NOT FOR PUBLICATION

Online Appendix to "Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia"

November 3, 2009

A Formal Model

In this section, we present a simple model of Clorin purchase and use, allowing for sunk-cost effects using Eyster's (2002) framework. At the end of the section, we discuss how our results would generalize to alternative models of sunk-cost psychology.

A.1 Material Payoffs

We will consider household behavior in some number of periods indexed by t. In each period t, each household i must decide whether to purchase Clorin and, if so, whether to use it in their drinking water. In period t, each household i can decide to purchase Clorin for an exogenous, possibly household-specific offer price $p_{it} \in (0, R]$, where $R > 0$ is the retail price of Clorin. If the household decides to buy Clorin, it may receive an exogenous discount, so that it will only have to pay the transaction price $\tilde{p}_{it} \in [0, p_{it}]$. We assume that households do not anticipate the possibility of a discount, or, equivalently, that they consider the probability of a discount after purchase to be negligible.

We experimentally manipulate the prices faced by households in a particular period t'. For simplicity, we will assume that in all previous and subsequent periods, Clorin retails at R, and there are no discounts. That is, we assume that, in any period $t \neq t'$, $p_{it} = \tilde{p}_{it} = R$. In period t', however, prices $p_{it'}$ and $\tilde{p}_{it'}$ are randomly assigned across households in a manner specified in section IIExperimental and Survey Designsection.2 of the paper. We will also assume for simplicity that the utility from using Clorin in a given period is independent of its use in previous or future periods.[1]

We will think of a period as approximately one month, about the time it takes to use up a bottle of Clorin in drinking water for a family of six. We assume that Clorin cannot be stored across periods. Chemically, Clorin is storable, but in practice the vast majority of households in our sample appear to have exhausted the Clorin we sold them within the six weeks of our follow-up period. This suggests that it is reasonable to focus on within-period use. Implicitly, we are

[1] Together, these assumptions mean that when we consider the effect of an increase in the offer price $p_{it'}$, we ignore cross-channel substitution (households buying from our door-to-door marketers who would otherwise have bought in a store) and intertemporal substitution (households buying from us who would otherwise have bought at a later time). The demand parameters we estimate in our experiment are therefore unlikely to generalize (quantitatively) to a permanent, market-wide change in the retail price R.

thereby assuming that if a household buys Clorin in some period and does not use it in drinking water, it is exhausted in some other way. Field interviews we conducted after our study suggest a plausible account is that households use Clorin to undertake household chores such as bleaching sheets, washing vegetables, and cleaning toilets (see section EAdditional Survey Datasubsection.2.5 of the paper).[2]

In each period, each household makes two decisions. First, the household must decide whether to buy Clorin. Then, if the household has purchased Clorin, it must decide whether to use Clorin in its drinking water. We let $b_{it} \in \{0, 1\}$ be an indicator for whether household i buys Clorin in period t. We let $d_{it} \in \{0, 1\}$ denote whether household i puts Clorin in its drinking water in period t, with $d_{it} = 0$ whenever $b_{it} = 0$. We assume these decisions are made sequentially, so that the decision to purchase is fixed when the household chooses whether to use Clorin in its drinking water. Consistent with the discussion above, we can think of $d_{it} = 0$ as performing household chores (other than purifying drinking water), which, as seems reasonable, we assume is possible with or without Clorin.

To apply Eyster's (2002) framework, we need to specify both material payoffs and utility, the latter incorporating the psychological desire to rationalize past choices. We begin by specifying the material payoff function, which we normalize so that the payoff from not buying Clorin is 0.

Households differ in the benefits and costs of using Clorin in their drinking water. We will write the net material payoff (in Kwacha) to household i from buying Clorin and using it in its drinking water in period t as $v_i + \varepsilon_{it} - \tilde{p}_{it}$. Here, v_i captures factors that are constant over time for a given household and are known at the time of purchase, and ε_{it} captures time-varying shocks unknown at the time of purchase.

Formally, we assume that v_i is distributed i.i.d. across households, that ε_{it} is distributed i.i.d across households and time periods, and that v_i and ε_{it} are independent of one another and of offer and transaction prices p_{it} and $\tilde{p}_{it}$ (the last condition being maintained by experimental randomization). Consistent with the interpretation of ε_{it} as an unanticipated shock, we will normalize $E(\varepsilon_{it}) = 0$. To simplify exposition, we assume that both v_i and ε_{it} are distributed according to (possibly different) differentiable CDFs with full support on the real line, and that v_i has finite mean.

The terms v_i and ε_{it} are general enough to accommodate a range of factors that affect a household's net return at the margin from using Clorin in its drinking water. For example, Clorin is used when the household obtains water from its local source, and hence its use requires attention from the female head of household at that time. In some households, the female head may have many other chores to complete when obtaining water, making it hard to put Clorin in at the right time (low v_i), whereas in others the female head may have few other responsibilities coincident with obtaining water (high v_i). In addition, for a given household, variation in the household's day-to-day needs may lead to higher (low ε_{it}) or lower (high ε_{it}) demands on the female head's attention when she obtains water, affecting the incentive to use Clorin.[3]

The term v_i can also capture variation across households in the return to purifying drinking

[2]If our data are misleading, and in fact many households store and defer use of Clorin in drinking water, this force would reduce the power of our experimental tests to detect sunk-cost effects. In the model below, sunk-cost effects arise because households view it as a mistake to have purchased Clorin when they do not use it in their drinking water. In a model with storability, sunk-cost effects of the sort we test for would instead require that households view it as a mistake to have purchased Clorin when they do not use it in their drinking water in the period immediately following purchase.

[3]Many other interpretations are possible. For example, households may differ in their general level of concern about water-borne illness, inducing variation in v_i. After purchase, some households may hear about an especially bad local incident of child diarrhea, leading to a high ε_{it}. Or they may hear that diarrhea episodes have been rare recently, leading to a low ε_{it}.

water by means *other* than Clorin. For example, if boiling water were the main alternative to Clorin, and if boiling water took more time to implement than using Clorin, then households whose female head has a greater opportunity cost of time would have a higher v_i, because at the margin Clorin is displacing a less attractive option for such households.

To complete the specification of payoffs we must specify the payoff to buying Clorin and then not using it in drinking water, using it instead for household cleaning. In that case, the benefit of Clorin is the market value of the standard household cleaners whose use is offset by Clorin. That interpretation suggests a small return from using Clorin for non-drinking-water purposes. For example, data from a 2006 survey of retail prices show that the sodium hypochlorite in a bottle of Clorin can be obtained from Jik, a more concentrated household bleach, for about 300 Kw.[4] Moreover, our field experience suggests that households find products like Jik to be more effective cleaners, suggesting that 300 Kw of Jik may deliver even more than one Clorin bottle's worth of cleaning power. Given the small benefit to Clorin not used in drinking water, we assume for simplicity that the material payoff to buying Clorin and not using it in drinking water is $-\tilde{p}_{it}$, implying that no household would buy Clorin if it were not possible to use it in drinking water. Although that assumption is extreme and unlikely to hold exactly for all households, it seems likely to be a reasonable approximation on average, based both on the calculations above and on our conversations with female heads of household.[5]

To summarize, we can specify an (*ex-post*) material payoff function $u\left(b_{it}, d_{it}; \tilde{p}_{it}, v_i, \varepsilon_{it}\right)$ that relates choices to payoffs as a function of household characteristics and the transaction price:

$$
\begin{aligned}
u\left(0,0;\tilde{p}_{it},v_i,\varepsilon_{it}\right) &= 0 \\
u\left(1,0;\tilde{p}_{it},v_i,\varepsilon_{it}\right) &= -\tilde{p}_{it} \\
u\left(1,1;\tilde{p}_{it},v_i,\varepsilon_{it}\right) &= v_i+\varepsilon_{it}-\tilde{p}_{it}.
\end{aligned}
$$

(Recall that $d_{it}=0$ whenever $b_{it}=0$, so $u\left(0,1;\tilde{p}_{it},v_i,\varepsilon_{it}\right)$ is undefined.)

A.2 Utility and Regret

To allow for psychological effects of prices, we will suppose that households have a taste for consistency, i.e. for taking actions in the present that rationalize their past choices. Following Eyster (2002), we implement this assumption by positing that realized household utility $U\left(b_{it}, d_{it}; \tilde{p}_{it}, v_i, \varepsilon_{it}\right)$ depends both on *ex-post* material payoff and on a regret function:

$$
U\left(b_{it},d_{it};\tilde{p}_{it},v_i,\varepsilon_{it}\right) \equiv u\left(b_{it},d_{it};\tilde{p}_{it},v_i,\varepsilon_{it}\right) - \rho r\left(b_{it},d_{it};\tilde{p}_{it},v_i,\varepsilon_{it}\right) \tag{1}
$$

where $r\left(\right)$ denotes regret and $\rho \geq 0$ indexes the importance of regret in the household's utility function.

We refer the reader to Eyster (2002) for a more careful exposition of the regret function. For-

[4] As of the 2006 retailer survey, a 750ml container of Jik, consisting of 3.5% sodium hypochlorite, retailed for a median of 7,000 Kw. A 250ml bottle of Clorin, consisting of 0.5% sodium hypochlorite, retailed for a median of 800 Kw. The amount of sodium hypochlorite in a container of Jik is therefore equivalent to the amount found in about 21 bottles of Clorin, so 333 Kw of Jik buys as much sodium hypochlorite as one bottle (800 Kw) of Clorin. Note that this calculation ignores the convenience of Clorin's smaller size (and hence lower price per sales unit). Recall, however, that a week's supply of cooking oil for a household costs about 4,800 Kw, suggesting that 7,000 Kw is not a prohibitive outlay.

[5] If the assumption were to fail, then the sunk-cost effects we derive below would operate only over the range of transaction prices $\tilde{p}_{it}$ *above* the cost savings from using Clorin as a household cleaner.

mally, it is defined as follows

$$r\left(b_{it}, d_{it}; \tilde{p}_{it}, v_i, \varepsilon_{it}\right) \equiv u\left(\tilde{b}\left(d_{it}; \tilde{p}_{it}, v_i, \varepsilon_{it}\right), d_{it}; \tilde{p}_{it}, v_i, \varepsilon_{it}\right) - u\left(b_{it}, d_{it}; \tilde{p}_{it}, v_i, \varepsilon_{it}\right) \tag{2}$$

where

$$\tilde{b}\left(d_{it}; \tilde{p}_{it}, v_i, \varepsilon_{it}\right) \equiv \arg\max_b \left\{u\left(b, d_{it}; \tilde{p}_{it}, v_i, \varepsilon_{it}\right)\right\}.$$

Informally, given an action pair, regret is how much better the household's material payoff would be if it could re-choose its first stage action taking its second stage action, and the realization of ε_{it}, as given.[6] Households prefer to avoid regret, i.e. to choose actions that limit the harm done by their past choices given their current ones.

The definition in (2) implies that

$$\begin{aligned} r\left(0, 0; \tilde{p}_{it}, v_i, \varepsilon_{it}\right) &= 0 \\ r\left(1, 0; \tilde{p}_{it}, v_i, \varepsilon_{it}\right) &= \tilde{p}_{it} \\ r\left(1, 1; \tilde{p}_{it}, v_i, \varepsilon_{it}\right) &= 0. \end{aligned} \tag{3}$$

That is, regret is experienced only when the household buys Clorin but does not use it in its drinking water, and is felt in proportion to the amount spent on Clorin.

In the model, then, if a household knew for sure that it would not use Clorin to purify its drinking water, it would not buy Clorin. Households buy Clorin because they may use it in drinking water, but if circumstances are such that Clorin is not used in drinking water, the household regrets its purchase, and the regret experienced is greater the more the household paid for Clorin. Households may therefore be willing to use Clorin in drinking water to avoid regret, i.e. to rationalize the past decision to buy.

A.3 Choice and Identification

We can now specify how the household chooses whether to purchase Clorin and, if so, whether to use it in drinking water. Beginning with the use decision, if the household has purchased Clorin, the use decision is given by $d^*_{it}\left(\tilde{p}_{it}, v_i, \varepsilon_{it}\right)$, where

$$d^*_{it}\left(\tilde{p}_{it}, v_i, \varepsilon_{it}\right) \equiv \arg\max_d U\left(1, d; \tilde{p}_{it}, v_i, \varepsilon_{it}\right).$$

This, in turn, implies that

$$\begin{aligned} d^*_{it}\left(\tilde{p}_{it}, v_i, \varepsilon_{it}\right) &= 1 \\ &\Longleftrightarrow \\ v_i + \varepsilon_{it} &\geq -\rho\tilde{p}_{it}. \end{aligned}$$

That is, the household will use Clorin in its drinking water if the net benefit of doing so exceeds the regret associated with not doing so. (Note that we have assumed for simplicity that ties are broken in favor of use.)

When households are deciding whether to purchase Clorin in the first place, they do not an-

[6] Note that, as we have defined it, regret depends on the realized transaction price $\tilde{p}_{it}$, which is technically not observed by non-purchasing households. However, because non-purchasing households do not make a use decision, and because households are unaware of the possibility of a discount at the time of purchase, this is purely a notational simplification, with no implications for behavior.

ticipate the discount, and do not know the realization of the time-varying shock ε_{it}. Therefore we write the expected-utility-maximizing purchase decision as $b_{it}^{*}(p_{it}, v_i)$, with

$$b_{it}^{*}(p_{it}, v_i) \equiv \arg\max_{b} EU\left(b, d_{it}^{*}(p_{it}, v_i, \varepsilon_{it}); p_{it}, v_i, \varepsilon_{it}\right)$$

where the expectation is taken over the distribution of the shocks ε_{it}. It follows that

$$\begin{array}{rcl} b_{it}^{*}(p_{it}, v_i) & = & 1 \\ & \Longleftrightarrow & \\ E\left[\max\left(v_i + \varepsilon_{it}, -\rho p_{it}\right) \mid v_i, p_{it}\right] & \geq & p_{it} \end{array}$$

where we adopt the arbitrary rule that ties are broken in favor of purchasing.

Note that the household conditions its decision on the (known) household-specific benefit parameter v_i and on the offer price p_{it}. In particular, there exists a real-valued cutoff $v^{*}(p_{it})$, strictly increasing in p_{it}, such that the household purchases if and only if $v_i \geq v^{*}(p_{it})$. The proof of this result is straightforward; economically, it follows from the fact that a higher v_i increases the anticipated benefit from buying Clorin and a higher p_i increases the anticipated cost (both financially and psychologically).

Empirically, we do not observe the choice functions $b_{it}^{*}(p_{it}, v_i)$ and $d_{it}^{*}(\tilde{p}_{it}, v_i, \varepsilon_{it})$, but rather the empirical frequencies of different choices as a function of the prices p_{it} and $\tilde{p}_{it}$. To develop the model's empirical predictions, then, it will be helpful to write out these probabilities:

$$\begin{array}{rcl} \Pr\left(b_{it}^{*}(p_{it}, v_i) = 1 \mid p_{it}, \tilde{p}_{it}\right) & = & \Pr\left(v_i \geq v^{*}(p_{it})\right) \\ \Pr\left(d_{it}^{*}(\tilde{p}_{it}, v_i, \varepsilon_{it}) = 1 \mid p_{it}, \tilde{p}_{it}\right) & = & \Pr\left(v_i + \varepsilon_{it} \geq -\rho\tilde{p}_{it} \mid v_i \geq v^{*}(p_{it})\right) \end{array}$$

where the second probability conditions on the decision to purchase Clorin.

In a non-experimental period $t \neq t'$, $p_{it} = \tilde{p}_{it} = R$, so there is no price variation and hence no comparative statics to test. (The model does imply restrictions on the relationship between prices in the experimental period and use in a non-experimental period, which we explore in online appendix section B.)

In the experimental period t', by contrast, the above equations have two important empirical implications about the relationship between prices and use.[7]

- The first, which we call the *screening effect*, is that, conditional on purchase, the probability of use is higher the greater is $p_{it'}$. This is because a higher $p_{it'}$ imposes a stricter (higher) cutoff $v^{*}(p_{it'})$ on anticipated benefits for buyers.

- The second, which we call the *sunk-cost effect*, is that, conditional on purchase, the probability of use is increasing in the transaction price $\tilde{p}_{it'}$ whenever $\rho > 0$. This is because a greater transaction price implies a greater desire to rationalize the purchase decision (or, equivalently, greater regret from not using Clorin in drinking water).

Observe that these two effects cannot be distinguished if transaction prices cannot vary independently of offer prices. If $p_{it'} = \tilde{p}_{it'}$ for all i in some period t, then a finding that higher prices causes more drinking water use conditional on purchase would be consistent with the presence of

[7] A third, more obvious, implication is that the greater is the price of Clorin, the fewer households will purchase it. Note that this result comes both from the traditional substitution effect, and from the fact that a higher price implies greater anticipated regret in the case in which Clorin is not used in drinking water.

either the sunk-cost effect ($\rho > 0$) or the screening effect (heterogeneity in v_i), or both.[8] Hence, the two-stage pricing design solves an identification problem that would be present in data with only a single price, even if that price were suitably exogenous.

It is worth noting that, although Eyster's model is the most fully articulated single-agent theory of sunk-cost effects of which we are aware, reasonable alternatives exist with possibly different empirical implications. For example, if sunk-cost effects were driven by a desire to justify the *ex ante* wisdom of one's decision, rather than *ex post* wisdom of one's decision as in Eyster's framework, it is possible that the offer price, rather than the transaction price, would influence use behavior.[9] Such effects would confound our tests.[10] On the other hand, Thaler's (1980) prospect-theoretic justification for sunk-cost effects hinges on a desire to avoid a feeling of loss experienced when one does not realize consumption gains from a past purchase. In such a model transaction prices likely would mediate the effect, suggesting that our tests may be valid under mechanisms other than the one we model explicitly.

B Intertemporal Implications of the Model

In the model, realized transaction prices $\tilde{p}_{it'}$ in period t' will not affect purchase or use behavior in any period $t \neq t'$: sunk-cost effects are localized to the period in which the transaction occurs. The model does, however, have implications for the relationship between offer prices $p_{it'}$ conditional on purchase in period t' and purchase and use behavior in periods $t \neq t'$. In this appendix we briefly summarize those implications. We also provide some empirical tests using data from our baseline survey, reported in online appendix table 1. The baseline period can be thought of as a reasonable approximation to period $t' - 1$, in that any household owning Clorin as of the baseline survey would have purchased it from a retailer (or health clinic). It is worth noting, however, that we did not design our study to test these intertemporal implications, so the evidence reported below should be thought of as preliminary.

Begin by considering the relationship between purchase behavior in the experiment and purchase behavior in non-experimental periods. The model predicts that households purchasing in period t' are more likely to buy at other periods $t \neq t'$, the greater is the offer price $p_{it'}$ in period t'. To see this, observe that, for households that purchase in the experimental period t', we can write the

[8] Note that the psychology of regret is also relevant for the magnitude of the screening effect, as the parameter ρ partially determines the degree of price sensitivity in the purchase decision.

[9] Sunk-cost effects could also operate through a desire to justify the purchase to another member of the household, rather than to oneself (Prendergast and Stole, 1996). In that case, whether the offer or transaction price would influence use behavior through the sunk-cost mechanism would depend on the informational conditions in the household. Because only the offer price is known to the buyer at the time of the purchase decision, a fully informed household member would judge the intelligence of the decision based on the offer price. However, because only the transaction price is actually implemented, it may be more observable to other members of the household than the offer price. In that case, higher transaction prices would lead to greater use due to a desire to justify the purchase decision to other household members, so our tests for sunk-cost effects would be valid.

[10] Another possible confound arises if use depends not on the transaction price itself but on the discount, i.e. the difference between the offer and transaction prices. We find no evidence in our data of a relationship between use and the relative size of the discount—that is, the difference between offer and transaction prices, divided by the offer price. If instead psychological effects operate through the absolute (as opposed to relative) size of the discount (offer price minus transaction price), the resulting model is collinear with those we estimate, and is therefore not identified. If the effect of a greater discount is to increase use, then our estimates will tend to overstate the screening effect and to understate the sunk-cost effect. If greater discounts tend to decrease use, our estimates will understate the screening effect and overstate the sunk-cost effect.

following for any non-experimental period $t \neq t'$:

$$\Pr\left(b^*_{it}\left(R, v_i\right) = 1 \mid b^*_{it'}\left(p_{it'}, v_i\right) = 1, R, p_{it'}\right) = \Pr\left(v_i \geq v^*\left(R\right) \mid v_i \geq v^*\left(p_{it'}\right)\right)$$

with $v^*\left(R\right) \geq v^*\left(p_{it'}\right)$. Because $v^*\left(\right)$ is increasing in $p_{it'}$, it follows directly that, conditional on purchase in period t', the probability of purchase in a period $t \neq t'$ is greater the greater is $p_{it'}$. Intuitively, the higher is a household's demonstrated willingness to pay in period t', the greater is its likelihood of purchasing at the retail price in other periods. We present a test of this implication in column (1) of online appendix table 1. We measure purchase at baseline with an indicator for whether our surveyor inventory indicates that the household had a non-empty Clorin bottle at the time of the baseline survey. We find a marginally statistically significant positive relationship between offer price and baseline purchase (conditional on purchase at the marketing stage), consistent with the model.

Consider next the relationship between purchase behavior in the experiment and use behavior in non-experimental periods. Perhaps somewhat surprisingly, the model predicts that, conditional on purchase in the experiment and in a non-experimental period, the experimental offer price will be unrelated to use in non-experimental periods. Formally, consider the expression for use conditional on purchase in both the experimental period t' and some non-experimental period $t \neq t'$:

$$\Pr\left(d^*_{it}\left(R, v_i, \varepsilon_{it}\right) = 1 \mid b^*_{it'}\left(p_{it'}, v_i\right) = 1, R, p_{it'}\right) = \Pr\left(v_i + \varepsilon_{it} \geq -\rho R \mid v_i \geq v^*\left(R\right) \wedge v_i \geq v^*\left(p_{it'}\right)\right)$$

That is, the set of households purchasing at both time t and time t' are those for whom both $v_i \geq v^*\left(R\right)$ and $v_i \geq v^*\left(p_{it'}\right)$. However, because $R \geq p_{it'}$, $v^*\left(R\right) \geq v^*\left(p_{it'}\right)$, so that the two conditions collapse to a single one:

$$\begin{aligned}
\Pr\left(d^*_{it}\left(R, v_i, \varepsilon_{it}\right) = 1 \mid b^*_{it'}\left(p_{it'}, v_i\right) = 1, R, p_{it'}\right) &= \Pr\left(v_i + \varepsilon_{it} \geq -\rho R \mid v_i \geq v^*\left(R\right)\right) \\
&= \Pr\left(d^*_{it}\left(R, v_i, \varepsilon_{it}\right) = 1 \mid R\right)
\end{aligned}$$

Therefore we predict no relationship between offer price at time t' and use behavior at time $t \neq t'$. Intuitively, this result comes about because anyone willing to pay R should also be willing to pay $p_{it'} \leq R$, so that offer prices in the experiment convey no new information about the composition of households purchasing in non-experimental periods.[11] Columns (2) and (3) of online appendix table 1 present tests of this implication. Among those who had Clorin in their households as of the baseline (our proxy for purchase at time $t' - 1$) and who purchased Clorin in the marketing stage (time t'), there is no relationship between offer price and self-reported or measured use at baseline, consistent with the model. The coefficients are small, negative, and statistically insignificant.

Finally, consider the reverse question of how purchase behavior in non-experimental periods affects our predictions for use in the experimental period. Following a logic parallel to the preceding argument, it is straightforward to show that, conditional on purchase in a non-experimental period, there should be no relationship between offer price in the experimental period and use in the experimental period, holding constant the transaction price. That is, the screening effect should be absent for households that purchase in non-experimental periods. In equations, this is because:

$$\begin{aligned}
\Pr\left(d^*_{it'}\left(\tilde{p}_{it'}, v_i, \varepsilon_{it}\right) = 1 \mid p_{it'}, \tilde{p}_{it'}, b^*_{it}\left(R, v_i\right) = 1\right) &= \Pr\left(v_i + \varepsilon_{it} \geq -\rho \tilde{p}_{it'} \mid v_i \geq v^*\left(R\right) \wedge v_i \geq v^*\left(p_{it'}\right)\right) \\
&= \Pr\left(v_i + \varepsilon_{it} \geq -\rho \tilde{p}_{it'} \mid v_i \geq v^*\left(R\right)\right).
\end{aligned}$$

[11] Our model assumes a constant willingness-to-pay over time. A model with time-varying shocks to v_i that are known in advance of purchase to the household could yield a positive relationship between use in non-experimental periods and experimental offer price. However, given that the non-experimental periods we study are close in time to the experimental period, constant willingness-to-pay over time may be a reasonable approximation.

By contrast, households that did not purchase in non-experimental periods will still display the screening effect, although with a different magnitude than the overall population:

$$\Pr\left(d^*_{it'}\left(\tilde{p}_{it'}, v_i, \varepsilon_{it}\right) = 1 \mid p_{it'}, \tilde{p}_{it'}, b^*_{it}\left(R, v_i\right) = 0\right) = \Pr\left(v_i + \varepsilon_{it} \geq -\rho \tilde{p}_{it'} \mid v^*\left(p_{it'}\right) < v_i \leq v^*\left(R\right)\right).$$

We test these implications in specifications reported in columns (4) through (7) of online appendix table 1. Consistent with the model, among those who purchased at the marketing stage and had Clorin at baseline, there is a statistically insignificant relationship between offer price and use at follow-up. (We note, however, that due to small sample size our confidence intervals cannot rule out nontrivial effects.) By contrast, among those who purchased at the marketing stage but did not have Clorin in the home as of the baseline, there is a positive and statistically significant relationship between offer price and use.

C Implications of Survey Attrition

In this section we address the correlates of survey attrition and the implications of attrition for our main findings.

Online appendix table 2 analyzes the determinants of attrition. Among those who bought Clorin in the marketing stage, the rate of contact at follow-up is 89 percent. Column (1) of online appendix table 2 shows that contact is somewhat related to observable household characteristics. For example, households with an older female head, households with a married female head, and households with more members are more likely to be reached at follow-up. An F-test of the joint significance of our full vector of observables cannot quite reject the null that the observables have no impact on the likelihood of contact ($p = 0.102$). Columns (2) and (3) of online appendix table 2 show that we cannot reject the null that the treatment variables (offer and transaction price) have no effect on the likelihood of contact.

Lee (2009) shows that if treatment does not affect attrition, standard estimates of treatment effects are valid, but are "local" to the population who are successfully contacted. Online appendix table 2 shows that we cannot rule out statistically that there is no effect of treatment on attrition. However, the power of these tests is finite, and it is therefore worthwhile asking whether effects of treatment on attrition on the order of the point estimates in columns (2) and (3) of online appendix table 2 are sufficient to affect the interpretation of our findings. To do this, we take the point estimates in columns (2) and (3) of online appendix table 2, and assume that the households removed (or added) from the follow-up sample by the treatment variable were either all users or all non-users of Clorin. The resulting upper and lower bounds are extremely tight:

Specification	Screening effect		Sunk-cost effect	
Use measure	Self-reported	Measured	Self-reported	Measured
Main tables	0.0373	0.0321	0.0097	-0.0071
Lower bound	0.0354	0.0302	0.0092	-0.0075
Upper bound	0.0399	0.0346	0.0112	-0.0055

Therefore, attrition due to treatment is unlikely to be a concern in our setting, and we can reliably conduct inference local to the population of households reached in the follow-up survey.

A steeper challenge is to conduct inference on the entire population reached at marketing, including those who are never observed. Below we produce upper and lower bounds for these effects computed following Horowitz and Manski (2000). In particular, we compute residuals of the treatment variables after partialling out all controls. For an upper bound we assign those attriters with positive residuals to be users and those with negative residuals to be non-users. We reverse the assignment to compute lower bounds. The bounds are wide:

Specification	Screening effect		Sunk-cost effect	
Use measure	Self-reported	Measured	Self-reported	Measured
Main tables	0.0373	0.0321	0.0097	-0.0071
Lower bound	0.0013	-0.0037	-0.0238	-0.0400
Upper bound	0.0650	0.0605	0.0423	0.0263

These bounds admit slightly negative values for the screening effect and reasonably large positive values for sunk-cost effects. It is therefore valuable to explore whether attrited households might plausibly be close to the worst-case scenarios defined by the above bounds. By definition we cannot directly analyze the use behavior of attrited households. However we can use observed variation within the sample to attempt some inference about those we do not observe. We approach this in two ways.

First, we estimate a logit analogue of column (1) of online appendix table 2. We then re-estimate our main specifications, weighting by the inverse of the predicted probability from the logit model. If all heterogeneity in treatment effects were related to observables, these estimates would be valid for the entire population of households reached at the marketing stage. Weighted and unweighted estimates are very similar:

Specification	Screening effect		Sunk-cost effect	
Use measure	Self-reported	Measured	Self-reported	Measured
Main tables	0.0373	0.0321	0.0097	-0.0071
Weighted	0.0384	0.0311	0.0068	-0.0079

If anything, sunk-cost effects are weaker in the weighted specifications, and there is no systematic direction of change for our estimates of screening effects.

Second, we examine heterogeneity in treatment effects related to how many attempts we had to make in order to reach the household at follow-up. If attrited households are more similar to those that were reached after the first attempt than to those we reached on the first attempt, comparing these two groups may help to narrow the range of plausible scenarios within the bounds we provide above. The results are below:

Specification	Screening effect		Sunk-cost effect	
Use measure	Self-reported	Measured	Self-reported	Measured
Main tables	0.0373	0.0321	0.0097	-0.0071
First attempt	0.0343	0.0284	-0.0032	-0.0116
Second or third attempt	0.0402	0.0366	0.0406	0.0129

The differences between the two groups are never statistically significant. If anything we find that screening effects are larger among the households that required a second or third attempt before contact, indicating that difficulty of contact is not associated with weaker screening effects. Directionally, difficulty of contact is associated with stronger sunk-cost effects, though the difference is economically much larger for self-reported than for measured use.

We take three lessons from the analysis in this section. First, there is no evidence of an effect of treatment on attrition, and even at the point estimates for those effects there is little reason to be concerned about the validity of our main estimates for the population of households reached in the follow-up. Second, worst-case bounds for our main specifications are wide, indicating that we cannot rule out reversals of our conclusions among the overall population if attrited households differ radically from those we observe. Third, an analysis of heterogeneity of treatment effects with respect to both observable characteristics and the difficulty of re-contact do not indicate such radical differences with respect to screening effects. If anything the analysis suggests we might expect somewhat stronger screening effects on the overall population. However, point estimates do indicate that sunk-cost effects may be larger for households that were more difficult to reach, indicating that some caution is warranted in interpreting the power of our sunk-cost tests.

D Additional Implications of Sunk-cost Psychology

In this section we empirically test empirical implications of extensions to the sunk-cost model that we formalize in section A.

In the model, we assume that the regret parameter ρ is identical across households. If there is heterogeneity in this parameter, then, all else equal, households with higher values of ρ should exhibit a lower propensity to purchase Clorin, because these households anticipate the possibility of regret if they purchase and do not use Clorin in drinking water. We proxy for the regret parameter ρ using households' sunk-cost status defined by their responses to hypothetical choices. We assume that sunk-cost status is orthogonal to other drivers of purchase behavior, conditional on our set of baseline controls. We can then test for an effect of ρ on purchase propensity by regressing a dummy for whether the household purchased Clorin on a dummy for sunk-cost status and our set of baseline controls. Column (1A) of online appendix table 3 shows that we find, consistent with this prediction, that sunk-cost households have a lower propensity to purchase Clorin than non-sunk-cost households, although the coefficient is statistically insignificant. As a further caveat, we note also that sunk-cost status is correlated with the baseline controls in the model. In a regression of sunk-cost status on baseline controls, an F-test of the null hypothesis that baseline controls jointly do not affect sunk-cost status, we can reject the null at conventional significance levels ($p = 0.0062$). This suggests a note of caution in interpreting the coefficient in column (1A), as sunk-cost status may be correlated with unobservable determinants of purchase propensity.

A related hypothesis is that sunk-cost households are more price-sensitive than non-sunk-cost households, as higher prices increase the risk of utility losses due to regret. Columns (2A) and (3A) show that price-sensitivity is fairly similar between sunk-cost and non-sunk-cost groups. The difference in the effect of a 100 Kw increase in price on the propensity to purchase Clorin is economically similar and statistically indistinguishable between the two groups.

In the model, feelings of regret arise in equilibrium because households are subject to ex-post shocks to the net value of using Clorin in drinking water. Although some shocks (hassle, forgetting, disease episodes, etc.) could in principle affect any household, other shocks, such as learning about the taste or other properties of Clorin itself, are more likely to affect households with less prior information about Clorin. Because our sample largely consists of households with past experience with Clorin, our main specifications might therefore miss important sunk-cost effects for those without such experience. In panel B of online appendix table 3 we ask whether the magnitude of sunk-cost effects is related to whether a household has ever used Clorin in the past. If anything, we find that sunk-cost effects are statistically significantly *lower* for households without prior experience with Clorin. Such households are at least marginally significantly less likely to use Clorin the higher is the transaction price, with confidence intervals that can rule out sizable positive effects of transaction price on use. Among those who have used Clorin in the past, we continue to find small and statistically insignificant effects of transaction price on use. Because these hypothesis tests were not part of our original design, future work will be needed to determine whether the surprising negative effects we find for never-users are robust and, if so, what mechanism underlies them. At the very least, these specifications suggest that our main specifications do not conceal important positive effects of transaction prices on use for inexperienced households.

Online Appendix Table 1 *Intertemporal implications of the model*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent variable	Have Clorin in household? (baseline; inventory)	Water currently treated with Clorin? (baseline; self-reported)	Drinking water contains free chlorine? (baseline; measured)	Water currently treated with Clorin? (follow-up; self-reported)		Drinking water contains free chlorine? (follow-up; measured)	
Sample: Clorin in home at baseline?	All	Yes	Yes	Yes	No	Yes	No
Offer price	0.0182	-0.0067	-0.0093	-0.0289	0.0512	-0.0046	0.0386
(100 Kw)	(0.0104)	(0.0245)	(0.0265)	(0.0288)	(0.0169)	(0.0294)	(0.0171)
No. of obs.	614	133	133	117	429	117	425

Notes: Standard errors in parentheses. See online appendix section B for details. Sample consists of households who purchased Clorin at marketing stage.

Online Appendix Table 2 *Determinants of sample attrition*

	(1)	(2)	(3)
Offer price (100 Kw)		-0.0042	
		(0.0090)	
Transaction price (100 Kw)			0.0063
			(0.0080)
Water currently treated with Clorin?	-0.0097	-0.0165	-0.0092
(baseline; self-reported)	(0.0336)	(0.0339)	(0.0338)
Drinking water contains free chlorine?	-0.0034	-0.0030	-0.0057
(baseline; measured)	(0.0271)	(0.0274)	(0.0274)
Use of soap before handling food	0.0028	0.0044	0.0004
(index)	(0.0481)	(0.0484)	(0.0487)
Use of soap after using toilet	-0.0194	-0.0249	-0.0231
(index)	(0.0491)	(0.0492)	(0.0496)
Attitude toward water purification	0.0882	0.0878	0.0838
(index)	(0.0686)	(0.0689)	(0.0690)
Age in years	0.0031	0.0033	0.0034
	(0.0015)	(0.0015)	(0.0015)
Ever attended school?	-0.0057	-0.0129	-0.0035
	(0.0552)	(0.0554)	(0.0556)
Years of completed schooling	0.0077	0.0083	0.0078
	(0.0059)	(0.0059)	(0.0059)
Currently married?	0.0775	0.0751	0.0811
	(0.0362)	(0.0364)	(0.0364)
Currently pregnant?	-0.0400	-0.0439	-0.0410
	(0.0428)	(0.0429)	(0.0430)
Ever given birth to any children?	-0.0146	-0.0108	-0.0242
	(0.0560)	(0.0562)	(0.0563)
No. of children in household under age 5	-0.0099	-0.0120	-0.0105
	(0.0177)	(0.0177)	(0.0178)
No. of people in household	0.0119	0.0126	0.0121
	(0.0057)	(0.0057)	(0.0058)
Share of durables owned	0.0040	-0.0037	0.0020
	(0.0866)	(0.0867)	(0.0870)
Locality fixed effects?	YES	YES	YES
Fixed effects for offer price?	NO	NO	YES
Fixed effects for transaction price?	NO	YES	NO
F-test that control coefficients are 0	1.45	1.54	1.50
p-value of F-test	0.1015	0.0712	0.0837
Number of observations	605	605	605

Notes: Standard errors in parentheses. All variables measured as of baseline survey. Offer price and transaction price variables excluded from F-tests. Sample includes all households who purchased Clorin at marketing stage and for whom all observables are nonmissing.

Online Appendix Table 3 *Additional implications of sunk-cost psychology*

Panel A: Take-up by sunk-cost status

	(1A)	(2A)	(3A)
Dependent variable	Purchased Clorin?		
Sample	All	Sunk-cost household?	
		Yes	No
Sunk-cost household?	-0.0606		
	(0.0396)		
Offer price		-0.0661	-0.0673
(100 Kw)		(0.0214)	(0.0113)
Difference	—	*0.0012*	
(sunk-cost vs. non-sunk-cost)		*(0.0242)*	
Baseline controls?	YES	YES	YES
Sample mean of dep. var.	0.6130	0.5567	0.6300
No. of obs.	876	203	673

Panel B: Sunk-cost effects by ever-use status

	(1B)	(2B)	(3B)	(4B)
Dependent variable	Water currently treated with Clorin? (follow-up; self-reported)		Drinking water contains free chlorine? (follow-up; measured)	
Sample	Ever used Clorin?		Ever used Clorin?	
	Yes	No	Yes	No
Transaction price	0.0204	-0.0648	0.0018	-0.0889
(100 Kw)	(0.0147)	(0.0382)	(0.0146)	(0.0378)
Difference	*0.0853*		*0.0907*	
(ever used vs. never used)	*(0.0410)*		*(0.0405)*	
Offer price fixed effects?	YES	YES	YES	YES
Baseline controls?	YES	YES	YES	YES
Sample mean of dep. var.	0.5308	0.4388	0.5459	0.4948
No. of obs.	439	98	436	97

Notes: Standard errors in parentheses. See online appendix section D for details.

References

[1] Erik Eyster. Rationalizing the past: A taste for consistency. *Nuffield College Mimeograph*, November 2002.

[2] Joel L. Horowitz and Charles F. Manski. Nonparametric analysis of randomized experiments with missing covariate and outcome data. *Journal of the American Statistical Association*, 95(449):77–84, 2000.

[3] David S. Lee. Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *Review of Economic Studies*, 76(3):1071–1102, 2009.

[4] Canice Prendergast and Lars Stole. Impetuous youngsters and jaded old-timers: Acquiring a reputation for learning. *Journal of Political Economy*, 104(6):1105–1134, December 1996.

[5] Richard Thaler. Toward a positive theory of consumer choice. *Journal of Economic Behavior and Organization*, 1(1):39–60, March 1980.