

Working paper



International
Growth Centre

Informal Monitoring Mechanisms in Public Service Delivery

Evidence from the
Public Distribution
System in India



Sriniketh Nagavarapu
Sheetal Sekhri

February 2013

DIRECTED BY



FUNDED BY



Informal Monitoring Mechanisms in Public Service Delivery: Evidence from the Public Distribution System in India *

Sriniketh Nagavarapu[†] and Sheetal Sekhri[‡]

January 31, 2013

Abstract

Informal monitoring and enforcement can increase the efficacy of public service delivery. We study the Targeted Public Distribution System of India and find that Scheduled Castes (SC) have a higher take-up of government subsidized food when facing SC delivery agents. We provide evidence suggesting that this effect works through increased informal monitoring and enforcement when the delivery agent is corrupt. We then estimate a structural model and show that the welfare that SC households would gain from lowering monitoring and enforcement costs – an amount equivalent to moving from a non-SC shopkeeper to a SC shopkeeper – are important, equaling approximately one-fifth of the average subsidy amount. Additionally, expanding the generosity of the program - as envisioned in the proposed National Food Security Bill - can perversely lower welfare for SCs and non-SCs due to increased incentives for black-marketing.

Keywords: Food Security, Corruption, Service Delivery

*We thank Florencia Borrescio-Higa, Paul Christian, Brian Cross, and David Glick for excellent research assistance. In addition, we have received valuable comments from Anna Aizer, Ken Chay, Leora Friedberg, Andrew Foster, Vernon Henderson, Anjini Kochar, Yongsuk Lee, Thomas MaCurdy, Aprajit Mahajan, Kaivan Munshi, Karthik Muralidharan, Paul Niehaus, Holger Sieg, Steve Stern, Nick Wilson, Frank Wolak, and seminar participants at Brown's Race and Inequality Seminar, the WBCA Colloquium, the Stanford Institute for Theoretical Economics, the Freeman Spogli Institute, the 21st BREAD Conference, the IGC-ISI Development Conference, and Tufts University. Funding from International Growth Centre (IGC) under CPR-INC-HUM-2012-CPP-35045 is acknowledged and greatly appreciated.

[†]Brown University, Email: ssn@brown.edu

[‡]University of Virginia, Email: ssekhi@virginia.edu

1 Introduction

The efficacy of public service delivery is undermined by large scale corruption in developing countries. This can disproportionately affect less powerful areas (Olken (2006)) or less powerful households (Broussard, Dercon, and Somanathan (2009)). As a result, the services often do not reach the intended beneficiaries. Monitoring and enforcement may improve take-up. In this paper, we examine whether the most vulnerable beneficiaries of public service delivery programs are differentially affected by corruption, whether informal monitoring through social networks affects outcomes, and whether improving monitoring and enforcement can yield quantitatively important welfare gains for vulnerable beneficiaries.

We focus on the distribution of subsidized goods in India under the Targeted Public Distribution System (TPDS). We explore whether the historically disadvantaged Scheduled Castes (SCs), who are much poorer than the general population and often the subject of discrimination,¹ have a lower take-up of the program when facing non-SC delivery agents. We establish that take-up is affected by the caste of the agent and provide evidence that this effect works through monitoring and enforcement. We then develop and estimate a structural model to quantify the welfare gains that SC households could obtain by having monitoring and enforcement costs lowered by an amount equivalent to their having a SC shopkeeper, rather than a non-SC shopkeeper.

India is a pertinent setting for examining these phenomena. The incidence of malnourishment among children in India is very high. By one estimate, India recently accounted for 40% of malnourished children in the world (von Braun, Ruel, and Gulati (2008)). As one of the government programs aimed at addressing malnutrition, the TPDS distributes grains and other goods through over 450,000 “Fair Price (FP) Shops,” where households are entitled to purchase rice, wheat, and other goods at below market rates from a locally appointed shopkeeper. The subsidy for grain is primarily targeted toward below-poverty-line (BPL) households, where BPL status is assigned by local elected officials. The scale of beneficiaries is massive: the most recent government Economic Survey asserted that 65.2 million BPL households are entitled to benefits from the TPDS (GoI (2011)). The financial scale of the program is equally impressive, with the food subsidy accounting for over Rs. 582 billion in FY 2009-2010 (approximately \$12.6 billion at December 2009 exchange rates), or 1.3% of GDP (GoI (2011)). However, despite this massive investment, take-up rates vary greatly geographically and are quite low in many states, leaving the promise of an entitlement to

¹The poverty rate among SCs in 2004-05 was 37%, as compared to a population-wide rate of 28% (Ahluwalia (2011)).

food unfulfilled. In part, this is because of massive leakage from the delivery system. Recent estimates for 2007-08 place the fraction of grain diverted at an astounding 43.9% across India, which is actually an improvement over previous years (Khera (2011); PEO (2005)).²

Our empirical analysis has three components. First, we use data from the state of Uttar Pradesh to examine whether SC households suffer disproportionately from this breakdown in local service delivery. We find strong evidence for this in the case of non-SC shopkeepers. Yet, take-up of grains is strikingly higher – more than 40 percentage points – for SC households (relative to non-SC households) when shopkeepers are SC. However, we do not observe the same differential in take-up of kerosene and sugar.

Second, we conduct empirical analysis to explain why the caste of the shopkeeper as well as the good in question matters for SC households. We demonstrate that the most plausible channel involves caste networks that play a role in monitoring and enforcement. Informal social networks may serve as a conduit of information and lower the cost of monitoring and enforcement. We hypothesize that when SCs face SC shopkeepers (or non-SCs face non-SC shopkeepers), the cost of monitoring and enforcement is lowered. The value of this effect matters most when the cost of monitoring and enforcement is high to begin with. In Uttar Pradesh, only a fraction of households in each village – those with a BPL card – are entitled to the grain subsidy at the time of our survey. On the other hand, since all households are entitled to the sugar and kerosene subsidy, information may flow more easily for these goods and collective punishment of a shopkeeper may be more forthcoming. Therefore, these goods will have a lower initial cost of monitoring and enforcement than grain will. In this way, the empirical patterns across different goods that we see in the data could emerge because of the different costs of monitoring and enforcement across goods and castes. We provide additional evidence supporting this interpretation, including data we recently collected in the survey villages with the specific intent of understanding the role of caste in monitoring and enforcement. Our reduced-form empirical analysis shows that our findings are not supportive of a series of potential alternative explanations, most notably taste-based discrimination against SCs

² The leakage comes in two forms, both of which are facilitated by the absence of effective monitoring and enforcement. The first involves inclusion and exclusion errors in the allocation of BPL cards by local officials. (see Niehaus, Atanassova, Bertrand, and Mullainathan (2011) for details). The second form of leakage comes from the fact that shopkeepers and agents in other portions of the supply chain have a strong incentive to sell goods on the black market (a specific example from Delhi is discussed in (Parivartan (2004))). Because shopkeepers have privileged access to the supply chain, it is difficult for citizens or elected officials to verify shopkeepers' claims of receiving only limited supplies for distribution.

and elite capture.

Finally, we develop and estimate a structural model of take-up that we use to quantify the welfare gains to SCs of improving monitoring and enforcement, as envisioned by the National Food Security Bill under consideration in India (NAC (2011)). Specifically, we examine the welfare gains to SCs of moving from the higher monitoring costs of facing a non-SC shopkeeper to the lower monitoring costs of facing a SC shopkeeper. In this exercise, we find that the average welfare gains are quantitatively important, approximately one-fifth the size of the average grain subsidy. Importantly, we also use the model to evaluate another aspect of the National Food Security Bill, increasing the effective subsidy rate. We find that, for the large majority of cases, a more generous subsidy rate can stimulate greater black marketing of goods and perversely lower take-up. This effect is greater for SCs than non-SCs.

The role of caste in transmitting information and enabling enforcement of agreements has been explored in several different contexts. Related to our work, Fisman, Paravasini, and Vig (2011) use data from a large state-owned bank in India to demonstrate convincingly that loan amounts are higher – and default rates lower – when a household and bank branch official share the same caste. Munshi and Rosenzweig (2006) examine the implications of caste-based occupational networks that channel members of the same caste into particular jobs. And in the context of local government, Munshi and Rosenzweig (2010) find evidence supporting the hypothesis that caste can serve as a disciplinary mechanism for politicians. This paper complements the previous papers by showing that not only the caste identity of the agents in a transaction matter, but the item being transacted does as well (as we found for grains, available to a subset of the local population, versus other subsidized items, available to all); moreover, our structural estimates allow us to quantify the welfare effects that result from differential access to the market and, in this case, the government subsidy in question.

Our results have immediate relevance to the TPDS. There is a lively debate in Indian policy circles about the future of the PDS, with two broad camps emerging. The first camp supports making changes to the PDS that are roughly reflected in the National Food Security Bill. The second camp of reformers supports moving toward a system based on food vouchers or outright cash transfers (Kotwal, Murugkar, and Ramaswami (2011) and Chaudhuri and Somanathan (2011)). The issue of how households induce delivery agents to properly transfer their full entitlements – present in either set of reforms – depends on both formal accountability mechanisms and informal mechanisms, as through social networks. And in either set of reforms, the welfare of the targeted beneficiaries

depends not only on whether a household takes up a public service, but on the monitoring costs the household must undertake to ensure take-up. Our estimates can be used to examine how changes in monitoring costs affect welfare, particularly for Scheduled Castes.

Our results also have broader implications for local public service delivery. A large literature has shown that the quality of service delivery may be compromised by elite capture (Bardhan and Mookherjee (2000), Bardhan and Mookherjee (2006)) or self-enrichment by authorities (Reinikka and Svensson (2004), Niehaus and Sukhtankar (2011)). Reservation of political seats or political strength within villages can still ensure that economically less powerful households enjoy public benefits (Pande (2003), Foster and Rosenzweig (2004)). Our paper shows that in the absence of strong formal checks and balances, the economically vulnerable may rely on informal social networks to strengthen accountability of public service delivery agents and ensure access to public services.

National-level monitoring and enforcement can effectively discipline local authorities in certain settings (Olken (2007)). But, this is not always forthcoming or cost-effective, and implementation often requires some accountability at the local level. Simple technological solutions, coupled with strictly enforced changes in incentive structures, can overcome obstacles of cost (Duflo, Hanna, and Ryan (2012)). However, in some cases, implementing the necessary technology or enforcing the proper incentive structures may be too expensive.³ In these cases, finding low-cost, local-level monitoring and enforcement mechanisms becomes more pressing. Bjorkman and Svensson (2009) demonstrate that community-level disciplinary actions, encouraged by NGOs, effectively improved service delivery by public health providers in the context of Uganda. Our study contributes to the understanding of how informal monitoring and enforcement mechanisms operate and what is the resulting benefit.

The rest of the paper is organized as follows. In Section 2, we discuss background on the TPDS. In Section 3, we describe the data and provide summary statistics. In Section 4, we conduct our benchmark empirical analysis of take-up and consider a number of mechanisms that could explain the results. We show evidence suggesting that monitoring and enforcement differences drive the empirical patterns. In Section 5, we develop a structural model of household demand

³For example, there has been discussion of technological improvements that force PDS shopkeepers to have a customer present in order to confirm the purchased quantity electronically, after which the information is transmitted to a higher-level bureaucrat who can compare a shop's sales to households with government sales to the shopkeeper in order to identify corruption. However, it is still possible that the shopkeeper forces households to confirm a larger quantity of goods than they actually purchased in order to receive anything. Whether the shopkeeper can do this ultimately depends on local monitoring and enforcement.

and shopkeeper supply for delivery of grain. In Section 6, we estimate the structural model and use the estimates to analyze the impact on SCs of improvements in monitoring and enforcement and increases in the effective subsidy rate, as proposed in the National Food Security Bill. Finally, Section 7 concludes.

2 The Public Distribution System

2.1 Mechanics of the Program

The Public Distribution System (PDS) is the government of India’s flagship program to combat hunger and malnutrition. The PDS is the government’s largest instrument for targeting the poor, consuming a large share of government resources. In 2002-03, program expenditures accounted for more than 5% of central government expenditures, more than twice the amount spent on education (Kochar (2005)). In the PDS, the central government procures grains from farmers and sells it to state governments at the “central issue price”. From this point, the state government is responsible for distributing grains (and potentially other select goods) to government-licensed “Fair Price Shops”, using either government agents or private agents. Eligible households are entitled to buy grains and other select items from these shops at a below market rate, where the rate for grains is set by the state government as a mark-up on the central issue price.⁴ Rice, wheat, sugar and kerosene are made available through this system, which currently has a network of over 450,000 shops across India. In 1997, the government modified the program, creating the present Targeted Public Distribution System (TPDS) which is an explicit anti-poverty strategy.

Under the TPDS, states are mandated by the central government to periodically identify below-poverty-line (BPL) families, subject to a centrally-imposed cap on the total number of BPL families in each state. The central government then allocates grains to the states based on the number of BPL families, with the central issue price for BPL allocations being roughly half that for above-poverty-line (APL) allocations (Rs 3.5/kg and Rs 2.5/kg for BPL rice and wheat, respectively). State governments can add no more than Rs 0.5 to these prices when selling to BPL households at the Fair Price Shop.

⁴The mark-up is intended to consider transportation costs, storage costs, etc. No mark-up is allowed in the more recently created sub-category of Antodaya Anna Yojana (AAY) households, a subset of BPL households.

2.2 Criticism of the Program

The TPDS has the critical objective of improving food security and nutrition, but critics have questioned its efficacy. A recent study suggests that for every Rs. 3.65 spent by the central government, only one rupee is eventually transferred to below-poverty-line (BPL) households (PEO (2005)). In 2003-04, families in Uttar Pradesh only consumed one-third of the total amount of grain purchased from the Food Corporation of India (FCI) by the state. A similar fraction applied in Madhya Pradesh, and Bihar faced an even lower fraction (Ajwad (2007)).

Despite the money spent on the PDS, malnourishment rates among Indian children and adults are very high. Families have found alternative ways to cope with food security (Tarozzi (2005)), but these ways may require much in the way of money or time. And these methods seem to be falling short. One study suggests that in 1999-2000, around 233 million people were undernourished in calorific and micro nutrient terms (Mane (2006)). Self-reported rates of hunger were also higher among SCs/STs than others, with a rate of 5% and 6% for rural SCs/STs, respectively, as compared with an overall rural rate of 3.3% (Mane (2006)).

2.3 Shopkeeper Decision-Making

Agents in the PDS supply chain face a strong incentive to black market PDS goods because of the limited monitoring and enforcement in the program. While corruption could occur anywhere in the supply chain, we focus on corruption at the level of the Fair Price Shops, which could lead to endogenous scarcity in the amount of grain available for beneficiaries.

There is a thriving black market for goods that are otherwise designated to be sold at Fair Price Shops, and widespread corruption in the system is well-documented. While it is often difficult to find evidence for corruption, the PDS requires record keeping that can be used to reveal substantial irregularities.⁵

According to a recent report involving shops in Delhi, a number of systematic discrepancies have been found in these records that suggest shopkeepers siphon off goods from these shops. On

⁵ The shop keepers are required to maintain (i) *Card Register* that lists the number, the name, and the type of card for all households in the market served by the shopkeeper; (ii) *Stock Register* that shows the delivery of stocks and total sales made; (iii) *Daily Sales Register* that reports the details of every sale including the date, card number, card holder's name, good purchased, and quantity purchased; (iv) *Inspection Book* that reports the inspection of stocks done by the food inspectors; and (v) *Cash Memo* that is used to issue receipts for daily transactions. In addition, every transaction has to be reported in the card that is issued to the beneficiary, including the date, the good transacted, and the quantity sold.

comparing the details recorded in records maintained by the shopkeepers with the cards held by the beneficiaries, several facts come to light. There are many cases where the shopkeepers' records have transaction details that do not appear in the relevant cards. The dates in the records and the cards do not match for many transactions. In some cases, the card-holding beneficiaries are not identified residents of the areas serviced by the shops. Sales are also reported for households who do not have a card issued to them. The detailed evidence of these discrepancies in various records, and copies of some of these records can be found in *Annexure* I, II, III, IV A, IV B, VIII, IX, X, and XI of Parivartan (2004).

Moreover, these irregularities are by no means unique to Delhi. Investigations by government authorities have uncovered large amounts of leakage in the system across India (see PEO (2005)). Prior work suggests that up to one-third of grains are diverted from the legal PDS channels into the black market (Mane (2006)), although recent work suggests that there have been substantial improvements in some states (Khera (2011)).

3 Data and Descriptive Statistics

We use the 1997-98 Uttar Pradesh-Bihar Survey of Living Conditions (SLC), modeled after the World Bank's Living Standard Measurement Study (LSMS) data. These data consist of household surveys and village surveys undertaken between December 1997 and March 1998, shortly after the initiation of the targeting in the TPDS. The SLC covers 120 villages in UP and Bihar, with 15 or 30 households sampled in each village. There are data on 2250 households in total. We restrict our analysis to the UP portion of the data, as implementation may have been lagging in Bihar at the time. In the primary analysis, we drop households that are Muslim and villages that have a Muslim shopkeeper. Moreover, we focus on households with BPL cards since only they are legally entitled to grain at the special below-market TPDS rates. This leaves 335 households in our main sample. We refer to the Fair Price Shops operated under the PDS as PDS shops in the rest of the paper.

The household and village surveys provide a rich set of information. Among other important characteristics, the data also contain information on whether the household is included on the village BPL list, the take-up of each of the four PDS goods (rice, wheat, kerosene, and sugar), prices paid at PDS shops, the distance to PDS shops, knowledge of entitlements, and crude self-reported measures of quality relative to market quality.⁶ From the village-level data, we use

⁶The household survey allows us to distinguish between households that are on the BPL list and households that,

information on caste composition, total number of households, the location of the PDS shop, the caste of the Pradhan, and the caste of the PDS shopkeeper. In villages without a PDS shop, we do not have any information on the PDS shopkeeper or the shop.⁷ Appendix A provides more details about the sample.

Summary statistics for the SLC are reported in Tables 1, 2, and Appendix Tables 1 and 2. Table 1 provides summary statistics for basic household characteristics used in the analysis. Approximately 60% of the sample consists of SC households, while there are very few Upper or Middle Caste (UMC) BPL cardholders. In results not shown here, we verify the results of Kriesel and Zaidi (1999), who show that even after conditioning on BPL criteria, SCs are much more likely to obtain BPL cards than other households. Almost 70% of the sample has a PDS shop in their village, and of those responding, only 42% have the shop in their tola (neighborhood). On average, households live approximately 21 minutes away from their PDS shop. The largest share of households face a Backward Caste (BC) shopkeeper, but about 25% faces a SC shopkeeper. A quarter of households face a shopkeeper of their own (broad) caste category.⁸

In Appendix Table 1, we report the quantities of PDS goods purchased. Three points are worth highlighting. First, grain purchases (and, to a lesser extent, sugar purchases) are less common than kerosene purchases. Second, grain and sugar are often deemed worse than market quality.⁹ Third, there is very little variation in prices charged, regardless of good.¹⁰ We also observe that the vast majority of households purchase the entire entitled quantities of 3 kg of rice and 7 kg of wheat in results reported in the appendix. Appendix Table 2 compares the characteristics of villages with SC shopkeepers and non-SC shopkeepers. Villages with SC shopkeepers tend to have a higher proportion of SCs in the village, are more likely to have a SC Pradhan, and are much more likely to

in addition, have their new BPL card. Not all households on the list have a card, because the process of handing out cards was in progress during the time of the survey. While it is difficult to know if it is more informative to focus on ownership of a card, rather than being on the list, we assume for the remainder of this paper that being on the list is the key attribute and having a card is more of a technicality. Kriesel and Zaidi (1999) also focus on the list, saying “In all villages we visited, the fact that a household’s name was included on the BPL list entitled it to make purchases from the PDS shop, regardless of whether or not it had actually received the card.”

⁷ If every village in a district were sampled, we could use the names and location of villages to infer information about the PDS shop frequented by every household in villages without a shop. However, due to the random sampling of villages, this is not necessarily possible in our case.

⁸We define caste categories as Scheduled Castes (SC), Backward Castes (BC), Upper and Middle Castes (UMC), and Muslim.

⁹Since kerosene is not legally sold on the market, it is not feasible to compare to market quality.

¹⁰While the standard deviations for quantities are higher, in the case of grain this is driven by extreme values.

have the largest share of households be SC. However, these differences are imprecisely estimated. An important observation is that villages with SC shopkeepers have more BPL card holders, and this difference is statistically significant at the 1% level.

In the empirical analysis, we focus on how outcomes differ when SC households face SC shopkeepers, as opposed to considering how outcomes differ by sub-caste (*jaati*). There are two reasons for this. First, in most villages there is only one dominant sub-caste from any single caste residing in a village. In Uttar Pradesh, there are only five villages – out of a total of 63 – where the first and second dominant sub-castes by population share are from the same caste. Given the limited heterogeneity within caste in each village, the social network is effectively proxied by our caste variables. Second, although we have information on household sub-caste, we do not have information on the sub-caste of shopkeepers.

Table 2 examines how SCs differ depending on whether they face a SC shopkeeper or a non-SC shopkeeper.¹¹ In the first four rows, there is some indication that SCs facing SC shopkeepers are poorer than those facing non-SC shopkeepers, but none of these differences are statistically significant. Moreover, those in SC shopkeeper villages actually appear to be more educated on average. The large differences emerge when comparing measures of SC economic or political dominance. In villages with SC shopkeepers, SC households are more likely to be the plurality caste, more likely to be the same caste as the village Pradhan, and more likely to have a SC sub-caste be the village’s largest sub-caste and land dominant sub-caste.

The remaining rows of the table speak to the households’ relationship to the PDS. The data contain information on the largest caste (by population) in the tola where the PDS shop is located. With a SC shopkeeper, SC households are much more likely to have SCs dominate the tola of the PDS shop. Moreover, SCs are more likely to be close to the PDS shop with a SC shopkeeper. Both of these facts begin to suggest that having a SC shopkeeper may confer monitoring advantages to SCs relative to other castes. The last three rows show that take-up of grain, sugar, and kerosene are all significantly higher with SC shopkeepers.

In order to better understand the role of caste in the take-up of the PDS, we conducted a follow-up survey of 300 households in the SLC villages in 2012, with the sample roughly divided in half between SC households and non-SC households. Figure 2 demonstrates that the tendency for friends, relatives, and fellow caste members to obtain better treatment from the shopkeeper has persisted over time. The top panel shows that, when respondents are asked to name which house-

¹¹Villages with Muslim shopkeepers and without information on shopkeeper caste are dropped.

holds face few problems with the PDS, and then rank their responses, the majority of respondents indicate friends and relatives. A significant number of households also report that members of the shopkeeper’s caste (jaati) do not face many problems (bottom two bars). Since it is likely that friends and relatives of the shopkeeper are of the same broad caste category as the shopkeeper, the bottom panel of the figure aggregates responses for the first pick of each respondent. There, we see that friends and relatives of the shopkeeper and members of the shopkeeper’s caste are selected by more than half the sample. Below, we use this follow-up survey to inform our interpretation of the empirical results from the SLC.

4 Empirical Analysis

We saw in Table 2 that take-up of all goods is higher for SCs when they face a SC shopkeeper. In this section, we conduct our benchmark analysis and use the emergent results to distinguish between competing models that could explain the influence of shopkeeper caste on SCs’ take-up of PDS goods. In comparing SCs who face SC shopkeepers versus those who face non-SC shopkeepers, the main empirical concern for interpretation is that the caste of the shopkeeper may be endogenously chosen. For instance, in villages where SCs are politically pivotal, leaders may direct lucrative bureaucratic posts (e.g. PDS shopkeeper) and available PDS goods toward SCs.¹² The 72nd amendment of the constitution reserved positions of village council head for SCs in 1992. In Uttar Pradesh, for the same villages, PDS shopkeeping was also reserved for SCs. Uttar Pradesh also had village council elections before the data were collected so that the council heads in villages reserved for SCs were officiating. While we do not have data on the tenure of the PDS shopkeepers, we could potentially use the determinants of reservation of council seats for SCs as instruments to address the endogeneity problem. As Anderson (2011) shows, the SC population in these villages has been historically very stable. In Table 2 above, we do observe that the plurality of SC population, which was the criterion used to reserve council seats, is positively related with having an SC shopkeeper. Also, having an SC Pradhan (or council head) is positively correlated with having a SC shopkeeper.

We do not use plurality of SC population as an instrument because in our application, the population share of SCs can have an independent effect on take-up. Villages with a larger presence of SCs may face higher monitoring and enforcement costs, if higher-level bureaucrats pay less attention to these villages. Queues may be longer in these villages, since SCs tend to have greater access to BPL cards and may have greater reliance on public services regardless. Thus, we control for

¹²We explore these and other mechanisms in the section on alternative channels below.

share of SC population in our structural relationship for take-up rather than using an instrumental variable approach.

In providing an explanation for our results that relies on informal monitoring and enforcement, our identifying assumption is that conditional on population of SCs, caste of the village Pradhan and the village's largest and land dominant sub-caste, the variation in the village PDS shopkeeper's caste is plausibly exogenous to unobservable components of households' preferences and costs of PDS access. We also explicitly address several endogeneity concerns in the discussion below, and cast doubt on alternative explanations for our results.

We illustrate our main findings in Figure 1. For each good – grains (rice or wheat), sugar, and kerosene – we graph the proportion of households that purchase the good in the prior month. The top panel refers to SC households, and the bottom panel refers to non-SC households. For each good, we show take-up when a household's PDS shopkeeper is SC and when the shopkeeper is not SC. We restrict ourselves to BPL households (who are legally entitled to grains at below market rates). It is clear that the take-up of grain is higher among SCs when they face an SC shopkeeper, increasing from about 49% with a non-SC shopkeeper to about 76% with a SC shopkeeper. The same is not true for non-SC households, who actually have lower take-up when facing a SC shopkeeper. SCs' differential benefit from SC shopkeepers is positive but smaller for sugar, and less than 4 percentage points for kerosene.¹³

We formalize these results with the following empirical specifications for take-up of grain, sugar, and kerosene.

$$Y_{iv} = \alpha + \beta S_{iv} + \gamma B_{iv} + \eta S_{iv} B_{iv} + \epsilon_{iv} \quad (1)$$

Where Y_{iv} is an index for take-up of the subsidized goods (grain , sugar or kerosene). S_v is an indicator which takes value 1 if the shopkeeper of the village belongs to a scheduled caste, B_{iv} is an indicator which equals to 1 if the buyer household is a Scheduled Caste household, and ϵ_{iv} is the error term. The coefficient η on the interaction term is the parameter of interest. We refer to this as the *caste pairing effect* or *caste interaction effect*. If we were to use a linear probability model, η would be a difference-in-difference coefficient. However, since take-up is a binary variable,

¹³These results do not use survey weights. Muslim households and households with Muslim shopkeepers are not included in order to focus on the Hindu caste pairings. In a SURE regression of a dummy for take-up of each good on a SC dummy, a SC shopkeeper dummy, and the interaction of the two dummies, the interaction effect is positive and significant for grain. While the effect is statistically indistinguishable from the interaction effect for sugar, it is statistically different from the interaction effect for kerosene.

we instead use a probit model to estimate Equation 1.¹⁴ Because of our non-linear specification, η cannot be interpreted as the usual difference-in-difference parameter. Therefore, we explicitly calculate and report the following difference for every specification we run:

$$\begin{aligned}\Delta &= [Pr(Y = 1|keepSC, SC) - Pr(Y = 1|keepNotSC, SC)] \\ &\quad - [Pr(Y = 1|keepSC, NotSC) - Pr(Y = 1|keepNotSC, NotSC)]\end{aligned}$$

where the first term equals the take-up of SC buyers when they face an SC shopkeeper minus their take-up when they face a non-SC shopkeeper, and the second term equals take-up of non-SC buyers with a SC shopkeeper minus their take-up with a non-SC shopkeeper. That is, Δ indicates SC households' differential gain from having a SC shopkeeper.

We carry out several robustness tests below. A robust finding is that, once controls are included, the caste interaction effect is positive and statistically significant for grain, but not sugar and kerosene. We provide evidence that caste-based monitoring/enforcement can explain the empirical patterns we observe. We consider a number of alternative hypotheses that could explain these results. In a variety of tests of these hypotheses, we do not find empirical support for these alternate channels.¹⁵

4.1 Results and Robustness Checks: Take-Up

Table 3 formalizes the results from Figure 1 and tests their robustness to adding controls and making methodological changes that account for a large class of omitted variables bias concerns. We estimate a number of specifications of Equation (1). Each panel contains empirical models for take-up of the three goods (grain, sugar, and kerosene), restricting the sample to BPL households in villages that have a PDS shop. Within each panel, the first two columns are standard probit models, the third column is a random effects probit model with random effects at the village level, and the fourth column is a fixed effects logit model with fixed effects at the SC Shopkeeper-SC Pradhan-Survey Stratum level. For the probit models, we include the implied estimate of Δ at the

¹⁴We have also run linear probability models, but we prefer the probit specifications because average take-up of kerosene is quite high, so that predicted values may be more likely to fall outside the unit interval.

¹⁵ Throughout, we display results that use survey weights and do not include Muslim households or villages with Muslim shopkeepers. However, in each section we comment briefly on how the results change when weights are dropped or Muslim households/shopkeepers are included.

bottom of each column.¹⁶ We highlight the suitability and limitations of each approach in the following discussion of our results.

Columns 1, 4, and 7 of Table 3 show the results for the most basic specification with only a SC buyer household dummy, the SC shopkeeper dummy, and the interaction of the two (the caste pairing effect). Standard errors are clustered at the village level. For both grains and sugar, the caste pairing effect is positive and significantly different from zero; for kerosene, it is not.¹⁷ This implies that for grain and sugar, SC buyers benefit more than non-SCs from having a SC shopkeeper. The bottom row shows the estimates of Δ , which suggest that the SC buyers' differential benefit from having a SC shopkeeper is very large for grain, approximately 54 percentage points.

Several endogeneity concerns might arise in interpreting this as the effect of having a SC shopkeeper on take-up. First, omitted household characteristics can be confounding the results. One possibility is that in areas where SCs are the poorest – and consequently most dependent on the PDS for food – they mobilize to ensure that there is a SC shopkeeper. In other words, wealth determines food take-up and is correlated with the interaction of interest. Second, caste-based residential segregation results in pairings of SC households with SC shopkeepers such that the distance to the PDS shop is reduced. Third, pairings of SC households with SC shopkeepers may be occurring in places where SCs are more politically powerful. In these examples, it is plausible that kerosene take-up is not positively associated with the omitted variable. In the first case, kerosene may be a luxury good relative to food items. In the second case, PDS entitlements of grain and sugar may be more difficult to carry than kerosene. And in the third case, SCs may have a differentially higher demand for food products because of their greater poverty.

Columns 2, 5, and 8 of Table 3 address such concerns. In addition to the caste variables, these specifications controls for a rich set of characteristics that might be correlated with the interaction term. These include land holdings, type of housing structure, education dummies, a dummy variable indicating whether the household is less than 0.5 km from the PDS shop, a dummy variable for whether the household's caste forms a plurality in the village, a dummy for whether the household is the same caste as the Pradhan, a dummy for whether the household's caste has a sub-caste that

¹⁶Note that here the other covariates matter for the size of the effect because of the non-linear models. In specifications including other covariates besides the household and shopkeeper caste, we evaluate Δ for a household with zero land, no assets, no pucca housing, primary schooling, more than 0.5 km from the PDS shop, not the same caste as the Pradhan, not the plurality caste, not a land dominant caste, and with the average share of SCs in the village.

¹⁷When Muslims are added with weights, the grain effect remains significant at the 1% level. Without weighting, the grain effect is positive and significant with a P-value of 0.02 (without Muslims) or 0.052 (with Muslims). For sugar, the effect becomes insignificant at the 10% level when weights are removed or when Muslims are added.

has the highest land ownership in the village, and the share of households in the village that are SC. For grain, we continue to find a positive interaction effect that is significant at the 1% level. As indicated in the bottom row, Δ remains similar, at around 55 percentage points. For sugar, the P-value of the interaction effect increases to 0.111 and Δ falls.¹⁸ For kerosene, the caste pairing effect is negative and insignificant at the 10% level.¹⁹

Another concern might be that take-up is strongly correlated among households within villages, perhaps because all households within a village see the same queues at the PDS shop, or they see the same market prices for goods. If villages with many SCs and an SC shopkeeper just happen to have higher market prices for grain in our sample (even though this may not be the case in the population), we could be erroneously attributing the effect to the caste pairing. We address this type of concern in Columns 3, 7, and 11 of Table 3. These columns use random effects probit models of take-up with random effects at the village level. The village level random effects would control for village specific heterogeneity that may account for the correlations described above. The results are robust to controlling for village level random effects. For grain, we see that the coefficient estimate for the interaction term remains positive and significant at the 1% level, though Δ falls slightly to about 52 percentage points. For sugar, the caste interaction effect is not significantly different from zero, and the same is true for kerosene.²⁰ The limitation of random effects models is that the random effect is assumed to be uncorrelated with the included covariates.

Another class of concerns is that SC shopkeepers are chosen for unobservable reasons involving village leaders' political desire to ensure households access to PDS grain. Since BPL cards are directed toward SC households as discussed earlier, it may be natural to direct the bureaucratic post serving BPL cardholders to a SC. To address this, we use fixed effects logit models with fixed effects at the SC shopkeeper-SC Pradhan-Survey stratum level. The results in Columns 4, 8, and 12 confirm the earlier results: The caste pairing effect is significantly different from zero at the 1% level for grain, but the effect is not statistically distinguishable from zero at conventional levels for

¹⁸With Muslims and/or without weights, the grain effect is still significant at the 5% level or better. For sugar, the interaction effect is not significantly different from zero at the 10% level in the case without weights and in the cases with Muslims.

¹⁹ Without weighting (and with or without Muslims) the effect is also statistically indistinguishable from zero for kerosene.

²⁰For grain, the P-value is 0.038 without weights. With Muslims added, the P-value is 0.093 (without weights) or 0.022 (with weights). The conclusions for sugar and kerosene are not affected by weights or the inclusion of Muslims.

sugar or kerosene.²¹²²

A final concern might be that the purchases of grain, sugar, and kerosene are not being treated jointly above. Many households may purchase these goods simultaneously.²³ Appendix Table 3 addresses this by estimating a model of take-up of all three goods that allows for the unobserved errors for each good to be correlated with one another. The basic patterns continue to hold up. For grain, the caste interaction effect remains positive and significant (still at the 1% level), and the magnitude of Δ is in the same range as before. The bottom line of the table presents the P-values from the test that Δ for each good is different from Δ for grain. The marginal effect for grain is significantly different from that for kerosene, though not from that of sugar. Nevertheless, the most important point is that even when we allow for the errors to be correlated across goods, the caste interaction effects are not significantly different from zero for sugar or kerosene.²⁴

These tables beg the question: Why is the take-up of grains larger for SC buyers when they face SC shopkeepers relative to non-SC shopkeepers, whereas having an SC shopkeeper does not affect take-up of grain in the same way for other castes? In the subsequent sections, we explore various possible mechanisms and provide evidence to show that caste plays out as a monitoring and enforcement device.

4.2 Role of Monitoring and Enforcement

The patterns observed in our data are consistent with a hypothesis that caste for SC households serves as an informal monitoring and enforcement mechanism. In this sub-section, we substantiate this hypothesis. In the following sub-section, we also explore a variety of other mechanisms that

²¹Without weights, the P-value for grain rises to just 0.060 (without Muslims) or 0.102 (with Muslims). With Muslims and weights included, the effect is significant at the 1% level. The conclusions for sugar and kerosene are not affected by weighting or the inclusion of Muslims.

²²We have tried fixed effects logit specifications with village fixed effects, and the caste pairing effect is no longer significant at conventional levels in this case. However, we find these specifications unreliable because they remove a substantial portion of the data, 10 villages in total. In particular this removes 10 of the 38 households that are SC with SC shopkeepers (7 from one village, and 3 from another), and all 10 of these households purchase grain. It also removes 32 of the 108 SC households that face a non-SC shopkeeper (spread over 8 villages), and only 9 of these households purchase PDS grain.

²³This may not always be the case, however. In one village observed by the authors in Uttar Pradesh, the PDS shop sold kerosene on a particular day without selling other goods that day.

²⁴Without weights, the grain caste interaction effect is significant at the 1% level, the sugar effect is significant at the 10% level, and the kerosene effect is insignificant. When Muslims are included, only the grain effect is significant at the 10% level (P-value of 0.020 without weights, P-value of 0.004 with weights).

could explain why we see a robust caste pairing effect for grain, but not one for sugar and kerosene. None of these alternate explanations is compelling.

To see why, consider the actionable evidence that a household would need to conclude that shopkeeper corruption had deprived it of goods. Specifically, the household needs to establish that it was denied its entitlement of a good when that good was available to the shopkeeper. Three sources of information could provide this “actionable evidence”: (1) a household’s own observation of goods being available but the household being denied beforehand; (2) members of the social network observing availability of and/or obtaining goods but the household being denied later; and (3) government sources describing delivery to the shopkeeper but the shopkeeper still reporting stock outs.

All three sources of information are costly. In the absence of a reliable public source of information about a shopkeeper’s malfeasance, a household must necessarily go to the shop and be denied goods itself in order to establish malfeasance. This removes the ability to free ride on the information gathering of others. Using a supplemental survey that we conducted in October 2012 in the SLC villages, we will demonstrate below that households do indeed expend effort and visit the shops. But the household can supplement its own efforts at direct observation with information acquired through its social network or directly from government sources.

A monitoring and enforcement mechanism can explain the fact that SC households benefit differentially from SC shopkeepers for grain, but not for sugar and kerosene. If a household has a shopkeeper of the same caste, then monitoring the shopkeeper to obtain “actionable evidence” about his performance is less costly for the household because the caste network can facilitate the acquisition of information. A household having an own caste shopkeeper can translate into a large, discernible monitoring/enforcement advantage relative to households of other castes if two conditions hold: 1) Information is important because many in the household’s caste care about the good in question; 2) Information is not easily acquired through alternative sources besides the caste network.

Condition 2 is relevant to grain, but not sugar and kerosene. This is because only a subset of the village – BPL card holders – are entitled to the steep subsidies for grain, whereas everyone is entitled to sugar and kerosene. Condition 1 is more relevant to SCs than other castes because BPL cards were handed out more generously to SC households than other households. Table 4 shows probit models for whether or not a household is on the BPL list. The first column is a standard probit, while the second includes village-level random effects. Both specifications include

controls for the four BPL eligibility criteria at the time, as described in the note under the table (Coefficients not shown). The key result is that even after controlling for the eligibility criteria, SCs are more likely to have BPL cards, whether or not the village leader (Pradhan) is a SC. Therefore, this monitoring/enforcement mechanism can explain the caste pairing effect for SCs in the case of grain, but not other goods.

We test this mechanism with a series of additional descriptive results in Table 5. This table augments the basic specification for grain take-up in a variety of ways to test specific implications of this mechanism. In doing so, the specifications use the same set of controls as the full set of covariates in Table 3, but we only report the coefficient estimates for the relevant subset.

First, we turn to columns 1 and 2 of Table 5. The monitoring/enforcement hypothesis suggests that having a higher proportion of households that are eligible for a good – and thus seeking information about it – should lower monitoring costs for everyone and thereby increase take-up. We examine this issue by including the fraction of sampled households in the village who hold BPL cards, since only BPL card holders are entitled to subsidized grains.²⁵ Column 1 uses the same sample as above, while Column 2 incorporates data from households that do not have a PDS shop in their village. In the latter case, we include a dummy variable that indicates if there is no PDS shop in the village, and if there is no shop, we set the value of the SC shopkeeper dummy to zero.

The results show that a higher fraction of BPL cardholders is associated with significantly higher take-up of grain.²⁶ Of course, a higher fraction of BPL households could also simply reflect better public service delivery more generally. To address this issue, we estimate analogous specifications for sugar and kerosene. We find that the estimated BPL share is never significantly different from zero for sugar. It is only significantly different from zero and positive for kerosene if weights are included, and even this is not true when Muslims are not included and all villages are used.

A second test of our posited mechanism involves distance from the PDS shop. If households enjoy low monitoring costs for grain to begin with by being close to the PDS shop, they are likely to benefit less from the caste interaction effect. Columns 3 and 4 of Table 5 explore this hypothesis. Column 3 interacts the SC shopkeeper-SC household interaction with the dummy variable for being within 0.5 km of the shop. Limited variation prevents us from employing the fully interacted model.

²⁵We are estimating the true share of BPL households using just 15 or 30 households, so that there is potentially sampling error in this measure.

²⁶This conclusion is robust to dropping weights, in which case the effect is statistically significantly different from zero at the 5% level (with baseline sample) or 10% level (with villages without a PDS shop). The effect is no longer statistically significant at conventional levels if Muslims are included.

Column 4 uses time it takes to go home from the PDS shop instead of distance, and includes the full set of interactions between the SC household dummy, the SC shopkeeper dummy, and time to the shop. In all columns, we additionally include the estimated share of BPL households in the village.

The results bear out the prediction. In Column 3 we see that households that are further than 0.5 km from the PDS shop have a significantly larger caste interaction effect. The same is not true for sugar.²⁷ In Column 4 we see that as households get further from the PDS shop in terms of time, the caste interaction effect becomes larger. We can evaluate the size and statistical significance of the caste interaction effect at the minimum, median, and maximum time to the PDS shop. Hypothesis tests using these times show that the caste interaction effect is not significantly different from zero at 2 minutes, but is significantly different from zero at 15 minutes (P-value of 0.001), and 70 minutes (P-value of 0.087), with larger magnitudes at the larger times.²⁸ The same is not true for sugar or kerosene, with the only exception being a positive caste pairing effect for sugar take-up at 70 minutes, but only when weights are included.

A third test of our posited mechanism involves the ease of information sharing. As the population share of SCs in a village increases, communicating with SCs becomes easier, making monitoring less costly. Therefore, we should observe larger caste pairing effects when the SC share is high, and small or non-existent caste pairing effects when the SC share is low. Column 5 of Table 5 shows the relevant results. There, we include a full set of interactions between the SC household dummy, the SC shopkeeper dummy, and the population share of SCs in the village. The results show that the caste interaction effect is indeed larger as the SC share rises. We can evaluate the caste interaction effect at the minimum, median, and maximum values of the SC share in the sample. The effect is negative and insignificant at a share of 0.04, positive and significant at 0.35 (P-value of 0.009), and positive and significant at 0.7 (P-value of 0.025).²⁹ Identification of the model is tenuous with kerosene, but we can estimate the model reliably for sugar, where the only positive and significant effects occur at a share of 0.7 when weights are used; otherwise, the same patterns as for grain

²⁷These results do not depend on weighting or the inclusion of Muslims. Unfortunately, lack of variation in the data prevents us from estimating the model for kerosene.

²⁸If weights are dropped, the caste pairing effect is significant at 15 minutes, but not 2 or 70 minutes. If Muslims are included, then regardless of weighting the caste interaction effect is not significant at 2 minutes, but becomes significant at conventional levels for 15 minutes and 70 minutes.

²⁹Without weights, a similar pattern holds, but the P-value at a share of 0.7 rises to 0.1099. When Muslims are included, the caste pairing effects are no longer statistically significant without weights, and the effect is significant at a share of 0.35 with weights.

do not arise. To summarize this sub-section, our results strongly suggest that informal monitoring along caste lines must be playing a role in PDS take-up.³⁰

Thus far, we have used results from the SLC to bolster our hypothesis that monitoring and enforcement are driving the caste interaction effects. Using the follow-up survey we conducted in the SLC villages in 2012, we can speak a bit more directly to the hypotheses. The situation in 2012 is not directly comparable to that in 1997; currently, the UP government has in place a program that fixes delivery dates for each village, so that the shopkeeper must be open on a particular set of days each month. Whether this actually occurs depends on the incentives for delivery date auditors to collude with the shopkeeper, but we show that some respondents expect fixed delivery dates by 2012.

Despite this difference, the follow-up survey reveals a few key points. First, the caste network is an important source of information about the delivery of goods. Figure 3 breaks the sample into SC and non-SC households, and asks respondents about their sources of information regarding delivery dates. It is evident from the results that a substantial fraction of respondents are in villages where fixed delivery dates have taken hold. At the same time, it is quite common to hear of delivery dates from a messenger from the shopkeeper and from neighbors of one's own caste (jaati).

Second, it is clear that the caste (jaati) of the shopkeeper matters for how information circulates and enforcement happens. Figure 4 shows this in two panels. The top panel asks the follow-up survey respondents whether information about the PDS is more accessible when the shopkeeper is of the same caste as them. Overwhelmingly, respondents confirm that the caste of the shopkeeper matters for information. The bottom panel asks whether who respondents speak to regarding concerns about the PDS depends on the caste of the shopkeeper. Again, respondents make clear that the shopkeeper's caste matters. We provide detail regarding who households approach with concerns in the top panel of Figure 5. There, households demonstrate that they discuss concerns heavily with own caste neighbors, friends and relatives. The bottom panel of the figure asks what would happen to shopkeepers if the shopkeeper refused to sell goods that are clearly present. The

³⁰ Enforcement may be part of the story, but we cannot separately identify the relative roles of monitoring and enforcement with our data. Shopkeeper decision-making will depend on the expected punishment for not serving an individual properly, and the expected punishment involves both monitoring (the probability of detection) and enforcement (the punishment, conditional on detection). Consequently, any attempt to disentangle monitoring and enforcement while using just information on grain take-up will depend almost entirely on functional form assumptions. Therefore, we do not attempt to distinguish clearly between monitoring and enforcement here. This is a very promising avenue for future research.

sample responding to this question is quite small since the fixed delivery dates solve some of the issues with this form of corruption. But the responses suggest that sanctions from the shopkeeper’s own caste form an important disciplinary mechanism. This is consistent with empirical results from the 1997 SLC. In Column 5 of Table 5, we show that the caste pairing effect is actually strongest for households with larger landholdings, households that are more likely to be able to provide the shopkeeper assistance.

Third, these key features of monitoring and enforcement do not seem to depend in important ways on whether a village has a SC shopkeeper or not. Figure 1 in the appendix has two panels. The top panel divides the survey respondents into two categories based on whether they are in villages with SC shopkeepers in 1997 or in villages with non-SC shopkeepers in 1997. We then examine what fraction of respondents believe that information about the PDS is more accessible if the shopkeeper is of the same jaati as the household. The graph demonstrates that the proportion of households who believe this is quite similar across the caste of the shopkeeper. The bottom panel of the figure performs the same exercise, except this time dividing the sample by whether or not the respondent has a SC shopkeeper at the time of the survey in 2012. While a higher fraction of those living in villages with a SC shopkeeper believe that jaati matters for information, the differences are not huge.

We can do a similar exercise for enforcement. Figure 2 of the appendix examines the consequences for the shopkeeper if the household has made a complaint. The top panel divides respondents into SC shopkeeper and non-SC shopkeeper villages by the 1997 shopkeeper caste, while the bottom panel uses the 2012 caste. In either case, we see similar patterns across SC shopkeeper and non-SC shopkeeper villages. A majority of respondents consistently report that a consequence is that the shopkeeper will not receive assistance from his jaati.

Therefore, there is suggestive evidence that caste facilitates both monitoring and enforcement in the survey villages, leading to the caste pairing effect we observe in the 1997 SLC data. Our posited mechanism relies critically on the fact that SC households were disproportionately allocated BPL cards in our 1997 survey villages, regardless of shopkeeper caste. If our explanation is true, we would expect that the caste pairing effect for grain take-up should be sensitive to the allocation of BPL cards. We examine this issue in Appendix Table 5 using household survey data for Uttar Pradesh from the latest round of the Rural Economic and Development Survey (REDS). These data were collected in 2008. We construct dependent and control variables in a manner analogous to our definitions for the 1997 SLC, insofar as possible. The first column of the table shows a probit model

for whether a household has a BPL card. The remaining columns show probit models for whether a household purchases any grain and any sugar (data on kerosene were not available). In all cases, we include village-level random effects. Column 1 shows that, while SCs in non-SC shopkeeper villages are more likely to obtain a BPL card than their non-SC counterparts, this is not true for SCs in SC shopkeeper villages. The caste interaction coefficient is negative and significant, and the size is large enough that one cannot reject the hypothesis that SCs in SC shopkeeper villages have the same likelihood of obtaining a BPL card as non-SCs in non-SC shopkeeper villages. Based on these results, and our monitoring and enforcement hypothesis, we should expect the caste pairing effect for grain take-up to be smaller or zero. The results in Columns 2 and 3 demonstrate that there is no significant caste pairing effect for grain or sugar.

4.3 Alternative Channels

We discuss a number of alternative channels that could potentially explain our basic results, but find the evidence for these channels less convincing than the evidence for the monitoring/enforcement channel.

4.3.1 Elite Capture

Given previous work in India, a compelling possibility is that elite capture is occurring: In villages where SCs are strong, a SC shopkeeper is selected and goods are channeled toward SCs; when SCs are weak, a non-SC shopkeeper is selected and goods are channeled toward non-SCs. The implication of this elite capture explanation is not that members of all castes in SC shopkeeper villages have a higher probability of take-up,³¹ but rather that SC shopkeepers direct goods to SC households to satisfy political pressures.

Given the results above, this explanation seems unlikely. We control for proxies of household political or economic power in the village including plurality of caste, and land dominance. Moreover, PDS grain is widely considered an inferior good, whereas sugar and kerosene may not be. The results in Table 3 pan this out: the coefficient on whether households have a pucca house (better quality) is insignificant for grain, but positive and significant for sugar and kerosene; and the coefficient on land in the equation for grain take-up is negative and significant. Thus, if elites were to capture a PDS good, it would likely be sugar or kerosene, and we have already seen that the caste interaction effect is not positive and significant for sugar and kerosene.

³¹This can be accounted for in part by using the random effects probit model.

The first two columns of Table 6 test this explanation in an alternative way, by including interactions of the SC household dummy with proxies for SC political power (whether or not the Pradhan is SC and share of the village that is SC). We discuss these tests, which do not show evidence that elite capture alone is driving our results, in the appendix.

4.3.2 Taste-Based Discrimination

Another potential explanation is that SC shopkeepers favor SCs and non-SC shopkeepers favor non-SCs in a simple case of taste-based discrimination. However, this by itself is insufficient to explain the patterns we are seeing in the data above. Taste-based discrimination by shopkeepers – with no other differences across castes – should result in positive caste pairing effects for all goods, not just grain. Even looking just at the raw data, it is clear that transactions between SC households and non-SC shopkeepers do occur. Of the SCs with BPL cards in villages with a PDS shop and a non-SC shopkeeper, 58 did not purchase grain. Of these, 36 still purchased kerosene and 17 still purchased sugar from the same shopkeeper that they did not purchase grain from. These facts are not consistent with a pure taste-based discrimination story.

4.3.3 Selective BPL Allocation

Another possibility is that in villages with SC (non-SC) shopkeepers, local authorities are better able to identify which SC (non-SC) households will make use of BPL grain benefits. Since we are conditioning on BPL status, this means that a positive caste pairing effect could emerge purely because of selection into BPL status. Villages run by a SC Pradhan may have better information about which SCs will use benefits and this may tend to occur in the same places where there are SC shopkeepers.

In order to address this, we add controls for whether or not the Pradhan is SC, and an interaction of this dummy with the SC household dummy in Column 1 of Table 6. The results show that the caste pairing effect continues to be significant. Moreover, the estimated effect of having a SC Pradhan on SCs is actually negative (though neither the SC Pradhan dummy or the interaction with the SC dummy are significantly different from zero).³² If selective BPL allocation were the explanation, we would expect the coefficient to be positive. An additional test that rules out this explanation is discussed in the appendix.

³²The same is true when Muslims are added or weights are not included.

4.3.4 Stigma

Since PDS grain is a relatively undesirable good, it may be the case that SCs face stigma when purchasing grain from non-SCs (and vice versa). This could explain the positive caste pairing effect for grain, but not for sugar and kerosene.

To explore this channel, we use a more detailed set of caste pairings in Column 3 of Table 6. The omitted category consists of UMC households. The notable observation is that the BC shopkeeper/BC household coefficient is not statistically significantly larger than the UMC shopkeeper/BC household coefficient in any specification. At the same time, the SC shopkeeper/SC household coefficient is statistically significantly larger than the UMC shopkeeper/SC household coefficient in every specification, with the magnitude of this difference larger than the gap between SC shopkeepers and BC shopkeepers for SC households. So for the stigma story to be true, it must be the case that SCs face stigma most intensely with UMC shopkeepers, and that BCs do not face stigma with UMC shopkeepers. We find this unlikely.³³

4.3.5 Credit Provision

Another possibility is that SCs are liquidity constrained and cannot come up with the cash necessary to purchase PDS goods. SC shopkeepers may alleviate this constraint by providing credit, with the caste ties substituting for the security of collateral. If this issue primarily affects the lowest income SCs, then we may not see the same patterns for sugar and kerosene, since these may be in less demand by the lowest income households. We estimate a random effects probit specification in which our dependent variable is whether or not households received credit for food from the shopkeeper. We find a positive and significant caste pairing effect for credit, so that SCs differentially obtain more credit for food from SC shopkeepers.³⁴ This lends support to the premise of this explanation.

Therefore, we assess whether credit is driving grain take-up differences. In Column 4 of Table 6, we run the same model for grain take-up as in Column 3 of Table 3, but now we control for whether or not the household has received credit to buy food from the shopkeeper. The caste pairing effect remains statistically significantly different from zero at the 1% level and Δ continues to be large,

³³Additionally, the results indicate that the BC shopkeeper/BC household coefficient is statistically indistinguishable from the SC shopkeeper/BC household coefficient, regardless of the use of weights or inclusion of Muslims. This is also inconsistent with the stigma explanation.

³⁴The effect is not statistically significant without weights, but is significant when Muslims are included (with and without weights).

at 0.507.³⁵ Therefore, it appears unlikely that credit is driving the caste pairing effect.

4.3.6 Knowledge of Entitlements

The SLC was administered only shortly after the PDS was revamped into the TPDS. For this reason, it is possible that many households are simply not aware of their grain entitlements at this point in time. If SCs are better informed in villages with SC shopkeepers – either because information moves more freely or because villages that select SC shopkeepers are more likely to have leaders that disseminate information about entitlements to SCs – then the caste pairing effect could emerge for grain, but not for sugar and kerosene.

In results not shown here, we find support for the premise of this argument. We estimate random effects probit models identical to Column 3 of Table 3, except with a dummy variable for whether or not the household correctly identifies the rice or wheat entitlement as the dependent variable.³⁶ We find that the SC shopkeeper-SC household interaction is positive and significant for rice and for wheat.³⁷

To explore how this is related to the caste pairing effect for grain take-up, we estimate the same models for take-up as before, but now include dummy variables controlling for whether or not the household is aware of the rice entitlement and the wheat entitlement. The results appear in Column 5 of Table 6. There, we see that take-up tends to be higher among those who correctly identify the entitlement. The caste pairing effect remains statistically significant, with a P-value of 0.005, and the size of Δ continues to be large.³⁸ This suggests that information about entitlements cannot explain the patterns we are seeing.

4.3.7 Household Production

It is possible that SC shopkeepers are chosen in villages where SCs have particularly limited home production opportunities. Even if SCs are not in charge of the village, the village leadership may find it politically useful to grant requests for a SC shopkeeper by SCs. To better understand the

³⁵Without weights, the P-value rises to 0.054. When Muslims are included, the P-value rises to 0.040 (with weights) or 0.127 (without weights).

³⁶Households that did not visit the PDS shop in the last 30 days were not asked these questions, so these households are dropped.

³⁷For rice, this is no longer true if weights are not used. For wheat, this is no longer true in the one case when Muslims are included and weights are not included.

³⁸Without weights, the P-value rises to 0.046. With Muslims, the P-value becomes 0.066 (without weights) or 0.024 (with weights).

issue of home production, Table 7 regresses a series of variables related to home production on the full set of covariates from Table 3. These variables are: the amount of grain (rice+wheat) consumed out of home production or received as in-kind payment (these two sources cannot be distinguished in the data); whether or not the primary source of household income is own farm activities; land planted with grain (rice or wheat); self-evaluated price of land per acre; and percent of land irrigated. We present only the household and shopkeeper caste coefficient estimates for the sake of brevity.

The estimate of the interaction effects in Column 1 suggest that SCs have especially low consumption of home produced or in-kind grain when they face a SC shopkeeper.³⁹ We observe SCs consuming less out of home production in villages with SC shopkeepers, but it is difficult to know whether this is because better access to PDS grain induces them to consume less out of home production, or because lower home production opportunities induce them to purchase more PDS grain.

Therefore, we examine characteristics of SCs' farming opportunities that are likely to be exogenous to PDS take-up. In Column 2, we see that SCs have particularly low rates of deriving income from own farm activities when facing SC shopkeepers.⁴⁰ On the other hand, Columns 3-5 show no evidence that SCs plant less acreage with grain, have lower land quality (as measured by price), or have lower irrigation access in villages with SC shopkeepers.⁴¹ Therefore, there is some evidence that SCs who face SC shopkeepers specialize less in own farm production, but no evidence that they have poorer land or less ability to produce grain. The specialization decision may instead be driven by factors unrelated to land quality.

In light of this, one potential explanation of our take-up results is that SC shopkeepers were chosen in areas where SCs did not routinely specialize in agriculture. To examine whether differences in involvement in own farm activities are driving our results for take-up, we present a random effects probit model for take-up including the own farm activities dummy variable as an additional control.⁴² If limited household production opportunities are driving the results, we would expect

³⁹The effect is only statistically significantly different from zero at conventional levels when weights are used.

⁴⁰The effect is only statistically significant at conventional levels when weights are used. The magnitudes are also very sensitive to the use of weights, going down sharply when weights are not used.

⁴¹This does not depend on the use of weights or the inclusion of Muslims.

⁴²We also estimate models that additionally include the amount of home-produced/in-kind receipt of grain as a covariate. Such a specification is difficult to interpret as discussed above. The quantity of home produced/in-kind grain consumed is a choice of the household, and could be caused by better access to the PDS. Nevertheless, the caste pairing effect remains statistically significantly different from zero, with a P-value of 0.003 and an implied value

the caste pairing effect to become statistically insignificant and the coefficient on the “own-farm” dummy to be negative and significant. The results appear in the last column of Table 6 and do not corroborate this story. The caste pairing effect remains positive and statistically significant with a P-value of 0.003, though the size of Δ falls to 0.496. Moreover, the coefficient on own farm activities is not statistically different from zero at conventional significance levels.⁴³ This suggests that the caste pairing effect is not simply being driven by selective placement of SC shopkeepers in places where SCs have limited household production opportunities.

4.4 Summary

To sum up, the alternate channels we have examined are inconsistent with the empirical evidence. We believe that only caste-based monitoring and enforcement can plausibly explain the evidence we have seen thus far. In the next section, we develop and estimate a structural model in order to quantify the welfare impact of this informal monitoring and enforcement device.

5 Structural Model and Estimation Procedure

The first sub-section describes the model, and the second presents the likelihood function used in maximum likelihood estimation.

5.1 Model

We model the household-shopkeeper interaction as a simple two-stage game. In the first stage of the game, the household chooses whether or not to buy grain from the PDS shop and how much effort to exert in monitoring the shop. As described in the section on monitoring and enforcement earlier in the paper, we assume that monitoring effort takes a very specific form. In order to find actionable evidence of shopkeeper malfeasance, the household needs to establish that a good is available after the household itself is denied the good. A household can establish the availability of the good in three ways: (1) a household’s own observation of goods being purchased by other households in the time following a household’s own attempted purchase (or direct observation of sufficient

of Δ of 0.484. Without weights, the P-value rises to 0.039. With Muslims, the P-value is 0.100 (without weights) or 0.022 (with weights). The amount of home produced/in-kind grain never carries a coefficient that is statistically significantly different from zero.

⁴³Without weights, the P-value of the caste pairing effect becomes 0.036. With Muslims, the P-value is 0.094 (without weights) and 0.022 (with weights). The coefficient on own farm activities is never significantly different from zero at conventional significance levels.

stocks being delivered to the shopkeeper initially); (2) members of the social network observing availability of goods in the time following a household’s attempted purchase; and (3) government sources describing the timing and amount of delivery to the shopkeeper. More monitoring effort translates into a higher probability that actionable evidence of shopkeeper malfeasance is available, and therefore a higher probability that the shopkeeper will be disciplined if the household is not served a good.

In the second stage of the game, the shopkeeper observes the chosen level of monitoring effort and the household decision of whether to purchase grain, and then decides whether to sell to the household. In making this decision, the shopkeeper trades off the gain from black marketing the good with the expected punishment of doing so. While the expected punishment may involve the revoking of a shopkeeper’s license, there is reason to believe that district and village officials generally use other mechanisms, including increases in requested bribes.⁴⁴ Our model focuses on these other mechanisms of punishment.

Beyond the structural assumptions about household utility functions and monitoring technology, the model uses two key simplifications: (1) The model considers only grain provision; and (2) Shopkeepers provide the full grain entitlement if they serve a household, and households demand the full grain entitlement if they purchase any PDS grain. The first simplification is due to our narrow goals for this exercise and data limitations; we leave a more general model for future work. The second simplification is motivated by a clear regularity in the data, already noted above.

In what follows, we first set out the household and shopkeeper objective functions. Then, we discuss the solution of this simple game.

5.1.1 Households’ Utility Maximization

Let x_m denote the quantity of grain purchased from the market, and x_s denote the quantity purchased from the PDS shop. D is a dummy variable taking a value of 1 if $x_s > 0$, and 0 otherwise. All non-grain expenditures are denoted by q . The price of grain on the market and in the PDS store, respectively, are given by p_m and p_s , and y denotes income. An additional net utility gain of purchasing PDS grain is κ . The household’s choice of monitoring effort is M , and the cost of monitoring is c . Total costs of making a purchase are given by c_T , and monitoring costs are one component of this. Finally, α , β , and ζ are parameters of the utility function and γ is the

⁴⁴Kriesel and Zaidi (1999) find that the only recorded instance of a license being revoked in their area was for a shopkeeper who claimed he was being punished for not black marketing enough goods.

entitled amount of grain.

The household solves the following optimization problem:

$$\max_{q, x_m, x_s, D, M} (x_m + \beta x_s - \zeta)^\alpha q^{1-\alpha} + (\kappa - c_T)D$$

subject to:

$$q + p_m x_m + p_s x_s = y$$

$$q, x_m, x_s \geq 0$$

$$x_s \leq \gamma$$

That is, PDS grain and market grain are assumed to be substitutes, with $\beta < 1$ meaning that PDS grain is of inferior quality. Under the assumptions that $\beta p_m > p_s$ and that households purchase at least γ kg of grain overall, households choose $x_s = \gamma$ if they make any PDS purchase. We maintain this assumption for everything that follows.⁴⁵

The term c_T plays an important role in the model. It captures monitoring costs, the cost of going to the shop and retrieving goods, any stigma from obtaining goods, and the cost of waiting in a queue or ensuring sufficient cash on hand to purchase goods. It is assumed to be heterogeneous across households, with a distribution and mean parameterized below.

5.1.2 Shopkeeper Profit Maximization

Shopkeepers must balance the incentive to sell a good on the black market with the expected cost of turning away beneficiaries that would have purchased that good. This cost comes in the form of a fine – this can be thought of as a bribe paid to village/district officials to retain the position or as remuneration to households in the form of properly delivered goods in the future.

The market and PDS price are p_m and p_s as before, and p_a is the price at which the shopkeeper must purchase the goods from the state government (Assume $p_a < p_s$). Let $\delta_i(M) \in [0, 1)$ be the probability that the shopkeeper is reprimanded from turning away household i , where $\delta_i(M)$ is increasing in monitoring effort M . Define f to be the fine conditional on being caught. Then we

⁴⁵Strictly speaking, it is possible that a household purchases no market grain and less than or equal to γ kg of PDS grain. One household in our data reports PDS purchases of grain and does not report any grain through the market or through home production/in-kind receipt. This household consumes the PDS entitlement for rice and wheat, and has missing values for market purchases and home production. Given the rarity of this event, we believe it is not driving the subsequent estimation results.

assume that the shopkeeper chooses whether to sell to household i , S_i , by solving the following maximization problem:

$$\max_{S_i \in \{0,1\}} p_m \gamma (1 - S_i) + p_s \gamma S_i - p_a \gamma - \delta_i(M)(1 - S_i)f$$

5.1.3 Subgame Perfect Equilibrium

To solve the two-stage game for the unique subgame perfect equilibrium, we proceed recursively. First, we solve the shopkeeper's problem and then use this solution to solve for the household's choice in the first stage. If the shopkeeper sells to household i , he receives $\gamma(p_s - p_a)$; if he instead black markets this quantity, he receives $\gamma(p_m - p_a) - \delta_i(M)f$ in expectation. Therefore, the shopkeeper is willing to sell household i 's good through the PDS channel if and only if:

$$\gamma(p_m - p_a) - \delta_i(M)f \leq \gamma(p_s - p_a) \quad (2)$$

If this does not hold, the shopkeeper instead sells the household's allocation on the black market.

Next, consider the household's decision, taking the shopkeeper's strategy as given. The household realizes that the minimum value of M necessary to make the shopkeeper willing to sell through the PDS is given by setting the left-hand side equal to the right-hand side in Inequality 2 above. Denote this minimum value M^* . The household's take-up decision is therefore between undertaking monitoring effort M^* and securing the PDS grain, or undertaking no monitoring effort and not buying PDS grain.

More formally, define $V_1(p_m, p_s, y, M^*)$ to be the household's indirect utility when $x_s = \gamma$ and $V_0(p_m, y)$ to be the indirect utility when $x_s = 0$. The household optimizes by choosing $D_i = 1$ if $V_1 - V_0 > 0$, and $D_i = 0$ otherwise. We assume that $\delta_i(M) = 1 - e^{-M}$. Then one can show that $M^* = \log \left[\frac{f}{f - \gamma(p_m - p_s)} \right]$ and:

$$V_1 - V_0 = \alpha^\alpha (1 - \alpha)^{1-\alpha} p_m^{-\alpha} [p_m \beta \gamma - p_s \gamma] + \kappa - c_T(M^*)$$

and that the household's choice of quantity of non-PDS grain is given by:

$$x_m = \alpha \frac{y}{p_m} + (1 - \alpha)\zeta - D \left[(1 - \alpha)\beta \gamma + \alpha \gamma \frac{p_s}{p_m} \right]$$

5.2 Likelihood Function and Estimation Procedure

We first note what outcomes we observe, and then we parameterize the model and make distributional assumptions. For each household, we observe whether or not the household has purchased

PDS grain in the past 30 days and we observe a potentially mis-measured estimate of monthly consumption of non-PDS grain. To construct this estimate, we take the product of (1/12) and the reported annual consumption of grain from outside purchases and from home production or in-kind receipts. There are reports of zero outside purchases by people who claim to have gone to the PDS shop for grain. This suggests that the outside purchases do not always include PDS purchases for the purposes of the survey, but we cannot know this for sure.⁴⁶ Denote our estimated consumption of non-PDS grain as x_m^* , whereas the true consumption in the past 30 days is x_m .

Let $G_i = 1$ if we observe the household purchase PDS grain, and $G_i = 0$ otherwise. We model the joint distribution of $(G_i, \log(x_{mi}^*))$, which is substantially simplified by the independence assumption below. We assume that $x_m^* = x_m e^{\theta G + \eta}$, where η is classical measurement error. This simple form allows for the survey reports of market consumption to mean different things on average for people who go to the PDS shop versus people who do not.

Next, we require distributional assumptions for the sources of error above. We assume η is distributed as $N(0, \sigma^2)$ and that c_T has an exponential distribution with mean $\frac{1}{\lambda}$ that depends on monitoring effort, monitoring costs, and other costs as specified below. We assume that c_T and η are independent.

Finally, we need to parameterize κ and $\frac{1}{\lambda}$. We assume:

$$\begin{aligned}\kappa &= \kappa_0 \log(1 + Land) + \kappa_1 SC + \kappa_2 SCShopkeep \\ c &= \exp(c_0(1 - SC)(1 - SCShopkeep) + c_1 SC(1 - SCShopkeep) + c_2(1 - SC)SCShopkeep \\ &\quad + c_3 SC * SCShopkeep + c_6 ShareSC + c_7 ShareBC)(Time)^{c_4} (ShareBPL)^{-c_5} \\ \frac{1}{\lambda} &= \delta_0 Far + c * \log \left[\frac{f}{f - \gamma(p_m - p_s)} \right]\end{aligned}$$

Here, *Far* indicates the household is further than 0.5 km from the shop, *Land* is land ownership in acres, *SC* is a dummy for the household being SC, *SCShopkeep* is a dummy for the shopkeeper being SC, *ShareBPL* is the estimated share of villagers who are BPL cardholders, *ShareSC* is the share of village households who are SC, and *ShareBC* is the share of village households who are BC. The market price and PDS price are denoted as before.

There are two important notes to make regarding the included covariates. First, we include the SC dummy and the SC shopkeeper dummy in κ to allow for greater demand by SC households and

⁴⁶For simplicity, we assume that the separation property holds, so that consumption out of home production is treated just like consumption from the private market.

the selection of SC shopkeepers in villages that are poorer (and have greater demand). Second, the role of the caste composition variables is to capture the fact that district officials may be more responsive to complaints from villages that have a particular caste composition, or that information about the PDS may flow differently depending on a village's caste composition (because, e.g., UMC households are generally without BPL cards and therefore not invested in PDS grain).

We now set out the likelihood function used for estimation, leaving the conditioning on covariates implicit. The likelihood function used in estimation is:

$$L = \prod_{i=1}^N [Pr(G_i = 1)f(\log(x_{mi}^*)|G_i = 1)]^{G_i} [Pr(G_i = 0)f(\log(x_{mi}^*)|G_i = 0)]^{1-G_i}$$

The conditional densities and the probability are given by the expressions:

$$\begin{aligned} f(\log(x_{mi}^*)|G_i) &= \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-1}{2\sigma^2}[\log(x_{mi}^*) - \log(A(G_i)) - \theta_{G_i}]^2} \\ Pr(G_i = 1) &= 1 - \exp(-\lambda B) \end{aligned}$$

where

$$\begin{aligned} A(G_i) &= \frac{\alpha y}{p_m} + (1 - \alpha)\zeta - G_i \left[(1 - \alpha)\beta\gamma + \alpha\gamma \frac{p_s}{p_m} \right] \\ B &= \alpha^\alpha (1 - \alpha)^{1-\alpha} p_m^{-\alpha} [p_m\beta\gamma - p_s\gamma] + \kappa \end{aligned}$$

We maximize the likelihood function using data from BPL households with a shop in their village. We do not include households with a Muslim shopkeeper or Muslim households. In total, this leaves us with 226 observations that have non-missing data for all covariates.

6 Structural Estimates, Validation, and Implications

The first sub-section below presents the estimation results. In the second sub-section, we examine the fit of the model to the Uttar Pradesh SLC survey data, and also conduct an out-of-sample validation exercise using SLC data from the neighboring state of Bihar. In the third sub-section, we use the estimates to quantify the welfare impact of caste-based monitoring and enforcement, as well as discuss the potential impact of expanding the generosity of the program.

6.1 Estimation Results

The estimation results appear in Table 8. In the first pass, we do not include standard errors. Each panel involves a different category of parameters, ranging from preferences to the fine used to

punish shopkeepers. First, we discuss preferences. The estimates suggest that the share parameter α is approximately 0.24, and that grain consumption of around 24 kg is the baseline that households meet. The estimates of β suggest that PDS grain is a close substitute to market grain, but slightly inferior. Those with greater land ownership have less utility from consuming PDS grain, all else equal. This should be expected, as land owners may be less reliant on outside purchases of grain in general. In the second column, we also see that SC households are estimated to have a higher utility gain from PDS consumption; the estimate is large, on the order of 70% of the average utility return (the indirect utility gain, minus the cost term). Finally, SC shopkeepers may be placed in areas with higher values from PDS goods. This is consistent with the fact that SC shopkeepers seemed to be in villages where more BPL cards were allocated in the section on descriptive statistics.

In terms of costs, monitoring costs are estimated to increase with time and decrease with the share of the village that has a BPL card. Villages with a higher share of SC population or BC population have higher monitoring costs. SCs face higher monitoring costs with non-SC shopkeepers, and non-SCs face higher costs with SC shopkeepers.

The remaining parameters relate to the fine and measurement error. The estimated fine is approximately Rs. 40. This should be compared to a value of entitlement that is on the order of Rs. 25, though the exact size depends on the village. The estimate of θ indicates that our measure of grain consumption from the survey systematically over-estimates actual market purchases for those who purchased PDS grain in the last month. This makes sense, if we think that some households choose to report their PDS purchases as market purchases.

6.2 Model Validation- Out of Sample Test

Next, we examine the fit of the model. We examine market grain consumption and PDS grain take-up in each of four cells defined by whether the household is SC and whether the shopkeeper is SC. The top portion of each figure presents the fit for UP. It is not surprising that the patterns fit the UP data well as the estimates were chosen to fit the UP data.

We perform an analogous exercise for Bihar as an out of sample check and report the results in the bottom panel. Bihar is a very different setting from UP, and take-up tends to be strikingly lower in Bihar. To account for this, we conduct the out-of-sample checks as follows: We increase the fine by 20% to ensure that it is large enough that households can choose some effort level to make the shopkeeper comply with regulations; and we add to the mean cost of take-up, a number that is sufficient to closely match the observed overall take-up of 25.8% (within 0.3 percentage points). We

leave all other components of the model unchanged. If the model accurately captures differences across SC household-SC shopkeeper combinations, we consider this to provide some faith in the validity of the model.

We use the first set of estimates and begin with Figure 6. The figure shows two columns for each cell; the first column shows the average of our measure of actual log grain consumption, while the second column shows the expected value of log grain consumption using the model’s parameter estimates and household covariates. In the top panel, the model predicts these averages quite well for UP. There are some differences for Bihar, but overall we miss the Bihar average by less than 5%. Next, we turn to Figure 7. Within each cell, the first column shows the share of households making a PDS purchase in the last month, while the second column shows the average of the predicted probabilities from using the model’s estimates. Again, the fit for UP looks reasonable. The fit for Bihar is surprisingly good, within 5-10 percentage points, though the model seems to over-predict take-up in places with SC shopkeepers. It is difficult to know if this is because the context for SC shopkeepers was different in Bihar or because the model is deficient, but this is worth investigating in the future.

6.3 Welfare Implications

Finally, in this sub-section we use the model to assess the welfare implications for SCs of lowering the cost of monitoring and enforcement, part of the reforms considered in the National Food Security Bill. We need to choose an amount by which to change monitoring costs. Specifically, we assess the implications of changing from a non-SC shopkeeper to a SC shopkeeper, allowing only an effect through the monitoring cost channel. The move to a SC shopkeeper changes the return to purchasing PDS grain; however, we ignore this portion in this exercise because this reflects selection of SC shopkeepers in particular areas rather than monitoring costs.

To calculate welfare consequences, we use an equivalent surplus concept in the following procedure. First, we randomly draw a cost for each household 100 times, sampling from an exponential distribution with mean given by the estimated parameters and a household’s covariates. Second, for each of the 100 draws, we calculate the household’s indirect utility when (A) purchasing grain from a SC shopkeeper; (B) purchasing grain from a non-SC shopkeeper; (C) not purchasing PDS grain. Importantly, the random component of cost is kept constant across (A) and (B) for any given household. Third, we calculate the amount of income that would have to be given to the household to leave it indifferent when the caste of the shopkeeper is changed from non-SC to SC. Note that

this number will be greater than or equal to zero. It will be identically zero for households that do not purchase PDS grain when facing either type of shopkeeper. Fourth, we calculate two summary measures of the 100 draws for each household, one that finds the mean equivalent surplus across the 100 draws and one that takes the 75th percentile of equivalent surplus across the draws.

Figures 8 and 9 show the results in two histograms, with the top one corresponding to the mean of equivalent surplus across the 100 draws and the bottom one corresponding to the 75th percentile of equivalent surplus across the 100 draws. In each case, the height of each bar corresponds to the percent of observations with values in the given bin. The weighted average of the distribution in Figure 8 is Rs. 4.64 per month, with a standard deviation of Rs. 1.47. Figure 9, by looking at the 75th percentile across draws for each household, makes clear that the welfare gains can be significantly higher for any given household, given a particularly high draw of random cost. The weighted average of this distribution is Rs. 7.73 per month, with a maximum value of over Rs. 13.⁴⁷ While these numbers are numerically small, we should keep in mind that the average subsidy amount (difference between market price and PDS price, multiplied by entitled quantity) is around Rs. 25. Therefore, for a large number of SC households, the welfare gains from reduced monitoring costs can be a significant fraction of the subsidy amount.

The National Food Security Bill envisions making the TPDS more generous. Given that we have estimated the structural model, it is possible to see what the effect would be of increasing the subsidy amount by reducing the PDS price further, and how this differs by shopkeeper caste and household caste. We use our estimates to calculate the semi-elasticity of the probability of take-up with respect to the PDS price and display percentiles of this semi-elasticity for each caste pairing in Table 9.

Two facts stand out. First, a large portion of the distribution (or all of the distribution) for each caste pairing is positive. This means that as the PDS price falls (i.e. the size of the subsidy increases), the take-up probability falls. This occurs because the increased incentive for black-marketing outweighs the increased incentive for households to demand goods. The size of the semi-elasticities are relatively large. For instance, 0.515 is the median value for the SC Shopkeeper/SC Household pairing. This means that as the PDS price decreases by about 10% (i.e. $\log(p_s)$ changes by 0.1), the take-up probability decreases by about 5 percentage points. The second key fact is that

⁴⁷Note that a fraction of households have a 75th percentile value of welfare gains that is identically zero, while no households have a mean value of welfare gains that is identically zero. Households with a 75th percentile value of welfare gains that are zero therefore have higher percentiles that are non-zero, leading to the non-zero mean for all households.

SCs are more negatively affected than non-SCs; one can see this by simply comparing across the caste pairings for any given percentile. This makes clear that simply expanding the PDS subsidy without concomitant increases in monitoring and enforcement can actually lead to lower take-up. In this case, households that continue to make PDS purchases may be better off than before, but households that no longer make purchases are worse off.

7 Conclusion

Subsidized food programs form an important aspect of domestic social programs and foreign aid programs across the world. While the most vulnerable beneficiaries may be prioritized by national governments or foreign aid donors, these beneficiaries may also be disproportionately affected by corruption. A pressing question concerns the availability and impact of monitoring/enforcement mechanisms in the case of these vulnerable populations. In this paper, we examine this question in the specific setting of India’s Targeted Public Distribution System, using data collected shortly after the introduction of the targeted program in 1997. First, we establish that caste networks may serve as a surrogate monitoring/enforcement mechanism that can ensure access to grain for a historically vulnerable population, SC households. Second, we use this insight to help build and estimate a model of household demand and shopkeeper supply that features the cost of monitoring/enforcement. Third, we use the model to quantify the welfare effects of lowering monitoring and enforcement costs for SC households, as well as the effects of expanding the generosity of the program. In this exercise, we find that relatively important welfare gains – approximately one-fifth the size of the total grain subsidy on average – result from lowering monitoring/enforcement costs by an amount equivalent to the difference between facing a non-SC shopkeeper and facing a SC shopkeeper. But expanding the generosity of the program can perversely lower take-up, with mixed implications for welfare. These results have implications for India’s National Food Security Bill, which proposes PDS reforms that increase monitoring and enforcement while increasing the generosity of the program.

Our work has several limitations. First, we do not have explicit exogenous variation in monitoring or enforcement, so we use variation in the number of beneficiaries and caste networks to isolate the mechanisms. Thus, we are able to provide only indirect evidence by refuting competing explanations for the caste pairing effect. Second, we have posited a very simple model of household demand. More complicated models may allow for larger increases in household demand as a result of the subsidy expansion, making the negative effects of the expansion on take-up more negligible.

Third, in the absence of any current nationally representative data which has fine details about caste of the households, PDS purchases, and PDS shopkeeper's caste, we use data that is from a particularly poor section of India in 1997. Fourth, our data cannot pin down whether corruption is happening at the level of the shopkeeper or if, instead, it is primarily higher in the supply chain, in which case the shopkeeper would act an intermediary conduit of information in the social network. Addressing some of these concerns are important avenues of future work.

We view this paper as an effort to illustrate the potential for local informal mechanisms to provide significant monitoring/enforcement, a first attempt at understanding how the National Food Security Bill may affect vulnerable populations in India, and a vehicle with which to think about the design of subsidized food programs more generally when corruption is a threat.

References

- AHLUWALIA, M. S. (2011): “Prospects and Policy Challenges in the Twelfth Plan,” *Economic and Political Weekly*, 46(21), 88–105.
- AJWAD, M. I. (2007): “Performance of Social Safety Net Programs in Uttar Pradesh,” *Social Protection Discussion Paper*, 0714, The World Bank.
- ANDERSON, S. (2011): “Caste as an Impediment to Trade,” *American Economic Journal: Applied Economics*, 3, 239–263.
- BARDHAN, P., AND D. MOOKHERJEE (2000): “Capture and Governance at Local and National Levels,” *American Economic Review Papers and Proceedings*, 90(2), 135–139.
- (2006): “Decentralisation and Accountability in Infrastructure Delivery in Developing Countries,” *The Economic Journal*, 116, 101–127.
- BJORKMAN, M., AND J. SVENSSON (2009): “Power to the People: Evidence from a Randomized Field Experiment on Community-Based Monitoring in Uganda,” *Quarterly Journal of Economics*, 124, 735–769.
- BROUSSARD, N., S. DERCON, AND R. SOMANATHAN (2009): “Aid and Agency in Africa: Explaining Food Disbursements Across Ethiopian Households, 1994-2004,” Working Paper.
- CHAUDHURI, A. R., AND E. SOMANATHAN (2011): “Impact of Biometric Identification-Based Transfers,” *Economic and Political Weekly*, 46(21), 77–80.
- DUFLO, E., R. N. HANNA, AND S. P. RYAN (2012): “Incentives Work: Getting Teachers to Come to School,” *American Economic Review*, *Forthcoming*.
- FISMAN, R., D. PARAVASINI, AND V. VIG (2011): “Social Proximity and Loan Outcomes: Evidence from an Indian Bank,” Working Paper.
- FOSTER, A., AND M. ROSENZWEIG (2004): “Democratization and the Distribution of Local Public Goods in a Poor Rural Economy,” Working Paper.
- GOI (2011): “Economic Survey 2010-2011,” Discussion paper, Government of India.
- KHERA, R. (2011): “Trends in Diversion of Grain from the Public Distribution System,” *Economic and Political Weekly*, 46(21), 106–114.

- KOCHAR, A. (2005): “Can Targeted Food Programs Improve Nutrition? An Empirical Analysis of India’s Public Distribution System,” *Economic Development and Cultural Change*, 54(1), 203–235.
- KOTWAL, A., M. MURUGKAR, AND B. RAMASWAMI (2011): “PDS Forever?,” *Economic and Political Weekly*, 46(21), 72–76.
- KRIESEL, S., AND S. ZAIDI (1999): “The Targeted Public Distribution System in Uttar Pradesh, India - An Evaluation,” Working Paper.
- MANE, R. (2006): “Targeting the Poor or Poor Targeting: A Case for Strengthening the Public Distribution System of India,” *Journal of Asian and African Studies*, 41(4), 299–317.
- MUNSHI, K., AND M. ROSENZWEIG (2006): “Traditional Institutions Meet the Modern World: Caste, Gender, and Schooling Choice in a Globalizing Economy,” *American Economic Review*, 96, 1225–1252.
- (2010): “Networks, Commitment, and Competence: Caste in Indian Local Politics,” Working Paper.
- NAC (2011): “Note on the Draft National Food Security Bill,” Discussion paper, National Advisory Council, Accessed June 22, 2011 at <http://www.nac.nic.in/foodsecurity/nfsb.pdf>.
- NIEHAUS, P., A. ATANASSOVA, M. BERTRAND, AND S. MULLAINATHAN (2011): “Targeting with Agents,” Working Paper.
- NIEHAUS, P., AND S. SUKHTANKAR (2011): “Corruption Dynamics: The Golden Goose Effect,” Working Paper.
- OLKEN, B. A. (2006): “Corruption and the costs of redistribution: Micro evidence from Indonesia,” *Journal of Public Economics*, 90, 853–870.
- (2007): “Monitoring Corruption: Evidence from a Field Experiment in Indonesia,” *Journal of Political Economy*, 115(2), 200–249.
- PANDE, R. (2003): “Can Mandated Political Representation Increase Policy Influence for Disadvantaged Minorities? Theory and Evidence from India,” *American Economic Review*, 93(4), 1132–1151.

- PARIVARTAN (2004): “Report of Jansunwai on Public Distribution System,” Discussion paper, Parivartan.
- PEO (2005): “Performance Evaluation of Targeted Public Distribution System (TPDS),” Discussion paper, Programme Evaluation Organization, Planning Commission, Government of India.
- REINIKKA, R., AND J. SVENSSON (2004): “Local Capture: Evidence from a Central Government Transfer Program in Uganda,” *Quart*, May, 679–705.
- TAROZZI, A. (2005): “The Indian Public Distribution System as provider of food security: Evidence from child nutrition in Andhra Pradesh,” *European Economic Review*, 49, 1305–1330.
- VON BRAUN, J., M. RUEL, AND A. GULATI (2008): “Accelerating Progress Toward Reducing Childhood Malnutrition in India,” Discussion paper, International Food Policy Research Institute.

Table 1: Summary Statistics: Household Characteristics

	N	Mean	S.D.	Min	Max
HOUSEHOLD CHARACTERISTICS					
Land (Acres)	335	0.94	1.40	0	10
≥ 1 Asset	335	0.11	0.31	0	1
Pucca House	335	0.29	0.45	0	1
Expenditures	334	1476.72	996.98	250	7683.667
HH Size	333	5.77	2.72	1	17
No Schooling	334	0.23	0.42	0	1
Primary or Less	334	0.38	0.49	0	1
SC	335	0.60	0.49	0	1
BC	335	0.34	0.48	0	1
UMC	335	0.06	0.24	0	1
HOUSEHOLDS' PDS SHOP					
PDS in Village	335	0.68	0.47	0	1
PDS in Tola	252	0.42	0.49	0	1
PDS 0-0.5 km	334	0.44	0.50	0	1
Time to PDS (Min)	334	21.22	15.17	2	90
SC Keeper	229	0.25	0.44	0	1
BC Keeper	229	0.44	0.50	0	1
UMC Keeper	229	0.31	0.46	0	1
HH/Keep Same Caste	335	0.25	0.43	0	1
Credit from Keeper	334	0.13	0.33	0	1

Table 2: Summary Statistics: SC Households' Characteristics by Shopkeeper Caste Group

Variable	Non-SC Keeper	SC Keeper	P-Value	Stat. Sig?
Land (Acres)	0.90	0.82	0.803	
≥ 1 Asset	0.08	0.03	0.131	
Pucca House	0.25	0.32	0.453	
Expenditures	1448.52	1265.95	0.304	
HH Size	5.55	5.79	0.674	
No School	0.30	0.24	0.475	
Prim School	0.42	0.32	0.267	
Plurality Caste	0.36	0.68	0.001	***
HH-Pradhan Same Caste	0.43	0.71	0.002	***
Caste of Village's Largest Sub-Caste	0.42	0.89	0	***
Caste of Land Dominant Sub-Caste	0.00	0.18	0.006	***
Caste Same as Largest in PDS Tola	0.23	0.92	0	***
PDS 0-0.5 km	0.54	0.84	0	***
Knows Rice Entitlement	0.38	0.53	0.128	
Knows Wheat Entitlement	0.48	0.61	0.178	
Bought PDS Grain	0.49	0.76	0.002	***
Bought PDS Sugar	0.50	0.66	0.09	*
Bought PDS Kerosene	0.77	0.95	0.001	***

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

Table 3: Choice to Purchase PDS Goods

Covariates	Grain				Sugar			Kerosene				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
SC	-0.099 (0.281)	0.451 (0.411)	0.438 (0.361)	0.460* (0.278)	-0.100 (0.268)	0.198 (0.296)	0.259 (0.299)	0.277 (0.414)	0.176 (0.379)	0.702 (0.445)	1.207** (0.557)	2.184* (1.164)
SC Shopkeeper	-0.628** (0.306)	-0.695** (0.315)	-0.416 (0.482)		-0.694* (0.371)	-0.790* (0.418)	-0.662* (0.375)		1.269** (0.529)	1.744*** (0.676)	2.517* (1.306)	
SC Keeper X SC	1.578*** (0.490)	1.743*** (0.552)	1.758*** (0.580)	1.999*** (0.154)	1.044** (0.497)	0.811 (0.509)	0.583 (0.471)	0.446 (0.370)	-0.300 (0.466)	-0.884 (0.576)	-0.709 (1.464)	-1.428 (1.629)
Land (Acres)		-0.154** (0.077)	-0.193** (0.084)	-0.190*** (0.070)		0.000 (0.080)	0.044 (0.073)	0.047 (0.110)		-0.024 (0.152)	0.087 (0.163)	0.044 (0.445)
Pucca House		0.129 (0.243)	0.052 (0.256)	0.015 (0.418)		0.546** (0.231)	0.494** (0.216)	0.963*** (0.364)		1.606*** (0.390)	1.975** (0.888)	3.730* (2.095)
No Schooling		-0.648* (0.332)	-0.730** (0.300)	-1.287** (0.625)		-0.543 (0.348)	-0.454* (0.259)	-0.817** (0.399)		-1.013** (0.405)	-1.020* (0.573)	-1.905 (1.261)
Primary or Less		0.157 (0.275)	0.372 (0.286)	0.441** (0.203)		-0.267 (0.226)	-0.138 (0.240)	-0.262 (0.250)		-0.844** (0.403)	-0.701 (0.519)	-1.513 (1.132)
PDS Dist < 5km		0.759*** (0.285)	0.963*** (0.294)	1.580*** (0.470)		0.468* (0.279)	0.407* (0.234)	0.819 (0.613)		0.202 (0.382)	-0.594 (0.548)	0.125 (1.034)
HH's Caste Forms Plurality		0.632* (0.363)	0.768** (0.341)	1.018* (0.535)		0.257 (0.233)	0.238 (0.258)	0.536 (0.648)		0.575*** (0.241)	0.707 (0.496)	1.196* (0.645)
Pradhan-HH Same Caste		-0.386 (0.323)	-0.447 (0.319)	-0.608*** (0.222)		-0.248 (0.178)	-0.244 (0.272)	-0.495* (0.263)		-0.484** (0.232)	-0.457 (0.519)	0.260 (0.830)
Caste w/ Land Dominant Sub-Caste		-0.073 (0.413)	-0.277 (0.373)	-1.270*** (0.458)		0.248 (0.308)	0.309 (0.306)	-0.027 (0.523)		-0.217 (0.439)	-0.394 (0.645)	-1.134 (1.465)
SC VillPop Share		-2.822*** (1.005)	-2.946*** (1.073)	-5.066*** (1.300)		-0.560 (0.741)	-0.378 (0.747)	-1.264 (0.994)		-1.905 (1.197)	-2.171 (1.856)	-1.639 (2.528)
Obs	229	228	228	221	229	228	228	228	229	228	228	196
Groups			34				34				34	
Correlation			0.302				0.143				0.545	
Δ	.543***	.550***	.519***		.402**	.231	.172		-.048	-.268	-.198	

Note: Dependent variable is a dummy variable for whether or not the relevant good was purchased from the PDS shop. Probit models (first two columns in each panel), village random effects probit models (third column in each panel) and fixed effects logit (SC shopkeeper-SC Pradhan-Survey Stratum fixed effect, fourth column in each panel). Constants not shown. Basic probit models allow for village-level clustering in standard errors. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 4: Factors Affecting BPL Card Allocation

Covariates	(1)	(2)
SC	0.906*** (0.170)	1.049*** (0.160)
SC Shopkeeper	0.516** (0.209)	0.602** (0.241)
SC Keeper X SC	-0.014 (0.353)	-0.099 (0.311)
Pradhan SC	-0.545*** (0.200)	-0.631* (0.335)
Pradhan SC X SC	0.769*** (0.288)	0.841** (0.359)
Obs	732	732
Groups		37
Correlation		0.154
Method	Probit	RE Probit

Note: Dependent variable is a dummy variable equal to 1 if household is on the BPL list. Additional controls include four dummy variables that capture the BPL eligibility criteria: whether the household has more land than the BPL cutoff, whether it has more income than the cutoff, whether it has a BPL-excluding asset, and whether it lives in a pucca or semi-pucca house. Random effects probit models include random effects at the village level. ; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 5: Choice to Purchase Grain: Role of Monitoring and Enforcement

Covariates	(1)	(2)	(3)	(4)	(5)	(6)
SC	0.559 (0.358)	0.864*** (0.308)	0.594 (0.362)	0.005 (0.544)	0.159 (0.739)	0.694 (0.424)
SC Shopkeeper	-1.068** (0.476)	-0.702 (0.446)	-1.247*** (0.476)	0.080 (0.873)	-0.184 (0.998)	-0.154 (0.641)
SC Keeper X SC	1.793*** (0.538)	1.418*** (0.535)	3.780*** (1.027)	0.728 (0.974)	-0.952 (1.349)	0.585 (0.756)
Share BPL in Village	4.105*** (1.497)	2.384** (0.942)	4.992*** (1.551)	2.371* (1.420)	4.612*** (1.467)	4.339*** (1.516)
No PDS in Vill.		-1.260*** (0.296)				
PDS Dist <.5km	1.150*** (0.288)	0.754*** (0.248)	1.516*** (0.338)		1.218*** (0.286)	1.221*** (0.291)
SC Keeper X SC X PDS Dist <.5km			-2.303** (0.965)			
Time to PDS (Min.)				-0.042 (0.028)		
SC X Time				0.022 (0.030)		
SC Keeper X Time				-0.070 (0.062)		
SC Keeper X SC X Time to PDS				0.100 (0.076)		
SC VillPop Share	-3.606*** (1.019)	-3.604*** (0.782)	-4.028*** (1.041)	-3.099*** (1.019)	-4.670*** (1.779)	-3.628*** (1.037)
SC X SC VillPop Share					1.123 (2.248)	
SC Keeper X ShareSC					-3.818 (4.157)	
SC Keeper X SC X ShareSC					8.191* (4.640)	
Land (Acres)	-0.169** (0.079)	-0.153** (0.078)	-0.170** (0.077)	-0.120 (0.079)	-0.105 (0.076)	-0.161 (0.156)
SC X Land						-0.027 (0.196)
SC Keeper X Land						-1.162** (0.555)
SC Keeper X SC X Land						1.308** (0.603)
Obs	228	333	228	228	228	228
Groups	34	55	34	34	34	34
Correlation	0.193	0.196	0.174	0.210	0.141	0.181

Note: Dependent variable is a dummy variable equal to 1 if household purchases grain from the PDS shop. All covariates from Columns 3, 7, and 11 of baseline table are included here as well, but not shown. Village random effects probit models ; ***

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

Table 6: Alternative Explanations for the Grain Purchase Decision

Covariates	EliteCap 1	EliteCap 2	Stigma	Credit	Knowledge	Home Prod
SC	0.547 (0.353)	-0.946 (0.601)		0.470 (0.365)	0.781 (0.562)	0.435 (0.362)
SC Shopkeeper	-0.528 (0.466)	-0.099 (0.507)		-0.388 (0.488)	-0.544 (0.641)	-0.378 (0.482)
SC Keeper X SC	1.757*** (0.593)	1.422** (0.604)	2.106*** (0.745)	1.660*** (0.596)	2.336*** (0.832)	1.747*** (0.584)
Pradhan SC	0.160 (0.759)					
Pradhan SC X SC	-0.117 (0.745)					
SC VillPop Share	-3.067*** (1.104)	-4.918*** (1.639)	-3.417*** (1.242)	-2.945*** (1.094)	-4.855*** (1.701)	-2.960*** (1.079)
SC X SC VillPop Share		4.101** (1.892)				
BC Keeper X SC			1.079* (0.649)			
UMC Keeper X SC			0.150 (0.601)			
SC Keeper X BC			-0.191 (0.713)			
BC Keeper X BC			-0.201 (0.637)			
UMC Keeper X BC			0.734 (0.678)			
Food Credit from Shopkeep				0.627 (0.469)		
Know Rice Entit.					1.109* (0.575)	
Know Wheat Entit.					1.516*** (0.555)	
Own Farm						-0.227 (0.247)
Obs	228	228	228	228	196	228
Groups	34	34	34	34	34	34
Correlation	0.295	0.328	0.383	0.314	0.400	0.302
Δ	.533***	.355*		.507***	.657***	.496***

Note: Village random effects probit models for whether household purchases grain. All specifications include a constant, caste plurality dummy, a caste land dominance dummy, land owned, housing structure, education dummies, distance from PDS shop, and share of SCs in the village. All specifications except the first column include a dummy for whether the household and Pradhan are from the same caste. All specifications except the second column include a dummy for being from a plurality caste. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 7: Home Production of Grain

Covariates	Home Prod./In-Kind Receipt (kg)	Prim Inc Own Farm (0/1)	Land Planted with Grain (Acres)	Land Value (Rs/Acre)	Percent Land Irrigated
SC	5.417 (8.409)	-0.049 (0.113)	-0.197 (0.262)	-1627.728 (12936.609)	-10.945 (8.363)
SC Shopkeeper	11.541 (11.155)	0.311** (0.129)	0.370 (0.270)	37401.202** (15506.197)	16.766 (10.210)
SC Keeper X SC	-27.565** (12.744)	-0.426** (0.173)	-0.500 (0.376)	-1.11e+04 (21523.504)	-8.319 (13.913)
Constant	45.153*** (9.115)	0.135 (0.116)	0.442 (0.268)	92901.695*** (19146.180)	65.781*** (11.278)
Obs	228	228	228	228	222

Note: All specifications estimated with OLS, with dependent variable as indicated in the top of the column. Grain includes rice and wheat only. Additional covariates from Columns 3, 7, and 11 of baseline take-up table are included here as well, but not shown ; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 8: Estimation Results

Parameter	Estimate
PREFERENCES	
β	0.957
α	0.237
ζ	23.339
κ_0 (Land)	-3.031
κ_1 (SC)	10.337
κ_2 (KeepSC)	2.679
COSTS	
δ_0 (Far)	0.967
c_0 (NonSC-KeepNonSC)	-1.534
c_1 (SC-KeepNonSC)	-0.820
c_2 (NonSC-KeepSC)	-0.267
c_3 (SC-KeepSC)	-1.236
c_4 (Time)	0.532
c_5 (ShareBPL)	0.769
c_6 (ShareSC)	3.417
c_7 (ShareBC)	0.929
FINE	
f	40.101
MEASUREMENT ERROR	
σ	0.377
θ	0.154

Table 9: Semi-Elasticity of Take-Up w.r.t. PDS Price

Pairing	Percentile				
	10th	25th	50th	75th	90th
KeepNSC/NSC	-0.546	-0.222	0.017	0.153	0.362
KeepNSC/SC	0.268	0.484	0.583	0.652	0.738
KeepSC/NSC	0.189	0.229	0.322	0.396	0.482
KeepSC/SC	0.236	0.322	0.515	0.694	0.921

Figure 1: PDS Purchases by Household and Shopkeeper Caste

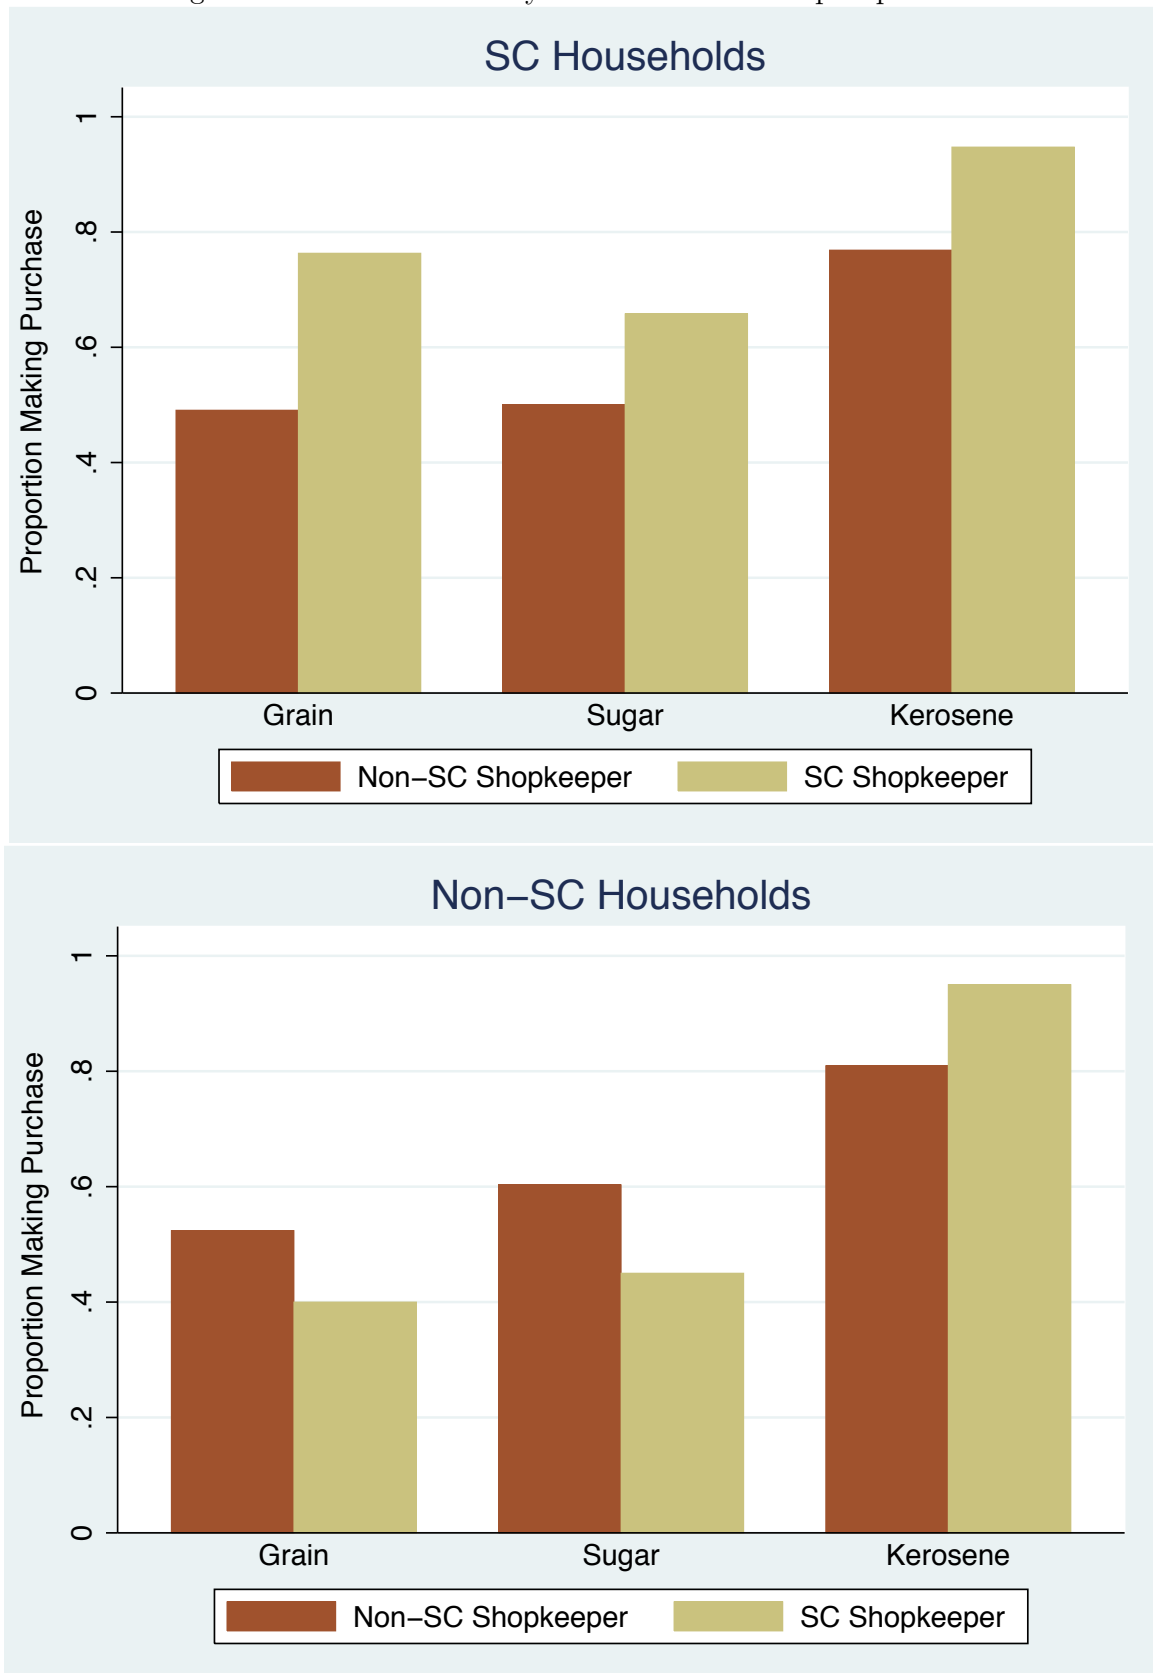


Figure 2: People Facing Few PDS Problems

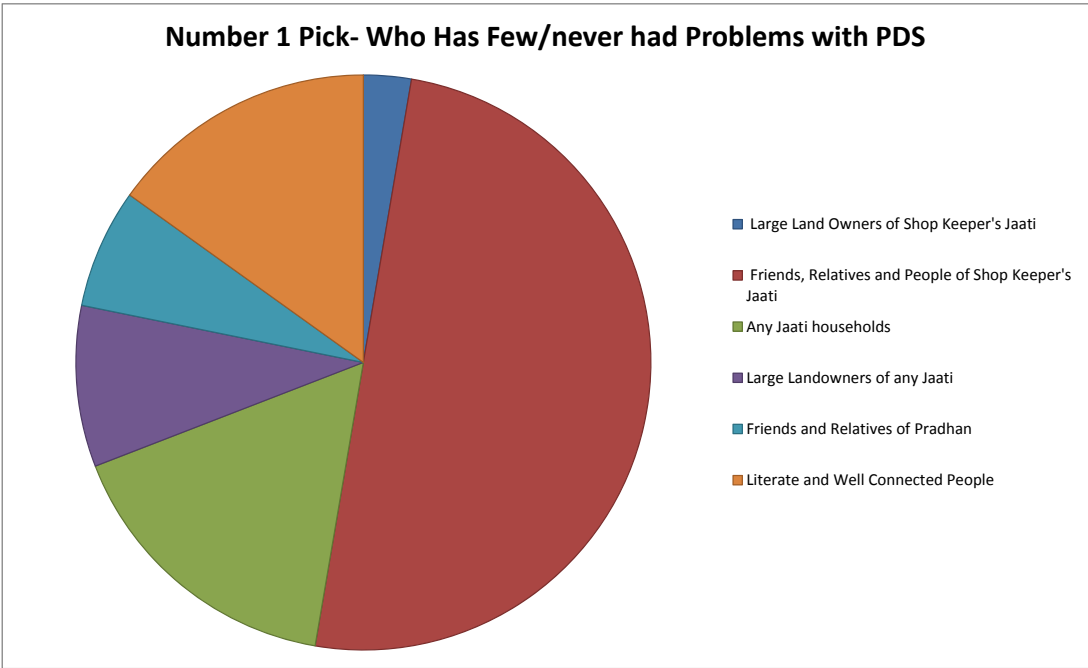
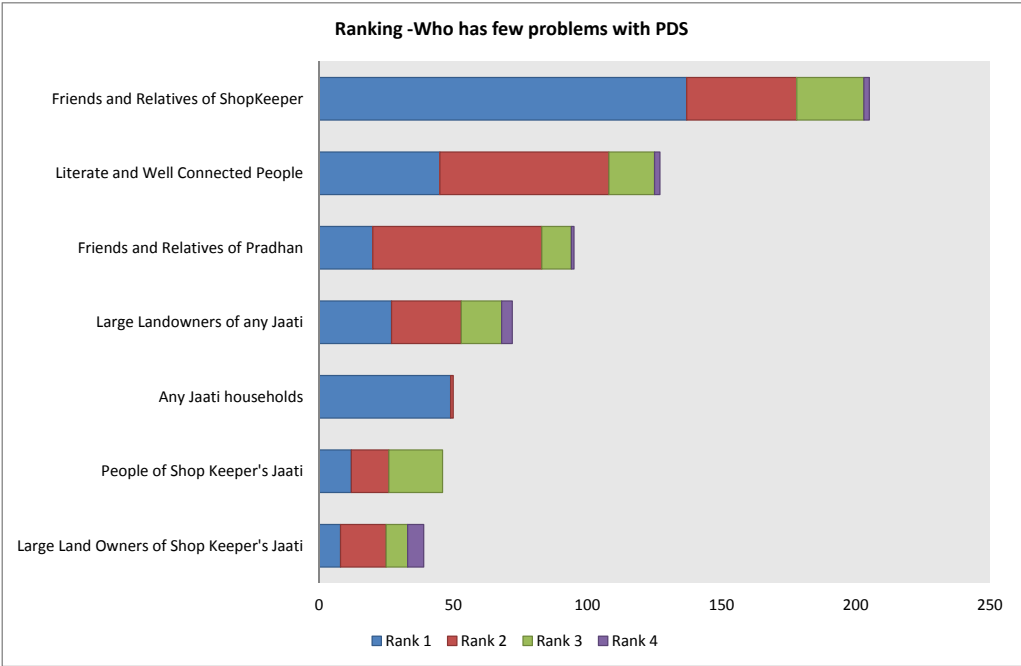


Figure 3: Information About Delivery Dates




Figure 4: Discovering Information About the PDS


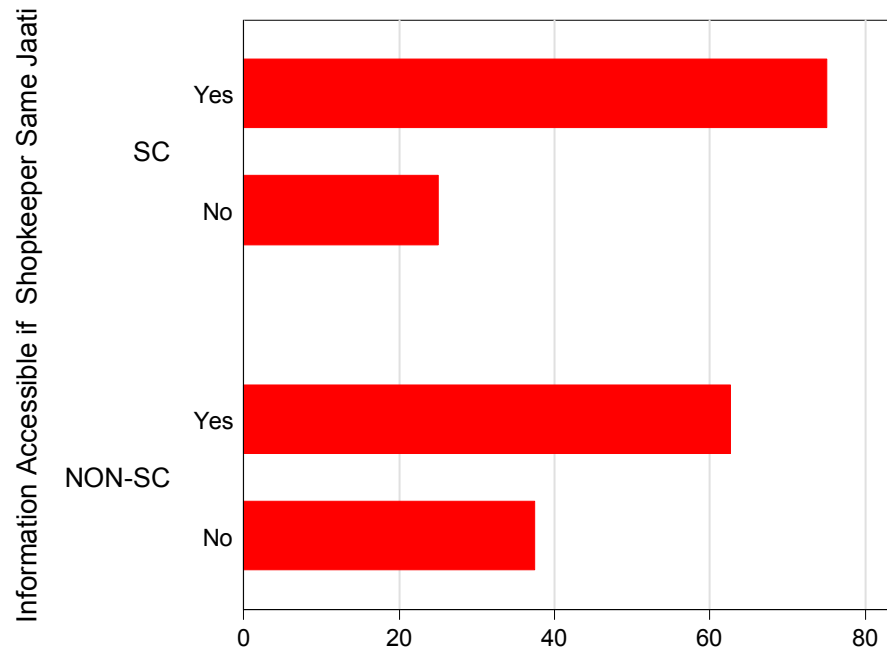


Figure 5: Concerns Over the PDS and Consequences

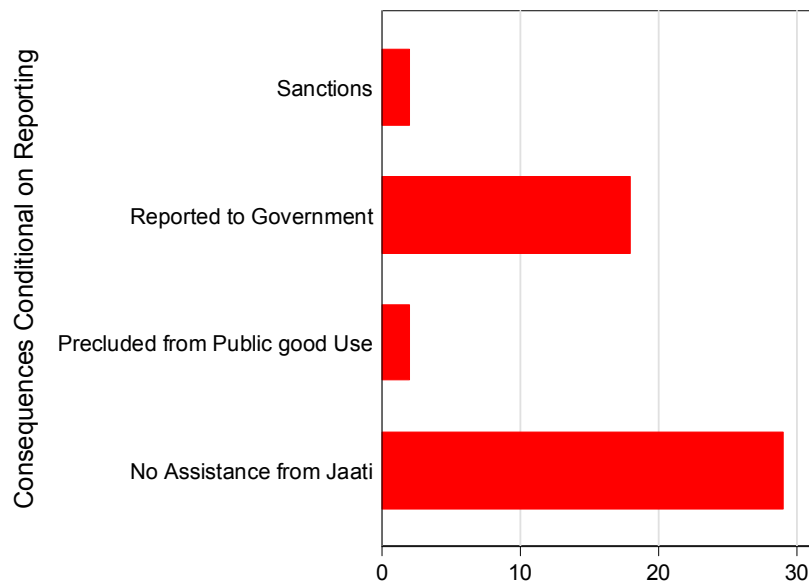
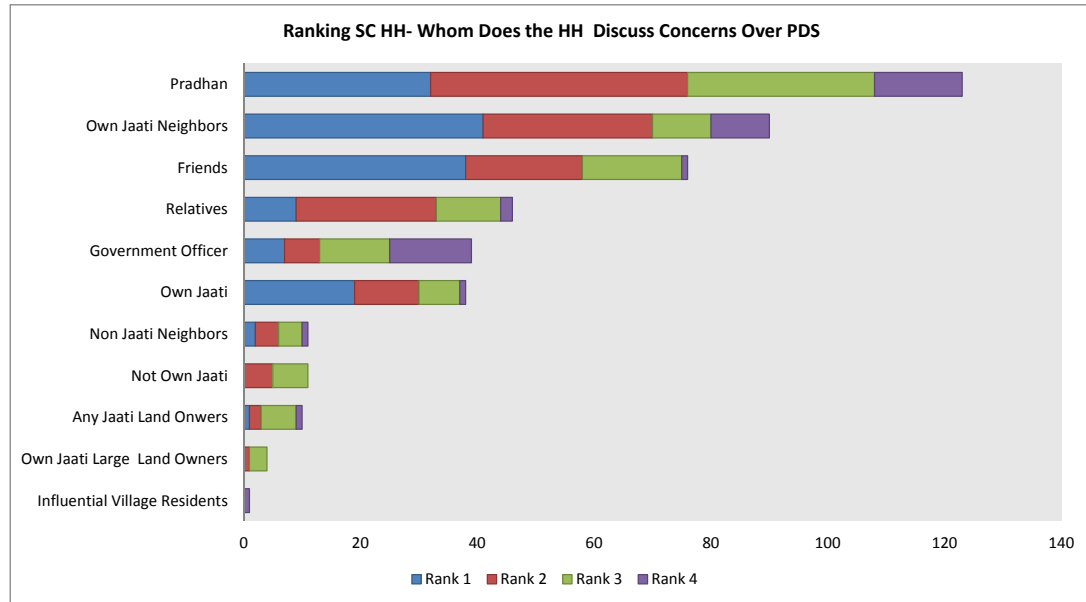


Figure 6: Fit: Log Grain Purchased from Market




Figure 7: Fit: Probability of PDS Grain Take-Up




Figure 8: SC Hholds' Welfare Gains from SC Shopkeepers (Mean Across Simulations)





Figure 9: SC Hholds' Welfare Gains from SC Shopkeepers (75th Ptile Across Simulations)



A Data Appendix

The household data contain basic demographic characteristics of the household members, including caste, age and education. The data also include extensive information on income, food and non-food expenditures, and farming and non-farming assets, caste of the shop keeper and PDS purchases. All villages do not necessarily have a PDS shop. Households in a village without a shop must acquire PDS goods from a neighboring village's shop, and we have information on these households' PDS purchases.

The sample used in the analysis is restricted to households with PDS shops in their village with the exception of one instance where we set variables regarding the PDS shop to zero and incorporate a dummy variable for whether the village has a PDS shop. We discuss this in the text. Not considering Muslim households and villages with Muslim shopkeepers, 229 households in 37 villages have BPL cards and information on the caste of the shopkeeper. There are 146 SC households, with 38 facing SC shopkeepers. There are 83 non-SC households, with 20 facing SC shopkeepers.

Table 1 covers households' PDS purchases for each good. Quality is a self-reported measure taking three values: 1, better than market quality; 2, same as market quality; 3, worse than market quality. Of those making PDS rice purchases, 91.3% of households purchase 3 kg (full entitlement). Of those making wheat purchases, 89.6% of households purchase 7 kg. Of the 173 households purchasing either rice or wheat, 121 (70%) purchase 10 kg of total grain, with an additional 22 households purchasing 7 kg of wheat and not purchasing any rice and 6 households purchasing 3 kg of rice and no wheat.

Table 2 examines how key village-level variables differ between villages with a SC shopkeeper, villages with a non-SC shopkeeper, and villages where the caste of the shopkeeper is not available. Villages with Muslim shopkeepers are not included.⁴⁸ Villages with SC shopkeepers tend to have a higher proportion of SCs in the village, are more likely to have a SC Pradhan, and are much more likely to have the largest share of households be SC. However, these differences are not statistically significantly different from zero. Only two variables have significantly different means across the village types: the proportion of villages with a UMC Pradhan and the estimated fraction of households with BPL cards. This estimate of BPL cardholders is simply the percentage of households in our sample who hold BPL cards in the village.⁴⁹ Therefore, SC shopkeepers are

⁴⁸The price of goods is the median price paid by a village's households for the good in question.

⁴⁹The test does not account for sampling variation in the estimate.

more likely to be observed in villages that have more BPL card holders. In the third panel, we see that villages without a PDS shop tend to be much smaller, as one might expect. They are also significantly more likely to have a BC Pradhan, and have significantly lower caste fractionalization.

B Further Evidence on Alternate Channels

B.1 Elite capture

Elite capture can operate in another form, as compared to the form stressed in the literature and in the main text. SCs are allocated BPL cards disproportionately and may have limited home production opportunities. Therefore, SCs may care about PDS grain more than other castes, and elite SCs may go out of their way to ensure poor SC households have access to PDS grain. If SCs do not care as much about kerosene and sugar, this explanation can generate the patterns we see.

To test this further, we turn to the first two columns of Table 6. In Column 1, we remove the dummy for whether the household is the same caste as the Pradhan, and instead add controls for a SC Pradhan and interaction of SC Pradhan with SC. The caste pairing effect is robust and the magnitude of Δ remains similar.⁵⁰ In Column 2 of Table 6, we represent SC political power in a slightly different way by adding two variables, the share of village population that is SC and the interaction of this share with a SC household dummy. The caste pairing effect remains statistically significantly different from zero (P-value of 0.019), though the magnitude of Δ falls somewhat.⁵¹

To test the elite capture hypothesis, we can also look at the price, quantity, and quality of goods purchased by households. If villages with SC shopkeepers are places where elite SCs capture PDS goods for themselves, we would expect that these elites would ensure that they obtain lower prices, larger quantities, or better quality. Table 4 shows the caste interaction effects and their standard errors from OLS regressions of price and quantity, and ordered probit regressions of quality, on the same set of covariates used in Column 3 of Table 3. Quality takes three values (Below market=3, Same as market=2, Better than market=1). Each cell shows the estimate and standard error for a particular dependent variable (listed along the left) and a particular good (listed along the top). For the sake of brevity, we do not show the estimates for all other covariates included.

Overall, there is little evidence that SC households obtain lower prices, higher quantities, or

⁵⁰Without weights, the P-value increases to 0.056. With Muslims included, the P-value is 0.136 (no weights) or 0.028 (with weights).

⁵¹Without weights, the caste pairing effect has a P-value of 0.048. When Muslims are included, the P-value rises to 0.128 (no weights) and 0.103 (with weights).

better quality of rice, wheat, sugar, or kerosene from SC shopkeepers. Regardless of whether Muslims are included and whether weights are used, prices, quantities, and quality are not significantly better for SCs facing SC shopkeepers for rice or wheat. In fact, the table shows that SCs obtain slightly worse quality from SC shopkeepers, though this result is not robust to dropping weights or including Muslims. Moreover, SCs seem to pay higher prices and purchase lower quantities of kerosene when facing a SC shopkeeper, and this does not depend on weighting or the inclusion of Muslims. There is some evidence that prices for sugar are lower for SCs facing SC shopkeepers, but this is sensitive to the use of weights and the magnitudes are very small. At the same time, there is no evidence that quantities of sugar and quality of sugar are higher for SCs facing SC shopkeepers. We should interpret these results carefully since they select on households that choose to make a purchase and we cannot observe the prices, quantities, and quality faced by households that did not make a purchase. Still, overall, the results are not consistent with a pure elite capture story.

B.2 Preference-Driven and Supply-Driven Differences in Take-Up Across Goods

One possibility is that SC shopkeepers favor their own caste but that SC households have a higher demand for grain than they do for sugar and kerosene. Another possibility is that SC households have similar preferences to other households, but that SC shopkeepers have a different ability to acquire sugar and kerosene than non-SC shopkeepers. These explanations are particularly relevant for kerosene, since the good is not sold on private markets and is therefore fundamentally different from other PDS goods. Either explanation could result in the caste pairing effect differing across goods.

However, these explanations are unlikely. Either story would have to explain the heterogeneity in the caste pairing effect that we have demonstrated. For instance, there is no reason to think that either explanation would lead to larger caste pairing effects for households that live further away from the PDS shop, or larger caste pairing effects for households with more land.

B.3 Taste Based Discrimination

A variation on the taste based discrimination channel is that there is an interaction between taste-based discrimination and black marketing incentives. If the incentives to black market grain are higher than those to black market sugar and kerosene, then shopkeepers could exercise taste-based discrimination in a manner that is differential across goods. However, this story is still not sufficient by itself to explain the fact that SCs do not see differentially better prices, quantities, or quality

with SC shopkeepers, as we saw above.

B.4 Selective BPL Allocation

A second approach to testing this explanation is to examine whether the gap in observables between APL and BPL households is larger for SC households facing SC shopkeepers relative to the SC households facing non-SC shopkeepers. That is, in villages with SC shopkeepers, are relatively poorer SC households being selected for BPL status? To do this, we regress key observables relating to wealth/income on BPL status, a SC household dummy, a SC shopkeeper dummy, and a full set of interactions of these three variables. The observables include: land ownership, non-durable expenditure, asset ownership, dummies for no schooling or for primary schooling only, and land price and percent of land irrigated (with 0 for the landless). The coefficient of interest is the coefficient on the triple interaction, which indicates whether there is differential selection along observables into BPL status for SCs who face SC shopkeepers. In results not shown here, we find that this coefficient is not significantly different from zero for any of the variables above, regardless of choices about weights or the inclusion of Muslims. This suggests that the gap in wealth between BPL households and APL households is no larger when SCs face SC shopkeepers.

B.5 Shop Level Stock Outs

One concern might be that the stock outs announced by the shop keeper are not a result of corruption at the shop level. The shop keeper may not get the goods. Specifically, the SC shopkeepers may not get grain but get kerosene. There are two reasons why we think this is unlikely. First, as we show in Table 5, the caste pairing effect is stronger for households living farther from the shop keeper and diminishes as the household gets closer to the shop. This result cannot be reconciled with genuine stock outs. In addition, we show that the caste pairing effect grows in size with land ownership. We replicate the standard specification for grain take-up, but include a full set of interactions between land owned, the SC household dummy, and the SC shopkeeper dummy. Column 5 of Table 5 reports the results. Wealthier households actually see a larger caste pairing effect. Households with zero land see a statistically insignificant caste pairing effect. The effect is statistically significant when land is 1 acre (P-value of 0.001), 3 acres (P-value of 0.001), and 5 acres (P-value of 0.005).⁵² If there was a genuine stock out at the shop, we would not expect to

⁵²Similar conclusions apply without weights. When Muslims are included, the caste pairing effect is significant at 1 acre only when weights are not used, and at 1 and 3 acres when weights are used.

see a differential caste pairing effect for households with for more land.

Table 1: [Appendix] Summary Statistics: Household PDS Purchases

	N	Mean	S.D.	Min	Max
RICE					
Purchase?	335	0.38	0.49	0	1
Price	127	4.37	0.27	3.5	5
Quantity	127	3.23	1.59	2	20
Quality	125	2.57	0.66	1	3
Know Entitlement?	293	0.45	0.50	0	1
WHEAT					
Purchase?	335	0.42	0.49	0	1
Price	141	3.43	0.22	3	4.5
Quantity	141	7.16	1.83	3	20
Quality	136	2.61	0.61	1	3
Know Entitlement?	293	0.51	0.50	0	1
SUGAR					
Purchase?	335	0.49	0.50	0	1
Price	164	12.33	0.78	4	13.5
Quantity	164	1.41	0.79	0.5	5
Quality	160	2.64	0.55	1	3
KEROSENE					
Purchase?	335	0.84	0.37	0	1
Price	282	4.14	0.61	3	9
Quantity	282	2.73	1.12	1	10

Table 2: [Appendix] Summary Statistics: Village Characteristics by Shopkeeper Caste Group

Variable	Non-SC	SC	P-Value	No Information	P-Value
Rice Price	6.90	6.71	0.421	6.47	0.016**
Wheat Price	5.84	5.75	0.760	5.46	0.036**
Sugar Price	14.01	13.79	0.691	14.10	0.789
Total Hholds	281.93	299.14	0.822	141.35	0.000***
Percent Landless	10.40	17.14	0.304	11.00	0.816
Percent BPL (Estimated)	0.270	0.452	0.001***	0.283	0.801
SC Share	0.29	0.36	0.390	0.29	0.986
BC Share	0.49	0.45	0.626	0.51	0.764
MUS Share	0.05	0.08	0.544	0.08	0.520
UMC Share	0.17	0.11	0.270	0.12	0.421
SC Pradhan	0.20	0.43	0.320	0.30	0.401
BC Pradhan	0.60	0.57	0.900	0.35	0.071*
UMC Pradhan	0.20	0.00	0.012**	0.26	0.613
SC Plurality	0.20	0.43	0.320	0.26	0.613
Caste Fractionalization Index	0.59	0.64	0.204	0.51	0.082*
Number of Villages	30	7		23	

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

Table 3: [Appendix] Choice to Purchase PDS Goods: Simultaneous Estimation

Covariates	Grain	Sugar	Kerosene
SC	0.432 (0.346)	0.196 (0.314)	0.476 (0.460)
SC Shopkeeper	-0.713* (0.366)	-0.788** (0.393)	1.747*** (0.542)
SC Keeper X SC	1.721*** (0.518)	0.829* (0.504)	-1.052 (0.652)
Land (Acres)	-0.136* (0.073)	0.004 (0.067)	-0.022 (0.076)
Pucca House	0.100 (0.236)	0.536** (0.252)	1.375*** (0.371)
No Schooling	-0.669** (0.284)	-0.563* (0.288)	-0.992*** (0.307)
Primary or Less	0.153 (0.255)	-0.258 (0.245)	-0.787*** (0.278)
PDS Dist <.5km	0.771*** (0.241)	0.465** (0.231)	0.530** (0.251)
HH's Caste Forms Plurality	0.564** (0.277)	0.224 (0.261)	0.649* (0.338)
Pradhan-HH Same Caste	-0.325 (0.272)	-0.244 (0.239)	-0.556* (0.315)
Caste w/ Land Dominant Sub-Caste	-0.066 (0.308)	0.277 (0.302)	-0.287 (0.417)
SC VillPop Share	-2.648*** (0.745)	-0.631 (0.727)	-0.978 (0.778)
Constant	0.411 (0.392)	0.065 (0.377)	1.260*** (0.397)
Obs	228		
Δ	.549***	.239	-.359
P-Value Diff From Grain	-	.137	.007

Note: Probit models allowing for correlation of three unobserved errors. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 4: [Appendix] Prices, Quantities, and Quality of PDS Goods (Condtl. on Purchase)

Dep. Variable	Rice	Wheat	Sugar	Kerosene
Price	-0.095 (0.147)	0.020 (0.125)	-0.320* (0.175)	0.449*** (0.169)
Quantity	-0.135 (0.180)	-0.079 (0.468)	-0.116 (0.282)	-1.067*** (0.388)
Quality	1.440** (0.695)	1.190* (0.695)	0.924 (0.574)	-0.044 (0.477)

Note: Price and quantity specifications estimated with OLS. Quality specifications estimated using ordered probit, where the categories are (1) Better than market quality (2) Same as market quality (3) Worse than market quality. Additional covariates from Columns 3, 7, and 11 of baseline take-up table are included here as well, but not shown ;

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

Table 5: [Appendix] Role of BPL Card Allocation in Caste Pairing Effect

	BPL Card	Grain Take-Up	Sugar Take-Up
SC/ST	0.846*** (0.186)	0.568 (0.428)	0.384 (0.318)
Keeper SC/ST	0.360 (0.287)	-0.671 (0.668)	-0.121 (0.705)
SC/ST X Keeper SC/ST	-0.637** (0.279)	-0.470 (0.642)	0.440 (0.554)
SC/ST Pradhan	0.071 (0.267)		
SC/ST X SC/ST Pradhan	0.560** (0.264)		
Land	-0.180*** (0.032)	-0.627*** (0.161)	-0.282** (0.129)
No Schooling	0.465** (0.195)	-0.421 (0.403)	-0.778** (0.348)
Below Primary	0.485*** (0.175)	-0.574 (0.371)	-0.526 (0.341)
PDS < 0.5 km		0.368 (0.332)	0.175 (0.332)
SC Share		1.633 (2.024)	0.994 (2.058)
Constant	-0.848*** (0.169)	1.141** (0.558)	0.196 (0.590)
Obs	802	220	220
Groups	24	24	24
Correlation	0.192	0.476	0.553

Note: All specifications estimated using random effects probit, with random effects at the village level; *** $p < 0.01$,

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

Figure 1: [Appendix] Discovering Information About the PDS

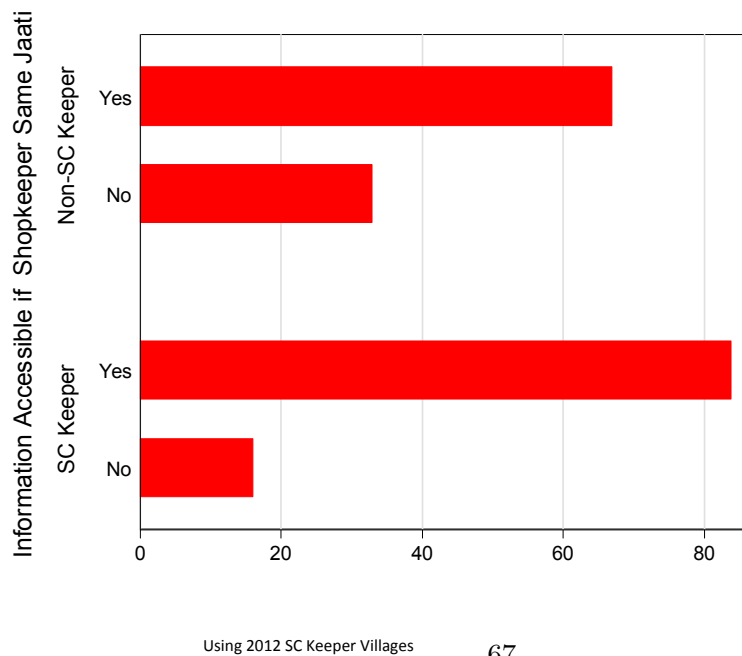
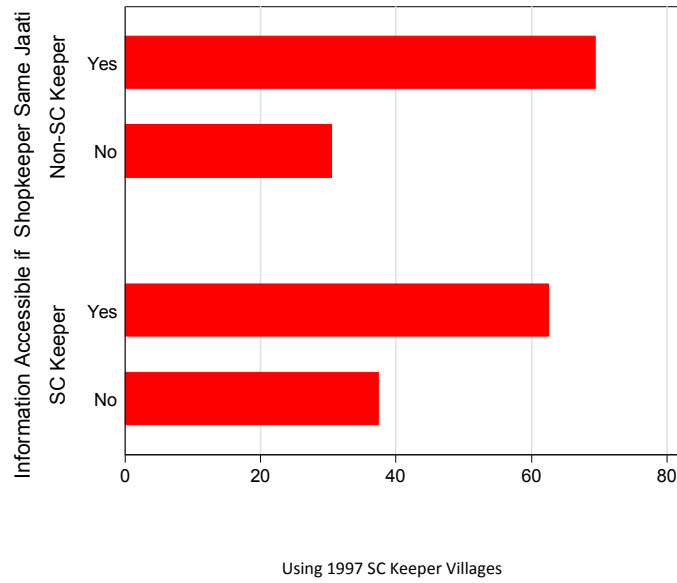
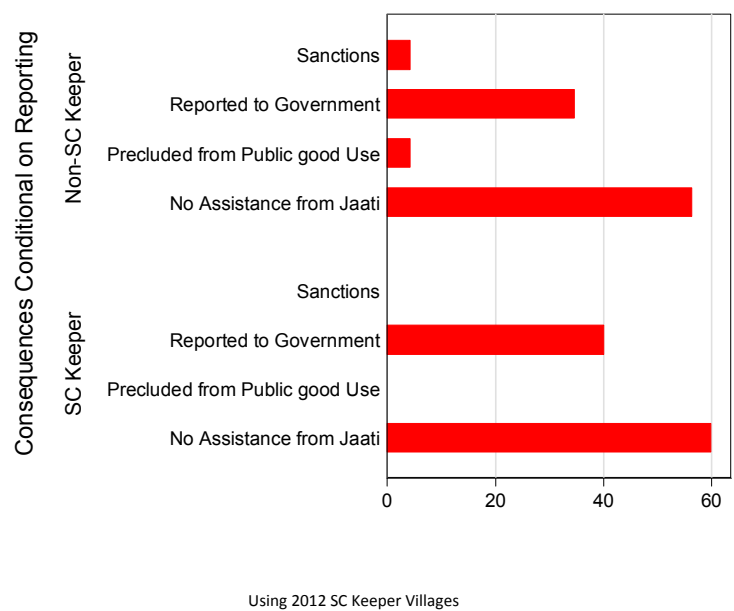
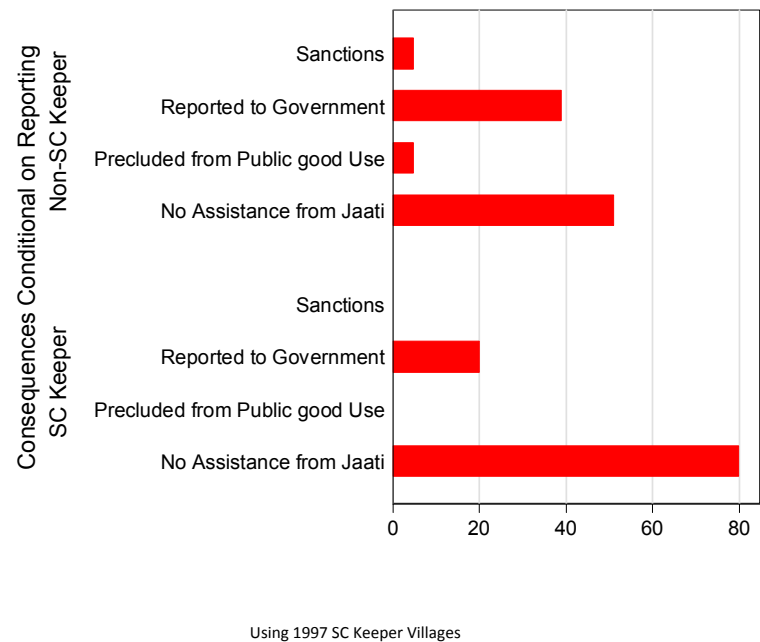


Figure 2: [Appendix] Consequences Conditional on Reporting



The International Growth Centre (IGC) aims to promote sustainable growth in developing countries by providing demand-led policy advice based on frontier research.

Find out more about
our work on our website
www.theigc.org

For media or communications
enquiries, please contact
mail@theigc.org

Subscribe to our newsletter
and topic updates
www.theigc.org/newsletter

Follow us on Twitter
[@the_igc](https://twitter.com/the_igc)

Contact us
International Growth Centre,
London School of Economic
and Political Science,
Houghton Street,
London WC2A 2AE

IGC
International
Growth Centre

DIRECTED BY



FUNDED BY



Designed by soapbox.co.uk