Do Microenterprises Maximize Profits?

A Nonrandomized Vegetable-Market Experiment in India

Abhijit Banerjee Greg Fischer Dean Karlan

Matt Lowe Benjamin N. Roth*

July 14, 2026

Abstract

We ran a nonrandomized, market-level experiment in Kolkata vegetable markets in which we subsidized vendors in some markets to sell additional produce. The vendors earned over 60% higher profits, excluding the value of the subsidy. Nevertheless, after the subsidy ended many vendors stopped selling the additional produce. Vendors knew about the profitable opportunity and demonstrated that they were capable of exploiting it without assistance. We conclude that their behavior meaningfully diverges from profit maximization when considering take-home pay, likely due to a combination of high marginal costs of effort and feared sanctions from breaking anti-competitive norms.

*Banerjee: M.I.T.; Fischer: London School of Economics and Boston Consulting Group; Karlan: Northwestern University and Innovations for Poverty Action; Lowe: University of British Columbia; Roth: Harvard Business School. IRB from IFMR #00007107. We thank Carolina Allende, Jacob Appel, Alex Nisichenko, Diego Santa Maria, and Rafael Tiara for project management and data analysis.

1 Introduction

Microenterprises often fail to grow (Hsieh and Olken, 2014). Economists have predominantly focused on external constraints as the key barriers to growth, such as a lack of capital (De Mel et al., 2008), labor (De Mel et al., 2019; Hardy and McCasland, 2023), managerial skill (Drexler et al., 2014), and information (Hanna et al., 2014). We identify a setting in which none of these constraints binds. Microentrepreneurs have an opportunity to materially increase their profits – by over 60% – for which they have access to the necessary capital, labor, skill, and information, and the opportunity poses little additional risk, yet it remains unexploited. We rule out typically hypothesized external constraints and are left with a stark result: these enterprises are simply not profit-maximizing with respect to take-home pay, either at the individual or group level. We emphasize that this is consistent with rational optimization once the physical toll of expansion and the risk of sanctions are counted, and is not evidence of irrationality or a taste for idleness.

We work with vegetable vendors in India. These vendors operate in densely populated markets and their capacity is often underutilized – common features of informal markets in developing countries (Lewis, 1954; Walker et al., 2024). We conducted a nonrandomized experiment in 20 Kolkata vegetable markets in which we subsidized vendors in three markets to expand their product offerings and utilize some of their spare capacity. We recorded prices and quantities for all vendors in all markets for three weeks. Then, in the three treatment markets we offered vendors three-week subsidies to procure and stock carrots and peas. We offered the carrot subsidy to all vendors and the pea subsidy only to those vendors who were not previously frequent pea sellers. The status quo prevailed in the remaining 17 markets. After the subsidy was removed, we recorded prices and quantities in all markets for a final two weeks. We use a difference-in-differences approach to estimate the contemporaneous and persistent effects of the subsidies, and, for small-N inference, we use the wild bootstrap and permutation tests.

Our first finding is that vendors in treatment markets stocked more peas and carrots during the subsidy period. Vendors tended to sell their expanded stock without cutting prices. As a result, the profits of vendors in treatment markets rose by more than 60%, not including the value of the subsidy; our experiment reveals the existence of a market-level profitable deviation. Importantly, vendors who received the subsidy had to provide the additional capital upfront and had to procure the additional produce on their own; they were only reimbursed for their purchases later in the day. Hence, by design, all vendors who exploited the opportunity to sell additional peas and carrots must have had access to the capital and knowledge necessary to do so. That vendors can, in partial equilibrium, significantly increase their profits by increasing their inventory is an important finding in its own right, given that there are over five million street vendors in India alone. Given its importance, we consider issues that might cause us to overestimate the effects on profits, including spillover effects, measurement error in prices, omitted costs, and misreporting by participants. We present a range of evidence against such issues, concluding that the profit effect is unlikely to be meaningfully mismeasured.

Strikingly, after the subsidy period concluded, many treated vendors reduced their procurement of peas and carrots to pre-intervention levels. That is, despite having experienced higher profits when they stocked the new products, and despite having the knowledge, capital, skill, and labor required to do so, many vendors reverted to their prior scale of operation and refrained from exploiting the profitable opportunity that they had just experienced.

Could vendors have failed to realize that profits increased despite having experienced higher profits from expanding their offerings? We view this as unlikely, because, in this setting it is straightforward to verify that profits have increased. As long as revenues from sales of a product exceed the cost of procurement, and as long as vendors have spare capacity to stock the additional products – two features satisfied by our environment – then stocking the additional products increases total profits. For a vendor to verify that they sold peas or carrots at a profit does not require complex counterfactual reasoning. We also rule out risk aversion and loss aversion as likely explanations. At the vendor-by-week level, stocking the additional peas and carrots results in lower profits less than 1% of the time, greatly reducing the specters of risk and loss.

Our experiment demonstrates the existence of a very profitable, low-risk deviation from current business practices that is not inhibited by external constraints, and that is not exploited even after vendors personally experienced it. We provide a simple theoretical framework to show that these results imply that vendors do not seek to maximize their profits, either at the individual- or group-level.

Two stories might explain the failure to maximize profits. First, the failure to maximize profits may reflect a form of separation failure (Benjamin 1992; LaFave and Thomas 2016). The effort required to procure and sell additional produce, and the stress of running a larger, more active business, may exceed the additional profits, while labor-market frictions may prevent vendors from hiring workers. Alternatively, departures from profit maximization may operate at the market level if norms or explicit agreements preclude people from expanding their inventories and selling too many products in direct competition with their nearby neighbors. Such collusive practices differ meaningfully from "classical" collusion, because they do not maximize joint profits at the market-level.

We report descriptive evidence for the two explanations from a survey of 391 vendors from the experimental markets, fielded about four years after the experiment. Vendors work long hours: 50 hours per week on average. To expand their operations, vendors most commonly say that they would need to work longer hours and work harder per hour – very few vendors would hire labor. Vendors perceive the additional effort as costly: 69% report that working one additional hour per day would be somewhat or very difficult, with the difficulty often due to the vendor's old age (they are 51 on average), poor health, or lack of energy. And while our profit effects remain meaningful when allowing for the treatment to have increased daily labor by a couple of hours, and valuing that labor at 60% of the market wage (following Agness et al. 2025), we note that the value of time at the margin may be much higher. After accounting for high marginal costs of effort, it may no longer be profitable for a vendor to expand in the absence of the subsidy. Perhaps reflecting this, when we asked vendors an open-ended question about their main hopes for their business, 85% said that they hoped to preserve the status quo, rather than expand their business.

The separation failure we describe – with the vendor's disutility of labor determining the scale of the vendor's business – is distinct from classic findings of rural households using inefficiently high labor on their farms (LaFave and Thomas 2016). In contrast to family farms, the markets we study are densely populated with many vendors operating side by side, procuring and selling similar produce. Exploiting the profitable opportunity we identify does not require access to a large and illiquid asset such as land, and thus the indivisibilities that characterize family farms do not appear to be operative in our setting. We conclude that, even if each vendor is privately optimizing, there must be an additional market failure that deters the exploitation of this profitable opportunity by a worker with low effort costs.

In addition to the separation failure (or cost of effort) channel, we find some descriptive support for collusive norms hindering expansion. Forty-five percent say that negative consequences would be likely if a vendor doubled their business over a few months, and 27% of vendors say that negative consequences would likely follow if a vendor who had never sold carrots or peas began to sell these products. The most common negative consequences reported are that other vendors would be angry, that they would spread information about the behavior to other vendors and markets, and that they would prevent the offending vendor from working at the market. Such perceived sanctions may then hinder expansion. While such collusive norms are well-documented in informal and agricultural markets (e.g. Breza et al., 2025; Bergquist and Dinerstein, 2020; Casaburi and Reed, 2022), our contribution is to highlight that in this setting, these norms appear to be in opposition to profit maximization, even at the market level. The norm explanation, however, has limits, as we fully discuss below. For one thing, it is not clear exactly how product-variety norms are enforced in a context in which a substantial number of vendors move in and out of selling peas and carrots.

The findings of our experiment also pose additional puzzles that we do not fully resolve. First, the subsidy intervention caused a large positive shock to a market's supply of carrots and peas, and yet we do not detect any reduction in retail prices. We find some support for a range of explanations, including rationing, highly elastic demand, and vendors inducing additional demand by, for example, staying longer at the market. Second, we find suggestive evidence that the subsidy intervention also increased the supply of non-subsidized vegetables. We find some evidence, albeit mixed, that this crowd-in effect is caused by the sale of vegetables complementary to carrots and peas.

Despite these unresolved puzzles, our core result is striking: given the magnitude of profits being left on the table – more than 60% of vendor profits, on average – these microentrepreneurs are not maximizing their take-home pay, even though declining to expand may be the right choice once the physical toll and the risk of sanctions are counted. This observation speaks to a longstanding puzzle within development economics. As early as 1954, Arthur Lewis noted the ubiquity of small firms operating side-by-side in densely packed urban markets, often seeming to operate below their capacity (Lewis, 1954), as is the case in the vegetable markets we study. Lewis conjectured that consumers would be no worse off if many traders left the market, leaving others to expand. Why this does not occur is a puzzle insofar as firms' natural desire is to grow their businesses and take market share, thereby driving some out of the market. However, our results may partially resolve this puzzle by casting doubt on the presumption that vendors uniformly seize opportunities to increase their scale and profitability.

Beyond the papers cited above, our paper is related to the literature that documents some firms' failure to maximize profits. Cho and Rust (2010), Atkin et al. (2017), and DellaVigna and Gentzkow (2019) document various failures of profit maximization due to the organizational complexity of large firms. In contrast, the microenterprises we study are overwhelmingly sole proprietorships. Beaman et al. (2014) and Gertler et al. (2025) document failures by small firms to adopt business practices that increase profits by 3 to 8%. These authors attribute the phenomenon to limited attention, memory, and trust. In contrast, we identify a failure to adopt business practices that are far more profitable as a percentage of baseline earnings, and failures of attention, memory, and trust are not plausible explanations.[1]

[1] Our paper also relates to the literature examining the extent to which micro- and small-business growth comes

2 An Experiment to Induce Vegetable Vendors to Increase Their Scale

2.1 Theoretical Framework

Before describing the experiment, we outline a stylized theoretical framework that we will use to interpret our experimental results. Suppose a vendor operates a business by choosing a scale k and producing profits $\pi(k)$, where $\pi(\cdot)$ may be a random variable. The choice variable k can be interpreted as capital, but, more generally, could represent any vector of inputs including labor, human capital, and so forth. The vendor maximizes

$$\max_{k \in \mathscr{C}} U\left(\hat{\pi}(k)\right) + \phi(k)$$

where $\mathscr{C}$ is an arbitrary set of constraints, $U(\cdot)$ is the vendor's expected utility from profits, assumed to be increasing, $\hat{\pi}(\cdot)$ is their perceived profit function, potentially differing from their true profit function $\pi(\cdot)$, and $\phi(\cdot)$ is a non-pecuniary cost of scale, representing the costs and benefits of scale not encompassed by profits. On the positive side, these factors could include the prestige of running a larger business, while on the negative side, they could reflect stress or effort costs, or social sanctions arising from violating norms that dictate the type and amount of inventory a vendor stocks.

Let k^* be the solution to the vendor's maximization problem. Our experiment seeks to address why vendors do not increase their scale to make more money and our framework yields a set of exhaustive explanations.

at the expense of competing businesses (De Mel et al., 2008; McKenzie and Woodruff, 2008; Drexler et al., 2014; Cai and Szeidl, 2024). Most closely related is McKenzie and Puerto (2021), which examines the impact of business training on female vendors in rural markets. In that setting, providing training to some entrepreneurs did not negatively impact competitors; rather, profits increased at the market level. Similarly, we find that our intervention causes profits to increase at the market level, and we find no evidence of negative spillovers on vendors who did not receive the pea subsidy.

1. $k^* = \underset{k \in \mathscr{C}}{\operatorname{argmax}}\, U(\pi(k))$. That is, the vendor is already choosing the scale that maximizes their expected utility from profits, given their constraints. In this case, increasing the vendor's scale would either not be feasible without outside intervention, or would reduce their profits.

2. $\hat{\pi}(\cdot) \neq \pi(\cdot)$. That is, the vendor may not be fully informed about their profit function, and may therefore not realize that increasing their scale would increase their profit.

3. $\phi(\cdot) \neq 0$. That is, the vendor chooses their business scale to maximize some combination of profits and other factors. In this case, we say that vendors do not maximize (only) profits.

Our experiment will rule out the first two explanations for why vendors do not increase their scale. We demonstrate that increasing scale is profitable and nearly risk-free. We further demonstrate that external constraints do not bind, in that vendors are capable of increasing their scale without outside assistance. Together these results rule out the first explanation.

We further demonstrate that despite knowledge of this profitable deviation from their business practices, vendors tend not to exploit it after our subsidies are removed. This casts doubt on the second explanation.

We are therefore left with the third explanation: vendors' objective functions significantly deviate from monetary profit maximization. This subsumes a broad class of explanations, given that $\phi(\cdot)$ may capture many considerations. In Section 5 we provide descriptive survey evidence for several factors in vendors' objective functions, other than profits, that influence their chosen business scale.

Regardless of the specific considerations captured by $\phi(\cdot)$, we demonstrate that vendors do not maximize their take-home pay; indeed, we find that vendors could increase their monetary profits by more than 60% on average, and despite knowledge of this opportunity and the ability to exploit it, they choose not to.

2.2 Experiment Design

Timeline and Market Selection. Our experiment took place from December 2018 to March 2019 in 20 vegetable markets around Kolkata. Because of the cost of market-wide subsidy interventions we could only intervene in three markets. With so few treated units, we did not randomize. Instead, we chose three markets deterministically based on two criteria. First, we chose markets of roughly medium-size relative to all 20 markets. Second, we chose markets with relatively little price volatility for peas. This reduces the possibility of idiosyncratic market-level shocks confounding the subsidy intervention.[2] Our three intervention markets are Charu Market ($n = 45$ vendors), Sarkar Bazar ($n = 73$), and Alam Bazar ($n = 85$).[3] Figure A1 presents a map of our treatment and control markets.

We break our analysis into three periods: pre-subsidy, subsidy, and post-subsidy. The pre-subsidy period lasted three weeks from December 15, 2018, to January 4, 2019; the subsidy period lasted three weeks from February 23, 2019 to March 15, 2019; and the post-subsidy period lasted two weeks from March 16, 2019 to March 31, 2019. In each period we collected daily data from all vendors in all 20 markets, with surveys starting at the beginning of the day in each market. Surveyors cycled through each vegetable present at the vendor's stall and administered the *daily price survey*, asking: (i) What quantity of this vegetable did you purchase from the wholesale market for today? (wholesale quantity$_t$), (ii) What quantity of this vegetable do you have left over from previous days? (left over$_{t-1}$), (iii) What price did you pay at the wholesale market for this vegetable that you are selling today (in Rs.)? (wholesale price$_t$), and (iv) What price are you charging for this vegetable today (in Rs.)? (price$_t$). We convert these price and quantity variables to a common unit: the unit most commonly reported for a given

[2] Given that there are more markets in our control group than our treatment group, idiosyncratic variation in our outcome variables is more likely to average out in our control group.

[3] Table A1 presents some descriptive statistics on each of our three intervention markets and our 17 control markets. Our intervention markets had 68 vendors on average while our control markets had an average of 86 vendors. Vendors in our intervention markets earned an average of Rs. 355/day compared to vendors in control markets with an average daily profit of Rs. 531/day. 57% (50%) of vendors in our intervention markets sold peas (carrots), while the corresponding number in control markets is 65% (59%).

vegetable, using vendor-reported unit conversion rates (e.g., number of pieces per kilogram).[4] For questions (i) and (ii), surveyors were instructed to visually verify the presence of the vegetable at the market stall, but they did not weigh each item to verify amounts stocked precisely. Answers to the questions above were not used to determine subsidy payouts, and thus participants did not have a monetary incentive to misreport. Subsidy payouts were determined by a separate survey, explained below.

We calculate the daily profit from selling a given vegetable as daily revenue for that vegetable minus its procurement cost. Importantly, this measure excludes the subsidy payouts received by treated vendors. The profit from vegetable v on day t is then

$$\pi_{vt} = \text{price}_{vt} \times \left(\text{wholesale quantity}_{vt} + \text{left over}_{vt-1} - \text{left over}_{vt}\right) \\ - \left(\text{wholesale price}_{vt} \times \text{wholesale quantity}_{vt}\right)$$

With this calculation, we are inferring a vendor's daily quantity sold from the difference between the amount of a vegetable the vendor held at the start of the day (wholesale quantity + left over from previous days) and the amount of the vegetable the vendor reports was left over at the start of the next day. We calculate the total daily profit for a vendor as the sum of profits across vegetables, i.e., $\pi_t = \Sigma_v \pi_{vt}$.[5]

Subsidy Intervention. The 17 control markets received no intervention during any of the periods. During the subsidy period, we offered a subsidy to all vendors in the three intervention markets

[4] In particular, we first define one main unit for each vegetable as the unit that is most commonly used. We then convert prices and quantities that were not reported using the main unit, by applying a conversion factor. We use vendor-specific reported conversion factors where available (these were not collected during the pre-subsidy period), and otherwise, we apply the market-level median conversion factor. If there are fewer than six observations for the unit-conversion factor, we code the observations as missing, given that the conversion factor data might be unreliable in these cases.

[5] There are several instances where we do not observe left over$_{vt}$. First, we did not include this question in our survey in the pre-subsidy period. Second, if a vendor was present in the market on day t but not $t+1$ then we do not observe left over$_{vt}$. When we do not observe left over$_{vt}$, we impute the value based on the vendor's average fraction of left over stock on days we do observe it. We note, however, that these imputations have little consequence, as vendors rarely carry over inventory from one day to the next. Control vendors report positive values of leftover peas (carrots) only 4.7% (8.2%) of the time during the subsidy period, and the average amount of leftover peas (carrots) is only 1.2% (2.5%) of the amount they had at the beginning of the day.

to procure carrots, and to infrequent pea sellers in the same three intervention markets to procure peas. Surveyors administered a *daily subsidy survey* to each treatment-market vendor, during which vendors weighed their stocks of carrots and peas to determine the cash subsidy to be paid at the end of the survey. The *daily subsidy survey* was administered before the *daily price survey* described above, with a different surveyor for each. The subsidy payout was determined based on how much of the carrots and peas vendors had at their stall at the time of surveyor verification, regardless of whether these were procured that morning or left over from the day before. We deliberately structured the subsidy in this manner to eliminate the incentive for vendors to claim leftover produce as newly procured.[6]

The carrot subsidy value was equal to Rs. 20/kg, which was the median procurement cost of carrots during the pre-subsidy period. The maximum subsidized quantity was randomized at the vendor level each week to be either low or high. Vendors who received the low subsidy were compensated for a maximum of 2kg of carrots, while the high quantity was set at the median of the distribution of daily wholesale purchases of carrots during the pre-subsidy period, for each market (7kg in Charu market, and 5kg in Sarkar Bazar and Alam Bazar).

We offered the pea subsidy to the roughly one-third of vendors who were present in the market for more than eight days (i.e. more than 50% of surveyed days) in the pre-subsidy period, but sold peas on fewer than eight days (the median). The pea subsidy value was equal to Rs. 30/kg, which was the median procurement cost of peas during the pre-subsidy period, and once again the subsidized volume was either low or high. Like carrots, the low quantity was set at 2kg, and the high quantity was set at the median of the distribution of daily wholesale purchases of peas during the pre-subsidy period, for each market (8kg in Charu market, 6kg in Sarkar Bazar, and 10kg in Alam Bazar).

[6]Within a given market, the order in which vendors were surveyed was reversed from day to day, primarily to achieve balance in vendor survey times. We note that as an auxiliary assurance, this also introduced considerable uncertainty among vendors about when they would be surveyed, which rendered some "cheating schemes" more difficult. For instance, one might worry that vendors would share the same stock of carrots, transporting it from one stall to the next between surveys, so that several vendors could receive the subsidy for the same produce. Not only did our surveyors never observe such behavior, it would be especially challenging in light of survey timing uncertainty.

We randomized the weekly maximum subsidized quantity to explore whether vendors with greater opportunity to stock the new produce would exhibit more persistent adoption of these products in the post-subsidy period. Unfortunately, our intervention did not induce enough variation in the number of weeks a vendor was exposed to the high subsidies, and as a result this analysis is underpowered. Therefore, we pool vendors who received a high versus low subsidy and focus only on the market-level variation in whether vendors were offered the subsidy.

Finally, we introduced one universal subsidy (carrots) and one targeted subsidy (peas) for two reasons. First, we wanted to ensure that all vendors in intervention markets received at least one subsidy to reduce the likelihood that vendors would feel unfairly treated. Second, the two different subsidies allow us to explore the effects of two margins of vendor expansion. The pea subsidy more closely approximates the "entry" of new vendors who previously did not sell the product and allows us to explore business-stealing or other spillover effects on incumbent vendors. However, note that eligible pea vendors include those who sold peas on up to seven of the days in the pre-subsidy period, so not all of these vendors should be considered new entrants to the market – see Figure A2 for the distribution of the number of days in which a vendor sold peas in the pre-subsidy period. In contrast, the carrot subsidy also induces incumbents to expand their inventory on the intensive margin, which illuminates whether vendors are effectively constraining the supply of even the goods they have chosen to sell.

Follow-up Survey. We revisited the 20 markets in October 2023 to gather descriptive evidence on possible mechanisms for the effects of the subsidies. We aimed to survey a random sample of 20 vendors in each market, ultimately completing 391 surveys. In Appendix A we provide the list of our survey questions and summary statistics. We further describe the survey questions when reviewing the evidence below.

2.3 Empirical Approach

Given our nonrandomized approach, we use a difference-in-differences strategy, checking throughout that pre-subsidy trends in key outcomes are parallel. We estimate the following specification:

$$y_{imt} = \alpha + \beta_1 During_t + \beta_2 Post_t + \beta_3 Treat_m + \gamma_1 During_t \times Treat_m + \gamma_2 Post_t \times Treat_m + \varepsilon_{imt} \quad (1)$$

where y_{imt} is the outcome of interest for vendor i in market m on day t. $During_t$ is a dummy taking a value of one if day t was during the subsidy period, $Post_t$ is a dummy taking a value of one if day t was after the subsidy period, and $Treat_m$ is a dummy taking a value of one if market m is one of the three intervention markets. This is a difference-in-differences model in which our coefficients of interest are γ_1, capturing the effect of our subsidies during the subsidy period, and γ_2, capturing the persistent effect of our subsidies after the subsidy period ended.

Because we have only 20 markets, three of them treated, traditional econometric inference based on large-sample asymptotics is unlikely to perform well in our setting. Instead, we report p-values and confidence intervals computed using the wild bootstrap (Cameron et al., 2008; Roodman et al., 2019), and p-values computed using Fisher's permutation test (Fisher, 1936; Young, 2019), both using markets as the relevant cluster unit. For the permutation test, with 20 markets, there are 1,140 possible combinations of three intervention markets. We re-run a given regression 1,140 times, each time using a different unique combination of hypothetical intervention markets. We then calculate p-values as the fraction of t-statistics at least as large in absolute value as the t-statistic from the original regression. While our tables report the results from both inference approaches, the two approaches largely coincide.

For any vegetable, we consider a vendor $\times$ day observation to be an outlier if (i) the values of wholesale price or quantity, retail price or quantity, or amount a vendor has left over at the end of the day are more than one standard deviation above the 99th percentile of its distribution, or more than one standard deviation below the 1st percentile of its distribution, or (ii) if the vegetable is reported using a unit (e.g. per piece) for which we have insufficient conversion-rate data to convert the report to a common unit. For analyses that only pertain to a single vegetable (Tables 1 and 3), we drop outliers relevant to that vegetable. For analyses that pertain to all vegetables (Table 2), we drop outliers that pertain to all vegetables. We also note that the sample size falls when we analyze effects on prices, as prices are only observed for vendors that bought or sold a given vegetable on a given day.

3 Results

Graphical Summary. Figure 1 summarizes our main findings graphically. First, outcomes in the intervention markets trend similarly to those in the control markets in the pre-subsidy period. We never reject the null that the trends are parallel (Table A2). Second, the subsidies had an important effect during the subsidy period. Vendors in intervention markets were more likely to sell peas and carrots and had higher average profits from pea and carrot sales. They also had higher overall profits during the subsidy period (and as a reminder, our profit measures exclude subsidy payouts). In contrast, prices in intervention markets trended similarly to those in control markets. Third, the effects of our subsidies largely disappeared in the post-subsidy period.

Before turning to the regression results, we note that the panels appear to show a discontinuity for many of the outcomes in control markets between the pre-subsidy and subsidy periods. This is because one and a half months elapsed between our pre-subsidy period and our subsidy period, and it reflects that the aggregate sales of both peas and carrots declined somewhat in that intervening period due to seasonal variation. Nevertheless, neither vegetable went out of season during our study period, with at least 18% of vendors selling each vegetable at all points throughout the study.

The Subsidy Period. During the subsidy period, vendors in intervention markets were 57 percentage points more likely to sell carrots on any given day (95% CI: 39pp – 73pp, $\hat{\gamma}_1$ in column

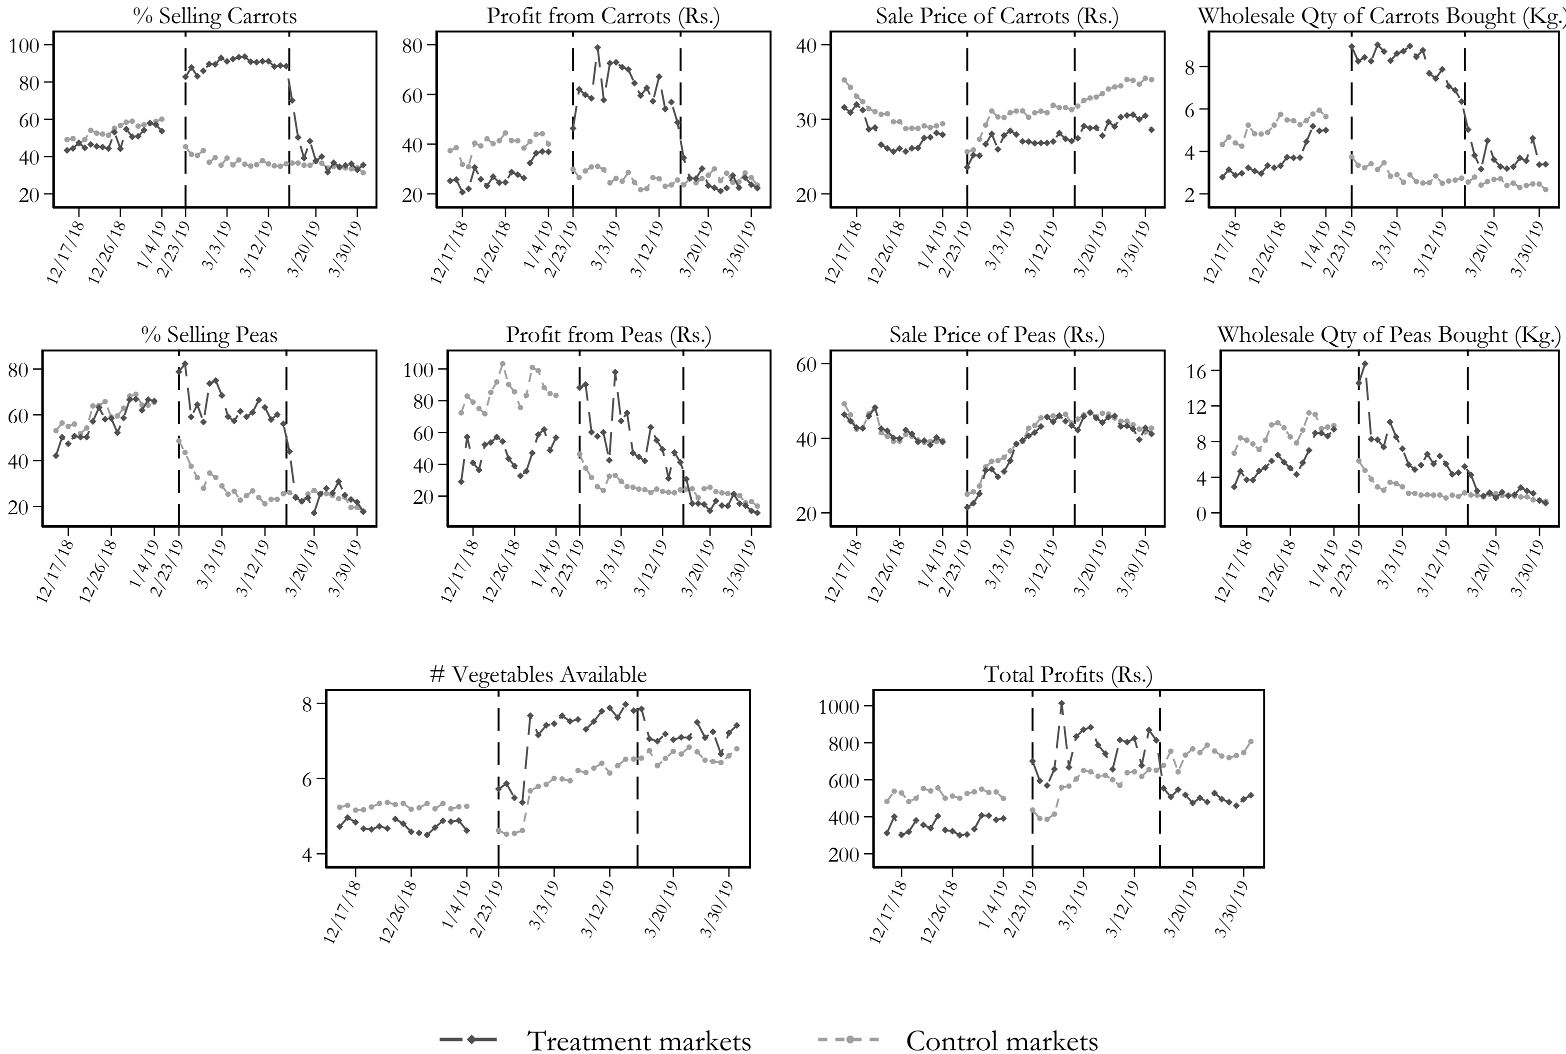

Figure 1: The first row plots the probability a vendor sells carrots, the daily profits accruing from the sales of carrots, the vendor's sale price of carrots, and the quantity of carrots procured in kilograms. The second row plots the same four outcomes for peas. The third row plots the number of types of vegetables a vendor stocks on a given day, and the vendor's daily total profits. The first vertical line demarcates the start of the subsidy period and the second line demarcates the end of the subsidy period. Profits are calculated as: (amount of vegetable at the start of the day - amount left over at the end of the day)* sale price - (amount procured at the start of the day * procurement cost). On days where the amount left over was not observed, we impute the vendor's average amount left over all days in which it was measured. Our measure of profit does not include the subsidy vendors received as part of our intervention.

1, Table 1) and 39 percentage points more likely to sell peas (95% CI: 14pp – 59pp, column 5). On average vendors in intervention markets procured an extra 6.0kg of carrots per day (95% CI: 3.7kg – 7.7kg, column 3) and an extra 6.7kg of peas (95% CI: 3.1kg – 10.7kg, column 7). These are substantial increases relative to average procurement volumes of 3.4kg of carrots and 5.9kg of peas in intervention markets in the pre-subsidy period.[7] Importantly, these increases are not explained by vendors misreporting quantities to claim the subsidy, as these reports from the daily price survey were not used to determine the subsidy amount. As explained in Section 2.2, the subsidy amount was instead determined through a separate survey in which a different surveyor weighed the vendor's produce.

Profits for vendors in intervention markets increased during the subsidy period, even though our profit measures exclude the value of the subsidy. Profits from carrots increased by Rs. 44.8 per day (95% CI: Rs. 20.8 – Rs. 57.3, column 4) compared to an average profit from carrots of Rs. 25.4 per day in the pre-subsidy period. Profits from peas increased by Rs. 59.7 per day (95% CI: Rs. 12.0 – Rs. 96.2, column 8) compared to an average profit from peas of Rs. 47.7 per day in the pre-subsidy period.

In addition, there is no evidence that our intervention caused sale prices for peas and carrots to decline (columns 2 and 6). This may indicate that vendors had not been meeting customers' full demand prior to our intervention, or that customers were diverted from nearby markets to treatment markets. However, our confidence intervals are consistent with our subsidy having caused a modest price reduction. We discuss these points further in Section 4.1.

After the Subsidy Ended. The impacts of the subsidy diminished after the subsidy period ended. There is no statistically significant increase in the likelihood that vendors in intervention markets

[7] Indeed, these point estimates suggest that vendors increased their average procurement of peas and carrots by more than the average subsidized quantity. This may be because once a vendor is induced to procure any positive quantity of peas or carrots (or induced to continue procuring peas or carrots if they would otherwise have ceased doing so), they find it worthwhile to procure more than the subsidized quantity. This would be reasonable behavior, for example, if they have negotiated a temporary exemption from a collusive norm during the experiment, and want to take full advantage of it. We discuss the phenomenon of going beyond the subsidy cap further below, where we interpret it as revealed preference evidence for a profitable deviation.

sell carrots or peas ($\hat{\gamma}_2$ in columns 1 and 5, Table 1). Vendors in intervention markets procured only an additional 1.9kg of carrots per day (95% CI: -0.1kg – 4.0kg) and only procured an additional 2.6kg of peas (95% CI: -0.8kg – 6.4kg). These are increases relative to the pre-period, but a roughly two-thirds drop relative to the subsidy period. Nevertheless, there are signs of persistent effects for some vendors, particularly at the intensive margin, in that the post-subsidy effects on wholesale quantities are marginally significant for both carrots and peas ($p = 0.07$ with the wild bootstrap). The impact on wholesale quantity also translates into a persistent positive effect on profits from peas ($p = 0.08$ with the wild bootstrap).

The fact that there are more persistent effects on the intensive margin (quantity procured) than on the extensive margin (whether selling at all) may be consistent with a story in which collusive norms apply more to product entry than to expansion after entry. In particular, entry is more easily monitored than expansion. We discuss this further in Section 5.2, when considering the possibility of collusive norms explaining our results.

We note that 100% (99%) of vendors in intervention markets who sold carrots (peas) experienced positive profits from sales of those vegetables during the subsidy period. Thus, the fact that many vendors stopped selling the additional produce is not driven by the possibility that a majority of vendors found it marginally unprofitable to sell peas and carrots and only those who experienced the profit increase continued selling peas and carrots. Rather, many vendors who experienced positive profits from sales of peas and carrots nevertheless chose to stop selling these vegetables after our intervention concluded.

Nevertheless, the results indicate that not *all* vendors stopped selling peas and carrots after the subsidy was withdrawn. What differentiates vendors who continued selling peas and carrots from the majority that did not? To explore this, we estimate a generalized random forest model to predict post-subsidy treatment effects using only pre-subsidy period vendor characteristics, in the spirit of Chernozhukov et al. (2018) and Athey et al. (2019) (for full details, see Appendix B). Our predicted heterogeneous post-subsidy treatment effects predict post-subsidy effects on pea profits, but not carrot profits, and so we focus on the former (Table A3). Accordingly, we

Table 1: Subsidy Impacts: Carrots and Peas

	Carrot				Peas			
	Prob. of selling (%) (1)	Sale price (Rs.) (2)	Wholesale qty. (kg) (3)	Profits (Rs.) (4)	Prob. of selling (%) (5)	Sale price (Rs.) (6)	Wholesale qty. (kg) (7)	Profits (Rs.) (8)
β_3 Treat	–0.05	–2.62	–1.40	–12.71	–0.04	–0.01	–2.53	–33.29
	[–0.31, 0.19]	[–6.88, 2.17]	[–3.74, 1.02]	[–28.63, 7.82]	[–0.33, 0.28]	[–3.51, 3.74]	[–6.83, 1.45]	[–74.99, 19.33]
	{ 0.647 }	{ 0.087 }	{ 0.293 }	{ 0.303 }	{ 0.548 }	{ 0.984 }	{ 0.076 }	{ 0.074 }
	⟨ 0.603 ⟩	⟨ 0.147 ⟩	⟨ 0.263 ⟩	⟨ 0.141 ⟩	⟨ 0.604 ⟩	⟨ 0.989 ⟩	⟨ 0.148 ⟩	⟨ 0.049 ⟩
γ_1 Treat × During Subs	0.57	–0.44	5.99	44.75	0.39	–0.81	6.72	59.67
	[0.39, 0.73]	[–3.03, 2.35]	[3.74, 7.70]	[20.81, 57.32]	[0.14, 0.59]	[–4.14, 2.73]	[3.13, 10.66]	[11.95, 96.16]
	{ 0.001 }	{ 0.633 }	{ < 0.001 }	{ 0.001 }	{ 0.021 }	{ 0.764 }	{ 0.014 }	{ 0.031 }
	⟨ < 0.001 ⟩	⟨ 0.684 ⟩	⟨ < 0.001 ⟩	⟨ < 0.001 ⟩	⟨ < 0.001 ⟩	⟨ 0.686 ⟩	⟨ < 0.001 ⟩	⟨ < 0.001 ⟩
γ_2 Treat × After Subs	0.10	–2.04	1.94	8.79	0.05	–0.87	2.58	26.52
	[–0.07, 0.26]	[–9.58, 4.70]	[–0.07, 3.97]	[–11.38, 30.80]	[–0.18, 0.25]	[–4.46, 1.63]	[–0.78, 6.38]	[–21.16, 60.69]
	{ 0.418 }	{ 0.152 }	{ 0.065 }	{ 0.238 }	{ 0.459 }	{ 0.561 }	{ 0.065 }	{ 0.083 }
	⟨ 0.354 ⟩	⟨ 0.188 ⟩	⟨ 0.191 ⟩	⟨ 0.200 ⟩	⟨ 0.546 ⟩	⟨ 0.532 ⟩	⟨ 0.069 ⟩	⟨ 0.055 ⟩
Pre-subsidy intervention market mean	0.49	27.91	3.43	25.41	0.57	41.80	5.94	47.73
Wild Bootstrap p-value: $\gamma_1 = \gamma_2$	<0.001	0.292	<0.001	0.007	0.002	0.960	0.004	0.019
Fisher p-value: $\gamma_1 = \gamma_2$	0.002	0.357	<0.001	<0.001	<0.001	0.976	0.003	0.022
Number of Vendors	1631	1470	1631	1631	1631	1489	1631	1631
Number of Observations	55218	25073	55218	55213	55243	22657	55243	55241

Notes: This table estimates specification 1 on our full sample. Coefficients for During and Post not shown. 95% wild bootstrap confidence intervals are in [], wild bootstrap p-value is in { }, and Fisher permutation p-value is in ⟨⟩. Columns 1 - 4 present outcomes for carrots, and 5 - 8 for peas. The outcome in columns 1 and 5 is whether the vendor sells carrots or peas on the given day, the outcome in columns 2 and 6 measure the vendor's anticipated sale price for the relevant vegetable, the outcome in columns 3 and 7 measure the wholesale quantity procured of the relevant vegetable, and the outcome in columns 4 and 8 measure the daily profits accrued from the relevant vegetables. Profits are calculated by computing (amount of vegetable at the start of the day - amount left over at the end of the day)*anticipated sale price - (amount procured at the start of the day * procurement cost). On days where amount left over was not observed, we impute the vendor's average amount left over all days in which it was measured. Our measure of profit does not include the subsidy vendors received as part of our intervention.

restrict this analysis to the vendors eligible for the pea subsidy (those that were infrequent pea sellers in the pre-subsidy period). Among these pea subsidy-eligible vendors, those who sold peas more frequently in the pre-period, as well as those that sold fewer types of vegetables and those with more total sales are predicted to have larger post-subsidy treatment effects on pea profits (Table A4). That vendors who sold peas more often in the pre-period have a larger post-subsidy treatment effect is consistent with the possibility of collusive norms that discourage vendors from entering new product categories. We return to this possibility in Section 5.2.[8]

Beyond Carrots and Peas. Table 2 presents the impact of our subsidies on aggregate vendor outcomes, rather than those corresponding to either carrots or peas. The aggregate picture is largely consistent with the results from the individual vegetables. During the subsidy period, total costs of wholesale purchases in intervention markets rose by Rs. 690 per day (95% CI: Rs. 218 – Rs. 1,120, column 1) compared to an average cost of wholesale purchases of Rs. 825 in intervention markets in the pre-subsidy period. Average vendor profits rose by Rs. 228 per day (95% CI: Rs. -103 – Rs. 478, column 3) compared to an average profit of Rs. 342 in the pre-subsidy period – a 67% increase. On average vendors stocked an additional 2.0 (95% CI: 0.5 – 3.3, column 4) types of vegetables during the subsidy period compared to an average of 4.7 products stocked per vendor in the pre-subsidy period. Once again, these effects diminish after our subsidy concluded, with no statistically significant increase on any of these outcomes.

Interestingly, the effect of our intervention on total profits is larger than the sum of the effects on the profits from sales of peas and carrots. This difference is only statistically significant at the 10% level when using the wild bootstrap, and not statistically significant at the 10% level using the Fisher permutation test. Similarly, the effect on the cost of total wholesale purchases is larger than the sum of the effects on the costs of purchases of peas and carrots. This difference

[8] A natural question is whether the vendors that are predicted to have large during-subsidy profit gains are also those that have larger post-subsidy treatment effects on profits. We explored this with the same generalized random forest approach, but unfortunately the resulting predicted during-subsidy treatment effects are not predictive of during-subsidy profits, whether for carrots or peas (Table A5).

Table 2: Subsidy Impacts: Vendor-Level Outcomes

	Aggregate			
	Total Cost of Wholesale Purchases (Rs.)	Sales (Rs.)	Profits (Rs.)	# vegetables available
	(1)	(2)	(3)	(4)
β_3 Treat	–448.49	–589.02	–140.53	–0.50
	[–1029.09, 49.28]	[–1291.86, 51.20]	[–353.59, 79.76]	[–2.56, 1.56]
	{ 0.057 }	{ 0.057 }	{ 0.120 }	{ 0.572 }
	⟨ 0.032 ⟩	⟨ 0.029 ⟩	⟨ 0.075 ⟩	⟨ 0.471 ⟩
γ_1 Treat × During Subs	689.60	917.98	228.32	1.97
	[218.32, 1119.62]	[201.65, 1524.12]	[–102.95, 477.54]	[0.51, 3.26]
	{ 0.025 }	{ 0.031 }	{ 0.073 }	{ 0.029 }
	⟨ < 0.001 ⟩	⟨ 0.005 ⟩	⟨ 0.025 ⟩	⟨ 0.013 ⟩
γ_2 Treat × After Subs	527.12	466.27	–60.84	1.16
	[–134.64, 1065.87]	[–325.96, 1086.38]	[–373.98, 248.20]	[–0.58, 2.65]
	{ 0.290 }	{ 0.354 }	{ 0.316 }	{ 0.323 }
	⟨ 0.268 ⟩	⟨ 0.342 ⟩	⟨ 0.403 ⟩	⟨ 0.229 ⟩
Pre-subsidy intervention market mean	825.12	1167.30	342.18	4.73
Wild Bootstrap p-value: $\gamma_1 = \gamma_2$	0.571	0.148	0.026	0.050
Fisher p-value: $\gamma_1 = \gamma_2$	0.583	0.139	0.002	0.059
Number of Vendors	1628	1628	1628	1628
Number of Observations	52898	52898	52898	52898

Notes: This table estimates specification 1 on our full sample. Coefficients for During and Post not shown. 95% wild bootstrap confidence intervals are in [], wild bootstrap p-value is in {}, and Fisher permutation p-value is in ⟨⟩. The outcome in column 1 is the total cost of wholesale purchases on a given day, the outcome in column 2 is the vendor's total revenues on a given day accruing from all produce, the outcome in column 3 is the daily profits accrued from all produce, and the outcome in column 4 is the number of distinct types of vegetables a vendor has available on a given day. Profits are calculated by computing (amount of vegetable at the start of the day - amount left over at the end of the day)*anticipated sale price - (amount procured at the start of the day * procurement cost). On days where amount left over was not observed, we impute the vendor's average amount left over all days in which it was measured.

is statistically significant at the 10% level using both the wild bootstrap and Fisher permutation tests. Thus, these results leave open the possibility that our subsidy "crowded in" the sale of complementary produce. We explore this possibility more fully in Section 4.2.

Permutation Test Robustness. While we have used two approaches to inference, one limitation of our permutation test is that it allows for treatment market combinations in which treated markets are not mid-sized or do not have low price volatility – failing to satisfy the two criteria we used when selecting the three treatment markets. To account for this, we re-run the permutation tests in Tables 1 and 2, this time constraining the re-randomization of treatment, such that eight of the markets are kept in the control group: the three largest, the three smallest, and the three with the largest pre-period volatility in pea prices (one market is in two of these groups, giving us eight excluded markets overall). Fortunately, the findings are similar with this alternative approach to inference (Tables A6 and A7).

Pea-Subsidy Impacts by Eligibility. Recall that while everyone in intervention markets received a carrot subsidy, only infrequent pea sellers were eligible for the pea subsidy. We now turn to the differential effects of the pea subsidy on vendors in intervention markets who did and did not receive the subsidy. These are presented in Table 3, which again uses Specification 1, but now splits the sample by pea subsidy eligibility.

The qualitative patterns for vendors eligible for the pea subsidy are the same as in the previous analyses. During the subsidy period, eligible vendors in intervention markets were 66 percentage points (95% CI: 54pp – 75pp) more likely to stock peas on any given day, they procured an extra 7.8kg of peas per day (95% CI: 4.1kg – 10.3kg), and earned an extra Rs. 65.6 per day (95% CI: Rs. 30.7 – Rs. 92.4) from the sale of peas. Unlike in the previous analysis, there is evidence of a price increase of Rs. 4.4/kg (95% CI: Rs. 0.9 – Rs. 9.2), statistically significant at the 5% level using the wild bootstrap, but not when using the permutation test. Qualitative evidence we collected suggests this may be because vendors substituted toward

Table 3: Subsidy Impacts: By Pea Subsidy Eligibility

	Eligible for Pea Subsidy (Infrequent Pea Seller)				Ineligible for Pea Subsidy (Frequent Pea Seller)			
	Prob. of selling (%)	Sale price (Rs.)	Wholesale qty. (kg)	Profits (Rs.)	Prob. of selling (%)	Sale price (Rs.)	Wholesale qty. (kg)	Profits (Rs.)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
β_3 Treat	–0.03	0.38	–1.00	–17.39	0.01	–0.05	–2.87	–37.71
	[–0.06, 0.04]	[–10.24, 14.98]	[–2.09, 0.33]	[–39.57, 20.13]	[–0.11, 0.12]	[–3.44, 3.54]	[–9.06, 2.37]	[–98.96, 25.16]
	{ 0.123 }	{ 0.831 }	{ 0.074 }	{ 0.100 }	{ 0.764 }	{ 0.948 }	{ 0.293 }	{ 0.075 }
	⟨ 0.053 ⟩	⟨ 0.909 ⟩	⟨ 0.060 ⟩	⟨ 0.033 ⟩	⟨ 0.793 ⟩	⟨ 0.961 ⟩	⟨ 0.268 ⟩	⟨ 0.104 ⟩
γ_1 Treat × During Subs	0.66	4.35	7.84	65.59	0.16	–2.00	5.37	50.90
	[0.54, 0.75]	[0.93, 9.20]	[4.11, 10.34]	[30.65, 92.36]	[–0.10, 0.46]	[–5.84, 1.37]	[0.86, 11.26]	[–11.48, 96.70]
	{ < 0.001 }	{ 0.033 }	{ < 0.001 }	{ 0.005 }	{ 0.078 }	{ 0.301 }	{ 0.044 }	{ 0.070 }
	⟨ 0.003 ⟩	⟨ 0.181 ⟩	⟨ 0.006 ⟩	⟨ 0.008 ⟩	⟨ 0.125 ⟩	⟨ 0.311 ⟩	⟨ 0.007 ⟩	⟨ 0.018 ⟩
γ_2 Treat × After Subs	0.09	0.25	1.74	19.36	–0.00	–0.56	2.66	27.51
	[–0.01, 0.21]	[–4.07, 3.39]	[0.30, 3.21]	[–15.90, 40.23]	[–0.23, 0.22]	[–4.53, 2.76]	[–2.28, 7.45]	[–40.23, 81.82]
	{ 0.065 }	{ 0.894 }	{ 0.033 }	{ 0.099 }	{ 0.943 }	{ 0.781 }	{ 0.247 }	{ 0.188 }
	⟨ 0.121 ⟩	⟨ 0.914 ⟩	⟨ 0.010 ⟩	⟨ 0.012 ⟩	⟨ 0.969 ⟩	⟨ 0.754 ⟩	⟨ 0.274 ⟩	⟨ 0.250 ⟩
Pre-subsidy intervention market mean	0.18	39.38	1.63	10.56	0.85	42.15	8.98	73.99
Wild Bootstrap p-value: $\gamma_1 = \gamma_2$	<0.001	0.235	<0.001	<0.001	0.041	0.622	0.040	0.061
Fisher p-value: $\gamma_1 = \gamma_2$	<0.001	0.539	0.006	0.009	0.025	0.554	0.043	0.060
Number of Vendors	562	480	562	562	1069	1009	1069	1069
Number of Observations	19763	3687	19763	19761	35480	18970	35480	35480

Notes: This table estimates specification 1 on our full sample. Coefficients for During and Post not shown. 95% wild bootstrap confidence intervals are in [], wild bootstrap p-value is in { }, and Fisher permutation p-value is in ⟨⟩. Columns 1 - 4 present outcomes for vendors eligible for the peas subsidy, and 5 - 8 for vendors ineligible for the subsidy. The outcome in columns 1 and 5 is whether the vendor sells peas on the given day, the outcome in columns 2 and 6 measure the vendor's anticipated sale price for peas, the outcome in columns 3 and 7 measure the wholesale quantity procured, and the outcome in columns 4 and 8 measure the daily profits accrued from the sale of peas. To be eligible for the pea subsidy, vendors must have been present in the market for more than eight days in the pre-subsidy period and have sold peas on fewer than eight days. Profits are calculated by computing (amount of peas at the start of the day - amount left over at the end of the day)*anticipated sale price - (amount procured at the start of the day * procurement cost). On days where amount left over was not observed, we impute the vendor's average amount left over all days in which it was measured. Our measure of profit does not include the subsidy vendors received as part of our intervention.

higher-quality peas. Consistent with this, Table A8 shows suggestive evidence that the procurement cost of peas also increased (column 1), such that *markups* for peas were stable or even declined somewhat in response to our subsidy (column 2). Once again, the effects in Table 3 diminish considerably after the subsidy is removed.

We find no evidence of business-stealing effects. In fact, the patterns for vendors who were ineligible for the pea subsidy are largely the same as the patterns for eligible vendors. These vendors procured more peas and earned higher profits from the sale of peas during the subsidy period, despite not having access to a pea subsidy. Qualitative evidence we collected after the intervention suggests that this is in part due to informal arrangements between vendors who sold peas prior to our intervention, typically larger vendors, and vendors who did not. Specifically, these larger vendors would procure and transport additional peas at the wholesale market and then sell them to vendors who received a subsidy. Note, however, that this remains consistent with our basic narrative. It is possible for many vendors to increase their sales volume and profits by purchasing and selling more carrots and peas, but even after directly experiencing these opportunities firsthand, they refrained from exploiting them after our intervention.

In part this may also be driven by the fact that, as we show in Section 4.2, carrots and peas are complementary vegetables. Vendors who were ineligible for the pea subsidy nevertheless received the carrot subsidy, inducing them to stock additional carrots, which may, in turn, have induced them to stock additional peas.

For completeness, in Appendix Table A9 we present analogous results from the carrot subsidy, disaggregated by vendors who were frequent or infrequent carrot sellers in the pre-subsidy period. The results are qualitatively the same. Both types of vendors experienced an increase in sales and profits of carrots during the subsidy period, and then these increases diminished in the post period.

Having described our core results, we consider three threats to our finding that our treatment substantially increased profits during the subsidy period: spillovers from treatment to control markets, mismeasurement of profits, and manipulation of procurement volumes.

Accounting for Potential Spillovers from Treatment to Control Markets. An identifying assumption of our difference-in-differences approach is that our intervention in treatment markets did not influence outcomes in our control markets. This assumption would be violated if there were demand-side spillovers, whereby vendors who expanded their scale in treated markets diverted customers from control markets, or supply-side spillovers, whereby vendors in treated markets purchased so many peas and carrots from wholesale markets that procurement prices for control vendors increased, reducing their ultimate scale of operation.

To rule out demand-side spillovers, we re-estimate Specification 1 after excluding the control markets that are most likely to be affected by our intervention in treatment markets. Specifically, within each market m we asked all vendors which nearby market customers would be most likely to shop from if they were not to shop at market m. For each treatment market m, we drop any market from our sample that is in the top three markets that are most frequently listed as likely competitors. This results in dropping only three markets from the analysis, as most of the frequently listed competitor markets are not within our experimental sample of 20 markets.

To rule out supply-side spillovers, we re-estimate Specification 1, but remove from the analysis the control vendors most likely to be affected through supply-side spillovers. Specifically, for each vendor, we know the primary wholesale market from which they procure their produce. 45% of vendors in our treatment markets primarily procure their produce from two wholesale markets, while only 7% of vendors in our treatment markets procure their produce from the next most common wholesale market. We drop from our analysis all control-market vendors served by these two wholesale markets, which removes 626 vendors.

The results are presented in Tables A10 and A11 for demand-side spillovers, and Tables A12 and A13 for supply-side spillovers. Importantly, none of the patterns are qualitatively altered. Vendors are significantly more likely to sell peas and carrots during the subsidy period and earn significantly higher profits from the sale of peas and carrots, as well as total profits. Again, all of these effects diminish significantly after the subsidy is removed.

Measurement Error in Profits. In this subsection we consider two types of measurement error in profits: measurement error in sale prices and omitted costs. We conclude this subsection with a revealed preference argument that selling peas and carrots was profitable, which does not require the direct measurement of profits.

Prices. Could our estimated treatment effects on profits be biased upward due to mismeasurement of sale prices? We measure each vegetable's sale price at only one point in the day (i.e. at the vendor-day level), rather than for each separate transaction. For each vegetable we asked, "What price are you charging for [Vegetable Name] today?" One concern is that prices change throughout the day. This would lead us to overestimate the profits of treated vendors if (i) treated vendors stay later in the retail market to sell their additional produce, and (ii) vendors lower prices later in the day, with our team more often measuring prices earlier in the day, before they had fallen. Against this concern, we find no evidence of markups falling throughout the day, using variation in the time of day at which a surveyor measured prices (Figure A3).

A related concern is that prices vary depending on the customer. We capture vendors' stated selling price, but it is possible that they offer discounts for some retail customers. Vendors might also consume some of their own produce at a nominal sale price of zero. And they might sell some of their produce to other vendors at a deep discount. We do not believe that these sources of measurement error are meaningfully biasing our results. If treatment market vendors were disproportionately offering discounts of the nature described above, they should also lower their stated selling price ascribed to the bulk of their customers, to minimize the volume of peas and carrots that they would need to sell at a deep discount. Yet we document above that our intervention had minimal impact on stated selling prices. Further, as we document below, vendors' procurement of peas and carrots very often exceeded the subsidized quantity. That they were willing to purchase marginal units at the unsubsidized price suggests they did not anticipate selling these marginal units at a deep discount. Nevertheless, we cannot fully rule out that some treated vendors offered frequent discounts on peas and carrots, yet reported a higher "ordinary sale price" to our surveyors. Ignoring such discounts would bias our profit effects upward, though the discounts would have to be substantial and frequent to fully account for the more than 60% increase in profits.

Omitted Costs. A second limitation of our profit measure is its omission of several important costs, such as the cost of renting a spot in a market, the cost of transporting vegetables from the wholesale market to the market stall, the cost of hired labor, and the opportunity cost of the vendor's own labor. Nevertheless, under reasonable assumptions, excluding these costs biases downward our main estimate that subsidizing the procurement of peas and carrots increases profits by more than 60%. We formalize this argument through a simple model.

Suppose that a vendor's business generates r_0 in revenue from c_0 expenses on vegetable procurement (i.e. the revenues and costs that we measure at baseline). Suppose further that, conditional on scaling their business to include additional peas and carrots it would generate $r_1 = (1+s)r_0$ revenue and $c_1 = (1+s)c_0$ expenses from vegetable procurement, for some scalar $s > 0$. Then our treatment effect corresponds to

$$\frac{(r_1 - c_1) - (r_0 - c_0)}{r_0 - c_0} = s \equiv \hat{\tau}.$$

Now consider any unmeasured fixed cost f – i.e. costs that do not scale with vegetable procurement – such as the rental expenses of a vendor's market stall.[9] Accounting for a fixed cost reduces the denominator of the treatment effect (measured profits are lower in the control group), but not the numerator (since the fixed cost cancels out when considering the treatment-control difference in profits). As a result, accounting for fixed costs would increase our estimate of the impact of scaling the vendor's business on their profits to

$$\frac{(r_1 - c_1) - (r_0 - c_0)}{r_0 - c_0 - f} > \hat{\tau}.$$

Having established that properly accounting for fixed costs would increase our estimated treat-

[9] In our survey of 391 vendors four years after the experiment concluded, the mean reported monthly cost of selling in the market is INR 6,459, or roughly 70 USD.

ment effect, we now assume them to be zero and turn to unmeasured variable costs v, which scale with the amount of produce procured. These would include the cost of transporting the additional produce as well as the opportunity cost of the vendor's labor required to procure and sell the additional produce. We begin by assuming that the total unmeasured variable cost scales proportionately with the costs of vegetable procurement – the variable cost is vc_0 at the baseline scale, and is vc_1 at the larger scale induced by our intervention. Then, accounting for these variable costs would not change the estimate of our impact. That is,

$$\frac{(r_1-(1+v)c_1)-(r_0-(1+v)c_0)}{r_0-(1+v)c_0}=\hat{\tau}$$

Accounting for proportional variable costs reduces the level of profits in both the treatment and control conditions by the same proportion, leaving the percentage difference between them unchanged.

Departing from our assumption that unmeasured variable costs scale proportionately with the costs of vegetable procurement, if instead the unmeasured variable costs scaled less than proportionately, accounting for the variable costs would only serve to increase our estimated treatment effect. Therefore, the only unaccounted-for expenses that could weaken our results are variable expenses that scale *more* than proportionately with the cost of vegetable procurement. We do not believe these types of variable monetary expenses are likely. For instance, in our 2023 survey of 391 vegetable vendors, the mean cost of transportation for one procurement trip is INR 198. When asked the cost of transportation if the procurement quantity was doubled, the average vendor reports a 67% increase in cost, and 97% of vendors report an increase of 100% or less. This suggests that transport costs scale less than proportionately with the cost of procurement, consistent with quantity discounts.

Regarding labor costs, we note that in our setting, hired labor is extremely rare. In our 2023 survey, we asked vendors how they would manage any extra work if they were to expand their business. Only 7% said that they would hire workers. Relatedly, no vendor reported that they would arrange for childcare. This leaves the opportunity cost of the vendor's labor, which we consider in Section 5.2. We note there that a vendor's disutility of effort may be convex, and thus potentially high at the margin, deterring business expansion.

A Revealed Preference Argument Selling Peas and Carrots was Profitable. We close this subsection with an argument that selling peas and carrots was profitable, which does not rely on measuring profits directly. For treatment-market vendors during the subsidy period, Figure A4 presents a histogram of the amounts of peas and carrots they stocked, as a fraction of their subsidized quantity (recall that the subsidized quantity was randomized at the vendor-by-week level for each of the three weeks of the subsidy period). For peas (carrots) this fraction exceeds one 72% (81%) of the time for vendors assigned to the low subsidy and 38% (61%) of the time for vendors assigned to the high subsidy. If, without the subsidy, selling these products was unprofitable, then vendors might purchase quantities up to the subsidized amount, as these would be reimbursed, but they would rarely procure more than the subsidized amount, such that the marginal produce was unsubsidized. The fact that procurement volume exceeds the low subsidy amount in most cases strongly indicates that, at least for low volumes, selling peas and carrots is profitable even without the subsidy.

Relatedly, one possible concern about our positive difference-in-differences effects on procurement quantities would be that these effects exist only because vendors procure up to the maximum subsidized quantity, not finding it profitable to go beyond this amount – i.e. not finding it profitable to procure additional vegetables in the absence of the subsidy. Against this interpretation, we also estimate positive effects on a dummy variable for procuring more than the maximum subsidized quantity of carrots and peas (Table A14).

Manipulation of Procurement Volumes. Could vendors have manipulated how our surveyors recorded their procurement volumes, in order to benefit from the procurement subsidy without expanding procurement? We believe this is unlikely. First, as explained earlier, we had one survey for determining the subsidy payout and a separate survey for measuring the quantities and prices used as inputs to our profit measure. Thus, there was no monetary incentive for participants to misreport the variables we use to measure profits.

Second, as part of the survey to determine subsidy payouts, surveyors were trained to observe vendors as they weighed their produce, ensuring that they had procured the requisite peas and carrots. One threat to the integrity of this process would be if vendors could sell or transfer peas or carrots to one another, so that each of them could receive the subsidy for the same set of produce. That is, one vendor could procure peas, have them present at their stall when our surveyor came to verify procurement and disburse the subsidy, and then the vendor could give or sell those same peas to another vendor who would also use them to satisfy the requirement for the subsidy. While technically possible, this manipulation is unlikely as our surveyors conducted the surveys for every vendor in a given market one right after the other. Thus, if such a manipulation were occurring, it would be very likely that our surveyors would detect it.

A related concern is that vendors had leftover peas or carrots from the prior day, but claimed to procure them in the morning to benefit from the subsidy. Anticipating this possibility, we designed the reimbursement procedure such that it applied to any peas and carrots that the vendor had in stock, regardless of when they were procured. Thus, vendors had no incentive to lie about whether produce was left over from the day before.

4 Unresolved Puzzles

In this section we discuss two puzzles raised by our results, and consider various explanations.

4.1 The Lack of Detectable Impact of the Subsidy on Prices

Our results indicate that vendors were able to expand their sales of peas and carrots without detectably lowering prices (columns 2 and 6, Table 1). How could this be?

One possibility is that in the absence of our intervention, vendors were effectively rationing peas and carrots, thereby not fully meeting demand at market prices. An implication would be that vendors frequently experienced stockouts of their produce, running out before the end of the day. We find some support for this implication. Though we do not have granular within-day data on inventory levels, we do observe end-of-day inventory levels. In the control markets during the subsidy period, 95% (92%) of vendors who stocked peas (carrots) ran out by the end of the day, indicating that within-day stock-outs are likely to be common.

If vendors were indeed not meeting demand at market prices, that would raise two further questions. First, why did vendors not stock more peas and carrots to meet demand and increase their profits? On this, we note that the core empirical observation of this study is that vendors do not seem to be exploiting precisely this type of opportunity to increase their profits, and we suggest possible explanations in Section 5. The second question concerns why vendors do not raise their prices to equate demand with supply. While we do not have a definite answer, it seems plausible to us that heuristics and norms constrain optimal pricing. Consistent with this, we note that there is almost no variation in carrot and pea prices within a market after conditioning on procurement costs (Figure A6). In addition, Brown and Tommasi (2025) report related findings from street food vendors in the same setting, Kolkata. They estimate precise null effects of a large sanitation-related capital infusion on prices, despite evidence of high consumer willingness to pay for improved sanitation practices. Their interpretation is that norms constrain price changes. In support of this, they find that 86% of vendors agreed that substantial deviations from prevailing prices could trigger social sanctions.

Aside from rationing, an alternative possibility is that vendors *were* meeting demand for peas and carrots at market prices, but that demand is sufficiently elastic that we could not detect downward pressure on prices. To explore this idea, we estimate the demand elasticity for peas and carrots using two independent sources of supply shocks: the shock induced by our experimental subsidy, and observational variation induced by the number of vendors choosing to sell peas and carrots on any given day in control markets. Two insights result from the exercise. First, the confidence interval on the estimated demand elasticity from our experiment is large, resulting from meaningfully large confidence intervals on both the price and quantity impacts of our subsidy. Second, both sources of variation result in similar estimates of the demand elasticity, with overlapping confidence intervals. This bolsters the credibility that our estimated price impacts from the experiment do not arise from a mismeasurement of prices, but rather reflect highly elastic demand.

To estimate the demand elasticity for peas and carrots using observational data, we restrict the sample to control markets during the subsidy period. We aim to estimate the market-level demand curve

$$ln(Price)_{mt} = \alpha + \beta ln(Quantity)_{mt} + \varepsilon_{mt} \tag{2}$$

where the unit of observation is the market × day. $ln(Price)_{mt}$ is the log of average stated selling price of either peas or carrots in market m on day t, and $ln(Quantity)_{mt}$ is the log of quantity of either peas or carrots in market m sold on day t. β is the inverse elasticity of demand; a small negative β indicates highly elastic demand.

To circumvent endogeneity concerns, we show estimates with and without market and day fixed effects, we instrument for quantity sold with the number of vendors present in market m on day t selling peas or carrots, and we control for the total number of vendors present in the market to account for general levels of activity. Our approach to identification allows for any market- or day-level confounds, but not for market-day-level confounds conditional on the total number of vendors present in the market.[10] Table A15 presents the first stage: each additional vendor selling either peas or carrots results in a 5–7 percent increase in sales of that vegetable, with an F-statistic ranging from 16 to 52 across specifications.

Our instrumented inverse elasticity estimates are presented in Table A16.[11] Our preferred

[10] One example would be if some markets cater to one religious group, with that group celebrating a festival at which carrots are consumed—such a case may mean a positive shock to both market supply of carrots (without necessarily changing the total number of vendors present) and market demand for carrots, whereas our instrument presumes only a shock to market supply.

[11] See Table A17 for an uninstrumented estimation of Specification 2.

estimates include market and day fixed effects, and indicate an inverse demand elasticity for peas of -.03 and for carrots of -.04, though neither is statistically significant (columns 3 and 6). Our estimates without day fixed effects are somewhat larger in magnitude, at -.27 for peas and -.12 for carrots, with both statistically significant at the 5% level (columns 2 and 5). The fact that removing day fixed effects leads to less elastic demand makes intuitive sense: without day fixed effects, we allow for identification to come from Kolkata-level supply shocks over time, including the fact that peas are gradually going out of season during the subsidy period (recall Figure 1). This means we estimate something closer to an aggregate city-level demand elasticity, which is more inelastic to price changes given a limited ability to substitute across vegetable markets. With both day and market fixed effects, we estimate a market-level elasticity, which is more elastic given the possibility of demand shifting from one market to another. The latter is the most relevant for our exercise, because it is the relevant demand elasticity for considering why a given market did not need to lower prices (much) in response to a market-level supply shock.

To compare our demand elasticity estimates to those based on our experimental variation, we re-estimate Specification 2, this time including both pre-subsidy and subsidy period data, as well as the data from both treatment and control markets, and we instrument for $ln(Quantity)_{mt}$ using our experimental variation. Specifically, our first stage is

$$ln(Quantity)_{mt} = \alpha + \beta Treat_m \times During_t + \gamma_1 Treat_m + \gamma_2 During_t + \varepsilon_{mt} \tag{3}$$

where the unit of observation is the market $\times$ day, and the variables are defined as above. In this case, our instrument is $Treat_m \times During_t$, and again, we estimate specifications with and without market and day fixed effects. Estimates are presented in Table A18, confirming the strength of our instrument, with the F-statistic ranging from 42 to 71.

Our instrumented inverse elasticity estimates are presented in Table A19. Here we augment Specification 2 with controls for $Treat_m$ and $During_t$. Our preferred estimates include market and day fixed effects, and indicate an inverse demand elasticity for peas of -.02, and for carrots of -.02, though neither is statistically significant (columns 3 and 6). The 95% confidence interval for the inverse demand elasticity for peas ranges from -.10 to .03 and for carrots from -.07 to .04, both of which contain the point estimates of the inverse demand elasticity from our preferred specification of the observational approach (Table A16). Taking the lower end of the confidence intervals, our estimates of the impact of our subsidy intervention on market prices are consistent with demand elasticities of -10 and -14. While these are quite elastic, they are in line with some other estimates of demand elasticities in markets for relatively undifferentiated goods; for instance, Ellison and Ellison (2009) finds that the demand elasticity for memory modules is as high as -33.

What accounts for the high elasticity we observe? One possibility is that even slight reductions in prices by vendors in our study allowed vendors to capture demand from larger grocery stores and other markets not included in our study. Our analysis above indicates that the demand for peas and carrots in treatment markets did not come from control markets within our study, but this does not rule out that treatment vendors captured demand from non-study markets and grocery stores. Indeed, as discussed above, most markets in which customers said they would shop, if not in our treatment market, were not part of our study at all.

A high demand elasticity poses another puzzle. Pre-subsidy markups are nontrivial: at 46% on average for carrots, and 34% for peas.[12] Why had vendors not already lowered their prices and captured greater demand? We suspect that the explanation lies in the same reasons that prevented vendors from stocking more carrots and peas prior to the intervention – a mixture of high effort costs of expansion and possible social sanctions from defying a collusive norm, in this case by cutting prices. We discuss these explanations in more detail in Section 5. An alternative explanation is that vendors use heuristics, rather than optimization, to determine pricing. For example, vendors might price a vegetable at the procurement cost plus a 25%

[12] Using the follow-up survey of vendors in 2023, the average daily profit margin is smaller, but still meaningful, at 19%. This number comes from dividing a vendor's self-reported typical daily profit by typical daily revenue.

profit margin, rounded to the nearest five rupees. Consistent with this, though also consistent with a collusive norm, vendors tend to set the same markup for a given vegetable on a given market-day (Figure A6).

A third and final explanation for the non-impact of the subsidy on prices is that treated vendors effectively *induced* additional demand by exerting more selling effort (e.g. by convincing customers to buy more than they had planned to), or by staying longer at the retail market, allowing them to reach customers that they would normally miss. While these forces might explain some of the greater sales in the absence of price cuts, we do not find it plausible that they could fully explain the puzzle. However, we lack measures of selling effort and hours present at the retail market, and so we are unable to test directly for effects on these behaviors.

4.2 Why Did Supply Increase for Non-Subsidized Vegetables?

Our analysis in Table 2 suggests that the effect of our subsidy on vendor profits exceeded the effect on profits from peas and carrots alone. Indeed, Table A20 estimates Specification 1 using as an outcome variable the number of vegetables a vendor stocks other than peas and carrots. The treatment effect is roughly one additional vegetable stocked, though it is marginally insignificant ($p = 0.13$ with the wild bootstrap and $p = 0.16$ with the permutation test).

While this result is not dispositive, it seems likely that the subsidy induced additional crowding-in of other vegetables beyond peas and carrots. Such an effect was not obvious ex ante – in particular, we might have expected other vegetables to be crowded out if vendors have limited shelf space to display products. We investigate two explanations. First, perhaps the pea and carrot subsidy induced vendors to take more trips to the wholesale market, and, once there, vendors restocked on additional vegetables. Table A21 investigates this possibility, re-estimating Specification 1 using as an outcome variable an indicator for whether the vendor visited a wholesale market on a given day. The estimate is close to zero and not statistically significant, likely due to a ceiling effect – vendors were already visiting the wholesale market almost every day they sold produce.

A second possibility is that some vegetables are complementary with peas and carrots, such that vendors commonly sell them alongside one another, and that the subsidy induced additional sales of these complementary vegetables. We find some support for this hypothesis. To identify the vegetables that are complementary with peas and carrots, we use the specification

$$y_{it} = \alpha_m + \alpha_t + \sum_j \beta_j Veg_{it}^j + \varepsilon_{it} \tag{4}$$

where the unit of observation is the vendor $\times$ day, the sample is all observations from the pre-subsidy period, and α_m and α_t are market and day fixed effects. y_{it} is an indicator taking a value of 1 if vendor i sells the focal vegetable (either carrots or peas). Veg_{it}^j is an indicator taking a value of 1 if vendor i sells vegetable j, where the sum incorporates each vegetable in our data other than the focal vegetable. Thus, β_j is the association between the vendor's decision to stock vegetable j and the decision to stock the focal vegetable, and we interpret higher levels of β_j as reflecting stronger complementarity of the two vegetables.

The results are presented in Table A22. Interestingly, the strongest complement for carrots is peas, and carrots are the second strongest complement for peas. This may partly explain why we found in Table 3 that vendors who received the carrot subsidy but not the pea subsidy (i.e. vendors in treatment markets who were ineligible for the pea subsidy) nevertheless stocked additional peas during the subsidy period.

Next, we estimate our subsidy's impact on the likelihood that vendors stock each vegetable other than peas and carrots, by re-estimating Specification 1 once for each vegetable other than peas and carrots. Figure A5 plots the relationship between the level of complementarity for a given vegetable – drawn from Table A22 – and the subsidy's impact on stocking that vegetable. Specifically, in panel (a), each point corresponds to a vegetable other than peas and carrots. The y-axis corresponds to the value of the vegetable's β_j estimated from Specification 4, averaged over the estimates for both peas and carrots, and the x-axis plots the estimate of the impact of our subsidies on the probability a vendor stocks that vegetable. The sample in this panel is all vendors in the study who were eligible for the pea subsidy, so that the intervention applied to those in the treated markets was the joint subsidy of peas and carrots. In this panel, we see a positive relationship between the level of complementarity of a vegetable and the subsidies' impact on the likelihood that vendors stock them.

Panel (b) presents a similar plot, but restricts the sample to only vendors who were ineligible for the pea subsidy, and hence those in treated markets that only received the carrot subsidy. Correspondingly, the y-axis corresponds to the value of the vegetable's β_j estimated from Specification 4, this time only estimated for carrots. The x-axis is the same as in panel (a). In this case, we do not find evidence of a positive relationship between the level of complementarity of a vegetable and the subsidies' impact on the likelihood they are stocked. In summary, we consider this analysis to be suggestive, but mixed, evidence in support of the theory that the subsidy induced extra sales of complementary vegetables.

5 Why Don't Vendors Exploit Their Opportunity for Increased Scale and Profits?

Thus far, we have established that vendors have an opportunity to substantially increase their profits – on average by more than 60% – yet they tend not to exploit it, either before or after our intervention. We now consider a number of explanations. We organize our analysis along the lines of the explanations provided by the theoretical framework outlined in Section 2.1.

5.1 (Perceived) Profit Maximization Subject to Constraints

Are vendors in fact choosing the scale that maximizes their expected utility from profits, given their constraints? To rule out this possibility, we must demonstrate that selling more peas and carrots is profitable, entails little risk, and is feasible without outside intervention. Indeed, we have already demonstrated that selling additional peas and carrots is highly profitable on average.

In principle, a sufficiently high degree of risk or loss aversion can explain any failure to adopt profitable business practices (e.g. Kremer et al., 2013). But we emphasize that stocking peas and carrots offered a high return with little risk. In treatment markets during the subsidy period, vendors who sold carrots earned positive profits from doing so on 96% of the vendor $\times$ days in which they stocked carrots, and 100% of the vendor $\times$ weeks. 100% of vendors earned positive profits over the full subsidy period. The analogous numbers for peas are 95%, 99%, and 99%.

Given these statistics, even an extremely risk-averse vendor ought to find it worthwhile to stock at least a small amount of peas and carrots, yet we found that the impact of our intervention on the probability a vendor stocked any carrots or peas fell to just 10 and 5 percentage points respectively after the subsidy was removed. Therefore, risk and loss aversion are unlikely to be the driving force for vendors' failure to continue stocking carrots and peas.

We can also rule out the possibility that external constraints bind. Our experimental design ensures that vendors who procured additional peas and carrots during the subsidy period had access to all of the necessary capital, labor, and skill required to do so. Specifically, each day the subsidy was delivered to vendors only after they procured the additional vegetables on their own. Therefore, a lack of any of these factors must not be the explanation for why vendors did not continue to stock peas and carrots after the intervention concluded.

Do vendors know that selling (more) peas and carrots would increase their profits? Our experiment strongly suggests that a lack of knowledge about the profitability of selling peas and carrots is not the inhibiting factor. Vendors in our treatment markets experienced higher profits from selling peas and carrots for three weeks, and despite this, they largely ceased selling the additional products once the subsidy was removed.

While in principle it is possible that vendors did not realize they were earning higher profits than in their unsubsidized counterfactual, we do not believe this is likely. This is a setting in which learning the profitability of selling a new product requires neither complicated accounting nor inference about a counterfactual scenario. So long as the revenues generated by selling peas and carrots exceed the cost of procuring them – a fact that is clearly satisfied in our setting – and so long as vendors have sufficient excess capacity to stock additional produce without removing any of their existing produce, then selling the additional products should result in increased profits. This latter fact is confirmed in Table 2, demonstrating that the sales of peas and carrots complemented, rather than displaced, sales of existing produce.

For these reasons, it is not likely that vendors ceased selling peas and carrots for lack of knowledge that doing so would be profitable.

If these arguments are correct, our theoretical framework in Section 2.1 leaves only one category of explanations: the objective function of vendors deviates substantially from monetary profit maximization.

5.2 Vendors' Objectives Deviate Substantially from Profit Maximization

Stress and effort costs of running a larger business. At the individual level, running a larger and more active business may entail substantially more stress and physical exertion, especially when the owner-operator expands without hiring additional labor. If these factors are sufficiently costly, this could deter vendors from exploiting opportunities to increase their profits. In one view, this might represent an unmeasured component of a vendor's profit function. According to this view, our evidence ought to be interpreted as finding that vendors do not maximize their *monetary profits* or take-home pay. In our preferred view, consistent with a financial accounting definition of profits, the non-monetary costs of a vendor's labor are a factor in the vendor's objective function outside of their profits. However, we emphasize that under either, the cost of the vendor's labor is captured by the ϕ term in our theoretical framework in Section 2.1. Moreover, we argue below that even if a vendor is privately optimizing given their own labor-cost function, our results suggest a broader failure at the market level to optimize profits.

While our primary measure of profits does not include the opportunity cost of a vendor's labor, existing estimates for the value of time for the self-employed cannot easily explain why our vendors do not exploit the profitable opportunity of increasing scale. Agness et al. (2025) suggests a rule of thumb of assigning 60% of the market wage as the opportunity cost of labor. We benchmark the "market wage" in our setting by dividing the average daily earnings of vendors in our treatment markets in the pre-subsidy period – Rs. 342 – by the average number of hours that vendors reported working per day in the descriptive 2023 survey – 7.8.[13] Table A23 then re-estimates the treatment effects of our subsidy on a vendor's total profits under the assumption that the subsidy induced treatment market vendors to work an additional one, two, or three hours per day, and deducts the corresponding cost of labor from their daily profits.[14] The effect of the intervention on profits falls to 54%, 42%, and 29%, as we assume one, two, and three additional hours worked, respectively. These effects on profits remain substantial, although they lose statistical significance with two and three additional hours worked (wild $p = 0.105$ and 0.171).

Importantly, Agness et al. (2025) estimates the value of time in a setting in which participants struggle to find work, such that they had only worked an average of 13 days in the prior three months. In our setting, vendors work much longer hours at baseline. It is possible that the marginal hour for our vendors is much more costly than the 60%-of-market-wage rule-of-thumb estimated in Agness et al. (2025). Our findings may then be consistent with a story in which vendors face a sharply upward-sloping cost of supplying labor at the margin.

One piece of evidence is inconsistent with this story. Namely, as discussed above, Figure A4 demonstrates that many vendors purchased more peas and carrots than the subsidized amount.

[13] An alternative benchmark is the local minimum wage for unskilled work in early-2019, Rs. 308 per day (Government of West Bengal, Office of the Labour Commissioner 2019). Given that Rs. 308 is slightly smaller than Rs. 342, the use of the minimum wage benchmark would slightly strengthen our effects on own-labor cost-adjusted profits.

[14] These would correspond to a 0.42-, 0.85-, and 1.27-standard deviation increase in daily working hours reported in our 2023 survey.

Since marginal units beyond the subsidy cap were unsubsidized, this suggests that, conditional on selling these products at all, vendors often found it profitable to sell additional units even without a subsidy. Under a model with only convex variable effort costs, this would imply that selling at least some additional peas and carrots should also have been privately optimal in the absence of the subsidy. However, this inference need not hold if selling peas or carrots involves fixed costs – for example, the need to engage with an additional wholesaler. In that case, the subsidy may have helped vendors overcome the extensive-margin cost of entry, after which additional unsubsidized units were profitable on the margin.

Our descriptive survey, fielded after our experiment, provides direct evidence for a cost-of-effort story. Vendors report working an average of 50 hours per week. When asked how they would manage any extra work if they were to expand their business, 48% of vendors say that they would arrive at the retail market earlier or stay later, while 30% say that they would work harder while at the market (as opposed to hiring additional labor). Vendors report that this additional exertion would be personally costly. When asked how easy it would be to work one additional hour per day, 69% say that it would be somewhat or very difficult (as opposed to somewhat or very easy). This increases to 79% when asked about working two additional hours per day. Of the vendors answering somewhat or very difficult to either question, 35% reply that it would be difficult due to old age (the mean age of our sample is 51), 24% say that it would be difficult due to health problems, 52% say that they do not have enough energy to work the additional hours, and 52% say that they would need to tend to household duties. Correspondingly, when asked "If the market still has customers shopping when you are planning to leave, and you still have produce to sell, do you stay longer or do you need to go home to tend to home duties regardless of how busy the market is?", 68% of vendors reported that they would need to go home. Consistent with high personal costs of expansion, when we ask vendors about their main hopes for their business over the coming year, without prompting answer categories, 85% reply that they hope to preserve the status quo, rather than grow their business.[15]

[15] Our qualitative evidence of informal arrangements between eligible and ineligible pea subsidy vendors – with

With a well-functioning labor market, vendors might take advantage of profitable expansion by hiring labor. In this setting, however, vendors instead have to bear the cost of expansion using their own labor, which even large profit increases might not justify (though the profits combined with the subsidy may be sufficient). This is a form of separation failure: total labor used in the vendor's business is shaped by the vendor's disutility of labor, rather than solely by the production function and the prevailing market wage (Singh et al. 1986; Benjamin 1992; LaFave and Thomas 2016). The canonical rejection of separation concerns family farms, with household demographic composition predicting total farm labor use (LaFave and Thomas 2016). Such a separation failure can be rationalized by a missing capital market and rationing of the quantity of labor household members can supply to the labor market (Brimble et al. 2026)—both plausible market frictions in poor, rural settings (Breza et al. 2021).

The separation failure in our setting is different. In contrast to family farms, the markets we study are densely populated with dozens of vendors operating side-by-side, selling similar goods, and using similar procurement channels. The relevant margin to exploit this profitable opportunity is not the use of a large, illiquid household asset such as land – which cannot be easily rented to other farmers – but the procurement and sale of additional inventory. The indivisibilities that make separation failures difficult to overcome in rural agriculture should therefore be less binding here.

Thus, even if a given vendor's marginal effort cost is high enough that selling more peas and carrots is not privately attractive, the persistence of the unexploited opportunity remains a puzzle at the market level. In a market with many nearby vendors, spare capacity, and observable demand, one might expect the opportunity to be exploited through hiring or entry. A vendor with especially high returns could hire a worker for part of the day, or a new entrant could operate an additional stall. The fact that these adjustments do not occur suggests that the relevant constraint is not merely that owner-operators have high marginal effort costs; there must also be a failure

the latter purchasing extra peas to then sell to the former – can also be seen as indirect evidence that the additional effort costs required to expand procurement are non-trivial. These costs might include the effort required to identify and buy from a new seller in the wholesale market.

deterring the exploitation of these profitable opportunities by workers with lower effort costs. We next turn to one such potential failure.

Inefficient collusion. At the group level, there may be norms or implicit or explicit agreements that discourage vendors from selling the same produce as their nearby competitors, or more generally, from increasing their scale. To the extent that such norms or agreements exist, they diverge significantly from "classical" collusion, in which vendors within a market jointly suppress sales volume to maximize their collective profits (e.g. Tirole, 1988) – because, when we induced some vendors to stock additional peas and carrots, profits went up for the entire market on average (column 3, Table 2). One caveat is that if norms or collusive agreements extend across several markets, they may indeed be profit-maximizing at the group level, as our results do not rule out the possibility that inducing vendors in treatment markets to expand the sales of peas and carrots harmed the profits of vendors in other (non-study) markets so much that aggregate profits declined.

If implicit or explicit collusion is driving our results, it must be sufficiently flexible to allow for the contemporaneous effects of our subsidies. Several factors may have made our intervention well-suited to relax collusive norms. Within treatment markets, it was public knowledge that our research team was providing subsidies for expansion, it was known that the subsidies were temporary, the subsidies were large – intended to cover the full cost of procurement – and the carrot subsidy was extended to all vendors within treatment markets.

To explore the role of collusive norms, we asked a series of questions about the acceptability of various expansionary business practices in our 2023 follow-up survey. We find some suggestive evidence of collusive norms, but not with the pervasiveness found in other contexts (e.g. Breza et al. 2025). Eighteen percent of vendors report that it would be unacceptable or highly unacceptable (as opposed to acceptable or highly acceptable) for a vendor who had never sold carrots or peas before to begin selling carrots or peas, and 27% say that negative consequences of such behavior would be somewhat or highly likely. Among the vendors who reported a negative consequence, the most common are that other vendors would be angry (30%), that vendors would spread information about the behavior to other vendors/markets (29%), that other vendors would prevent them from working at the market (29%), and that other vendors would steer customers away from them (14%).[16]

When we ask about a vendor who "worked hard to expand their business, stocking more produce, and almost doubling their business in size over a few months," 14% describe the behavior as unacceptable or highly unacceptable, and a far larger 45% say that negative consequences would be somewhat or highly likely. These numbers are similar when we ask about price cuts, a behavior that our intervention did not induce. When asked how a vendor would be perceived if they were to sell a vegetable at 10% below the market price, 30% of vendors described this behavior as unacceptable or highly unacceptable, and 40% said that negative consequences would be somewhat or highly likely.[17][18]

Of the three behaviors, 32% of vendors report that at least one is unacceptable, and 54% of vendors report that at least one would lead to negative consequences. While fairly widely held, there is a limit to the extent to which these norms can explain our results. First, nearly half of vendors report not perceiving any such norms. In principle, these vendors should be willing to expand without fear of social consequences. Second, while roughly one-third of vendors always or never sell carrots or peas in the pre-subsidy period, a substantial number of vendors move in and out of selling carrots and peas (Figure A7). This might signal that norms preventing entry are not strictly enforced. Relatedly, roughly half of vendors were already selling carrots or peas

[16]Far fewer report social sanctions of stopping talking (7%) or stopping lending them money (4%), or stopping drinking/playing cards/visiting their home (3%).

[17]For each of the three behaviors we asked about, more vendors said that negative consequences would likely follow than said that the behavior was unacceptable. This might reflect that on average, vendors overestimate how many other vendors would disapprove of these behaviors (Bursztyn et al. 2020), though we did not ask about this directly. It may also be that a small minority that deems a behavior unacceptable may nevertheless go out of their way to sanction others, ensuring that negative consequences are common.

[18]Consistent with the presence of collusive norms, Figure A6 demonstrates a striking consistency in pricing choices across vendors. Specifically, within any market on a given day, holding fixed procurement costs, vendors choose the same markup for peas in more than 89% of cases, and the same markup for carrots in more than 84% of cases. Of course, this could also be consistent with individual optimization, where similar market conditions lead vendors to choose the same prices.

on a typical pre-subsidy day, ruling out a scenario in which a handful of sellers carefully protect their market share for a niche product. Third, while product entry is plausibly observable, it seems less plausible that norm enforcers could easily monitor expansion conditional on entry – for example, another vendor increasing their procurement of carrots by 20%. As we noted earlier, this might explain the evidence in Table 1 of more persistent effects on the intensive margin of procurement than on the extensive margin. It may be that norms concerning product entry are more strictly enforced than those concerning expansion conditional on entry – vendors who experienced profits from adding carrots and peas to their product bundle could not continue doing so, due to expected sanctions, but those that experienced profits from increasing their existing stock of carrots and peas were able to continue doing so (or at least, some were).

If collusive norms constrain profits, a separate question remains: why do such inefficient norms persist, and why were they not revised after our intervention increased aggregate profits? One possibility is that treated vendors noticed that their own profits increased, but did not realize that the same was true of aggregate profits. They might then wrongly believe that attempts to expand their business would harm others in the market. An alternative possibility is that vendors do in fact have complete information, and collusive norms exist not to maximize profits, but to ensure somewhat equitable sharing of total market demand. In this case, the norm might intentionally not maximize joint profits, sacrificing some market-level profit in the service of equity. In line with this, we do see that the intervention increases the inequality in profits across vendors within a market, by substantially increasing the profits of right-tail vendors, while having little effect on the profits of left-tail vendors (Table A24).

It is well-documented that collusive norms like these exist in informal and agricultural markets (e.g. Bergquist and Dinerstein, 2020; Breza et al., 2025; Breza and Kaur, 2026). While these norms are conventionally assumed to support higher profits at the market level, the novelty of our results is that they demonstrate that these norms may in fact strongly conflict with profit-maximization, both at the individual and market levels.

Identity and kinship taxation. We do not see evidence for two alternative explanations related to departures from profit maximization. First, we might hypothesize that vendors stopped selling the additional produce because their identity as a vendor is tied to the types of produce they procure and sell (Akerlof and Kranton, 2000; Oh, 2023). However, even preexisting carrot sellers tend to return to stocking their unsubsidized levels after the subsidy was removed (Appendix Table A9). Presumably selling additional carrots would not challenge their identity given that they already stocked carrots, and therefore identity alone cannot explain why vendors ceased selling the additional produce.

Second, it could be that the vendor's own return on effort is diluted through kinship taxation – the social pressure to share profits with family and friends (Jakiela and Ozier 2016; Squires 2024; Carranza et al. 2025). Under this explanation, although it would be profitable to expand their business, vendors choose not to because they do not internalize the benefits of these profits. To measure kinship taxation pressures in our descriptive survey of 391 vendors, we ask the unincentivized income-hiding question introduced in Squires (2024): "Imagine that I offer you INR 500 today. Now what if I gave you the choice of not telling your extended family and friends that I gave you money. Then they would not know that you received any money from me. If you could choose either, (1) I give you INR 500 and I announce this to your extended family and friends, or (2) I give you INR 400 and do not tell your extended family and friends, which would you prefer?" 19% of the Kenyan microenterprises studied in Squires (2024) selected option (2) – preferring to receive less money, if it means being able to conceal the money from family and friends. In contrast, only 5% of Kolkata vegetable vendors prefer the hidden income. Kinship taxation is therefore unlikely to explain vendors' unwillingness to remain at a larger scale following the removal of our subsidies.

In conclusion, our experiment demonstrates that vendors' objectives deviate from profit maximization with respect to take-home pay, and we provide descriptive evidence that this may be partly due to high marginal costs of effort and perceived social sanctions in response to product entry and expansion.

6 Discussion and Conclusion

We conducted an experiment demonstrating that vegetable vendors in Kolkata could stock additional peas and carrots and increase their profits by over 60%. No external constraint prevented vendors from exploiting this opportunity. Even after we subsidized vendors to stock additional produce, allowing them to directly verify how much doing so would increase their profits, many vendors reduced or stopped the sale of this produce altogether after the subsidy was removed.

Our results are best explained by the fact that vendors do not maximize their take-home pay, at the individual or group level. At the individual-level, this may be attributable to the additional stress or effort required to sustain a larger business (a separation failure), though we have argued that in this case an additional market failure must be present to preclude a worker with lower effort costs from exploiting the profitable opportunity. At the group-level, our results may be attributable to norms that discourage vendors from maintaining an inventory that is too similar to that of their close neighbors. In this case, we emphasize that while these norms may facilitate a form of collusion, it is one that induces vendors to leave significant profits on the table, even as a group. We provide descriptive evidence consistent with both the individual preferences and group norms that induce important deviations from profit-maximizing behavior.

These results have important implications for development research and policy. Development economists have long noted the ubiquity of small firms operating side-by-side in densely packed urban markets. Lewis (1954) observes that petty retail trading in developing countries is dominated by crowded markets and traders making only a few sales each. According to Lewis, consumers would be no worse off if many of these traders left the market, leaving others to expand. The non-pecuniary costs of business expansion that we identify may be a principal explanation for why many markets can sustain so many microenterprises, seemingly operating significantly below capacity, without "natural forces" inducing the most efficient of them to grow and the least efficient to exit. It may be that the modal microentrepreneur's objectives are sufficiently misaligned with monetary profit-maximization that the prospect of growing their business and competing their neighbors out of the market may not be worth the physical toll or the social risk that growth would require. Indeed, our study takes place precisely in such a market. Our intervention induced vendors to compete more aggressively with their neighbors, it induced vendors to earn significantly higher profits (without even imposing any apparent cost on their neighbors), and after our subsidy ended, vendors revealed that maintaining this expansion was unattractive. When asked about their main hopes for their business over the coming year, the vast majority of vendors replied that they hoped to preserve the status quo. Future research should be wary of presuming that microentrepreneurs maximize profits.

References

Agness, Daniel, Travis Baseler, Sylvain Chassang, Pascaline Dupas, and Erik Snowberg, "Valuing the time of the self-employed," *Review of Economic Studies*, 2025, *92* (6), 3471–3503.

Akerlof, George A and Rachel E Kranton, "Economics and Identity," *The Quarterly Journal of Economics*, 2000, *115* (3), 715–753.

Athey, Susan, Julie Tibshirani, and Stefan Wager, "Generalized Random Forests," *The Annals of Statistics*, 2019, *47* (2), 1148–1178.

Atkin, David, Azam Chaudhry, Shamyla Chaudry, Amit K Khandelwal, and Eric Verhoogen, "Organizational barriers to technology adoption: Evidence from soccer-ball producers in Pakistan," *The Quarterly Journal of Economics*, 2017, *132* (3), 1101–1164.

Beaman, Lori, Jeremy Magruder, and Jonathan Robinson, "Minding small change among small firms in Kenya," *Journal of Development Economics*, 2014, *108*, 69–86.

Benjamin, Dwayne, "Household composition, labor markets, and labor demand: testing for separation in agricultural household models," *Econometrica: Journal of the Econometric Society*, 1992, pp. 287–322.

Bergquist, Lauren Falcao and Michael Dinerstein, "Competition and entry in agricultural markets: Experimental evidence from Kenya," *American Economic Review*, 2020, *110* (12), 3705–47.

Breza, Emily and Supreet Kaur, "Psychology and Development: Applications from Cognitive and Social Psychology," Technical Report, National Bureau of Economic Research 2026.

_ , _ , and Nandita Krishnaswamy, "Social norms as a determinant of aggregate labor supply," Technical Report, National Bureau of Economic Research 2025.

— , — , and Yogita Shamdasani, "Labor rationing," *American Economic Review*, 2021, *111* (10), 3184–3224.

Brimble, Paul, Robert Garlick, and Kate Orkin, "How Capital Can Increase the Efficiency of Production," 2026. Working paper, May 2026 version.

Brown, Caitlin and Denni Tommasi, "Quality Upgrading in the Street Food Market: Is Better Equipment and Training Sufficient?," IZA Discussion Paper 18328, IZA Institute of Labor Economics, Bonn, Germany 2025.

Bursztyn, Leonardo, Alessandra L González, and David Yanagizawa-Drott, "Misperceived social norms: Women working outside the home in Saudi Arabia," *American Economic Review*, 2020, *110* (10), 2997–3029.

Cai, Jing and Adam Szeidl, "Indirect effects of access to finance," *American Economic Review*, 2024, *114* (8), 2308–2351.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller, "Bootstrap-based improvements for inference with clustered errors," *The Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Carranza, Eliana, Aletheia Donald, Florian Grosset-Touba, and Supreet Kaur, "The social tax: Redistributive pressure and labor supply," *Econometrica*, 2025, *93* (6), 2273–2308.

Casaburi, Lorenzo and Tristan Reed, "Using individual-level randomized treatment to learn about market structure," *American Economic Journal: Applied Economics*, 2022.

Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins, "Double/Debiased Machine Learning for Treatment and Structural Parameters," *The Econometrics Journal*, 2018, *21* (1), C1–C68.

Cho, Sungjin and John Rust, "The flat rental puzzle," *The Review of Economic Studies*, 2010, *77* (2), 560–594.

De Mel, Suresh, 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.

_ , _ , and _ , "Labor drops: Experimental evidence on the return to additional labor in microenterprises," *American Economic Journal: Applied Economics*, 2019, *11* (1), 202–235.

DellaVigna, Stefano and Matthew Gentzkow, "Uniform pricing in US retail chains," *The Quarterly Journal of Economics*, 2019, *134* (4), 2011–2084.

Drexler, Alejandro, Greg Fischer, and Antoinette Schoar, "Keeping it simple: Financial literacy and rules of thumb," *American Economic Journal: Applied Economics*, 2014, *6* (2), 1–31.

Ellison, Glenn and Sara Fisher Ellison, "Search, obfuscation, and price elasticities on the internet," *Econometrica*, 2009, *77* (2), 427–452.

Fisher, Ronald Aylmer, "Design of experiments," *British Medical Journal*, 1936, *1* (3923), 554.

Gertler, Paul, Sean Higgins, Ulrike Malmendier, and Waldo Ojeda, "Do Behavioral Frictions Prevent Firms from Adopting Profitable Opportunities?," Working Paper 33387, National Bureau of Economic Research 2025.

Government of West Bengal, Office of the Labour Commissioner, "Circular No. 35/703/Stat/2RW/29/2016/LCS/JLC: Minimum Rates of Wages for Establishments under the Shops & Establishments Act, 1963," 2019. Effective 1 January 2019 to 30 June 2019.

Hanna, Rema, 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.

Hardy, Morgan and Jamie McCasland, "Are small firms labor constrained? experimental evidence from Ghana," *American Economic Journal: Applied Economics*, 2023, *15* (2), 253–284.

Hsieh, Chang-Tai and Benjamin A Olken, "The missing "missing middle"," *Journal of Economic Perspectives*, 2014, *28* (3), 89–108.

Jakiela, Pamela and Owen Ozier, "Does Africa need a rotten kin theorem? Experimental evidence from village economies," *The Review of Economic Studies*, 2016, *83* (1), 231–268.

Kremer, Michael, Jean Lee, Jonathan Robinson, and Olga Rostapshova, "Behavioral biases and firm behavior: Evidence from Kenyan retail shops," *American Economic Review*, 2013, *103* (3), 362–68.

LaFave, Daniel and Duncan Thomas, "Farms, families, and markets: New evidence on completeness of markets in agricultural settings," *Econometrica*, 2016, *84* (5), 1917–1960.

Lewis, W. Arthur, "Economic Development with Unlimited Supplies of Labour," *The Manchester School*, 1954, *22* (2), 139–191.

McKenzie, David and Christopher Woodruff, "Experimental evidence on returns to capital and access to finance in Mexico," *The World Bank Economic Review*, 2008, *22* (3), 457–482.

_ **and Susana Puerto**, "Growing markets through business training for female entrepreneurs: A market-level randomized experiment in Kenya," *American Economic Journal: Applied Economics*, 2021, *13* (2), 297–332.

Oh, Suanna, "Does identity affect labor supply?," *American Economic Review*, 2023, *113* (8), 2055–2083.

Roodman, David, Morten Ørregaard Nielsen, James G MacKinnon, and Matthew D Webb, "Fast and wild: Bootstrap inference in Stata using boottest," *The Stata Journal*, 2019, *19* (1), 4–60.

Singh, Inderjit, Lyn Squire, and John Strauss, *Agricultural Household Models: Extensions, Applications, and Policy* World Bank Research Publication, Baltimore, MD: Johns Hopkins University Press for the World Bank, 1986.

Squires, Munir, "Kinship taxation as an impediment to growth: Experimental evidence from Kenyan microenterprises," *The Economic Journal*, 2024, *134* (662), 2558–2579.

Tirole, Jean, *The theory of industrial organization*, MIT press, 1988.

Walker, Michael W, Nachiket Shah, Edward Miguel, Dennis Egger, Felix Samy Soliman, and Tilman Graff, "Slack and economic development," Technical Report, National Bureau of Economic Research 2024.

Young, Alwyn, "Channeling Fisher: Randomization tests and the statistical insignificance of seemingly significant experimental results," *The Quarterly Journal of Economics*, 2019, *134* (2), 557–598.

A Descriptive Survey Details

The table below summarizes the variables used from our 2023 descriptive survey of 391 vendors. The table is followed by descriptions of the survey questions underlying each variable.

	N	Mean	SD	Min	Median	Max
Male	391	.78	.41	0	1	1
Age of vendor (in years)	391	51	12.67	19	52	82
Years selling in this market	391	25	14.46	.25	25	60
Typical daily revenue (INR)	391	3,624	3248.16	100	3,000	40,000
Typical daily profit (INR)	391	578	445.41	25	500	5,000
Typical weekly hours	391	50	18.42	0	49	119
Typical minutes waiting per hour	391	24	8.25	1	20	60
Monthly market cost (INR)	391	6,459	6856.8	0	5,000	91,245
Transport cost per trip (INR)	391	198	535.29	0	150	10,000
Transport cost for double procurement (INR)	391	314	800.2	0	200	15,000
To expand: arrive earlier	391	.36	.48	0	0	1
To expand: stay later	391	.38	.49	0	0	1
To expand: extra visits	391	.013	.11	0	0	1
To expand: work harder per hour	391	.3	.46	0	0	1
To expand: arrange for childcare	391	0	0	0	0	0
To expand: hire people	391	.069	.25	0	0	1
To expand: other	391	.23	.42	0	0	1
How easy to work extra hour	391	1.8	1	0	2	3
How easy to work two extra hours	391	2.1	1.03	0	2	3
Why difficult: old age	311	.35	.48	0	0	1
Why difficult: health problem	311	.24	.43	0	0	1
Why difficult: not enough energy	311	.52	.5	0	1	1
Why difficult: household duties	311	.52	.5	0	1	1
Why difficult: other	311	.11	.31	0	0	1
Hopes: status quo	391	.85	.36	0	1	1
Hopes: want to expand	391	.34	.47	0	0	1
Hopes: want to contract	391	.12	.32	0	0	1
Hopes: want to move	391	.018	.13	0	0	1
Hopes: want to hire	391	.049	.22	0	0	1
Hopes: want new job	391	.013	.11	0	0	1
Hopes: want to retire	391	.023	.15	0	0	1
Hopes: other	391	.013	.11	0	0	1
New seller: acceptable?	391	1.8	.57	0	2	3
New seller: negative consequences?	391	1.2	.8	0	1	3
New seller reaction: become angry	110	.6	.49	0	1	1

	N	Mean	SD	Min	Median	Max
New seller reaction: spread info	110	.45	.5	0	0	1
New seller reaction: prevent from working	110	.46	.5	0	0	1
New seller reaction: steer customers away	110	.25	.43	0	0	1
New seller reaction: stop talking	110	.1	.3	0	0	1
New seller reaction: stop drinking/cards/visit	110	.036	.19	0	0	1
New seller reaction: stop lending	110	.064	.25	0	0	1
New seller reaction: other	110	.0091	.1	0	0	1
Vendor expand: acceptable?	391	1.9	.46	0	2	3
Vendor expand: negative consequences?	391	1.4	.84	0	1	3
Vendor undercut: acceptable?	391	1.7	.66	0	2	3
Vendor undercut: negative consequences?	391	1.3	.91	0	1	3
Vendor same price: acceptable?	391	1.9	.48	0	2	3
Kinship tax: prefer private money?	391	.046	.21	0	0	1
Agree with: want to expand	391	.48	.5	0	0	1
Prefer wage employment?	391	.16	.36	0	0	1
Prefer wage employment, same pay?	391	.16	.37	0	0	1

- Years selling in the market derives from the survey question “How long have you been selling in this market?”
- For daily revenue and profit, we ask “What is your typical daily revenue? (i.e. how much total money do you receive from your customers on a typical day)” and “What is your typical daily profit? (i.e. your typical daily revenue less your typical daily costs)”
- For typical weekly hours, we first ask vendors which days they typically work on, then we ask them how many hours they typically work on each of these days. We sum up these numbers across the seven days of the week.
- For typical minutes waiting, we ask “For how many minutes of your typical working hour do you spend waiting for customers (e.g. sitting idly, chatting with friends nearby, etc.)?”
- For monthly market cost, we ask “What do you pay to be able to sell in this market? (e.g. rent, payments to market association, to police, etc.)”

- For transport cost per trip, we ask "What is the typical procurement cost of the vegetables you would buy in one trip to the wholesale market?" followed by "How much does it cost to transport this amount of vegetables to the market?" Transport cost per trip is the answer to the latter question.

- For transport cost for double procurement, we ask "If you wanted to buy double the vegetables that you usually buy, how much would the transportation cost then?"

- For the variables beginning "To expand:", we ask "If you wanted to expand your business (by selling 50% more produce), how would you manage the extra work?" The variables in the table are indicator variables for each of the possible answer options: (1) I would arrive at the retail market earlier, (2) I would stay at the retail market later, (3) I would do extra visits to sell at the retail market (e.g. an extra day at the weekend, or an extra evening during the week), (4) I would have to work harder for each hour that I am at the market, to make sure I sell the extra produce, (5) I would arrange for childcare so that I can work when I would usually be looking after my children, (6) I would hire people to work for me, and (7) other.

- For how easy to work extra hour, we ask "How easy would it be for you to work for 1 additional hour each day, if you wanted to? For example, suppose that customers were still shopping at the end of your work day, so that you could still earn your average hourly profit in that extra hour." The answer options are 0 = very easy, 1 = somewhat easy, 2 = somewhat difficult, and 3 = very difficult.

- For how easy to work two extra hours, we ask "How easy would it be for you to work for 2 additional hours each day, if you wanted to? For example, suppose that customers were still shopping at the end of your work day, so that you could still earn your average hourly profit in the two extra hours." The answer options are 0 = very easy, 1 = somewhat easy, 2 = somewhat difficult, and 3 = very difficult.

- For those that respond that it would be somewhat or very difficult to work one or two extra hours, we ask "Why would it be difficult?" The answer options are (1) old age, (2) health problem, (3) not enough energy to work additional hours, (4) need to tend to household duties, and (5) other. Indicator variables for each of these answers are included in the table with names "Why difficult: ..."

- "Hopes:" variables derive from the question "What are your main hopes for your business over the coming year?" We did not read out answer categories to the vendors for this question. Surveyors selected from the following categories: (1) no hopes/status quo, (2) want to expand/sell more, (3) want to contract/sell less, (4) want to move to a different market, (5) want to hire people to work for me, (6) want to stop selling vegetables and shift into another job, (7) want to retire from working, and (8) other.

- For new seller: acceptable?, we asked "Suppose a vendor in this market who had never sold carrots or peas before began to sell carrots or peas. Do you think this behavior is acceptable?" The answer categories are: 0 = highly unacceptable, 1 = unacceptable, 2 = acceptable, and 3 = highly acceptable.

- For new seller: negative consequences?, we asked "Suppose a vendor in this market who had never sold carrots or peas before began to sell carrots or peas. How likely is it that he would face negative consequences from other vendors in this market?" The answer categories are: 0 = highly unlikely, 1 = somewhat unlikely, 2 = somewhat likely, and 3 = highly likely.

- We then ask "Suppose a vendor in this market who had never sold carrots or peas before began to sell carrots or peas. What reactions will he face from the other vendors in the market?" 110 vendors reported at least one reaction. Among these vendors, we include indicator variables in the table above for each type of reaction reported: (1) become angry, (2) spread information to other vendors/markets, (3) prevent from working at market,

(4) steer customers away, (5) social sanctions – stop talking, (6) social sanctions – stop drinking/playing cards/visiting home, (7) stop lending money, and (8) other.

- For vendor expand: acceptable?, we asked "Suppose a vendor in this market worked hard to expand their business, stocking more produce, and almost doubling their business in size over a few months. Do you think that this behavior is acceptable?" The answer categories are: 0 = highly unacceptable, 1 = unacceptable, 2 = acceptable, and 3 = highly acceptable.

- For vendor expand: negative consequences?, we asked "Suppose a vendor in this market worked hard to expand their business, stocking more produce, and almost doubling their business in size over a few months. How likely is it that he would face negative consequences from other vendors in this market?" The answer categories are: 0 = highly unlikely, 1 = somewhat unlikely, 2 = somewhat likely, and 3 = highly likely.

- For vendor undercut, acceptable?, we asked "Suppose a vendor in this market were to sell similar produce to yours but at 10% lower price. Do you think this behavior is acceptable?" The answer categories are: 0 = highly unacceptable, 1 = unacceptable, 2 = acceptable, and 3 = highly acceptable.

- For vendor undercut, negative consequences?, we asked "Suppose a vendor in this market were to sell similar produce to yours but at 10% lower price. How likely is it that he would face negative consequences from other vendors in this market?" The answer categories are: 0 = highly unlikely, 1 = somewhat unlikely, 2 = somewhat likely, and 3 = highly likely.

- For vendor same price: acceptable?, we asked "Suppose a vendor in this market were to sell similar produce to yours at a similar price. Do you think this behavior is acceptable?" The answer categories are: 0 = highly unacceptable, 1 = unacceptable, 2 = acceptable, and 3 = highly acceptable.

- For kinship tax: prefer private money?, we asked "Imagine that I offer you INR 500 today. Now what if I gave you the choice of not telling your extended family and friends that I gave you money. Then they would not know that you received any money from me. If you could choose either, (1) I give you INR 500 and I announce this to your extended family and friends, or (2) I give you INR 400 and do not tell your extended family and friends, which would you prefer?" The variable in the table is an indicator variable equal to one if the vendor selected the second option.

- For agree with: want to expand, we asked "I'm now going to read out a list of options for your hopes for your business. We want to know if you agree or disagree with each option as one of your hopes for the business." One of the options was "want to expand/sell more." The indicator variable in the table is equal to one if the vendor agreed with this option.

- For prefer wage employment?, we asked "Suppose someone offered you wage employment at the prevailing wage, for a task that you are capable of doing. Would you prefer this job to your current job selling vegetables?"

- For prefer wage employment, same pay?, we asked "How about if the job was not at the prevailing wage, but paying the same amount that you earn selling vegetables. Would you prefer this job to selling vegetables?"

B Estimation of Heterogeneous Treatment Effects

We estimate heterogeneous treatment effects across vendors using a machine learning approach based on generalized random forests. Specifically, we estimate conditional average treatment effects (CATEs) of the subsidy as a function of pre-subsidy vendor characteristics, following the framework of Chernozhukov et al. (2018) and Athey et al. (2019).

Let Y_i denote the outcome of interest (change in profits), W_i indicate treatment status, and X_i denote a vector of baseline covariates. Our object of interest is:

$$\tau(X_i) = \mathbb{E}[Y_i(1) - Y_i(0) \mid X_i]. \tag{5}$$

We estimate $\tau(X_i)$ using a flexible, doubly robust procedure that combines nonparametric prediction of outcomes and treatment assignment. In practice, we implement a random forest-based estimator that first estimates conditional outcome functions and propensity scores, and then constructs individual-level treatment effect estimates using a doubly robust transformation. This approach allows for high-dimensional and non-linear relationships between baseline characteristics and treatment effects while remaining robust to misspecification of either component.

All covariates are constructed using only pre-subsidy data and aggregated to the vendor level, ensuring that each observation corresponds to a single vendor. The covariates include measures of product mix (e.g., share of days selling peas or carrots), prices, wholesale quantities, and profits by product, overall business outcomes (total costs, sales, and profits), operating characteristics (e.g., travel time, working hours, shop size), and vendor demographics (gender, age, and experience).

The approach produces a vendor-specific predicted treatment effect $\hat{\tau}_i$ for each of peas and carrots, and for both the during-subsidy and post-subsidy period (and note that for peas, we only use the data from the vendors eligible for the pea subsidy). We then validate these predictions by regressing

$$y_{im} = \alpha + \beta_1 Treat_m + \beta_2 \hat{\tau}_i + \gamma Treat_m \times \hat{\tau}_i + \varepsilon_{im} \tag{6}$$

where y_{im} represents the difference between the average profits of selling the focal vegetable (either peas or carrots) in the during-subsidy (or post-subsidy) and pre-subsidy periods for vendor i in market m, $\hat{\tau}_i$ is the predicted treatment effect on profits from the focal vegetable for vendor i, and as above, $Treat_m$ is an indicator for the market-level subsidy treatment. The results are presented in Tables A3 and A5, and indicate that our model only has predictive power for the post-subsidy treatment effects on pea profits. Hence we focus on the vendor characteristics that predict post-subsidy effects on pea profits in the main text. To do so, we regress predicted post-subsidy pea profit treatment effects $\hat{\tau}_i$ on pre-treatment covariates (Table A4).

C Appendix Figures and Tables

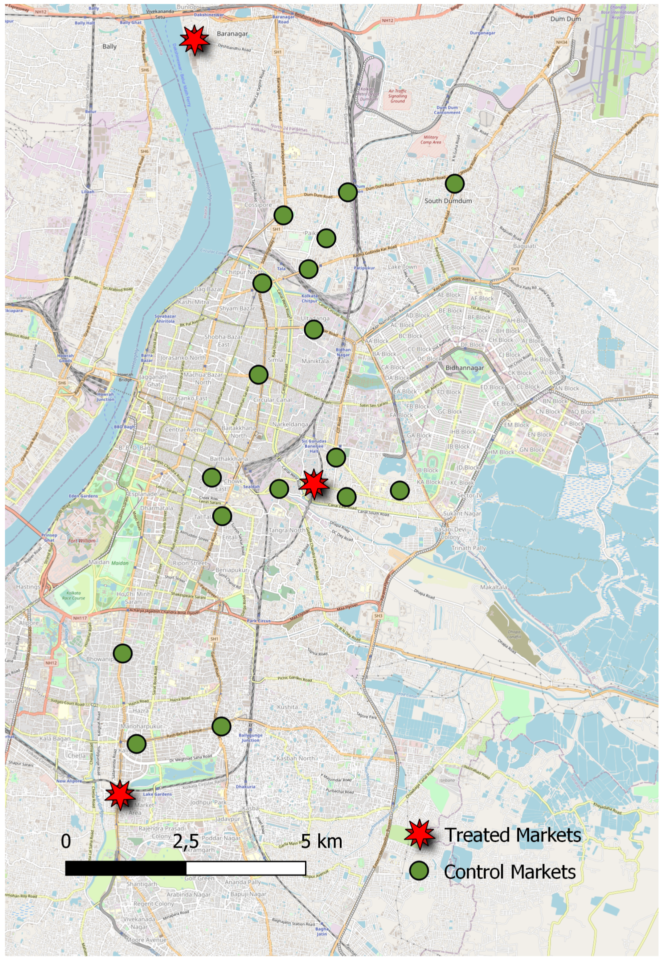

Figure A1: This map shows the location of the treated markets (in red) and control markets (in green) in our Kolkata sample.

Figure A2: Distribution of the Number of Days in the Pre-Subsidy Period Selling Peas and Carrots

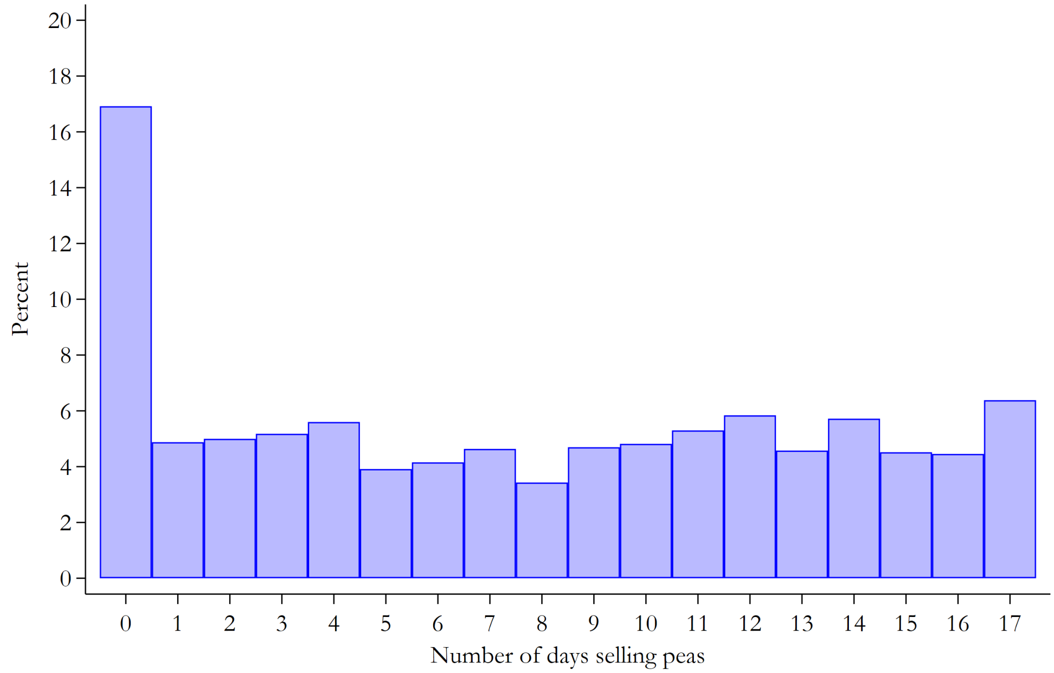

(a) Peas

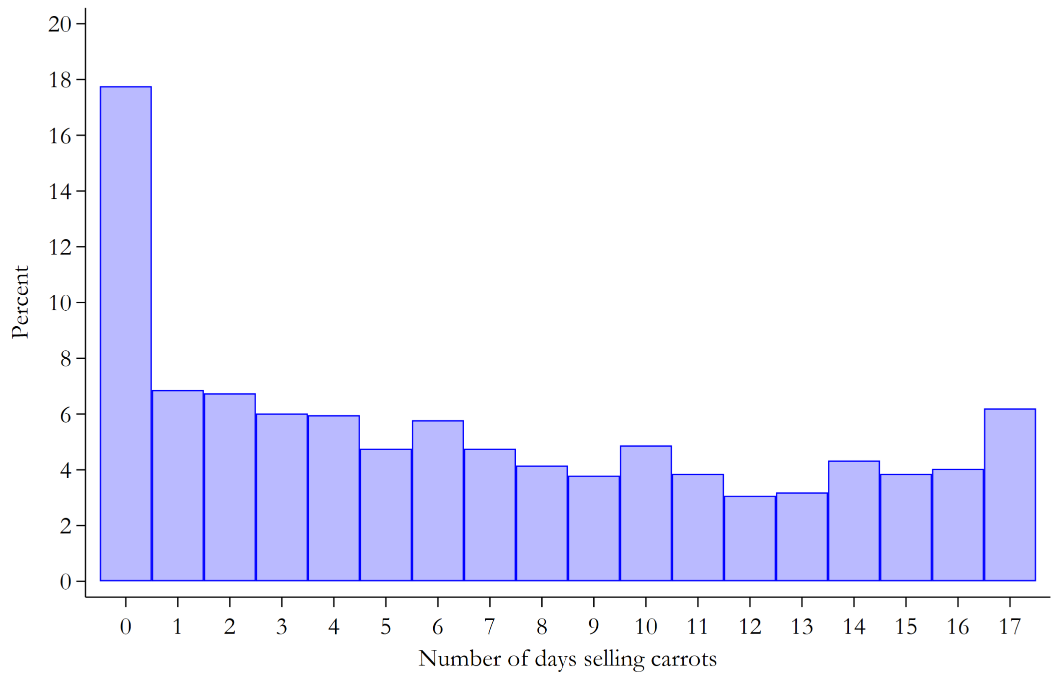

(b) Carrots

Notes: This histogram shows the distribution of the number of days in the pre-subsidy period in which a vendor sold peas and carrots.

Figure A3: Vegetable Markup by Hour

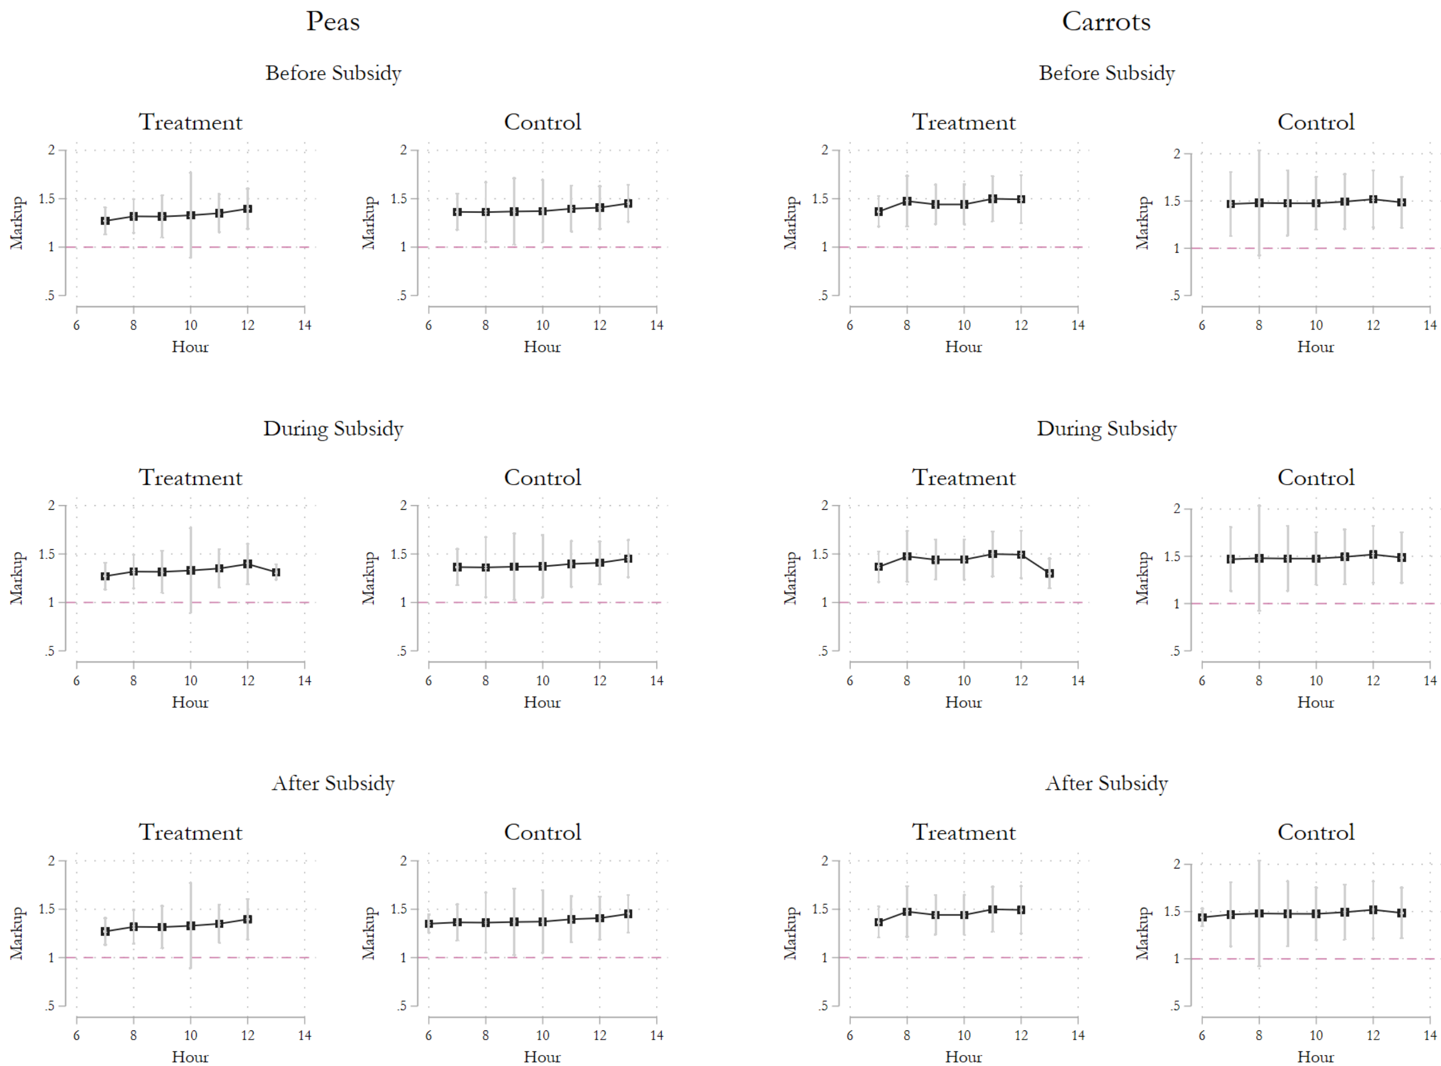

Notes: The figure illustrates the markup of peas and carrots during different survey hours. The markup is calculated as the ratio of retail price to wholesale market cost. The first row displays the markup evolution in the pre-subsidy period, the second row shows the subsidy period, and the third row represents the post-period. Each period is further divided into different plots for both the treatment and control groups.

Figure A4: Stock Volumes Compared to Subsidy Sizes

(a) Peas

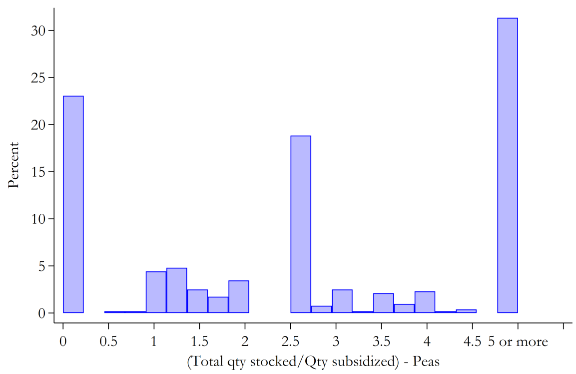

Low Subsidy Amount

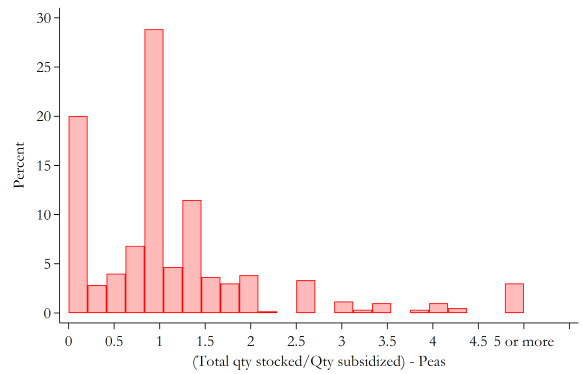

High Subsidy Amount

(b) Carrots

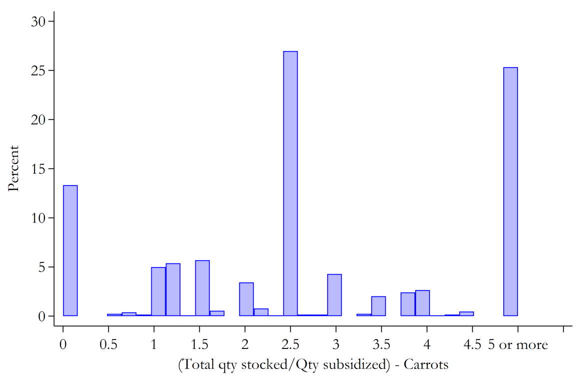

Low Subsidy Amount

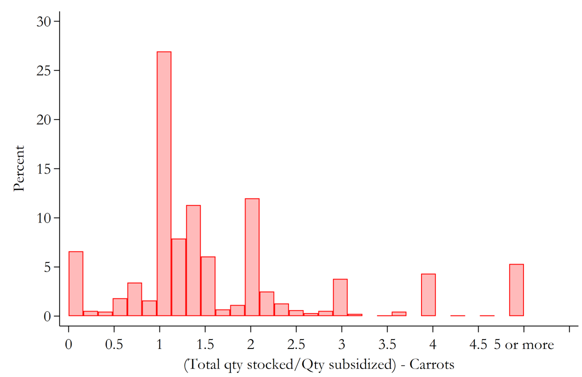

High Subsidy Amount

Notes: The figure illustrates the distribution of stock volumes for peas and carrots compared to the size of the subsidy (in terms of kgs subsidized) a vendor received for a given vegetable. Each observation is a vendor x day in the treatment markets during the subsidy period.

Figure A5: Impact of the Subsidy on Vegetables Stocking Versus Complementarity of the Vegetable with Peas and Carrots

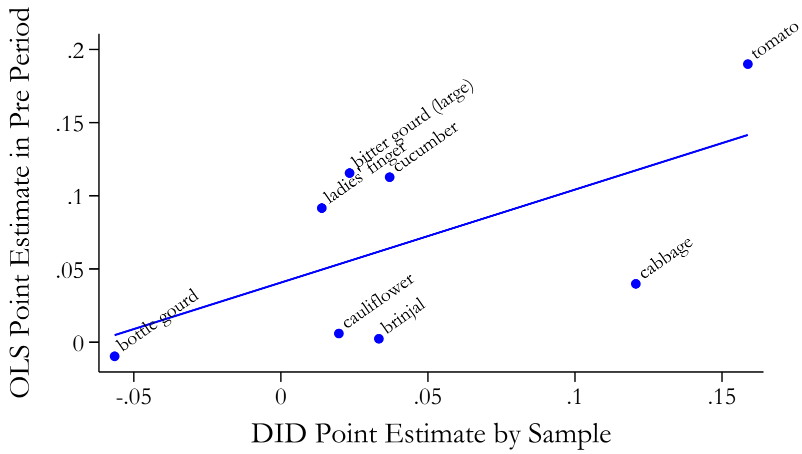

(a) Eligible for Peas Subsidy

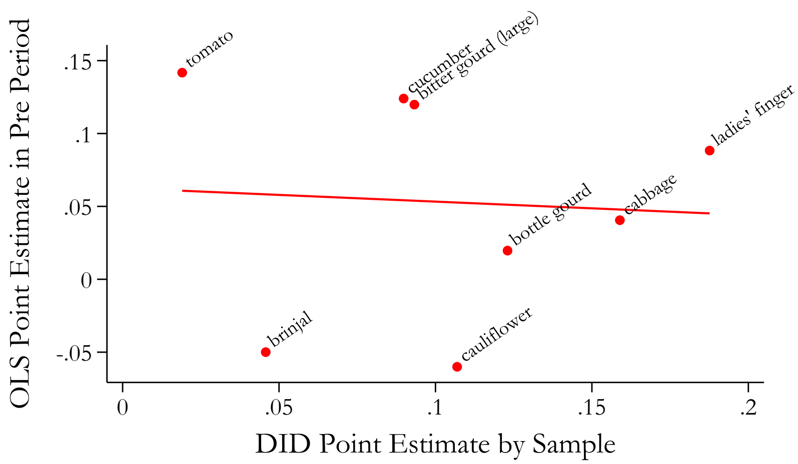

(b) Ineligible for Peas Subsidy

Notes: The figure illustrates the relationship between the impact of the subsidy on vegetable sales and each vegetable's pre-subsidy complementarity with peas and carrots. Dots represent vegetables; the fitted line is a linear best fit. The x-axis plots the DID estimate of the probability of selling each vegetable based on specification (1). The y-axis reports pre-subsidy complementarity coefficients from columns (1) and (2) of Table A22, depending on the panel. Panel A plots the DID point estimates on $During_t \times Treat_m$ from specification (1) for vendors eligible for the pea subsidy, and the average of the coefficients from columns (1) and (2) of Table A22. Panel B plots the DID point estimates on $During_t \times Treat_m$ from specification (1) for vendors ineligible for the pea subsidy, and the coefficients of column (1) of Table A22.

Figure A6: Distribution of Markups for Peas and Carrots

(a) Peas

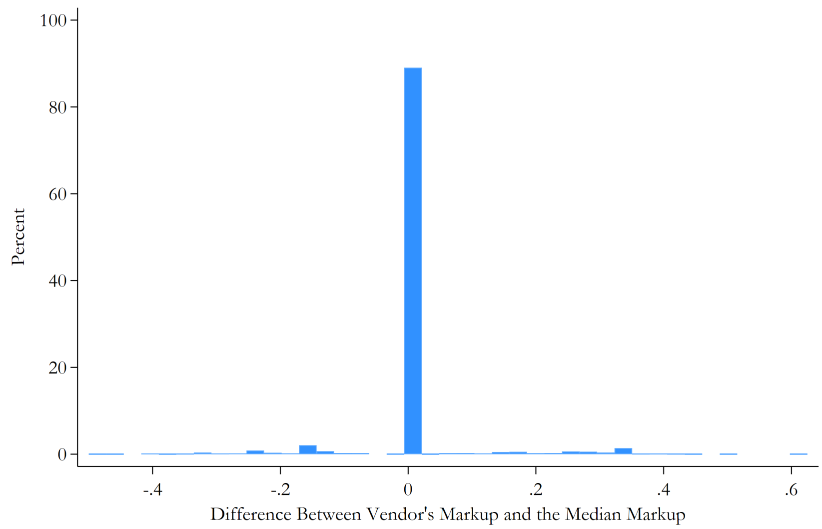

(b) Carrots

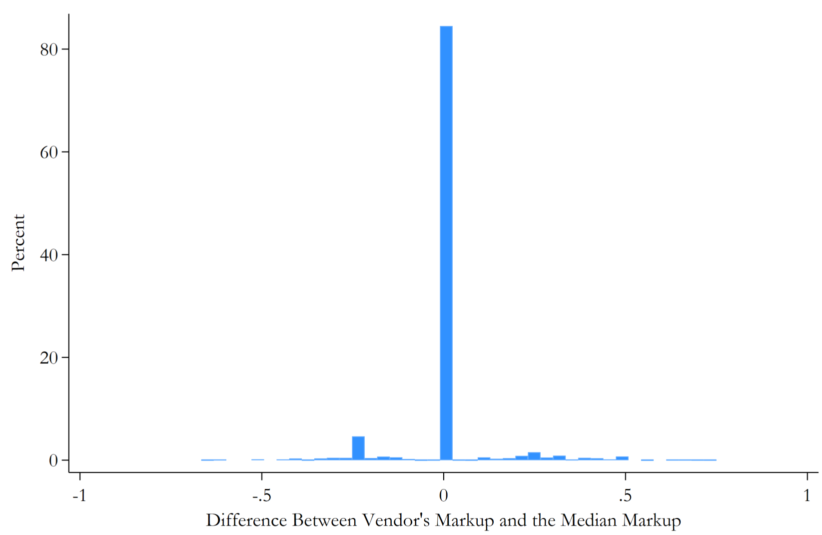

Notes: The figure illustrates within-market variation in markups for peas and carrots. Each observation is a vendor–day. Markups are computed as $(p^{sell} - p^{ws})/p^{ws}$, where p^{sell} denotes the vendor's retail price and p^{ws} the procurement price. The x-axis reports the difference between each vendor's markup and the median markup among vendors operating in the same market, on the same day, and facing the same procurement price. The sample excludes the top and bottom 1 percent of the markup distribution and restricts the analysis to market–day–procurement-price groups with at least five observations.

Figure A7: Percentage of Days in the Pre-Subsidy Period Selling Peas and Carrots

(a) Peas

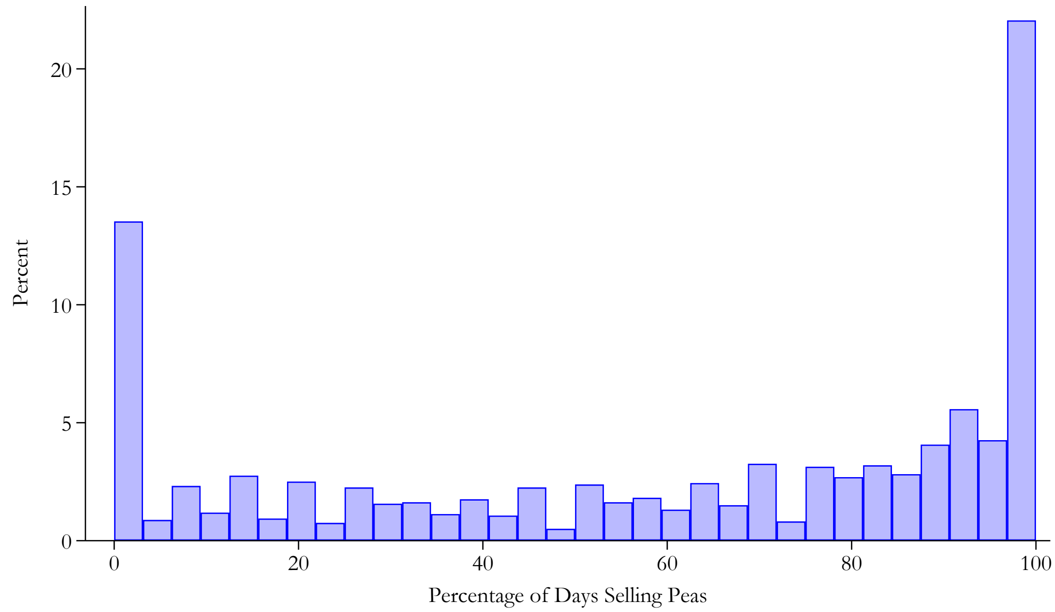

(b) Carrots

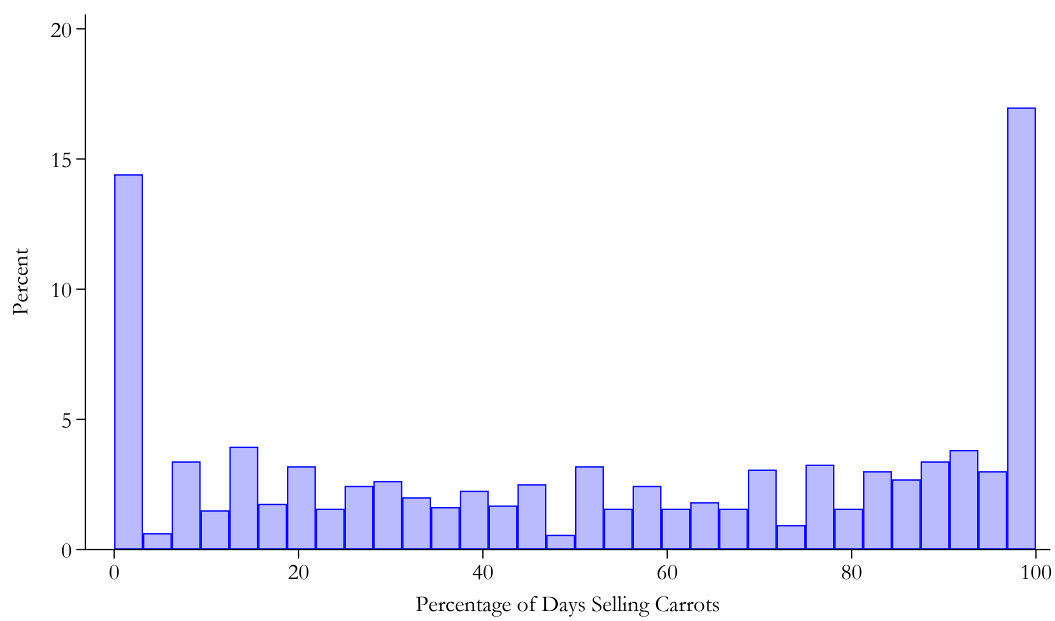

Notes: The figure shows the distribution of the percentage of days during the pre-subsidy period in which vendors sold peas and carrots. For each vendor, this percentage is computed as the number of days the vegetable was available divided by the total number of days the vendor was surveyed during the pre-subsidy period, multiplied by 100. Because vendors did not necessarily respond to the survey every day, the denominator may vary across vendors. Vendors for whom the entire pre-subsidy period is missing are excluded from the sample.

Table A1: Descriptive Statistics for Intervention and Control Markets

	Charu Market (1)	Sarkar Bazar (2)	Alam Bazar (3)	Control markets (4)
Mean # vendors in census	45.0	73.0	85.0	85.8
Mean # vendors present per day	34.5	57.4	71.7	63.1
Mean profits per vendor (Rs.)	448.9	344.9	314.6	531.0
Mean # vegetables available per vendor	5.7	5.1	4.0	5.6
% of present vendors selling peas	62.9	59.9	51.8	64.4
% of present vendors selling carrots	65.1	52.2	39.1	59.4
Mean total cost of daily purchases (Rs.)	969.0	894.4	748.3	1,336.8
Mean value of Sales (Rs.)	1,417.9	1,239.4	1,062.9	1,868.3
Mean number of years selling in market	22.0	22.9	26.6	24.3
Mean age	44.2	48.7	49.1	48.0
% of female vendors	33.3	47.9	12.9	25.1

Notes: This table presents summary statistics averaged for each intervention market in columns 1 - 3 and averaged over all control markets in column 4. All statistics are calculated using data from the pre-subsidy period. For control markets in column 4, statistics are first computed at the market level and then averaged across markets. Profits are calculated by computing (amount of vegetable at the start of the day - amount left over at the end of the day)*anticipated sale price - (amount procured at the start of the day * procurement cost). On days where amount left over was not observed, we impute the vendor's average amount left over all days in which it was measured.

Table A2: Testing for Parallel Trends in the Pre-Subsidy Period

	Carrot				Peas			
	Prob. of selling (%) (1)	Sale price (Rs.) (2)	Wholesale qty. (kg) (3)	Profits (Rs.) (4)	Prob. of selling (%) (5)	Sale price (Rs.) (6)	Wholesale qty. (kg) (7)	Profits (Rs.) (8)
Panel A: Carrot and Peas								
β_3 Treat	–0.05	–3.21	–1.51	–13.06	–0.08	–0.44	–3.92	–32.10
	[–0.30, 0.19]	[–7.66, 0.57]	[–3.39, 0.44]	[–30.74, 7.34]	[–0.35, 0.20]	[–5.10, 5.35]	[–7.89, –0.53]	[–61.64, 1.96]
	{ 0.681 }	{ 0.083 }	{ 0.113 }	{ 0.288 }	{ 0.564 }	{ 0.763 }	{ 0.042 }	{ 0.053 }
	⟨ 0.605 ⟩	⟨ 0.180 ⟩	⟨ 0.218 ⟩	⟨ 0.182 ⟩	⟨ 0.504 ⟩	⟨ 0.789 ⟩	⟨ 0.012 ⟩	⟨ 0.021 ⟩
γ_1 Treat × Day	–0.00	0.08	0.04	0.17	0.00	0.06	0.19	–0.22
	[–0.01, 0.01]	[–0.35, 0.53]	[–0.07, 0.16]	[–1.40, 2.22]	[–0.00, 0.01]	[–0.28, 0.36]	[–0.13, 0.54]	[–3.14, 3.50]
	{ 0.984 }	{ 0.687 }	{ 0.466 }	{ 0.716 }	{ 0.574 }	{ 0.697 }	{ 0.633 }	{ 0.746 }
	⟨ 0.992 ⟩	⟨ 0.625 ⟩	⟨ 0.529 ⟩	⟨ 0.747 ⟩	⟨ 0.446 ⟩	⟨ 0.605 ⟩	⟨ 0.518 ⟩	⟨ 0.838 ⟩
Pre-subsidy intervention market mean	0.49	27.88	3.40	24.95	0.57	41.75	5.71	44.50
Number of Vendors	1591	1361	1591	1591	1591	1373	1591	1591
Number of Observations	20040	10675	20040	20040	20040	12053	20040	20040
	Cost of wholesale purchases (Rs.)	Sales (Rs.)	Profits (Rs.)	# vegetables available				
Panel B: Aggregate								
β_3 Treat	–550.19	–689.20	–139.01	–0.40				
	[–1154.05, 3.17]	[–1345.47, –75.36]	[–322.99, 61.70]	[–2.78, 1.82]				
	{ 0.049 }	{ 0.041 }	{ 0.219 }	{ 0.630 }				
	⟨ 0.069 ⟩	⟨ 0.051 ⟩	⟨ 0.154 ⟩	⟨ 0.645 ⟩				
γ_1 Treat × Day	13.02	12.88	–0.14	–0.01				
	[–9.05, 33.74]	[–14.51, 40.32]	[–10.23, 9.62]	[–0.08, 0.07]				
	{ 0.513 }	{ 0.629 }	{ 0.968 }	{ 0.357 }				
	⟨ 0.367 ⟩	⟨ 0.516 ⟩	⟨ 0.970 ⟩	⟨ 0.454 ⟩				
Pre-subsidy intervention market mean	825.12	1167.30	342.18	4.73				
Number of Vendors	1591	1591	1591	1591				
Number of Observations	20040	20040	20040	20040				

Notes: This table estimates the following specification: $y_{it} = \alpha + \beta_1 Day_t + \beta_2 Treat_m + \beta_3 Day_t \times Treat_m + \varepsilon_{imt}$ on our sample during the pre-subsidy period. Coefficients for Day not shown. 95% wild bootstrap confidence intervals are in [], wild bootstrap p-value is in {}, and Fisher permutation p-value is in ⟨⟩. In Panel A, outcomes are specific to peas or carrots. The outcome in columns 1 and 5 is whether the vendor sells carrots or peas on the given day, the outcome in columns 2 and 6 measure the vendor's anticipated sale price for the relevant vegetable, the outcome in columns 3 and 7 measure the wholesale quantity procured of the relevant vegetable, and the outcome in columns 4 and 8 measure the daily profits accrued from the relevant vegetables. Profits are calculated by computing (amount of vegetable at the start of the day - amount left over at the end of the day)*anticipated sale price - (amount procured at the start of the day * procurement cost). On days where amount left over was not observed, we impute the vendor's average amount left over all days in which it was measured. Our measure of profit does not include the subsidy vendors received as part of our intervention. In Panel B the outcomes correspond to aggregate measures. The outcome in column 1 is the total cost of wholesale purchases on a given day, the outcome in column 2 is the vendor's total revenues on a given day accruing from all produce, the outcome in column 3 is the daily profits accrued from all produce, and the outcome in column 4 is the number of distinct types of vegetables a vendor has available on a given day.

Table A3: Heterogeneous treatment effects in the post-subsidy period by predicted post-period treatment effects

	Peas	**Carrots**
	Profits	Profits
	(1)	(2)
β_1 Treat	0.698	2.987
	(1.378)	(5.115)
β_2 Predicted HTEs	-1.106***	0.531
	(0.073)	(0.623)
γ Treat $\times$ Predicted HTEs	1.625***	1.510
	(0.403)	(0.936)
Observations	543	1529
R-squared	0.681	0.044

Notes: The table reports difference-in-differences estimates from regressions interacting the predicted conditional average treatment effect ($\hat{\tau}$) with the treatment indicator. The predicted $\hat{\tau}$ is estimated using a causal forest (AIPW) with pre-subsidy vendor characteristics. The outcome is the change in profits, defined as the difference between average profits in the post- subsidy period and average profits in the pre-subsidy period. Column 1 corresponds to peas, where the prediction is estimated using only vendors eligible for the pea subsidy, and column 2 corresponds to carrots. All regressions are estimated at the vendor level with standard errors clustered at the market level reported in parentheses.

Table A4: Correlates of predicted treatment effect $\hat{\tau}_i$ for the post-subsidy period

	Peas	**Carrots**
	Predicted HTE ($\hat{\tau}$)	Predicted HTE ($\hat{\tau}$)
	(1)	(2)
Share of days selling vegetable	92.551***	0.175
	(13.261)	(1.529)
Average selling price	0.425	-0.591***
	(0.368)	(0.103)
Average total sales	0.012***	-0.003***
	(0.004)	(0.000)
Average number of vegetable sold	-5.755***	-0.391
	(1.282)	(0.268)
Average travel time to wholesale market	0.010	-0.020***
	(0.040)	(0.006)
Average arrival time at market	0.779	-1.195***
	(1.496)	(0.179)
Average daily working hours	0.132	-0.164*
	(0.197)	(0.085)
Share of days taking a break	3.852	-0.563
	(2.533)	(0.610)
Vendor age	0.044	0.017
	(0.050)	(0.021)
Years selling in the market	0.054	0.074***
	(0.072)	(0.021)
Male	1.974	-0.864
	(2.222)	(0.503)
Observations	562	1661
R-squared	0.439	0.602

Notes: The table presents OLS regressions of the predicted Individual Average Treatment Effect (CATE) on baseline characteristics, where the unit of observation is the vendor. The predicted CATE is defined as the difference in average profits post- subsidy period relative to the pre-subsidy period. The sample includes vendors with non-missing baseline information, and standard errors are clustered at the market level. Column 1 uses the change in profits for peas as the outcome, while Column 2 reports the corresponding specification for carrots.

Table A5: Heterogeneous treatment effects during the subsidy period by predicted subsidy-period treatment effects

	Peas	**Carrots**
	Profits	Profits
	(1)	(2)
β_1 Treat	21.846	26.833
	(44.690)	(20.252)
β_2 Predicted HTEs	-1.026***	1.061***
	(0.109)	(0.198)
γ Treat $\times$ Predicted HTEs	0.555	0.395
	(0.961)	(0.627)
Observations	548	1549
R-squared	0.591	0.242

Notes: The table reports difference-in-differences estimates from regressions interacting the predicted conditional average treatment effect ($\hat{\tau}$) with the treatment indicator. The predicted $\hat{\tau}$ is estimated using a causal forest (AIPW) with pre-subsidy vendor characteristics. The outcome is the change in profits, defined as the difference between average profits during the subsidy period and average profits in the pre-subsidy period. Column 1 corresponds to peas, where the prediction is estimated using only vendors eligible for the pea subsidy, and column 2 corresponds to carrots. All regressions are estimated at the vendor level with standard errors clustered at the market level reported in parentheses.

Table A6: Subsidy Impacts: Carrots and Peas (Alternative Permutation Test)

	Carrot				Peas			
	Prob. of selling (%) (1)	Sale price (Rs.) (2)	Wholesale qty. (kg) (3)	Profits (Rs.) (4)	Prob. of selling (%) (5)	Sale price (Rs.) (6)	Wholesale qty. (kg) (7)	Profits (Rs.) (8)
β_3 Treat	–0.05	–2.62	–1.40	–12.71	–0.04	–0.01	–2.53	–33.29
	⟨ 0.614 ⟩	⟨ 0.118 ⟩	⟨ 0.177 ⟩	⟨ 0.195 ⟩	⟨ 0.627 ⟩	⟨ 0.991 ⟩	⟨ 0.077 ⟩	⟨ 0.059 ⟩
γ_1 Treat × During Subs	0.57	–0.44	5.99	44.75	0.39	–0.81	6.72	59.67
	⟨ 0.005 ⟩	⟨ 0.650 ⟩	⟨ 0.005 ⟩	⟨ 0.009 ⟩	⟨ < 0.001 ⟩	⟨ 0.682 ⟩	⟨ 0.005 ⟩	⟨ 0.005 ⟩
γ_2 Treat × After Subs	0.10	–2.04	1.94	8.79	0.05	–0.87	2.58	26.52
	⟨ 0.327 ⟩	⟨ 0.223 ⟩	⟨ 0.145 ⟩	⟨ 0.205 ⟩	⟨ 0.495 ⟩	⟨ 0.627 ⟩	⟨ 0.045 ⟩	⟨ 0.082 ⟩
Pre-subsidy intervention market mean	0.49	27.91	3.43	25.41	0.57	41.80	5.94	47.73
Fisher p-value: $\gamma_1 = \gamma_2$	0.005	0.395	0.005	<0.001	0.009	0.968	0.005	0.082
Number of Vendors	1631	1470	1631	1631	1631	1489	1631	1631
Number of Observations	55218	25073	55218	55213	55243	22657	55243	55241

Notes: This table estimates specification 1 on our full sample. Coefficients for During and Post not shown. Fisher permutation p-value is in ⟨⟩. Columns 1 - 4 present outcomes for carrots, and 5 - 8 for peas. The outcome in columns 1 and 5 is whether the vendor sells carrots or peas on the given day, the outcome in columns 2 and 6 measure the vendor's anticipated sale price for the relevant vegetable, the outcome in columns 3 and 7 measure the wholesale quantity procured of the relevant vegetable, and the outcome in columns 4 and 8 measure the daily profits accrued from the relevant vegetables. Profits are calculated by computing (amount of vegetable at the start of the day - amount left over at the end of the day)*anticipated sale price - (amount procured at the start of the day * procurement cost). On days where amount left over was not observed, we impute the vendor's average amount left over all days in which it was measured. Our measure of profit does not include the subsidy vendors received as part of our intervention. Fisher permutation p-values are obtained by reassigning treatment status across markets. The three smallest and three largest markets (by number of vendors), as well as the three markets with the highest pre-period pea price variability—defined as the range between the 5th and 95th percentiles—are restricted to always remain in the control group. Because one market falls into more than one exclusion category, eight unique markets are excluded in total.

Table A7: Subsidy Impacts: Vendor-Level Outcomes (Alternative Permutation Test)

	Aggregate			
	Total Cost of Wholesale Purchases (Rs.)	Sales (Rs.)	Profits (Rs.)	# vegetables available
	(1)	(2)	(3)	(4)
β_3 Treat	–448.49	–589.02	–140.53	–0.50
	$\langle$ 0.023 $\rangle$	$\langle$ 0.036 $\rangle$	$\langle$ 0.091 $\rangle$	$\langle$ 0.527 $\rangle$
γ_1 Treat $\times$ During Subs	689.60	917.98	228.32	1.97
	$\langle$ < 0.001 $\rangle$	$\langle$ 0.009 $\rangle$	$\langle$ 0.095 $\rangle$	$\langle$ 0.018 $\rangle$
γ_2 Treat $\times$ After Subs	527.12	466.27	–60.84	1.16
	$\langle$ 0.268 $\rangle$	$\langle$ 0.377 $\rangle$	$\langle$ 0.418 $\rangle$	$\langle$ 0.286 $\rangle$
Pre-subsidy intervention market mean	825.12	1167.30	342.18	4.73
Fisher p-value: $\gamma_1 = \gamma_2$	0.555	0.186	0.005	0.041
Number of Vendors	1628	1628	1628	1628
Number of Observations	52898	52898	52898	52898

Notes: This table estimates specification 1 on our full sample. Coefficients for During and Post not shown. Fisher permutation p-value is in $\langle\rangle$. The outcome in column 1 is the total cost of wholesale purchases on a given day, the outcome in column 2 is the vendor's total revenues on a given day accruing from all produce, the outcome in column 3 is the daily profits accrued from all produce, and the outcome in column 4 is the number of distinct types of vegetables a vendor has available on a given day. Profits are calculated by computing (amount of vegetable at the start of the day - amount left over at the end of the day)*anticipated sale price - (amount procured at the start of the day * procurement cost). On days where amount left over was not observed, we impute the vendor's average amount left over all days in which it was measured. Fisher permutation p-values are obtained by reassigning treatment status across markets. The three smallest and three largest markets (by number of vendors), as well as the three markets with the highest pre-period pea price variability—defined as the range between the 5th and 95th percentiles—are restricted to always remain in the control group. Because one market falls into more than one exclusion category, eight unique markets are excluded in total.

Table A8: Subsidy Impacts on Wholesale Prices and Markups

	Peas				Carrots	
	Eligible for Pea Subsidy (Infrequent Pea Seller)		Ineligible for Pea Subsidy (Frequent Pea Seller)		All Vendors	
	Wholesale price (Rs.)	Markup (%)	Wholesale price (Rs.)	Markup (%)	Wholesale price (Rs.)	Markup (%)
	(1)	(2)	(3)	(4)	(5)	(6)
β_3 Treat	2.66	–0.14	–0.54	0.00	–2.69	0.05
	[–7.66, 18.26]	[–0.52, 0.17]	[–6.22, 2.35]	[–0.15, 0.28]	[–7.74, 2.00]	[–0.21, 0.36]
	{ 0.186 }	{ 0.164 }	{ 0.364 }	{ 0.943 }	{ 0.076 }	{ 0.508 }
	⟨ 0.347 ⟩	⟨ 0.279 ⟩	⟨ 0.478 ⟩	⟨ 0.963 ⟩	⟨ 0.112 ⟩	⟨ 0.546 ⟩
γ_1 Treat $\times$ During Subs	3.84	–0.04	0.24	–0.06	1.23	–0.10
	[–1.03, 11.87]	[–0.32, 0.26]	[–3.62, 5.14]	[–0.23, 0.09]	[–1.91, 4.20]	[–0.30, 0.06]
	{ 0.142 }	{ 0.690 }	{ 0.864 }	{ 0.252 }	{ 0.269 }	{ 0.110 }
	⟨ 0.498 ⟩	⟨ 0.819 ⟩	⟨ 0.914 ⟩	⟨ 0.266 ⟩	⟨ 0.345 ⟩	⟨ 0.192 ⟩
γ_2 Treat $\times$ After Subs	2.36	–0.05	3.38	–0.13	1.34	–0.18
	[–4.85, 7.51]	[–0.45, 0.78]	[–2.54, 7.73]	[–0.31, 0.08]	[–5.79, 6.69]	[–0.36, 0.04]
	{ 0.312 }	{ 0.577 }	{ 0.317 }	{ 0.078 }	{ 0.417 }	{ 0.061 }
	⟨ 0.418 ⟩	⟨ 0.699 ⟩	⟨ 0.142 ⟩	⟨ 0.017 ⟩	⟨ 0.419 ⟩	⟨ 0.029 ⟩
Pre-subsidy intervention market mean	31.36	0.26	31.93	0.33	19.12	0.49
Wild Bootstrap p-value: $\gamma_1 = \gamma_2$	0.629	0.878	0.151	0.171	0.965	0.320
Fisher p-value: $\gamma_1 = \gamma_2$	0.878	0.906	0.212	0.202	0.958	0.250
Number of Vendors	480	480	1009	1009	1470	1470
Number of Observations	3631	3631	18807	18807	24789	24789

Notes: This table estimates specification 1 on our full sample. Coefficients for During and Post not shown. 95% wild bootstrap confidence intervals are in [], wild bootstrap p-value is in { }, and Fisher permutation p-value is in ⟨⟩. Columns 1 and 2 present outcomes for vendors eligible for the pea subsidy, columns 3 and 4 for vendors ineligible for the subsidy, and columns 5 and 6 for all vendors. The outcome in columns 1, 3 and 5 is the procurement cost on the given day, the outcome in columns 2, 4 and 6 is the markups computed as $(p^{sell} - p^{ws})/p^{ws}$, where p^{sell} denotes the vendor's retail price and p^{ws} the procurement price. To be eligible for the pea subsidy, vendors must have been present in the market for more than eight days in the pre-subsidy period and have sold peas on fewer than eight days.

Table A9: Subsidy Impacts on Carrots: By Pre-Period Carrot Sales

	Infrequent Carrot Seller				Frequent Carrot Seller			
	Prob. of selling (%) (1)	Sale price (Rs.) (2)	Wholesale qty. (kg) (3)	Profits (Rs.) (4)	Prob. of selling (%) (5)	Sale price (Rs.) (6)	Wholesale qty. (kg) (7)	Profits (Rs.) (8)
β_3 Treat	–0.04	0.15	–0.99	–8.07	0.00	–3.23	–1.24	–13.00
	[–0.12, 0.07]	[–6.46, 8.34]	[–2.25, –0.00]	[–19.72, 0.11]	[–0.15, 0.12]	[–8.23, 1.35]	[–6.00, 2.00]	[–42.41, 16.46]
	{ 0.235 }	{ 0.929 }	{ 0.049 }	{ 0.050 }	{ 0.957 }	{ 0.069 }	{ 0.217 }	{ 0.292 }
	⟨ 0.158 ⟩	⟨ 0.932 ⟩	⟨ 0.011 ⟩	⟨ 0.034 ⟩	⟨ 0.949 ⟩	⟨ 0.082 ⟩	⟨ 0.301 ⟩	⟨ 0.251 ⟩
γ_1 Treat × During Subs	0.70	–1.72	6.03	46.40	0.43	0.17	5.85	43.16
	[0.60, 0.78]	[–4.37, 1.82]	[4.72, 6.76]	[29.49, 56.36]	[0.21, 0.64]	[–3.50, 4.17]	[3.62, 8.43]	[16.74, 66.15]
	{ < 0.001 }	{ 0.522 }	{ < 0.001 }	{ 0.002 }	{ 0.014 }	{ 0.802 }	{ 0.004 }	{ 0.016 }
	⟨ < 0.001 ⟩	⟨ 0.373 ⟩	⟨ < 0.001 ⟩	⟨ 0.002 ⟩	⟨ < 0.001 ⟩	⟨ 0.854 ⟩	⟨ < 0.001 ⟩	⟨ 0.003 ⟩
γ_2 Treat × After Subs	0.11	–1.80	1.81	10.71	0.09	–2.05	1.91	7.17
	[0.00, 0.19]	[–8.01, 2.94]	[0.28, 2.73]	[0.65, 19.47]	[–0.19, 0.31]	[–11.59, 7.10]	[–0.73, 5.22]	[–33.49, 49.15]
	{ 0.044 }	{ 0.150 }	{ 0.017 }	{ 0.039 }	{ 0.422 }	{ 0.205 }	{ 0.091 }	{ 0.336 }
	⟨ 0.096 ⟩	⟨ 0.330 ⟩	⟨ 0.032 ⟩	⟨ 0.015 ⟩	⟨ 0.481 ⟩	⟨ 0.258 ⟩	⟨ 0.304 ⟩	⟨ 0.489 ⟩
Pre-subsidy intervention market mean	0.18	27.99	1.02	8.26	0.81	27.90	5.94	43.21
Wild Bootstrap p-value: $\gamma_1 = \gamma_2$	<0.001	0.978	<0.001	<0.001	0.017	0.191	0.017	0.062
Fisher p-value: $\gamma_1 = \gamma_2$	<0.001	0.976	<0.001	<0.001	0.004	0.207	<0.001	<0.001
Number of Vendors	698	586	698	698	933	884	933	933
Number of Observations	24383	5374	24383	24381	30835	19699	30835	30832

Notes: This table estimates specification 1 on our full sample. Coefficients for During and Post not shown. 95% wild bootstrap confidence intervals are in [], wild bootstrap p-value is in { }, and Fisher permutation p-value is in ⟨⟩.Columns 1–4 report outcomes for vendors who were present in the market for more than eight days during the pre-subsidy period and sold carrots on fewer than eight days in that period (analogous to the pea subsidy eligibility criterion). Columns 5–8 report outcomes for vendors who were also present for more than eight days in the pre-subsidy period but sold carrots on eight or more days. The outcome in columns 1 and 5 is whether the vendor sells carrots on the given day, the outcome in columns 2 and 6 measure the vendor's anticipated sale price for carrots, the outcome in columns 3 and 7 measure the wholesale quantity of carrots procured, and the outcome in columns 4 and 8 measure the daily profits accrued from carrots. Profits are calculated by computing (amount of carrots at the start of the day - amount left over at the end of the day)*anticipated sale price - (amount procured at the start of the day * procurement cost). On days where amount left over was not observed, we impute the vendor's average amount left over all days in which it was measured. Our measure of profit does not include the subsidy vendors received as part of our intervention.

Table A10: Subsidy Impacts: Carrots and Peas, Removing Control Markets Most Likely to be Impacted by Demand Spillovers

	Carrot				Peas			
	Prob. of selling (%) (1)	Sale price (Rs.) (2)	Wholesale qty. (kg) (3)	Profits (Rs.) (4)	Prob. of selling (%) (5)	Sale price (Rs.) (6)	Wholesale qty. (kg) (7)	Profits (Rs.) (8)
β_3 Treat	–0.04	–2.41	–1.29	–10.66	–0.04	0.40	–2.53	–30.33
	[–0.30, 0.21]	[–6.84, 2.29]	[–4.20, 1.11]	[–26.03, 9.89]	[–0.32, 0.28]	[–3.36, 3.50]	[–7.17, 1.44]	[–72.24, 25.11]
	{ 0.683 }	{ 0.095 }	{ 0.288 }	{ 0.317 }	{ 0.549 }	{ 0.599 }	{ 0.096 }	{ 0.096 }
	⟨ 0.666 ⟩	⟨ 0.131 ⟩	⟨ 0.312 ⟩	⟨ 0.168 ⟩	⟨ 0.618 ⟩	⟨ 0.656 ⟩	⟨ 0.151 ⟩	⟨ 0.065 ⟩
γ_1 Treat × During Subs	0.58	0.00	5.99	44.95	0.40	–0.59	6.84	59.02
	[0.40, 0.73]	[–2.11, 2.04]	[3.73, 7.69]	[20.49, 57.72]	[0.11, 0.62]	[–3.90, 2.96]	[3.10, 10.93]	[3.60, 102.35]
	{ 0.002 }	{ 1.000 }	{ < 0.001 }	{ 0.003 }	{ 0.021 }	{ 0.795 }	{ 0.017 }	{ 0.036 }
	⟨ 0.001 ⟩	⟨ 0.999 ⟩	⟨ 0.001 ⟩	⟨ 0.001 ⟩	⟨ < 0.001 ⟩	⟨ 0.774 ⟩	⟨ < 0.001 ⟩	⟨ 0.001 ⟩
γ_2 Treat × After Subs	0.11	–1.08	2.01	11.29	0.06	–0.66	2.58	25.22
	[–0.06, 0.26]	[–4.20, 1.54]	[–0.02, 4.03]	[–4.09, 25.80]	[–0.24, 0.30]	[–4.17, 1.65]	[–0.95, 6.79]	[–29.61, 69.36]
	{ 0.333 }	{ 0.287 }	{ 0.056 }	{ 0.198 }	{ 0.452 }	{ 0.632 }	{ 0.075 }	{ 0.111 }
	⟨ 0.297 ⟩	⟨ 0.312 ⟩	⟨ 0.176 ⟩	⟨ 0.100 ⟩	⟨ 0.504 ⟩	⟨ 0.601 ⟩	⟨ 0.094 ⟩	⟨ 0.074 ⟩
Pre-subsidy intervention market mean	0.49	27.91	3.43	25.41	0.57	41.80	5.94	47.73
Wild Bootstrap p-value: $\gamma_1 = \gamma_2$	<0.001	0.394	<0.001	<0.001	0.002	0.965	0.004	0.014
Fisher p-value: $\gamma_1 = \gamma_2$	0.001	0.459	0.001	<0.001	0.003	0.974	<0.001	0.007
Number of Vendors	1477	1329	1477	1477	1477	1351	1477	1477
Number of Observations	50042	22173	50042	50037	50060	20334	50060	50058

Notes: This table replicates Table 1 excluding control markets that were frequently cited as a likely substitute for each treatment market. Substitute control markets were defined as any of top three responses by vendors in treatment markets to this question: 'If customers were not buying from this market, where would they buy?' Relative to our full sample, this sample excludes three control markets, as the majority of responses to this question were markets that are not in our sample. Coefficients for During and Post not shown. 95% wild bootstrap confidence intervals are in [], wild bootstrap p-value is in {}, and Fisher permutation p-value is in ⟨⟩. Columns 1 - 4 present outcomes for carrots, and 5 - 8 for peas. The outcome in columns 1 and 5 is whether the vendor sells carrots or peas on the given day, the outcome in columns 2 and 6 measure the vendor's anticipated sale price for the relevant vegetable, the outcome in columns 3 and 7 measure the wholesale quantity procured of the relevant vegetable, and the outcome in columns 4 and 8 measure the daily profits accrued from the relevant vegetables. Profits are calculated by computing (amount of vegetable at the start of the day - amount left over at the end of the day)*anticipated sale price - (amount procured at the start of the day * procurement cost). On days where amount left over was not observed, we impute the vendor's average amount left over all days in which it was measured. Our measure of profit does not include the subsidy vendors received as part of our intervention.

Table A11: Subsidy Impacts: Aggregate, Removing Control Markets Most Likely to be Impacted by Demand Spillovers

	Aggregate			
	Total Cost of Wholesale Purchases (Rs.)	Sales (Rs.)	Profits (Rs.)	# vegetables available
	(1)	(2)	(3)	(4)
β_3 Treat	–439.42	–563.24	–123.82	–0.46
	[–1113.57, 101.23]	[–1402.15, 17.62]	[–360.95, 85.94]	[–2.61, 1.68]
	{ 0.062 }	{ 0.052 }	{ 0.165 }	{ 0.590 }
	⟨ 0.050 ⟩	⟨ 0.046 ⟩	⟨ 0.107 ⟩	⟨ 0.491 ⟩
γ_1 Treat × During Subs	712.69	952.08	239.33	2.06
	[316.76, 1139.68]	[319.81, 1458.50]	[–55.53, 516.72]	[0.64, 3.21]
	{ 0.026 }	{ 0.028 }	{ 0.070 }	{ 0.031 }
	⟨ < 0.001 ⟩	⟨ 0.004 ⟩	⟨ 0.021 ⟩	⟨ 0.013 ⟩
γ_2 Treat × After Subs	561.93	542.41	–19.52	1.26
	[–93.44, 1100.95]	[–208.41, 1145.89]	[–297.12, 216.91]	[–0.45, 2.72]
	{ 0.249 }	{ 0.301 }	{ 0.698 }	{ 0.307 }
	⟨ 0.219 ⟩	⟨ 0.304 ⟩	⟨ 0.776 ⟩	⟨ 0.200 ⟩
Pre-subsidy intervention market mean	825.12	1167.30	342.18	4.73
Wild Bootstrap p-value: $\gamma_1 = \gamma_2$	0.587	0.241	0.011	0.050
Fisher p-value: $\gamma_1 = \gamma_2$	0.596	0.201	0.001	0.051
Number of Vendors	1474	1474	1474	1474
Number of Observations	48098	48098	48098	48098

Notes: This table replicates Table 2 excluding control markets that were frequently cited as a likely substitute for each treatment market. Substitute control markets were defined as any of top three responses by vendors in treatment markets to this question: 'If customers were not buying from this market, where would they buy?'. Relative to our full sample, this sample excludes three control markets, as the majority of responses to this question were markets that are not in our sample. Coefficients for During and Post not shown. 95% wild bootstrap confidence intervals are in [], wild bootstrap p-value is in {}, and Fisher permutation p-value is in ⟨⟩. The outcome in column 1 is the total cost of wholesale purchases on a given day, the outcome in column 2 is the vendor's total revenues on a given day accruing from all produce, the outcome in column 3 is the daily profits accrued from all produce, and the outcome in column 4 is the number of distinct types of vegetables a vendor has available on a given day. Profits are calculated by computing (amount of vegetable at the start of the day - amount left over at the end of the day)*anticipated sale price - (amount procured at the start of the day * procurement cost). On days where amount left over was not observed, we impute the vendor's average amount left over all days in which it was measured.

Table A12: Subsidy Impacts: Carrots and Peas, Removing Control Vendors Most Likely to be Impacted by Supply Spillovers

	Carrot				Peas			
	Prob. of selling (%) (1)	Sale price (Rs.) (2)	Wholesale qty. (kg) (3)	Profits (Rs.) (4)	Prob. of selling (%) (5)	Sale price (Rs.) (6)	Wholesale qty. (kg) (7)	Profits (Rs.) (8)
β_3 Treat	0.03	–1.54	–0.59	–5.66	0.03	0.20	–1.33	–18.72
	[–0.14, 0.25]	[–4.33, 1.14]	[–2.08, 1.65]	[–14.89, 14.87]	[–0.22, 0.31]	[–2.26, 3.06]	[–4.28, 1.37]	[–44.58, 8.68]
	{ 0.637 }	{ 0.253 }	{ 0.546 }	{ 0.348 }	{ 0.631 }	{ 0.834 }	{ 0.384 }	{ 0.093 }
	⟨ 0.713 ⟩	⟨ 0.295 ⟩	⟨ 0.586 ⟩	⟨ 0.488 ⟩	⟨ 0.676 ⟩	⟨ 0.869 ⟩	⟨ 0.527 ⟩	⟨ 0.286 ⟩
γ_1 Treat × During Subs	0.56	0.29	5.59	43.38	0.39	0.55	6.13	52.33
	[0.39, 0.72]	[–1.58, 2.31]	[3.33, 7.36]	[18.85, 55.97]	[0.22, 0.56]	[–2.83, 4.01]	[3.97, 8.32]	[30.44, 73.30]
	{ < 0.001 }	{ 0.755 }	{ < 0.001 }	{ < 0.001 }	{ 0.006 }	{ 0.747 }	{ 0.004 }	{ 0.009 }
	⟨ 0.006 ⟩	⟨ 0.801 ⟩	⟨ 0.011 ⟩	⟨ 0.012 ⟩	⟨ 0.015 ⟩	⟨ 0.833 ⟩	⟨ 0.005 ⟩	⟨ < 0.001 ⟩
γ_2 Treat × After Subs	0.11	–1.51	1.75	10.36	0.04	–0.82	1.89	18.47
	[–0.07, 0.26]	[–3.66, 0.38]	[–0.16, 3.83]	[–4.06, 20.37]	[–0.16, 0.27]	[–4.43, 1.29]	[–0.53, 4.68]	[–6.56, 41.22]
	{ 0.395 }	{ 0.145 }	{ 0.078 }	{ 0.186 }	{ 0.549 }	{ 0.603 }	{ 0.065 }	{ 0.064 }
	⟨ 0.396 ⟩	⟨ 0.214 ⟩	⟨ 0.219 ⟩	⟨ 0.162 ⟩	⟨ 0.726 ⟩	⟨ 0.605 ⟩	⟨ 0.275 ⟩	⟨ 0.183 ⟩
Pre-subsidy intervention market mean	0.49	27.91	3.43	25.41	0.57	41.80	5.94	47.73
Wild Bootstrap p-value: $\gamma_1 = \gamma_2$	<0.001	0.241	<0.001	<0.001	0.001	0.557	0.001	0.001
Fisher p-value: $\gamma_1 = \gamma_2$	0.004	0.344	<0.001	0.004	0.012	0.704	0.025	0.028
Number of Vendors	1013	888	1013	1013	1013	898	1013	1013
Number of Observations	33581	13556	33581	33576	33615	12386	33615	33613

Notes: This table replicates Table 1 excluding vendors in control markets who frequently buy in the two most popular wholesale markets among treated vendors in the sample, Kole Market and Sealdaha. Relative to our full sample, this sample excludes 626 vendors. Coefficients for During and Post not shown. 95% wild bootstrap confidence intervals are in [], wild bootstrap p-value is in { }, and Fisher permutation p-value is in ⟨⟩. Columns 1 - 4 present outcomes for carrots, and 5 - 8 for peas. The outcome in columns 1 and 5 is whether the vendor sells carrots or peas on the given day, the outcome in columns 2 and 6 measure the vendor's anticipated sale price for the relevant vegetable, the outcome in columns 3 and 7 measure the wholesale quantity procured of the relevant vegetable, and the outcome in columns 4 and 8 measure the daily profits accrued from the relevant vegetables. Profits are calculated by computing (amount of vegetable at the start of the day - amount left over at the end of the day)*anticipated sale price - (amount procured at the start of the day * procurement cost). On days where amount left over was not observed, we impute the vendor's average amount left over all days in which it was measured. Our measure of profit does not include the subsidy vendors received as part of our intervention.

Table A13: Subsidy Impacts: Aggregate, Removing Control Vendors Most Likely to be Impacted by Supply Spillovers

	Aggregate			
	Total Cost of Wholesale Purchases (Rs.)	Sales (Rs.)	Profits (Rs.)	# vegetables available
	(1)	(2)	(3)	(4)
β_3 Treat	–349.25	–459.82	–110.57	–0.13
	[–797.23, 67.41]	[–1050.17, 120.66]	[–289.15, 65.07]	[–1.97, 1.66]
	{ 0.071 }	{ 0.084 }	{ 0.193 }	{ 0.851 }
	⟨ 0.148 ⟩	⟨ 0.136 ⟩	⟨ 0.243 ⟩	⟨ 0.904 ⟩
γ_1 Treat × During Subs	687.91	923.20	235.23	2.06
	[390.76, 1086.14]	[530.97, 1315.43]	[77.80, 369.36]	[0.76, 3.02]
	{ 0.002 }	{ 0.004 }	{ 0.021 }	{ 0.008 }
	⟨ 0.015 ⟩	⟨ 0.009 ⟩	⟨ 0.026 ⟩	⟨ 0.039 ⟩
γ_2 Treat × After Subs	531.50	499.16	–32.34	1.30
	[–137.03, 1078.55]	[–258.28, 1097.24]	[–176.49, 124.00]	[–0.45, 2.70]
	{ 0.273 }	{ 0.307 }	{ 0.507 }	{ 0.292 }
	⟨ 0.318 ⟩	⟨ 0.414 ⟩	⟨ 0.620 ⟩	⟨ 0.269 ⟩
Pre-subsidy intervention market mean	825.12	1167.30	342.18	4.73
Wild Bootstrap p-value: $\gamma_1 = \gamma_2$	0.568	0.212	0.006	0.070
Fisher p-value: $\gamma_1 = \gamma_2$	0.689	0.299	0.010	0.134
Number of Vendors	1011	1011	1011	1011
Number of Observations	32170	32170	32170	32170

Notes: This table replicates Table 2 excluding vendors in control markets who frequently buy in the two most popular wholesale markets among treated vendors in the sample, Kole Market and Sealdaha. Relative to our full sample, this sample excludes 626 vendors. Coefficients for During and Post not shown. 95% wild bootstrap confidence intervals are in [], wild bootstrap p-value is in {}, and Fisher permutation p-value is in ⟨⟩. The outcome in column 1 is the total cost of wholesale purchases on a given day, the outcome in column 2 is the vendor's total revenues on a given day accruing from all produce, the outcome in column 3 is the daily profits accrued from all produce, and the outcome in column 4 is the number of distinct types of vegetables a vendor has available on a given day. Profits are calculated by computing (amount of vegetable at the start of the day - amount left over at the end of the day)*anticipated sale price - (amount procured at the start of the day * procurement cost). On days where amount left over was not observed, we impute the vendor's average amount left over all days in which it was measured.

Table A14: The Subsidy Intervention Induces Vendors to Stock Beyond the Subsidy Cap: Carrots and Peas

	Carrot	Peas
	> 7kg Carrots in Stock	> 10kg Peas in Stock
	(1)	(2)
β_3 Treat	–0.13	–0.16
	[–0.27, 0.05]	[–0.31, –0.04]
	{ 0.288 }	{ 0.032 }
	⟨ 0.143 ⟩	⟨ 0.002 ⟩
γ_1 Treat × During Subs	0.26	0.17
	[0.14, 0.40]	[0.05, 0.30]
	{ 0.020 }	{ 0.034 }
	⟨ < 0.001 ⟩	⟨ 0.002 ⟩
γ_2 Treat × After Subs	0.10	0.13
	[–0.11, 0.26]	[–0.02, 0.29]
	{ 0.327 }	{ 0.055 }
	⟨ 0.332 ⟩	⟨ 0.026 ⟩
Pre-subsidy intervention market mean	0.32	0.30
Wild Bootstrap p-value: $\gamma_1 = \gamma_2$	0.035	0.258
Fisher p-value: $\gamma_1 = \gamma_2$	0.015	0.289
Number of Vendors	1661	1661
Number of Observations	79418	79443

Notes: This table estimates specification 1 on our full sample. Coefficients for During and Post not shown. 95% wild bootstrap confidence intervals are in [], wild bootstrap p-value is in { }, and Fisher permutation p-value is in ⟨⟩. The outcome in column 1 is a dummy equal to one if the vendor stocked more than 7 kg of carrots on a given day, while in column 2 it is a dummy equal to one if the vendor stocked more than 10 kg of peas on a given day.

Table A15: Number of Vendors Instrument: First Stage

	Peas			Carrots		
	ln(Quantity Sold)			ln(Quantity Sold)		
	(1)	(2)	(3)	(4)	(5)	(6)
Number of vendors selling vegetable *i*	0.066***	0.064***	0.051***	0.054***	0.053***	0.050***
	(0.009)	(0.012)	(0.013)	(0.011)	(0.008)	(0.008)
Total Number of vendors selling anything	0.006***	-0.003	-0.001	0.007**	-0.006**	-0.003
	(0.001)	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)
Observations	302	302	302	304	304	304
Market FE		X	X		X	X
Day FE			X			X
F-Stat	51.6	27.0	16.2	24.4	42.6	38.7

Notes: The table presents the first-stage relationship between the quantity sold and the total number of vendor in each market selling carrots and peas during the subsidy period in control markets. The table reports estimates at the market-day level. The dependent variable is the natural logarithm of the quantity sold in the market, and the main regressor of interest is the number of vendors in the market on a given day selling peas or carrots. All estimations use standard errors clustered at the market level. Columns 1 and 4 do not include fixed effects, columns 2 and 5 include market fixed effects, and columns 3 and 6 include market and day fixed effects.

Table A16: Number of Vendors Instrument: Elasticity of Demand Estimates

	Peas			Carrots		
	ln(Price)			ln(Price)		
	(1)	(2)	(3)	(4)	(5)	(6)
ln(Quantity sold)	-0.081*	-0.269***	-0.025	0.044	-0.120***	-0.040
	(0.044)	(0.039)	(0.035)	(0.060)	(0.037)	(0.040)
Total Number of vendors in each Market	0.001	0.008***	0.002	-0.002	0.002	0.000
	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	302	302	302	304	304	304
Market FE		X	X		X	X
Day FE			X			X

Notes: This table presents estimates of the inverse demand elasticity (the coefficient on ln(Quantity Sold)) of carrots and peas during the subsidy for control markets. Observations are at the market-day level. The coefficients are estimated using an instrumental variables approach, where the dependent variable is the log of the vegetable's daily price (in rupees) and the endogenous regressor is the log of the quantity sold in the market, instrumented by the total number of vendors in each market selling a vegetable each day. Standard errors are clustered at the market level. Columns 1 and 4 do not include fixed effects, columns 2 and 5 include market fixed effects, and columns 3 and 6 include market and day fixed effects.

Table A17: OLS Regression of Log Price on Log Quantity

	Peas			Carrots		
	ln(Price)			ln(Price)		
	(1)	(2)	(3)	(4)	(5)	(6)
ln(Quantity sold)	-0.067	-0.192***	-0.005	-0.004	-0.071*	-0.024
	(0.039)	(0.051)	(0.024)	(0.043)	(0.036)	(0.028)
Total Number of vendors in each Market	0.001	0.007***	0.002	-0.001	0.001	0.000
	(0.001)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)
Observations	302	302	302	304	304	304
Market FE		X	X		X	X
Day FE			X			X

Notes: The table presents the OLS regression of the log of average market price (in rupees) on the log of quantity sold at the market x day level for peas and carrots. The sample is all market x days for control markets during the subsidy period. Standard errors are two-way clustered by market and by day. Columns 1 and 4 do not include fixed effects; columns 2 and 5 include market fixed effects; columns 3 and 6 include market and day fixed effects.

Table A18: Subsidy Instrument: First Stage

	Peas			**Carrots**		
	ln(Quantity Sold)			ln(Quantity Sold)		
	(1)	(2)	(3)	(4)	(5)	(6)
Treat × During Subs	1.36	1.38	1.38	1.38	1.39	1.39
	[0.57, 2.06]	[0.58, 2.04]	[0.58, 2.02]	[0.89, 1.90]	[0.88, 1.91]	[0.89, 1.91]
	{0.018}	{0.017}	{0.018}	{0.002}	{0.002}	{0.002}
	⟨ 0 ⟩	⟨ 0 ⟩	⟨ 0 ⟩	⟨ 0 ⟩	⟨ 0 ⟩	⟨ 0 ⟩
Observations	696	696	696	698	698	698
Market FE		X	X		X	X
Day FE			X			X
F-Stat	70.16	71.17	68.59	45.19	43.96	41.86

Notes: This table presents the first-stage relationship between the log quantity sold of peas and carrots and the subsidy intervention. The main regressor of interest is the interaction between the treatment market dummy and the subsidy period (Treatment × During), which is the only coefficient reported across all specifications. In addition to the interaction term, Columns 1 and 4 include the main effects for Treatment and During; Columns 2 and 5 include During and market fixed effects; and Columns 3 and 6 include market and day fixed effects. 95% wild bootstrap confidence intervals are in [], wild bootstrap p-value is in {}, and Fisher permutation p-value is in ⟨⟩. The table reports estimates at the market-day level. The dependent variable is the log quantity of peas or carrots sold in the market, and the main regressor of interest is the interaction between whether a market was a treatment market and whether the subsidy intervention was taking place. Standard errors are clustered at the market level. Columns 1 and 4 do not include fixed effects, columns 2 and 5 include market fixed effects, and columns 3 and 6 include market and day fixed effects.

Table A19: Subsidy Instrument: Elasticity of Demand Estimates

	Peas			Carrots		
	ln(Price (Rs.))			ln(Price (Rs.))		
	(1)	(2)	(3)	(4)	(5)	(6)
ln(Quantity sold)	-0.021	-0.020	-0.021	-0.018	-0.018	-0.017
	[-0.10, 0.03]	[-0.10, 0.03]	[-0.10, 0.03]	[-0.07, 0.04]	[-0.07, 0.04]	[-0.07, 0.04]
	{0.716}	{0.722}	{0.698}	{0.488}	{0.482}	{0.488}
	⟨ .276 ⟩	⟨ .279 ⟩	⟨ .263 ⟩	⟨ .296 ⟩	⟨ .307 ⟩	⟨ .287 ⟩
Observations	696	696	696	698	698	698
Market FE		X	X		X	X
Day FE			X			X

Notes: This table presents estimates of the demand elasticity of carrots and peas during the subsidy for control markets. Observations are at the market-day level. In addition to the interaction term, Columns 1 and 4 include the main effects for Treatment and During; Columns 2 and 5 include During and market fixed effects; and Columns 3 and 6 include market and day fixed effects. Main effects for Treatment and During are not individually reported. 95% wild bootstrap confidence intervals are in [], wild bootstrap p-value is in { }, and Fisher permutation p-value is in ⟨⟩. The table reports estimates at the market-day level. The coefficients are estimated using an instrumental variables approach, where the dependent variable is the log of vegetable's daily price and the endogenous regressor is the log of the quantity sold in the market, instrumented by the interaction between whether a market was a treatment market and whether the subsidy intervention was taking place. Standard errors are clustered at the market level. Columns 1 and 4 do not include fixed effects, columns 2 and 5 include market fixed effects, and columns 3 and 6 include market and day fixed effects.

Table A20: Effects on the stocking of non-targeted vegetables

	Total number of vegetables available
	(1)
β_3 Treat	–0.41
	[–2.06, 1.25]
	{ 0.559 }
	$\langle$ 0.464 $\rangle$
γ_1 Treat $\times$ During Subs	1.00
	[–0.23, 2.13]
	{ 0.131 }
	$\langle$ 0.156 $\rangle$
γ_2 Treat $\times$ After Subs	1.00
	[–0.47, 2.37]
	{ 0.322 }
	$\langle$ 0.239 $\rangle$
Pre-subsidy intervention market mean	3.68
Wild Bootstrap p-value: $\gamma_1 = \gamma_2$	0.987
Fisher p-value: $\gamma_1 = \gamma_2$	0.988
Number of Vendors	1628
Number of Observations	52898

Notes: This table estimates specification 1 on our full sample. Coefficients for During and Post not shown. 95% wild bootstrap confidence intervals are in [], wild bootstrap p-value is in { }, and Fisher permutation p-value is in $\langle\rangle$. Column 1 reports the total number of vegetables a vendor has available for sale on a given day, excluding carrots and peas.

Table A21: Impacts on Wholesale Market Visits

	Visited Wholesale Market (=0/1)
	(1)
β_3 Treat	–0.01
	[–0.05, 0.02]
	{ 0.846 }
	$\langle$ 0.771 $\rangle$
γ_1 Treat $\times$ During Subs	0.01
	[–0.03, 0.05]
	{ 0.848 }
	$\langle$ 0.847 $\rangle$
γ_2 Treat $\times$ After Subs	–0.00
	[–0.06, 0.05]
	{ 0.947 }
	$\langle$ 0.968 $\rangle$
Pre-subsidy intervention market mean	0.97
Wild Bootstrap p-value: $\gamma_1 = \gamma_2$	0.104
Fisher p-value: $\gamma_1 = \gamma_2$	0.046
Number of Vendors	1628
Number of Observations	52898

Notes: This table estimates specification 1 on our full sample. Coefficients for During and Post not shown. 95% wild bootstrap confidence intervals are in [], wild bootstrap p-value is in {}, and Fisher permutation p-value is in $\langle\rangle$. The outcome for column 1 is a dummy variable equal to one if the vendor visited the wholesale market on a given day.

Table A22: Complementary Goods Analysis

	Sells Carrots	Sells Peas
	(1)	(2)
Sells Ladies' Fingers	0.09***	0.09***
	(0.01)	(0.01)
Sells Bitter Gourds (Large)	0.12***	0.11***
	(0.01)	(0.01)
Sells Cucumbers	0.12***	0.10***
	(0.01)	(0.01)
Sells Brinjals	-0.05***	0.05***
	(0.02)	(0.02)
Sells Tomatoes	0.14***	0.24***
	(0.01)	(0.02)
Sells Bottle Gourds	0.02*	-0.04***
	(0.01)	(0.01)
Sells Cabbages	0.04***	0.04***
	(0.01)	(0.01)
Sells Cauliflowers	-0.06***	0.07***
	(0.01)	(0.01)
Sells Peas	0.20***	
	(0.01)	
Sells Carrots		0.18***
		(0.01)
Mean of outcome	0.534	0.601
Observations	20848	20833
Market FE	X	X
Day FE	X	X

Notes: The table reports the association between vendors selling different vegetables and whether a vendor sells carrots or peas during the pre-subsidy period. All regression variables are indicators equal to 1 if the vendor sells the corresponding vegetable. Each column presents estimates from linear regressions at the vendor-day level, where the dependent variable is an indicator equal to 1 if the vendor sells the main good of each panel (carrots in column 1 and peas in column 2). Standard errors are clustered at the vendor level. Columns 1 and 2 include day and market fixed effects.

Table A23: The Subsidy's Impact on Profits Under Varying Assumptions About Labor Supply

	Profits adjusted for hours worked			
	Not adjusted	1 hour	2 hours	3 hours
	(1)	(2)	(3)	(4)
β_3 Treat	–140.53	–140.53	–140.53	–140.53
	[–353.59, 79.76]	[–353.59, 79.76]	[–353.59, 79.76]	[–353.59, 79.76]
	{ 0.120 }	{ 0.120 }	{ 0.120 }	{ 0.120 }
	⟨ 0.075 ⟩	⟨ 0.076 ⟩	⟨ 0.075 ⟩	⟨ 0.075 ⟩
γ_1 Treat × During Subs	228.32	184.93	141.55	98.16
	[–102.95, 477.54]	[–146.34, 434.15]	[–189.73, 390.76]	[–233.11, 347.37]
	{ 0.073 }	{ 0.084 }	{ 0.105 }	{ 0.171 }
	⟨ 0.025 ⟩	⟨ 0.039 ⟩	⟨ 0.104 ⟩	⟨ 0.210 ⟩
γ_2 Treat × After Subs	–60.84	–60.84	–60.84	–60.84
	[–373.98, 248.20]	[–373.98, 248.20]	[–373.98, 248.20]	[–373.98, 248.20]
	{ 0.316 }	{ 0.316 }	{ 0.316 }	{ 0.316 }
	⟨ 0.404 ⟩	⟨ 0.403 ⟩	⟨ 0.403 ⟩	⟨ 0.404 ⟩
Pre-subsidy intervention market mean	342.18	342.18	342.18	342.18
Wild Bootstrap p-value: $\gamma_1 = \gamma_2$	0.026	0.042	0.055	0.069
Fisher p-value: $\gamma_1 = \gamma_2$	0.002	0.002	0.005	0.016
Number of Vendors	1628	1628	1628	1628
Number of Observations	52898	52898	52898	52898

This table estimates specification 1 on the full sample and shows the Subsidy's Impact on Profits after accounting for the cost of hours worked. Column (1) reports the unadjusted daily profits, while Columns (2), (3), and (4) report daily profits net of one, two, and three hours of labor, respectively. During the subsidy period, profits adjusted for labor are obtained by subtracting the cost of the first one, two, or three hours of work from daily profits for treatment-group vendors. Daily profits accrued from all produce are computed as revenues minus procurement costs. Revenues are measured as the quantity sold during the day-defined as the difference between the quantity held at the start of the day and the quantity left over at the end of the day-multiplied by the anticipated sale price. Procurement costs are calculated as the quantity procured at the start of the day multiplied by the procurement cost. When leftover quantities are unobserved, we impute the vendor's average leftover quantity across days in which it is observed. This profit measure excludes the subsidy received as part of the intervention. To account for labor costs, we impute an average cost per hour defined as the ratio of average pre-period daily profits to average hours worked per day. Average pre-period daily profits are computed for treatment-group vendors using the same vendor-day aggregation as in Table 2. Average hours worked per day are calculated using vendor-day observations from the 2023 follow-up survey, conditioning on days in which the vendor worked. Coefficients for During and Post are not shown. 95% wild bootstrap confidence intervals are reported in [ ], wild bootstrap p-values in { }, and Fisher permutation p-values in ⟨⟩.

Table A24: Subsidy Impacts: Market-Level Outcomes

	Aggregate			
	Standard deviation (1)	Gap between 95th and 5th percentiles (2)	95th percentile (3)	5th percentile (4)
β_3 Treat	–38.11	–128.19	–180.19	–52.00
	[–211.67, 127.50]	[–691.97, 406.09]	[–717.96, 415.85]	[–134.77, 51.09]
	{ 0.562 }	{ 0.550 }	{ 0.411 }	{ 0.145 }
	⟨ 0.514 ⟩	⟨ 0.546 ⟩	⟨ 0.444 ⟩	⟨ 0.178 ⟩
γ_1 Treat $\times$ During Subs	188.36	535.02	537.40	2.38
	[17.54, 330.05]	[–23.80, 965.63]	[–149.84, 1061.76]	[–129.79, 138.96]
	{ 0.046 }	{ 0.052 }	{ 0.058 }	{ 0.967 }
	⟨ 0.009 ⟩	⟨ 0.025 ⟩	⟨ 0.033 ⟩	⟨ 0.978 ⟩
γ_2 Treat $\times$ After Subs	–74.98	–210.86	–219.31	–8.44
	[–328.36, 163.75]	[–1031.45, 601.13]	[–1117.59, 586.05]	[–183.55, 149.70]
	{ 0.304 }	{ 0.266 }	{ 0.237 }	{ 0.836 }
	⟨ 0.317 ⟩	⟨ 0.323 ⟩	⟨ 0.296 ⟩	⟨ 0.868 ⟩
Pre-subsidy intervention market mean	280.24	899.92	957.12	57.20
Wild Bootstrap p-value: $\gamma_1 = \gamma_2$	0.034	0.044	0.044	0.775
Fisher p-value: $\gamma_1 = \gamma_2$	<0.001	<0.001	0.002	0.839
Number of Markets	20	20	20	20
Number of Observations	959	959	959	959

Notes: This table estimates specification 1 on our full sample. The unit of observation is market-day. Coefficients for During and Post not shown. 95% wild bootstrap confidence intervals are in [], wild bootstrap p-value is in {}, and Fisher permutation p-value is in ⟨⟩. The outcome in column 1 is the standard deviation of profits across vendors within a market-day, column 2 is the gap between the 95th and 5th percentiles of profits across vendors within a market-day, column 3 is the 95th percentile of profits, and column 4 is the 5th percentile of profits. Profits are calculated by computing (amount of vegetable at the start of the day - amount left over at the end of the day)*anticipated sale price - (amount procured at the start of the day * procurement cost). On days where amount left over was not observed, we impute the vendor's average amount left over all days in which it was measured.