month – to test whether households with higher monthly cash on hand are less affected by an increase in the over-the-counter refill price. Using the information on monthly income and expenditure in the PLFS for rural India (see Section 2.4.3), we predict average cash-on-hand (income minus expenditure) at the caste-occupation-month level as a proxy for household liquidity.³² These measures are merged into the household Survey data by matching cash-on-hand to household at the caste-occupation-month level. Table 2.8 shows the association of our measure of log monthly cash-on-hand and its interaction with log over-the-counter price, with monthly refill purchases of rural households in our sample between November 2017 to December 2019. (See Table A3 for summary statistics.) Unsurprisingly, more cash on hand is associated with higher LPG refill purchases (Column 1, Table 2.8). The coefficients on over-the-counter price remain negative across all specifications (Columns 2 and 3). Importantly, the positive coefficient on the interaction of over-the-counter price with cash on hand (Column 3) indicates that more cash on hand significantly attenuates the negative effect of an over-the-counter price increase on LPG purchases, consistent with liquidity constraints driving the observed behavior.

Third, we analyze a plausibly exogenous driver of liquidity: favorable rainfall during the agricultural sowing season. High rainfall during a sowing season results in better crop yields and therefore generally greater liquidity during the associated harvest season. Thus if liquidity constraints drive the results, good growing season rainfall should attenuate the negative effect of price on refill purchases in the harvest season.

To test this mechanism, we examine how refill consumption during harvest months responds to LPG prices, conditional on rainfall in the preceding sowing season. Note that the two main crops in Indore are soy (a monsoon or rainy season crop, sown during June–September and harvested in October–December) and wheat (a winter crop sown in October–January and harvested in February–April). We restrict the analysis to harvest-season months (February–April and October–December) and define an indicator variable ‘High Rain Harvest’ that equals 1 if any month in the preceding sowing season received normal, excess, or highly excess rainfall, based on IMD classifications ([Indore District Website](#); [ICRISAT](#)).³³

Results in Table 2.9 show that while for non-PMUY households the expected sensitivity of LPG refill purchases to prices after low rainfall does not hold, it does so for the PMUY (i.e., low-income) households, who are more likely to face binding liquidity constraints. For

³²We use 4 occupation categories (salaried, self-employed on-farm, self-employed off-farm and casual labor) and 4 caste groups (SC, ST, OBC and General or higher-caste categories). Ideally we would construct this measure for each household in our Survey, but the needed income and expenditure data are not available. The monthly variation in our PLFS measure of household liquidity aligns with the harvest cycle of the three major crops grown in Indore - wheat, chickpea and soy. As depicted in Figure A2, average cash on hand for rural households exhibits peaks during the harvest months for wheat and chickpea (Feb - May), and soy (Oct - Dec).

³³The IMD (India Meteorological Department) categorizes rainfall in each month relative to long-period averages for each district. Rainfall $\geq 60\%$ is classified as ‘Large Excess’, 20% to 59% as ‘Excess’, -19% to $+19\%$ as ‘Normal’, -59% to -20% as ‘Deficient’, and -99% to -60% as ‘Large Deficient’. 57% of harvest months in our study period are coded as ‘High Rain Harvest’. Source: [IMD](#).

the PMUY households total and relative purchases are much less sensitive to changes in over-the-counter prices after a good rainfall. In the full-sample specification controlling for year fixed effects (Column 2, Panel A), a 1% increase in the over-the-counter price leads PMUY households to purchase 0.004 fewer LPG refills in harvest months after a good rainfall but 0.010 fewer refills after low rainfall. Compared to non-PMUY households, the PMUY households respond more negatively to a 1% refill price increase – 0.003 fewer refills after a good rainfall compared to 0.012 fewer refills after poor rainfall, implying four times lower consumption. The results suggest that better rainfall during the sowing season, by boosting liquidity, mitigates the adverse impact of higher LPG prices on refill consumption by PMUY households during the harvest season.

Since it is plausible that rainfall shocks exogenously affect agricultural households’ cash flow, these results provide helpful evidence in favor of liquidity constraints as the mechanism that explains our main results.³⁴ However, we do acknowledge that our empirical analysis here may suffer from low power since we restrict the sample to only the harvest months (14 of 26 months in our main data) and that high rainfall is seasonal (i.e., perfectly correlated in our sample with the monsoon crops).

2.8.2 Alternative explanations

In this section, we assess and rule out mechanisms other than liquidity constraints that could explain our results on sensitivity to price.

Macro-level factors – General inflation and unemployment

Increases in the LPG price may be causing, or symptomatic of, a general increase in prices. Inflation could lower households’ real wealth, leading them to reduce consumption of LPG by cooking less and/or substituting toward biomass-based solid fuels; such real impacts are likely regressive, affecting poorer households more. Indeed, in the theory, even without liquidity constraints, LPG purchases fall with a drop in wealth past some likely low threshold (Section 2.3). Appendix Table A16 shows the correlation of domestic LPG refill price with each component of the Consumer Price Index (CPI) in our study period. Other than a marginally significant coefficient ($p < 0.10$) on the ‘Fuel and Light’ Price Index, all other components (including overall CPI) are uncorrelated with the changes in the over-the-counter price of LPG refill. Thus LPG price changes do not appear to be simply proxying for overall inflation. Further, since our data capture high frequency, generally small variations in the LPG refill price (standard deviation of INR 62.45 per refill during our study period), LPG price fluctuations are unlikely to be causing inflationary trends. We thus rule out the story that price changes actually proxy for general inflation-based wealth effects as an explanation of our results.

A slightly different macroeconomic explanation for our results could be that the international price of gas acts as a supply shock that affects real macroeconomic activity, and thus households’ economic prospects. This would again lead to LPG purchases responding to the LPG price via a wealth effect rather than a price or liquidity effect.

³⁴Rainfall can also affect wealth or net worth at harvest, but absent liquidity constraints, a wealth shock should not alter the effect of over-the-counter price on purchases, since the net-of-subsidy price is held fixed.

However, this seems *a priori* unlikely, since the Indian macroeconomy is not generally significantly impacted by the international LPG price, and the price changes in our data are moderate. To further cast doubt on this story, we control for the monthly unemployment rate in the specification used for our main results (Table 2.3). As shown in Appendix Table A20, the effects of over-the-counter price and cash-back subsidy remain very similar for PMUY households after controlling for the monthly unemployment rate.

Subsidy salience

A potential alternative explanation for the observed effects of over-the-counter refill price is subsidy salience. If households are not fully aware of the subsidy and the fact that net-of-subsidy price is held constant, their response may be due to a pure price effect rather than a liquidity effect. Afridi et al. (2021) find a 13% increase in the monthly demand for LPG refills when information on the refill subsidy is bundled with improving awareness of the adverse health effects of using solid fuels. Chetty et al. (2009) show that the grocery demand goes down by 8% when tax-inclusive prices are displayed. In our setting, lack of salience of the rebate could lead consumers to mistakenly respond to over-the-counter price changes as if the net-of-subsidy price is also changing.

We explore this alternative explanation using survey data on households' understanding of the LPG subsidy program. Household knowledge is measured using self-reported responses to five statements about the program in the second round of household survey (October - December 2019). The household was asked to agree or disagree with each statement, some of which were true and others false in random order.³⁵ The variable for each statement is coded 1 iff the respondent correctly responds to the statement; thus each variable is an indicator of the household having accurate knowledge of the LPG subsidy program. Appendix Table A17 shows estimates of heterogeneous price impacts by each household knowledge indicator. We find no significant difference in effect of over-the-counter price based on different knowledge of the program for non-PMUY and PMUY households, with one exception: non-PMUY households who state they believe the subsidy is only for PMUY households are more responsive to over-the-counter price than those who state that the subsidy is for all. However, these results offer no supporting evidence that subsidy salience is driving our main results.

Trust in government

A related possible explanation of our findings is that households do not trust the government to transfer the cash-back subsidy into their bank accounts after they purchase the LPG refill. Consequently, they respond to the market price as the final price, or to a probability-weighted average of the market price and the net-of-subsidy price. Our observed results could then be interpreted as a price effect rather than a liquidity effect.

³⁵The survey collects self-reported information on respondents' knowledge of the program. Conditional on having an LPG account, about 33% of respondents did not agree with the statement that the government deposits a subsidy in their bank account after they purchase an LPG cylinder, and a little over half agreed that their net out-of-pocket expenditure was less than the over-the-counter market price of an LPG refill. About 32% of respondents stated that only PMUY customers are eligible to receive the subsidy. The specific statements are listed in the notes to Appendix Table A17.

A priori this seems unlikely: the DBT scheme was implemented across the country much before our study period, in 2014, while the PMUY program began in 2016. The fact that the government was indeed depositing the subsidy into millions of consumers’ bank accounts was well-established.

Nevertheless, we address this issue in two ways. First, we directly asked our survey households (at baseline) whether they trusted information provided by the government-appointed accredited social health activist (ASHA) of their village; about 83% respond affirmatively. (This is a community health worker responsible for encouraging households to immunize newborns and ensure ante and post-natal care for mothers, among other health information and assistance.) We test whether the response to over-the-counter price varies by the household’s reported trust in ASHA as a proxy for trust in government. There is no detectable difference by trust level for either non-PMUY or PMUY households in their responsiveness to price, as seen in Appendix [Table A18](#). Second, we analyze a subsample of the administrative data representing relatively experienced LPG users, i.e. only households who had obtained LPG accounts by the beginning of 2017, well before the beginning of our study period (November 2017). These households are more likely to be aware of the reliability of the cash-back subsidy design, since they have had longer exposure to the historically reliable DBT scheme. If our main results are due to lack of trust in government, then our coefficients of interest should be smaller or insignificant for this sample. The estimates, however, are similar in both size and significance to those of our main results; see Appendix [Table A19](#). Lack of trust in the government is, therefore, unlikely to be the primary mechanism.

2.9 Conclusion

Our results show that clean fuel usage decreases and solid fuel usage increases, due to the liquidity constraints amplified by delayed transfer of the subsidy amount. The associated health and environmental externalities, especially for the women and children who spend more time indoors and near solid-fuel cooking stoves, are likely to be significant, although our analysis is limited (by lack of data) to health outcomes of newborns and children. Overall, our analysis of India’s cooking gas subsidy program highlights the interaction of the cash-back subsidy design with households’ liquidity constraints, which acts counter to the policy intent behind subsidizing clean fuel take-up of low-income households.


More broadly, these findings indicate that the administrative design of subsidy and conditional cash transfer programs, especially the *timing* of transfers, matters. While cash transfers are replacing in-kind transfer programs across the world, it is important to understand the welfare implications of their design. India, alone, runs more than 300 public programs where the government provides cash benefits to citizens through direct bank transfers. Policy reforms aiming to improve the efficiency of welfare delivery by replacing in-kind or point-of-sale subsidies with delayed bank transfers may impose an unintended regressive cost on households, when any time lag embedded in a cash transfer program interacts with households’ liquidity constraints.

Our findings carry implications for designing subsidy transfers that soften or eliminate this tradeoff. In the context of the program we study, for example, the results suggest that

the length of reimbursement delay matters, so minimizing delay subject to clean verification is an implicit recommendation. Beyond that, allowing at least the poorer consumers to use program-backed credit to purchase the good, at least the subsidized portion, could alleviate the concerns here. Alternatively, a digital banking approach that provides electronic vouchers to be used at the point of sale can help low-income households.³⁶ Another approach could be to encourage purchase of the subsidized good in smaller amounts (and at lower upfront price) or provide a higher subsidy for savings- and liquidity-constrained households ([Abubakari et al. \(2024\)](#)). While these alternative program designs may come with new tradeoffs (viz. increased administrative costs), this paper advances the discussion by illuminating a drawback of existing programs and pointing out directions for policy improvements.


³⁶[Banerjee et al. \(2023\)](#) show that Indonesia’s electronic voucher-for-food program provides better targeting and is cost-effective. A similar electronic voucher system may help alleviate liquidity-related concerns in cash-back subsidy designs while avoiding some of the problems associated with point-of-sale subsidy designs. An alternative could be to encourage LPG distributors to extend credit to households. However, in the presence of strong religion and caste identities, such an approach may have its limitations ([Alhorr et al., 2025](#)).

Figure 2.1: Over-the-counter Price, Effective Price and Subsidy on LPG
(November 2017 - December 2019)



Notes: LPG prices obtained from the OMC administrative dataset. Over-the-counter price of LPG is the price paid by consumers upfront to purchase LPG. Effective price is over-the-counter price minus the cash-back subsidy to LPG to consumers. The cash-back subsidy floated in tandem with the over-the-counter price. The Y-axis shows all three prices in real (2012) INR for 1 kg of LPG. The X-axis shows the corresponding month, from November 2017 to December 2019 (26 months). All prices are averaged across all markets and dealers.

Figure 2.2: Correlation between International LPG Price and Over-the-counter Domestic LPG Price (November 2017 - December 2019)



Notes: Saudi Aramco LPG price (defined as price of butane (60%) and price of propane (40%)) obtained from [here](#). The domestic price of LPG obtained from the OMC administrative dataset. Each observation on the Y-axis is the monthly price of Saudi Armco (lagged by 1 period). The X-axis shows the domestic price of LPG spanning 26 months from November 2017 to December 2019. All prices are in nominal INR per KG of LPG. The dashed line shows the coefficient of correlation between the lagged international price and domestic market price, which is 0.99 ($p\text{-val} = 0.0004$).




Figure 2.3: Households with wealth higher than $\widehat{W}(q)$ purchase fuel, while those with lower wealth do not. The wealth cutoff would be the dash-dotted line absent any liquidity constraints.

Table 2.1: LPG Consumer Characteristics (Administrative Data)

	Obs	Mean	SD	Minimum	Maximum
Panel A: Attributes					
PMUY	911,454	0.074	0.262	0.000	1.000
Female	911,099	0.320	0.466	0.000	1.000
Rural	908,650	0.572	0.495	0.000	1.000
Panel B: LPG Refills					
All	23,697,794	0.578	0.638	0.000	16.000
PMUY	1,754,298	0.252	0.495	0.000	12.000
Non-PMUY	21,943,496	0.604	0.640	0.000	16.000
Rural	13,522,165	0.576	0.643	0.000	16.000
PMUY	1,242,540	0.259	0.503	0.000	12.000
Non-PMUY	12,279,625	0.608	0.646	0.000	16.000
Urban	10,102,717	0.580	0.631	0.000	13.000
PMUY	510,276	0.235	0.475	0.000	11.000
Non-PMUY	9,592,441	0.598	0.633	0.000	13.000

Notes: This table summarizes the OMCs administrative data from Nov 2017 to Dec 2019. **Panel A** summarizes the demographics of all consumers in the district of Indore. PMUY is a dummy which equals 1 if the consumer has a PMUY LPG connection, and 0 otherwise. Female is a dummy which equals 1 if the LPG account is in the name of a woman and 0 otherwise. Rural is a dummy which equals 1 if the consumer resides in a rural area. Some observations could not be classified as female or rural, hence the total observations vary. **Panel B** summarizes LPG refill transactions from Nov 2017 to Dec 2019 overall, by rural/urban, and PMUY status of the LPG consumers.

Table 2.2: Household Characteristics (Survey Data)

	Obs	Mean	SD	Minimum	Maximum
Panel A: Household Attributes (Baseline)					
PMUY	2,091	0.391	0.488	0.000	1.000
Asset Index	2,091	1.620	0.748	-0.150	3.965
General Caste	2,091	0.150	0.357	0.000	1.000
Salaried HH Head	2,091	0.106	0.308	0.000	1.000
Agri Self-Employed HH Head	2,091	0.305	0.461	0.000	1.000
Agri Casual Laborer HH Head	2,091	0.391	0.488	0.000	1.000
Land Owner	2,091	0.522	0.500	0.000	1.000
Land Owner/Leaser	2,091	0.535	0.499	0.000	1.000
HH Head Education	2,091	0.411	0.492	0.000	1.000
Latrine in HH	2,091	0.879	0.326	0.000	1.000
Panel B: Household Solid Fuel					
Baseline & Endline					
Explicit firewood expenditure (Rs.)	4,150	56.791	282.530	0.000	6,000.000
Implicit firewood expenditure (Rs.)	4,150	543.258	833.181	0.000	4,480.000
Explicit dung expenditure (Rs.)	4,150	151.177	470.125	0.000	6,000.000
Implicit dung expenditure (Rs.)	4,150	503.106	626.538	0.000	9,800.000
Panel C: Household Refills Matched to Admin Data					
All Refills	54,366	0.411	0.576	0.000	7.000
PMUY refills	21,268	0.285	0.505	0.000	7.000
Non-PMUY refills	33,098	0.492	0.604	0.000	6.000

Notes: **Panel A** provides information on household characteristics collected in 2018 (at baseline). Asset Index is the first component of a principal component analysis over several indicators measuring the economic status of a household including, ownership of land and farm animals, *pucca* house, and a list of consumer durables. Education of the head of the household is measured by an indicator that takes value one for above primary education and zero otherwise. General Caste is a dummy for household that is neither SC/ST nor OBC category. Salaried, Self-Employed, Casual Laborer, Land Owner, Land Owner/ Leaser are dummies equal to 1 if the household head belongs to the respective category. Latrine in HH is a dummy if the household has a pit or flush toilet built inside house. **Panel B** summarizes the explicit and implicit expenditure on firewood and dung. Implicit expenditure is calculated as total number of hours household spent collecting/making the solid fuel in a week times the minimum hourly wage in Madhya Pradesh for unskilled labor (i.e. INR 35) and then times 4 to get a monthly estimate of implicit expense on solid fuels. Explicit firewood and dung expenditure per month is reported in the survey. **Panel C** summarizes the LPG refill transactions for the sampled households over 26 months by matching them to the administrative records.

Table 2.3: Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills
(Administrative Data)

	(1)	(2)	(3)	(4)
Panel A: Full Sample				
Over-the-counter price	-0.091*** (0.026)	-0.043 (0.034)		
PMUY $\times$ Over-the-counter price	-0.349*** (0.049)	-0.347*** (0.053)		
Cash-back subsidy			-0.025*** (0.007)	-0.007 (0.009)
PMUY $\times$ Cash-back subsidy			-0.110*** (0.012)	-0.110*** (0.013)
Mean of Dependent Var.	0.578	0.578	0.578	0.578
Mean of Dependent Var. for PMUY	0.252	0.252	0.252	0.252
Mean of Dependent Var. for Non-PMUY	0.604	0.604	0.604	0.604
Observations	23,697,794	23,697,794	23,697,794	23,697,794
Panel B: Rural Sample				
Over-the-counter price	-0.098*** (0.026)	-0.046 (0.035)		
PMUY $\times$ Over-the-counter price	-0.374*** (0.041)	-0.372*** (0.046)		
Cash-back subsidy			-0.028*** (0.007)	-0.008 (0.009)
PMUY $\times$ Cash-back subsidy			-0.116*** (0.011)	-0.115*** (0.012)
Observations	13,522,165	13,522,165	13,522,165	13,522,165
Panel C: Urban Sample				
Over-the-counter price	-0.083*** (0.026)	-0.041 (0.034)		
PMUY $\times$ Over-the-counter price	-0.275*** (0.073)	-0.275*** (0.074)		
Cash-back subsidy			-0.023*** (0.007)	-0.006 (0.008)
PMUY $\times$ Cash-back subsidy			-0.094*** (0.018)	-0.093*** (0.018)
Observations	10,102,717	10,102,717	10,102,717	10,102,717
Month FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes
Household FE	Yes	Yes	Yes	Yes

Notes: This table shows the impact on monthly refills of changes in log of “Over-the-counter price” – the price paid upfront per kg of LPG – and log of “Cash-back subsidy” – the subsidy per kg of LPG that the consumers get back after purchase – using the OMC administrative data. Panel A shows the estimates for the full sample, Panel B for the rural sub-sample, and Panel C for the urban sub-sample. The dependent variable is the total monthly purchases of 14.2 kg LPG cylinder refills. The unit of observation is consumer-month-year, leading to a total of 23.7 million observations, corresponding to monthly refills for 26 months of the 911,454 unique included consumers. Refill is coded zero if the consumer does not have a gas connection or did not purchase in a given month. PMUY is a dummy that equals 1 if the consumer has a PMUY connection, 0 if it's a regular gas connection. Standard errors, clustered at the market-month-year level (676 clusters in Full Sample), reported in parentheses. Significance at * 10%, ** 5% and *** 1% level. The sum of the coefficient for price or subsidy and of the coefficient for its interaction with PMUY is significant at 1% in every specification.

Table 2.4: Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills
(Survey Data)

	(1)	(2)	(3)	(4)
Over-the-counter price	-0.172*** (0.033)	-0.082** (0.041)		
PMUY $\times$ Over-the-counter price	-0.255*** (0.051)	-0.254*** (0.051)		
Cash-back subsidy			-0.046*** (0.010)	-0.015 (0.012)
PMUY $\times$ Cash-back subsidy			-0.079*** (0.016)	-0.078*** (0.016)
Mean of Dependent Var.	0.411	0.411	0.411	0.411
Mean of Dependent Var. for PMUY	0.285	0.285	0.285	0.285
Mean of Dependent Var. for Non-PMUY	0.492	0.492	0.492	0.492
Observations	54,366	54,366	54,366	54,366
Month FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes
Household FE	Yes	Yes	Yes	Yes

Notes: This table shows the impact on monthly refills of changes in log of “Over-the-counter price” – the price paid upfront per kg of LPG – and log of “Cash-back subsidy” – the subsidy per kg of LPG that the consumers get back after purchase – using the household survey data. The dependent variable is the total monthly purchases of 14.2 kg LPG cylinder refills. The unit of observation is household-month-year, leading to a total of 54,366 observations, corresponding to monthly refills for 26 months of the 2091 unique included consumers. Refill is coded zero if the consumer does not have a gas connection or did not purchase in a given month. PMUY is a dummy that equals 1 if the consumer has a PMUY connection, 0 if it's a regular gas connection. Standard errors, clustered at the market-month-year level (598 total clusters), reported in parentheses. Significance at * 10%, ** 5% and *** 1% level. The sum of the coefficient for price or subsidy and of the coefficient for its interaction with PMUY is significant at 1% in every specification.

Table 2.5: Impact of Over-the-counter Price and Cash-back Subsidy on Solid Fuel Expenditure (Survey Data)

	Explicit Expenditure		Implicit Expenditure		Total Expenditure	
	(1)	(2)	(3)	(4)	(5)	(6)
Over-the-counter price	78.088** (34.149)		706.871*** (80.369)		784.959*** (89.208)	
PMUY $\times$ Over-the-counter price	181.004** (86.126)		1,002.342*** (140.107)		1,183.346*** (184.085)	
Cash-back subsidy		25.145** (10.659)		230.826*** (21.657)		255.971*** (23.465)
PMUY $\times$ Cash-back subsidy		58.804** (28.121)		329.690*** (46.411)		388.494*** (60.514)
Mean of Dependent Var.	208.975	208.975	1,047.526	1,047.526	1,256.501	1,256.501
Observations	4,118	4,118	4,118	4,118	4,118	4,118
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	No	No	No	No	No	No

Notes: This table shows the impact on monthly expenditure on solid fuels of changes in log of over-the-counter price and log of cash-back subsidy using the survey data. “Over-the-counter price” is the log of the price paid upfront per kg of LPG. “Cash-back subsidy” is the log of the subsidy per kg of LPG that the consumers get back after purchase. Dependent variable in columns 1 and 2 is the explicit expenditure on solid fuel, which is the amount the household reports to have spent out-of-pocket on firewood and dung (combined) in the previous month. Dependent variable in columns 3 and 4 is the implicit expenditure on solid fuel which is defined as number of times person responsible in a households went to collect solid fuel (firewood and dung) in the previous month multiplied by number of hours spent in collection multiplied by the hourly wage of INR 35, to get implicit expenditure per month. Columns 5 and 6 show the impact on the sum of implicit and explicit expenditure on solid fuel. All specifications include month and household fixed effects. Unit of observation is consumer-month-year. In the survey, each household is asked information on solid fuel collection for only one month, hence, this dataset is a panel of 2-month observations (baseline survey in 2018 and endline survey in 2019) for each LPG-using household. PMUY is a dummy, which is 1 if the consumer has a PMUY connection and 0 if it’s a regular gas connection. Standard errors, clustered at the market-month-year level are reported in parentheses. Significant at * 10%, ** 5% and ***1% level.

Table 2.6: Impact of Over-the-counter Price and Cash-back Subsidy on Birth and Post-birth Child Health Outcomes

	Birth Weight		Cough		Rapid Breathing	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full Sample						
Over-the-counter price	-0.463		0.089*		0.106**	
	(0.333)		(0.044)		(0.040)	
Cash-back subsidy		-0.021		0.005		0.029***
		(0.026)		(0.010)		(0.005)
Mean of Dependent Var.	2.784	2.784	0.129	0.129	0.065	0.065
Observations	16,801	16,801	223,764	223,764	223,862	223,862
FDR-adjusted p -value	[0.102]	[0.65]	[0.066]	[0.65]	[0.057]	[0.001]
Panel B: Rural Sample						
Over-the-counter price	-0.497		0.085		0.135**	
	(0.323)		(0.054)		(0.056)	
Cash-back subsidy		-0.025		0.005		0.030***
		(0.027)		(0.010)		(0.009)
Mean of Dependent Var.	2.780	2.780	0.130	0.130	0.066	0.066
Observations	13,566	13,566	177,951	177,951	178,035	178,035
FDR-adjusted p -value	[0.099]	[0.557]	[0.099]	[0.677]	[0.094]	[0.015]
Panel C: Urban Sample						
Over-the-counter price	-0.479		0.083		0.081**	
	(0.574)		(0.070)		(0.038)	
Cash-back subsidy		-0.013		-0.003		0.054**
		(0.040)		(0.029)		(0.019)
Mean of Dependent Var.	2.800	2.800	0.124	0.124	0.059	0.059
Observations	3,235	3,235	45,813	45,813	45,827	45,827
FDR-adjusted p -value	[0.379]	[1]	[0.339]	[1]	[0.164]	[0.035]
Mother and Child Controls	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
State $\times$ Month-Year Trend	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows the impact of changes in log of over-the-counter LPG price and log of cash-back subsidy on child health outcomes. Over-the-counter prices and Cash-back subsidy is from the administrative dataset for BPCL Indore (June 2018 - May 2021); health outcomes from the Demographic and Health Survey (DHS) Round 5, 2019-21 (previous 12 months is the reference period for interviews conducted during June 2019 to May 2021) for both rural and urban sample. 'Birth Weight' is weight at birth measured in kilograms. LPG price and subsidy for birth weight is the average price and subsidy during the gestation period, with the sample restricted to births whose conception was within a year prior to the DHS survey (pregnancies conceived from June 2018 to Dec 2020). Gestational length includes both month of conception and month of birth/outcome. 'Cough' is a dummy that equals 1 if the child (under 5 years of age) has been ill with cough in last two weeks and 'Rapid breathing' equals 1 if the child (under 5 years of age) suffered from rapid breathing in last two weeks. LPG price and subsidy for cough and rapid breathing is as of the month of the interview. Time controls correspond to conception month-year of child in columns 1 and 2 and interview month-year of household in columns 3 - 6. Child controls include gender, whether the child is part of a multiple birth and birth order of the child. For cough and rapid breathing, age of child (and age squared) are additional controls. Mother controls include mother's age at birth (including age squared), mother's education level, if household is located in rural/urban area, religion, caste and wealth index. State $\times$ Month-Year trend is 'state $\times$ c.month-year' of conception/interview. p -values of coefficients reported for Multiple Hypothesis Testing using the Anderson's sharpened q -values to control for the false discovery rate. Standard errors clustered at month-year of conception in columns 1 and 2 and at month-year of interview in columns 3 - 6. Significant at * 10%, ** 5% and ***1% level.

Table 2.7: Impact of Over-the-counter Price on Monthly Refills: Heterogeneity by Household Attributes (Survey Data)

	Z					
	Asset Index		General Caste		Casual Laborer	
	(1)	(2)	(3)	(4)	(5)	(6)
Over-the-counter price	-0.300*** (0.057)		-0.197*** (0.036)		-0.152*** (0.038)	
Over-the-counter price $\times$ Z	0.073*** (0.026)		0.103* (0.058)		-0.076* (0.039)	
Cash-back subsidy		-0.079*** (0.018)		-0.049*** (0.011)		-0.037*** (0.012)
Cash-back subsidy $\times$ Z		0.021*** (0.008)		0.031* (0.018)		-0.019 (0.012)
Mean of Dependent Var.	0.411	0.411	0.411	0.411	0.411	0.411
Observations	54,366	54,366	54,366	54,366	54,366	54,366
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows the impact on monthly refills of changes in log of over-the-counter price and log of cash-back subsidy interacted with household attributes at baseline using the survey data. “Over-the-counter price” is the log of the price paid upfront per kg of LPG. “Cash-back subsidy” is the log of the subsidy per kg of LPG that the consumers get back after purchase. The dependent variable is refills, which is the total monthly refills of a 14.2 kg LPG cylinder. Each estimate comes from a separate regression. Columns 1 - 2 show the impact on refills of changes in the log of over-the-counter price and cash-back subsidy interacted with the asset index. Columns 3 - 4 show the impact on refills of changes in log of over-the-counter price and cash-back subsidy interacted with a dummy indicating whether household belongs to general (upper) caste. Columns 5 - 6 show the impact on refills of changes in log of over-the-counter price and cash-back subsidy interacted with a dummy indicating whether household head is a casual laborer. Unit of observation is household-month-year, leading to a total of 54,366 observations, corresponding to monthly refills for 26 months of the 2091 unique consumers. Refill is coded zero if the consumer does not have a gas connection or did not purchase in a given month. All specifications include month, year and household fixed effects. PMUY is a dummy which is 1 if the consumer has a PMUY connection, 0 if it’s a regular gas connection. Standard errors, clustered at the market-month-year level, reported in parentheses. Significant at * 10%, ** 5% and ***1% level.

Table 2.8: Impact of Over-the-counter Price and Cash on Hand (COH) on Monthly Refills (Survey Data)

	(1)	(2)	(3)	(4)	(5)
COH	0.087*** (0.006)	0.087*** (0.006)	0.087*** (0.006)	0.086*** (0.006)	0.086*** (0.006)
Over-the-counter price		-0.230*** (0.057)	-0.234*** (0.057)		
Over-the-counter price $\times$ COH			0.102* (0.053)		
Cash-back subsidy				-0.067*** (0.018)	-0.067*** (0.018)
Cash-back subsidy $\times$ COH					0.015 (0.016)
Mean of Dependent Var.	0.410	0.410	0.410	0.410	0.410
Observations	48,880	48,880	48,880	48,880	48,880
Month FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes

Notes: This table shows the impact on monthly refills of changes in cash on hand (COH) and log of over-the-counter price and log of cash-back subsidy using the survey data. COH is constructed from 3 rounds of the Periodic Labor Force Survey (PLFS), 2017-18, 2018-19 and 2019-20. We average the COH data at the occupation-caste-month level; with 12 months, 4 occupations and 4 caste groups, this gives 192 COH values. We merge COH into the RCT Survey household data by matching at the same caste-occupation-month level. This matching drops about 10% of household-month observations that cannot be assigned a valid caste $\times$ occupation cell in the RCT or cannot be matched to a corresponding PLFS cell for that month. “Over-the-counter price” is the log of the price paid upfront per kg of LPG. “Cash-back subsidy” is the log of the subsidy per kg of LPG that the consumers get back after purchase. The dependent variable is refills, which is the total monthly refills of a 14.2 kg LPG cylinder. Standard errors, clustered at the month-year level, are reported in parentheses. Significant at * 10%, ** 5% and ***1% level.


Table 2.9: Effect of High Rainfall in Sowing Season on Harvest Season Refills

	(1)	(2)
Panel A: Full Sample		
Price	0.025 (0.167)	0.276* (0.149)
PMUY $\times$ Price	-1.235*** (0.211)	-1.240*** (0.254)
High Rain Harvest=1 $\times$ Price	-0.157 (0.172)	-0.374** (0.186)
PMUY $\times$ High Rain Harvest=1 $\times$ Price	0.929*** (0.221)	0.936*** (0.265)
Mean of Dependent Var.	0.586	0.586
Mean of Dependent Var. for PMUY	0.243	0.243
Mean of Dependent Var. for Non-PMUY	0.613	0.613
Observations	12,760,353	12,760,353
Panel B: Rural Sample		
Price	0.043 (0.179)	0.293* (0.154)
PMUY $\times$ Price	-1.227*** (0.172)	-1.232*** (0.201)
High Rain Harvest=1 $\times$ Price	-0.184 (0.183)	-0.392** (0.191)
PMUY $\times$ High Rain Harvest=1 $\times$ Price	0.901*** (0.181)	0.908*** (0.211)
Observations	7,281,167	7,281,167
Panel C: Urban Sample		
Price	-0.001 (0.156)	0.253* (0.146)
PMUY $\times$ Price	-1.253*** (0.428)	-1.257*** (0.451)
High Rain Harvest=1 $\times$ Price	-0.120 (0.160)	-0.350* (0.182)
PMUY $\times$ High Rain Harvest=1 $\times$ Price	1.010** (0.439)	1.014** (0.462)
Observations	5,439,925	5,439,925
Month FE	Yes	Yes
Year FE	No	Yes
Household FE	Yes	Yes

Notes: This table estimates the effect of high rainfall in the sowing season on harvest season refill consumption. The sowing season for winter crops is Oct-Jan, and the corresponding harvest season is Feb-Apr. The sowing season for monsoon crops is Jun-Sep, and the corresponding harvest season is Oct-Dec. The sample is restricted to harvest months (Feb-Apr & Oct-Dec). “High Rain Harvest” = 1 for each harvest season if the preceding sowing season experienced normal, excess or large excess rainfall. “PMUY” and (Over-the-counter) “Price” are defined as in Table 2.3. Standard errors, clustered at the market-month-year level, reported in parentheses. Significant at * 10%, ** 5% and ***1% level.


Appendix 2.A Online Appendix: Additional Results

Figure A1: Delay in LPG Subsidy Transfer




Notes: This figure shows the mean number of days it takes for the subsidy to be transferred after refill cylinder delivery, for refill transactions from BPCL, within a monthXyear. ‘Days’ is defined as the difference between the date of transfer of LPG subsidy to the consumer’s bank account and the date of delivery of the LPG cylinder to the consumer. Subsidy date is available for the majority of BPCL transactions, but for neither other OMC. Data are restricted to 14.2kg cylinder deliveries between Nov 2017 to Dec 2019 in Indore. The average delay in this period is 6.89 days (green line) while the median delay is 5.

Figure A2: Average Cash on Hand (Rural Households)




Notes: This figure shows how Cash on Hand (COH) varies by month during 26 months period (Nov 2017 to Dec 2019) for rural households in India, using data from the Periodic Labor Force Survey 2017-18, 2018-19 and 2019-20 rounds. The correlation coefficient of COH for 12 (collapsed) calendar months during this period between rural households in India and Madhya Pradesh is 0.655 ($p=0.021$).

Figure A3: Correlation of Over-the-counter LPG Price Across Districts (November 2017 - December 2019)



Notes: The domestic price of LPG is obtained from the administrative dataset of BPCL for 10 districts across 5 states - Telangana, Andhra Pradesh, Karnataka, Punjab, Madhya Pradesh and Uttar Pradesh (2 districts per state). The X-axis denotes the duration of study period (spanning 26 months from November 2017 to December 2019). The red line represents Over-the-counter LPG Price for Indore. The grey lines represent the Over-the-counter LPG Price for each of the remaining 9 districts. The blue line represents the average Over-the-counter LPG Price for 9 districts (excluding Indore). All prices are in nominal INR per KG of LPG. The coefficient of correlation between the over-the counter domestic per kg price for Indore and the average of the 9 districts (excluding Indore), is 0.9981 ($p\text{-val} = 0.0000$).

Figure A4: Market-level Over-the-counter and Effective LPG Prices
(November 2017 - December 2019)



Notes: This figure depicts the monthly over-the-counter and effective LPG prices (obtained from the OMC administrative dataset) in each of the 26 markets and averaged across markets in each month. “Over-the-counter price” is the price paid upfront per kg of LPG. “Effective price” is the over-the-counter price minus the cash-back subsidy. The Y-axis shows prices in real (2012) INR for 1 kg of LPG. The X-axis shows the corresponding month, from November 2017 to December 2019 (26 months).

Table A1: Average Cash-back Subsidy, Over-the-counter Price, and Effective Price

	Obs	Mean	SD	Minimum	Maximum
Cash-back Subsidy	676	169.31	62.11	46.90	343.32
Over-the-counter Price	676	520.82	62.45	410.43	702.60
Effective Price	676	351.51	4.98	338.59	374.01

Notes: This table summarizes the cash-back subsidy, over-the-counter price, effective price in the administrative data. Over-the-counter price of LPG is the price paid by consumers upfront to purchase a standard 14.2 kg LPG cylinder. The government kept the effective price (over-the-counter price minus cash-back subsidy) of LPG to consumers relatively constant by letting the cash-back subsidy float in tandem with the over-the-counter price. All prices are in real 2012 INR and vary at market-month level. There are 26 months in our sample (November 2017-December 2019), and 26 markets.

Table A2: Characteristics of Non-PMUY and PMUY Households (Survey Data)

	Non-PMUY			PMUY			Difference
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Obs	Mean	SD	Obs	Mean	SD	(PMUY - Non-PMUY)
Asset Index	1,273	1.85	0.74	818	1.27	0.61	-0.580***
General Caste Category	1,273	0.18	0.39	818	0.10	0.30	-0.088***
Salaried Head of HH	1,273	0.12	0.33	818	0.08	0.27	-0.046***
Self Employed Head of HH	1,273	0.36	0.48	818	0.22	0.41	-0.142***
Casual Laborer Head of HH	1,273	0.31	0.46	818	0.51	0.50	0.198***
Land Owning HH	1,273	0.57	0.50	818	0.45	0.50	-0.116**
Land Owning or Leasing HH	1,273	0.58	0.49	818	0.46	0.50	-0.116**
Above Primary Level Edu of HH Head	1,273	0.47	0.50	818	0.32	0.47	-0.151***
Has Latrine in House	1,273	0.93	0.26	818	0.81	0.39	-0.117***

Notes: This table compares Non-PMUY and PMUY households along several dimensions using the household survey data. All the variables are as defined in Table 2.2. Columns 1 - 3 (4 - 6) show the number of observations, mean and standard deviation of the characteristics of Non-PMUY (PMUY) households. Column 7 is the difference between Columns 5 and 2. Standard errors are clustered at the market-month-year level. Significant at * 10%, ** 5% and ***1% level.

Table A3: Cash on Hand (COH)

	Obs	Mean	SD	Minimum	Maximum
Overall (level)	192	3,664.26	2,144.41	1,029.15	9,577.55
Overall (logs)	192	8.05	0.55	6.94	9.17
Occupation Groups (level)					
Self-employed on farm	48	2,349.13	341.53	1,761.03	3,138.82
Self-employed on non-farm	48	3,718.04	558.97	2,571.21	4,760.78
Casual labor	48	1,720.68	367.62	1,029.15	2,701.50
Salaried	48	6,869.20	1,416.26	4,352.54	9,577.55
Caste Groups (level)					
Scheduled Caste (SC)	48	2,995.24	1,458.85	1,070.15	6,044.01
Scheduled Tribe (ST)	48	4,246.91	2,611.22	1,351.51	9,577.55
Other Backward Caste (OBC)	48	3,359.98	1,746.54	1,029.15	7,120.67
General	48	4,054.93	2,374.04	1,259.29	8,890.91

Notes: This table summarizes the monthly cash-on-hand (COH), in levels (INR) and logs, constructed from the Periodic Labor Force Survey (PLFS), rounds 2017-18, 2018-19 and 2019-20. We attribute the COH data to households in the survey data by collapsing the COH data to the occupation-caste-month level. With 12 months, 4 occupations (Salaried, Self-employed on farm, Self-employed on non-farm, Casual labor) and 4 caste groups (SC, ST, OBC, and General) we have 192 total observations.

Table A4: Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills
(Administrative Data)

	(1)	(2)	(3)	(4)
Panel A: Full Sample				
Over-the-counter price	-0.003*** (0.001)	-0.001 (0.001)		
PMUY $\times$ Over-the-counter price	-0.009*** (0.001)	-0.009*** (0.001)		
Cash-back subsidy			-0.003*** (0.001)	-0.001 (0.001)
PMUY $\times$ Cash-back subsidy			-0.009*** (0.001)	-0.009*** (0.002)
Mean of Dependent Var.	0.578	0.578	0.578	0.578
Mean of Dependent Var. for PMUY	0.252	0.252	0.252	0.252
Mean of Dependent Var. for Non-PMUY	0.604	0.604	0.604	0.604
Observations	23,697,794	23,697,794	23,697,794	23,697,794
Panel B: Rural Sample				
Over-the-counter price	-0.003*** (0.001)	-0.002 (0.001)		
PMUY $\times$ Over-the-counter price	-0.010*** (0.001)	-0.010*** (0.001)		
Cash-back subsidy			-0.003*** (0.001)	-0.002 (0.001)
PMUY $\times$ Cash-back subsidy			-0.010*** (0.001)	-0.010*** (0.001)
Observations	13,522,165	13,522,165	13,522,165	13,522,165
Panel C: Urban Sample				
Over-the-counter price	-0.002*** (0.001)	-0.001 (0.001)		
PMUY $\times$ Over-the-counter price	-0.007*** (0.002)	-0.007*** (0.002)		
Cash-back subsidy			-0.002*** (0.001)	-0.001 (0.001)
PMUY $\times$ Cash-back subsidy			-0.008*** (0.002)	-0.007*** (0.002)
Observations	10,102,717	10,102,717	10,102,717	10,102,717
Month FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes
Household FE	Yes	Yes	Yes	Yes

Notes: Same as in Table 2.3 in the main paper but with over-the-counter price and cash-back subsidy in levels.

Table A5: Effect of LPG Subsidy Delay on Refill Demand, by PMUY status
(Administrative Data)

	(1)	(2)		
Over-the-counter price	-0.030 (0.029)	-0.023 (0.033)		
Over-the-counter price $\times$ Delay	-0.119** (0.051)	-0.084** (0.035)		
PMUY $\times$ Over-the-counter price	-0.275*** (0.055)	-0.273*** (0.059)		
PMUY $\times$ Over-the-counter price $\times$ Delay	-0.001 (0.093)	-0.002 (0.094)		
Cash-back subsidy			-0.008 (0.009)	0.001 (0.009)
Cash-back subsidy $\times$ Delay			-0.022* (0.013)	-0.019** (0.009)
PMUY $\times$ Cash-back subsidy			-0.100*** (0.017)	-0.099*** (0.018)
PMUY $\times$ Cash-back subsidy $\times$ Delay			0.026 (0.024)	0.026 (0.025)
Mean of Dependent Var.	0.578	0.578	0.578	0.578
Mean of Median Delay Days	4.808	4.808	4.808	4.808
Observations	23,697,794	23,697,794	23,697,794	23,697,794
Month FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes
Household FE	Yes	Yes	Yes	Yes

Notes: “Over-the-counter price” is the log of the price paid upfront per kg of LPG. “Cash-back subsidy” is the log of the subsidy per kg of LPG that the consumers get back after purchase. The dependent variable is refills, which is the total monthly refills of a 14.2 kg LPG cylinder. Unit of observation is consumer-month-year, leading to a total of 23.69 million observations, corresponding to monthly refills for 26 months of the 911,454 unique consumers. Refill is coded zero if the consumer does not have a gas connection or did not purchase in a given month. “PMUY” is a dummy that equals 1 if the consumer has a PMUY connection, 0 if it’s a regular gas connection. “Delay” (in days) is the month-year median number of days it takes for the cash-back subsidy to be transferred to the consumer (date of subsidy transferred minus date of refill delivery), demeaned across time. Delay data are available only from BPCL and assigned to each consumer (irrespective of their dealer’s OMC) who bought refills in the corresponding month-year. Delay and PMUY $\times$ Delay are included in all regressions but not reported. Standard errors, clustered at the market-month-year level, reported in parentheses. Significant at * 10%, ** 5% and ***1% level.

Table A6: IV Estimation: Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills (Administrative Data)

	(1)	(2)	(3)	(4)
Panel A: Full Sample				
Over-the-counter price	-0.104*** (0.026)	-0.044 (0.033)		
PMUY $\times$ Over-the-counter price	-0.346*** (0.051)	-0.346*** (0.056)		
Cash-back subsidy			-0.031*** (0.008)	-0.014 (0.010)
PMUY $\times$ Cash-back subsidy			-0.102*** (0.014)	-0.102*** (0.015)
Mean of Dependent Var.	0.578	0.578	0.578	0.578
Mean of Dependent Var. for PMUY	0.252	0.252	0.252	0.252
Mean of Dependent Var. for Non-PMUY	0.604	0.604	0.604	0.604
Observations	23,697,794	23,697,794	23,697,794	23,697,794
Panel B: Rural Sample				
Over-the-counter price	-0.111*** (0.027)	-0.046 (0.034)		
PMUY $\times$ Over-the-counter price	-0.368*** (0.043)	-0.369*** (0.048)		
Cash-back subsidy			-0.033*** (0.009)	-0.014 (0.011)
PMUY $\times$ Cash-back subsidy			-0.109*** (0.012)	-0.109*** (0.013)
Observations	13,522,165	13,522,165	13,522,165	13,522,165
Panel C: Urban Sample				
Over-the-counter price	-0.095*** (0.026)	-0.042 (0.033)		
PMUY $\times$ Over-the-counter price	-0.280*** (0.080)	-0.280*** (0.082)		
Cash-back subsidy			-0.028*** (0.008)	-0.013 (0.010)
PMUY $\times$ Cash-back subsidy			-0.082*** (0.021)	-0.082*** (0.022)
Observations	10,102,717	10,102,717	10,102,717	10,102,717
Month FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes
Household FE	Yes	Yes	Yes	Yes

Notes: This table shows the impact on monthly refills of changes in log of over-the-counter price and log of cash-back subsidy. **Panel A** shows the estimates for the full OMC administrative sample, **Panel B** shows the estimates for rural sub-sample, and **Panel C** shows the estimates for the urban sub-sample. “Over-the-counter price” is the log of the price paid upfront per kg of LPG. “Cash-back subsidy” is the log of the subsidy per kg of LPG that the consumers get back after purchase. 1 month lag of international LPG price is used as an IV for over-the-counter price and cash-back subsidy. The dependent variable is refills, which is the total monthly refills of a 14.2 kg LPG cylinder. Unit of observation is consumer-month-year, leading to a total of 23.69 million observations, corresponding to monthly refills for 26 months of the 911,454 unique consumers. Refill is coded zero if the consumer does not have a gas connection or did not purchase in a given month. PMUY is a dummy which equals 1 if the consumer has a PMUY connection, 0 if it's a regular gas connection. Standard errors, clustered at the market-month-year level, reported in parentheses. Significant at * 10%, ** 5% and ***1% level.

Table A7: IV Estimation: Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills (Survey Data)

	(1)	(2)	(3)	(4)
Over-the-counter price	-0.165*** (0.034)	-0.069 (0.042)		
PMUY $\times$ Over-the-counter price	-0.259*** (0.053)	-0.260*** (0.053)		
Cash-back subsidy			-0.049*** (0.010)	-0.022* (0.013)
PMUY $\times$ Cash-back subsidy			-0.078*** (0.016)	-0.078*** (0.016)
Mean of Dependent Var.	0.411	0.411	0.411	0.411
Mean of Dependent Var. for PMUY	0.285	0.285	0.285	0.285
Mean of Dependent Var. for Non-PMUY	0.492	0.492	0.492	0.492
Observations	54,366	54,366	54,366	54,366
Month FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes
Household FE	Yes	Yes	Yes	Yes

Notes: This table shows the impact on monthly refills of changes in log of over-the-counter price and log of cash-back subsidy using the survey data. “Over-the-counter price” is the log of the price paid upfront per kg of LPG. “Cash-back subsidy” is the log of the subsidy per kg of LPG that the consumers get back after purchase. 1 month lag of international LPG price is used as an IV for over-the-counter price and cash-back subsidy. The dependent variable is refills, which is the total monthly refills of a 14.2 kg LPG cylinder. Unit of observation is household-month-year, leading to a total of 54,366 observations, corresponding to monthly refills for 26 months of the 2091 unique consumers. Refill is coded zero if the consumer does not have a gas connection or did not purchase in a given month. PMUY is a dummy which equals 1 if the consumer has a PMUY connection, 0 if it’s a regular gas connection. Standard errors, clustered at the market-month-year level, reported in parentheses. Significant at * 10%, ** 5% and ***1% level.

Table A8: Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills
(Administrative Data with Bootstrapped Standard Errors)

	(1)	(2)	(3)	(4)
Panel A: Full Sample				
Over-the-counter price	-0.091** [0.012]	-0.043** [0.036]		
PMUY $\times$ Over-the-counter price	-0.349*** [0.000]	-0.347*** [0.000]		
Cash-back subsidy			-0.025** [0.018]	-0.007* [0.070]
PMUY $\times$ Cash-back subsidy			-0.110*** [0.000]	-0.110*** [0.000]
Mean of Dependent Var.	0.578	0.578	0.578	0.578
Mean of Dependent Var. for PMUY	0.252	0.252	0.252	0.252
Mean of Dependent Var. for Non-PMUY	0.604	0.604	0.604	0.604
Observations	23,697,794	23,697,794	23,697,794	23,697,794
Panel B: Rural Sample				
Over-the-counter price	-0.098** [0.021]	-0.046* [0.064]		
PMUY $\times$ Over-the-counter price	-0.374*** [0.000]	-0.372*** [0.000]		
Cash-back subsidy			-0.028** [0.025]	-0.008 [0.151]
PMUY $\times$ Cash-back subsidy			-0.116*** [0.000]	-0.115*** [0.000]
Observations	13,522,165	13,522,165	13,522,165	13,522,165
Panel C: Urban Sample				
Over-the-counter price	-0.083** [0.046]	-0.041* [0.058]		
PMUY $\times$ Over-the-counter price	-0.275*** [0.000]	-0.275*** [0.000]		
Cash-back subsidy			-0.023* [0.064]	-0.006** [0.024]
PMUY $\times$ Cash-back subsidy			-0.094*** [0.000]	-0.093*** [0.000]
Observations	10,102,717	10,102,717	10,102,717	10,102,717
Month FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes
Household FE	Yes	Yes	Yes	Yes

Notes: This table shows the impact on monthly refills of changes in log of over-the-counter price and log of cash-back subsidy for the administrative sample. Panel A shows the estimates for the full OMC administrative sample, Panel B shows the estimates for the rural sub-sample, and Panel C shows the estimates for the urban sub-sample. “Over-the-counter price” is the log of the price paid upfront per kg of LPG. “Cash-back subsidy” is the log of the subsidy per kg of LPG that the consumers get back after purchase. The dependent variable is refills, which is the total monthly refills of a 14.2 kg LPG cylinder. Unit of observation is consumer-month-year, leading to a total of 23.69 million observations, corresponding to monthly refills for 26 months of the 911,454 unique consumers. Refill is coded zero if the consumer does not have a gas connection or did not purchase in a given month. PMUY is a dummy which equals 1 if the consumer has a PMUY connection, 0 if it's a regular gas connection. Standard errors are bootstrapped at market level, and significance corresponds to bootstrapped p -values reported in square brackets. Significant at * 10%, ** 5% and ***1% level.

Table A9: Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills
(Survey Data with Bootstrapped Standard Errors)

	(1)	(2)	(3)	(4)
Over-the-counter price	-0.172*** [0.003]	-0.082 [0.188]		
PMUY $\times$ Over-the-counter price	-0.255*** [0.000]	-0.254*** [0.000]		
Cash-back subsidy			-0.046*** [0.005]	-0.015 [0.391]
PMUY $\times$ Cash-back subsidy			-0.079*** [0.001]	-0.078*** [0.001]
Mean of Dependent Var.	0.411	0.411	0.411	0.411
Mean of Dependent Var. for PMUY	0.285	0.285	0.285	0.285
Mean of Dependent Var. for Non-PMUY	0.492	0.492	0.492	0.492
Observations	54,366	54,366	54,366	54,366
Month FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes
Household FE	Yes	Yes	Yes	Yes

Notes: This table shows the impact on monthly refills of changes in log of over-the-counter price and log of cash-back subsidy for the survey data. “Over-the-counter price” is the log of the price paid upfront per kg of LPG. “Cash-back subsidy” is the log of the subsidy per kg of LPG that the consumers get back after purchase. The dependent variable is refills, which is the total monthly refills of an LPG cylinder (which is an indivisible object of 14.2 kg). Unit of observation is household-month-year, leading to a total of 54,366 observations, corresponding to monthly refills for 26 months of the 2091 unique consumers. Refill is coded zero if the consumer does not have a gas connection or did not purchase in a given month. Columns 1 and 3 include month and household (i.e. customer) fixed effects, and columns 2 and 4 add year fixed effects. PMUY is a dummy which is 1 if the consumer has a PMUY connection, 0 if it's a regular gas connection. Standard errors are bootstrapped at market level, and significance corresponds to bootstrapped p-values reported in square brackets. Significant at * 10%, ** 5% and ***1% level.

Table A10: Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills
(Administrative Data) - Poisson Pseudo-Maximum Likelihood Estimation

	(1)	(2)	(3)	(4)
Panel A: Full Sample				
Over-the-counter price	-0.151*** (0.041)	-0.062 (0.057)		
PMUY $\times$ Over-the-counter price	-1.473*** (0.178)	-1.462*** (0.183)		
Cash-back subsidy			-0.042*** (0.011)	-0.008 (0.014)
PMUY $\times$ Cash-back subsidy			-0.426*** (0.041)	-0.426*** (0.043)
Mean of Dependent Var.	0.578	0.578	0.578	0.578
Mean of Dependent Var. for PMUY	0.252	0.252	0.252	0.252
Mean of Dependent Var. for Non-PMUY	0.604	0.604	0.604	0.604
Observations	23,667,665	23,667,665	23,667,665	23,667,665
Panel B: Rural Sample				
Over-the-counter price	-0.161*** (0.041)	-0.064 (0.057)		
PMUY $\times$ Over-the-counter price	-1.552*** (0.140)	-1.540*** (0.147)		
Cash-back subsidy			-0.045*** (0.012)	-0.008 (0.015)
PMUY $\times$ Cash-back subsidy			-0.436*** (0.035)	-0.437*** (0.036)
Observations	13,505,577	13,505,577	13,505,577	13,505,577
Panel C: Urban Sample				
Over-the-counter price	-0.137*** (0.041)	-0.059 (0.058)		
PMUY $\times$ Over-the-counter price	-1.239*** (0.312)	-1.231*** (0.313)		
Cash-back subsidy			-0.038*** (0.011)	-0.008 (0.014)
PMUY $\times$ Cash-back subsidy			-0.391*** (0.070)	-0.391*** (0.071)
Observations	10,089,228	10,089,228	10,089,228	10,089,228
Month FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes
Household FE	Yes	Yes	Yes	Yes

Notes: This table shows the impact on monthly refills of changes in log of over-the-counter price and log of cash-back subsidy estimated using Poisson Pseudo Maximum Likelihood Estimation. Panel A shows the estimates for the full OMC administrative sample, Panel B shows the estimates for rural sub-sample, and Panel C shows the estimates for the urban sub-sample. "Over-the-counter price" is the log of the price paid upfront per kg of LPG. "Cash-back subsidy" is the log of the subsidy per kg of LPG that the consumers get back after purchase. The dependent variable is refills, which is the total monthly refills of a 14.2 kg LPG cylinder. Unit of observation is consumer-month-year, leading to a total of 23.69 million observations, corresponding to monthly refills for 26 months of the 911,454 unique consumers. Refill is coded zero if the consumer does not have a gas connection or did not purchase in a given month. PMUY is a dummy which equals 1 if the consumer has a PMUY connection, 0 if it's a regular gas connection. Standard errors, clustered at the market-month-year level, reported in parentheses. Significant at * 10%, ** 5% and ***1% level.

Table A11: Birth and Post-Birth Child Health Outcomes

	Obs	Mean	SD	Minimum	Maximum
Birth Weight	16,801	2.784	0.568	0.500	9.000
Cough	223,764	0.129	0.335	0.000	1.000
Rapid Breathing	223,862	0.065	0.246	0.000	1.000

Notes: This table summarizes the birth and post-birth child health outcomes from June 2018 to May 2021. All health outcomes are from the DHS Round 5, 2019-21 (previous 12 months is the reference period for interviews conducted during June 2019 to May 2021). DHS data used in the analysis is for all India (rural and urban) sample. ‘Birth Weight’ is weight at birth measured in kilograms. The sample is restricted to births whose conception was within a year prior to the DHS survey (pregnancies conceived from June 2018 to Dec 2020). ‘Cough’ is a dummy that equals 1 if the child (under 5 years of age) has been ill with cough in last two weeks and ‘Rapid breathing’ equals 1 if the child (under 5 years of age) suffered from rapid breathing in last two weeks.

Table A12: Impact of Over-the-counter Price and Cash-back Subsidy on Birth and Post-birth Child Health Outcomes (with Sampling Weights)

	Birth Weight		Cough		Rapid Breathing	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full Sample						
Over-the-counter price	-0.370 (0.334)		0.055 (0.043)		0.055** (0.020)	
Cash-back subsidy		-0.005 (0.026)		-0.002 (0.016)		0.019*** (0.005)
Mean of Dependent Var.	2.784	2.784	0.129	0.129	0.065	0.065
Observations	16,801	16,801	223,764	223,764	223,862	223,862
FDR-adjusted p -value	[0.226]	[1]	[0.226]	[1]	[0.039]	[0.006]
Panel B: Rural Sample						
Over-the-counter price	-0.510 (0.308)		0.044 (0.043)		0.079** (0.035)	
Cash-back subsidy		-0.017 (0.027)		0.003 (0.014)		0.018* (0.010)
Mean of Dependent Var.	2.780	2.780	0.130	0.130	0.066	0.066
Observations	13,566	13,566	177,951	177,951	178,035	178,035
FDR-adjusted p -value	[0.127]	[1]	[0.195]	[1]	[0.127]	[0.272]
Panel C: Urban Sample						
Over-the-counter price	0.067 (0.591)		0.035 (0.095)		0.049 (0.031)	
Cash-back subsidy		0.047 (0.041)		-0.040 (0.041)		0.025 (0.020)
Mean of Dependent Var.	2.800	2.800	0.124	0.124	0.059	0.059
Observations	3,235	3,235	45,813	45,813	45,827	45,827
FDR-adjusted p -value	[1]	[0.513]	[1]	[0.513]	[0.745]	[0.513]
Mother and Child Controls	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
State $\times$ Month-Year Trend	Yes	Yes	Yes	Yes	Yes	Yes

Notes: As shown in Table 2.6 in the main paper. Observations in the regressions are weighted using the DHS household sampling weights.

Table A13: Effect of Over-the-counter Price and Cash-back Subsidy on Post-birth Child Health Outcomes

	Cough		Rapid Breathing	
	(1)	(2)	(3)	(4)
Panel A: Full Sample				
Over-the-counter price	0.089*		0.106***	
	[0.081]		[0.006]	
Cash-back subsidy		0.005		0.029**
		[0.593]		[0.011]
Mean of Dependent Var.	0.129	0.129	0.065	0.065
Observations	223,764	223,764	223,862	223,862
Panel B: Rural Sample				
Over-the-counter price	0.085		0.135***	
	[0.200]		[0.003]	
Cash-back subsidy		0.005		0.030*
		[0.674]		[0.053]
Mean of Dependent Var.	0.130	0.130	0.066	0.066
Observations	177,951	177,951	178,035	178,035
Panel C: Urban Sample				
Over-the-counter price	0.083		0.081	
	[0.422]		[0.202]	
Cash-back subsidy		-0.003		0.054
		[0.915]		[0.292]
Mean of Dependent Var.	0.124	0.124	0.059	0.059
Observations	45,813	45,813	45,827	45,827
Mother and Child Controls	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
State $\times$ Month-Year Trend	Yes	Yes	Yes	Yes

Note: This table is analogous to Table 2.6 (columns 3 to 6) in the main paper. Standard errors are bootstrapped at Month-Year of interview, and significance corresponds to bootstrapped p -values reported in square brackets. Significant at * 10%, ** 5% and ***1% level.

Table A14: Effect of Over-the-counter Price and Cash-back Subsidy on Child Anthropometric Outcomes

	Height for Age Z-Score		Weight for Age Z-Score	
	(1)	(2)	(3)	(4)
Over-the-counter price	0.017 [0.932]		-0.041 [0.800]	
Cash-back subsidy		0.015 [0.891]		-0.150 [0.120]
Mean of Dependent Var.	-1.312	-1.312	-1.395	-1.395
Observations	206,025	206,025	210,524	210,524
Mother and Child Controls	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
State $\times$ Month-Year Trend	Yes	Yes	Yes	Yes

Notes: Log of over-the-counter LPG price and cash-back subsidy are obtained from the administrative dataset for BPCL Indore, while the health outcomes are from the DHS Round 5, 2019-21 (June 2019 to May 2021). DHS data used in the analysis is for all India (rural and urban) sample. Ht for Age and Wt for Age Z-scores are calculated using the new Child Growth Standards released by WHO in 2006, to develop a new international standard for assessing the physical growth nutrition status and motor development in all children from birth to age five. All time controls correspond to interview month-year of household. Standard errors are bootstrapped at Month-Year of interview, and significance corresponds to bootstrapped p -values reported in square brackets. Child controls include sex of child, if child is part of a multiple birth and birth order of the child in the household, age of child (and age squared). Mother controls include mother's age at birth (including age squared), mother's education level, if household is located in rural/urban area, religion, caste and wealth index. State $\times$ Year-Month Trend is 'i.state $\times$ c.month-year' of interview. Significant at * 10%, ** 5% and ***1% level.

Table A15: Impact of Over-the-counter Price on Monthly Refills: Heterogeneity by Asset Ownership (Survey Data)

	(1)	(2)
Over-the-counter price	-0.220*** (0.040)	-0.274*** (0.050)
Asset Above median $\times$ Over-the-counter price	0.077** (0.037)	
Asset Quartile 2 $\times$ Over-the-counter price		0.108** (0.054)
Asset Quartile 3 $\times$ Over-the-counter price		0.140** (0.055)
Asset Quartile 4 $\times$ Over-the-counter price		0.122** (0.056)
Constant	1.064*** (0.125)	1.064*** (0.125)
Mean of Dependent Var.	0.411	0.411
Observations	54,366	54,366
Month FE	Yes	Yes
Year FE	Yes	Yes
Household FE	Yes	Yes

Notes: This table shows the impact on monthly refills of changes in log of over-the-counter price interacted with median asset and asset quartiles for the survey data. “Over-the-counter price” is the log of the price paid upfront per kg of LPG. The dependent variable is refills, which is the total monthly refills of a 14.2 kg LPG cylinder. Households are divided into 2 and 4 groups based on their baseline asset index. Asset Half 1 and Asset Quartile 1 (lowest) are the omitted categories. Unit of observation is household-month-year, leading to a total of 54,366 observations, corresponding to monthly refills for 26 months of the 2091 unique consumers. Refill is coded zero if the consumer does not have a gas connection or did not purchase in a given month. All specifications include month, year and household fixed effects. PMUY is a dummy which equals 1 if the consumer has a PMUY connection, 0 if it’s a regular gas connection. Standard errors, clustered at the market-month-year level, reported in parentheses. Significant at * 10%, ** 5% and ***1% level.

Table A16: Correlation of Consumer Price Index with Nominal Over-the-counter Price of LPG

	(1)	(2)	(3)	(4)	(5)
	Food & Beverage	Pan, Tobacco & Intoxicants	Clothing & Footwear	Fuel & Light	All
Over the counter LPG price (Nominal terms)	-0.175 (0.139)	-0.227 (0.158)	-0.001 (0.046)	0.164* (0.081)	-0.098 (0.086)
Constant	150.658*** (7.303)	172.076*** (8.300)	150.104*** (2.429)	137.670*** (4.256)	109.858*** (4.502)
Observations	26	26	26	26	26

Notes: This table regresses the monthly rural Consumer Price Index (CPI) on the (nominal) average monthly Over-the counter LPG price (per kg) between November 2017 - December 2019 (26 months). In Column 1 the dependent variable is the food and beverages CPI, in Column 2 it is pan, tobacco and intoxicants CPI, in Column 3 it is clothing and footwear CPI, and Column 4 is fuel and light CPI. Column 5 is a combination of all the four CPIs using on the following weights 0.54, 0.074, 0.0326, and 0.079, respectively as per [Goyal et al. \(2021\)](#). Standard errors reported in parentheses. Significant at * 10%, ** 5% and ***1% level.

Table A17: Impact of Over-the-counter Price on Monthly Refills: Heterogeneity by Knowledge About Subsidy (Survey Data)

	(1)	(2)	(3)	(4)	(5)
Over-the-counter price	-0.102* (0.058)	-0.110** (0.056)	-0.083* (0.049)	-0.090 (0.058)	0.002 (0.048)
PMUY $\times$ Over-the-counter price	-0.211*** (0.069)	-0.244*** (0.082)	-0.286*** (0.062)	-0.230*** (0.068)	-0.335*** (0.071)
Govt. Deposits Subsidy $\times$ Over-the-counter price	0.024 (0.060)				
PMUY $\times$ Govt. Deposits Subsidy $\times$ Over-the-counter price	-0.066 (0.083)				
Subsidy Amt. Same $\times$ Over-the-counter price		0.038 (0.056)			
PMUY $\times$ Subsidy Amt. Same $\times$ Over-the-counter price		-0.014 (0.082)			
Out-of-Pocket<Paid $\times$ Over-the-counter price			-0.001 (0.056)		
PMUY $\times$ Out-of-Pocket<Paid $\times$ Over-the-counter price			0.067 (0.077)		
Subsidy SMS Alert $\times$ Over-the-counter price				0.009 (0.058)	
PMUY $\times$ Subsidy SMS Alert $\times$ Over-the-counter price				-0.037 (0.082)	
Subsidy only for PMUY $\times$ Over-the-counter price					-0.124*** (0.045)
PMUY $\times$ Subsidy only for PMUY $\times$ Over-the-counter price					0.118 (0.075)
Mean of Dependent Var.	0.412	0.412	0.412	0.412	0.412
Observations	53,534	53,534	53,534	53,534	53,534
Month FE	Yes	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes

Notes: The dependent variable is the total number of monthly refills of a 14.2 kg LPG cylinder consumed by a household in the survey data. LPG subsidy knowledge equals 1 if the household correctly agreed/disagreed with each of the five statements: (1) ‘govt. deposits subsidy’ to your bank account (2) deposited ‘subsidy amount remains same’ every time (3) ‘out-of-pocket expense’ on LPG cylinder is less than the market price paid (4) govt. sends an ‘SMS alert’ about subsidy once deposited (5) LPG refill ‘subsidy is only for PMUY’ customers. Unit of observation is consumer-financial month- financial year. Refill is missing if consumer does not have a gas connection, and refills are non-missing once the consumer avails the connection (hence, the panel is unbalanced). Standard errors, clustered at the market-month-year level, reported in parentheses. Significant at * 10%, ** 5% and ***1% level.

Table A18: Impact of Over-the-counter Price on Monthly Refills: Heterogeneity by Trust in ASHA Worker (Survey Data)

	(1)
Over-the-counter price	-0.073 (0.081)
PMUY $\times$ Over-the-counter Price	-0.250** (0.103)
Trust info from ASHA $\times$ Over-the-counter Price	-0.010 (0.085)
PMUY $\times$ Trust info from ASHA $\times$ Over-the-counter Price	-0.004 (0.108)
Mean of Dependent Var.	0.411
Observations	54,366
Month FE	Yes
Household FE	Yes
Year FE	Yes

Notes: The dependent variable is the total monthly refills of a 14.2 kg LPG cylinder consumed by a household in the survey data. ASHA is an acronym for ‘Accredited Social Health Activist’, a community health worker. "Trust info from ASHA" is the response to the survey question "Does your household trust information given by ASHA worker in your village", which equals 1 if the household reports saying they do trust the information from ASHA worker, and 0 otherwise. "Over-the-counter price" is the log of the price paid upfront per kg of LPG. Unit of observation is household-month-year, leading to a total of 54,366 observations, corresponding to monthly refills for 26 months of the 2091 unique consumers. Refill is coded zero if the consumer does not have a gas connection or did not purchase in a given month. All specifications include month, year and household fixed effects. PMUY is a dummy which equals 1 if the consumer has a PMUY connection, 0 if it's a regular gas connection. Standard errors, clustered at the market-month-year level, reported in parentheses. Significant at * 10%, ** 5% and ***1% level.

Table A19: Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills
(Administrative Data)
(Sample Restricted to experienced LPG Users)

	(1)	(2)	(3)	(4)
Panel A: Full Sample				
Over-the-counter price	-0.001 (0.029)	-0.073** (0.035)		
PMUY $\times$ Over-the-counter price	-0.312*** (0.043)	-0.312*** (0.040)		
Cash-back subsidy			0.003 (0.007)	-0.017* (0.009)
PMUY $\times$ Cash-back subsidy			-0.087*** (0.012)	-0.087*** (0.011)
Mean of Dependent Var.	0.641	0.641	0.641	0.641
Observations	18,541,172	18,541,172	18,541,172	18,541,172
Panel B: Rural Sample				
Over-the-counter price	-0.010 (0.028)	-0.080** (0.036)		
PMUY $\times$ Over-the-counter price	-0.327*** (0.042)	-0.327*** (0.040)		
Cash-back subsidy			0.001 (0.007)	-0.019* (0.010)
PMUY $\times$ Cash-back subsidy			-0.091*** (0.012)	-0.091*** (0.012)
Mean of Dependent Var.	0.642	0.642	0.642	0.642
Observations	10,463,582	10,463,582	10,463,582	10,463,582
Panel C: Urban Sample				
Over-the-counter price	0.010 (0.031)	-0.063* (0.035)		
PMUY $\times$ Over-the-counter price	-0.216*** (0.055)	-0.216*** (0.053)		
Cash-back subsidy			0.006 (0.007)	-0.015* (0.009)
PMUY $\times$ Cash-back subsidy			-0.062*** (0.014)	-0.062*** (0.014)
Mean of Dependent Var.	0.640	0.640	0.640	0.640
Observations	8,014,064	8,014,064	8,014,064	8,014,064
Month FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes
Household FE	Yes	Yes	Yes	Yes

Notes: This table reproduces Table 2.3 in the main paper to show the impact of changes in log of over-the-counter price and cash-back subsidy on monthly consumption of LPG refills. The sample is restricted to only those consumers who had an LPG account by the beginning of 2017. **Panel A** shows the estimates for the full OMC administrative sample, **Panel B** shows the estimates for rural sub-sample, and **Panel C** shows the estimates for the urban sub-sample. “Over-the-counter price” is the log of the price paid upfront per kg of LPG. “Cash-back subsidy” is the log of the subsidy per kg of LPG that the consumers get back after purchase. The dependent variable is refills, which is the total monthly refills of a 14.2 kg LPG cylinder. Unit of observation is consumer-month-year. Refill is coded zero if consumer did not purchase in a given month. PMUY is a dummy which equals 1 if the consumer has a PMUY connection, 0 if it's a regular gas connection. Standard errors, clustered at the market-month-year level, reported in parentheses. Significant at * 10%, ** 5% and ***1% level.

Table A20: Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills (Administrative Data) including Unemployment as a Control

	(1)	(2)	(3)	(4)
Panel A: Full Sample				
Over-the-counter price	-0.049 (0.032)	-0.024 (0.033)		
Unemployment	0.015*** (0.003)	-0.010 (0.007)	0.015*** (0.003)	-0.015* (0.008)
PMUY $\times$ Over-the-counter price	-0.348*** (0.053)	-0.347*** (0.052)		
Cash-back subsidy			-0.011 (0.009)	-0.001 (0.008)
PMUY $\times$ Cash-back subsidy			-0.110*** (0.013)	-0.110*** (0.013)
Mean of Dependent Var.	0.578	0.578	0.578	0.578
Mean of Dependent Var. for PMUY	0.252	0.252	0.252	0.252
Mean of Dependent Var. for Non-PMUY	0.604	0.604	0.604	0.604
Observations	23,697,794	23,697,794	23,697,794	23,697,794
Panel B: Rural Sample				
Over-the-counter price	-0.071** (0.030)	-0.018 (0.036)		
Unemployment (Rural)	0.020*** (0.004)	-0.014 (0.011)	0.020*** (0.005)	-0.019* (0.012)
PMUY $\times$ Over-the-counter price	-0.372*** (0.046)	-0.372*** (0.045)		
Cash-back subsidy			-0.017** (0.009)	0.002 (0.009)
PMUY $\times$ Cash-back subsidy			-0.115*** (0.012)	-0.115*** (0.011)
Observations	13,522,165	13,522,165	13,522,165	13,522,165
Panel C: Urban Sample				
Over-the-counter price	-0.040 (0.033)	-0.028 (0.032)		
Unemployment (Urban)	0.012*** (0.003)	-0.007 (0.004)	0.012*** (0.003)	-0.009* (0.005)
PMUY $\times$ Over-the-counter price	-0.275*** (0.074)	-0.275*** (0.074)		
Cash-back subsidy			-0.008 (0.009)	-0.003 (0.007)
PMUY $\times$ Cash-back subsidy			-0.094*** (0.018)	-0.093*** (0.018)
Observations	10,102,717	10,102,717	10,102,717	10,102,717
Month FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes
Household FE	Yes	Yes	Yes	Yes

Notes: This table shows the impact on monthly refills of changes in log of over-the-counter price and log of cash-back subsidy using the OMC administrative data, controlling for the monthly unemployment rate. **Panel A** shows the estimates for the full OMC administrative sample, **Panel B** shows the estimates for rural sub-sample, and **Panel C** shows the estimates for the urban sub-sample. “Over-the-counter price” is the log of the price paid upfront per kg of LPG. “Cash-back subsidy” is the log of the subsidy per kg of LPG that the consumers get back after purchase. “Unemployment” is the monthly unemployment rate for all-India (Panel A) and rural (Panel B) and urban (Panel C) regions, calculated (from CMIE data) as the total number of men aged 15 and above who are unemployed and willing to work as a percentage of the labour force (men aged 15 and above who are employed or unemployed and willing to work). The dependent variable is refills, which is the total monthly refills of a 14.2 kg LPG cylinder. Unit of observation is consumer-month-year, leading to a total of 23.69 million observations, corresponding to monthly refills for 26 months of the 911,454 unique consumers. Refill is coded zero if the consumer does not have a gas connection or did not purchase in a given month. PMUY is a dummy that equals 1 if the consumer has a PMUY connection, 0 if it’s a regular gas connection. Standard errors, clustered at the market-month-year level, reported in parentheses. Significant at * 10%, ** 5% and ***1% level.

Table A21: Impact of Over-the-counter Price and Cash-back Subsidy on Solid Fuel Expenditure (Survey Data) (Adding Year FE)

	Explicit Expenditure		Implicit Expenditure		Total Expenditure	
	(1)	(2)	(3)	(4)	(5)	(6)
Over-the-counter price	-207.356 (708.855)		1,786.951 (1,562.332)		1,579.595 (1,762.133)	
PMUY $\times$ Over-the-counter price	179.255** (86.475)		1,008.961*** (138.637)		1,188.216*** (182.887)	
Cash-back subsidy		-24.642 (65.481)		166.540 (147.326)		141.897 (166.834)
PMUY $\times$ Cash-back subsidy		57.566** (28.266)		328.092*** (45.690)		385.657*** (60.065)
Mean of Dependent Var.	208.975	208.975	1,047.526	1,047.526	1,256.501	1,256.501
Observations	4,118	4,118	4,118	4,118	4,118	4,118
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Same as Table 2.5 in the main paper but with Year FE.

Appendix 2.B Supplementary Results: Reproduction of analyses in the main paper with LPG Refills recoded as missing until consumer obtains an LPG account during our study period (Nov 2017 - Dec 2019)

Table B1: LPG Consumer Characteristics (Administrative Data)

	Obs	Mean	SD	Minimum	Maximum
Panel A: Attributes					
PMUY	910,778	0.074	0.262	0.000	1.000
Female	910,423	0.320	0.466	0.000	1.000
Rural	907,975	0.572	0.495	0.000	1.000
Panel B: LPG Refills					
All	21,898,168	0.625	0.640	0.000	16.000
PMUY	1,299,503	0.340	0.549	0.000	12.000
Non-PMUY	20,598,665	0.643	0.642	0.000	16.000
Rural	12,468,102	0.624	0.646	0.000	16.000
PMUY	936,682	0.343	0.554	0.000	12.000
Non-PMUY	11,531,420	0.647	0.648	0.000	16.000
Urban	9,360,226	0.626	0.633	0.000	13.000
PMUY	361,737	0.332	0.535	0.000	11.000
Non-PMUY	8,998,489	0.637	0.634	0.000	13.000

Notes: as shown in Table 2.1 in the main paper.

Table B2: Household Characteristics (Survey Data)

	Obs	Mean	SD	Minimum	Maximum
Panel A: Household Attributes (Baseline)					
PMUY	2,085	0.391	0.488	0.000	1.000
Asset Index	2,085	1.620	0.748	-0.150	3.965
General Caste	2,085	0.150	0.357	0.000	1.000
Salaried HH Head	2,085	0.106	0.309	0.000	1.000
Agri Self-Employed HH Head	2,085	0.305	0.461	0.000	1.000
Agri Casual Laborer HH Head	2,085	0.390	0.488	0.000	1.000
Land Owner	2,085	0.522	0.500	0.000	1.000
Land Owner/Leaser	2,085	0.536	0.499	0.000	1.000
HH Head Education	2,085	0.412	0.492	0.000	1.000
Latrine in HH	2,085	0.879	0.326	0.000	1.000
Panel B: Household Solid Fuel					
Baseline & Endline					
Explicit firewood expenditure (Rs.)	4,150	56.791	282.530	0.000	6,000.000
Implicit firewood expenditure (Rs.)	4,150	543.258	833.181	0.000	4,480.000
Explicit dung expenditure (Rs.)	4,150	151.177	470.125	0.000	6,000.000
Implicit dung expenditure (Rs.)	4,150	503.106	626.538	0.000	9,800.000
Panel C: Household Refills Matched to Admin Data					
All Refills	50,578	0.442	0.586	0.000	7.000
PMUY refills	18,335	0.331	0.530	0.000	7.000
Non-PMUY refills	32,243	0.505	0.606	0.000	6.000

Notes: as shown in Table 2.2 in the main paper.

Table B3: Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills
(Administrative Data)

	(1)	(2)	(3)	(4)
Panel A: Full Sample				
Over-the-counter price	0.015 (0.029)	-0.060* (0.035)		
PMUY $\times$ Over-the-counter price	-0.230*** (0.035)	-0.227*** (0.035)		
Cash-back subsidy			0.008 (0.007)	-0.013 (0.009)
PMUY $\times$ Cash-back subsidy			-0.059*** (0.010)	-0.058*** (0.010)
Mean of Dependent Var.	0.625	0.625	0.625	0.625
Mean of Dependent Var. for PMUY	0.340	0.340	0.340	0.340
Mean of Dependent Var. for Non-PMUY	0.643	0.643	0.643	0.643
Observations	21,894,452	21,894,452	21,894,452	21,894,452
Panel B: Rural Sample				
Over-the-counter price	0.008 (0.028)	-0.067* (0.037)		
PMUY $\times$ Over-the-counter price	-0.259*** (0.034)	-0.256*** (0.033)		
Cash-back subsidy			0.006 (0.007)	-0.015 (0.010)
PMUY $\times$ Cash-back subsidy			-0.067*** (0.010)	-0.065*** (0.010)
Observations	12,466,067	12,466,067	12,466,067	12,466,067
Panel C: Urban Sample				
Over-the-counter price	0.025 (0.031)	-0.052 (0.034)		
PMUY $\times$ Over-the-counter price	-0.150*** (0.053)	-0.146*** (0.055)		
Cash-back subsidy			0.010 (0.007)	-0.011 (0.008)
PMUY $\times$ Cash-back subsidy			-0.038** (0.016)	-0.036** (0.016)
Observations	9,358,549	9,358,549	9,358,549	9,358,549
Month FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes
Household FE	Yes	Yes	Yes	Yes

Notes: as shown in Table 2.3 in the main paper.

Table B4: Impact of Over-the-counter Price and Cash-back Subsidy on Monthly Refills
(Survey Data)

	(1)	(2)	(3)	(4)
Over-the-counter price	-0.143*** (0.034)	-0.146*** (0.042)		
PMUY $\times$ Over-the-counter price	-0.190*** (0.052)	-0.189*** (0.052)		
Cash-back subsidy			-0.037*** (0.010)	-0.035*** (0.013)
PMUY $\times$ Cash-back subsidy			-0.054*** (0.016)	-0.053*** (0.016)
Mean of Dependent Var.	0.442	0.442	0.442	0.442
Mean of Dependent Var. for PMUY	0.331	0.331	0.331	0.331
Mean of Dependent Var. for Non-PMUY	0.505	0.505	0.505	0.505
Observations	50,577	50,577	50,577	50,577
Month FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes
Household FE	Yes	Yes	Yes	Yes

Notes: as shown in Table 2.4 in the main paper.

Table B5: Impact of Over-the-counter Price on Monthly Refills: Heterogeneity by Household Attributes (Survey data)

	Z					
	Asset Index		General Caste		Casual Laborer	
	(1)	(2)	(3)	(4)	(5)	(6)
Over-the-counter price	-0.328*** (0.061)		-0.232*** (0.038)		-0.187*** (0.040)	
Over-the-counter price $\times$ Z	0.069** (0.028)		0.108* (0.060)		-0.075* (0.040)	
Cash-back subsidy		-0.085*** (0.018)		-0.059*** (0.012)		-0.048*** (0.013)
Cash-back subsidy $\times$ Z		0.019** (0.008)		0.031* (0.018)		-0.017 (0.012)
Mean of Dependent Var.	0.442	0.442	0.442	0.442	0.442	0.442
Observations	50,577	50,577	50,577	50,577	50,577	50,577
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: as shown in Table 2.7 in the main paper.

Chapter 3

Climate Shocks and the Limits of Labor Market Adaptation

Abstract. Cities in developing countries are critical engines of growth but face heightened vulnerability to climate change, particularly from increasingly frequent extreme rainfall events. Rapid urban expansion has often outpaced infrastructure investment, raising concerns that extreme weather events could exacerbate existing inequalities. We study how urban workers adapt to extreme rainfall using high-frequency labor market data from a location-based platform in Delhi. Earnings fall by 4 percent on rain shock days, driven by declines in labor supply rather than customer demand. Workers partially smooth income losses through intertemporal substitution but exhibit no spatial adaptation. The effects are highly unequal: workers in peripheral, poorer neighborhoods experience income losses three times larger than those in wealthier areas. Infrastructure quality, particularly drainage and metro connectivity, is positively associated with neighborhood resilience. Our findings suggest that without deliberate adaptation, climate change will reinforce existing spatial inequalities and limit the economic potential of developing cities.

3.1 Introduction

Cities generate over 80 percent of global GDP and are widely regarded as engines of growth, particularly in the developing world (World Bank, 2009; Glaeser, 2012). Yet rapid, unplanned urban expansion has left many cities highly vulnerable to climate change. Extreme rainfall, in particular, threatens to overwhelm infrastructure, disrupt mobility, and exacerbate existing spatial inequalities (IPCC, 2023). However, the economic consequences of extreme rainfall in developing cities remain understudied (Kocornik-Mina et al., 2020; Gandhi et al., 2022). Cities are often overlooked in the literature on climate and weather impacts in developing countries (Kala et al., 2023; Bryan et al., 2025), while climate risk remains peripheral in much of the research on developing cities (Bryan et al., 2020).

We study the economic costs of extreme rainfall, the scope for adaptation, and the distributional consequences of weather shocks in the context of location-based platform workers in Delhi, the second most populous city in the world. Location-based platform work¹ is a particularly useful labor market technology for weather adaptation, as it offers temporal and spatial adaptation options that are typically unavailable to informal workers in developing countries (Bobba et al., 2025). In addition to this flexibility, these systems are designed to optimize spatial localization, potentially reducing the disadvantages faced by workers in peripheral areas (Ghili et al., 2025). In this sense, the platform represents a best-case scenario for adaptation. If meaningful negative effects persist even in this context, it would suggest that deeper structural constraints – such as residential segregation, infrastructure gaps, or uneven demand – remain binding. Studying this labor market also provides access to high-frequency, spatially granular data that is rarely available in developing city contexts.

We find that worker earnings decline by approximately 4% on days with extreme rainfall, with effects persisting for up to two days after the shock. These losses are primarily driven by a reduction in labor supply rather than customer demand: the number of job offers remains unchanged or increases slightly, while acceptance rates fall. Workers partially offset these losses by reallocating labor across time – completing more jobs in the days before and after heavy rain. Summing the dynamic effects yields point estimates of a 3% reduction in earnings and a 4% decline in jobs over a ten-day window, though these estimates are imprecise and not statistically significant. Workers’ increased labor supply in the days preceding extreme rainfall events suggests anticipatory behavior in response to short-term forecasts or early warnings. However, we find no evidence of spatial adaptation as workers do not shift toward closer jobs on rain shock days.

The average effects mask substantial spatial heterogeneity. To quantify this, we estimate rainfall impacts separately for each residential neighborhood, defined by the grid cell where workers live. We then compare the distribution of these estimated effects across space. Income losses for workers in neighborhoods at the 75th percentile of the impact distribution are 3.5 times larger than those in the 25th percentile. The most severely affected workers tend to live in peripheral areas that are poorer, less dense, less connected by metro transit, and have weaker drainage infrastructure. For these workers, rainfall reduces daily income

¹Location-based platform work refers to tasks mediated by digital platforms that must be performed at a specific physical location, typically involving services such as delivery, transportation, or on-site household tasks (ILO, 2021)

by over 6%, compared to a 2% decline among workers living near the urban core. We find suggestive evidence that metro accessibility may help mitigate these losses, underscoring the role of local infrastructure in shaping economic resilience to weather shocks.

Taken together, our results show that even in a flexible, platform-mediated labor market, extreme rainfall leads to meaningful and unequal income losses, with limited scope for individual adaptation. While workers are able to partially adjust through intertemporal substitution, they are unable to adapt spatially. The concentration of losses in poorer, less connected areas highlights how rainfall shocks amplify existing urban inequalities. Targeted infrastructure investments, particularly in peripheral neighborhoods, may offer disproportionate gains in both resilience and equity.

This paper contributes to several strands of literature. First, it adds to the limited body of work on climate and weather costs and adaptation in urban areas of developing countries (Kocornik-Mina et al., 2020; Gandhi et al., 2022). While much of the existing literature focuses on rural settings or agricultural outcomes (Kala et al., 2023), far less is known about how weather shocks affect labor markets in cities characterized by informal employment and infrastructural constraints (Bryan et al., 2025). Most empirical studies of rainfall shocks emphasize rural impacts and rural-to-urban migration as an adaptation strategy (Black et al., 2011), leaving urban labor market dynamics themselves relatively understudied.

Second, the paper contributes to the literature on inequality within cities (Brueckner et al., 1999; Glaeser et al., 2009; Fogli and Guerrieri, 2019; Couture et al., 2024), and specifically to research on how policies and infrastructure shape spatial inequality in developing cities (Balboni et al., 2020; Tsivanidis, 2022). The findings show that pre-existing disparities in infrastructure quality amplify the economic costs of weather shocks, suggesting that climate change is likely to exacerbate spatial inequality in urban labor markets. The paper also highlights how the emergence of digital labor platforms provides new spatially granular data that can be leveraged to study these dynamics and distributional consequences in settings where such data is typically scarce.

Third, the paper contributes to the literature on labor market responses to environmental shocks. While a growing body of research examines how workers adjust to extreme heat, pollution, and other environmental risks, evidence on adaptation to rainfall shocks in urban labor markets remains limited. Additionally, existing work has largely focused on sectors such as construction or agriculture (Graff Zivin and Neidell, 2014; Hanna and Oliva, 2015; Hoffmann and Rud, 2024; Papp, 2024). This paper focuses on rainfall and studies its effects on service-sector workers in a large developing city – a setting that remains understudied. The high frequency of the data allows us to trace short-run earnings dynamics and uncover patterns of intertemporal substitution at the daily level. We also document a novel adaptation margin, potentially enabled by the digital infrastructure of the platform: anticipatory labor supply responses based on weather forecasts or early warnings.

Finally, the paper contributes to the empirical literature on gig-economy labor markets and the role of flexibility in alternative work arrangements. Prior work has primarily focused on platform pricing, allocation efficiency, and surplus distribution in response to shocks and reputational dynamics (Dubé and Misra, 2023; Hall et al., 2023; Papp, 2024). We shift the focus to the spatial organization of platform work and examine whether flexible, platform-mediated labor arrangements can overcome the structural features of urban environments when optimizing for spatial localization (Ghili et al., 2025). The results also highlight

a new dimension of flexible work arrangements (Mas and Pallais, 2017): the ability to intertemporally substitute labor supply in response to short-term disruptions. More broadly, the findings shed light on the benefits and limits of gig work as a form of climate adaptation, an increasingly relevant issue as the share of gig workers grows globally (Tubaro et al., 2025), particularly in developing countries (NITI-Aayog, 2022).

Our paper is organized as follows. Section 3.2 provides background on the platform labor market we study and the broader metropolitan context of Delhi. Section 3.3 describes the data sources used in the analysis. Section 3.4 presents descriptive statistics on worker activity, rainfall shocks, and the spatial organization of economic activity in the city. Section 3.5 outlines our empirical strategy. Section 3.6 presents the main results, and Section 3.7 analyzes how impacts vary across space and examines the geographic, infrastructural, and socioeconomic factors associated with the magnitude of income losses. Section 3.8 concludes.

3.2 Context

Our worker data comes from a leading digital platform for on-demand home services in India (“the platform”). It connects customers with independent service professionals such as electricians, plumbers, beauticians, and cleaners through an app-based interface. All services are performed indoors, so extreme weather is expected to affect workers primarily through transportation and mobility constraints, rather than through job performance conditional on arrival. As of 2024, this platform operates in over 50 cities, engaging more than 50,000 workers and serving approximately 12 million customers. Its business model reflects a broader global trend toward platform-mediated labor, particularly in developing countries, where digital intermediation is increasingly used to organize informal labor (NITI-Aayog, 2022; Datta et al., 2023).

Jobs on the platform are initiated by customer requests, which are routed through a centralized matching system. Workers receive these offers through a mobile application and can choose whether to accept them. There is no dynamic pricing; prices are fixed by service type and occasionally adjusted across time or cities by the platform itself. This pricing structure removes concerns about price endogeneity influencing labor demand and supply fluctuations. The platform also monitors key performance metrics such as acceptance rates and customer ratings. Workers are expected to maintain minimum thresholds on these metrics, and repeated cancellations or prolonged inactivity can lead to temporary restrictions on platform access. For example, exceeding a set number of cancellations in a given month may result in a temporary suspension. The platform company also provides onboarding and extensive training for workers.

The platform’s use of digital infrastructure allows for high-frequency observation of worker location, job acceptance, and movement across the city, enabling the empirical analyses of spatial and temporal patterns of labor supply adjustment to weather shocks such as rainfall.

3.2.1 Comparisons to the broader urban Indian labor market

To assess the external validity of our findings, we situate the workers in our sample within the broader context of India’s urban labor market. We proceed in two steps. First, we present evidence showing that the income levels, work intensity, and commuting patterns of workers in our sample are broadly similar to those of non-platform urban workers engaged in traditional occupations. Second, we document that the platform-based service occupations we study constitute a growing and economically relevant share of urban employment.

A defining institutional feature of platform-mediated gig work is the nature of the contractual relationship between workers and the platform. Workers are typically classified as independent contractors rather than employees, which shapes both their rights and their exposure to risk. They typically lack formal contracts, job security, or access to employment-linked social protections such as paid leave or health insurance – core features that define informal work in these labor markets (Canagarajah and Sethuraman, 2001). Thus, the working conditions of workers in our sample mirror that of informal employment, which accounts for 66% of the India’s urban workforce (Waghmare, 2024).

Workers in our dataset earn a median monthly income of approximately \$230 (Rs. 20,000), placing them between the 50th and 75th percentiles of the urban earnings distribution based on nationally representative data (PLFS, 2021). Workers complete an average of 1.17 jobs per day, suggesting commutes to and from one place of work similar to traditional occupations. Conditional on having joined the platform before the start of the year, the median worker does at least one job on the platform 174 days per year. Note that this reflects a combination of labor demand and supply factors. While this is lower than typical work days for urban workers in India, our results are consistent if we restrict our sample to workers who do at least one job 70% of the days annually.

The geographic scope of work also aligns closely with broader commuting patterns in developing cities. The median distance traveled per job is 5.6 kilometers, which is comparable to nationally representative estimates of urban travel distance to work from the NSS 76th round (Bhatt et al., 2020) and to commute distances documented in other large cities in developing countries (Kreindler and Miyauchi, 2023; Kreindler, 2024). Although platform workers face variability in their daily work locations compared to traditional fixed workplaces, repeated assignments within familiar neighborhoods and reliance on digital navigation tools significantly reduce the associated coordination and search costs (Phuangsuwan et al., 2024).

Additionally, gig work itself is a meaningful segment of the labor market. As of 2020–21, approximately 7.7 million workers were engaged in gig work in India, comprising 7.6 percent of the urban labor force (NITI-Aayog, 2022; Roy et al., 2022). This figure is projected to rise to 23.5 million by 2029–30, reflecting the increasing role of digital platforms in mediating service-sector employment. The types of services provided by workers in our data are similar to those found in other urban service occupations, including transport, delivery, and home-based repairs. These occupations collectively account for a 2% share of urban employment (PLFS, 2021).

In sum, the workers in our sample reflect the broader characteristics of India’s urban labor force, where informal employment remains the dominant form. Their contractual arrangements, income levels, work intensity, and commuting patterns are all broadly consistent with those observed in nationally representative data. Moreover, the services

they provide fall within a segment of the urban economy that is expanding rapidly. As such, the responses we document offer meaningful insights into how a large and growing share of urban workers in developing cities may experience and adapt to climate-related disruptions (Papp, 2024).

3.2.2 Delhi-National Capital Region (Delhi-NCR)

Delhi-NCR is India’s capital and most populated region, with a metropolitan population exceeding 30 million. It is a major economic hub, contributing over 4% of national GDP, and exhibits many of the defining characteristics of rapidly growing megacities in developing countries such as high population density, large informal labor markets, and sharp spatial inequalities. The city’s physical growth has been marked by extensive low-density horizontal expansion (Jain et al., 2013).

Residents of the area face numerous severe challenges during periods of heavy rainfall, transforming the city’s landscape into a hazardous environment. The national capital routinely experiences widespread waterlogging within 20–30 minutes of heavy rain. The drainage system remains outdated, as it was designed in 1976 for a population of 6 million and rainfall of up to 50 mm, now inadequately serving over 25 million residents experiencing rainfall exceeding 150 mm (The Sunday Guardian Live, 2025). The consequences are multifaceted and often tragic, with reports of multiple fatalities during heavy rain events, including drownings in waterlogged drains and underpasses (The Economic Times, 2025). In one instance, at least 11 people died from rain and flooding in a single week, with four drowning in submerged underpasses (CNN, 2024). Transportation networks collapse under these conditions, with severe traffic congestion paralyzing movement throughout the city, while vehicles break down on submerged roads, leading to long delays for commuters (Business Standard, 2025).

3.3 Data

We combine multiple high-resolution data sources to examine the labor market effects of rainfall shocks on platform workers in Delhi-NCR. The analysis is conducted at the worker-date level, using detailed records from a large location-based service platform, combined with geospatial, meteorological, and infrastructure data at the neighborhood grid level. Below, we describe each data source and how they are integrated for analysis.

3.3.1 Worker Data

We use high-frequency transaction-level data from the platform’s internal systems to construct a panel of worker activity at the daily level. The data capture all job-level interactions between workers and customers, including total earnings, number of jobs completed, and distance traveled each day. Each worker is linked to a fixed home location and a primary occupation category. For the purposes of our analysis, we treat home location as time-invariant. Using supplemental GPS data collected by the platform, we confirm that residential relocation is rare over the study period. Worker locations are recorded

using Google Plus Codes, which we map to $0.05^\circ \times 0.05^\circ$ grid cells to merge with spatial covariates and rainfall data.

We aggregate the data to the worker-by-date level, which forms the primary unit of analysis for much of the empirical work in this paper. To capture both participation and intensity margins of labor supply, we fill in zeros for days on which a worker performs 0 jobs between their first and last observed job of the year. This approach results in a complete panel of worker activity over time. The richness and granularity of the data allow us to observe spatial and temporal responses to weather shocks with precision, and to study heterogeneity in those responses by worker location and job type.

Our analysis focuses on the period from January 1, 2023 to November 1, 2024. While platform transaction data are available for earlier periods, we restrict attention to this window for two main reasons. First, demand-side data—specifically, job offers and acceptance rates—are only available beginning in 2023. Second, key outcome variables such as distance traveled per job were not systematically recorded by the platform prior to this period. To ensure consistency across all components of the analysis and maintain a common sample for worker, demand, and spatial data, we limit our study to this two-year window.

3.3.2 Platform Demand Data

In addition to realized jobs, we use platform-generated records on job offers to observe the demand side of the marketplace. These data consist of customer requests that were algorithmically assigned to individual workers. Each record corresponds to a worker-customer interaction, allowing us to observe the number of offers received and whether those offers were accepted.

By aggregating these interactions to the worker-by-date level, we construct measures of job offers (a proxy for labor demand) and acceptance rates (a proxy for labor supply responsiveness). These measures are critical for disentangling whether observed declines in completed jobs during rainfall shocks are driven by reduced worker availability or changes in labor demand.

3.3.3 Rainfall Data

To measure rainfall exposure, we use daily gridded precipitation estimates from CHIRPS v2.0, a high-resolution (0.05°) global rainfall dataset (Funk et al., 2015). CHIRPS combines satellite imagery with in-situ rain gauge observations to produce spatially and temporally consistent rainfall estimates, particularly in data-sparse regions such as South Asia. Each worker is assigned rainfall exposure based on the grid in which they reside.

We define a rainfall shock at the grid-day level as any instance in which precipitation exceeds 10 millimeters. We adopt this absolute threshold rather than defining shocks as deviations from long-term means, as is common in the literature, for several reasons. First, while rainfall deviations from historical averages are relevant for modeling agricultural productivity, which is the focus of most of the existing literature studying rainfall shocks (Jayachandran, 2006), in urban settings an absolute threshold is more appropriate for capturing the strain placed on infrastructure. Second, the use of a fixed threshold ensures comparability across grids and over time, allowing for a consistent definition of shocks. In

our data, the 10mm threshold corresponds approximately to the 90th percentile of daily rainfall for grids, capturing moderate to heavy precipitation events. These are the types of events most likely to disrupt mobility and reduce economic activity in dense urban environments such as Delhi (Gupta, 2017).² We check for robustness by adopting different definitions of the shock and find results consistent with our current analysis.

3.3.4 Infrastructure Data

To measure the quality and density of local infrastructure, we extract infrastructure data from OpenStreetMap as of January 1, 2023 (OpenStreetMap Contributors, 2023), and rasterize it to estimate infrastructure characteristics at the grid level. Specifically, we focus on three infrastructure measures relevant to our analysis: (i) drainage density, defined as the total length of mapped drains within a grid; (ii) road density, calculated as the cumulative length of roads in each grid; and (iii) metro accessibility, measured by the number of metro stations located within the grid. We examine whether variation in these infrastructure indicators is associated with differences in the severity of rainfall shock impacts.

3.3.5 Geographic and Topographic Data

We calculate each grid’s radial distance from the geometric center of Delhi to capture spatial structure. To account for natural variation in elevation, we incorporate digital elevation data from the Shuttle Radar Topography Mission (SRTM) (NASA Shuttle Radar Topography Mission, 2013). These variables are used to analyze spatial gradients in vulnerability and to control for geographic features that may influence both infrastructure placement and exposure to rainfall shocks.

3.3.6 Wealth Index Data

We incorporate the Meta Relative Wealth Index (RWI), a satellite- and connectivity-based estimate of household wealth at high spatial resolution (Chi et al., 2022). Since each grid contains multiple RWI observations, we compute the grid-level wealth measure as the average of all available points within the grid. We standardize the RWI across our sample and use it to capture spatial variation in socioeconomic status. This variable allows us to assess whether poorer neighborhoods experience greater vulnerability to rainfall shocks and whether pre-existing inequalities are amplified by weather-related disruptions.

3.3.7 Weather Forecast Data

To assess the availability and accuracy of advance weather information, we incorporate historical rainfall forecast data from OpenMeteo (Zippenfenig, 2023). These forecasts provide daily predictions of precipitation at various lead times, which we match to realized rainfall at grid points in the Delhi-NCR region. Coverage is limited to the year 2024 and available

²In principle, one could validate this threshold using localized inundation or flood incidence data. However, such high-frequency, spatially disaggregated data are not currently available for Delhi (Patel, 2023).

at a coarser spatial resolution than CHIRPS. This forecast data allows us to get suggestive evidence related to the informational environment facing workers and evaluate whether labor supply responses differ for shocks that were anticipated versus those that were not.

3.4 Descriptive Statistics

Before presenting the main results, we summarize key characteristics of the platform workers in our sample, their spatial distribution across Delhi, and the infrastructure and socioeconomic conditions of the neighborhoods where they reside. This descriptive evidence provides important context for interpreting our results.

3.4.1 Labor Market

Table 3.1 presents summary statistics on worker activity, aggregated at the daily, weekly, and monthly levels. On average, workers complete 1.19 jobs per day, or 1.97 jobs conditional on working. The average distance traveled per job is 7.14 kilometers, and the mean earnings per job are Rs. 749, resulting in a daily income of Rs. 770.58. Workers receive approximately 2.29 job offers per day, with an average acceptance rate of 68 percent. Same-day offers, defined as jobs where the request and delivery occur on the same calendar day, account for 1.18 offers per day and contribute 47 percent of daily earnings.

At the weekly level, workers earn an average of Rs. 5,142 across 7.83 jobs, working 3.7 days per week. Monthly aggregates indicate average earnings of Rs. 21,237 and 32 completed jobs over 15.3 active days. The sample includes 9,779 unique workers.

Figure 3.1 shows the distribution of workers across service categories. Throughout the paper, we use the terms “service category” and occupation interchangeably. Salon services for women and cleaning account for the largest shares, each representing about 30 percent of workers. The remaining 40 percent is distributed across categories such as spa services, men’s grooming, appliance repair, electricians, carpenters, plumbers, massage therapy, party makeup, and other personal services. Some services are gender-specific; for example, salon services for women and men’s grooming are typically provided by same-gender workers.

3.4.2 Rainfall Shocks

As discussed earlier, we define a rainfall shock at the grid-day level as any instance in which precipitation exceeds 10 millimeters. This provides a consistent definition of shocks across grids. In theory, one could estimate the effect of rainfall shocks using a two-way fixed effects regression with grid and time fixed effects, comparing outcomes for workers in shocked versus non-shocked grids. However, rainfall patterns in Delhi offer limited cross-sectional variation across grids. As shown in Figure 3.2, the proportion of grids receiving a shock on any given day tends to oscillate between 0 and 1 – that is, most grids are either simultaneously shocked or not shocked.³ Additionally, 78% of all shock-day observations fall

³This contrasts with the more spatially dispersed rainfall patterns observed in other Indian cities, such as Mumbai (Bearpark et al., 2025).

within the monsoon months of July, August, and September. Overall, approximately 10% of worker-day observations in our data are classified as rainfall shocks.

3.4.3 Economic Geography of the City: Distance, Infrastructure, and Inequality

The spatial organization of economic activity within cities plays a central role in shaping both the feasibility of adaptation to climate shocks and the distributional incidence of those shocks. Spatial structure determines not only access to infrastructure and services, but also the proximity to economic opportunities (Duranton and Puga, 2004). In this subsection, we examine how this spatial structure operates in Delhi, and how it is reflected for workers on the platform labor market. Because platform work allows for decentralized task matching and flexible mobility, it is often viewed as a promising technology for mitigating spatial frictions. Studying its operation in a city like Delhi therefore offers insight into how far such systems can go in overcoming structural constraints. To characterize urban space, we use distance from the city center—a foundational concept in urban economics and the core organizing principle behind classic models of spatial structure and urban form (Krugman, 1991).

In Delhi, urban development has followed a pattern of rapid, low-density horizontal expansion, widely described as urban sprawl (Jain et al., 2013). This pattern is reflected in a steep decline in employment density with increasing distance: employment within 10 kilometers of the city center is roughly twice as dense as in the 20-kilometer band, and nearly ten times that in the 40-kilometer band (Mathews et al., 2023). While the platform’s demand is more likely to be based on residential rather than employment density due to the nature of the services, we find similar patterns. The density of customer requests within 10 kilometers of the center is 1.25 times that in the 20-kilometer band, and three times that in the 40-kilometer band.

To systematically characterize the spatial environment in which workers operate, we divide the grid cells where workers reside into terciles based on their radial distance from the city center. Workers in Tercile 1 live closest to the urban core, while those in Tercile 3 are located in the most peripheral areas. Because the data include the customer location associated with each job offer, we are able to map both worker residences and their corresponding service destinations. Figure 3.3 displays these patterns, disaggregated by worker residence tercile. While job locations are generally closer to workers’ home grids within each tercile (as indicated by the red areas on the customer heatmaps), the figure reveals that workers must still travel across a wide spatial range to complete jobs. Despite the platform’s emphasis on spatial optimization, underlying residential and demand patterns constrain the degree of localization achievable in practice.

We next examine work patterns across the distance terciles. Table 3.2 reports average labor market outcomes by worker location. Workers in Terciles 2 and 3 complete slightly more jobs per day than those in Tercile 1, which may reflect higher economic need since these workers tend to live in poorer neighborhoods, as shown in Table 3.3. In terms of mobility, workers in Tercile 3 travel the longest distance per job (6% more than those in Tercile 1) indicating higher baseline commuting burdens. In contrast, workers in Tercile 2 travel the

least and earn the highest income per job, potentially reflecting a locational advantage from being situated between the dense demand of the core and less saturated markets in the periphery, or variation in service types or neighborhood characteristics.

Table 3.3 reports average neighborhood characteristics across distance terciles. For each grid, we standardize variables and compute the mean z-score within each tercile. Peripheral neighborhoods tend to have lower relative wealth, as measured by the Meta Relative Wealth Index (RWI) (Chi et al., 2022), indicating lower living standards on average. Infrastructure quality also declines with distance from the city center. This pattern holds across all three infrastructure measures used in our analysis—drainage density, road density, and metro accessibility (defined in Section 3.3.4), suggesting that spatial disparities in public investment mirror underlying differences in socioeconomic status.

Overall, these patterns suggest that workers living in peripheral areas face a constellation of structural disadvantages that may constrain their ability to adapt to rainfall shocks. They travel longer distances to reach jobs, live in poorer neighborhoods with fewer private resources, and are embedded in areas with systematically weaker infrastructure. These disadvantages may amplify the impact of mobility disruptions during heavy rainfall, both by raising the cost of work and by limiting feasible adaptation strategies. While platform work introduces flexibility in when and where individuals work, its ability to mitigate the unequal impacts of climate shocks across neighborhoods ultimately depends on whether it can overcome the spatial and infrastructural constraints that disproportionately affect workers in more disadvantaged areas. We now formally estimate the causal effects of rainfall shocks on labor market outcomes and assess whether these impacts vary across workers in different parts of the city.

3.5 Empirical Strategy

Our objective is to estimate the effect of rainfall shocks on labor market outcomes. To isolate plausibly exogenous variation in daily realized rainfall, our empirical specification includes a rich set of controls such as worker fixed effects, occupation-by-week fixed effects, and day-of-week fixed effects that flexibly absorb time-invariant worker characteristics, seasonal variation in platform activity by service type, and calendar-based trends (Papp, 2024; Hoffmann and Rud, 2024).

Our baseline empirical specification estimates the impact of daily rainfall shocks on same day daily labor market outcomes using the following regression:

$$Y_{it} = \alpha + \beta \cdot \text{Shock}_{g(i)t} + \omega_i + \delta_{c(i)w(t)} + \phi_{d(t)} + \varepsilon_{it} \quad (3.1)$$

In this equation, Y_{it} denotes a labor market outcome for worker i on day t , such as the number of jobs completed, total earnings, or distance traveled per job. $\text{Shock}_{g(i)t}$ is a binary indicator equal to one if the geographic grid $g(i)$ where worker i resides experienced a rainfall shock on day t , defined as precipitation exceeding 10mm in total daily accumulation.

The coefficient of interest, β , captures the average effect of a rainfall shock on the outcome variable. We include worker fixed effects (ω_i) to absorb all time-invariant heterogeneity across individuals, such as occupation, baseline productivity, or residential location. To flexibly

control for seasonal and service-specific variation in demand, we include occupation-by-week fixed effects ($\delta_{c(i)w(t)}$). Day-of-week fixed effects ($\phi_{d(t)}$) account for systematic differences across weekdays and weekends. Standard errors are clustered at the geographic grid level, which is the level at which we measure rainfall, to account for spatial correlation in treatment assignment and potential serial correlation within workers living in the same grid over time.

We assess heterogeneity in labor supply and demand by the time sensitivity of the service provided by the worker. We classify workers who work in the categories of Electricians, Plumbers, Water Filter Repair, and Party Makeup as ‘Urgent’ workers, and other categories as ‘Non-Urgent’. We estimate the regression:

$$Y_{it} = \alpha + \beta_1 \cdot \text{Shock}_{g(i)t} + \beta_2 \cdot \text{Shock}_{g(i)t} \cdot \text{NU}_i + \omega_i \cdot \text{NU}_i + \delta_{c(i)w(t)} \cdot \text{NU}_i + \phi_{d(t)} \cdot \text{NU}_i + \varepsilon_{it} \quad (3.2)$$

where $\text{NU}_i = 1$ if worker i ’s occupation category is classified as ‘Non-Urgent’. The fixed effects are interacted with NU_i to obtain consistent estimates.

To examine the degree of intertemporal substitution in labor supply around the day of the rainfall shock, we estimate the following dynamic specification:

$$Y_{it} = \alpha + \sum_{\tau=-9}^9 \beta_{\tau} \text{Shock}_{g,t-\tau} + \omega_i + \delta_{c(i)w(t)} + \phi_{d(t)} + \epsilon_{it} \quad (3.3)$$

This model replaces the single shock indicator with a series of binary variables indicating whether a rainfall shock occurred in the worker’s home grid on each day τ relative to date t , where $\tau \in [-9, 9]$. The coefficient β_{τ} captures the change in the outcome variable τ days before or after a shock, relative to non-shock days. As in the baseline specification, the model includes individual fixed effects, occupation-by-week fixed effects, and day-of-week fixed effects. This setup allows us to trace the temporal dynamics of labor supply, testing for anticipatory behavior prior to shocks and compensatory behavior in the days following.

To assess spatial heterogeneity in the effects of rainfall shocks, we estimate Equation 3.1 separately for each geographic grid in the Delhi-NCR region. This allows us to recover a distribution of shock coefficients $\hat{\beta}_g$, capturing localized responses to rainfall across space. We restrict the analysis to grids with at least five workers who reside in the grid during the estimation period to avoid estimating effects in grids with limited worker representation. Standard errors in this grid-level analysis are clustered at the individual level. We then correlate the estimated grid-specific coefficients with observable characteristics of the grid, including infrastructure quality, average wealth, geographic features, and land use classifications, to identify the factors that predict the magnitude of the impacts of rainfall shocks.

3.6 Results

Table 3.4 presents estimates from Equation 3.1, which quantifies the average effect of daily rainfall shocks on daily labor market outcomes for platform workers. Across all specifications, we include worker fixed effects to account for time-invariant heterogeneity in productivity or preferences, as well as occupation-by-week and day-of-week fixed effects to flexibly control

for seasonality in labor demand and supply.

Column (1) shows that earnings are 30.1 rupees lower on rain shock days, a decline of 3.8 percent relative to the mean of Rs. 798 of average income on days without extreme rain. This estimate is precise and statistically significant at the 1 percent level. Column (2) reveals that the reduction in earnings is accompanied by a 4.5 percent decline in the number of jobs completed, indicating a contraction in labor supply or disruptions to task completion. Column (3) tests whether workers are able to adapt spatially by shifting toward jobs closer to home on shock days. The estimate is small and precisely estimated, allowing us to rule out any economically meaningful adaptation along this margin.

While the results above establish that platform workers experience reduced labor market activity on heavy rain days, they do not by themselves reveal whether these effects are driven by lower customer demand or reduced worker participation. Distinguishing between these two channels is critical for interpreting the nature of vulnerability in this labor market. If adverse weather primarily depresses demand, the economic losses are likely driven by customer-side constraints. In contrast, if the effects are rooted in lower labor supply, this would suggest that workers face higher disutility from working in times of high rain. This disutility may stem from increased commuting costs due to congestion and deteriorated driving conditions, or from a direct reduction in the utility of working during rain, such as discomfort (Farber, 2015).

To shed light on this distinction, Table 3.5 presents effects of rainfall shocks on two intermediate outcomes: the number of job offers a worker receives, and the share of those offers they accept. These outcomes serve as reduced-form proxies for the behavior of customers and workers, respectively, and allow us to assess whether the observed decline in labor activity reflects a supply-side or demand-side response.

Column (1) shows that rainfall shocks lead to a small and precisely estimated (0.45%) increase in the number of job offers. This pattern is inconsistent with a substantial contraction in customer-side demand. In contrast, column (2) documents a 1.2% decline in the offer acceptance rate, from a control mean of 68 percent. These results indicate that the decline in completed jobs is not due to a drying up of demand, but rather reflects a contraction in labor supply. Workers appear less able or less willing to accept tasks on days with rainfall shocks, suggesting that the adjustment costs of participation remain nontrivial despite access to the platform technology.

3.6.1 Heterogeneity by Time Sensitivity of Occupation Category

To better understand the labor demand and supply responses, we examine heterogeneity by the time sensitivity of the service provided by the worker, distinguishing between “urgent” occupations, such as party makeup and plumbing, and “non-urgent” services like home cleaning or beauty treatments. This distinction is motivated by the idea that both customer demand and worker behavior may respond differently depending on the immediacy of service needs. In particular, customers may be less willing to postpone urgent services, and workers may have stronger incentives to complete such jobs despite adverse weather (Papp, 2024).

Results from Equation 3.2 are presented in Table 3.6. The coefficient in the first row of each column captures the effect of rainfall shocks for workers in urgent service occupations, while the second row reports the differential effect for those in non-urgent categories. We find

that the number of jobs completed declines similarly across both groups (Column (1)), despite an increase in job offers for urgent occupations (Column (2)). This divergence between demand and realized outcomes indicates that labor supply constraints remain binding even when demand persists. Acceptance rates fall in both groups, but the decline is attenuated among workers in urgent occupations, consistent with an endogenous supply response to elevated demand. Taken together, the results suggest that while demand conditions vary by occupation type, rainfall shocks induce a contraction in labor supply across the board.

3.6.2 Intertemporal Substitution

We next examine whether workers adjust their labor supply intertemporally in response to rainfall shocks. Intertemporal substitution is a well-established adaptation mechanism in the context of environmental shocks (Graff Zivin and Neidell, 2014; Hanna and Oliva, 2015; Aragón et al., 2017; Hoffmann and Rud, 2024). Compared to traditional jobs, platform work may offer a stronger scope for such adjustment, as it enables high-frequency scheduling flexibility with minimal administrative or contractual frictions. This makes it a particularly promising technology for intertemporal adaptation. We therefore assess whether the earnings losses documented in the previous section are partially offset by labor supply adjustments before or after the shock, providing evidence on the extent to which platform-mediated work supports intertemporal responses to weather variability.

Figure 3.4 plots the dynamic coefficients from Equation 3.3, estimating the impact of rainfall shocks on daily earnings (Figure 3.4a) and labor supply (Figure 3.4b) in the days before and after the event. Earnings decline sharply on the day of the shock, replicating the contemporaneous effects documented earlier. In the two to four days preceding the shock, we observe a statistically significant increase in earnings and jobs done. These anticipatory adjustments suggest that workers are able to reallocate labor supply in advance, likely in response to short-term weather forecasts or early indicators of deteriorating conditions. We return to this mechanism and provide supporting evidence below.

Negative impacts persist for one to two days following the rainfall event, consistent with pluvial flooding and prolonged waterlogging observed in Delhi after episodes of heavy rain (Mehtab and Kamal, 2023). After this period, we find evidence of workers compensating for lost earnings and labor supply in the following days.

Overall, the estimates point to meaningful intertemporal substitution: workers shift labor supply across days to partially smooth earnings in the face of rain-related disruptions. This behavior underscores the potential of platform work to facilitate adaptation at the daily margin, consistent with its technological design and low adjustment frictions.

Anticipatory Effects and Forecasts

While compensating for labor supply after a weather related shock has been documented previously (Connolly, 2008; Hoffmann and Rud, 2024), the anticipatory response we find before the effect is novel, especially for high frequency shocks (Weill, 2023). This raises an important question: how do workers anticipate upcoming rainfall shocks in our setting? In the absence of institutional scheduling or managerial oversight, any systematic increase in labor supply prior to rain shocks must originate from workers' ability to predict or perceive

adverse conditions. In this section, we present three pieces of evidence supporting the hypothesis that anticipatory responses are driven by weather forecasts or early signals of rain. First, we find that days with extreme rain tend to be preceded by light rain, which possibly serves as an early warning sign. Second, we document that the hit rate for different thresholds of rainfall gets fairly accurate three to four days before the shock. Lastly, for a subset of our data, we are able to classify our shocks as forecasted or unforecasted using the forecast data, and we find that the anticipatory response is largely absent for unforecasted shocks.

Figure 3.5a provides suggestive evidence that workers receive timely information about impending heavy rainfall. The figure estimates the probability of any rain occurring on a given day, conditional on the rain shock (defined as rain $\geq 10\text{mm}$) occurring on day 0. In the three to four days before the shock, the likelihood of light rain increases. This pattern indicates that workers are exposed to gradual increases in rainfall prior to the event, which may serve as a warning signal and prompt forward-looking behavior.

Complementing this, Figure 3.5b assesses the quality of rainfall forecasts using historical forecast data matched to realized outcomes for locations around Delhi using data from OpenMeteo (Zippenfenig, 2023). The hit rate is high even for moderate thresholds: forecasts issued three days in advance correctly predict rainfall above 10mm approximately 60 percent of the time, and above 5mm approximately 80 percent of the time with accuracy improving closer to the event. These data indicate that rain forecasts are not only available but reasonably informative, providing a plausible channel through which workers may adjust behavior in advance.

Finally, Figure 3.6 compares dynamic treatment effects (Equation 3.3) on jobs delivered for shocks that were forecasted versus those that were not for a subset of our data.⁴ For forecasted shocks, the drop in jobs on the shock day is preceded by a clear rise in job completion in the preceding days, consistent with anticipatory reallocation. By contrast, unforecasted shocks exhibit no such pre-shock increase and a more severe drop in labor supply on the event day. This divergence provides direct evidence that forecast quality shapes labor supply dynamics, reinforcing the interpretation that forward-looking behavior is a function of information availability.

Together, these results suggest that short-term forecasts and early warning signals enable workers to adjust labor supply in anticipation of adverse weather conditions. These findings complement recent evidence on the role of long-range seasonal forecasts in agricultural settings (Burlig et al., 2024; Innovation Commission, 2024), by demonstrating that high-frequency, short-horizon forecasts can also support adaptive behavior in urban labor markets.

⁴Open-Meteo forecast data is only available for 2024 and is less granular than the CHIRPS data we use for our main analysis. Thus, we are only able to link it for one region, Noida, for which the Open Meteo actual data matches CHIRPS data

3.7 Spatial Heterogeneity in the Impacts of Rainfall Shocks

A growing body of research highlights that the economic costs of climate change may exacerbate spatial inequality. However, most existing work has examined distributional consequences at the national or regional level (Desmet and Rossi-Hansberg, 2024), with far less attention to how these impacts vary across neighborhoods within developing cities.⁵ Rainfall shocks may have particularly pronounced distributional consequences within cities due to the mobility disruptions they create through waterlogging and flash flooding. These effects are mediated by substantial heterogeneity in infrastructure quality: poorer neighborhoods often have weaker drainage systems, limited road connectivity, and inadequate flood mitigation infrastructure (Cutter et al., 2014; Kocornik-Mina et al., 2020). In addition, workers residing in peripheral areas may face larger disruptions due to longer commutes and higher baseline travel costs, making them more exposed to the consequences of flooding and congestion (Kilgarriff et al., 2019). As a result, the spatial incidence of rainfall-induced economic losses may systematically disadvantage already vulnerable populations within urban labor markets.

Platform work may be uniquely situated to mitigate spatial inequalities in exposure to rainfall shocks. Because spatial localization is a central design feature of many location-based gig platforms (Ghili et al., 2025), including the one analyzed in this paper, these systems can enable hyperlocal service delivery and reduce the spatial frictions that typically disadvantage workers in peripheral or poorly connected areas. However, the extent to which this flexibility enhances resilience depends on both the geographic distribution of demand and the mobility constraints workers face. In the presence of infrastructure deficits such as poor drainage, limited transit access, or congested roads, platform work may not fully insulate workers from location-based vulnerabilities. This section examines whether the impacts of rainfall shocks vary systematically across space and explores the geographic, infrastructural, and socioeconomic factors that account for this heterogeneity.

To quantify spatial heterogeneity in the impacts of rainfall shocks, we begin by estimating Equation 3.1 separately for each $0.05^\circ \times 0.05^\circ$ grid cell where workers in our data reside across the Delhi-NCR region. This yields a distribution of grid-specific coefficients $\hat{\beta}_g$, capturing localized earnings responses to rainfall. We then examine how these estimated impacts vary with observable grid-level characteristics, including distance from the city center, measures of infrastructure quality, socioeconomic status, and land use. This approach allows us to isolate geographic patterns in vulnerability and assess whether infrastructure and spatial disadvantage systematically predict the magnitude of economic losses from rainfall shocks.

3.7.1 Grid Level Impacts

Figure 3.7 presents a spatial distribution of the estimated grid-level coefficients, $\hat{\beta}_g$, across the Delhi-NCR region. Grids with negative point estimates are divided into terciles based on the magnitude of the estimated effect: the most negative tercile is shown in red, followed by yellow and green for progressively smaller losses. Grids with positive point estimates

⁵Hsu et al. (2021); Heilmann et al. (2021) study the unequal effects of extreme heat within US cities.

are shown in blue. Appendix Figure A1 demonstrates that when restricting to grids with estimates statistically significant at the 10 percent level, the positive coefficients disappear. This suggests that the observed positive point estimates are likely driven by noise.

We uncover substantial spatial heterogeneity. At the lower end of the distribution, the most exposed grids exhibit earnings losses exceeding Rs. 40 on rain days, with the most negative effect reaching nearly Rs. 270. By contrast, the least-affected grids experience small or near-zero effects. The interquartile range spans roughly Rs. 12 to Rs. 42 in daily earnings losses, indicating substantial dispersion in exposure to similar shocks.

The dispersion in estimated effects points to stark differences in how workers experience rainfall shocks based on where they live within the city, despite the platform’s emphasis on spatial optimization and hyperlocal task allocation. That such variation persists even in this setting suggests that deeper structural constraints such as the spatial distribution of demand, residential sorting, and infrastructure quality continue to shape exposure and limit adaptive responses. In the sections that follow, we examine the extent to which these factors predict the magnitude of income losses from rainfall shocks.

3.7.2 Role of Geography: Distance from City Center

Motivated by the spatial structure of economic activity described in Section 3.4.3, and by the heterogeneity evident in Figure 3.7, we examine whether the labor market impacts of rainfall shocks vary systematically with radial distance from the city center. We also examine if there is heterogeneity in intertemporal adaptation patterns documented earlier.

To assess whether the effects of rainfall shocks vary systematically with location within the city, we divide the grid cells where workers in our data reside into terciles based on their radial distance from the city center, as described in Section 3.4.3. We then estimate Equation 3.1 separately for each tercile to recover the average earnings loss associated with rainfall shocks for workers at different distances from the urban core. Because average earnings may differ across distance terciles - for example, workers in more central areas may have higher average incomes - we also show effects as a percentage of baseline income of workers in that tercile.

Figure 3.8 presents the results. Figure 3.8a shows the estimated losses in levels, while Figure 3.8b expresses these losses as a percentage of average daily income on non-shock days. The magnitude of income loss increases monotonically with distance from the center. Workers living closest to the city center (Tercile 1) lose Rs. 18.5 on rain days, whereas those living farthest from the center (Tercile 3) lose Rs. 51.3. When scaled by average earnings, the gradient remains: workers in Tercile 3 lose roughly 6 percent of daily income, compared to 2.8 percent in Tercile 1.

To test whether intertemporal substitution patterns for labor supply differ for workers in different distance terciles, we estimate Equation 3.3 for each tercile separately. In Figure 3.9, we find that these patterns are broadly similar for the different distance terciles. These patterns suggest that workers residing farther away from the city center face larger income losses on days of rainfall shocks, while their intertemporal adaptation patterns are similar.

3.7.3 Do Rainfall Shocks Exacerbate Economic Inequality Within Cities?

We test whether workers living in poorer neighborhoods experience larger income losses from rainfall shocks. To measure neighborhood wealth, we use the Meta Relative Wealth Index (RWI) (Chi et al., 2022), which combines satellite imagery and Facebook connectivity data, and has been validated against ground-truth estimates from the Demographic and Health Surveys. We standardize the index across our sample and estimate the following cross-grid regression:

$$\beta_g = \alpha + \gamma \text{RWI}_g + \epsilon_g \quad (3.4)$$

where β_g is the estimated grid-level income effect of rainfall shocks from Equation 3.1, and RWI_g is the standardized wealth index for that grid.

Figure 3.10 shows that a one standard deviation increase in grid-level relative wealth is associated with a Rs. 8 smaller income loss from rainfall shocks.⁶ In other words, workers living in poorer grids experience meaningfully larger economic losses on rain days. This finding suggests that rainfall shocks compound existing patterns of urban inequality in developing cities (Bryan et al., 2020).

3.7.4 What kind of infrastructure might mitigate income losses?

A central motivation of this paper is to examine to what extent the flexibility offered by platform work allows workers to adapt to rainfall shocks. The preceding sections show that, despite this flexibility, large disparities remain in how workers experience rainfall shocks. One plausible explanation is that physical infrastructure imposes binding constraints on mobility and access, especially during periods of waterlogging or localized flooding. In this section, we examine whether variation in local infrastructure is associated with differential vulnerability to rainfall shocks. We focus on three types of infrastructure that are likely to affect mobility and service access: drainage density, road density, and metro accessibility. These indicators are constructed at the grid level using OpenStreetMap data.⁷

To evaluate the relative importance of each infrastructure dimension, we estimate a “horse race” regression in which the grid-level rainfall shock coefficient $\hat{\beta}_g$ is regressed on tercile indicators for each infrastructure variable (Table 3.7). Column (2) of Table 3.7 additionally includes topographical controls - elevation and proximity to water bodies - that influence both infrastructure placement and flood exposure (Government of Delhi, 2018).

⁶This is equivalent to about 0.2 standard deviations of the distribution of $\hat{\beta}_g$.

⁷Ideally, we would identify the causal effects of infrastructure improvements using natural experiments from infrastructure construction, but we were unable to identify a suitable event or setting for this purpose.

$$\begin{aligned}\hat{\beta}_g = & \alpha + \sum_{j=2}^3 \gamma_j^{\text{Drain}} \cdot \mathbf{1}\{\text{Drain Tercile} = j\} + \sum_{j=2}^3 \gamma_j^{\text{Road}} \cdot \mathbf{1}\{\text{Road Tercile} = j\} \\ & + \sum_{j=2}^3 \gamma_j^{\text{Metro}} \cdot \mathbf{1}\{\text{Metro Tercile} = j\} + \varepsilon_g\end{aligned}\tag{3.5}$$

We find suggestive evidence that higher infrastructure quality is associated with smaller rainfall-induced earnings losses. In particular, workers living in grids with more extensive drainage networks and better metro accessibility experience lower income losses during rainfall shocks. These findings support the view that localized infrastructure, especially drainage and metro transit in this context, can play a critical role in reducing economic losses due to rainfall shocks (Gandhi et al., 2022).

3.7.5 Policy Implications

Distance from the center is a useful way to characterize a neighborhood’s location within the city, but it is not itself a policy lever. Policymakers cannot change how far a neighborhood is from the urban core. What matters for policy is addressing the structural disadvantages that often correlate with peripheral location. Combined with the description of Delhi’s peripheral areas in Section 3.4.3, the evidence presented in this section shows that workers in these areas face consistently greater exposure to rainfall shocks, despite the spatial optimization and flexibility of the platform.

These vulnerabilities arise from multiple overlapping constraints. First, peripheral areas tend to have lower wealth (Table 3.2) making adaptation to shocks harder (Bellon and Massetti, 2022). Second, these areas tend to have lower infrastructure quality, including weaker drainage and limited metro connectivity, which is associated with larger economic losses during rainfall events (Table 3.7). Third, the potential for location-based gig platforms to deliver hyperlocal work is limited in practice by existing residential and demand patterns (Figure 3.3), and workers in peripheral areas have to travel longer distances per job (Table 3.2), possibly increasing their exposure to disruption.

These findings suggest that, given existing spatial inequalities and the increasing frequency of weather shocks, targeted place-based investments (Balboni et al., 2020) in peripheral neighborhoods may play an important role in supporting more equitable urban resilience.

3.8 Conclusion

This paper studies how urban workers adapt to climate shocks in a setting that offers unusually high flexibility: location-based platform work in Delhi. Despite the temporal and spatial margins of adjustment available to workers, extreme rainfall leads to meaningful and unequal income losses. Workers partially adapt through intertemporal substitution, reallocating labor across days, but we find no evidence of spatial adaptation. The persistence

of large losses, concentrated among poorer and more peripheral neighborhoods, highlights the limits of individual flexibility when underlying structural constraints remain binding.

These findings underscore a broader lesson: while digital platforms can create new margins for adaptation – particularly by facilitating intertemporal flexibility—they cannot overcome deep-rooted spatial inequalities within cities. Technology alone cannot substitute for physical infrastructure, accessible transit, or the spatial distribution of demand. In developing cities where growth has outpaced public investment, climate shocks are likely to reinforce existing disparities unless adaptation strategies explicitly target the geographic patterns of vulnerability.

We find suggestive evidence that targeted infrastructure investments could improve resilience to climate risks. Improving drainage, expanding transit connectivity, and enhancing access to economic hubs could not only reduce the direct costs of weather shocks but also help mitigate the broader amplification of inequality across urban space. While our setting focuses on a single developing city and a specific type of service-sector platform work, the mechanisms we document, such as mobility constraints, infrastructural bottlenecks, and spatially uneven exposure are likely to be salient across many rapidly growing cities facing the challenges of climate change. Future research should focus on evaluating the impacts of specific infrastructure improvements on climate resilience.

More broadly, understanding how extreme weather interacts with urban spatial structure will be critical for preserving the economic promise of cities. Building cities that are both productive and resilient will require addressing not only the quantity of infrastructure but also the geography of opportunity. As climate risks intensify, ensuring that adaptation efforts are spatially inclusive will be essential for sustaining the role of cities as engines of growth and upward mobility in the developing world.






Figure 3.1: Distribution of Providers Across Occupation

Notes: This figure shows the distribution of service providers across occupational categories on the platform. Certain categories are grouped based on worker-overlap.




(a) Distribution by Date




(b) Histogram

Figure 3.2: Proportion of Grids Shocks


Notes: This figure summarizes the cross-sectional distribution of rainfall shocks across grid cells in Delhi. Panel (a) plots the daily share of grids experiencing a rainfall shock – defined as precipitation exceeding 10mm – from January 2023 to early 2025. Panel (b) shows a histogram of this daily proportion. The figure highlights that shocks tend to affect large parts of the city simultaneously, with most shock days involving either very few or nearly all grid cells. This limited cross-sectional variation motivates our empirical design.




(a) **Tercile 1: Workers**




(b) Customers of **Tercile 1** Workers




(c) **Tercile 2: Workers**



(d) Customers of **Tercile 2** Workers




(e) **Tercile 3: Workers**




(f) Customers of **Tercile 3** Workers

Figure 3.3: Spatial distribution of workers and customers by tercile

Notes: This figure displays the spatial distribution of workers and customer job locations by distance terciles. The left column shows worker residences, while the right column plots heatmaps of customer job locations corresponding to those workers. Terciles are constructed based on the radial distance of worker residences from the city center. Warmer colors (red) / darker areas indicate higher density.




(a) Earnings




(b) Services Delivered

Figure 3.4: Intertemporal Substitution of Earnings and Labor Supply

Notes: This figure shows event study estimates of daily effects around a rainfall shock, with day 0 representing the day of shock. Panel (a) plots average effects on daily earnings; Panel (b) shows effects on number of jobs completed. Each point reflects the estimated coefficient for the corresponding day relative to the shock, with 95% confidence intervals. Estimates come from Equation 3.3. The total effect over the full window is reported above each panel. Standard errors are clustered at the grid level, the unit at which rainfall is measured.



(a) Warnings before Rain Shocks



(b) Forecast Hit Rate

Figure 3.5: Warnings or Forecasts Before Heavy Rain

Notes: Panel (a) plots the estimated probability of receiving a rain warning in the days before a rainfall shock, where the dependent variable is an indicator for any rain and the independent variables are leads and lags of the rain shock indicator ($\text{rain} > 10\text{mm}$). Panel (b) shows forecast hit rates by day and threshold using publicly available data from India's meteorological service. Lines correspond to different rain thresholds: any rain, 5mm, and 10mm.




Figure 3.6: Forecasted vs Unforecasted shocks

Note: This analysis is restricted to Noida, the one region for which we were able to accurately link forecast data with shocks data used for our analysis. Estimates are from estimating equation 3.3 in two separate regressions: one in which only shock days classified as forecasted and non-shock days are included, one in which only shock days classified as unforecasted and non-shock days are included.





Figure 3.7: Spatial heterogeneity in shock effects across grids

Notes: Figure shows grid-level heterogeneity in estimated effects of rainfall shocks on daily worker earnings. We estimate a separate shock coefficient for each grid using a saturated regression specification. Grids with negative point estimates are divided into terciles and shaded red (most negative), yellow (middle), and green (least negative). Grids with positive point estimates are shown in blue. Estimates in blue are not statistically significant at the 10 percent level.



(a) Loss in Rupees

(b) Loss as a % of Average Income in the Tercile

Figure 3.8: Daily Income Loss Due to Rainfall Shocks, by Distance from Center

Notes: Figure plots the estimated daily income loss due to rainfall shocks by worker residence tertile. Panel (a) shows losses in rupees; Panel (b) scales these losses by average daily income for each tertile. Estimates are from separate regressions for each tertile that control for worker, day-of-week, and week-by-occupation fixed effects. Tertiles are defined by radial distance from the city center based on worker residence location.




Figure 3.9: Dynamic Adaptation, by Distance from Center

Notes: This figure shows dynamic estimates of rainfall shock effects on services delivered, separately for workers in each tercile of distance from the city center. Estimates are from separate regressions by tercile based on Equation 3.3, with controls for worker, day-of-week, and week-by-occupation fixed effects. Shaded bands indicate 95% confidence intervals.




Figure 3.10: Correlation between Relative Wealth and Income Loss in the Grid

Notes: This figure plots the estimated grid-level earnings impact of rainfall shocks against the standardized Relative Wealth Index (RWI), with a fitted linear regression line and 95% confidence interval. Each point represents a grid cell, and the slope γ is from a cross-sectional regression of grid-level impacts on standardized RWI. The positive coefficient indicates that workers in poorer neighborhoods experience larger income losses on shock days.

Table 3.1: Labor Market Summary Statistics

Outcome	Mean / SD
Daily Activity	
Jobs Done	1.19 (0.83)
Jobs Done (Cond. on Working)	1.97 (0.80)
Distance per Job (km)	7.14 (4.37)
Earnings per Job (Rs.)	749.17 (514.05)
Daily Income (Rs.)	770.58 (573.59)
Jobs Offered	2.29 (1.26)
Acceptance Rate	0.68 (0.18)
Same-Day Jobs Offered	1.18 (0.79)
Earnings Share from Same-Day Jobs	0.47 (0.19)
Weekly Aggregates	
Weekly Income (Rs.)	5142.04 (3382.88)
Weekly Jobs	7.83 (4.61)
Active Days per Week	3.73 (1.38)
Monthly Aggregates	
Monthly Income (Rs.)	21237.17 (14565.10)
Monthly Jobs	32.09 (19.63)
Active Days per Month	15.32 (6.28)
Number of Workers	9779

Notes: This table reports summary statistics for worker labor market outcomes. Each row shows the mean and standard deviation (in parentheses) of the outcome across all worker-day observations. ‘Jobs Done (Cond. on Working)’ is conditional on the worker accepting at least one job on that day. ‘Acceptance rate’ is the ratio of jobs accepted to total job offers. ‘Active days’ refer to the number of days a worker completes at least one job. ‘Same-Day Jobs Offered’ refers to offers for jobs requested and scheduled on the same calendar day. ‘Earnings Share from Same-Day Jobs’ denotes the proportion of daily income earned from such jobs. Weekly and monthly aggregates are computed by summing daily outcomes for each worker. The sample includes 9,779 unique workers observed between January 2023 and November 2024.

Table 3.2: Statistics for the Median Worker, by Distance from Center Terciles

	Tercile 1 (1)	Tercile 2 (2)	Tercile 3 (3)
Daily Average Jobs Done	1.20	1.30	1.29
Daily Average Earnings	735.82	831.26	826.70
Average Distance Per Job (km)	7.30	6.95	7.88
Average Earnings Per Job	727.42	740.92	720.54
Observations	5321	3981	1099

Notes: This table reports median labor market outcomes for workers, stratified by distance-from-center terciles based on worker residence. Tercile 1 includes workers living closest to the urban core; Tercile 3 includes workers in the most peripheral areas. Daily averages are calculated over active days (days when at least one job was completed). Average Distance Per Job refers to the total distance traveled by the worker on the day divided by the total number of jobs. Average Earnings Per Job is computed as daily earnings divided by number of jobs.

Table 3.3: Wealth and Infrastructure of Neighborhoods, by Distance from Center Terciles

	Tercile 1	Tercile 2	Tercile 3
	(1)	(2)	(3)
Meta Relative Wealth Index (Z-Score)	0.53	0.14	-0.65
Drain Density (Z-Score)	0.64	-0.25	-0.37
Road Density (Z-Score)	0.57	0.05	-0.61
Metro Stations (Z-Score)	0.89	-0.26	-0.60
Observations	32	33	33

Notes: This table reports average z-scores of neighborhood-level (grid-level) characteristics by distance-from-center terciles. Variables are standardized at the grid-cell level before aggregation. The Meta Relative Wealth Index is sourced from Meta ([Chi et al., 2022](#)). Infrastructure measures include drain density, road density, and number of metro stations within each grid, as defined in Section 3.3.4 from OpenStreetMap ([OpenStreetMap Contributors, 2023](#)).

Table 3.4: Rainfall Shock Impact

	Earnings (Rs.) (1)	Jobs Delivered (2)	Distance Per Job (km) (3)
Rain Shock	-30.09*** (3.523)	-0.05*** (0.005)	0.02* (0.010)
DV Control Mean	798.44	1.17	6.92
% Impact	-3.77	-4.47	0.26
Observations	4,207,073	4,207,073	2,358,406
Unique Workers	9,779	9,779	9,744
Unique Grids	157	157	157
Worker FE	Yes	Yes	Yes
Category x Week FE	Yes	Yes	Yes
Day-of-Week FE	Yes	Yes	Yes

Notes: This table reports estimates from regressions of daily worker outcomes on a rainfall shock indicator. The dependent variables are daily earnings (column 1), number of jobs delivered (column 2), and average distance per job in kilometers (column 3). A rainfall shock is defined as more than 10mm of precipitation in a worker's grid on a given day. All regressions include worker fixed effects, day-of-week fixed effects, and service category-by-week fixed effects. Standard errors are clustered at the grid level, the unit at which rainfall is measured. The percent impact is computed as the coefficient divided by the mean of the dependent variable on days without the shock. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.5: Labor Demand vs Supply

	Offers (1)	Acceptance Rate (2)
Rain Shock	0.011** (0.005)	-0.008*** (0.001)
DV Control Mean	2.37	0.68
% Impact	0.45	-1.24
Observations	4,207,073	2,925,025
Unique Workers	9,779	9,613
Unique Grids	157	157
Worker FE	Yes	Yes
Category x Week FE	Yes	Yes
Day-of-Week FE	Yes	Yes

Notes: This table estimates the impact of rainfall shocks on labor demand and labor supply measures. Column 1 reports the effect on the number of job offers received, while Column 2 shows the effect on the acceptance rate (defined as the share of offered jobs that were accepted). Each regression includes worker fixed effects, day-of-week fixed effects, and service category-by-week fixed effects. Standard errors are clustered at the grid level, the unit at which rainfall is measured. The percent impact is computed as the coefficient divided by the mean of the dependent variable on days without the shock. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.6: Heterogeneity: Urgent vs Non-Urgent

	Jobs Done (1)	Offers (2)	Acceptance Rate (3)
Rain Shock	-0.036*** (0.011)	0.055*** (0.017)	-0.005** (0.002)
Rain Shock * Non-Urgent	-0.018* (0.010)	-0.049*** (0.017)	-0.004* (0.002)
DV Control Mean	1.17	2.37	0.68
Observations	4207073	4207073	2925025
Unique Workers	9,779	9,779	9,613
Unique Grids	157	157	157
Worker FE	Yes	Yes	Yes
Category x Week FE	Yes	Yes	Yes
Day-of-Week FE	Yes	Yes	Yes

Notes: This table reports the heterogeneous effects of rainfall shocks by job urgency. “Urgent” services include occupations such as party make-up and plumbing, while “non-urgent” services include home cleaning and beauty treatments (details in Section 3.5). The classification is based on the expected immediacy of service needs. Columns 1–3 show effects on jobs completed, job offers received, and acceptance rate, respectively. All regressions include worker fixed effects, day-of-week fixed effects, and service category-by-week fixed effects, interacted with the dummy for ‘non-urgent’. Standard errors are clustered at the grid level, the unit at which rainfall is measured. The percent impact is computed as the coefficient divided by the mean of the dependent variable on days without the shock. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.7: Infrastructure and Rain Impacts

		$\beta_g = \text{Earnings Loss}$	
		(1)	(2)
Drain Density	Tercile 2	6.98 (13.326)	2.20 (12.802)
Drain Density	Tercile 3	14.61* (7.530)	12.22 (9.938)
Road Density	Tercile 2	3.37 (13.569)	-3.21 (17.783)
Road Density	Tercile 3	-2.62 (10.678)	-6.82 (13.603)
Metro Stations	Tercile 2	22.12* (12.829)	21.82 (13.152)
Metro Stations	Tercile 3	21.81*** (8.163)	22.12** (9.148)
N		89	89
Topographical Controls			X

Notes: This table reports coefficients from cross-sectional regressions of estimated grid-level rainfall impacts ($\hat{\beta}_g$) on tercile indicators for infrastructure quality. The dependent variable is the grid-level effect of rainfall shocks on earnings, estimated from the main specification. Drain, road, and metro variables reflect standardized measures of infrastructure density constructed from OpenStreetMap data. Column (2) adds topographic controls including elevation and proximity to water bodies. All coefficients are relative to the omitted first tercile (lowest infrastructure availability). Standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix 3.A Appendix




Figure A1: Grid Level Estimates with P-Value less than 0.1

General Conclusion

This dissertation contributes to the understanding of how frictions shape welfare outcomes for workers and households in developing economies. The three essays offer a coherent argument: frictions that appear modest in absolute terms, such as a higher upfront price or a day of heavy rain, can have large and persistent welfare consequences when they interact with binding liquidity constraints or infrastructure deficits that are pervasive among low-income households, and their costs fall disproportionately on those least equipped to absorb them.

Chapters 1 and 2 identify the same mechanism, namely purchase-time liquidity constraints, operating in two institutionally distinct settings: employer-facilitated capital investment and government energy subsidies. The consistency of the finding across such different contexts strengthens confidence in the mechanism and has clear policy implications. Program and employer designs that front-load costs will systematically under-reach the poor, while interventions that ease the upfront burden, whether through employer subsidies, purchase-point credit, or electronic vouchers, can substantially raise adoption and welfare at potentially low fiscal cost.

Chapter 3 complements the financial-frictions perspective with an analysis of environmental frictions. Even in a labor market explicitly designed for flexibility, infrastructure gaps impose real and regressive income costs on workers in the most vulnerable neighborhoods. The results suggest that adaptation to climate shocks in rapidly urbanizing developing cities cannot rely on market flexibility alone, and that spatially targeted investments in drainage, transit, and connectivity are a necessary complement.

Several directions for future research follow naturally from this work. The worker-owned capital framework developed in Chapter 1 raises questions about the generality of the mechanism across other occupations and asset types, and about the sources of the subsidy gap itself: understanding whether employer reluctance to subsidize worker capital differs between new and experienced workers would shed light on whether the friction stems primarily from incomplete appropriability, from asymmetric information about worker quality, or from other features of the employment relationship.

The findings of Chapter 2 suggest that reducing the upfront cost of clean fuel at the point of purchase, rather than reimbursing it later, could be substantially more effective at shifting households away from polluting energy sources. Randomizing the timing of subsidy delivery would allow researchers to identify the causal consequences of upfront cost reduction more cleanly and to quantify how much of the adoption gap is attributable to payment architecture alone.

Chapter 3 opens questions about how to actively mitigate the consequences of

environmental shocks for low-income workers, including through weather-indexed insurance products that smooth income in the wake of extreme events, and through better information provision that helps workers anticipate shocks and adjust effort in advance. Understanding how insurance and information interact with physical infrastructure in shaping workers' ability to absorb climate shocks is an important direction for future work on urban poverty and resilience in developing cities.

Bibliography

- Abubakari, Sulemana, KP Asante, Misbath Daouda, B Kelsey Jack, Darby Jack, Flavio Malagutti, and Paulina Oliva**, “Targeting subsidies through price menus: Menu design and evidence from clean fuels,” Technical Report, Working paper 2024.
- Acemoglu, Daron and Jörn-Steffen Pischke**, “The structure of wages and investment in general training,” *Journal of political economy*, 1999, *107* (3), 539–572.
- Adams, William, Liran Einav, and Jonathan Levin**, “Liquidity constraints and imperfect information in subprime lending,” *American Economic Review*, 2009, *99* (1), 49–84.
- Adhvaryu, Achyuta, Namrata Kala, and Anant Nyshadham**, “Returns to on-the-job soft skills training,” *Journal of Political Economy*, 2023, *131* (8), 2165–2208.
- Afridi, Farzana, Sisir Debnath, and Eswaran Somanathan**, “A breath of fresh air: Raising awareness for clean fuel adoption,” *Journal of Development Economics*, 2021, *151*, 102674.
- , —, **Taryn Dinkelman, and Komal Sareen**, “Time for clean energy? Cleaner fuels and women’s time in home production,” *The World Bank Economic Review*, 2023, *37* (2), 283–304.
- Aguila, Emma, Arie Kapteyn, and Francisco Perez-Arce**, “Consumption smoothing and frequency of benefit payments of cash transfer programs,” *American Economic Review*, 2017, *107* (5), 430–435.
- Alhorr, Layane, Kartik Srivastava, and Alp Sungu**, “Between Trust and Trade: on Informal Credit Networks in India,” *Available at SSRN 5392415*, 2025.
- Aragón, Fernando M., Juan Jose Miranda, and Paulina Oliva**, “Particulate matter and labor supply: The role of caregiving and non-linearities,” *Journal of Environmental Economics and Management*, 2017, *86*, 295–309.
- Autor, David H**, “Why do temporary help firms provide free general skills training?,” *The Quarterly Journal of Economics*, 2001, *116* (4), 1409–1448.
- Balboni, Clare, Gharad Bryan, Simon Galle, and Gautam Rao**, “Transportation, Gentrification, and Urban Mobility: The Inequality Effects of Place-Based Policies,” 2020. Preliminary Draft 3.

- Bandiera, Oriana, Robin Burgess, Narayan Das, Selim Gulesci, Imran Rasul, and Munshi Sulaiman**, “Labor markets and poverty in village economies,” *The Quarterly Journal of Economics*, 2017, 132 (2), 811–870.
- Banerjee, Abhijit, Dean Karlan, and Jonathan Zinman**, “Six randomized evaluations of microcredit: Introduction and further steps,” *American Economic Journal: Applied Economics*, 2015, 7 (1), 1–21.
- , **Esther Duflo, Rachel Glennerster, and Cynthia Kinnan**, “The miracle of microfinance? Evidence from a randomized evaluation,” *American economic journal: Applied economics*, 2015, 7 (1), 22–53.
- , **Rema Hanna, Benjamin A Olken, Elan Satriawan, and Sudarno Sumarto**, “Electronic food vouchers: Evidence from an at-scale experiment in Indonesia,” *American Economic Review*, 2023, 113 (2), 514–547.
- Barbanchon, Thomas Le, Roland Rathelot, and Alexandra Roulet**, “Gender differences in job search: Trading off commute against wage,” *The Quarterly Journal of Economics*, 2021, 136 (1), 381–426.
- Barnwal, Prabhat**, “Curbing leakage in public programs: Evidence from India’s direct benefit transfer policy,” *American Economic Review*, 2024, 114 (12), 3812–3846.
- Barua, S. K. and Sobhesh Kumar Agarwalla**, “Lighting up lives through cooking gas and transforming society,” IIMA Working Papers WP 2018-12-05, Indian Institute of Management Ahmedabad, Research and Publication Department December 2018.
- Bazzi, Samuel, Sudarno Sumarto, and Asep Suryahadi**, “It’s all in the timing: Cash transfers and consumption smoothing in a developing country,” *Journal of Economic Behavior & Organization*, 2015, 119, 267–288.
- Beaman, Lori, Dean Karlan, Bram Thuysbaert, and Christopher Udry**, “Selection into credit markets: Evidence from agriculture in Mali,” *Econometrica*, 2023, 91 (5), 1595–1627.
- Bearpark, Tom, Ashwin Rode, and Archana Patankar**, “The Urban Mortality Consequences of Rainfall and Sea Level Rise,” 2025. Unpublished working paper.
- Becker, Gary S**, “Human Capital New York,” *Columbia University for the National Bureau of Economic Research*, 1964.
- Bellon, M. and E. Massetti**, “Economic Principles for Integrating Adaptation to Climate Change into Fiscal Policy,” Technical Report 2022/001, International Monetary Fund 2022. Retrieved April 23, 2025.
- Berkouwer, Susanna B and Joshua T Dean**, “Credit, attention, and externalities in the adoption of energy efficient technologies by low-income households,” *American Economic Review*, 2022, 112 (10), 3291–3330.

- **and** –, “Private actions in the presence of externalities: The health impacts of reducing air pollution peaks but not ambient exposure,” Technical Report, National Bureau of Economic Research 2023.
- Bhatt, Vasavi, S. Chandrasekhar, and Ajay Sharma**, “Regional Patterns and Determinants of Commuting Between Rural and Urban India,” *The Indian Journal of Labour Economics*, 2020, *63* (4), 1041–1063.
- Bibra, Ekta Meena, Elizabeth Connelly, Marine Gerner, Christopher Lowans, Leonardo Paoli, Jacopo Tattini, and Jacob Teter**, “Global EV Outlook 2021: Accelerating ambitions despite the pandemic,” 2021.
- Black, Richard, Sonya Bennett, Sophie Thomas et al.**, “Migration as Adaptation,” *Nature*, 2011, *478*, 447–449.
- Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts**, “Does management matter? Evidence from India,” *The Quarterly journal of economics*, 2013, *128* (1), 1–51.
- Bobba, Matteo, Lucie Gadenne, Mariaflavia Harari, and Gabrielle Ulyssea**, “Informality,” *VoxDevLit*, April 2025, *6* (2).
- Borker, Girija et al.**, “Safety first: Perceived risk of street harassment and educational choices of women,” 2021.
- Brueckner, Jan K., Jacques-François Thisse, and Yves Zenou**, “Why Is Central Paris Rich and Downtown Detroit Poor?: An Amenity-Based Theory,” *European Economic Review*, 1999, *43* (1), 91–107.
- Brune, Lasse, Eric Chyn, and Jason Kerwin**, “Pay me later: Savings constraints and the demand for deferred payments,” *American Economic Review*, 2021, *111* (7), 2179–2212.
- Bryan, Gharad, Edward Glaeser, and Nick Tsivanidis**, “Cities in the Developing World,” *Annual Review of Economics*, 2020, *12* (1), 273–297.
- , **Shyamal Chowdhury, and Ahmed Mushfiq Mobarak**, “Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh,” *Econometrica*, 2014, *82* (5), 1671–1748.
- Bryan, Gharad T, Kyra Frye, and Melanie Morten**, “Spatial Economics for Low- and Middle-Income Countries,” Working Paper 33453, National Bureau of Economic Research February 2025.
- Burlig, Fiona, Amir Jina, Erin M. Kelley, Gregory Lane, and Harshil Sahai**, “Long-range forecasts as climate adaptation: Experimental evidence from developing-country agriculture,” 2024. October.

- Business Standard**, “How Delhi Went from Extreme Heat to Heavy Floods; All You Need to Know,” <https://www.business-standard.com/india-news/how-delhi-went-from-extreme-heat-to-heavy-floods> 2025. Accessed April 23, 2025.
- Canada Revenue Agency**, “Medical Expenses 2024: Refundable Medical Expense Supplement,” 2025. Details the refundable tax credits available for medical expenses in Canada.
- Canagarajah, Sudharshan and S.V. Sethuraman**, “Social Protection and the Informal Sector in Developing Countries: Challenges and Opportunities,” Social Protection Discussion Paper 0130, World Bank 2001.
- Casaburi, Lorenzo and Jack Willis**, “Time versus state in insurance: Experimental evidence from contract farming in Kenya,” *American Economic Review*, 2018, 108 (12), 3778–3813.
- Cefala, Luisa, Pedro Naso, Michel Ndayikeza, and Nicholas Swanson**, “Under-training by employers in informal labor markets: Evidence from burundi,” Technical Report, Working Paper 2024.
- Cesur, Resul, Erdal Tekin, and Aydogan Ulker**, “Air pollution and infant mortality: evidence from the expansion of natural gas infrastructure,” *The economic journal*, 2017, 127 (600), 330–362.
- Chetty, Raj, Adam Looney, and Kory Kroft**, “Salience and taxation: Theory and evidence,” *American Economic Review*, 2009, 99 (4), 1145–77.
- Chi, G., H. Fang, S. Chatterjee, and J. E. Blumenstock**, “Microestimates of Wealth for All Low- and Middle-Income Countries,” *Proceedings of the National Academy of Sciences of the United States of America*, 2022, 119 (3), e2113658119.
- CNN**, “Delhi Suffers Extreme Weather Whiplash as Heat Waves Give Way to Floods,” <https://www.cnn.com/2024/07/01/climate/delhi-extreme-weather-floods> July 2024. Accessed April 23, 2025.
- Coffman, Lucas C, John J Conlon, Clayton R Featherstone, and Judd B Kessler**, “Liquidity affects job choice: Evidence from Teach for America,” *The Quarterly Journal of Economics*, 2019, 134 (4), 2203–2236.
- Connolly, Marie**, “Here Comes the Rain Again: Weather and the Intertemporal Substitution of Leisure,” *Journal of Labor Economics*, 2008, 26 (1), 73–100.
- Couture, Victor, Cecile Gaubert, Jessie Handbury, and Erik Hurst**, “Income Growth and the Distributional Effects of Urban Spatial Sorting,” *Review of Economic Studies*, 2024, 91 (2), 858–898.

- Crépon, Bruno, Mohamed El Komi, and Adam Osman**, “Is it who you are or what you get? Comparing the impacts of loans and grants for microenterprise development,” *American Economic Journal: Applied Economics*, 2024, 16 (1), 286–313.
- Cutter, Susan L., William Solecki, Nancy Bragado, JoAnn Carmin, Michail Fragkias, Matthias Ruth, and Thomas J. Wilbanks**, “Urban Systems, Infrastructure, and Vulnerability,” in Jerry M. Melillo, Terese (T.C.) Richmond, and Gary W. Yohe, eds., *Climate Change Impacts in the United States: The Third National Climate Assessment*, U.S. Global Change Research Program, 2014, pp. 282–296.
- Daher, Chaza Abou, Erica M Field, Kendal M Swanson, and Kate H Vyborny**, “Drivers of change: How intra-household preferences shape employment responses to gender reform,” Technical Report, National Bureau of Economic Research 2023.
- Datta, Ayona, Mark Graham et al.**, “Informal Labor, Platform Work, and Digital Intermediation in the Global South,” *Global Networks*, 2023, 23 (1), 123–145.
- Department of Agriculture, Republic of the Philippines**, “Memorandum Order No. 31, s.2020: Implementing Guidelines on the Provision of Fertilizer Assistance under the Rice Resiliency Project (RRP),” Office of the Secretary, Department of Agriculture April 2020. Fertilizer assistance through reimbursement scheme during COVID-19, allowing rice farmers to buy from accredited dealers and claim reimbursement with receipts.
- Desai, Sonalde, Reeve Vanneman, and National Council of Applied Economic Research**, “India Human Development Survey - II (IHDS-II) ICPSR36151-v2,” 2011.
- Desmet, Klaus and Esteban Rossi-Hansberg**, “Climate Change Economics over Time and Space,” *Annual Review of Economics*, 2024, 16 (1), 271–304.
- Doornik, Bernardus Van, Armando Gomes, David Schoenherr, and Janis Skrastins**, “Financial access and labor market outcomes: Evidence from credit lotteries,” *American Economic Review*, 2024, 114 (6), 1854–1881.
- Dubé, Jean-Pierre and Sanjog Misra**, “Personalized Pricing and Consumer Welfare,” *Journal of Political Economy*, 2023.
- Duflo, Esther, Michael Kremer, and Jonathan Robinson**, “How high are rates of return to fertilizer? Evidence from field experiments in Kenya,” *American economic review*, 2008, 98 (2), 482–488.
- , —, and —, “Nudging farmers to use fertilizer: Theory and experimental evidence from Kenya,” *American economic review*, 2011, 101 (6), 2350–2390.
- Duranton, Gilles and Diego Puga**, “Micro-foundations of Urban Agglomeration Economies,” in J. Vernon Henderson and Jacques-François Thisse, eds., *Handbook of Regional and Urban Economics*, Vol. 4, Elsevier, 2004, pp. 2063–2117.

- Fafchamps, Marcel, David McKenzie, Simon Quinn, and Christopher Woodruff**, “Microenterprise growth and the flypaper effect: Evidence from a randomized experiment in Ghana,” *Journal of development Economics*, 2014, 106, 211–226.
- Farber, Henry S.**, “Why you can’t find a taxi in the rain and other labor supply lessons from cab drivers,” *The Quarterly Journal of Economics*, 2015, 130 (4), 1975–2026.
- Floess, E., A. Grieshop, E. Puzzolo, D. Pope, N. Leach, C. J. Smith, A. Gill-Wiehl, K. Landesman, and R. Bailis**, “Scaling up gas and electric cooking in low- and middle-income countries: Climate threat or mitigation strategy with co-benefits?,” *Environmental Research Letters*, 2023, 18.
- Fogli, Alessandra and Veronica Guerrieri**, “The End of the American Dream? Inequality and Segregation in US Cities,” Working Paper 26143, National Bureau of Economic Research August 2019.
- Funk, Chris, Pete Peterson, Martin Landsfeld, Diego Pedreros, James Verdin, Shraddhanand Shukla, Gregory Husak, James Rowland, Laura Harrison, Andrew Hoell, and Joel Michaelsen**, “The Climate Hazards Infrared Precipitation with Stations—A New Environmental Record for Monitoring Extremes,” *Scientific Data*, 2015, 2, 150066.
- Gandhi, Sahil, Matthew E Kahn, Rajat Kochhar, Somik Lall, and Vaidehi Tandel**, “Adapting to Flood Risk: Evidence from a Panel of Global Cities,” Working Paper 30137, National Bureau of Economic Research June 2022.
- Garlick, Robert, Erica Field, and Kate Vyborny**, “Women’s Mobility and Labor Supply: Experimental Evidence from Pakistan,” 2025.
- Gelman, Michael, Shachar Kariv, Matthew D Shapiro, Dan Silverman, and Steven Tadelis**, “How individuals respond to a liquidity shock: Evidence from the 2013 government shutdown,” *Journal of Public Economics*, 2020, 189, 103917.
Gender Equality in Transport: Toolkit
- Gender Equality in Transport: Toolkit, Technical Report, Asian Development Bank 2021.*
- Ghili, Soheil, Vineet Kumar, and Fei Teng**, “Spatial Distribution of Access to Service: Theory and Evidence from Ridesharing,” *Technical Report, SSRN January 2025. SSRN Working Paper No. 4915262.*
- Glaeser, Edward**, *Triumph of the City: How Our Greatest Invention Makes Us Richer, Smarter, Greener, Healthier, and Happier*, New York: Penguin Books, 2012.
- Glaeser, Edward L., Matt Resseger, and Kristina Tobio**, “Inequality in Cities,” *Journal of Regional Science*, 2009, 49 (4), 617–646.

- Gould, Carlos F, Rob Bailis, Kalpana Balakrishnan, Marshall Burke, Sebastián Espinoza, Sumi Mehta, Samuel B Schlesinger, José R Suarez-Lopez, and Ajay Pillarisetti**, “In praise of cooking gas subsidies: transitional fuels to advance health and equity,” *Environmental Research Letters*, 2024, 19 (8), 081002.
- Government of Delhi**, “Drainage Master Plan for NCT of Delhi: Final Report,” *Technical Report, Department of Civil Engineering, Indian Institute of Technology Delhi July 2018. Prepared for the Department of Irrigation & Flood Control, Government of Delhi.*
- Goyal, Ananya, Radhika Pandey, and Renuka Sane**, “Consumption baskets of Indian households: Comparing estimates from the CPI, CES and CPHS,” *CES and CPHS*, 2021.
- Gross, Tal, Timothy J Layton, and Daniel Prinz**, “The liquidity sensitivity of healthcare consumption: Evidence from social security payments,” *American Economic Review: Insights*, 2022, 4 (2), 175–190.
- Gupta, Sanjay**, “Impact of Floods in Delhi,” *Technical Report, Climatrans, New Delhi, India 2017.*
- Hall, Jonathan V., John J. Horton, and Daniel T. Knoepfle**, “Ride-sharing Markets Re-equilibrate,” *NBER Working Paper w30883, National Bureau of Economic Research 2023.*
- Hanna, Rema and Paulina Oliva**, “The effect of pollution on labor supply: Evidence from a natural experiment in Mexico City,” *Journal of Public Economics*, 2015, 122, 68–79.
- , **Sendhil Mullainathan, and Joshua Schwartzstein**, “Learning through noticing: Theory and evidence from a field experiment,” *The Quarterly Journal of Economics*, 2014, 129 (3), 1311–1353.
- Heilmann, Kilian, Matthew E. Kahn, and Cheng Keat Tang**, “The urban crime and heat gradient in high and low poverty areas,” *Journal of Public Economics*, 2021, 197, 104408.
- Hoffmann, Bridget and Juan Pablo Rud**, “The Unequal Effects of Pollution on Labor Supply,” *Econometrica*, 2024, 92 (4), 1063–1096.
- Hsu, Angel, Glenn Sheriff, Tirthankar Chakraborty, and Diego Manya**, “Disproportionate exposure to urban heat island intensity across major US cities,” *Nature Communications*, 2021, 12, 2721.
- ILO**, “World Employment and Social Outlook: The Role of Digital Labour Platforms in Transforming the World of Work,” *Technical Report, International Labour Organisation, Geneva 2021.*
- Imelda**, “Cooking that kills: Cleaner energy access, indoor air pollution, and health,” *Journal of Development Economics*, 2020, 147.
- India—Making Public Transport and Urban Spaces Safer and Inclusive for Women: Policy Brief

India—Making Public Transport and Urban Spaces Safer and Inclusive for Women: Policy Brief, Technical Report, World Bank 2022.

IPCC, “Cities, Settlements and Key Infrastructure,” in “Climate Change 2022 – Impacts, Adaptation and Vulnerability: Working Group II Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change,” Cambridge University Press, 2023, pp. 907–1040.

Jack, B Kelsey, Seema Jayachandran, Namrata Kala, and Rohini Pande, “Money (not) to burn: payments for ecosystem services to reduce crop residue burning,” *American Economic Review: Insights*, 2025, 7 (1), 39–55.

Jacoby, Hanan G and Ghazala Mansuri, “Land tenancy and non-contractible investment in rural Pakistan,” *The Review of Economic Studies*, 2008, 75 (3), 763–788.

Jain, M., S. Siedentop, H. Taubenböck, and S. Namperumal, “From Suburbanization to Counterurbanization? Investigating Urban Dynamics in the National Capital Region Delhi, India,” *Environment and Urbanization ASIA*, 2013, 4 (2), 247–266.

Jayachandran, Seema, “Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries,” *Journal of Political Economy*, 2006, 114 (3), 538–575. Accessed April 23, 2025.

Jeuland, Marc, Manish A Desai, Elizabeth F Bair, Nafeesa Mohideen Abdul Cader, Durairaj Natesan, Wilson Jayakaran Isaac, Sankar Sambandam, Kalpana Balakrishnan, Gurusamy Thangavel, and Harsha Thirumurthy, “A randomized trial of price subsidies for liquefied petroleum cooking gas among low-income households in rural India,” *World Development Perspectives*, 2023, 30, 100490.

Kala, Namrata, Clare Balboni, and Shweta Bhogale, “Climate Adaptation,” *VoxDevLit*, June 2023, 7 (1).

Kansikas, Carolina, Anandi Mani, and Paul Niehaus, “Customized cash transfers: financial lives and cash-flow preferences in rural Kenya,” Working Paper 30930, National Bureau of Economic Research 2023.

Kar, Abhishek, Shonali Pachauri, Rob Bailis, and Hisham Zerriffi, “Using sales data to assess cooking gas adoption and the impact of India’s Ujjwala programme in rural Karnataka,” *Nature Energy*, 2019, 4 (9), 806–814.

Katz, Eliakim and Adrian Ziderman, “Investment in general training: The role of information and labour mobility,” *The economic journal*, 1990, 100 (403), 1147–1158.

Kilgarrieff, Paul, Thomas K. J. McDermott, Amaya Vega, Karyn Morrissey, and Cathal O’Donoghue, “The impact of flooding disruption on the spatial distribution of commuter’s income,” *Journal of Environmental Economics and Policy*, 2019, 8 (1), 48–64.

- Kleven, Henrik Jacobsen and Wojciech Kopczuk**, “Transfer program complexity and the take-up of social benefits,” *American Economic Journal: Economic Policy*, 2011, 3 (1), 54–90.
- Kocornik-Mina, Adriana, Thomas K. J. McDermott, Guy Michaels, and Ferdinand Rauch**, “Flooded Cities,” *American Economic Journal: Applied Economics*, 2020, 12 (2), 35–66.
- Kreindler, Gabriel**, “Peak-Hour Road Congestion Pricing: Experimental Evidence and Equilibrium Implications,” *Econometrica*, 2024, 92 (4), 1233–1268.
- Kreindler, Gabriel E. and Yuhei Miyauchi**, “Measuring Commuting and Economic Activity inside Cities with Cell Phone Records,” *The Review of Economics and Statistics*, 2023, 105 (4), 899–909.
- Krugman, Paul**, “Increasing Returns and Economic Geography,” *Journal of Political Economy*, 1991, 99 (3), 483–499.
- Laffont, Jean-Jacques and Mohamed Salah Matoussi**, “Moral hazard, financial constraints and sharecropping in El Oulja,” *The Review of Economic Studies*, 1995, 62 (3), 381–399.
- Levine, David I, Theresa Beltramo, Garrick Blalock, Carolyn Cotterman, and Andrew M Simons**, “What impedes efficient adoption of products? Evidence from randomized sales offers for fuel-efficient cookstoves in Uganda,” *Journal of the European Economic Association*, 2018, 16 (6), 1850–1880.
- Mas, Alexandre and Amanda Pallais**, “Valuing Alternative Work Arrangements,” *American Economic Review*, 2017, 107 (12), 3722–3759.
- Mathews, Rejeet, Amitabh Kundu, Paramjeet Chawla, Tintu Sebastian, Raj Bhagat Palanichamy, and Madhav Pai**, “Morphology of Delhi National Capital Region’s Economic Geography and Its Implications for Planning,” Technical Report, World Resources Institute India 2023.
- Meager, Rachael**, “Understanding the average impact of microcredit expansions: A Bayesian hierarchical analysis of seven randomized experiments,” *American Economic Journal: Applied Economics*, 2019, 11 (1), 57–91.
- Mehtab, Tauheed and Mohammad Arif Kamal**, “Investigating the Impact of Flood on Low Lying Settlements of Delhi, India: A Planning Perspective,” *Architecture Engineering and Science*, 2023, 4 (3), 136.
- Mel, Suresh De, David McKenzie, and Christopher Woodruff**, “Returns to capital in microenterprises: evidence from a field experiment,” *The quarterly journal of Economics*, 2008, 123 (4), 1329–1372.
- Mishra, Twesh**, “Ujjwala connections get three refills annually on an average,” The Hindu BusinessLine 2019.

- Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar**, “Identity verification standards in welfare programs: Experimental evidence from India,” *The Review of Economics and Statistics*, 2020, pp. 1–46.
- NASA Shuttle Radar Topography Mission**, “Shuttle Radar Topography Mission (SRTM) Global,” Distributed by OpenTopography 2013. Accessed: April 23, 2025.
- NITI-Aayog**, “India’s Booming Gig and Platform Economy: Perspectives and Recommendations on the Future of Work,” Technical Report June 2022.
- of Chicago Innovation Commission, University**, “Weather and Seasonal Forecasting,” Research Brief, Development Innovation Lab, University of Chicago 2024. Accessed April 22, 2025.
- OpenStreetMap Contributors**, “Planet Dump,” <https://planet.osm.org> 2023. Available at <https://www.openstreetmap.org>, Last accessed: February, 2025.
- Pandey, Kundan**, “Govt admits refilling LPG cylinders under Ujjwala a challenge, plans a new scheme,” Down To Earth 2019.
- Papp, Anna**, “Who Bears Climate Change Damages? Evidence from the Gig Economy,” November 2024. Job Market Paper, Columbia University.
- Patel, Dev**, “Floods,” November 2023. Available at SSRN: <https://ssrn.com/abstract=4636828>.
- Phuangsuwan, Penpim, Supaprawat Siripipatthanakul, Pongsakorn Limna, and Nuttharin Pariwongkhuntorn**, “The Impact of Google Maps Application on the Digital Economy,” *Corporate & Business Strategy Review*, 2024, 5 (1), 192–203.
- PLFS**, “Periodic Labour Force Survey (PLFS) Annual Report: July 2021 – June 2022,” Technical Report, National Statistical Office, Ministry of Statistics and Programme Implementation, Government of India February 2021.
- PPAC Report**, “LPG Profile : Data on LPG marketing as on 01.01.2018, Petroleum Planning and Analysis Cell (PPAC),” 2018.
- Reddy, Priya Adhishesha, Nathan Congdon, Graeme MacKenzie, Parikshit Gogate, Qing Wen, Catherine Jan, Mike Clarke, Jordan Kassalow, Ella Gudwin, Ciaran O’Neill et al.**, “Effect of providing near glasses on productivity among rural Indian tea workers with presbyopia (PROSPER): a randomised trial,” *The Lancet Global Health*, 2018, 6 (9), e1019–e1027.
- Roy, Devesh, Sushil Saroj, and Manoj Pradhan**, “Nature of employment and outcomes for urban labor: evidence from the latest labor force surveys in India,” *Indian Economic Review*, 2022, 57, 165–221.
- Shaban, Radwan Ali**, “Testing between competing models of sharecropping,” *Journal of Political Economy*, 1987, 95 (5), 893–920.

- Shah, Sonal et al.**, “Women and Transport in Indian Cities,” Technical Report, ITDP India 2017.
- Tarozzi, Alessandro, Aprajit Mahajan, Brian Blackburn, Dan Kopf, Lakshmi Krishnan, and Joanne Yoong**, “Micro-loans, insecticide-treated bednets, and malaria: evidence from a randomized controlled trial in Orissa, India,” *American Economic Review*, 2014, 104 (7), 1909–1941.
- The Economic Times**, “Delhi Rains: Latest Update, Delhi Rains Forecast,” <https://economictimes.indiatimes.com/news/delhi-rains-latest-update> 2025. Accessed April 23, 2025.
- The Sunday Guardian Live**, “Outdated Drainage System Worsens Delhi’s Waterlogging Crisis,” <https://sundayguardianlive.com/news/outdated-drainage-system-worsens-delhis-waterlogging-crisis> 2025. Accessed April 23, 2025.
- Thivillon, Thomas**, “Saving lives with cooking gas? Unintended effects of targeted LPG subsidies in Peru,” 2023.
- Tsivanidis, Nick**, “Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogota’s TransMilenio,” 2022. Unpublished manuscript.
- Tubaro, Paola, Antonio A. Casilli, Maxime Cornet, Clément Le Ludec, and Juana Torres Cierpe**, “Where Does AI Come From? A Global Case Study Across Europe, Africa, and Latin America,” Post-Print hal-04933816, HAL 2025.
Urban Transport: Policy Note Series
- Urban Transport: Policy Note Series, Technical Report, WRI India 2021.*
- Waghmare, Abhishek**, “Measuring employment and unemployment,” March 2024. *Data for India*.
- Weill, Joakim A.**, “Flood Risk Mapping and the Distributional Impacts of Climate Information,” Technical Report 2023-066, Board of Governors of the Federal Reserve System 2023.
- World Bank**, *World Development Report 2009: Reshaping Economic Geography*, World Bank, 2009. License: CC BY 3.0 IGO.
- , “How to Help Ensure Safe and Inclusive Public Spaces and Public Transport for Women in India,” 2023.
- Zeldes, Stephen P.**, “Consumption and liquidity constraints: an empirical investigation,” *Journal of political economy*, 1989, 97 (2), 305–346.
- Zippenfenig, Patrick**, “Open-Meteo.com Weather API,” 2023. Computer software.
- Zivin, Joshua Graff and Matthew Neidell**, “Temperature and the Allocation of Time: Implications for Climate Change,” *Journal of Labor Economics*, 2014, 32 (1), 1–26.