Targeting subsidies through price menus: Menu design and evidence from clean fuels*

Sulemana Abubakari† K.P. Asante† Misbath Daouda‡
B. Kelsey Jack‡ Darby Jack§ Flavio Malagutti¶ Paulina Oliva‖

August 17, 2025

Abstract

We demonstrate that design-based non-linear subsidies can be more cost effective than linear subsidies. We consider a policy maker who wants to increase the use of clean cooking fuel while keeping subsidy spending low. Fuel is sold in cylinders (packages) of different sizes, allowing for size-differentiated pricing, but not for limits on the frequency of purchase. We demonstrate implementation in two stages: (1) we gather the empirical inputs needed for a sufficient statistic based approach to designing the non-linear subsidy, then (2) we test the optimal subsidy against counterfactual linear subsidies. Relative to the counterfactual, allocating a larger per-unit subsidy to purchases in the smaller package size, which is taken up by relatively more price sensitive individuals, achieves the same level of clean fuel demand with 30% less in public expenditure. We show, both empirically and theoretically, that intuitive conditions – namely transaction costs and savings constraints – generate frictions that sustain a separating equilibrium when the smaller package is more heavily subsidized. Because poorer households have a preference for small-sized purchases, this non-linear price menu is also progressive.

*We thank Iddrisu Seidu, Alex Appiah, Sule Awuni and the field officers at the Kintampo Health Research Centre for collaboration, and Alexander Abajian, Akanksha Arora and Linnea Graham for research assistance. Numerous seminar and conference audiences provided useful feedback. We are grateful for financial support from Columbia World Projects and the King Climate Action Initiative at J-PAL. This RCT was registered in the American Economic Association Registry under number AEARCTR-0008400. Authors are in alphabetical order; Abubakari is listed in memoriam. Additional details are provided in an online supplement available here.

†Kintampo Health Research Centre
‡UC Berkeley
§Columbia University
¶UC Santa Barbara
‖University of Southern California

1 Introduction

Governments in nearly every country subsidize goods and services that support redistribution goals or generate positive externalities (World Bank, 2022). Uniform subsidies are costly. In many cases, better targeting can lead to the same policy objectives at a lower cost. Perhaps unsurprisingly, large literatures in public and development economics study alternative targeting approaches (e.g., Nichols and Zeckhauser, 1982; Alderman and Lindert, 1998; Hanna and Olken, 2018). We explore non-linear pricing – unit prices differentiated by purchase size – as a means of targeting product subsidies. The focus on non-linear pricing is motivated by the observation that consumers appear to sort into different purchase or package sizes in a range of contexts. For example, low-income consumers have a well-documented tendency to purchase in small package sizes, despite the costs associated with small or high-frequency purchases, like forgone bulk discounts and transaction costs (Pires and Salvo, 2015; Attanasio and Pastorino, 2020; Dillon et al., 2021). In addition to demand heterogeneity, frictions such as transaction costs and liquidity constraints could explain this behavior.[1] Firms take advantage of this consumer heterogeneity: around 855 billion single-use sachets of food and personal care products are sold annually, despite unit prices that often exceed the cost of the same product in a larger package (Geddie and Brock, 2022). We show that non-linear pricing strategies can be adapted for the public sector to improve the targeting of subsidies for socially desirable consumption, such as clean technologies.

In principle, targeting subsidies through non-linear pricing is attractive to a policy maker concerned about cost effectiveness. Unlike targeting on observables, such as proxy means tests (e.g., Alatas et al., 2012), it does not require observing characteristics of all potential recipients. Unlike other approaches to targeting on unobservables, such as ordeal mechanisms or bidding (e.g., Jack, 2013; Alatas et al., 2016; Dupas et al., 2016), it can be rolled out using existing market structures. However, designing the non-linear price menu requires knowledge

[1] The standard case for second degree price discrimination separates consumers on heterogeneity in demand: rents are extracted from low demand types who would rather pay less in total even if that means paying more per unit (e.g., Maskin and Riley, 1984). However, heterogeneity in transaction costs across consumers has been highlighted by Mussa and Rosen (1978) and Parks (1974) as an alternative source of separating equilibria. We develop a theoretical model that shows that heterogeneity in frictions can lead to a separating equilibrium under non-linear pricing in contexts where demand heterogeneity alone would lead to pooling on package size.

of how demand and the subsidy budget will be affected by different menus of price-quantity bundles. In this paper we develop a broadly applicable sufficient statistic framework that economizes the information needed to produce the policy maker's optimal price menu. We show that all of the relevant consumer heterogeneity that governs selection across package sizes is captured by the difference in willingness to pay (WTP) across sizes. Thus, eliciting this sufficient statistic within a representative sample of the population, and observing how this correlates with purchases at randomly assigned prices and sizes (i.e., size-specific price sensitivity), provides the necessary information for menu design. We demonstrate the feasibility of this approach by empirically testing an optimally designed price menu against a linear subsidy in the context of clean cooking fuel in Ghana.

Our setting is well suited to testing non-linear pricing for several reasons. First, the status-quo polluting fuel used by most households in our sample is charcoal, which low-income households purchase in small quantities multiple times per week. This is an initial indication of heterogeneity in the choice of purchase size for the incumbent cooking fuel. Second, while a tax is presumably the first-best policy to address the negative consumption externalities from charcoal, effective taxation is difficult both for practical and political reasons. Instead, Ghana, like many countries around the world, promotes cleaner substitutes, namely liquid petroleum gas (LPG), which we calculate generates roughly one-fifth the greenhouse gas emissions as charcoal.[2] In contrast to charcoal, LPG prices are set by a national regulator and all transactions are through formal markets. Third, households purchase LPG in cylinders of fixed sizes, creating the opportunity for quantity-differentiated pricing. We introduce study-operated LPG depots, where study participants can exchange cylinders at prices that we control.[3]

We implement our approach in two stages: (1) design, then (2) test. Each stage uses a stratified random sample of charcoal users drawn from the same population. In stage 1, we gather the data necessary to design the non-linear price menu that maximizes the policy maker's objective function, which we specify as demand, conditional on a fixed subsidy

[2] Subsidies for "greener" substitute goods are common around the world, in both developed and developing countries (Fowlie and Meeks, 2021; Blanchard et al., 2023; Sallee, 2025).

[3] Households obtain cylinders on deposit, and stoves are often subsidized, so we suppress the fixed cost of acquiring durables in our theory and implementation.

budget. Specifically, we use a multiple price list (MPL) elicitation format to measure WTP for LPG in cylinders of different sizes. This allows us to quantify the difference in WTP across the two sizes for each household – the sufficient statistic for selection into size. Conditional on responses to the MPL, participants receive a randomly selected size-price pair. They then have three months to purchase LPG (i.e., exchange empty cylinders for full ones) at the selected price with no restriction on the frequency of purchase. The resulting demand data allow us to estimate price elasticities for different consumer types when exchanging each of the two cylinder sizes. After imposing a budget constraint, we solve for the non-linear subsidized prices that maximize the policy objective function (i.e., demand). This objective function balances the health and environmental costs of charcoal use against the cost of public spending on subsidies.[4] The resulting optimal price menu places a higher subsidy per kilogram (kg) of LPG on the smaller cylinder compared to the larger cylinder – i.e., a bulk premium instead of a bulk discount. This result emerges from the relationship between the price elasticity of LPG demand and the sufficient statistic for selection into sizes: households more likely to choose the large cylinder are less price sensitive.

In stage 2, we test the optimal non-linear price menu against linear prices (i.e., prices that do not vary with cylinder size) that are budget-neutral in expectation. We randomly assign households to a non-linear price treatment or a linear price control. Households choose their preferred cylinder size at treatment-specific prices, and – like in stage 1 – purchase LPG by exchanging their cylinders at a study-run depot for a period of three months. We pre-specify that the treatment will increase the cost effectiveness of subsidies; either the policy maker will get more demand for her budget, or spend less for at least as much demand.[5]

We find that the non-linear price treatment delivers the same overall demand as linear pricing, but at a roughly 30% lower budget. As predicted by the stage 1 empirical model, large cylinder types are relatively insensitive to price, so the higher per kg price under non-linear pricing does little to dampen their demand, while reducing public expenditures. Small cylinder types, on the other hand, increase their demand in response to the subsidy, while

[4]A broader welfare maximization objective would also account for consumer surplus and non-monetary costs. We also provide a back of the envelope calculation of the overall welfare effects of non-linear pricing.

[5]Our test could have imposed budget neutrality or demand neutrality; we chose the former as the basis for our design given our fixed sample size, though our results end up closer to the latter. We pre-specified testing for both alternatives.

costing the policy maker comparatively less because small cylinder types consume less LPG. The additional subsidy on the small cylinder draws only 5% more households to the small cylinder compared to linear pricing. We see no difference on the extensive margin – i.e., choosing any cylinder vs. choosing none. Although we cannot measure the causal effect of LPG subsidies on charcoal displacement, we observe a drop in charcoal use following LPG adoption consistent with a displacement rate greater than 1:1. We calculate a cost of USD 10 per ton of GHG reduction through non-linear subsidies, well below even conservative estimates of the social cost of carbon (SCC).[6]

Our implementation offers a proof of concept for feasible design of non-linear pricing, with potential applications in a range of settings where consumer good subsidies can be applied to discrete package sizes, including health products, educational materials, and other "green" goods. However, the results present something of a puzzle under standard models of quantity differentiated pricing. Specifically, the standard theoretical model assumes that a monopolist can effectively restrict the purchase of a given price-quantity bundle to exactly one bundle per period per consumer. This is a seemingly innocuous assumption in the case of, for example, health insurance, where policies have a fixed one-year duration and there are strong disincentives to hold multiple policies at once (Rothschild and Stiglitz, 1976). However, for LPG and many other products sold in discrete package sizes, there is no practical way of restricting the number of purchases per period. Yet, a separating equilibrium is sustained under bulk discounts for many goods (including charcoal) and under the bulk premium for LPG that we implement. In light of this, we extend the theory to a more general case that allows for a continuous number of purchases per consumer over an arbitrary period of time. This allows individuals to consume large packages slowly (similar to Hendel and Nevo (2013)) or replicate the quantities in large packages through repeated purchase of small packages (Alger, 1999). To sustain separating equilibria, while relaxing the one-bundle-per-period assumption, we add two types of frictions that are plausible in a wide range of settings: saving costs and transaction costs (e.g., opportunity cost of time or hassle costs). The existence of saving costs rationalizes why some consumers opt for small packages, even when

[6]The displacement rate would have to be less than 0.38 for the cost of GHG reductions to fall below an SCC of USD 51 per ton. This ignores other benefits of reduced charcoal consumption, including ambient and indoor air quality.

they can buy large packages at bulk discounts and consume them very slowly. Transaction costs, which are proportional to the number of purchases, work against pooling on the small package size even in the face of bulk premia (like in our implemented optimal policy).[7] Our model shows that heterogeneity in either friction is sufficient to sustain separation under bulk discounts and bulk premia. While others have considered heterogeneity in transaction costs (Parks, 1974) or storage costs – which also limit pooling into large packages (Hendel et al., 2014) – our model is the first to bring these two frictions together to explain why we observe sustained separation under a wide range of non-linear pricing schemes.

Our stage 2 results are consistent with the model predictions. First, we find that most households prefer the large cylinder, even when the cost per kg of LPG is higher. This is consistent with the existence of strictly positive transaction costs on average. Second, households with higher savings costs are more likely to take up the small cylinder, but we observe no correlation between transaction costs and choice of size. This is consistent with heterogeneity in savings costs generating separation. In other words, all households would prefer to avoid the transaction costs associated with the small cylinder, but those with particularly high savings costs cannot. We find similar patterns in our stage 1 sample: the sufficient statistic for consumer type is negatively correlated with savings costs and positively correlated with wealth. An implication of this selection pattern is that – because the optimal schedule places a higher per kg subsidy on LPG in small cylinders – the non-linear pricing effectively targets a larger share of the total subsidy value to poorer households.[8]

We contribute to a large literature that investigates different approaches to targeting subsidies and transfers (Coady et al., 2004; Alatas et al., 2012; Cohen et al., 2015; Johnson and Lipscomb, 2022; Rafkin et al., 2023). Most of these papers either test alternative targeting strategies in the field (e.g., Jack, 2013; Alatas et al., 2016) or use existing data to

[7]In our setting, additional hassle costs of small LPG cylinders – e.g., the cost of interrupting cooking, disconnecting and empty cylinder and reconnecting a full cylinder – may also work against pooling in the small package. These costs will exist independent of the number of trips made. Nevertheless, in our experimental setting, we restrict households to exchange a single cylinder at a time, which mimics a "one discount per transaction" policy. We discuss the trade-offs in restrictions on continuous purchases in Section 6.

[8]Our sufficient statistic approach is agnostic about the determinants of selection into different package sizes or heterogeneity in demand elasticities, and may therefore result in regressive non-linear price schedules. However, the frictions that given rise to a separating equilibrium in our theoretical model suggest that bulk premia may generally be more progressive than bulk discounts (see also Pires and Salvo (2015); Attanasio and Pastorino (2020); Dillon et al. (2021)).

design theoretically optimal targeting schemes based on either on observable or unobservable characteristics (e.g., Ito et al., 2023; Ida et al., 2022; Langer and Lemoine, 2022; Haushofer et al., 2022; Higbee, 2024). Because these papers stop at the design stage, they implicitly assume that the policy environment is fixed and that the model used for policy design is correctly specified. To our knowledge, the only papers both gather the relevant information and test the resulting policy are Johnson and Lipscomb (2022) and Dubé and Misra (2023). In both of these, the relationship between observables and willingness to pay is measured in a first stage, and subsidies are targeted based on observables in a second stage (i.e., third degree price discrimination). We, instead, develop an elicitation procedure that reveals unobservables and use it to design and test a targeting strategy that does not require observing the characteristics of all individuals in the population for implementation (i.e., we use second degree rather than third degree price discrimination).

We follow the literature on second degree price discrimination in setting up the price menu design as a screening problem, with heterogeneous consumers.[9] A small literature examines second degree price discrimination as a policy tool for improving the cost effectiveness of public programs. Most notably, Dizon-Ross and Zucker (2023) test quantity differentiated incentive contracts for completing a target number of steps per day against a flat incentive contract.[10] We complement their empirical evidence with a methodology for how to design the optimal price menu, and apply it to product pricing.[11] In doing so, we extend the standard model of discrete price-quantity bundles by adding an intensive margin of demand, i.e., that a consumer can purchase her preferred bundle any non-negative number of times in an arbitrary period, and allowing for richer sources of consumer heterogeneity. These extensions allow theory on menu design to apply to a much wider set of empirical settings

[9]This falls into a much larger class of non-linear pricing problems (see Armstrong (2016) for a review). Public utilities have long used non-linear tariffs for both cost recovery and redistribution (Borenstein, 2012; Szabó, 2015; Nauges and Whittington, 2017).

[10]The setting in Dizon-Ross and Zucker (2023) resembles a larger literature on non-linear incentive contracts designed and tested in the personnel literature, many of which have social goals in the sense that they are applied to jobs like teaching or health care (e.g., Duflo et al., 2012).

[11]Using a mechanism design approach, Kang (2022) studies the case of non-linear pricing to address externalities from driving through a tax on vehicle miles traveled. He shows that the optimal price menu could result in either quantity discounts or mark-ups depending on the relationship between price sensitivity and the externality, but does not test the optimal menu. To our knowledge, other than Dizon-Ross and Zucker (2023), Levitt et al. (2016) offer the only other field test of second degree price discrimination, and find no effect on demand or profits in an online retail setting.

(including ours).

Finally, our application links us to literatures on household energy use in development and environmental economics and in public health. Most of the energy used by households in low-income settings is for cooking, and much of the prior work has focused on encouraging the take up of improved biomass cookstoves (e.g., Smith et al., 2011; Bensch and Peters, 2015; Hanna et al., 2016; Berkouwer and Dean, 2022). More recently, the emphasis has shifted toward cleaner fuels, namely LPG, where findings from large scale government programs suggest both widespread behavior change and some positive health impacts (Imelda, 2020; Verma and Imelda, 2023; Gould et al., 2024). While these public programs have achieved some success, they have done so through large public expenditures and limited success in targeting subsidies to marginal users, despite the fact that several countries follow our strategy of differentially subsidizing small cylinders, including Indonesia and Senegal (Kojima, 2021).[12] We demonstrate a novel and scalable approach to targeting LPG subsidies to lower public expenditures while meeting clean energy goals.

2 Background and context

We start by offering additional background on the cooking technologies and the policy problem that underlies our empirical strategy and our theoretical model.

2.1 Pollution from household energy use

Pollution resulting from the burning of solid fuels or biomass (firewood, dung and charcoal) for cooking and heating is a leading cause of morbidity and mortality around the world (World Health Organization, 2023), and contributes to climate change through the degradation of forests and the release of greenhouse gases and black carbon (Bailis et al., 2015). Replacing biomass with cleaner burning fossil fuels, including LPG, has the potential to reduce emissions of both local pollutants and greenhouse gases (Floess et al., 2023).[13] LPG

[12] In these cases, subsidies on small cylinders have been large and resulted in pooling on the small size; experimentation around how to set the subsidy appears limited.

[13] While renewable energy sources offer an even cleaner alternative, current renewables technology cannot meet global cooking energy demand at a feasible cost.

offers a scalable alternative to biomass, since it relies on decentralized distribution networks, and is relatively easy to transport and use (Gould et al., 2024).

Many governments have clean cooking targets that rely on the scale up of LPG as a cooking fuel, and many use subsidies to help reach these targets. For example, the World LPG Association estimates that India, Indonesia and Morocco together spent over $10 billion on residential LPG subsidies in 2022 (World LPG Association, 2023). Some countries, including Indonesia and Ivory Coast, have tried to target subsidies by cylinder size, similar to the approach tested here. In Indonesia, large price differentials – around a 70% subsidy on LPG in the small cylinder – have resulted in pooling (Thoday et al., 2018). A combination of subsidies and economic growth have contributed to the rise of LPG as a domestic cooking fuel: in the late 2000s, gas fuels overtook biomass as the leading cooking fuel in low- and middle-income countries (Stoner et al., 2021). Between 2010 and 2020, 75% of the households that transitioned to cleaner fuels switched from solid fuels to LPG (Floess et al., 2023).

2.2 LPG and charcoal in Ghana

In Ghana, the majority of households cook with biomass. According to the 2021 census, 54% of households used charcoal or firewood as their primary fuel, while 37% used LPG (Ghana Statistical Service, 2022). The Ghanaian government has a target of 50% LPG use by 2030, alongside a commitment to keep subsidization to a minimum.

In Ghana, and elsewhere, LPG is consumed out of cylinders connected to a burner or stove. Thus, LPG adoption requires both a one-time purchase of a stove and cylinder, and recurring expenditures of both money and time to acquire fuel.[14] To encourage LPG adoption, Ghana is transitioning from a refill model, in which households bring their cylinder to a filling station and choose how many kg to refill, to a recirculation model, in which households exchange empty cylinders for full ones. The cylinder recirculation model was announced in 2017 and launched in September 2023, at the end of our study. We design our

[14] A 2-burner stove costs roughly the same as 14.5 kg of LPG, or the amount in a large cylinder. Numerous donors offer subsidized or free stoves, so we choose to ignore this aspect of the adoption problem and focus on the challenge of repeated purchases and use of clean fuels.

data collection around a pilot version of cylinder recirculation (see Section 4.1).

LPG cylinders come in a range of sizes. For home cooking, 3, 6 and 14.5 kg cylinders are the most common. Under cylinder recirculation, suppliers in Ghana own the cylinders and are responsible for their maintenance; customers pay a one-time deposit for the cylinder and purchase gas by exchanging an empty cylinder for a full one. This represents an important change in how households pay for fuel: requiring purchase of a full cylinder reduces the flexibility in expenditures relative both to a refill model and to charcoal. However, the lumpiness in the quantities of LPG available for purchase also represents an opportunity for pricing that is differentiated on cylinder size.

We gathered data on status quo charcoal purchasing behavior from our study sample. Households reported purchasing charcoal in small quantities and at high frequency, on average (Table 1). Around 40% of the sample reported buying charcoal in small bags of 1-2 kg. The first column shows the average of each characteristic for households that bought in large quantities and the second column shows the estimated coefficient on purchasing in small quantities, in separate regressions for each characteristic. Households that bought in large bags bought an average of just over five bags per month, versus 32 times more per month among those who purchased in small quantities. Households that purchased in small quantities spent less money on charcoal and on travel to buy it, but spent almost three times more time buying it. They were also less wealthy, less likely to have formal savings and are more credit constrained. We observe no evidence of measurable differences in opportunity cost of time or access to transportation across purchasing types. These patterns may be shaped by other factors, including on the supply side, but suggest a revealed preference for small, high frequency purchases among poorer households, in spite of the time and hassle costs involved.[15]

2.3 Externalities from charcoal and LPG

Both charcoal and LPG generate externalities when burned. The argument for policy intervention to switch households from charcoal to LPG hinges on higher relative externalities

[15]This echoes patterns of high frequency, small quantity purchasing behavior among low-income households in other settings (Dillon et al., 2021; Pires and Salvo, 2015).

Table 1: Baseline charcoal purchasing patterns (stage 2)

	Large (mean)	Small (coef/se)
Number of monthly purchases	5.143	32.377***
	(16.063)	(1.217)
Implied monthly expenditure (GH¢)	315.316	-122.470***
	(997.093)	(44.842)
Implied monthly travel cost (GH¢)	6.305	-4.037
	(39.809)	(2.875)
Implied monthly travel time (hours)	0.773	1.689***
	(3.356)	(0.256)
Wealth quintile	4.140	-0.254***
	(1.195)	(0.092)
Any formal savings	0.825	-0.136***
	(0.419)	(0.034)
Days to get 150 GH¢	5.771	0.434**
	(2.514)	(0.180)
High opportunity cost of time	0.330	0.023
	(0.537)	(0.038)
Any no cost transport access	0.312	-0.007
	(0.527)	(0.037)
Household size	5.615	-0.118
	(3.725)	(0.294)

Note: Column 1 shows means and standard deviations in parentheses. Column 2 shows coefficients and robust standard errors from separate linear regressions of each row variable on a small purchases indicator (39% of households bought in small quantities at baseline). All statistics are weighted by stage 1 strata. * p<0.1, ** p<0.05, *** p<0.01

from charcoal. We do not measure these directly in our data collection, and so rely on the literature for a back of the envelope calculation. We focus on damages external to the household, though intrahousehold externalities or misinformation might also imply that some of the health damages accruing to household members are also inefficient.

Cooking with either charcoal or LPG emits both local air pollutants and greenhouse gases. The former imposes health damages on other households within an airshed, the magnitude of which depends on ambient pollution levels and the vulnerability of the exposed population. Given the complexity in these calculations, we construct a lower bound on the external damages, based on greenhouse gas emissions only. We are primarily interested in the relative emissions from the two fuels. Floess et al. (2023) calculate the CO_2 equivalent emissions (CO_2e) per MJ of usable energy for both charcoal and LPG, from both production and combustion.[16] Per MJ, charcoal emits roughly 5 times the CO_2e as LPG, including production, transportation and combustion.[17] Using a social cost of carbon range from USD 51 to 204, the avoided damages range from USD 1.7 to 6.8 per kg of LPG that replaces the equivalent energy quantity of charcoal.

The first best policy intervention to address external damages from charcoal would target charcoal directly, for example by taxing purchases at the marginal damage from charcoal. However, in spite of efforts to formalize the charcoal supply chain in Ghana, much of the market remains informal. This creates practical challenges for the feasibility of taxing the "bad." Instead, a second best policy, feasible in the nationally regulated LPG market, places a subsidy on the "good." The efficient level of the subsidy will depend on a number of factors, including the substitution rate between charcoal and LPG, which we do not observe. We set up a policy maker's problem that focuses on cost effectiveness rather than efficiency, by imposing a budget constraint on subsidies for LPG. This objective is consistent with a high marginal cost of public funds in our setting and political constraints on subsidies for LPG.

3 Non-linear price menu design

The descriptive results in Table 1 suggest that rich heterogeneity across consumers drives sorting into different charcoal package sizes. It is plausible that similar heterogeneity can determine whether consumers purchase LPG in a small or a large cylinder. If that is the case, a non-linear price menu may induce a separating equilibrium and increase the cost effectiveness of subsidies relative to a linear price (i.e., a price that does not depend on the size of the cylinder).

In this section, we first establish the objective function of the policy maker. This exercise illuminates the information needed to (a) confirm that separation under non-linear price menus can be sustained empirically,[18] and (b) find the optimal price menu. This informs our

[16]The ratio of usable energy per weight is around 1.6: for each kg of LPG used in cooking, around 1.6 kg of charcoal would be needed to generate the same energy.

[17]Calculations can be found here. We adjust for the Floess et al. (2023) Ghana-specific assumed fraction of non-renewable biomass used in charcoal production of 27.7%.

[18]Given that we show that separation can be sustained empirically, Section 5 discusses sufficient conditions

data collection strategy. We then develop a sufficient statistic approach that summarizes all the heterogeneity relevant for LPG cylinder size selection, and that can be elicited using an MPL exercise. Besides simplifying the evaluation of the objective function, this approach ensures that we are not missing or mis-measuring relevant sources of heterogeneity (e.g., characteristics not in Table 1).

3.1 Policy maker's optimization problem

We assume that the policy maker cares about increasing demand for LPG (in kg) at the lowest public cost possible. Thus, the policy objective function can either maximize demand subject to a given budget, or minimize the budget, subject to a given demand target. Next, we define the demand and budget functions in order to discuss the trade-offs faced by the policy maker when changing prices in each of the sizes of cylinder. We limit the menu to two sizes: the largest (14.5 kg) and the smallest (3 kg). [19]

We summarize all consumer heterogeneity into a single dimension θ. We assume a continuum of types, where $F(\theta)$ and $f(\theta)$ are the CDF and PDF of θ. Faced by quantity-specific exchange prices $p(q_S)$ and $p(q_L)$ (i.e., the price for exchanging an empty cylinder for a full one), type-θ agent chooses the small or large cylinder size $q \in \{q_L, q_S\}$, which leads to exchange frequency x. Define $\theta' = h(p(q_L), p(q_S))$ as the type that is indifferent between the two sizes.

for separation according to theory.

[19]Determining the optimal number of options in the menu is not a trivial exercise. The theoretical result that gets closest to our application is Nikzad (2023). However, his result applies only to settings where there is a single allocation (quantity) and constraints are linear, which is not our case.

Per-household average demand for LPG in kg is given by

$$
\begin{aligned}
Q\left(p(q_L), p(q_S)\right) \quad = \quad & \underbrace{F(\theta')}_{\substack{\text{share that} \\ \text{chooses } S}} \quad \underbrace{q_S}_{\substack{\text{converts} \\ \text{into kg}}} \quad \underbrace{\int_{-\infty}^{\theta'} x^*(p(q_S); \theta, q_S) \frac{f(\theta)}{F(\theta')} d\theta}_{\substack{\text{avg. \# of } S \text{ exchanges} \\ \text{for those that select into } S}} \\
+ \quad & \underbrace{(1 - F(\theta'))}_{\substack{\text{share that} \\ \text{chooses } L}} \quad \underbrace{q_L}_{\substack{\text{converts} \\ \text{into kg}}} \quad \underbrace{\int_{\theta'}^{\infty} x^*(p(q_L); \theta, q_L) \frac{f(\theta)}{1 - F(\theta')} d\theta}_{\substack{\text{avg. \# of } L \text{ exchanges} \\ \text{for those that select into } L}}
\end{aligned} \tag{1}
$$

where $F(\theta)$ and $f(\theta)$ are the CDF and PDF of θ.[20] The average public expenditure per household is given by an expression with a similar structure:

$$
\begin{aligned}
B\left(p(q_L), p(q_S)\right) \quad = \quad & F(\theta') \quad (r \times q_S - p(q_S)) \quad \int_{-\infty}^{\theta'} x^*(p(q_S); \theta, q_S) \tfrac{f(\theta)}{F(\theta')} d\theta \\
+ \quad & (1 - F(\theta')) \quad (r \times q_L - p(q_L)) \quad \int_{\theta'}^{\infty} x^*(p(q_L); \theta, q_L) \tfrac{f(\theta)}{1-F(\theta')} d\theta
\end{aligned} \tag{2}
$$

where r is the cost per kg of LPG to the policy maker.

Equations (1) and (2) illustrate the trade-offs that the policy maker balances when choosing $p(q_L)$ and $p(q_S)$ optimally. For example, assume a fixed public budget. As more of that budget is allocated towards subsidizing LPG in the small cylinder relative to the large cylinder, two margins will change. First, the number of large cylinder exchanges will fall, while the number of small cylinder exchanges will increase. If the low θ types are more price-responsive than the high θ types when exchanging small cylinders, this will increase average demand. Second, more individuals will select into the small cylinder (θ' will increase). The sign of the effect of this margin on demand depends on the sensitivity of the switchers to

[20] Note that we limit the household decision to choosing either a small or large, and do not consider the decision to choose none. Given that this extensive margin decision is accounted for by the corner solution, $x^*(p(q_S); \theta, q_S) = 0$, this is an innocuous omission, provided that the cost of the cylinder and stove is zero. In practice, we find that even when requiring a deposit for the cylinder and stove (which did not vary with size of the cylinder), take up of a cylinder of either size in the multiple price list was 96%. As shown in Table A.1, around 19% of households could not be found or were not able to complete the deposit at the time of delivery. Coding delivery failures as zero exchanges, 44% of households that took up had no exchanges over the following 3 months. This suggests the simplified choices implicit in (1) are a good approximation of household behavior.

prices, transaction costs and savings costs. For example, if switchers have high transaction costs, their demand level will fall when they switch to a small cylinder.

An advantage of restricting the policy maker to care about (1) and (2) is that both of these quantities are empirically observable outcomes in stage 2 of our study. However, these functions ignore the consumer surplus and deadweight loss (DWL) associated with a given menu of prices. In Section 6 we provide a back of the envelope estimate of the broader welfare impacts of non-linear pricing.

3.2 A sufficient statistic approach

Finding the $p(q_L)$ and $p(q_S)$ that solve the policy maker's optimization problem described by (1) and (2) based on observable individual characteristics is complicated both by the potentially numerous sources of heterogeneity summarized in θ and by the difficulty in measuring them. Thus, instead of estimating a fully specified demand model, our empirical approach relies on a sufficient statistic and is agnostic about the sources of heterogeneity that drive cylinder size choices. A disadvantage of this approach, relative to a fully-specified structural model, is that it cannot be used to evaluate whether separation is guaranteed for a given population based on its sample characteristics alone. It also limits the counterfactual analyses we can perform. An advantage of this approach is that it can capture unmodeled and unmeasurable drivers of separation.[21] To better understand the drivers of separation in our setting, and to explore external validity, Section 5 shows a plausible microfoundation for consumer heterogeneity that is consistent with our empirical results.

Our empirical approach relies on two shortcuts. First, we directly elicit a sufficient statistic for the relevant heterogeneity θ that governs selection into cylinder size. Second, we estimate a flexible average demand function conditional on selection. This approach allows us to evaluate equations (1) and (2) for each price pair, $(p(q_L), p(q_S))$, and conduct a numerical search for the optimal menu. Here, we summarize the theory behind the measurement of θ in the first shortcut (see Appendix C for additional detail). Specifically, we show how all the relevant information for selection across cylinder sizes can be summarized by an intuitive

[21]In our setting, we show that the bulk of heterogeneity that drives separation is not captured by the household characteristics we measured (see Section 4.2 under Step 1).

economic quantity: the difference in maximum willingness to pay between the two sizes. Our strategy for reducing the dimensionality of the relevant heterogeneity is similar to the aggregation technique described by Rochet and Stole (2003). We describe implementation of both shortcuts in Section 4.2.

An individual of type θ gets indirect utility $v(p(q);\theta,q)$ from purchasing LPG in a cylinder of size $q \in \{q_L, q_S\}$. The condition that governs selection into the large size is then given by

$$v(p(q_L);\theta,q_L) - v(p(q_S);\theta,q_S) > 0 \tag{3}$$

The sufficient statistic for consumer type emerges from the selection condition in equation (3). We approximate the left hand side of this inequality using a Taylor expansion, which delivers:[22]

$$-\frac{\left.\frac{\partial v(p(q);\theta,q))}{\partial q}\right|_{q_S}}{\left.\frac{\partial v(p(q);\theta,q))}{\partial p(q)}\right|_{q_S}} > \frac{p(q_L)-p(q_S)}{q_L-q_S}. \tag{4}$$

This condition is intuitive: it states that for the individual to choose the large cylinder, the marginal willingness to pay (MWTP) evaluated at the quantity in the small cylinder needs to be larger than the per kg difference in prices between the two sizes.

Next, we show that the MWTP in (4) can be expressed as the difference in WTP for exchanging the large and the small cylinders. This yields an even simpler approximation to the choice condition in (3):

$$p(\theta,q_L) - p(\theta,q_S) > p(q_L) - p(q_S). \tag{5}$$

Inequality (5) states that individuals will choose to purchase LPG in the large cylinder when this choice delivers a larger surplus compared to the surplus from choosing the small. Given this, the share of individuals who choose the large cylinder under (exchange) prices $p(q_L), p(q_S)$ is given by

$$S_L = \int_{-\infty}^{\infty} \mathbf{1}\left(p(\theta,q_L) - p(\theta,q_S) > p(q_L) - p(q_S)\right) f(\theta)\theta.$$

[22] A second-order Taylor expansion delivers similar empirical results so we implement a simpler first-order expansion.

Note that this expression depends on θ only through the difference in WTP across sizes, $p(\theta, q_L) - p(\theta, q_S)$. Thus, this difference is a sufficient statistic for selection. To simplify the discussion below, we denote this sufficient statistic as

$$\tilde{\theta} \equiv \tilde{\theta}(\theta) = p(\theta, q_L) - p(\theta, q_S).$$

4 Two-stage field experiment

Our empirical application proceeds in two stages: design (stage 1) and evaluation (stage 2). We first provide background on the sample and data collection. We then turn to the two stages of our field experiment.

4.1 Sample, data and implementation

In this section we discuss the details of implementation, many of which are common to both stages of our field experiment.

Sample Our study is implemented in Techiman, a city of around 250,000 in the Bono East region of central Ghana. Techiman is the site of a Health and Demographic Surveillance System (KHDSS) run by the Kintampo Public Health Research Centre (KHRC). KHRC is one of three national health research centers under the Ghana Health Services, Ministry of Health. The KHDSS covers over 500,000 individuals across six districts, and collects regular (approximately every six months) household data.

We drew stratified random samples for both stages using KHDSS data. First, we constructed a sampling catchment around four selected LPG cylinder exchange depot sites and included households more than 150 meters and less than 1 kilometer away from the depot. Second, we restricted eligibility to households that were primary charcoal users. Third, we stratified the sample on wealth (a durable assets index) and exchange depot, and assigned households to the two stages of the experiment.[23]

[23]The original stratification was on wealth quintile and depot. After revised power calculations based on data from stage 1 indicated that we required an expanded sample size for stage 2, we repeated the stratification combining quintiles 1 and 2 and quintiles 3 and 4 into larger wealth groupings. This was

During the baseline surveys (one per stage), we further excluded households that met any of the following screening criteria: (a) financial decision maker not available, (b) planned to move in next six months, (c) household contained more than 9 members, (d) household already used LPG as their primary fuel, or (e) household prepared food commercially.

Data and implementation Study data came from the following sources.

KHDSS household data. We used 2021 KHDSS data for sampling and randomization. The dataset is a census of households in Techiman that includes information on household size, a wealth index, and primary cooking fuel.

Household surveys. During each stage of the study, we gathered baseline survey data as part of the enrollment process. Surveys were conducted with the head of household or another financial decision maker. The questionnaire was administered by a trained enumerator, and covered household demographics, finances, time and travel costs, and fuel use and expenditures. Cylinder size choices were elicited at the end of the household survey.

At the end of the three month exchange window in stage 2, we administered a short endline survey to measure self-reported fuel use and expenditures.

Cylinder size choices. The enrollment process involved a decision of whether to take up an LPG "starter kit" (stove, empty cylinder with an ID code and regulator/hose). Households were required to make a deposit for these items, which did not vary with cylinder size.[24] In stage 1, households were randomly assigned to a deposit of GH¢ 50 or 100; in stage 2, it was set to GH¢ 50 for everyone.[25] At the time of the take up decision, participants received information about the cylinder exchange program and their subsidized LPG offer. They were also told that they would have a chance to switch cylinder sizes at the end of the 3-month

justified by data from stage 1 that showed a relationship between wealth and LPG demand only in the top quintile. Thus, stage 1 has 20 strata (4 depots and 5 wealth quintile) while stage 2 has only 12 (4 depots and 3 wealth groups). To correct for any potential sample differences across stages, we reweight our stage 2 results to match our stage 1 strata weights.

[24] In the market, at the start of stage 1, an empty 3 and 14.5 kg cylinders cost GH¢ 145 and 245, respectively. A 2-burner stove cost around GH¢ 122. We collected a deposit that did not depend on size to avoid confounding selection based on durable good costs with selection based on repeated expenditures.

[25] The variation in stage 1 was implemented as a source of exogenous variation in cash on hand liquidity, but we observe little effect on cylinder size preferences. The lower deposit increased compliance with the take up decision (stove delivery success), so we used it for all offers in stage 2. We also randomly assigned households in stage 1 to either their closest or second closest depot to generate variation in transaction costs; this too was dropped in stage 2 in favor of assigning everyone to their closest depot.

exchange window if they wished to keep the starter kit and forgo their deposit.[26] Households were given up to two weeks to gather their deposit before the starter kit was delivered.[27]

Stage 1 choices. Designing the optimal non-linear subsidy requires estimates of demand, accounting for selection, at each size-specific (small, large) price pair. We gathered the necessary data for this exercise from two main sources. First, a baseline survey for stage 1 included a multiple price list (MPL) WTP elicitation that also introduced random variation into cylinder sizes and prices. Second, we observed household level LPG purchases for three months, at assigned sizes and prices.

The MPL was implemented as follows (see Online Supplement I for additional detail). Each participant completed three different MPLs: Small cylinder vs Nothing (MPL A), Large cylinder vs Nothing (MPL B) and Small cylinder vs Large cylinder vs Nothing (MPL C). MPLs A and B each included eight binary choices between one of the cylinders and nothing, where the choices varied in the price of a cylinder exchange. The first four choices were the same for all participants; responses determined values for a second set of four choices, which used smaller intervals for a more precise measure of WTP. Choices in MPLs A and B determined the content of MPL C, which elicited preferences across three options: small, large, and nothing. Prices of the large cylinder exchange were anchored at 80-100% of the maximum WTP from MPL B and the exchange price of the small cylinder varied around the maximum WTP from MPL A. MPL C included only four rows.

After participants made all choices in all three MPLs, one row was drawn for implementation. This introduced random variation in the cylinder size and price combination, conditional on choices. When combined with the exchange data, this variation allows us to estimate demand at a wide range of price pairs, accounting for selection.

We observe a high degree of non-switching behavior on the extensive margin: 43% of participants prefer a cylinder over nothing in all choices. This is not altogether surprising;

[26]Allowing participants to switch sizes at the end of the exchange window ensured that choices were based on preferences over sizes at study prices, rather than on the long run value of the asset or preferences over sizes at future market prices.

[27]The fact that households could renege on their initial cylinder choices at the time of delivery (by not being home, for example) has the potential to interfere with the study design. In both stages, small cylinders had a slightly higher likelihood of delivery failure. The randomly assigned LPG price (stage 1) and treatment (stage 2) are statistically unrelated to delivery failures (see Table A.1). Our main results are robust to including or excluding households that reneged on their choices.

the price variation in the MPL was constrained by the market price of LPG since filling stations continued to operate during the LPG exchange phase of the project.

Stage 2 choices. Data for stage 2 are considerably simpler than for stage 1. Cylinder selection decisions comprised a single choice between the large and the small cylinder, where the cylinder exchange prices differed by treatment, but other aspects of the choice did not.

LPG cylinder exchanges. Households that acquired a cylinder could access LPG through their assigned exchange depot at their assigned price for a period of three months. Depots were supplied by our LPG partner, Andev, and consisted of a metal cage where both full and empty cylinders were stored. Depot managers completed exchanges, collected payments, tracked inventory and liaised with the study team. Managers used tablets to record purchase details, including participant study ID, cylinder ID, price and time of the transaction.[28]

During the three month study window (per stage), households could make as many exchanges as they wished, paying cash to the depot manager. They were, however, limited to exchanging a single cylinder at a time and to the cylinder size that they selected during enrollment.[29] Our analysis assumes that we observed all LPG demand during the three month exchange window. Households were not prohibited from visiting filling stations (prohibitions would have been unenforceable). We conducted random spot checks to assess filling outside of the exchange depots. In 79 unannounced visits, we observe no evidence of unmeasured LPG demand (see Supplement III).

LPG prices proved to be very volatile during implementation of stage 2. The design of the optimal price menu relied on prices used in stage 1. By the time stage 2 launched, in February 2023, LPG prices had risen by around 25%. At the start of stage 2, we therefore inflated the prices from stage 1 to account for both real and nominal price increases (see Supplement III for additional detail on our adjustments and implications). Alternative strategies for adjusting for inflation have little effect on the optimal price menu. After stage 2 launched, prices fell to the point that the price of LPG in the large cylinder in the treatment arm exceeded the price in the market. To mitigate the risk of purchases outside of the exchange

[28] To check adherence to study protocols, we conducted "mystery shopper" visits to the exchange depots. We found no cases in which depot managers allowed someone to purchase at a price that differed from their assigned price.

[29] These implementation details increased our ability to generate separation in the stage 2 treatment group. We discuss the potential to place restrictions on cylinder ownership and exchange in Section 6.

depots, we deflated prices back to stage 1 levels, following prices in the LPG market. In our analysis, these changes primarily affect the subsidy budget, so we show robustness to several alternative cost calculations. Our main analysis uses "nominal" prices as experienced by participants.[30] In Section 6, we discuss how a policy maker could accommodate price volatility and other time varying determinants of demand and costs.

Summary statistics and randomization We next describe the sample of households assigned to stages 1 (N = 543) and 2 (N = 915). Over half of households in the sample were female headed and the mean age of the household head was around 49 years. Households contained an average of 4.9 individuals. These characteristics were all similar across the two stages (see Table A.2). Wealth and wealth-related characteristics (liquidity constraints and baseline charcoal purchasing patterns) did, however, differ between the two stages, with a slightly higher average wealth in the second stage due to the change in how strata were defined between stages (see footnote 23).

Next, we examine randomization outcomes from stages 1 and 2. In stage 1, cylinder size and LPG prices were assigned randomly through the MPL. By design, the MPL choice set updated dynamically. We thus focus our balance test on the subset of choices viewed by all participants, which we refer to as the "static" choices. We see imbalance in one pairwise comparison of choices in MPL A, but no other significant differences (pairwise or joint) in the randomization outcomes in stage 1 (see Figure A.1).[31] In stage 2, households were randomly assigned to treatment (non-linear price menu) and control (linear price menu), stratified by LPG exchange depot and wealth. Conditional on strata fixed effects, all differences in means are small and insignificant (see Table A.3).

[30]We multiply stage 1 prices by 1.338 and apply to all transactions from February to June, then return to stage 1 prices for all stage 2 purchases from July to September. This affects both the prices charged to consumers and the calculation of the subsidy expenditure per exchange.

[31]Figure A.2 shows the frequency distribution of all prices and cylinder sizes, which is a function of both the randomized MPL draws and participant preferences. We conduct robustness checks on the demand elasticities resulting from the price variation to assess the effect of selection in the MPL.

4.2 Stage 1: Design

In the first stage of the experiment, we designed the optimal non-linear price menu following six steps. Notation follows Section 3, except that, from here on, we denote $p(q_S)$, the exchange price for the small cylinder, as p_S and $p(q_L)$ as p_L.

Step 1: Compute the sufficient statistic for each household. In the first step, we use the MPL data to approximate the difference in willingness to pay across cylinder sizes, which is the sufficient statistic for selection, $\tilde{\theta}_i = p(\theta_i, q_L) - p(\theta_i, q_S)$, for each household in the data. We denote this approximation as $\hat{\tilde{\theta}}_i$. Specifically, when we observe switching in MPLs A and B and/or MPL C, the WTP is interval-identified: $p(\theta_i, q_S)$ from MPL A, $p(\theta_i, q_L)$ from MPL B and $p(\theta_i, q_L) - p(\theta_i, q_S)$ from MPL C. In other words, we have two separate ways of identifying $\tilde{\theta}_i$. This redundancy is important in cases when there was non-switching behavior in either MPL A or MPL B, leaving $p(\theta_i, q_S)$ or $p(\theta_i, q_L)$ unidentified. In these cases we approximate $\tilde{\theta}_i$ using data from MPL C. See Supplement I.3 for additional detail.

This procedure yields $\hat{\tilde{\theta}}_i$ for all observations, including those with never switching behavior or choices that were inconsistent across MPL modules.[32] Figure 1 shows the distribution of $\hat{\tilde{\theta}} - (p_L - p_S)$, or the approximation of $p(\theta_i, q_L) - p(\theta_i, q_S) - (p_L - p_S)$, by MPL C choices. $\hat{\tilde{\theta}} - (p_L - p_S)$ also represents the relative surplus associated with choosing the large cylinder for different price-type combinations in MPL C. The figure shows a clear separation: most large cylinder choices are associated with a positive surplus and small cylinder choices with a negative surplus. Given that we use MPL A and B to proxy $p(\theta_i, q_L) - p(\theta_i, q_S)$ in cases where $p(\theta_i, q_L) - p(\theta_i, q_S)$ from MPL C is also available, this result is not mechanical.

All else equal, households who gravitated towards the large cylinder (have a higher $\hat{\tilde{\theta}}$) are younger and wealthier, less liquidity constrained, and have a lower baseline likelihood of purchasing charcoal in small quantities (see Table A.4). However, even with detailed survey data and geographic fixed effects, we can explain only around 10% of the variation in the WTP responses as measured by the R-squared from an OLS regression.[33] This limits

[32] We omit from the analysis 5 households for which the WTP difference is not available from MPLs A and B and who presented odd switching behavior in MPL C. Jack et al. (2022) provide a more general discussion of diagnosing and accommodating subject errors and truncated values in MPL data, consistent with the approaches adopted here.

[33] This echoes the inability of observables to explain preferences over cash transfer timing in rural Kenya

Figure 1: $\hat{\tilde{\theta}} - (p_L - p_S)$ and MPL choices

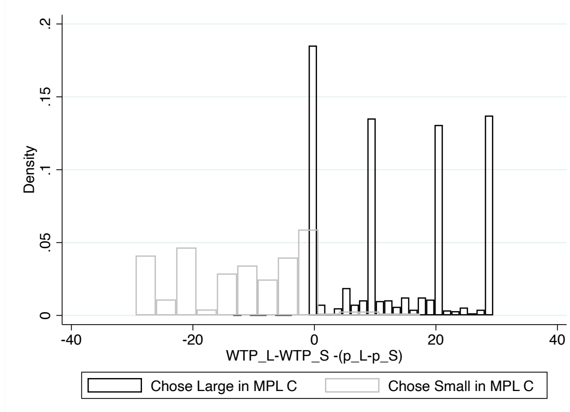

Note: Horizontal axis is the surplus from choosing the large cylinder size, or $\hat{\tilde{\theta}} - (p_L - p_S)$. Black bars are participants who chose the large cylinder in MPL C at the relevant price difference; gray bars are participants who chose the small cylinder.

the feasibility of targeting on observables using, for example, a two-stage approach similar to Johnson and Lipscomb (2022). It also suggests that a fully structural approach, which would rely on measuring different observable sources of heterogeneity, would perform poorly relative to our sufficient statistic approach.

Step 2: Estimate logit for choice of size. We observe 31 cases where choices were inconsistent either within or between MPLs. To allow for errors, and rationalize these cases, we add a random disturbance to the true indirect utility. We use a logit model to approximate the probability that a given individual chooses the large cylinder, conditional on his or her sufficient statistic for selection type, $\tilde{\theta}$, and the price menu she faces, (p_L, p_S). We fit this logit model to the observed choices in MPL C, under the assumption that the marginal utility of income is constant. The probability of choosing the large cylinder can then be written as:

$$\Pr\left(v(p_L;\theta,q_L) - v(p_S;\theta,q_S) + \varepsilon_L - \varepsilon_S > 0|\hat{\tilde{\theta}}_i\right) = \frac{\exp\left(\gamma\left(\hat{\tilde{\theta}}_i - (p_L - p_S)\right)\right)}{1 + \exp\left(\gamma\left(\hat{\tilde{\theta}}_i - (p_L - p_S)\right)\right)} \qquad (6)$$

where $\gamma = \frac{\partial v(p(q);\theta,q))}{\partial p(q)}\Big|_{q_S} \frac{1}{\sigma_\varepsilon}$.

(Kansikas et al., 2023). We also explored prediction using a random forest model, which offered little improvement (see Figure A.3).

Step 3: For a grid of potential price pairs, create selection weight, $\hat{w}_i$. We create selection weights to reflect the probability of choosing the large cylinder as a function of $\hat{\hat{\theta}}_i$ over a grid of price pairs. Averaging these weights over the whole sample generates a smooth predicted selection function as $(p_L - p_S)$ varies.

Specifically, using the estimate for γ that accounts for deviations in the indirect utility over time, and the formula in (6), we create weights for each observation in the sample, for each value of $(p_L - p_S)$ in the grid:

$$\hat{w}_i\left(p_L - p_S\right) = \frac{\exp\left(\hat{\gamma}\left(\hat{\hat{\theta}}_i - (p_L - p_S)\right)\right)}{1 + \exp\left(\hat{\gamma}\left(\hat{\hat{\theta}}_i - (p_L - p_S)\right)\right)} \tag{7}$$

Figure 2a shows the selection into cylinder size as a function of the price difference $(p_L - p_S)$ by averaging (7) over the whole sample. As the exchange price of the large cylinder grows, relative to the small (i.e., moving right on the horizontal axis), a smaller share of the sample selects into the large cylinder.

Figure 2: Selection and elasticities

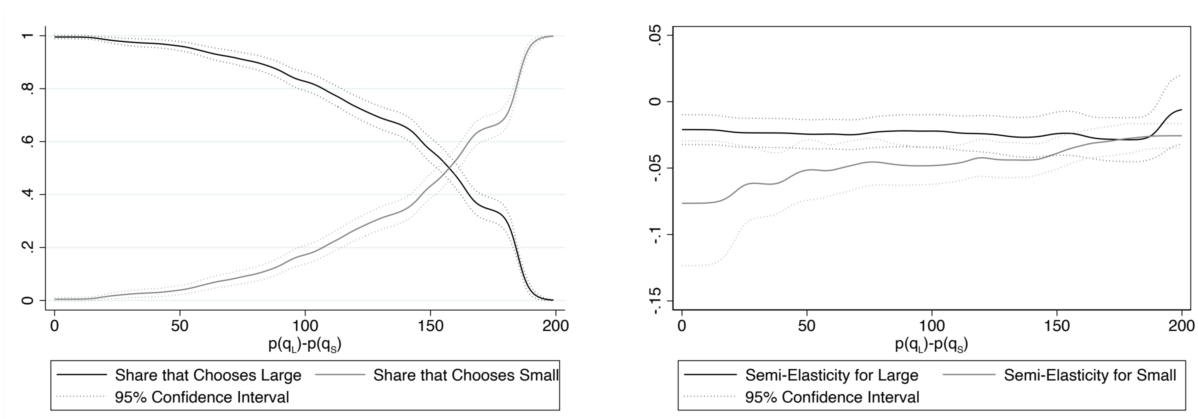

(a) Weighted selection probabilities (b) Predicted semi-elasticities of demand

Note: Panel A shows the predicted selection probabilities as a function of prices. The vertical axis is the share of the sample choosing each cylinder size. Panel B shows the predicted semi-elasticities of demand, incorporating selection. The vertical axis is the semi-elasticity of kg per month with respect to the price, for each cylinder size, evaluated for the sample that selects into each cylinder size at the corresponding price difference on the horizontal axis. 95% confidence intervals computed using bootstrapped standard errors from 1,000 random samples.

Step 4: Use WLS to estimate LPG demand at each price pair. Next, we use these weights to incorporate selection into predicted demand at each price pair, assuming a quadratic fit, consistent with our data (see Figure A.4).[34]

Denote the number of observations for which we observe x_{Si} as N_S and the number of observations for which we observe x_{Li} as N_L. We index each value of $(p_L - p_S)$ in the grid by k. We then estimate a different set of demand parameters for the large cylinder for each k, $\boldsymbol{\beta}_k^L = [\beta_{0k}^L, \beta_{1k}^L, \beta_{2k}^L]$, by weighted least squares (WLS):

$$\min_{\boldsymbol{\beta}_k^L} \sum_{i=1}^{N_L} \hat{w}_{ik}(x_{Li} - \beta_{0k}^L - \beta_{1k}^L p_{Li} - \beta_{2k}^L p_{Li}^2)^2.$$

We use an analogous procedure to estimate the demand parameters of the small cylinder, $\boldsymbol{\beta}_k^S$, where weights are given by $1 - \hat{w}_{ik}$.

WLS should deliver a consistent estimate of the average of price coefficients across the selected sample (Solon et al., 2015).[35] The average predicted demand for cylinder exchanges of size $Z \in \{S, L\}$, reflecting the selection weights, is then

$$\hat{x}_k^Z(p_Z) = \hat{\beta}_{0k}^Z + \hat{\beta}_{1k}^Z p_Z + \hat{\beta}_{2k}^Z p_Z^2$$

The first thing to notice in Figure 2b is that semi-elasticities of demand are lower (more price sensitive) for the small cylinder than for the large. This is true for all price differences, but particularly so when the exchange price difference between them is small, i.e., a higher price per kg in the small cylinder. The small number of participants who still chose the small cylinder were very price sensitive. As $(p_L - p_S)$ increases (moves right along the horizontal axis), the share of the sample choosing the large cylinder declines and becomes

[34]Interpreting this as a price elasticity and imposing linearity in demand gives $\varepsilon_p = -1.12$ in the small cylinder and $\varepsilon_p = -0.97$ in the large cylinder. Recall that the price and size were randomly assigned conditional on WTP. Table A.5 decomposes the demand elasticities into the extensive and intensive margins. We test whether price sensitivity in LPG demand depends on whether the price is endogenous (prices from the dynamic part of the MPL). Neither average price sensitivity nor size-specific price sensitivity depends on whether the participant received their cylinder through the static or dynamic portion of the MPL. This implies that the additional selection imposed by the dynamic design does not meaningfully affect the demand estimates used in the menu design.

[35]Solon et al. (2015) specify that a condition for consistency is that the variance of the independent variables is constant across selected samples. In principle, this should be the case given that prices were randomly assigned. We examine this in Table A.6 and discuss implications in Section 4.3.

more selected. The relatively flat line plotting the semi-elasticity among those selecting into the large suggests that selection has a relatively minor effect on semi-elasticities of demand for the large cylinder. On the other hand, moving right along the horizontal axis moves a larger share of less price sensitive people into preferring the small cylinder, decreasing price sensitivity, on average, among those selecting into the small. This indicates that subsidizing the small cylinder, up to a point, may increase demand because of the greater price sensitivity among small cylinder types for most price pairs. The total quantity demanded at each price pair also depends on how cylinder size-specific non-monetary costs affect demand levels.

Step 5: Evaluate the policy maker's objective function and budget constraint at each price pair. We evaluate the policy maker's objective function (aggregate demand) for each point of the $[(p_L - p_S), p_S]$ grid as

$$\left(\left(\frac{1}{N}\sum_{i=1}^{N}\hat{w}_{ik}\right) \times \hat{x}_k^L(p_L) \times q_L\right) + \left(\left(\frac{1}{N}\sum_{i=1}^{N}(1-\hat{w}_{ik})\right) \times \hat{x}_k^S(p_S) \times q_S\right) \tag{8}$$

Then, we evaluate the policy maker's budget at each value of the $[(p_L - p_S), p_S]$ grid by substituting q_L with $(q_L r - p_L)$ and q_S with $(q_S r - p_S)$ in (8).

Figure 3 shows how cylinder exchange prices affect the subsidy budget and objective function. As prices increase (move away from the origin), the budget (heat map) and demand (contour lines) both decrease, on average. The figure shows that (a) at high budgets (low prices) the policy maker typically cannot do better than the linear price menu (contours and budget shading are parallel) and (b) with smaller budgets, there exist some points where a non-linear menu can deliver the same budget, but higher predicted demand. The small budget (modest demand increases) scenario is more relevant for our specific context.

Step 6: Choose non-linear price menu and linear prices that deliver same budget level and calculate sample size needed to detect differences in demand. The non-linear price menu we implemented is labeled on Figure 3. Two counterfactual linear price menus are also shown. Linear price A shows a budget-neutral alternative to our chosen non-linear price menu. At the same budget, it is predicted to deliver lower demand (contour

Figure 3: Welfare and budget by price pair

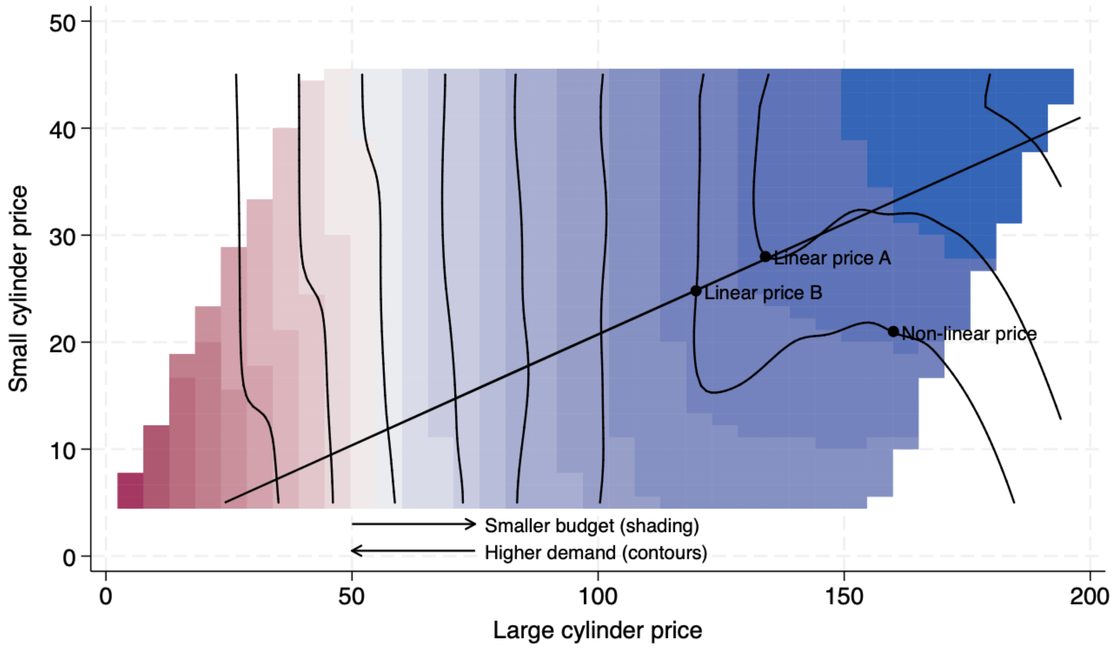

Note: Shaded heatmap shows subsidy budget, which falls as prices move away from the origin. Contour lines show the objective function – total demand – which is decreasing as contours move outward from the origin. The straight line shows the linear price (equal price per kg) points. Linear pricing A and B show alternative counterfactuals for non-linear pricing that hold fixed the budget and demand, respectively.

lies to the right). This is the linear price "control" that we implemented in stage 2. Linear price B is predicted to deliver the same level of demand (same contour) as the non-linear price menu, but at a higher budget.[36]

Our chosen menu corresponds to a budget of approximately GH¢ 14 per participant.[37] The predicted mean difference in LPG demand per household per month is 0.59 kg. Power calculations also depend on the variance of demand, accounting for selection, under each treatment arm. Online Supplement II provides the details on how we estimate the variance using an algorithm similar to that used for demand. Detecting a 0.59 kg difference in demand

[36] Results from stage 1 predict that non-linear pricing increases demand by 17.5% or reduces the budget by 29.4% relative to linear pricing counterfactuals A and B, respectively.

[37] The objective function and budget evaluation were carried out using a grid search over price combinations. As such, for budget neutrality we required that the per-household subsidy budget for the linear and non-linear price menus were within 0.50 GH¢ of each other. Within each budget group – e.g., price menus that fall within this maximum budget difference – we selected the non-linear menu that maximized the objective function and the linear menu that minimized the difference in the per kg price in the small and large cylinder. Finally, to facilitate exchanges using cash, exchange prices were set to whole numbers, resulting in a small difference in the kg price across cylinder sizes in the linear price control: 9.33 vs 9.24 for the small and large cylinder, respectively.

Table 2: Stage 2 study arms and menus

Study Arm	Cylinder	Exchange Price	Unit Price
Control (Linear Price)	3kg	GH¢ 28	GH¢ 9.3*
Control (Linear Price)	14.5kg	GH¢ 134	GH¢ 9.2*
Treatment (Non-Linear Price)	3kg	GH¢ 21	GH¢ 7.0
Treatment (Non-Linear Price)	14.5kg	GH¢ 160	GH¢ 11.0

Note: Unsubsidized per kg price = GH¢ 13.38; 1 USD = 8 GH¢. *To facilitate implementation, exchange prices were set to whole numbers, resulting in a small difference in the per kg price offered to participants in the linear price control: 9.33 vs 9.24 for the small and large cylinder, respectively.

requires a sample size of 828 households. We overshot this target with 915 households in our stage 2 sample to ensure adequate power.

4.3 Stage 2: Test

Stage 2 tests the effect of a non-linear price menu for LPG, relative to linear pricing designed to be budget neutral in expectation. From the policy maker's perspective, the objective of the separating contract is to maximize LPG demand, conditional on a fixed budget. Our primary outcomes are therefore (1) LPG demand per month and (2) subsidy budget per month, which is a function of demand and the exchange prices.[38]

Since the policy maker is interested in cost effectiveness, and since misspecification of the demand model or differences in implementation, sample or time may break the budget-neutrality of the design, we define two joint hypothesis tests for the effect of non-linear price treatment relative to the linear price control:

	Budget	Demand
H1:	$= 0$	> 0
H2:	< 0	≥ 0

[38]In our demand measure, we include all LPG purchased during the exchange window, regardless of whether it was consumed. This implicitly advantages the demand outcome for large cylinder owners since each exchange is approximately 5x the quantity of the small cylinder, leaving large cylinder owners with a higher stock of LPG (conditional on exchange date) at the end of their eligibility window.

Specifically, we predict that non-linear pricing will increase cost effectiveness either through (H1) an increase in demand, holding the budget fixed, or (H2) a decrease in the budget, with the same or greater demand.[39] Table 3 implements the hypotheses tests in an OLS regression specification, with a single cross sectional observation per participant.[40]

Table 3: Stage 2 subsidy budget and LPG demand

	Subsidy budget (GH¢/month)			LPG demand (kg/month)		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment effect	-4.008***	-4.080***	-3.712***	-0.223	-0.243	-0.119
	(1.003)	(0.998)	(0.985)	(0.320)	(0.317)	(0.314)
Control mean	11.585	11.585	11.585	3.089	3.089	3.089
Control weighted mean	12.681	12.681	12.681	3.385	3.385	3.385
Strata FE		×	×		×	×
Surveyor FE			×			×
Enrollment week FE			×			×
Observations	915	915	915	915	915	915

Note: Columns 1-3 show treatment effect on subsidy per month; columns 4-6 on LPG demand per month. Columns 2 and 5 include strata fixed effects; columns 3 and 6 include strata, surveyor and enrollment week fixed effects. Regressions are weighted by stage 1 strata. Unweighted version in Table A.7. * $p<0.1$, ** $p<0.05$, *** $p<0.01$

Columns 1-3 show the effect of treatment assignment on the per household-month subsidy budget and columns 4-6 show the effect on kg of LPG purchased. We observe outcomes consistent with H2: similar levels of demand in treatment and control, but a public budget 31% smaller when prices were differentiated by quantity (column 1, no controls).[41]

Next, we analyze treatment effects on LPG demand by cylinder size, which is a function of both selection and demand. Stage 1 results predict that households that select into the small cylinder will be more price sensitive regardless of treatment. The selection response to the non-linear prices will dampen this effect – high transaction costs of households on the

[39] Our predictions follow our pre-analysis plan and are asymmetric in the strictness of equality. In particular, in H1 we test the stage 1 model predictions, while H2 is more flexible. We implement these tests with and without strata fixed effects and controls for implementation details.

[40] We show results that include all households regardless of starter kit delivery outcomes (see Table A.1); delivery failures are coded as zeros both for demand and the subsidy budget. This is in contrast to the menu design in stage 1, which omits delivery failures. Table A.8 shows that results are largely unaffected by whether these households are in the sample or not.

[41] The treatment effect on demand is also negative but much smaller in magnitude (<7% lower than the control group mean) with a 95% confidence interval of (-0.85, 0.41) kg in column 4.

margin will lead to lower demand frequencies – but the overall effect will still be positive. Our model predicts a positive (but near-zero) effect on the demand for large exchanges, as the (positive) selection effect very slightly dominates the negative price response. As shown in Table 4, our empirical results are consistent with these predictions: LPG demand among small cylinder owners was around 47% higher in treatment than control (column 1); among large cylinder owners, demand was statistically unaffected by treatment. Also, consistent with stage 1 predictions, treatment increased take-up of the small cylinder by around 6.4 percentage points, or 26%.

Table 4: Stage 2 demand by cylinder size and size selection

	LPG demand (kg/month)				Small cylinder	
	(1)	(2)	(3)	(4)	(5)	(6)
Sample:	Small	Small	Large	Large	All	All
Treatment effect	0.531**	0.574**	-0.344	-0.119	0.064**	0.064**
	(0.224)	(0.235)	(0.415)	(0.419)	(0.033)	(0.032)
Control mean	1.105	1.105	3.854	3.854	0.266	0.266
Control weighted mean	1.115	1.115	4.158	4.158	0.245	0.245
Strata FE	×	×	×	×	×	×
Surveyor FE		×		×		×
Survey week FE		×		×		×
Observations	267	267	641	641	915	915

Note: Columns 1-4 show results from linear regression of LPG demand per month on indicator for treatment assignment. Columns 1-2 show results for sample of households with small cylinder, columns 3-4 for sample with large cylinder. Households that chose no cylinder (N=7) are omitted from columns 1-4. Columns 5-6 show results from linear regression of indicator for small cylinder take up on indicator for treatment assignment on full sample. All columns include stage 2 strata fixed effects, even columns also include surveyor and survey week fixed effects. Regressions are weighted by stage 1 strata. Unweighted version in Table A.9. * $p<0.1$, ** $p<0.05$, *** $p<0.01$

While the results support our main hypothesis that non-linear pricing can increase the cost effectiveness of subsidies, the pattern of results – a decrease in public spending, with no change in demand – deviates from the stage 1 predictions. We compare the predictions with the results in Table 5 to better understand the discrepancies.[42]

By design, stage 1's predicted subsidy per household per month is the same in treatment

[42]To facilitate a cleaner comparison, Table 5 includes only households with a successful delivery and stage 2 exchange data pre-price adjustment, i.e., imposes the same sample restriction as used in the stage 1 menu design. See also Table A.8.

Table 5: Model predictions vs stage 2 results

	Control			Treatment		
	(1) Model	(2) S2	(3) Diff	(4) Model	(5) S2	(6) Diff
Subsidy per month	13.867	14.183	0.316	14.220	10.275	-3.945
	(1.788)	(1.120)	[0.778]	(1.216)	(0.764)	[0.000]
Small cylinder share	0.195	0.245	0.051	0.340	0.306	-0.034
	(0.018)	(0.022)	[0.023]	(0.022)	(0.024)	[0.166]
Small exchanges per month	0.789	0.380	-0.409	1.206	0.610	-0.596
	(0.121)	(0.049)	[0.000]	(0.127)	(0.061)	[0.000]
Large exchanges per month	0.248	0.285	0.037	0.284	0.284	0.000
	(0.036)	(0.023)	[0.114]	(0.038)	(0.029)	[0.993]

Notes. Columns (1) and (4) present model estimates. Columns (2) and (5) present data from stage 2, weighted by stage 1 strata. Columns (3) and (6) present results from a two-sided test comparing the stage 2 mean with the model prediction point estimate. Reported values are the mean difference and the p-value for the probability that stage 2 result equals the model mean (with 95% confidence). See footnote 42 for details on the sample.

and control (see footnote 37 for an explanation of minor deviations from equality). In stage 2, the subsidy budget is significantly lower than predicted by the model for the treatment group only (column 6).[43] This is our main result, as shown in Table 3, but it deviates from the stage 1 prediction. The rest of Table 5 helps explain this deviation. First, in the stage 1 model predictions, non-linear pricing increases small cylinder exchanges both by increasing take up of the small cylinder and increasing demand among those who take up the small cylinder. In the stage 2 results, the effect on take up is muted because of higher than predicted take up of the small cylinder in the control (column 3). For the small cylinder, observed demand is proportionately lower – by 48% and 50% in control and treatment, respectively – than predicted demand. Consequently, the predicted demand boost from the small cylinder is largely undone. However, this substantially reduces the subsidy budget in the treatment group without negatively impacting the overall demand for LPG. Large cylinder exchanges per month more closely match predictions, though the slight demand boost predicted in large cylinder exchanges (due to the selection effect slightly dominating the price response)

[43]Our calculation of the subsidy is based on the nominal prices paid by study participants, with model predictions inflated to stage 2 values; these closely follow LPG prices at the pump (see Section 4.1). We consider three alternative calculations of the subsidy budget (see Table A.10): LPG market retail prices at the time of purchase, the time varying refinery price, and inflation-adjusted nominal prices. Treatment effects are of similar or larger magnitudes with these alternative calculations.

is not observed in the results from stage 2.

While we cannot determine the precise cause of the divergence between predicted and observed take up and demand in the small cylinder, we discuss several possibilities. First, inflation between stages 1 and 2 may have increased demand for the small cylinder in the control group and depressed the demand of low-income consumers, who were more likely to select the small cylinder. Second, sample differences between the two stages (see Section 4.1 and Table A.2), beyond what we adjust for based on stratification, may also have contributed to the divergence. Selection may also differ because of differences in how households chose their cylinder size (i.e., in an MPL in stage 1 and in a take-it-or-leave-it offer in stage 2). Finally, as noted in footnote 35, consistency of the demand parameters via WLS requires that the variance of prices is constant across types. However, Table A.6 shows higher variance of prices for higher types – a deviation from the assumptions needed for consistency under Solon et al. (2015) – which might contribute to a divergence between the stage 1 predictions and stage 2 results.

To assess robustness to variation in the price menu chosen in stage 1 (Figure A.5). Specifically, we vary the price per kg in the small and large cylinder to assess whether small deviations from the selected price menu result in large predicted differences in demand or the subsidy budget. Prices close to the ones we implemented always result in a non-linear price menu that outperforms a comparable linear one. We also look in a small window (+/- 10%) around the implemented objective function and budget to see what prices would have been chosen. In this exercise, the small cylinder is subsidized more heavily 85% of the time, and price menus that more heavily subsidize the large cylinder perform worse on predicted cost effectiveness. These results suggest that cost effectiveness is robust to small errors in the choice of prices for implementation.

The public budget savings in the treatment group would support expanding program coverage by roughly 30%, without increasing the LPG subsidy budget. In other words, "demand neutrality" can be converted into budget neutrality by expanding program coverage. However, without a pure control group that received subsidized durables but no subsidized LPG, we cannot assess the effect of program expansion on overall LPG demand.

5 Theory

Our empirical results indicate that a separating equilibrium under bulk premia, i.e., when the small cylinder is discounted relative to the large, can be sustained. This is something of a puzzle given that households can make as many cylinder exchanges as they want within the study period: why don't households pool on the small cylinder in that case? Can the sources of heterogeneity that drive selection into sizes in the charcoal market (where bulk discounts exist) be at play in the LPG market? Inspired by the sources of heterogeneity that correlate with size choice in Table 1, we next establish conditions that are sufficient to sustain a separating equilibrium under bulk premia and under bulk discounts.

Note that our empirical approach to designing the contracts we test in stage 2 does not rely on the specific functional form assumptions that we introduce in this section and is robust to additional sources of heterogeneity not discussed here. However, after presenting the model, we show suggestive empirical evidence that the sources of heterogeneity that drive separation in theory also contribute to separation in stage 2 of our empirical test.

5.1 Separating equilibrium with unconstrained purchase frequency

We ground our theory in our study setting, and assess the potential for lumpy purchases (i.e., a discrete number of fixed package sizes) of LPG to support targeting of public subsidies. Note, however, that lumpiness does not constrain the amount of LPG that can be purchased in a time period of a given length. Individuals that commit to a cylinder size can exchange that cylinder as many times as they want. This is, in fact, a general characteristic of non-perishable consumption goods that are sold in lumpy quantities. A household can consume the same amount of toilet paper per week by either buying a large family package that will be consumed over four weeks or a single small package each week. This possibility of making a continuum of purchases (including a fraction, e.g., a quarter of a family package if time is measured in weeks) over a single period of time rules out that the seller can constrain quantities purchased, a necessary assumption in most of the classic contract theory literature for a separating equilibrium. However, quantity discounts persist in many markets even in the absence of feasible constraints on the frequency of purchase. In their conclusion, Mussa and Rosen (1978) mention heterogeneity in transaction costs and access to capital markets as plausible drivers of separation in durables markets, citing Parks (1974). In line with this insight, we propose a model that incorporates transaction costs and savings costs, and yields a separating equilibrium in the absence of constraints on purchase frequency.

We start with a parameterization of the consumer problem that allows us to incorporate sufficient dimensions of heterogeneity while retaining some tractability. We summarize the model here, and provide details and formal proofs in Appendix B. Modeling savings and purchases of durables (a full cylinder may last for several weeks) generally requires a dynamic framework. For simplicity, we set up a static consumer problem that accounts for the lumpiness of LPG purchases and the presence of savings costs that affect the total costs of consuming LPG. Importantly, we reflect the freedom consumers have in buying any amount of LPG in a given period by adjusting the frequency of exchanges. Thus, we allow the number of exchanges in a given period to be continuous, as opposed to binary.[44]

The consumer's maximization problem in a given period is given by:

$$\max_{c,q,x} u(c, qx; \theta) \tag{9}$$

$$\text{s.t. } c + p(q)x + D(\theta)x + 1(x > 0)\frac{p(q)}{x}A(\theta) = y(\theta)$$

where $q \in \{q_S, q_L\}$ is the size of the LPG cylinder in kg, c denotes consumption measured in money, x is the number of exchange trips and can take any non-negative value (including those between 0 and 1), $p(q)$ is the price of an exchange as a function of cylinder size, $D(\theta)$ is the transaction cost per exchange, $1(x > 0)\frac{p(q)}{x}A(\theta)$ is the savings cost (which we explain in detail below), and $y(\theta)$ is income. Assume that type θ determines income $y(\theta)$, savings costs through $A(\theta)$, and transaction costs through $D(\theta)$. Type may also determine preferences for LPG, such that $\frac{\partial u}{\partial qx \partial \theta} > 0$.

The term that governs the savings cost contains the following components:

$$\underbrace{1(x > 0)}_{\substack{\text{if positive} \\ \text{exchanges}}} \quad \underbrace{p(q)}_{\substack{\text{amount to} \\ \text{hold on to}}} \quad \underbrace{\frac{1}{x}}_{\substack{\text{\# of periods} \\ \text{until exchange}}} \quad \underbrace{A(\theta)}_{\substack{\text{cost per period} \\ \text{per dollar}}}$$

[44]Consistent with our empirical setting, we abstract from the decision to acquire LPG-related durables and focus on cylinder size choice and demand.

The logic behind this parameterization of savings costs is as follows: when the frequency of purchase is less than one (e.g., $x = 0.5$), then the household has to hold on to the quantity of money $p(q)$ for multiple periods until it is time to buy the next cylinder. $A(\theta)$ represents the cost of holding on to one cedi (GH¢) for a full period without spending it. Thus, if the frequency of purchase is $x = 0.5$ per period, the number of periods you have to hold on to that cedi in order to purchase the cylinder is $\frac{1}{x} = \frac{1}{0.5} = 2$. If the cost of saving one cedi per period is 30 cents ($A(\theta) = 0.3$), then the additional cost that is incurred is $p(q) \times \frac{1}{x} \times A(\theta) = p(q) \times 2 \times 0.30$.

In order to make this problem tractable, we first analyze the optimization problem conditional on q, and then analyze the size choice across the discrete number of sizes. Consider the maximization problem conditional on q:

$$\begin{aligned} &\max_{c,x} u(c, qx; \theta, q) \\ \text{s.t. } & c + px + D(\theta)x + 1(x > 0)\tfrac{p}{x}A(\theta) = y(\theta) \end{aligned} \tag{10}$$

The first order conditions to (10) imply

$$\frac{u_2 q}{u_1} = p + D(\theta) - \frac{pA(\theta)}{x^2}, \tag{11}$$

where u_1 is the partial derivative of the utility function with respect to the first argument (non-LPG consumption), and u_2 is the partial derivative of the utility function with respect to the second argument (LPG in kg). Thus, the expression in (11) shows that the optimal number of exchanges is such that the marginal utility of an exchange (left hand side) is equal to the marginal cost of an exchange (right hand side). Note that the marginal cost of an exchange depends on the number of exchanges because of the non-linearity introduced by savings costs. Denote the solution to (10) as

$$x^*(p; \theta, q), \;\; c^*(p; \theta, q). \tag{12}$$

We now turn to the optimal choice across sizes. To do so, we plug (12) back to the utility function in (9) in order to obtain the indirect utility function conditional on size:

$$v(p(q); \theta, q) \equiv u(c^*(p(q); \theta, q), x^*(p(q); \theta, q)q).$$

The condition that governs the choice of the large cylinder over the small one is given by[45]

$$v(p(q_L);\theta, q_L) - v(p(q_S);\theta, q_S) > 0. \tag{13}$$

Single-crossing and separating equilibrium

The two conditions we need for establishing the existence of a separating equilibrium are the satisfaction of the single crossing property (SCP) and indifference between sizes by individuals characterized by a type within the continuum of types.

Proposition 1 We assume that the marginal utility of consumption, u_c, is constant, that $x^*(p;\theta,q)$ is non-decreasing in q, and that $p(q)$ is increasing in q.[46] The last assumption amounts to assuming that exchanging a larger cylinder is more expensive than exchanging a small one. We also assume that $A(\theta) \geq 0$, $D(\theta) \geq 0$, $\frac{\partial A(\theta)}{\partial \theta} = A'(\theta) \leq 0$ and $\frac{\partial D(\theta)}{\partial \theta} = D'(\theta) \geq 0$. Then, $v(p(q);\theta,q)$ is single crossing in (q,θ) and $q^*(\theta)$ is non-decreasing.

The proof can be found in Appendix B. The intuition for the role of costs $D(\theta)$ and $A(\theta)$ is as follows: as θ increases, purchasing in large quantities becomes more attractive because transaction costs are larger ($D'(\theta) \geq 0$) and savings costs are lower ($A'(\theta) \leq 0$).

The SCP guarantees that indifference conditions cross at most once. If they never cross, then there is a pooling equilibrium (where everyone prefers a single size). Proposition 2 discusses how the existence of both a transaction cost, $D(\theta) > 0$, and a savings cost, $A(\theta) > 0$, partitions the support of θ by preference over cylinder size.

Proposition 2 In addition to the assumptions under Proposition 1, assume that agents cannot borrow and that income is non-decreasing in type ($y'(\theta) \geq 0$). Assume also that there is a finite lower bound to transaction costs $D(\theta)$ and a zero lower bound to savings costs $A(\theta)$. Finally, assume that exchanging the large cylinder is unaffordable for some types in the lower part of the support of θ, who can still afford the small cylinder.[47] Then, a separating equilibrium is guaranteed to exist under linear prices or

[45] Note that this selection equation also appears as equation (3) in Section 3, and is used to motivate the sufficient statistic approach.

[46] In Lemma 1 of Appendix B we show that a $x^*(q,\theta)$ that is non-decreasing in q follows from the utility function being single crossing in $(x,-q)$.

[47] These assumptions are formally described in Appendix B.

bulk discounts. Existence under bulk premia requires in addition that the transaction cost, $D(\theta)$, is sufficiently large.

The proof of this proposition can be found in Appendix B. Again, it is helpful for intuition to think about the role of transaction and savings costs. First, note that in the absence of transaction and savings costs – i.e., $D(\theta) = A(\theta) = 0$ – individuals will be indifferent between the large and the small cylinder under linear pricing. This is because the price per kg of LPG is the same for both sizes and there is no other cost of purchasing LPG. Next, introduce a transaction cost, $D(\theta) > 0$, still under linear pricing. Now, everyone experiences a higher cost of purchasing a small cylinder compared to a large cylinder. This would result in a pooling equilibrium where everyone chooses the large cylinder, even if there is heterogeneity in $D(\theta)$. Now, introduce $A(\theta) > 0$. This cost penalizes the purchases of the large cylinders, and more so for those with small θ. Because this cost moves in the opposite direction as $D(\theta)$ as θ increases, there will be a type, θ', for whom the transaction costs and savings costs will exactly balance out. Because of SCP, we know that types above this indifferent type, θ', will prefer the large cylinder, and types below θ' will prefer the small cylinder.

Although this intuition is simplest for the case of linear prices, it also applies to non-linear prices. Once we have established the indifferent type for linear prices, changes in pricing that yield non-linear schedules will change the value of the indifferent type, θ'. E.g., if we make the large cylinder more expensive, θ' will decrease relative to the case with linear pricing.

5.2 Empirical support for the importance of savings and transaction costs

In our theoretical model, savings costs and transaction costs are sufficient for separation. Following our pre-analysis plan, we construct three indices, proxying for transaction costs (stated opportunity cost of time), savings or liquidity costs, and small purchases/wealth (measured by durable assets), along with an aggregate index that combines all three.[48] All are scored such that a higher index value is associated with a theoretically lower θ (low

[48]We deviate from the pre-analysis plan in one way: we exclude access to transport from the opportunity cost of time, since it has ambiguous effects: it reduces the hassle cost of a trip but also makes transportation of bulky cylinders less difficult. Table A.11 shows results for each pre-specified characteristic.

opportunity cost of time, high savings cost and lower wealth). Index construction is described in greater detail in Supplement I.

Table 6: Stage 2 targeting on observables

Index:	(1) Low time cost	(2) High savings cost	(3) Small purchases/ wealth	(4) Aggregate
Panel A: Outcome: Small cylinder				
Treatment effect	0.062*	0.058*	0.055	0.060*
	(0.033)	(0.033)	(0.034)	(0.034)
Index	-0.027	0.092***	0.058**	0.065***
	(0.023)	(0.021)	(0.023)	(0.022)
Treatment × Index	0.071**	-0.048	-0.015	-0.001
	(0.032)	(0.042)	(0.034)	(0.033)
Control weighted mean	0.245	0.245	0.245	0.245
Observations	915	915	915	915
Panel B: Outcome: LPG demand (kg/month)				
Treatment effect	-0.226	-0.215	-0.245	-0.188
	(0.321)	(0.317)	(0.277)	(0.302)
Index	0.088	-0.423**	-0.675***	-0.524***
	(0.206)	(0.193)	(0.230)	(0.202)
Treatment × Index	0.073	0.039	-0.019	0.097
	(0.306)	(0.234)	(0.322)	(0.276)
Control weighted mean	3.385	3.385	3.385	3.385
Observations	915	915	915	915
Panel C: Outcome: Subsidy budget (GH¢/month)				
Treatment effect	-4.020***	-3.961***	-3.411***	-3.607***
	(1.005)	(0.994)	(0.910)	(0.958)
Index	0.386	-1.443*	-2.586***	-1.881**
	(0.765)	(0.738)	(0.858)	(0.750)
Treatment × Index	0.238	0.791	1.352	1.310
	(0.933)	(0.811)	(1.015)	(0.889)
Control weighted mean	12.68	12.68	12.68	12.68
Observations	915	915	915	915

Note: Columns 1-4 show treatment effects for all households. Each column controls for index specified in column header and its interaction with treatment predictor. All regressions are weighted by stage 1 strata. Unweighted version in Table A.12 * $p<0.1$, ** $p<0.05$, *** $p<0.01$

The results in Table 6 align broadly with the theoretical model. First, recall that Proposition 1 requires heterogeneity in either $A(\theta)$ or $D(\theta)$. We observe no clear relationship between the time cost index and take up of the small cylinder in the control group (panel A, column 1), but households with higher savings costs were more likely to take up the small cylinder (column 2). This is consistent with savings costs $A(\theta)$ rather than transaction costs $D(\theta)$ being the more relevant source of heterogeneity in this context.[49] Then, for separation, we need transaction costs to be positive for all households and savings costs positive for some. In other words, all households would have preferred to avoid the transaction costs associated with the small cylinder, but those with particularly high savings costs could not. In the non-linear price menu, the price in the large cylinder went up, pushing some households who otherwise would have taken up the large cylinder to choose the small. The positive coefficient on the interaction term in panel A, column 1 suggests that these households tended to have a relatively low opportunity cost of time.

Panel B of Table 6 shows heterogeneity in LPG demand, inclusive of selection. Note that we do not have a specific prediction for this outcome; we show it for completeness and to facilitate interpretation of the results in Panel C. Panel C tests how the non-linear price menu targeted the subsidy budget based on observable characteristics, through a combination of cylinder size choices and demand responses. The coefficient on the index term in the control group is negative for three out of the four indices, consistent with households with theoretically lower θ having lower overall demand for LPG (and therefore receiving a smaller share of the linear per kg subsidy). However, the non-linear price menu increased their share of the pie: the interaction terms are all positive, though all imprecisely estimated. The targeting through non-linear prices fully cancels the regressivity of linear price subsidies when analyzed using the aggregate heterogeneity index (column 4).

6 Discussion and conclusion

Before concluding, we discuss extensions and caveats to our main findings, including policy implementation, and the broader welfare implications of non-linear pricing.

[49] It is also consistent with our time cost index measure not fully capturing true transaction costs.

6.1 Cost effectiveness and scalability

Our main results offer a proof of concept that non-linear pricing can make subsidies more cost effective; however, scale up requires several additional considerations.

Supply costs, durables and limits on exchange Several implementation details pose potential departures from a scaled-up version of non-linear pricing. First, we suppressed size-specific differences in supply costs. One argument for bulk discounts is that supplying goods in smaller packages is more expensive per unit. In our setting, the cost per kg to supply gas via a small cylinder is nearly 2.5 times the cost to supply it in a large cylinder. If we add this cost to the policy maker's subsidy budget, the total subsidy cost per household per month in the non-linear price treatment increases from GH¢ 8.3 to GH¢ 14.4 per household per month, but the magnitude of the treatment effect changes only slightly, from -3.21 to -2.70 GH¢ (note that these comparisons rely on unweighted results). We calculate that the supply cost difference between a small and large cylinder would need to more than triple (relative to the current difference) to eliminate the budget savings from non-linear pricing.

Second, to avoid selection based on differential cost of durables, we held the starter kit deposit price fixed across cylinder sizes.[50] Subsidized durables may have contributed to high participation in the exchange program (see footnote 24). Asking households to pay market prices would presumably change selection, both into LPG use and across cylinder sizes. Analysis of these effects is beyond the scope of this study; however, we note that cost effectiveness gains from non-linear pricing may be undone by large fixed costs, which may reduce heterogeneity in size preferences among those who take up any LPG.

Third, we limited participants to a single project cylinder, and required that their previous cylinder was returned (empty) in order to acquire a new (full) cylinder. Thus, households could not exchange multiple cylinders at the same time.[51] This falls between two extremes, one in which households are limited to a single exchange per week or month (similar to many settings with non-linear pricing such as phone plans or health insurance), and an-

[50]Recall that we did vary the deposit price randomly in stage 1. Lower deposit prices increased delivery success (see Table A.13) but had little effect on WTP in the MPL.

[51]Relative to other goods, this concern may be less important for LPG, where some of the hassle cost of small package sizes is associated with replacing a cylinder in the home.

other in which households can hold as many cylinders as they wish and exchange multiple cylinders per transaction. Intuitively, more flexibility in quantities makes separation harder to sustain, while greater restrictions may erode overall demand. Our solution of allowing customers unlimited exchanges but limiting their cylinder exchanges to a single cylinder at a time is therefore something of a compromise. This approach is scalable, as demonstrated by "one offer per customer" restrictions on discounts or sales in other contexts.[52] More broadly, the level of quantity restriction can be seen as an additional free parameter for the policy maker that can be adjusted to different contexts and policy goals; we leave exploration of this choice for future work.

Policy implementation Our two-stage approach to designing optimal non-linear pricing is, in principle, scalable. The policy designer would need to collect experimental data on cylinder size preferences and price elasticities for a representative subset of the population of interest and plug these into the sufficient statistic based price menu. These prices could then be rolled out to a much larger population. However, this approach presents at least two important challenges.

First, holding the population fixed, package size choices are likely to change over time for a variety of reasons including financial access, income growth, transportation costs, and other factors. In addition, macroeconomic conditions – including inflation and LPG prices – affect both consumer prices and the subsidy cost. Accommodating inflation, when LPG prices increase in lockstep with other inflation indicators, is straightforward. Conversely, when LPG prices change in ways that diverge from inflation, the subsidy cost may grow at a different pace than nominal prices. This would result in a new optimal price menu. If consumer demand is assumed to be unchanged, this can be identified using the same model and data as used to solve the policy maker's problem prior to the price change. Changes to consumer demand are more challenging and require new data inputs to solve for optimal prices. Note that these issues arise in any study that uses existing data to design or calibrate a policy.

Second, collecting experimental data to design optimal price menus means that the rel-

[52]Alternatively, cylinders could be registered per household to limit ownership, just as many countries do with SIM cards.

evant variation is available by design. However, it requires on-the-ground infrastructure for implementation and is costly if the results have limited external validity over time or across populations. In principle, one could use non-experimental sources of variation in cylinder size selection and LPG demand. Recall that the sufficient statistic approach requires observing $p(\theta_i, q_L) - p(\theta_i, q_S)$. While this may be identifiable with cross sectional or panel data from a market with naturally occurring price variation and observations of cylinder size choices, separately identifying selection and demand elasticities would require additional variation – e.g., in the size-specific price of durables. We leave for future work the exercise of extending the menu design framework to accommodate non-experimental data. Alternatively, structural estimation of a demand model could be used to design a generalizable price menu, based on modeled relationships between observables, size preferences and demand. Our results caution against this approach. In practice, this mapping may be unstable over time or across populations. In our data, we also find that it provides low explanatory power even in-sample.

6.2 Welfare effects of non-linear pricing

To move from our evaluation results to a more comprehensive assessment of the effects of non-linear price menus for welfare in our context, we first examine evidence that more cost effective LPG subsidies result in more cost effective displacement of charcoal use. We then turn to a back of the envelope calibration of the overall welfare effects of non-linear prices relative to linear prices.

Inframarginal subsidies and impacts on charcoal use The primary justifications for subsidizing LPG use – reductions in particulate matter and greenhouse gas emissions – depend on displacement of charcoal use. With neither a treatment effect on LPG demand, nor a pure control group that received subsidized durables but no LPG subsidy, measuring displacement is difficult. Here, we present what we do observe about counterfactual LPG and charcoal use.

We start with counterfactual LPG demand, relative to a control group with no subsidies. By construction, the study sample consists of households that did not use LPG as a primary cooking fuel at baseline, which helps mitigate concerns about inframarginal subsidies. We consider three different approaches to assessing counterfactual LPG demand. First, we use the empirical model with data from stage 1 to simulate demand at zero subsidy. Based on observed demand and the quadratic functional form used in the stage 1 model, demand is predicted to be zero in the absence of any LPG subsidy even if durables are subsidized, i.e., none of the subsidies are inframarginal. Second, we use stage 2 survey data at baseline and endline to measure stated use of LPG before and after the introduction of subsidies. At baseline, 23% of the sample reported using any LPG in the past three months. At endline, 61% of the sample reported using any LPG in the past three months (which coincides with the exchange window).[53] Without information on the intensive margin of demand for users who already had LPG durables at baseline, these survey responses offer a conservative upper bound of 38 to 80% inframarginal subsidies.[54] Third, we reference a separate project that worked with a smaller sample of households from the same sampling frame, where households were randomly assigned to the non-linear price treatment (N=69) or to a pure control that received neither subsidized durables nor subsidized LPG (N=80). At endline, only 10% (1.3%) of the control group reported any (primary) LPG use versus 94% (74%) of the treatment group. This suggests a much lower degree of inframarginality (10.6%) than the before-and-after comparison.

Next, we examine charcoal displacement, starting with self-reported before-and-after comparisons. Between stage 2 baseline and endline, average household monthly spending on charcoal fell from around GH¢ 400 to GH¢ 127, a drop of more than 50% (prices were stable over this period).[55] Given that the average amount of LPG purchased under the program is less than 50% of a typical household's cooking needs, the implied substitution rate is >1. Alternatively, we can quantify how much displacement is needed to justify subsidizing LPG,

[53] This latter number is higher than the share of the sample that made any exchanges at the study depots because some households presumably used their own cylinders outside of the study.

[54] Table A.14 shows how demand varies with ownership of LPG durables. Households that owned LPG durables at baseline did have higher overall demand (column 5) but were not differentially responsive to treatment, suggesting that the upper bound is likely too high.

[55] When calculated separately by cylinder size, displacement rates are higher among types that selected into the small cylinder, per kg of LPG purchased, and differentially so in the non-linear price treatment. This mitigates potential concerns that directing additional subsidies toward LPG in the small cylinder has lower impact on charcoal use than the linear price benchmark.

using the calculation of the external costs of greenhouse gas emissions presented in Section 2. We calculate the cost per ton of avoided CO_2e under the linear price control and non-linear price treatment, for different charcoal displacement rates (see Figure A.6). If displacement is assumed to be 1:1, the subsidies in both treatments are highly cost effective at USD 10 and 14 per ton under non-linear and linear pricing, respectively. As the displacement rate falls, the social cost of carbon (i.e., benefit of avoided emissions) required to justify the subsidies (i.e., cost of avoided emissions) increases, and targeting using non-linear prices can be justified under a lower charcoal displacement rate or a lower social cost of carbon. Net emissions are positive for displacement rates less than 0.22 (i.e., each kg of LPG purchased displaces less than 0.22 energy equivalent units of charcoal). This degree of fuel stacking is highly unlikely at the subsidy levels that we test because the income effect of the subsidy is likely to be very small.

Welfare As discussed in Section 3, the policy maker's objective function is narrow. Specifically, it ignores: (1) consumer surplus, (2) the level and distribution of non-monetary costs, i.e., transaction costs and savings costs, (3) the DWL from subsidies, (4) the marginal cost of public funds and (5) external costs or benefits from pollution. We conduct an empirical back of the envelope exercise using data from stage 2 to gauge the magnitude of these effects (see Table A.15).[56]

First, the difference in consumer surplus among "stayers" (those who always chose small or large) can be approximated from the area under the demand curve for cylinders of each size, as prices go from the linear price to the size-specific non-linear price. We add bounds on the non-monetary costs to "switchers" (those who preferred the large cylinder under the linear price menu and the small cylinder under the non-linear price menu).[57] The population weighted average change in consumer surplus are bounded by -3.65 and -3.44. Second, the DWL can be calculated as the complement of the consumer surplus result. Because the increase in subsidy goes to the relatively price sensitive group (small stayers), their DWL

[56]The stage 1 model can also be used to calibrate consumer welfare, but given the divergence between stage 1 and 2 results, we focus on the stage 2 data here.

[57]We cannot observe demand among switchers so we assume either that their demand goes to zero when they switch from the large to small cylinder, or that their demand stays at the level observed in the large cylinder under the linear price menu.

increases by more per household than the reduction for large cylinder types, where negative values correspond to an increase in the DWL of subsidies.[58] The population weighted average change in DWL is bounded by -0.33 and 0.045.

Finally, we consider two scenarios for the external benefits of non-linear pricing. First, holding the population fixed, the policy maker saves GH¢ 3.34 in subsidy budget per household-month. At a marginal cost of public funds of GH¢ 1.17 (Auriol and Warlters, 2012), this implies a public benefit of GH¢ 7.23 per household-month. Second, if instead the policy maker uses the savings to expand program coverage by 29%, the additional households covered under non-linear pricing provide additional greenhouse gas reductions, following the estimates in Section 2. We use the lower end of the social cost of carbon (USD 51 per ton) and assume full displacement of charcoal per kg of LPG purchased. This implies a GH¢ 11.49 external benefit from avoided climate damages per household-month.[59]

Combining the impact on consumer surplus, DWL and external benefits gives an overall welfare estimate under each scenario. Under scenario 1, this ranges from GH¢ 3.25 to 3.84 per household-month. Under scenario 2, it ranges from GH¢ 6.36 to 7.11 per household-month. Both calculations ignore the redistributional effect of targeting, which increases the subsidy allocation to relatively poor households (see Table 6). Finally, by restricting the pollution benefits to avoided climate damages, our estimates do not account for any health gains from improved air quality.

6.3 Conclusion

We contribute to a large literature on the targeting of subsidies for social goals. By leveraging sources of heterogeneity in the population, we show that non-linear prices can improve the cost effectiveness of subsidies, increasing demand for a clean fuel per dollar spent. Our two-stage experiment to design and test non-linear pricing against a counterfactual of linear prices offers a proof of concept for our approach. While further refinement is needed to make

[58]For switchers, we bound the DWL by zero (if their demand was unchanged) or by multiplying the difference in demand observed in the large cylinder in the linear price treatment and the demand observed in the small cylinder in the non-linear price treatment by the difference in price.

[59]Specifically, this assumes that each household under non-linear pricing provides 1.29 times the climate benefits as does a household under linear pricing. We also multiply the consumer surplus and DWL results by 1.29 in this scenario.

the approach scalable, our results demonstrate both the potential for second degree price discrimination to increase welfare and the value of experimental validation of theoretical predictions of optimal contracts and mechanisms.

Our application is to clean fuels in Ghana, but our approach can be extended to a range of consumer goods and services where demand is socially valuable, including health products and other environmental goods. Whether the population exhibits the heterogeneity necessary to induce separation and support non-linear pricing is an empirical question. While our theoretical model establishes sufficient conditions for separation using very specific sources of heterogeneity, preferences over small versus large purchase quantities are (anecdotally) common and leveraging these for targeting subsidies can be both progressive and efficient. We leave as future work generalizing our model to derive the necessary conditions for separation under a more general structure, as well as extending our empirical framework to accommodate non-experimental data.

References

Alatas, V., A. Banerjee, R. Hanna, B. A. Olken, and J. Tobias (2012). Targeting the poor: evidence from a field experiment in Indonesia. *American Economic Review 102*(4), 1206–1240.

Alatas, V., R. Purnamasari, M. Wai-Poi, A. Banerjee, B. A. Olken, and R. Hanna (2016). Self-targeting: Evidence from a field experiment in Indonesia. *Journal of Political Economy 124*(2), 371–427.

Alderman, H. and K. Lindert (1998). The potential and limitations of self-targeted food subsidies. *The World Bank Research Observer 13*, 213–229.

Alger, I. (1999). Consumer strategies limiting the monopolist's power: multiple and joint purchases. *The RAND Journal of Economics*, 736–757.

Armstrong, M. (2016). Nonlinear pricing. *Annual Review of Economics 8*, 583–614.

Attanasio, O. and E. Pastorino (2020). Nonlinear pricing in village economies. *Econometrica 88*, 207–263.

Auriol, E. and M. Warlters (2012). The marginal cost of public funds and tax reform in africa. *Journal of Development Economics 97*(1), 58–72.

Bailis, R., R. Drigo, A. Ghilardi, and O. Masera (2015). The carbon footprint of traditional woodfuels. *Nature Climate Change 5*, 266–272.

Bensch, G. and J. Peters (2015). The intensive margin of technology adoption: Experimental evidence on improved cooking stoves in rural Senegal. *Journal of Health Economics 42*, 44–63.

Berkouwer, S. B. and J. T. Dean (2022). Credit, attention, and externalities in the adoption of energy efficient technologies by low-income households. *American Economic Review 112*, 3291–3330.

Blanchard, O., C. Gollier, and J. Tirole (2023). The portfolio of economic policies needed to fight climate change. *Annual Review of Economics 15*(1), 689–722.

Borenstein, S. (2012). The redistributional impact of nonlinear electricity pricing. *American Economic Journal: Economic Policy 4*, 56–90.

Coady, D., M. Grosh, and J. Hoddinott (2004). Targeting outcomes redux. *World Bank Research Observer 19*, 61–85.

Cohen, J., P. Dupas, and S. Schaner (2015). Price subsidies, diagnostic tests, and targeting of malaria treatment: Evidence from a randomized controlled trial. *American Economic Review 105*, 609–645.

Dillon, B., J. De Weerdt, and T. O'Donoghue (2021). Paying more for less: Why don't households in Tanzania take advantage of bulk discounts? *The World Bank Economic Review 35*(1), 148–179.

Dizon-Ross, R. and A. Zucker (2023). Mechanism design for personalized policy: A field experiment incentivizing exercise. *Working Paper*.

Dubé, J.-P. and S. Misra (2023). Personalized pricing and consumer welfare. *Journal of Political Economy 131*(1), 131–189.

Duflo, E., R. Hanna, and S. P. Ryan (2012). Incentives work: Getting teachers to come to school. *American Economic Review 102*, 1241–1278.

Dupas, P., V. Hoffmann, M. Kremer, and A. P. Zwane (2016). Targeting health subsidies through a nonprice mechanism: A randomized controlled trial in Kenya. *Science 353*, 889–895.

Floess, E., A. Grieshop, E. Puzzolo, D. Pope, N. Leach, C. J. Smith, A. Gill-Wiehl, K. Landesman, and R. Bailis (2023). Scaling up gas and electric cooking in low-and middle-income countries: climate threat or mitigation strategy with co-benefits? *Environmental Research Letters 18*(3), 034010.

Fowlie, M. and R. Meeks (2021). The economics of energy efficiency in developing countries. *Review of Environmental Economics and Policy 15*(2), 238–260.

Geddie, J. and J. Brock (2022). Explainer: Plastic sachets: As big brands cashed in, a waste crisis spiraled. *Reuters*.

Ghana Statistical Service (2022). Ghana 2021 population and housing census - volume 3k: Housing characteristics. Technical report, Ghana Statistical Service.

Ghana Statistical Service (2022-2023). Price indices: Consumer price index bulletin. Accessed on January 23, 2023.

Gould, C. F., R. Bailis, K. Balakrishnan, M. Burke, S. Espinoza, S. Mehta, S. B. Schlesinger, J. R. Suarez-Lopez, and A. Pillarisetti (2024). In praise of cooking gas subsidies: transitional fuels to advance health and equity. *Environmental Research Letters*.

Hanna, R., E. Duflo, and M. Greenstone (2016). Up in smoke: The influence of household behavior on the long-run impact of improved cooking stoves. *American Economic Journal: Economic Policy 8*, 80–114.

Hanna, R. and B. A. Olken (2018). Universal basic incomes versus targeted transfers: Anti-poverty programs in developing countries. *Journal of Economic Perspectives 32*, 201–226.

Haushofer, J., P. Niehaus, C. Paramo, E. Miguel, and M. W. Walker (2022). Targeting impact versus deprivation.

Hendel, I., A. Lizzeri, and N. Roketskiy (2014). Nonlinear pricing of storable goods. *American Economic Journal: Microeconomics 6*(3), 1–34.

Hendel, I. and A. Nevo (2013). Intertemporal price discrimination in storable goods markets. *American Economic Review 103*(7), 2722–2751.

Higbee, S. D. (2024). Experimental design for policy choice. *Working paper*.

Ida, T., T. Ishihara, K. Ito, D. Kido, T. Kitagawa, S. Sakaguchi, and S. Sasaki (2022). Choosing who chooses: Selection-driven targeting in energy rebate programs. *Working Paper*.

Imelda (2020). Cooking that kills: Cleaner energy access, indoor air pollution, and health. *Journal of Development Economics 147*.

Ito, K., T. Ida, and M. Tanaka (2023). Selection on welfare gains: Experimental evidence from electricity plan choice. *American Economic Review 13*(4), 2937–2973.

Jack, B. K. (2013). Private information and the allocation of land use subsidies in Malawi. *American Economic Journal: Applied Economics 5*, 113–135.

Jack, B. K., K. McDermott, and A. Sautmann (2022). Multiple price lists for willingness to pay elicitation. *Journal of Development Economics 159*, 102977.

Johnson, T. R. and M. Lipscomb (2022). Pricing people into the market: Targeting through price discrimination. *Working Paper*.

Kang, Z. Y. (2022). Optimal indirect regulation of externalities. *Working Paper*.

Kansikas, C., A. Mani, and P. Niehaus (2023). Customized cash transfers: financial lives and cash-flow preferences in rural kenya. Technical report, National Bureau of Economic Research.

Kojima, M. (2021). Subsidizing bottled gas: Approaches and effects on household use. *World Bank ESMAP*.

Langer, A. and D. Lemoine (2022). Designing dynamic subsidies to spur adoption of new technologies. *Journal of the Association of Environmental and Resource Economists 9*(6), 1197–1234.

Levitt, S. D., J. A. List, S. Neckermann, and D. Nelson (2016). Quantity discounts on a virtual good: The results of a massive pricing experiment at King digital entertainment. *Proceedings of the National Academy of Sciences of the United States of America 113*, 7323–7328.

Maskin, E. and J. Riley (1984). Monopoly with incomplete information. *Source: The RAND Journal of Economics 15*, 171–196.

Mussa, M. and S. Rosen (1978). Monopoly and product quality. *Journal of Economic theory 18*(2), 301–317.

National Petroleum Agency (2022-2023). Indicative prices. Accessed on January 23, 2023.

Nauges, C. and D. Whittington (2017). Evaluating the performance of alternative municipal water tariff designs: Quantifying the tradeoffs between equity, economic efficiency, and cost recovery. *World Development 91*, 125–143.

Nichols, A. L. and R. J. Zeckhauser (1982). Targeting transfers through restrictions on recipients. *The American Economic Review 72*(2), 372–377.

Nikzad, A. (2023). Constrained majorization: Applications in mechanism design. *Working Paper*.

Parks, R. W. (1974). The demand and supply of durable goods and durability. *The American Economic Review 64*(1), 37–55.

Pires, T. and A. Salvo (2015). Cash-constrained households and product size. *Working Paper*.

Rafkin, C., A. Solomon, and E. Soltas (2023). Self-targeting in US transfer programs. *Working Paper*.

Rochet, J.-C. and L. A. Stole (2003). The economics of multidimensional screening. *Econometric Society Monographs 35*, 150–197.

Rothschild, M. and J. Stiglitz (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics 90*, 629–649.

Sallee, J. M. (2025). The trouble with green subsidies. *National Tax Journal 78*(1), 000–000.

Smith, K. R., J. P. McCracken, M. W. Weber, A. Hubbard, A. Jenny, L. M. Thompson, J. Balmes, A. Diaz, B. Arana, and N. Bruce (2011). Effect of reduction in household air pollution on childhood pneumonia in Guatemala (RESPIRE): A randomised controlled trial. *The Lancet 378*, 1717–1726.

Solon, G., S. J. Haider, and J. M. Wooldridge (2015). What are we weighting for? *Journal of Human resources 50*(2), 301–316.

Stoner, O., J. Lewis, I. L. Martínez, S. Gumy, T. Economou, and H. Adair-Rohani (2021). Household cooking fuel estimates at global and country level for 1990 to 2030. *Nature communications 12*(1), 5793.

Szabó, A. (2015). The value of free water: Analyzing South Africa's free basic water policy.

Econometrica 83, 1913–1961.

Thoday, K., P. Benjamin, M. Gan, and E. Puzzolo (2018). The mega conversion program from kerosene to lpg in indonesia: Lessons learned and recommendations for future clean cooking energy expansion. *Energy for Sustainable Development 46*, 71–81.

Verma, A. P. and Imelda (2023). Clean energy access: gender disparity, health and labour supply. *The Economic Journal 133*(650), 845–871.

World Bank (2022). *Poverty and shared prosperity 2022: Correcting course.* The World Bank.

World Health Organization (2023). Global health observatory, household air pollution data. Technical report, World Health Organization.

World LPG Association (2023). Defining and measuring LPG subsidies. Technical report, WLPGA.

Appendix to:

Targeting subsidies through price menus

A Figures and tables

Figure A.1: Balance check for MPL draws

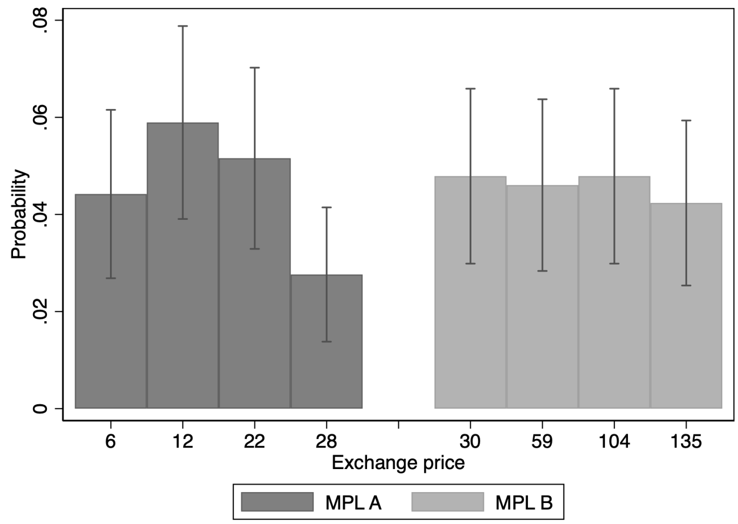

Note: Probability of draws in the "static" choices in the MPL. These choices were given to all participants. MPL A offered a choice between the small cylinder and nothing; MPL B offered a choice between the large cylinder and nothing.

Figure A.2: Distribution of prices by cylinder size

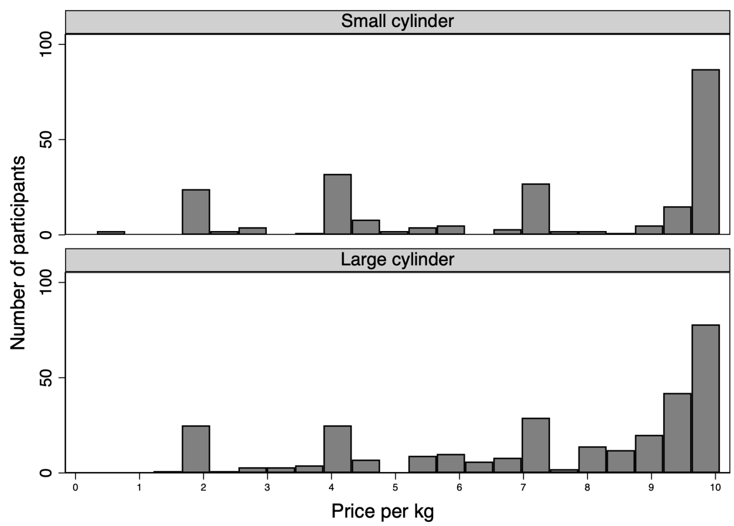

Note: Distribution of prices per kg of LPG by cylinder size, resulting from the MPL. Assignment of size and price were random conditional on participant choices.

Figure A.3: Predictions of $\hat{\hat{\theta}} - (p_L - p_S)$ using observables

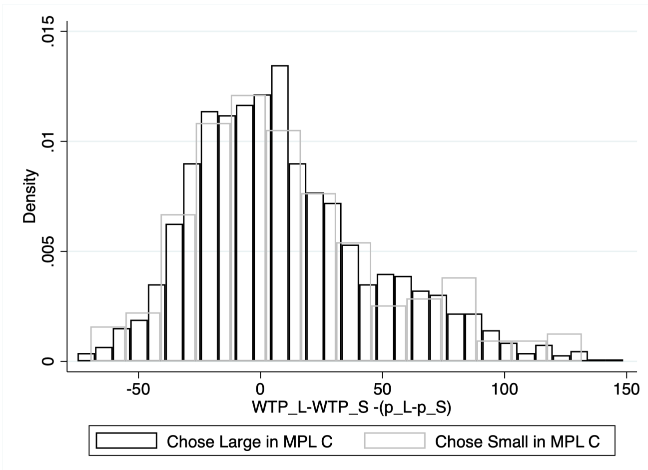

(a) OLS

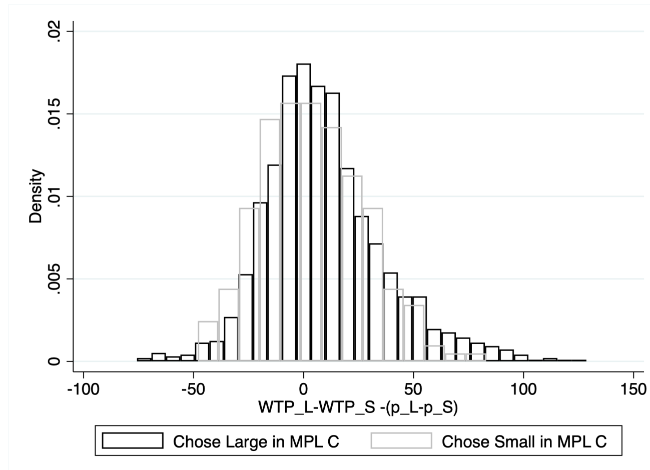

(b) Random forest

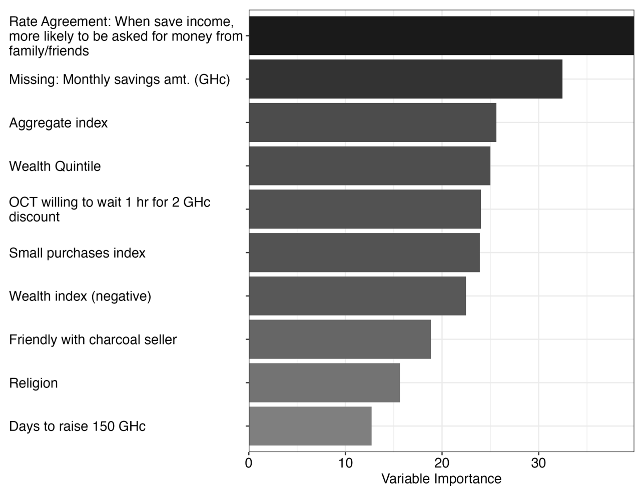

(c) Top 10 important variables

Note: Panels (a) and (b) show results after predicting $\hat{\hat{\theta}}$ using OLS regression (a) or random forest model (b). Horizontal axis is the surplus from choosing the large cylinder size, or $\hat{\hat{\theta}} - (p_L - p_S)$. Black bars are participants who chose the large cylinder in MPL C at the relevant price difference; gray bars are participants who chose the small cylinder. OLS regression is performed on the top 10 important variables identified from the random forest. Panel (c) shows these variables and their ranking according to their permutation importance.

Figure A.4: LPG exchanged per month by price per kg (Stage 1)

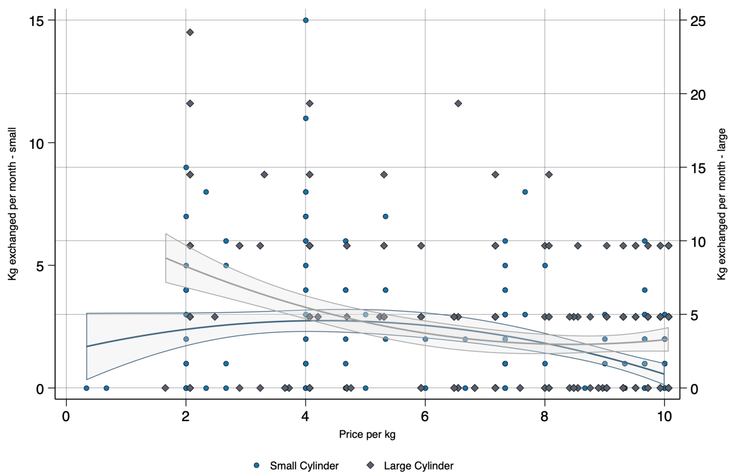

Note: Each scatterplot point is a participant with a randomly assigned cylinder size and price (conditional on MPL choices). Lines are fitted quadratics with confidence intervals. For small cylinder, main coefficient: 0.58 with p-value: 0.06; quadratic coefficient: -0.07 with p-value: 0.01. For large cylinder, main coefficient: -2.16 with p-value: 0.00; quadratic coefficient: 0.13 with p-value: 0.01.

Figure A.5: Robustness to variation around the selected prices, objective function and budget

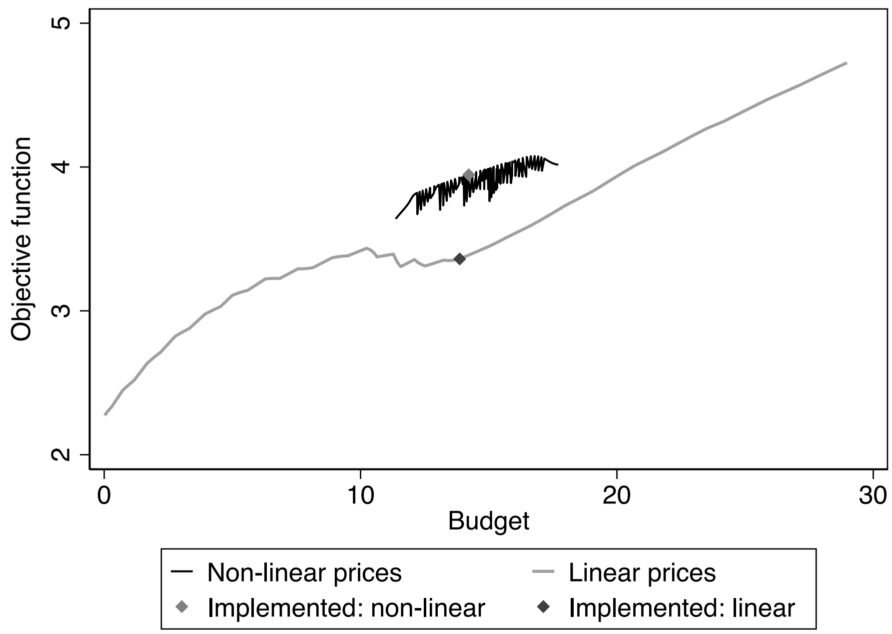

(a) Demand and budget by price combinations

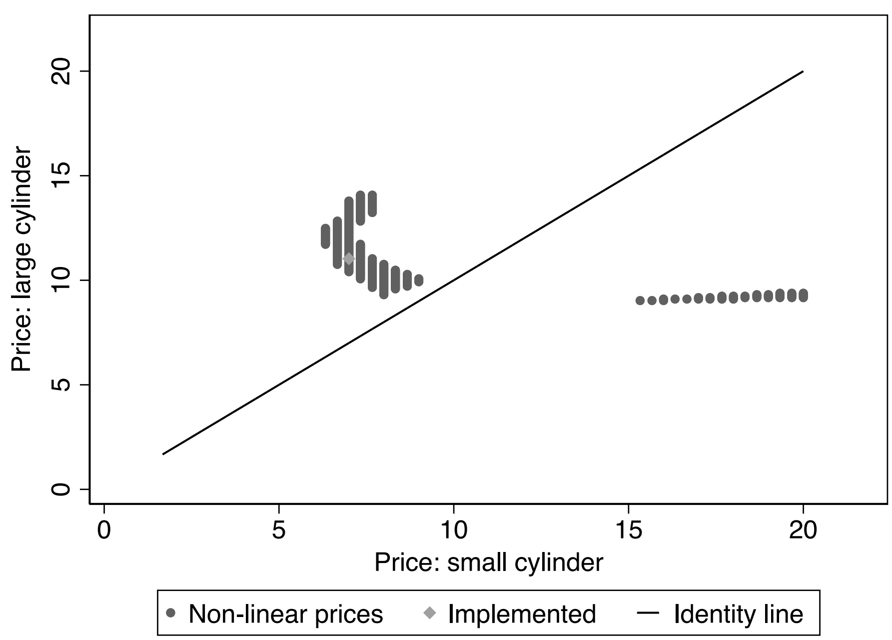

(b) Alternative demand and budget

Note: Panel A displays the objective function (demand) and social planner's budget for non-linear prices that are within 10% of the implemented prices. Panel B displays the per kg prices for menus with predictions within 10% of the objective function and budget of the implemented non-linear menu.

Figure A.6: Cost per ton of avoided CO_2e

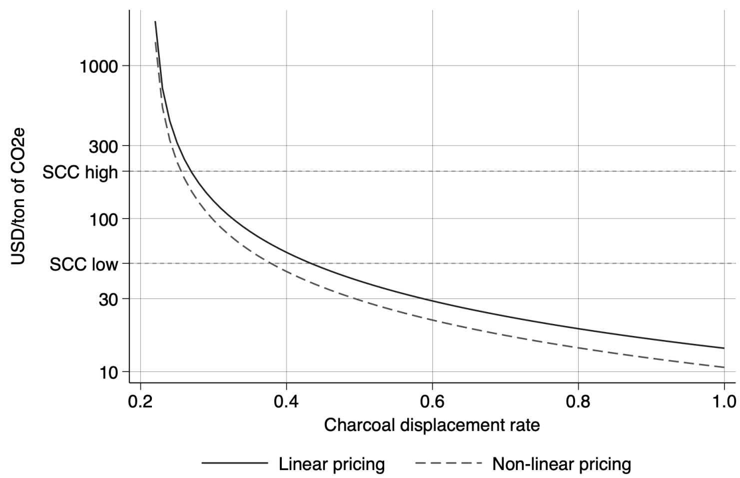

Note: Estimated cost per ton of avoided CO_2e, based on data from Floess et al. (2023). The emissions from LPG and charcoal include CO_2, CH_4, N_2O and black carbon from both production and consumption, and adjust for a Ghana-specific fraction of non-renewable biomass used for charcoal production. Assumed charcoal displacement rates use an energy ratio of 1.6 for LPG to charcoal.

Table A.1: Probability of starter kit delivery failure

	(1)	(2)	(3)
Panel A: Stage 1			
Small cylinder	0.107***		0.131
	(0.035)		(0.094)
Price per kg		-0.003	-0.000
		(0.006)	(0.007)
Small cylinder × Price per kg			-0.003
			(0.012)
Control mean	0.14	0.19	0.14
Observations	525	525	525
Panel B: Stage 2			
Small cylinder	0.036		0.057
	(0.023)		(0.036)
Treatment		0.009	0.017
		(0.018)	(0.019)
Small cylinder × Treatment			-0.040
			(0.045)
Control weighted mean	0.045	0.051	0.045
Observations	908	908	908

Note: Each column is a separate linear regression of indicator for delivery failure on row variables and strata fixed effects. Analysis is conditional on drawing a cylinder (i.e., omits households that chose no cylinder). Stage 2 results are weighted by stage 1 strata. Standard errors are in parentheses. * $p<0.1$, ** $p<0.05$, *** $p<0.01$

Table A.2: Summary statistics by stage

	Stage 1	Stage 2	Difference
Panel A: HDSS			
Female household head	0.530	0.537	0.027
	(0.500)	(0.499)	(0.027)
Age of household head	48.549	49.198	-0.934
	(17.727)	(16.843)	(0.930)
Household size	4.567	5.046	-0.003
	(2.776)	(3.143)	(0.147)
Wealth quintile	2.860	3.483	0.618***
	(1.393)	(1.400)	(0.076)
Panel B: Baseline			
Respondent is household head	0.580	0.502	-0.080***
	(0.494)	(0.500)	(0.027)
Age of respondent	43.055	43.844	-0.183
	(16.235)	(15.585)	(0.881)
Time cost index		0.011	
		(1.105)	
– Willing to wait for discount	-1.739	-1.698	0.059
	(1.145)	(1.113)	(0.063)
– Lack transport access		-0.417	
		(0.755)	
Savings cost index	0.049	-0.029	-0.058
	(1.035)	(0.978)	(0.054)
– Time to raise money	17.280	16.699	-0.405*
	(4.243)	(4.611)	(0.245)
– Time to repay loan	1.061	1.164	0.038
	(9.478)	(7.497)	(0.456)
Small purchases index	0.284	-0.169	-0.183***
	(1.001)	(0.961)	(0.039)
– Buys small charcoal bag	0.507	0.386	-0.099***
	(0.500)	(0.487)	(0.028)
– Wealth index (negative)	-2.860	-3.483	-0.618***
	(1.393)	(1.400)	(0.076)

Note: Columns 1 and 2 show means and standard deviations for the sample in each stage of the study. Column 3 shows the coefficient for the correlation of each variable with an indicator for stage 2 controlling for stage 2 strata fixed effects – except for the estimate on wealth which includes depot fixed effects. * $p<0.1$, ** $p<0.05$, *** $p<0.01$

Table A.3: Stage 2 balance

	Control	Treatment	Difference
Panel A: HDSS			
Female household head	0.528	0.545	0.014
	(0.500)	(0.498)	(0.033)
Age of household head	48.703	49.702	0.861
	(16.569)	(17.122)	(1.089)
Household size	4.946	5.148	0.186
	(2.969)	(3.310)	(0.187)
Wealth quintile	3.472	3.494	0.024
	(1.401)	(1.401)	(0.093)
Panel B: Baseline			
Respondent is household head	0.513	0.490	-0.025
	(0.500)	(0.500)	(0.033)
Age of respondent	43.866	43.821	-0.134
	(15.968)	(15.201)	(1.026)
Time cost index	-0.019	0.040	0.063
	(1.139)	(1.069)	(0.073)
– Willing to wait for discount	-1.714	-1.682	0.028
	(1.120)	(1.107)	(0.073)
– Lack transport access	-0.435	-0.400	0.042
	(0.781)	(0.727)	(0.050)
Savings cost index	-0.026	-0.032	-0.003
	(1.006)	(0.950)	(0.065)
– Time to raise money	16.694	16.703	-0.001
	(4.689)	(4.535)	(0.307)
– Time to repay loan	1.208	1.119	-0.040
	(7.781)	(7.204)	(0.495)
Small purchases index	-0.172	-0.165	0.019
	(0.955)	(0.967)	(0.045)
– Buys small charcoal bag	0.379	0.393	0.017
	(0.486)	(0.489)	(0.033)
– Wealth index (negative)	-3.472	-3.494	-0.024
	(1.401)	(1.401)	(0.093)

Note: Columns 1 and 2 show means and standard deviations for the sample assigned to the stage 2 control and treatment groups, respectively. Column 3 shows the coefficient for a regression of each variable on an indicator for treatment assignment. All estimates include stage 2 strata fixed effects except for the estimate on wealth which includes depot fixed effects. * $p<0.1$, ** $p<0.05$, *** $p<0.01$

Table A.4: Approximated $\hat{\theta}$ and stage 1 observables

	(1)	(2)	(3)
Panel A: HDSS			
Female household head	-3.847	-0.896	-1.662
	(3.198)	(3.180)	(3.112)
Age of household head	-0.163	-0.289***	-0.317***
	(0.101)	(0.107)	(0.108)
Household size	0.432	-0.096	0.038
	(0.545)	(0.594)	(0.605)
Wealth quintile	4.939***	5.945***	5.982***
	(1.123)	(1.288)	(1.286)
Depot FE			X
R^2		0.066	0.087
Panel B: Baseline			
Respondent is household head	-0.279	3.491	3.593
	(3.261)	(3.286)	(3.322)
Age of respondent	-0.170	-0.280**	-0.299***
	(0.110)	(0.110)	(0.111)
Willing to wait for discount	0.785	1.379	1.004
	(1.571)	(1.503)	(1.490)
Time to raise money	-1.026***	-0.654*	-0.560
	(0.353)	(0.347)	(0.346)
Time to repay loan	0.032	-0.033	-0.041
	(0.058)	(0.056)	(0.058)
Buys small charcoal bag	-9.047***	-7.766**	-7.226**
	(3.223)	(3.135)	(3.137)
Wealth index (negative)	-4.939***	-5.173***	-5.354***
	(1.123)	(1.129)	(1.118)
Depot FE			X
R^2		0.089	0.105

Note: Column 1 shows estimates from separate univariate linear regressions of the approximated difference in maximum willingness to pay between cylinder sizes ($p(\theta_i, q_L) - p(\theta_i, q_S)$) on each of the row variables. A higher difference results in a higher probability of choosing the large cylinder, all else equal. Column 2 shows estimates from multiple linear regression of maximum willingness to pay on all row variables. Column 3 shows estimates from regression in column 2 with depot fixed effects. Sample consists of all stage 1 households who paid a deposit for whom approximated $\hat{\theta}$ is observed. * p<0.1, ** p<0.05, *** p<0.01

Table A.5: Stage 1 price sensitivity of demand

	Any exchange		Ln(kg/month)	
	(1)	(2)	(3)	(4)
Price per kg	-0.018**	-0.033*	-0.096***	-0.090***
	(0.008)	(0.018)	(0.011)	(0.019)
Small cylinder	-0.007	-0.002	-1.068***	-0.778***
	(0.044)	(0.154)	(0.066)	(0.262)
Small cylinder × Price per kg		-0.005		-0.051
		(0.027)		(0.044)
Dynamic MPL		-0.294*		-0.212
		(0.171)		(0.196)
Dynamic MPL × Price per kg		0.039		0.033
		(0.024)		(0.025)
Dynamic MPL × Small cylinder		0.188		0.412
		(0.267)		(0.405)
Dynamic MPL × Small cylinder × Price per kg		-0.016		-0.036
		(0.037)		(0.055)
Observations	525	525	292	292

Note: Columns 1 and 2 present coefficients from a linear regression of indicator of any LPG cylinder exchange during the program on price per kg. Columns 3 and 4 present coefficients from a log-linear regression of LPG demand per month on price per kg. Each regressor is interacted with an indicator for the MPL assigned price being selected from a dynamic MPL row (where participant choices affect randomization outcomes). All regressions include stage 1 strata fixed effects. Samples in columns 1-2 include all households with small or large cylinder, and columns 3-4 drop households with no exchanges. * p<0.1, ** p<0.05, *** p<0.01

Table A.6: Stage 1 exchange price variance by $\hat{\hat{\theta}}$ group

	Small (mean/sd)	Large (mean/sd)
Low	21.236	79.459
	(7.710)	(29.984)
Medium	20.236	120.047
	(8.919)	(31.187)
High	22.255	121.459
	(9.700)	(39.947)
Observations	165	255

Note: Columns 1-2 show the mean and (standard deviation) for the exchange price drawn in stage 1 for the small and large cylinder, respectively. Rows indicate groups of $\hat{\hat{\theta}}$ in equal size bins, binned separately for large and small cylinders. Households falling between the medium and high bins for the large cylinder were randomly assigned to bins 100 times and the median of the mean/sd across randomizations is displayed.

Table A.7: Stage 2 subsidy budget and LPG demand (unweighted)

	Subsidy budget (GH¢/month)			LPG demand (kg/month)		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment effect	-3.333***	-3.391***	-3.209***	-0.153	-0.166	-0.105
	(0.837)	(0.824)	(0.824)	(0.259)	(0.254)	(0.253)
Control mean	11.585	11.585	11.585	3.089	3.089	3.089
Strata FE		×	×		×	×
Surveyor FE			×			×
Enrollment week FE			×			×
Observations	915	915	915	915	915	915

Note: This table replicates Table 3 with the exception that the regressions are not weighted by stage 1 strata. Columns 1-3 show treatment effect on subsidy per month; columns 4-6 on LPG demand per month. Columns 2 and 5 include strata fixed effects; columns 3 and 6 include strata, surveyor and enrollment week fixed effects. All regressions are not weighted by stage 1 strata. * p<0.1, ** p<0.05, *** p<0.01

Table A.8: Robustness of menu design and treatment effects

Panel A: Menu design, stage 1				
	Linear price		Non-linear price	
	Small	Large	Small	Large
Cylinder takers only	9.333	9.241	7.000	11.034
Intent to treat	9.667	9.724	7.667	11.931

Panel B: Treatment effects, stage 2						
	Subsidy budget (GH¢/month)			LPG demand (kg/month)		
	(1)	(2)	(3)	(4)	(5)	(6)
Cylinder takers only	-4.204***	-4.219***	-3.805***	-0.219	-0.225	-0.084
	(1.041)	(1.041)	(1.031)	(0.334)	(0.331)	(0.328)
Intent to treat	-4.008***	-4.080***	-3.712***	-0.223	-0.243	-0.119
	(1.003)	(0.998)	(0.985)	(0.320)	(0.317)	(0.314)
Control weighted mean	12.68	12.68	12.68	3.39	3.39	3.39
Observations	915	915	915	915	915	915
Strata FE		×	×		×	×
Surveyor FE			×			×
Enrollment week FE			×			×

Note: Robustness of menu design (stage 1, Panel A) and treatment effects (stage 2, Panel B). Rows labeled *Cylinder takers only* include only households with a successful cylinder delivery. Rows labeled *Intent to treat* include all households and set both the budget and demand to zero for households without a successful delivery. * p<0.1, ** p<0.05, *** p<0.01

Table A.9: Stage 2 demand by cylinder size and size selection (unweighted)

	LPG demand (kg/month)				Small cylinder	
	(1)	(2)	(3)	(4)	(5)	(6)
Sample:	Small	Small	Large	Large	All	All
Treatment effect	0.532**	0.548**	-0.310	-0.171	0.054*	0.055*
	(0.218)	(0.238)	(0.337)	(0.339)	(0.030)	(0.030)
Control mean	1.105	1.105	3.854	3.854	0.266	0.266
Strata FE	×	×	×	×	×	×
Surveyor FE		×		×		×
Survey week FE		×		×		×
Observations	267	267	641	641	915	915

Note: This table replicates Table 4 with the exception that the regressions are not weighted by stage 1 strata. Columns 1-4 show results from linear regression of LPG demand per month on indicator for treatment assignment. Columns 1-2 show results for sample of households with small cylinder, columns 3-4 for sample with large cylinder. Households that chose no cylinder (N=7) are omitted from columns 1-4. Columns 5-6 show results from linear regression of indicator for small cylinder take up on indicator for treatment assignment on full sample. All columns include stage 2 strata fixed effects, even columns also include surveyor and survey week fixed effects. All regressions are not weighted by stage 1 strata. * p<0.1, ** p<0.05, *** p<0.01

Table A.10: Stage 2 subsidy robustness checks

	Subsidy budget (GH¢/month)			
	(1)	(2)	(3)	(4)
Treatment effect	-3.712***	-3.987***	-3.551***	-3.954***
	(0.985)	(0.911)	(0.473)	(1.029)
Control weighted mean	12.681	11.749	1.638	13.344
Strata FE	×	×	×	×
Surveyor FE	×	×	×	×
Enrollment week FE	×	×	×	×
Observations	915	915	915	915

Note: Column 1 shows treatment effect on subsidy per month, column 2 on subsidy per month using pump prices, column 3 on subsidy per month using refinery prices, and column 4 using a CPI adjustment. LPG pump and refinery prices from National Petroleum Agency (2023). CPI data from Ghana Statistical Service (2023). All regressions are weighted by stage 1 strata. * p<0.1, ** p<0.05, *** p<0.01

Table A.11: Stage 2 targeting on observables

Index input:	(1) Low time cost	(2) Lacks transport	(3) Low liquidity	(4) Time to repay	(5) Small purchases	(6) Low wealth
Panel A: Outcome: Small cylinder						
Treatment effect	0.062*	0.062*	0.061*	0.062*	0.065*	0.056*
	(0.033)	(0.033)	(0.033)	(0.033)	(0.034)	(0.032)
Index input	-0.027	0.001	0.080***	0.059***	0.052**	0.032
	(0.023)	(0.021)	(0.018)	(0.020)	(0.024)	(0.022)
Treatment × Index input	0.071**	-0.028	-0.004	-0.058*	-0.012	-0.009
	(0.032)	(0.034)	(0.029)	(0.030)	(0.035)	(0.033)
Control weighted mean	0.270	0.270	0.270	0.270	0.270	0.270
Observations	915	915	909	915	862	915
Panel B: Outcome: LPG demand (kg/month)						
Treatment effect	-0.226	-0.226	-0.218	-0.217	-0.294	-0.167
	(0.321)	(0.321)	(0.317)	(0.321)	(0.327)	(0.249)
Index input	0.088	0.125	-0.457**	-0.154	-0.253	-0.772***
	(0.206)	(0.201)	(0.220)	(0.121)	(0.230)	(0.203)
Treatment × Index input	0.073	-0.347	0.086	-0.088	-0.119	0.108
	(0.306)	(0.313)	(0.299)	(0.135)	(0.324)	(0.286)
Control weighted mean	3.090	3.090	3.090	3.090	3.090	3.090
Observations	915	915	909	915	862	915
Panel C: Outcome: Subsidy budget (GH¢/month)						
Treatment effect	-4.020***	-4.009***	-3.971***	-3.997***	-4.105***	-3.258***
	(1.005)	(1.005)	(0.997)	(1.005)	(1.040)	(0.828)
Index input	0.386	0.355	-1.585*	-0.496	-1.119	-2.788***
	(0.765)	(0.766)	(0.841)	(0.472)	(0.866)	(0.762)
Treatment × Index input	0.238	-1.241	1.119	-0.049	0.693	1.354
	(0.933)	(0.968)	(0.988)	(0.513)	(1.040)	(0.918)
Control weighted mean	11.59	11.59	11.59	11.59	11.59	11.59
Observations	915	915	909	915	862	915

Panel A shows the treatment effect on small cylinder take up for all households, panel B on monthly LPG demand, and panel C on subsidy per month. Each column controls for index input variable specified in column header and its interaction with treatment indicator. The regressions are weighted by stage 1 strata. * p<0.1, ** p<0.05, *** p<0.01

Table A.12: Stage 2 targeting on observables (unweighted)

Index:	(1) Low time cost	(2) High savings cost	(3) Small purchases/wealth	(4) Aggregate
Panel A: Outcome: Small cylinder				
Treatment effect	0.050*	0.051*	0.044	0.049
	(0.030)	(0.030)	(0.031)	(0.030)
Index	-0.020	0.106***	0.070***	0.082***
	(0.021)	(0.020)	(0.022)	(0.020)
Treatment × Index	0.067**	-0.043	-0.041	-0.013
	(0.030)	(0.039)	(0.031)	(0.030)
Control mean	0.270	0.270	0.270	0.270
Observations	915	915	915	915
Panel B: Outcome: LPG demand (kg/month)				
Treatment effect	-0.157	-0.156	-0.118	-0.132
	(0.258)	(0.256)	(0.246)	(0.252)
Index	0.126	-0.403**	-0.715***	-0.528***
	(0.174)	(0.163)	(0.195)	(0.181)
Treatment × Index	0.027	-0.011	0.180	0.133
	(0.247)	(0.215)	(0.268)	(0.243)
Control mean	3.090	3.090	3.090	3.090
Observations	915	915	915	915
Panel C: Outcome: Subsidy budget (GH¢/month)				
Treatment effect	-3.353***	-3.319***	-3.012***	-3.157***
	(0.836)	(0.831)	(0.806)	(0.822)
Index	0.497	-1.401**	-2.745***	-1.942***
	(0.646)	(0.606)	(0.720)	(0.663)
Treatment × Index	0.225	0.691	1.820**	1.506*
	(0.784)	(0.714)	(0.866)	(0.795)
Control mean	11.59	11.59	11.59	11.59
Observations	915	915	915	915

Note: This table replicates Table 6 with the exception that the regressions are not weighted by stage 1 strata. Columns 1-4 show treatment effects for all households. Each column controls for index specified in column header and its interaction with treatment predictor. All regressions are not weighted by stage 1 strata. * $p<0.1$, ** $p<0.05$, *** $p<0.01$

Table A.13: Stage 1 effect of deposit discount and depot assignment

	WTP difference		Paid deposit		Demand (kg/month)	
	(1)	(2)	(3)	(4)	(5)	(6)
Deposit discount	-1.594		0.099***		0.188	
	(3.417)		(0.036)		(0.339)	
Far depot	-0.234		0.030		-0.186	
	(3.895)		(0.042)		(0.377)	
Wealth Q2		2.847		0.053		0.485
		(5.790)		(0.058)		(0.472)
Wealth Q3		16.901***		0.113**		0.597
		(5.047)		(0.055)		(0.470)
Wealth Q4		11.671**		0.105*		0.714
		(5.319)		(0.058)		(0.541)
Wealth Q5		20.545***		0.170***		1.687***
		(4.908)		(0.055)		(0.580)
Observations	420	420	543	543	543	543

Note: Effect of stage 1 variation on main outcomes. Odd columns show show treatment effects of a randomized deposit discount of GH¢ 50 and randomized assignment to the next-closest depot. Even columns show effects of wealth strata. Odd columns include strata (wealth quintile × depot) fixed effects and even columns include depot fixed effects. * p<0.1, ** p<0.05, *** p<0.01

Table A.14: Stage 2 take up and demand by baseline ownership

	Enrolled		Any exchange		Demand (kg/month)
	(1)	(2)	(3)	(4)	(5)
Treatment effect	-0.002	0.006	0.002	0.017	-0.029
	(0.018)	(0.023)	(0.037)	(0.052)	(0.445)
Owns LPG (at baseline)		0.009		0.073	0.730*
		(0.024)		(0.052)	(0.434)
Treatment × Owns LPG		-0.016		-0.032	-0.187
		(0.034)		(0.074)	(0.626)
Control weighted mean	0.943	0.943	0.551	0.551	3.385
Observations	915	915	915	915	915

Note: Heterogeneous treatment effects by baseline LPG ownership. "Owns LPG (at baseline)" equals one if the household owned any LPG equipment including a stove or cylinder at the time of the baseline survey. All columns include fixed effects for stage 2 strata, surveyor and baseline week. The regressions are weighted by stage 1 strata. * p<0.1, ** p<0.05, *** p<0.01

Table A.15: Welfare effect of non-linear pricing

		Small stayers	Large stayers	Switchers	Average, fixed pop.	Average, fixed budget
ΔCS	Lower	3.112	-6.633	0.000	-3.652	-4.703
	Upper	3.199	-7.101	8.638	-3.441	-4.431
ΔDWL	Lower	-0.621	0.190	-4.968	-0.326	-0.420
	Upper	-0.534	0.278	0.000	0.045	0.059
ΔExt benefit					7.232	11.485
ΔWelfare	Lower				3.254	6.362
	Upper				3.836	7.112

Note: Back of the envelope calculation of the welfare effects of non-linear pricing relative to the linear pricing benchmark, using unweighted results from stage 2. "Stayers" refers to the share of the population that always chose the small or large cylinder, 27 and 68% of the population, respectively. "Switchers" refers to the 5.8% of the population that was induced to switch from the large to small cylinder by treatment. The average, fixed population is the welfare calculation assuming the same population across treatments, which results in a lower budget for the policy maker, valued at a marginal cost of public funds of GH¢ 1.17. Average, fixed budget is a counterfactual in which the subsidy covers 29% more household under non-linear pricing. The lower and upper bounds reflect different assumptions about the price elasticity of switchers and the non-monetary (transaction and savings) costs they incur when purchasing gas in small and large cylinder. See text for additional detail.

B Theory

Consumer's Problem

The consumer's maximization problem in a given period is given by:

$$\max_{c,q,x} u(c, qx; \theta) \tag{1}$$

$$\text{s.t. } c + p(q)x + D(\theta)x + 1(x > 0)\frac{p(q)}{x}A(\theta) = y(\theta)$$

where $q \in \{q_S, q_L\}$ denotes the size of the LPG cylinder in kg, c is consumption measured in cedis, x is the number of exchange trips in a given period, $p(q)$ is the price of an exchange as a function of its size, $D(\theta)$ is the transaction cost per exchange, $1(x > 0)\frac{p(q)}{x}A(\theta)$ is the savings costs, and $y(\theta)$ denotes income. Assume that type θ determines income $y(\theta)$, savings costs $A(\theta)$, and transaction costs $D(\theta)$. Type may also determine preferences for LPG, such that $\frac{\partial u}{\partial q \partial \theta} > 0$.

Savings costs are the product of four components:

$$\underbrace{1(x > 0)}_{\substack{\text{if positive} \\ \text{exchanges}}} \quad \underbrace{p(q)}_{\substack{\text{amount to} \\ \text{hold on to}}} \quad \underbrace{\frac{1}{x}}_{\substack{\text{\# of periods} \\ \text{until exchange}}} \quad \underbrace{A(\theta)}_{\substack{\text{cost per period} \\ \text{per cedi}}}$$

The logic behind this parameterization is the following: when the frequency of purchase is less than one (e.g., $x = 0.5$), the agent has to hold on to a quantity of money $p(q)$ for multiple periods before buying the next cylinder. $A(\theta)$ represents the cost of holding on to one cedi for a full period of time without spending it. Thus, if the frequency of purchase is $x = 0.5$ per period, the number of periods she has to save in order to purchase the cylinder is $\frac{1}{x} = \frac{1}{0.5} = 2$. If the cost of saving one cedi per period is about 30 cents ($A(\theta) = 0.3$), then the additional cost that is incurred is $p(q) \times \frac{1}{x} \times A(\theta) = p(q) \times 2 \times 0.30$.

Proofs of Propositions 1 and 2

A type-θ agent chooses a cylinder size $q \in \{q_L, q_S\}$, exchange frequency x, and outside consumption level c that maximizes her utility over consumption of LPG and outside commodity subject to her budget constraint and non-negativity constraints on her refill frequency:

$$\max_{q,x,c} u\Big(c\,,\ xq \mid \theta\Big) \tag{2}$$

$$\text{s.t.} \qquad \begin{cases} y(\theta) = c + p(q)x + D(\theta)x + p(q)A(\theta)x^{-1} \\ x \geq 0 \end{cases} \tag{3}$$

Let

$$\varphi(\theta) = \max_{q,x,c}\ u\big(c, xq|\theta\big) \ \text{ s.t. (3)}$$

be a type-θ agent's value function associated with the problem in (2)-(3) and let $q^*(\theta), c^*(\theta), x^*(\theta)$ be her optimal policy functions for cylinder size, outside consumption, and refill frequency. Write her indirect utility function conditional on type θ and cylinder size choice q as

$$v\big(\theta|q\big) = \max_{x} u\Big(c\,,\ xq \mid q, \theta, p(q)\Big) \ \text{ s.t. (3)} \equiv \max_{x} f(x|q,\theta)$$

where the last equivalence stems from substituting in the budget constraint into (2). Her value function can then be written as her choice of cylinder, which provides the maximal value of indirect utility conditional on cylinder choice:

$$\varphi(\theta) = \max_{q \in q_S, q_L} v\big(\theta|q\big)$$

Define $x^*(q,\theta)$ and $c^*(q,\theta)$ as the number of trips taken and level of outside consumption that maximizes her utility subject to (3) conditional on a pair of cylinder size and type (q,θ).

Lemma 1. *Assume $f(x|q,\theta)$ is supermodular $(x,-q)$ such that for $x' > x$ we have that*

$$f(x'|q,\theta) - f(x|q,\theta)$$

is decreasing in q. Then, the policy function $x^*(q,\theta)$ is decreasing in q.

Proof. The proof follows directly from Topkis's Monotonicity Theorem. □

Proposition 1. *Assume for all θ on its support where $q' > q$ and $v(q',\theta) \geq v(q,\theta)$ the following hold:*

1. $u_c\big(c, xq|\theta\big)$ *is constant in q, x, θ, and c*
2. $f(x|q,\theta)$ *is supermodular in $(x,-q)$*
3. $p(q)$ *is increasing in q*
4. $A(\theta) \geq 0, D(\theta) \geq 0$
5. $A'(\theta) \leq 0,\ D'(\theta) \geq 0$

Then $v(\theta;q)$ is single-crossing in (q,θ) and $q^*(\theta)$ is non-decreasing.

Proof. To show the single crossing property (SCP) holds we must show that for $q' > q$ and for all $\theta' > \theta$ we have that

$$v(\theta|q') \geq v(\theta|q) \ \text{ implies that } \ v(\theta'|q') \geq v(\theta'|q) \tag{C1}$$

and

$$v(\theta|q') > v(\theta|q) \ \text{ implies that } \ v(\theta'|q') > v(\theta'|q) \tag{C2}$$

Assuming differentiability of $f(x|q,\theta)$ in θ, the total derivative of $v(\theta|q)$ in θ is, by the envelope theorem

$$dv(\theta|q) = \frac{\partial f(x^*|q,\theta)}{\partial \theta} d\theta$$

which by the functional form of utility is

$$\frac{\partial f(x^*|q,\theta)}{\partial \theta} = \frac{\partial u}{\partial c}\frac{\partial c^*(q,\theta)}{\partial \theta} = K\frac{\partial c^*(q,\theta)}{\partial \theta}$$

as the marginal utility of outside consumption is assumed constant. For the SCP to hold a sufficient condition is then

$$\frac{\partial c^*(q',\theta)}{\partial \theta} - \frac{\partial c^*(q,\theta)}{\partial \theta} \geq \ 0 \tag{4}$$

for $q' > q$ whenever $v(\theta|q') \geq v(\theta|q)$. Imposing the agent's budget constraint, optimal consumption conditional on cylinder size as a function of optimal frequency $x^*(q,\theta)$ is

$$c^*(q,\theta) = y(\theta) - p(q)x^*(q,\theta) - D(\theta)x^*(q,\theta) - p(q)A(\theta)x^*(q,\theta)^{-1}$$

so the partial derivative of $c^*(q,\theta)$ in θ is

$$\frac{\partial c^*(q,\theta)}{\partial\theta} = y'(\theta) - D'(\theta)x^*(q,\theta) - \frac{p(q)A'(\theta)}{x^*(q,\theta)}$$

A sufficient condition for the inequality in equation (4) to hold is then

$$\underbrace{y'(\theta) - D'(\theta)x^*(q_L,\theta) - \frac{p(q')A'(\theta)}{x^*(q',\theta)}}_{=\frac{\partial c(q',\theta)}{\partial\theta}} - \underbrace{y'(\theta) - D'(\theta)x^*(q,\theta) - \frac{p(q)A'(\theta)}{x^*(q,\theta)}}_{=\frac{\partial c(q,\theta)}{\partial\theta}} \geq 0 \tag{5}$$

which simplifies to

$$-D'(\theta)\Big[x^*(q',\theta) - x^*(q,\theta)\Big] - \frac{p(q')A'(\theta)}{x^*(q',\theta)} + \frac{p(q)A'(\theta)}{x^*(q,\theta)} \geq 0$$

When $x^*(q',\theta) \leq x^*(q,\theta)$, and $p(q) \leq p(q')$, this term is non-negative as long as $D'(\theta) \geq 0$ and $A'(\theta) \leq 0$. As Lemma 1 ensures $x^*(q',\theta) < x^*(q,\theta)$ the inequality holds weakly. This ensures condition (5) is met which establishes both conditions (C1) and (C2) and in turn that $v(q,\theta)$ is single-crossing in (q,θ). By the Milgrom-Shannon Monotonicity Theorem, $q^*(\theta)$ is non-decreasing. □

Proposition 2. *Assume the following additional conditions hold*

1. *Agents cannot borrow:* $c \geq 0$.
2. *Income is non-decreasing by type:* $y'(\theta) \geq 0$
3. *There exist types* $\underline{\theta}, \theta', \theta''$, *and* $\bar{\theta}$ *on the interior of the support of* θ *such that:*
   (a) *For all* $\theta < \theta'$ *we have*

   $$D'(\theta) = 0$$

   This condition amounts to assuming that there is a finite lower bound to the transaction cost in the lower tail of the distribution of types. Given Assumption (4) under Proposition 1, this lower bound is non-negative.

   (b) *For all agents who face the lowest possible transaction cost* ($D'(\theta) = 0$), *there are some types* $\underline{\theta} \leq \theta'$ *who cannot afford the large cylinder*

   $$c(q_L,\underline{\theta}) = y(\theta) - p(q_L)x - D(\theta)x - p(q_L)A(\theta)x^{-1} < 0\ , \qquad \forall x \geq 0$$

but can afford the small cylinder: $\exists x > 0$ s.t.

$$c(q_S, \underline{\theta}) = y(\theta) - p(q_S)x - D(\theta)x - p(q_S)A(\theta)x^{-1} \geq 0$$

(c) *For all* $\theta \geq \theta''$ *we have:* $\exists x > 0$ s.t.

$$c(q_L, \theta) = y(\theta) - p(q_L)x - D(\theta)x - p(q_L)A(\theta)x^{-1} > 0$$

and

$$A(\theta) = 0$$

This condition assumes types at the upper tail of the distribution have no savings costs and that some agents have a high enough income to purchase the large cylinder at some positive frequency.

Then a separating equilibrium is guaranteed to exist when unit costs, $p(q)/q$*, are non-increasing in cylinder size (i.e., linear prices or bulk discounts). In addition, when*

$$D(\theta)\left(\frac{1}{q_S} - \frac{1}{q_L}\right) \geq \frac{p(q_L)}{q_L} - \frac{p(q_S)}{q_S}$$

for a nonzero measure of agents of type $\bar{\theta} \geq \theta''$*, a separating equilibrium exists even under bulk premia.*

Proof. We can show that assumptions (1)-(3) above are sufficient conditions for a separating equilibrium to exist under non-increasing unit costs by showing that the following two conditions hold. First, we must show that all types below $\underline{\theta}$ choose the small cylinder. Second, we must show that all agents with $\theta = \bar{\theta}$ choose the large cylinder. Proposition 1 then guarantees that all types $\theta > \bar{\theta}$ also choose q_L, which ensures that a positive measure of agents choose each type of cylinder.

Show $q(\underline{\theta}) = q_S$***:*** Fix (q, θ). Consider the level of x which maximizes c, which solves

$$\min_{x \in \mathbb{R}_+} \left\{ p(q)x + D(\theta)x + p(q)A(\theta)x^{-1} \right\} \tag{6}$$

and satisfies the FOC of

$$p(q) + D(\theta) - \frac{p(q)A(\theta)}{2x^2} = 0 \tag{7}$$

Let

$$\hat{x}(q, \theta) = \sqrt{\frac{p(q)A(\theta)}{p(q) + D(\theta)}} \tag{8}$$

denote the solution to Equation (7), or the consumption-maximizing frequency for an agent of type θ, fixing their cylinder size. Note that

$$\frac{\partial \hat{x}(q,\theta)}{\partial \theta} \leq 0 \tag{9}$$

as $D'(\theta) \geq 0$ and $A'(\theta) \leq 0$. Plugging $\hat{x}$ into the budget constraint, the maximum achievable consumption level conditional on (q,θ) is

$$\hat{c}(q,\theta) = y(\theta) - p(q)\hat{x} - D(\theta)\hat{x} - p(q)A(\theta)\hat{x}^{-1}$$

Examine how $\hat{c}(q,\theta)$ changes with q:

$$\frac{\partial \hat{c}(q,\theta)}{\partial q} = -\frac{\partial p(q)}{\partial q}\hat{x} - p(q)\frac{\partial \hat{x}}{\partial p(q)} - D(\theta)\frac{\partial \hat{x}}{\partial p(q)} - \frac{\partial p(q)}{\partial q}A(\theta)\hat{x}^{-1} + \frac{\partial \hat{x}}{\partial p(q)}\frac{p(q)A(\theta)}{2\hat{x}^2}$$

which simplifies to

$$\frac{\partial \hat{c}(q,\theta)}{\partial q} = -\frac{\partial p(q)}{\partial q}\hat{x} - \frac{\partial p(q)}{\partial q}A(\theta)\hat{x}^{-1} \tag{10}$$

because at the consumption-maximizing value of x, $\hat{x}$, we have

$$p(q) + D(\theta) - \frac{p(q)A(\theta)}{2\hat{x}^2} = 0$$

by equation (7). As $p(q)$ is increasing in q by Assumption 3 in Proposition 1, maximum feasible outside consumption is strictly decreasing in cylinder size as equation (10) gives that $\frac{\partial \hat{c}(q,\theta)}{\partial q} < 0$. Consider now how changing θ affects maximum feasible consumption conditional on $q = q_L$:

$$\frac{\partial \hat{c}(q_L,\theta)}{\partial \theta} = y'(\theta) - D'(\theta)\hat{x} - p(q_L)A'(\theta)\hat{x}^{-1}$$

where the $\frac{\partial x(\theta,q)}{\partial \theta}$ terms again drop because of the envelope theorem. Then for $\theta < \theta'$, when $D'(\theta) = 0$, $\hat{c}$ is nondecreasing in θ as $y'(\theta) \geq 0$ and $A'(\theta) \leq 0$ by the assumptions of Proposition 1. This ensures for all types $\theta < \underline{\theta}$ where

$$\hat{c}(q_L,\underline{\theta}) = y(\theta) - p(q_L)\hat{x} - D(\theta)\hat{x} - p(q_L)A(\theta)\hat{x}^{-1} \leq 0$$

the large cylinder is not within the agent's budget constraint. By Assumption (3b) in Proposition 2 these agents can only only afford the small cylinder. Their choice set is then a singleton and it is optimal for them to choose the smaller cylinder.

Show $q(\bar{\theta}) = q_L$***:*** It is sufficient to show that for all levels of $xq > 0$ is at least as cheap (in terms of lost outside consumption) for some types $\theta > \theta''$ to buy the large cylinder as $\bar{\theta} \geq \theta''$. Fix some level of gas consumption $xq = G$. The expenditure required for an individual of type $\theta \geq \theta''$ who achieves this level of gas consumption using the small cylinder is

$$X_S = \frac{p(q_S)G}{q_S} + D(\theta)\frac{G}{q_S} + \frac{p_S q_S A(\theta)}{G} = \frac{p(q_S)G}{q_S} + D(\theta)\frac{G}{q_S}$$

where the last equality follows from the assumption that $A(\theta) = 0$ for $\theta \geq \theta''$. The analogous expenditure for the larger cylinder becomes

$$X_L = \frac{p(q_L)G}{q_L} + D(\theta)\frac{G}{q_L}$$

For total outside consumption to be at least as large with the large cylinder as it is with the small cylinder, for a given type θ and for all $G > 0$, we require that

$$X_L - X_S \leq 0 \tag{11}$$

This inequality simplifies to

$$D(\theta)\left(\frac{1}{q_S} - \frac{1}{q_L}\right) + \frac{p(q_S)}{q_S} - \frac{p(q_L)}{q_L} \geq 0, \tag{12}$$

which always holds for any difference in unit costs such that

$$\frac{p(q_S)}{q_S} \geq \frac{p(q_L)}{q_L},$$

This is satisfied when unit prices are non-increasing in q (e.g., the cases of linear pricing or bulk discounts). Then as $D(\theta) \geq 0$ and $D'(\theta) \geq 0$ this inequality always holds for $\theta > \theta''$ and will also always hold for any quantity of gas $G > 0$. When unit prices are non-increasing this implies it is always optimal for types $\theta \geq \theta''$ to choose the large cylinder and $q(\theta'') = q_L$.

The inequality (12) also establishes the conditions under which separation can occur under bulk premia. If there exists $\bar{\theta} \geq \theta''$ such that for all levels of $xq > 0$ it is cheaper to buy the large cylinder

$$D(\theta)\left(\frac{1}{q_S} - \frac{1}{q_L}\right) \geq \frac{p(q_L)}{q_L} - \frac{p(q_S)}{q_S} \tag{13}$$

even in the presence of bulk premia where:

$$\frac{p(q_S)}{q_S} < \frac{p(q_L)}{q_L}$$

then a separating equilibrium exists in the presence of bulk premia. This completes the proof, as we have shown given $\exists \underline{\theta}, \theta', \theta'', \bar{\theta}$, on the interior of the type support when unit costs are non-increasing in cylinder size, or when inequality (13) is met under bulk premia, a separating equilibrium exists.

Note the critical importance of frictions in the model. Return to the crossing condition in (12) for the case of agents of type $\theta < \theta''$ where $A(\theta) > 0$. A sufficient condition for agents choosing the large cylinder is that

$$D(\theta)\left(\frac{1}{q_S} - \frac{1}{q_L}\right) + \frac{p(q_S)}{q_S} - \frac{p(q_L)}{q_L} \geq A(\theta)\frac{p_L q_L - p_S q_S}{G^2} \tag{14}$$

holds for *all* gas consumption levels $G > 0$. In the absence of savings frictions ($A(\theta) = 0$), all types will choose the larger cylinder whenever unit costs are non-increasing in cylinder size. The presence of savings frictions leads some agents to choose the smaller cylinder in order to avoid large savings costs. Likewise, the presence of transportation costs ensures some types choose the large cylinder. If pricing was linear, $A(\theta) > 0$, and transaction costs were zero $D(\theta) = 0$, no agents would ever choose the large cylinder. The presence of both frictions allows for both cylinder types to be optimal for certain agents.

□

C Sufficient statistic approach

Selection condition

The left hand side of the inequality in (13), the condition that governs the choice of the large cylinder over the small one, can be approximated using a Taylor expansion:[A.1]

$$v(p(q_L);\theta,q_L) - v(p(q_S);\theta,q_S) \approx \left(\left.\frac{\partial v(p(q);\theta,q))}{\partial q}\right|_{q_S} + \frac{\partial v(p(q);\theta,q))}{\partial p(q)}\left.\frac{\partial p(q)}{\partial q}\right|_{q_S}\right) \times (q_L - q_S)$$

$$= \left.\frac{\partial v(p(q);\theta,q))}{\partial q}\right|_{q_S} \times (q_L - q_S) + \left.\frac{\partial v(p(q);\theta,q))}{\partial p(q)}\right|_{q_S} \frac{p(q_L) - p(q_S)}{q_L - q_s} \times (q_L - q_S)$$

$$= \left.\frac{\partial v(p(q);\theta,q))}{\partial q}\right|_{q_S} \times (q_L - q_S) + \left.\frac{\partial v(p(q);\theta,q))}{\partial p(q)}\right|_{q_S} (p(q_L) - p(q_S))\,.$$

So, (13) approximately holds if

$$\left.\frac{\partial v(p(q);\theta,q))}{\partial q}\right|_{q_S} \times (q_L - q_S) + \left.\frac{\partial v(p(q);\theta,q))}{\partial p(q)}\right|_{q_S} (p(q_L) - p(q_S)) > 0$$

or if

$$-\frac{\left.\frac{\partial v(p(q);\theta,q))}{\partial q}\right|_{q_S}}{\left.\frac{\partial v(p(q);\theta,q))}{\partial p(q)}\right|_{q_S}} > \frac{p(q_L) - p(q_S)}{q_L - q_S}. \tag{15}$$

[A.1] For robustness, we have also implemented a second order expansion. We obtained very similar results on selection, so we kept the first-order expansion. (Results not in the paper.)

The approximate condition for selection in (15) is very intuitive. It states that for the individual to choose the large cylinder, the marginal willingness to pay (MWTP) evaluated at the quantity in the small cylinder needs to be larger than the per kg price difference in prices between the two sizes.

Next, we show that the MWTP in (15) can be expressed as the difference in willingness to pay (WTP) for exchanging the large and the small cylinders. The MWTP for a given size is obtained by solving the following optimization problem:

$$\max p \text{ s.t. } v(p;\theta,q) = v_0(\theta) \tag{16}$$

The constraint in (16) defines an implicit function, $p(\theta, q)$. We can solve for the derivative of this function with respect to q by differentiating both sides of the equality with respect to p and q

$$\frac{\partial v(p;\theta,q)}{\partial p}dp + \frac{\partial v(p;\theta,q)}{\partial q}dq = 0;$$

thus

$$-\frac{\frac{\partial v(p;\theta,q)}{\partial q}}{\frac{\partial v(p;\theta,q)}{\partial p}} = \frac{\partial p(\theta,q)}{\partial q}. \tag{17}$$

We can approximate the right hand side of (17) as the difference in WTP across the two sizes:

$$\left.\frac{\partial p(\theta,q)}{\partial q}\right|_{q_S} \times (q_L - q_S) \approx p(\theta,q_L) - p(\theta,q_S). \tag{18}$$

Substituting (18) into (17), we get

$$\left.-\frac{\frac{\partial v(p;\theta,q)}{\partial q}}{\frac{\partial v(p;\theta,q)}{\partial p}}\right|_{q_S} \times (q_L - q_S) \approx p(\theta,q_L) - p(\theta,q_S)$$

$$\left.-\frac{\frac{\partial v(p;\theta,q)}{\partial q}}{\frac{\partial v(p;\theta,q)}{\partial p}}\right|_{q_S} \approx \frac{p(\theta,q_L) - p(\theta,q_S)}{q_L - q_S}. \tag{19}$$

Thus, we can approximate the choice condition in (13) with

$$\frac{p(\theta,q_L) - p(\theta,q_S)}{q_L - q_S} > \frac{p(q_L) - p(q_S)}{q_L - q_S}$$

or

$$p(\theta,q_L) - p(\theta,q_S) > p(q_L) - p(q_S). \tag{20}$$

The share of individuals who will choose the large cylinder under (exchange) prices p_L, p_S is given by

$$S_L = \int_{-\infty}^{\infty} \mathbf{1}\left(p(\theta, q_L) - p(\theta, q_S) > p(q_L) - p(q_S)\right) f(\theta)\theta$$

Note that this expression depends on θ only through the difference in WTP across sizes, $p(\theta, q_L) - p(\theta, q_S)$. Thus, this difference is a sufficient statistic for selection. To simplify the discussion below, we denote this sufficient statistic as

$$\tilde{\theta} \equiv \tilde{\theta}(\theta) = p(\theta, q_L) - p(\theta, q_S).$$

Adding a random disturbance to the indirect utility

Although the selection rule in (13) is straightforward, it is likely that our elicitation mechanism for WTP (explained in Section 4) will result in response errors. To allow for errors, we add a random disturbance to the true indirect utility. An individual chooses the large cylinder if

$$v(p(q_L); \theta, q_L) - v(p(q_S); \theta, q_S) + \varepsilon_L - \varepsilon_S > 0. \tag{21}$$

We can approximate the left hand side of inequality (21) using a first order Taylor expansion

$$\begin{aligned} v(p(q_L); \theta, q_L) - v(p(q_S); \theta, q_S) + \varepsilon_L - \varepsilon_S \approx & \\ & \left(\left.\frac{\partial v(p(q);\theta,q))}{\partial q}\right|_{q_S} + \frac{\partial v(p(q);\theta,q))}{\partial p(q)} \left.\frac{\partial p(q)}{\partial q}\right|_{q_S} \right) \\ & \times (q_L - q_S) + \varepsilon_L - \varepsilon_S. \end{aligned}$$

Rearranging terms, we have that (21) approximately holds if

$$-\frac{\left.\frac{\partial v(p(q);\theta,q))}{\partial q}\right|_{q_S}}{\left.\frac{\partial v(p(q);\theta,q))}{\partial p(q)}\right|_{q_S}} > \frac{(p(q_L) - p(q_S))}{(q_L - q_S)} + \frac{(\varepsilon_S - \varepsilon_L)}{-\left.\frac{\partial v(p(q);\theta,q))}{\partial p(q)}\right|_{q_S} \times (q_L - q_S)}. \tag{22}$$

Substituting (19) in (22) yields

$$\frac{p(\theta, q_L) - p(\theta, q_S)}{q_L - q_S} > \frac{p(q_L) - p(q_S)}{q_L - q_S} + \frac{(\varepsilon_S - \varepsilon_L)}{-\left.\frac{\partial v(p(q);\theta,q))}{\partial p(q)}\right|_{q_S} \times (q_L - q_S)}$$

or

$$p(\theta, q_L) - p(\theta, q_S) > p(q_L) - p(q_S) - \frac{(\varepsilon_S - \varepsilon_L)}{\left.\frac{\partial v(p(q);\theta,q))}{\partial p(q)}\right|_{q_S}} \tag{23}$$

which is similar to (5), except that the term $-\frac{(\varepsilon_S - \varepsilon_L)}{\left.\frac{\partial v(p(q);\theta,q))}{\partial p(q)}\right|_{q_S}}$ needs to be subtracted from the difference in prices across sizes. If we assume ε_S and ε_L are *iid* Extreme Value, we can use a logit model to approximate the probability that a given individual chooses the large cylinder, conditional on her sufficient statistic for selection type, $\tilde{\theta}$, and the price menu she faces, p_L, p_S.

$$\Pr\left(\left.\frac{\partial v(p(q);\theta,q))}{\partial p(q)}\right|_{q_S}\left(\tilde{\theta}-(p(q_L)-p(q_S))\right)>\varepsilon_L-\varepsilon_S\right)=\frac{\exp\left(\left.\frac{\partial v(p(q);\theta,q))}{\partial p(q)}\right|_{q_S}\frac{1}{\sigma_\varepsilon}\left(\tilde{\theta}-(p(q_L)-p(q_S))\right)\right)}{1+\exp\left(\left.\frac{\partial v(p(q);\theta,q))}{\partial p(q)}\right|_{q_S}\frac{1}{\sigma_\varepsilon}\left(\tilde{\theta}-(p(q_L)-p(q_S))\right)\right)} \tag{24}$$

Equation (24) is the basis of a logit model we estimate using observed cylinder size choices under different price pairs. In this logit estimation we treat $\gamma=\left.\frac{\partial v(p(q);\theta,q))}{\partial p(q)}\right|_{q_S}\frac{1}{\sigma_\varepsilon}$ as the only parameter to be estimated, while $p(q_L),p(q_S)$ and $\tilde{\theta}$ are treated as data. Section 4.1 explains the data collection procedure on cylinder choices and $\tilde{\theta}$. Once we estimate γ, we can produce a different set of probability weights for each household, where each element of the set corresponds to a different value of $(p(q_L)-p(q_S))$. These estimated weights take the form

$$\frac{\exp\left(\hat{\gamma}\left(\tilde{\theta}-(p(q_L)-p(q_S))\right)\right)}{1+\exp\left(\hat{\gamma}\left(\tilde{\theta}-(p(q_L)-p(q_S))\right)\right)} \tag{25}$$

In other words, conditional on observing $\tilde{\theta}$ for each household, (24) will approximate the probability that the household picks the large cylinder for a given pair of prices, $(p(q_L),p(q_S))$. These probability weights have two roles in the design of an optimal price menu.[A.2] First, we use them to account for selection in the demand estimation step (Step 4 of Section 4.2). Second, we use them to evaluate the policy objective function and resource constraint given in (1) and (2). Note that in these two functions, the policy maker cares about the demand and public expenditure averages, which account for selection into size based on type, θ. The probability weights in (24) will allow us to compute the empirical version of these averages.

[A.2] We also use these probability weights for the power calculation corresponding to our stage 2 Test. Online supplement II describes this power calculation in detail.

Online Supplement to:

Targeting subsidies through price menus

I Implementation and measurement

I.1 Depots

Depot managers only accepted a cylinder for an exchange if the following conditions were met:

- The person desiring to exchange a cylinder needs to present their participant ID card. This ID card contains the participant's depot as well as the cylinder size and exchange price in their contract.
- The information on the ID card must match that in the tablet.
- The person must be attempting to exchange at their assigned depot, as specified on the ID card, not any other depot.
- The person must be attempting to exchange for the same size cylinder as on the ID card.
- The exchange price paid must match the ID card, and tablet information and must be paid in full.
- The cylinder should be study-branded.
- The cylinder ID should match the cylinder ID previously exchanged by that participant.
- The participant must still be within their 3 month window (not past their program expiration date).

I.2 MPL implementation

Implementation of the MPLs involved several steps: first, a "mock MPL" was completed, followed by the LPG MPLs. The mock MPL was a simpler MPL involving choices to purchase two low value items: a package of 25 tea bags and a packet of milk powder. The MPL had four rows - two were binary choices between purchasing the milk powder (tea bags) and purchasing nothing and two were choices between purchasing milk powder, tea bags, or nothing. Prices of the goods depended on the row, but were the same across all participants. At the end, one scenario was randomly drawn to implement. Respondents were informed of the outcome and asked to pay the specified purchase price for the good, if they had selected to buy the good in that scenario. Respondents who didn't agree with the outcome (reneged) were allowed to change their choices. After gaining an understanding of the MPL exercise through the mock MPL, the main MPL exercise was implemented.

Surveyors followed specified scripts in their tablets to conduct the LPG MPLs. Surveyors introduced the respondents to the LPG "starter kit", explained the deposit system, and described how exchanges work. Visual aids assisted the surveyors in describing the LPG starter kit and depot exchanges. A series of comprehension questions were asked to assess understanding of the key features of the study program. If a respondent did not know the answer or misunderstood, the surveyor clarified before proceeding.

There were three LPG MPLs; each was composed of a series of choices between similar options where only the exchange price (subsidy) was varied. MPL A (B) rows involved a choice between purchasing the small (large) cylinder or purchasing nothing, at varying exchange subsidies. In MPL C the choice was between purchasing the small cylinder, large cylinder, or nothing. The exchange subsidy for the large cylinder was held constant within MPL C, but the exchange subsidy for the small varied across rows.

The LPG MPLs were dynamic, where exchange prices presented in some rows depended on choices in prior rows. MPL A and B were each composed of a static and a dynamic portion. The first part was the static portion where all participant saw the same four rows.[i] The second part was the dynamic portion where the exchange prices depended on choices in the static portion. This design enables a more precise estimate of WTP.

The dynamic prices were determined automatically in the tablet for MPLs A and B. First, the highest price the participant was willing to purchase the small cylinder at was determined ($P_{max,A}$). This was treated as the lower bound on their willingness to pay to purchase the cylinder.[ii] Second, the next highest price in the static rows of MPL A was determined. For those that always chose to purchase the cylinder, this was the market price[iii] Third, the integer intervals for the dynamic rows were determined by dividing the range in five to get the increments (d_A). A similar calculation followed for MPL B. See Table I for the static and dynamic portions of MPLs A and B.

MPL C also was dynamic; the prices for purchasing the small and large cylinders in each row depended on prior decisions in MPLs A and B, respectively. The exchange price of the large cylinder was fixed across all four scenarios within Table C while the exchange price of the small cylinder varied. The price of the small (large) cylinder was based on the highest price the participant purchased the cylinder in the full MPL A (B). The price of the large in MPL C was fixed across rows at 80, 90, or 100% of this maximum WTP.[iv] The exchange prices of the small cylinder in MPL C were calculated as three prices in even intervals between zero and the participant's max WTP in MPL A and included the maximum WTP. Participants were randomized to see the rows in descending/ascending order in the exchange price of the small cylinder.

At the end of the LPG MPL, one of the 20 scenarios was randomly chosen to implement on the spot. If the respondent agreed to follow through with the decision, and they drew a scenario in which they chose a small or large cylinder, then they would schedule a delivery. If a respondent

[i]While exchange prices in the rows were the same in the static portion of MPLs A and B, the order of the scenarios in each MPL (ascending/descending in exchange price) was randomized on-the-spot within the tablet. Table C also had a randomized order (ascending/descending in the exchange price of the 3 kg cylinder).

[ii]For those that always made the choice to purchase nothing in the static rows, this was treated as 0.

[iii]Based on a market per kg price of GH¢ 10.3.

[iv]Originally this was fixed at 80%, but was modified to 90% for those who exhibited switching behavior and 100% for those who exhibited never-switching behavior, always selecting the large cylinder.

Table I: MPL Tables A and B

Panel A: MPL A			
	Scenario	Option A: 3 kg cylinder	Option B: Nothing
Static	1	GH¢ 6	0
	2	GH¢ 12	0
	3	GH¢ 22	0
	4	GH¢ 28	0
Dynamic	5	GH¢ $P_{max,A} + d_A$	0
	6	GH¢ $P_{max,A} + 2d_A$	0
	7	GH¢ $P_{max,A} + 3d_A$	0
	8	GH¢ $P_{max,A} + 4d_A$	0
Panel B: MPL B			
	Scenario	Option A: 14.5 kg cylinder	Option B: Nothing
Static	9	GH¢ 30	0
	10	GH¢ 59	0
	11	GH¢ 104	0
	12	GH¢ 135	0
Dynamic	13	GH¢ $P_{max,B} + d_B$	0
	14	GH¢ $P_{max,B} + 2d_B$	0
	15	GH¢ $P_{max,B} + 3d_B$	0
	16	GH¢ $P_{max,B} + 4d_B$	0

Note: Option A to purchase the small (large) cylinder in MPL A (B) was presented as including paying a deposit price and indicated that the price was the per exchange price. Option B to purchase nothing had no cost but also no access to the exchanges. Participants were randomly assigned to see the MPL rows in either ascending or descending order of the exchange price for MPLs A and B, respectively.

reneged on her decision, this was recorded.[v]

I.3 MPL data corrections

We approximate $p(\theta, q_S)$ from MPL A using the maximum price at which individuals take up the small cylinder and $p(\theta, q_L)$ as the maximum price at which individuals take the large cylinder in MPL B. For 333 participants, we are missing proxies of either $p(\theta, q_L)$ or $p(\theta, q_S)$ or both because there was no switching behavior. In these cases, we proxy $\tilde{\theta}_i$ with either $p(\theta_i, q_L) - p(\theta_i, q_S)$ from MPL C, when there was switching behavior in MPL C, or the lower/upper bound from MPL C, when there was no switching behavior in MPL C. For example, if a participant chose the large size for every single price pair in MPL C, then the highest price difference between p_L and p_S in MPL C becomes the lower bound for $\tilde{\theta}_i$. Note that this bound differs across individuals based on implementation of MPL C.[iv]

In 67 cases, $p(\theta_i, q_L) - p(\theta_i, q_S)$ from MPLs A and B is lower than the lower bound from MPL C. In these cases we also proxy $\tilde{\theta}_i$ with the lower bound from MPL C. Similarly, in 9 cases, $p(\theta_i, q_L) - p(\theta_i, q_S)$ from MPLs A and B is higher than the upper bound from MPL C. We use the upper bound from MPL C in these cases.

I.4 Heterogeneity indices

The motivation for screening on cylinder size stems from heterogeneity in the opportunity cost of time and demand for liquidity in the population. The preference for the small cylinder should be higher among households with a lower opportunity cost of time and/or a higher demand for liquidity.

We collected baseline data to construct indices as follows, where a higher score in the index reflects a greater likelihood of preferring the small cylinder:

- **Opportunity cost of time and transaction costs**:
  - *Would you wait in line for a discount?*
    Willingness to wait for a discount (of different sizes) classifies households into low or high time cost. We score these such that households with a lower opportunity cost of time receive a higher score in the index.
  - *Access to transportation?*
    Access to transportation may mitigate the time cost of exchange, both by reducing the time needed for exchange and also the hassle cost. The benefit of transportation is higher for the large cylinder. This variable enters the index negatively; access to

[v] The respondent was not allowed to change any of their responses (as they were in the mock MPL); instead, they were not allowed to enter the program. Only 1 respondent reneged in the LPG MPL.

transport offsets the effect of a low opportunity cost of time in preferring the small cylinder.

- **Liquidity constraints**:
  - *How long would it take you to gather cash?*
    The time required to gather specified amounts of cash classifies households into low, medium and high liquidity constraints.
  - *How long would it take you to pay back a loan?*
    Similarly, a longer time to pay back a loan is associated with greater liquidity constraints.
- **Small purchase index**:
  - *Do you purchase small or large bags of charcoal?*
    Households purchase other energy in ways that mimic LPG; households with a preference for purchasing small bags of charcoal rather than large bags should also show this preference in LPG purchases.
  - *Wealth index*
    The sampling data from the KHDSS includes a wealth index based on assets. Higher wealth households are likely to be less liquidity constrained and have a higher opportunity cost of time. Wealth (and other inputs to these indices) may also be correlated with heterogeneity in unobserved preferences for LPG.
- **Aggregate index**: We construct a final, aggregate index that combines the three and reflects an overall preference for the small cylinder.

II Power calculations

Shorthand $p_L = p(q_L)$ and $p_S = p(q_S)$ and recall that $q_S = 3$ and $q_L = 14.5$. Also, denote the difference in prices between the two cylinders as $\Delta p = p_L - p_S$. The power calculation requires the variance of household demand for LPG:

$$\begin{aligned} Q(p_L, p_S) := 3 \quad & \Big(x(p_S; \theta_i, \epsilon_{Si}, q_S) \times \mathbf{1}\left[\varepsilon_i > \gamma(\theta_i - (p_L - p_S))\right] \Big) \\ & + 14.5 \Big(x(p_L; \theta_i, \epsilon_{Li}, q_L) \times \mathbf{1}\left[\varepsilon_i \leq \gamma(\theta_i - (p_L - p_S))\right] \Big) \end{aligned} \tag{i}$$

where $\varepsilon_i \equiv \varepsilon_{Si} - \varepsilon_{Li}$ is the difference in the Type 1 Extreme Value errors associated with each cylinder choice and is distributed logistic. Note that (i) is the sum of two correlated random variables as θ_i and ε_i appear in the two summands. Assume that we can write exchange demand for the small cylinder as

$$\begin{aligned} x(p_S; \theta_i, \epsilon_{Si}, q_S) & = \beta_0(\theta_i, q_S) + \beta_1(\theta_i, q_S)p_S + \beta_2(\theta_i, q_S)p_S^2 + \epsilon_{Si} \\ & = \hat{x}(p_S; \theta_i, q_S) + \epsilon_{Si} \end{aligned} \tag{ii}$$

Further, we will assume that

$$\begin{aligned} \mathbb{E}(\epsilon_{Si}|\theta_i) = \mathbb{E}(\epsilon_{Si}) = 0 \\ \mathbb{E}(\epsilon_{Li}|\theta_i) = \mathbb{E}(\epsilon_{Li}) = 0 \end{aligned} \tag{iii}$$

and

$$\varepsilon_i \perp \theta_i, \epsilon_{Si}, \epsilon_{Li} \tag{iv}$$

Let

$$\pi(\theta_i, \Delta p) \equiv \frac{1}{1 + \exp(\gamma(\theta - (p_L - p_S)))} = \mathbb{E}\Big[\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i - (p_L - p_S))\right] \Big| \theta_i\Big]$$

be the conditional probability a type-θ_i agent chooses the small cylinder and $(1 - \pi(\theta_i, \Delta p))$ be its compliment for the large cylinder. Note these probabilities are independent of x conditional on θ_i, as $\varepsilon_i \perp \theta_i, \epsilon_{Si}, \epsilon_{Li}$

We can derive the variance of (i) by using the formula $\mathrm{Var}(aX + bY) = a^2\mathrm{Var}(X) + b^2\mathrm{Var}(Y) + 2ab\mathrm{Cov}(X, Y)$. The variance of (i) is equal to

$$\begin{aligned}\text{Var}(Q) = 3^2\times \quad & \text{Var}\Big(x(p_S;\theta_i,\epsilon_{Si},q_S)\Big] \times \mathbf{1}\left[\varepsilon_i > \gamma(\theta_i - (p_L - p_S))\right] \Big) \\ + \quad & 14.5^2 \times \text{Var}\Big(x(p_L;\theta_i,\epsilon_{Li},q_L)\Big] \times \mathbf{1}\left[\varepsilon_i \leq \gamma(\theta_i - (p_L - p_S))\right] \Big) \\ + \quad & 87 \times \text{Cov}\Big(x(p_S;\theta_i,\epsilon_{Si},q_S) \times \mathbf{1}\left[\varepsilon_i > \gamma(\theta_i - (p_L - p_S))\right] \ , \\ & x(p_L;\theta_i,\epsilon_{Li},q_L) \times \mathbf{1}\left[\varepsilon_i \leq \gamma(\theta_i - (p_L - p_S))\right] \Big)\end{aligned} \tag{Expanded $V(Q)$}$$

First, let us examine the covariance term in (Expanded $V(Q)$). We will use the covariance identity, $\text{Cov}(X,Y) = \mathbb{E}(XY) - \mathbb{E}(X)\mathbb{E}(Y)$. Thus,

$$\begin{aligned}&\text{Cov} \quad \Big(x(p_S;\theta_i,\epsilon_{Si},q_S)\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i - (p_L - p_S))\right] \ , \ x(p_L;\theta_i,\epsilon_{Li},q_L)\big(\mathbf{1}\left[\varepsilon_i \leq \gamma(\theta_i - (p_L - p_S))\right]\big)\Big) \\ &= \mathbb{E}\Big[x(p_S;\theta_i,\epsilon_{Si},q_S)\big(\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i - (p_L - p_S))\right]\big) x(p_L;\theta_i,\epsilon_{Li},q_L)\big(\mathbf{1}\left[\varepsilon_i \leq \gamma(\theta_i - (p_L - p_S))\right]\big)\Big] \\ &-\mathbb{E}\Big[x(p_S;\theta_i,\epsilon_{Si},q_S)\big(\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i - (p_L - p_S))\right]\big)\Big]\mathbb{E}\Big[x(p_L;\theta_i,\epsilon_{Li},q_L)\big(\mathbf{1}\left[\varepsilon_i \leq \gamma(\theta_i - (p_L - p_S))\right]\big)\Big]\end{aligned}$$

The first expectation is equal to zero as the expectation of the product of random variables $\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i - (p_L - p_S))\right]$ and $\mathbf{1}\left[\varepsilon_i \leq \gamma(\theta_i - (p_L - p_S))\right]$ is always zero for all realizations of ε_i and θ_i. Thus

$$\begin{aligned}&\text{Cov} \quad \Big(x(p_S;\theta_i,\epsilon_{Si},q_S)\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i - (p_L - p_S))\right] \ , \ x(p_L;\theta_i,\epsilon_{Li},q_L)\big(\mathbf{1}\left[\varepsilon_i \leq \gamma(\theta_i - (p_L - p_S))\right]\big)\Big) = \\ &-\mathbb{E}\Big[x(p_S;\theta_i,\epsilon_{Si},q_S)\big(\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i - (p_L - p_S))\right]\big)\Big]\mathbb{E}\Big[x(p_L;\theta_i,\epsilon_{Li},q_L)\big(\mathbf{1}\left[\varepsilon_i \leq \gamma(\theta_i - (p_L - p_S))\right]\big)\Big]\end{aligned}$$

Apply the LIE in θ_i to the first RHS term above to yield

$$\begin{aligned}\mathbb{E} \quad & \Big[x(p_S;\theta_i,\epsilon_{Si},q_S)\big(\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i - (p_L - p_S))\right]\big)\Big] = \\ &= \mathbb{E}\Big[\mathbb{E}\Big[x(p_S;\theta_i,\epsilon_{Si},q_S)\big(\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i - (p_L - p_S))\right]\big)\Big|\theta_i\Big]\Big] \\ &= \mathbb{E}\Big[\pi(\theta_i,\Delta p)\mathbb{E}\Big[x(p_S;\theta_i,\epsilon_{Si},q_S)\Big|\theta_i\Big]\Big] \\ &= \mathbb{E}\Big[\pi(\theta_i,\Delta p)\hat{x}(p_S;\theta_i,q_S)\Big]\end{aligned} \tag{LIE}$$

where the last two lines follow from the independence of $\varepsilon_i \perp \theta_i, \epsilon_{S,i}, \epsilon_{L,i}$ in assumption (iv) and the substitution of $\hat{x}$ for x follows from the mean independence assumption in (iii) which gives $\mathbb{E}[\epsilon_{\cdot,i}|\theta_i] = 0$. An identical procedure for the case of the large cylinder as done in (LIE) gives

$$\mathbb{E}\Big[x(p_L;\theta_i,\epsilon_{Li},q_L)\big(\mathbf{1}\left[\varepsilon_i \le \gamma(\theta_i-(p_L-p_S))\right]\big)\Big] = \mathbb{E}\Big[\big(1-\pi(\theta_i,\Delta p)\big)\hat{x}(p_L;\theta_i,q_L)\Big]$$

We may then write the covariance above as

$$\begin{aligned}
&\text{Cov}\quad\Big(x(p_S;\theta_i,\epsilon_{Si},q_S)\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i-(p_L-p_S))\right]\ ,\ x(p_L;\theta_i,\epsilon_{Li},q_L)\big(\mathbf{1}\left[\varepsilon_i \le \gamma(\theta_i-(p_L-p_S))\right]\big)\Big) = \\
&= -\mathbb{E}\left[\hat{x}(p_S;\theta_i,q_S)\big(\pi(\theta_i,\Delta p)\big)\right]\mathbb{E}\left[\hat{x}(p_L;\theta_i,q_L)\big(1-\pi(\theta_i,\Delta p)\big)\right] < 0
\end{aligned} \tag{v}$$

as π and $\hat{x}$ must be non-negative. This implies we can write (Expanded $V(Q)$) as

$$\begin{aligned}
\text{Var}(Q) = 9 \quad &\times \text{Var}\Big(x(p_S;\theta_i,\epsilon_{Si},q_S)) \times \mathbf{1}\left[\varepsilon_i > \gamma(\theta_i-(p_L-p_S))\right]\Big) \\
+ \quad &210.25 \times \text{Var}\Big(x(p_L;\theta_i,\epsilon_{Li},q_L) \times \mathbf{1}\left[\varepsilon_i \le \gamma(\theta_i-(p_L-p_S))\right]\Big) \\
- \quad &87 \times \left[\mathbb{E}\Big[\pi(\theta_i,\Delta p)\hat{x}(p_S;\theta_i,q_S)\Big]\mathbb{E}\Big[\big(1-\pi(\theta_i,\Delta p)\big)\hat{x}(p_L;\theta_i,q_L)\Big]\right]
\end{aligned} \tag{$V(Q)'$}$$

Next, we turn to the variance terms in ($V(Q)'$). The first variance term (concerning small cylinder demand) can be expressed as

$$\begin{aligned}
\text{Var}\quad &\Big(x(p_S;\theta_i,\epsilon_{Si},q_S)\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i-(p_L-p_S))\right]\Big) \\
&= \mathbb{E}\left[\Big(x(p_S;\theta_i,\epsilon_{Si},q_S)\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i-(p_L-p_S))\right]\Big)^2\right] \\
&- \mathbb{E}\left[x(p_S;\theta_i,\epsilon_{Si},q_S)\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i-(p_L-p_S))\right]\right]^2
\end{aligned} \tag{vi}$$

Apply the LIE to the second RHS term above. Like with the result in (LIE), recalling $\varepsilon_i \perp \epsilon_{\cdot,i}$, we can substitute $\hat{x}$ for x and rewrite the variance as

$$\begin{aligned}
\text{Var}\quad &\Big(x(p_S;\theta_i,\epsilon_{Si},q_S)\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i-(p_L-p_S))\right]\Big) \\
&= \mathbb{E}\left[\Big(x(p_S;\theta_i,\epsilon_{Si},q_S)\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i-(p_L-p_S))\right]\Big)^2\right] \\
&- \mathbb{E}\left[\hat{x}(p_S;\theta_i,q_S)\pi(\theta_i,\Delta p)\right]^2
\end{aligned} \tag{vii}$$

Expanding the first RHS term in $x(p_S;\theta_i,\epsilon_{Si},q_S) = \hat{x}(p_S;\theta_i,q_S) + \epsilon_{Si}$ gives

$$
\begin{aligned}
\text{Var} \quad & \Big(x(p_S;\theta_i,\epsilon_{Si},q_S)\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i-(p_L-p_S))\right]\Big) \\
& = \mathbb{E}\ \Bigg[\hat{x}(p_S;\theta_i,q_S)^2\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i-(p_L-p_S))\right]^2 \\
& +2\hat{x}(p_S;\theta_i,q_S)\epsilon_{S,i}\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i-(p_L-p_S))\right]^2 + \epsilon_{S,i}^2\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i-(p_L-p_S))\right]^2\Bigg] \\
& -\mathbb{E}\ \Big[\hat{x}(p_S;\theta_i,q_S)\pi(\theta_i,\Delta p)\Big]^2
\end{aligned}
\tag{viii}
$$

and applying the LIE again to each element of the expansion in the first set of brackets yields (by an analogous argument as in LIE)

$$
\begin{aligned}
\text{Var} \quad & \Big(x(p_S;\theta_i,\epsilon_{Si},q_S)\mathbf{1}\left[\varepsilon_i > \gamma(\theta_i-(p_L-p_S))\right]\Big) \\
& = \ \mathbb{E}\ \Big[\hat{x}(p_S;\theta_i,q_S)^2\pi(\theta_i,\Delta p)\Big] + \underbrace{0}_{by \text{ by } \mathbb{E}[\epsilon_i|\theta_i]=0} + \sigma_{S,i}^2\mathbb{E}\ \Big[\pi(\theta_i,\Delta p)\Big] \\
& - \ \mathbb{E}\ \Big[\hat{x}(p_S;\theta_i,q_S)\pi(\theta_i,\Delta p)\Big]^2
\end{aligned}
\tag{ix}
$$

As for a Bernoulli random variable we have

$$
\mathbb{E}\ \Big[\mathbf{1}\left[\varepsilon_i \leq \gamma(\theta_i-(p_L-p_S))\right]^2 \Big|\theta_i\Big] = \pi(\theta_i,\Delta p)
$$

By an analogous argument, the second variance term (over large cylinders) can be expressed as

$$
\begin{aligned}
\text{Var} \quad & \Big(x(p_L;\theta_i,\epsilon_{Li},q_L)\mathbf{1}\left[\varepsilon_i \leq \gamma(\theta_i-(p_L-p_S))\right]\Big) \\
& = \ \mathbb{E}\ \Big[\hat{x}(p_L;\theta_i,q_L)^2\big(1-\pi(\theta_i,\Delta p)\big)\Big] + \sigma_{L,i}^2\mathbb{E}\ \Big[\big(1-\pi(\theta_i,\Delta p)\big)\Big] \\
& - \ \mathbb{E}\ \Big[\hat{x}(p_L;\theta_i,q_L)\big(1-\pi(\theta_i,\Delta p)\big)\Big]^2
\end{aligned}
\tag{x}
$$

Putting this all together we may write equation ($V(Q)'$) as

$$
\begin{aligned}
& \text{Var}(Q) = \\
& 9\Bigg(\mathbb{E}\ \Big[\hat{x}(p_S;\theta_i,q_S)^2\pi(\theta_i,\Delta p)\Big] + \sigma_{S,i}^2\mathbb{E}\ \Big[\pi(\theta_i,\Delta p)\Big] - \mathbb{E}\ \Big[\hat{x}(p_S;\theta_i,q_S)\pi(\theta_i,\Delta p)\Big]^2\Bigg) \\
& +210.25\Bigg(\mathbb{E}\ \Big[\hat{x}(p_L;\theta_i,q_L)^2\big(1-\pi(\theta_i,\Delta p)\big)\Big] + \sigma_{L,i}^2\mathbb{E}\ \Big[\big(1-\pi(\theta_i,\Delta p)\big)\Big] - \mathbb{E}\ \Big[\hat{x}(p_L;\theta_i,q_L)\big(1-\pi(\theta_i,\Delta p)\big)\Big]^2\Bigg) \\
& -87 \times \Bigg[\mathbb{E}\Big[\pi(\theta_i,\Delta p)\hat{x}(p_S;\theta_i,q_S)\Big]\mathbb{E}\Big[\big(1-\pi(\theta_i,\Delta p)\big)\hat{x}(p_L;\theta_i,q_L)\Big]\Bigg]
\end{aligned}
\tag{xi}
$$

By a similar procedure, the variance of the average budget is given by

$$
\begin{aligned}
&\mathrm{Var}(B) = \\
&(3r - p_S)^2 \times \left(\mathbb{E}\left[\hat{x}(p_S;\theta_i,q_S)^2\pi(\theta_i,\Delta p)\right] + \sigma^2_{S,i}\mathbb{E}\left[\pi(\theta_i,\Delta p)\right] - \mathbb{E}\left[\hat{x}(p_S;\theta_i,q_S)\pi(\theta_i,\Delta p)\right]^2 \right) \\
&+(14.5r - p_L)^2 \times \left(\mathbb{E}\left[\hat{x}(p_L;\theta_i,q_L)^2\big(1-\pi(\theta_i,\Delta p)\big)\right] + \right. \\
&\left. \sigma^2_{L,i}\mathbb{E}\left[\big(1-\pi(\theta_i,\Delta p)\big)\right] - \mathbb{E}\left[\hat{x}(p_L;\theta_i,q_L)\big(1-\pi(\theta_i,\Delta p)\big)\right]^2 \right) \\
&2 \times (3r - p_S) \times (14.5r - p_L)\left[\mathbb{E}\Big[\pi(\theta_i,\Delta p)\hat{x}(p_S;\theta_i,q_S)\Big]\mathbb{E}\Big[\big(1-\pi(\theta_i,\Delta p)\big)\hat{x}(p_L;\theta_i,q_L)\Big]\right]
\end{aligned} \tag{xii}
$$

We can estimate each of the terms in (xi) and (xii) through weighted regressions using the same logit weights that we used to estimate the average elasticities under selection for each price difference.

Estimation of the variance

The variances for the average demand and the average budget are functions of terms of the form

$$\mathbb{E}\left(\pi(\Delta p;\theta_i)\hat{x}(p_S;\theta_i,q_S)\right) \tag{xiii}$$

and

$$\mathbb{E}\left(\pi(\Delta p;\theta_i)\hat{x}(p_S;\theta_i,q_S)^2\right) \tag{xiv}$$

We can estimate (xiii) by WLS as described in Section 4.1.[vi] In addition, because

$$
\begin{aligned}
x(p_S;\theta_i,\epsilon_{Si},q_S)^2 &= \beta^2_{0i} + 2\left(\beta_{0i}\beta_{1i}\right)p_S + \left(2\beta_{0i}\beta_{2i} + \beta^2_{1i}\right)p^2_S + 2\left(\beta_{1i}\beta_{2i}\right)p^3_S + \beta^2_{2i}p^4_S \\
&\quad + 2\beta_{0i}\epsilon_{Si} + 2\beta_{1i}p_S\epsilon_{si} + 2\beta_{2i}p^2_S\epsilon_{Si} + \epsilon^2_{Si} \\
&= \eta_{i0} + \eta_{i1}p_S + \eta_{i2}p^2_S + \eta_{i3}p^3_S + \eta_{i4}p^4_S + \epsilon^2_{Si}
\end{aligned} \tag{xv}
$$

When taking expectations of (xv) conditional on θ_i, we get

$$
\begin{aligned}
&\mathbb{E}\left(x(p_S;\theta_i,\epsilon_{Si},q_S)^2 \middle| \theta_i\right) \\
&\qquad = \left(\eta_{i0} + \sigma^2_{\epsilon_S}\right) + \eta_{i1}p_S + \eta_{i2}p^2_S + \eta_{i3}p^3_S + \eta_{i4}p^4_S \\
&\qquad = \hat{x}(p_S;\theta_i,q_S)^2 + \sigma^2_{\epsilon_S}
\end{aligned}
$$

[vi] In the code, we actually rescale the weights to sum to one, as

$$\mathbb{E}\left(\pi(\Delta p;\theta_i)\right)\mathbb{E}\left(\frac{\pi(\Delta p;\theta_i)}{\mathbb{E}\left(\pi(\Delta p;\theta_i)\right)}\hat{x}(p_S;\theta_i,q_S)\right)$$

where $\mathbb{E}\left(\frac{\pi(\Delta p;\theta_i)}{\mathbb{E}(\pi(\Delta p;\theta_i))}\right) = 1$ and we can estimate the weights, $\frac{\pi(\Delta p;\theta_i)}{\mathbb{E}(\pi(\Delta p;\theta_i))}$, as $\frac{1}{\exp\left(\hat{\gamma}\left(\hat{\hat{\theta}}_i - \Delta p\right)\right)} \times \frac{1}{\sum_i \pi(\Delta p,\theta_i}$. This reweighting does not affect our estimates, as Stata normalizes weights by default.

This suggests we can estimate $\hat{x}(p_S; \theta_i, q_S)^2 + \sigma^2_{\epsilon_S}$ by regressing x_S^2 on a fourth degree polynomial in price and weighting observations to match appropriately selected sample. To estimate $\hat{x}(p_S; \theta_i, q_S)^2$, we subtract $\sigma^2_{\epsilon_S}$, which is available from the WLS results for estimating the average demand.

III Robustness and implementation checks

III.1 LPG prices and inflation

We provide additional detail on the LPG price fluctuations and consumer inflation during study implementation, along with more information on how we construct our robustness checks. Note that our pre-analysis plan describes only the adjustment prior to the start of stage 2, when both inflation and LPG prices were increasing. The adjustment on June 29, 2023, over halfway through stage 2, was unanticipated when we wrote the pre-analysis plan, but necessitated by the fact that study logistics required that the prices offered at the exchange depots remained below the prices at the pump.

Price adjustment prior to stage 2 LPG prices were relatively flat during stage 1 of the study, which ran from April to mid October 2022 (see Figure I). However, between the end of stage 1 and the point (December 2022) at which we used the stage 1 data to select prices for the stage 2 treatments, LPG prices increased by 43% (at the pump) to 46% (at the refinery).[vii] Consumer inflation, on the other hand, rose steadily between April and December, with a cumulative national CPI increase of 34%. In designing the stage 2 prices, we therefore considered the fact that we needed prices to remain below pump prices, to minimize the incentive to purchase LPG outside of the study. We also anticipated that LPG prices would start to come down. We therefore adjusted the prices in the optimal price selection by the inflation rate, i.e., multiplied stage 1 prices by 1.338. These were the prices given to households when stage 2 rolled out, and also the prices we initially used to calculate the subsidy budget.

In the main text, we show "nominal" prices at each stage, i.e., the prices shown in Section 4.2 (Step 6) are based on the prices in stage 1 inflated by 34%.

Price adjustment during stage 2 As shown in Figure I, LPG prices did begin to decline shortly after we selected the subsidy levels for stage 2. They continued to fall and eventually fell below the adjusted pump price and then below the stage 1 unadjusted pump price. In late June, it was clear that households with the smallest per kg subsidy (treatment group, large cylinder) had an incentive to take the cylinder to the filling station rather than exchange at a study depot, so we chose to adjust prices in the study to follow pump prices. Specifically, we simply undid the 34% increase in prices that we had imposed before launching stage 2, and reverted back to stage 1 prices. Households were informed of this by phone, and the exchange depot software was updated to adjust the charge to households. Any household that could not be reached by phone received an in-person visit.

[vii]Pump prices are regulated by the National Petroleum Agency, which smooths consumer (pump) prices to some degree, as evidenced by the price series.

Effect on demand, subsidy budget and robustness checks The main implications of these adjustments are for how we calculate the subsidy budget. We choose to present the results in the main text following what consumers observed, which roughly follows LPG pump prices, and therefore offers a proxy for the subsidy budget. However, we could have also measured the policy cost as following exactly from the pump price, which changed at higher frequency than the subsidized price paid by households in our study. We also could have determined the subsidy budget by the refinery price, which was even more volatile. Finally, given that inflation rose steadily, making the real cost of LPG fall by an even greater amount during stage 2, we could calculate the subsidy budget in real terms, adjusting for inflation. We present each of these alternative set of assumptions in the main text, in Table A.10. Impacts on the subsidy budget are small.

For consumer demand, we can only observe what we implemented but we can speculate on how demand might have been impacted. First, in stage 2, as prices fell, and prior to the June 29 price deflation, households in the treatment group with a large cylinder might have preferred to purchase from the filling station. This would have reduced demand in the treatment group relative to the control group, potentially contributing to the lower-than-predicted demand that we observe. However, as evidence that this behavior was no more than minimal, we targeted spot checks to households most likely to be visiting filling stations during this time period and see no evidence of filling outside of the study. Second, high inflation may have tightened households' budget constraints, but also led to LPG prices declining relative to the CPI. This potentially increased both price sensitivity and overall demand, which may have contributed to our results.

Figure I: LPG price index and inflation over stage 1 and 2

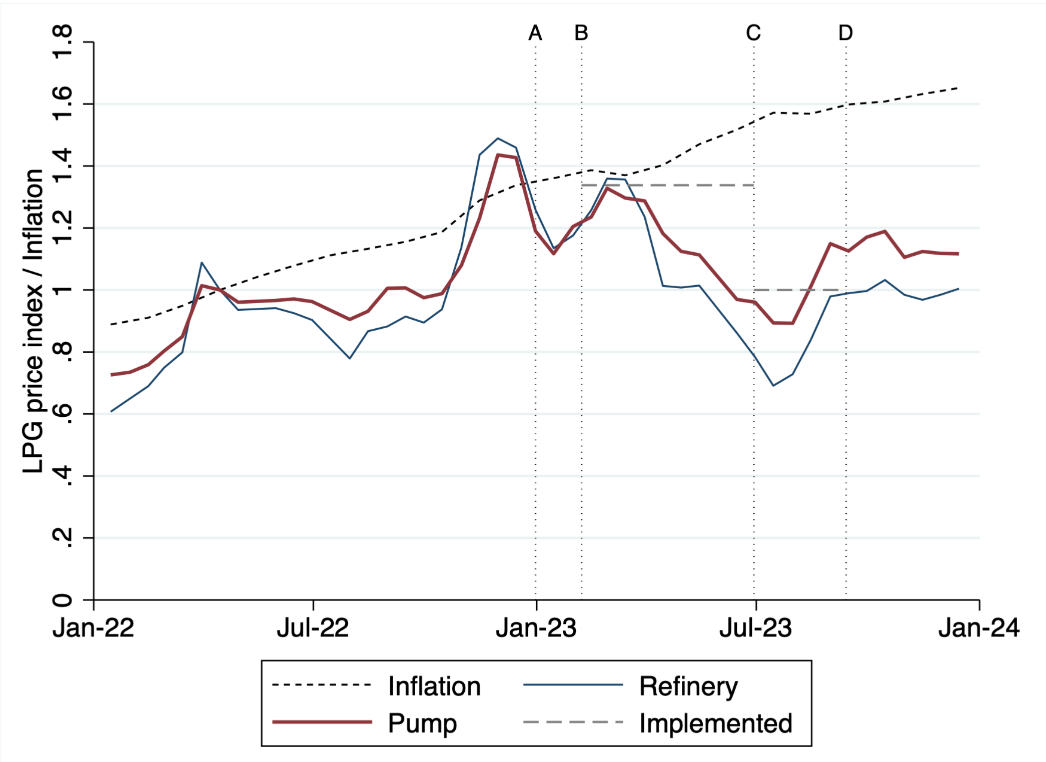

Note: The LPG price index is normalized to stage 1 prices (April 2022 = 1). Vertical lines indicate the following from left to right: A) stage 1 data analysis and menu design, B) start of stage 2 data collection, C) date of price adjustment and D) last date of stage 2 exchanges. LPG pump and refinery prices from National Petroleum Agency (2023). CPI data from Ghana Statistical Service (2023).

III.2 Spot checks

The primary objective of the spot checks was to detect LPG purchases not captured by the data collection protocols, i.e., by households filling their cylinders at filling stations. Households became eligible for spot checks three weeks after enrollment conditional on having never exchanged their cylinder. During the spot check, enumerators used a scale to weigh the study branded cylinder.

In stage 2, two rounds of spot checks were conducted. During the first round, the households were sampled with replacement from the eligible list each week. During the second round, to better balance the sample, eligible households were stratified by treatment and cylinder size. Enumerators conducting spot checks would visit the households on the sampled list each week. If the cylinder was at home, then the cylinder was weighed with an analog scale and its ID was recorded. If the cylinder was not at home, that was recorded.

A total of 79 spot checks occurred. Most were conducted during the first round (N=69). Of the spot checks attempted, 77 weighed the household's study cylinder and 2 reported cylinder was not at home. Households reported having made a recent exchange in three spot checks. Two of these three were confirmed depot exchanges - an exchange was made in the week prior for the same cylinder ID. The distribution of weights is shown in Figure II.

Figure II: Distribution of cylinder weight by size

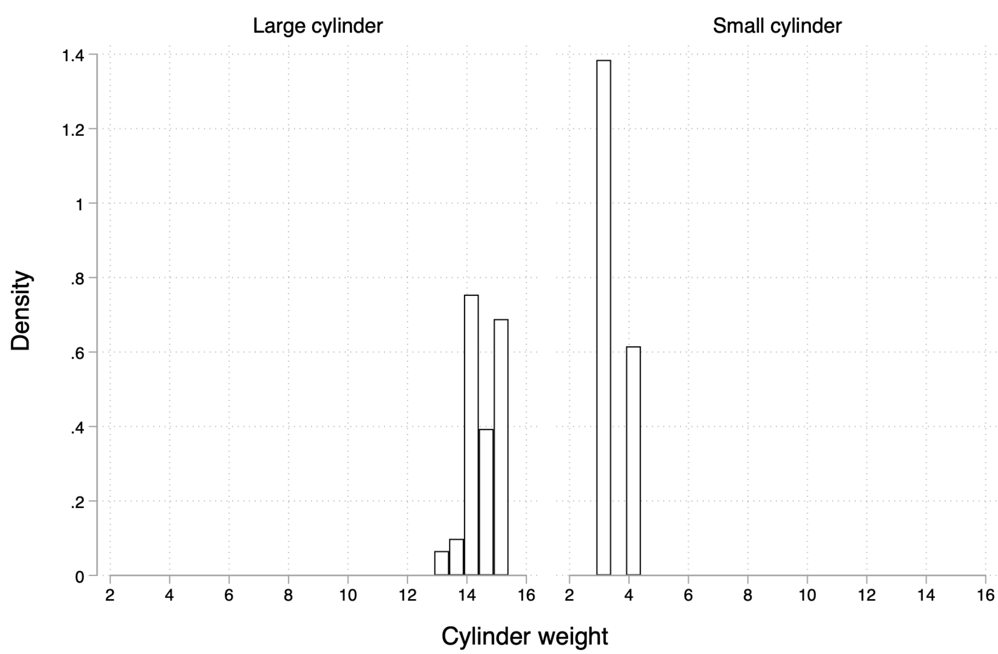

Note: Histogram shows distribution of LPG cylinder weights measured during spot checks, conditional on no recent cylinder exchange. All observations are within 22% of the expected weight.

III.3 Mystery shoppers

To assess adherence to study protocols by the depot managers, KHRC conducted mystery shopper visits during stage 2 of implementation. Mystery shopper visits were designed to test plausible scenarios that could generate errors in the data. Five scenarios were tested:

1. Ineligible study ID
2. Ineligible cylinder ID
3. Successful exchange
4. Past expiration date
5. Incorrect price

During the mystery shopper visit, KHRC staff, who were unknown to the depot managers, would visit the study depots with a fake participant ID card and an empty study cylinder. They would attempt to complete an exchange. The ID card and cylinder utilized in each mystery shopper visit matched the scenario being tested.

A total of 16 mystery shopper visits were conducted. Each depot was visited four times. Exchanges that should not have been allowed (1, 2, 4 and 5) did not proceed. Exchanges that should have been allowed (3) did proceed, except one which may have been an error by the mystery shopper. In a few cases, minor issues such as tablet sync delays were uncovered and rectified.