SEWA Wage Insurance Experiment: Pre-Analysis Plan

Anup Malani[*], Aprajit Mahajan[†], Grant Miller[‡], Morgen Miller[§], Pietro Tebaldi[¶] and Alex Torgovitsky[‖]

July 15, 2022

Abstract

We present the motivation, design and analysis plan for the SEWA Wage Insurance Experiment (SWIE). The SWIE is a randomized controlled trial examining the demand for and impacts of providing so-called "hospi-cash" insurance that provides a per-diem indemnity payment while a beneficiary is hospitalized. We brand this a wage insurance product because SEWA intends it to cover lost wages due to hospitalization, even though it can cover other costs and consequences of hospitalization as money is fungible. The SWIE offers roughly 200 villages access to two hospi-cash policies. The SWIE randomizes the prices to which villages are offered access to these two plans. The SWIE has two methodological innovations. First, it identifies sharp non-parametric bounds on demand using the approach in Tebaldi et al. (2019). Second, we develop a new approach to experimental design in order to choose prices so as to minimize the width of these non-parametric bounds. The SWIE also contributes to our understanding of the value of indemnity health insurance policies. The SWIE examines impacts of prices on hospi-cash insurance and health care insurance uptake and the impact of prices and uptake of hospi-cash on hospi-cash insurance claims and health care utilization. The evidence on uptake also estimates demand for indemnity insurance. The SWIE also estimates adverse selection into this wage insurance. Finally, the experiment also examine whether the hospi-cash insurance product improves risk-sharing within treatment villages.

[*]University of Chicago, amalani@uchicago.edu.
[†]University of California-Berkeley, aprajit@gmail.com.
[‡]Stanford University, ngmiller@stanford.edu.
[§]University of Chicago, mmmiller@uchicago.edu.
[¶]Columbia University, p.tebaldi@columbia.edu.
[‖]University of Chicago, atorgovitsky@gmail.com.

Contents

1 Introduction

Context. Each year, some 150 million people worldwide face financial catastrophe due to spending on health. According to a 2010 study, more than one third of them live in India (Shahrawat and Rao, 2012). The number of Indians falling below the poverty line (BPL) due to health spending may run as high as 63 million people: almost 7% of the nation's population (Berman et al., 2010).

Medical expenses are not the only financial cost people face due to treatment for sickness. One of these is lost income due to treatment.[1] In the US, it has been estimated that lost income associated with (treated) sickness is estimated to be at least triple the cost of medical care (Dobkin et al., 2018). Using data from NSS (71st round, 2014), we estimate that, in India, hospitalization is associated with INR 800 less income (see Table 1).[2] In theory, the loss of income following illness may be attributable to an inability to work due to continued sickness and/or to the need to miss work for medical treatment. In this project we study the latter component.

Loss of income due to the inability to work during treatment may impact the value of health insurance and the amount of health care utilization. The full cost of medical care has at least four components: the medical bill from providers, the cost of transportation to the provider, the cost of financing medical bills and transportation, and the loss of income during treatment and recovery. Health insurance typically only assists with the first three items at the time of medical care. Thus, like deductibles and co-pays, the lost wages during treatment are a form of co-insurance that increases the price of care at the margin, and thereby deter. These type of losses can reduce the value of health insurance if demand for insurance is decreasing in level of co-insurance. They can also reduce care at the margin relative to having no wage loss during treatment. For example, according to the Self-Employed Women's Association (SEWA), an NGO in India, rural women who are members of that organization report foregoing hospital care because they do not want to forego their daily income.

Despite the importance of lost wages during medical treatment (which we will call the wage cost of treatment), health insurance does not in general cover these wage costs. In the US, there are plans, such as AFLAC, that provide indemnity payments in the case of medical care rather than payment for medical care. Likewise, in India, there are so-called "hospi-cash" policies that provide a fixed indemnity payment for each day a beneficiary is hospitalized, up to some annual maximum.

Most relevant for our purposes, SEWA, a cooperative that, among other things, sells its self-employed women members financial products to support their financial well-being, has offered its

[1] There are other costs as well. For example, there are costs associated with transportation to the hospital. However, many insurance policies in India cover travel expenses to a point. For example, Rastriya Swasthya Bima Yojana, until 2018 the main public health insurance plan in India that covered hospital care, provided INR 100 reimbursement for travel expenses per visit. Therefore, we do not discuss travel expenses here.

[2] Health care costs are a lower percentage of income in India than in the U.S. India spends roughly 3.89% of its GDP on health (World Health Organization, 2019).

Table 1: *Income lost associated with hospitalization.*

	Amount of lost income		
	(1)	(2)	(3)
Hospitalized	811.4***	802.5***	800.9***
	(9.60)	(9.52)	(9.53)
Constant	174.4***	59.88***	72.52***
	(14.88)	(6.82)	(3.66)
N	37282	37277	37277
Severity controls	no	yes	yes
Asset controls	no	no	yes

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: Results from a linear regression on a sample who reported that they had a sickness in the last 15 days. The dependent variable is income lost due to sickness. We report the coefficient on an indicator for whether the person was hospitalized. Different columns add controls for the severity of sickness (e.g., number of days confined to a bad (not length of stay)) and measures of total asset (e.g., number of rooms in home). We estimate robust standard errors.

members a hospi-cash policy, which they call Saral Suraksha Yojana (SSY), since 2012.[3] Specifically, SSY provides any adult member of a household INR 200 per day they or a covered child is hospitalized, with a cap of 15 covered days per year. Prior to our study, SEWA sold two variants of this policy: a member could purchase coverage for either 2 parents and 2 children (at a premium of INR 375) or for 2 parents and 4 children (at a premium for INR 475). There is a 30 day waiting period for coverage and pre-existing conditions are excluded for 1 year.[4] Hospi-cash policies such as SSY can buffer against wage loss during treatment.[5]

Our study. In this study, called the SEWA Wage Insurance Experiment (SWIE), we estimate

[3] SEWA is not the only organization that offers a hospi-cash product in the market. Some insurance companies offered hospi-cash to high income households. Moreover, some microfinance companies offered such a product in conjunction with loans.

[4] SSY also includes life insurance coverage for deaths associated with a personal accident, with coverage of INR 100,000 for the (female) SEWA member and INR 50,000 for her husband. Through its financial services affiliate, VimoSEWA, SEWA also offers its members a hospital insurance product, a life insurance product and a credit insurance product. The hospital insurance product covers a household's in-patient treatments. The health insurance product has no deductibles or co-pays and an annual cap of INR 15,000 - 50,000, depending on the policy; the premium ranges from INR 1,130-1,350 depending on the annual cap and whether the policy reimburses the beneficiary for her expenses or pays the providers directly (i.e., is "cashless").

[5] Because the indemnity payment is cash and cash is fungible, hospi-cash policies can also cover medical expenses not covered by standard health insurance plans. In India, this includes medications. That said, SEWA's motivation in added hospi-cash coverage was to cover lost wages. Specifically, when surveying its members on obstacles to getting hospital care, members cited not only the cost of medical care but the lost wages during a hospital visit.

the demand for a new version of the SSY hospi-cash product[6] with different coverage limits.[7] We will also estimate various impacts of offering hospi-cash insurance.

Our study has three stages. First, we conduct pilot studies to help SEWA develop two hospi-cash policies and determine experimental prices for these policies. These policies replace the two pre-existing hospi-cash policies that SEWA was selling in the sample villages.[8] Second, we randomize SEWA members, at the village level, to 4 different price combinations for these 2 new hospi-cash policies. We use this assignment mechanism to produce pointwise causal estimates of the fraction of SEWA members that purchase each of the 2 products at each of the 4 possible price combinations. Additionally, following Tebaldi et al. (2019), we estimate sharp nonparametric bounds on the demand for these new products at 10 unobserved prices using uptake at the 4 observed price combinations. Importantly, the 4 price combinations used in the second stage are chosen optimally to obtain narrow bounds around the demand for hospi-cash products over the 10 unobserved prices. Third, we generate complier average treatment effect (CATE) or treatment on treated (TOT) estimates of the causal impact of hospi-cash uptake on uptake of health insurance and utilization of hospital care using the price offers as instrumental variables. This design also estimates adverse selection into hospi-cash insurance, following the approach outlined in Einav et al. (2010a). Finally, we estimate CATE of hospi-cash insurance on employment and household financial planning, viz. days of work, daily wages and savings, and asset portfolios.

Contributions. This experiment makes three major contributions. First, from a policy perspective, it helps estimate the demand for and value of a hospi-cash policy. It will help SEWA decide whether to continue offering hospi-cash policies and what sorts of hospi-cash policies VimoSEWA should offer to its 1.4 million members in Gujarat. This experiment also has national relevance. Given that the details of the current national public health insurance plan, Pradhan Mantri Jan Arogya Yojana (PMJAY), are still being worked out, the experiment can help inform the National Health Agency, which administers PMJAY, whether PMJAY should include a hospi-cash component.

Second, the experiment contributes to the health and development economics literature by determining whether, in low-to-middle-income countries (LMIC's) like India, medical treatment has important costs, specifically lost wages, that are not covered by health insurance. This will net out as demand for hospi-cash policies and the impact that such policies have on hospital utilization.

Third, the experiment contributes to the econometric literature on demand estimation and experimental design. First, instead of producing point estimates of a parametric demand function, we

[6] Uptake of the existing SSY product has been low for two reasons. First, SEWA did not market the product significantly. Specifically, their agents did not share information about the product with poorer members and they did not market the product in all villages in which they operated. Second, SSY's daily reimbursement was low relative to lost wages, especially in recent years and in non-rural areas. Our product and marketing will address both concerns.

[7] The version we study is also stripped of SSY's life insurance component.

[8] After the study is completed, SEWA hopes to scale these two new policies to other villages.

estimate bounds on a non-parametric demand function using the approach of Tebaldi et al. (2019), without requiring that the unobservable error terms take the probit or logit form. We improve on that paper by randomly assigning prices, facilitating causal interpretation of our nonparametric bounds. Second, we develop a new method to choose the product prices over which we randomize, so as to minimize the width of the bounds we estimate on demand. This moves the frontier on optimal experimental design by considering optimal design for an estimator able to characterize sharp nonparametric identified sets.

2 Funding, ethical approval, and trial registration

The SWIE was funded by the Tata Trusts via a grant to the Tata Centre for Development (TCD) at the University of Chicago, the Weiss Fund, JPAL and the University of Chicago Center for Global Health.

The IRB home for the study is the University of Chicago, though IRB filings are made at each investigator's home institution to the extent that they require access to data and in India via the survey agency (JPAL, housed at IFMR, Chennai, India). The study received IRB clearance from IFMR on June 10, 2019, and at the University of Chicago (IRB19-0219) on June 18, 2019.

The SWIE was registered with the American Economics Association (AEA) Registry (RCT ID: AEARCTR-0004240) on June 20, 2019 and with the ISRCTN Registry (RCT ID: ISRCTN63980830) on February 6, 2020.

3 Hypotheses

The SWIE was designed to address the following questions.

1. What is the demand curve for a hospi-cash policy in India?
2. Is there adverse selection into uptake of hospi-cash insurance?
3. Does purchase of a hospi-cash policy lead to increase utilization of care eligible for payouts (moral hazard)?
4. How does offering or purchase of a hospi-cash policy affect demand for or uptake of health insurance?
5. How does uptake of a hospi-cash policy affect health?
6. How does uptake of a hospi-cash policy affect employment and household financial status?
7. How does uptake of a hospi-cash policy affect employment and household financial status?
8. How does uptake of hospi-cash policy affect household consumption?

9. How does hospi-cash affect consumption differently for households that experience illness episodes?

10. Does hospi-cash uptake alter the pattern of risk-sharing within villages as reflected in changes in village level measures of consumption inequality.

In addition to these substantive questions, we are interested in the following methodological question.

7. Is it possible to optimally choose experimental prices to maximize the precision of sharp nonparametric estimates of a demand curve measuring purchase decisions at those prices?

4 Intervention

4.1 Hospi-cash policies

We examine demand for two hospi-cash policies. A hospi-cash policy is a financial product that pays beneficiaries a fixed indemnity payment for each day that a covered person undergoes a covered medical procedure at a hospital.

The hospi-cash policies we consider cover admission for an inpatient treatment or an outpatient surgery for certain family members. The household receives a fixed indemnity payment for each night that a covered member spends in a hospital. The policies have no deductible (i.e., minimum number of hospitalized days before compensation starts). Because these are indemnity policies, there is no co-pay per se. There is a cap of 15 days per household per annum on the total number of hospital stays covered by indemnity payments.

In order to receive compensation, a beneficiary must apply for benefits. The application requires submission of a hospital bill showing the number of days hospitalized or outpatient surgery at a hospital.[9] Reimbursement is placed in the SEWA member's JDY bank account after the application's acceptance.

These hospi-cash policies are offered on a "mutual-basis", meaning VimoSEWA, SEWA's insurance arm, serves as the insurer. The policies were not purchased from an outside insurance company. This allowed us greater control over product features and pricing.

The two hospi-cash policies in our study differ along 2 dimensions: the number and type of household members covered and the amount of indemnity per night hospitalized (Table 2).

[9] This application is submitted by phone or via a designated SEWA agent for each village. However, SEWA is in the process of developing an app that allows beneficiaries or agents to submit claims via smart phone or computer.

Table 2: *Two hospi-cash policies offered in study.*

	Policy 1	Policy 2
Covered persons	1 adult & 0 children	2 adults & up to 3 children
Indemnity amount	INR 500	INR 600

Notes: Policies can only be sold to households that have a SEWA member. (SEWA members are all female.) The first adult covered by a policy must be a female. This restriction is imposed because our partner, SEWA, is committed to improving female welfare. The second adult covered must be the female adult's husband. All adults covered by the policy must be between the ages of 18 and 54. SEWA imposed this restriction to stop adverse selection of high cost policy holders. The three children permitted under the 2+3 plan must all be younger than 18 years of age and must be the children of the primary adult and her husband. The household can decide which individuals to cover, as long as they meet the above eligibility criteria.

These two hospi-cash policies were chosen so as to ensure the two policies offered in sample villages are substantially different, that the two policies maximize consumer surplus,[10] and that the two policies are likely to generate equal shares of people enrolling in policy 1, policy 2 and neither policy. Appendix A describes the pilot study and methodology we employed to choose these two products.

Critical to implementation of the treatment is how VimoSEWA markets its hospi-cash products. SEWA stopped marketing its existing hospi-cash products in the 214 combined pilot and study villages. SEWA also hired and trained new sales staff for our sample villages. This allows us to review their training materials to ensure it was consistent with this study.

4.2 Pricing conditions

The SWIE randomizes sample villages to four different price combinations for the two policies above. The four price combinations are given in Table 3:

Table 3: *Four price conditions in the study.*

	Prices (INR/policy)	
Condition	1+0 plan	2+3 plan
1	150	500
2	250	700
3	400	950
4	550	700

Our study uses the methods in Tebaldi et al. (2019) to construct bounds on the demand for each hospi-cash policy at different prices. The prices in the table above were chosen optimally to minimize those bounds at 10 alternative price combinations using priors determined from a pilot

[10]For this purpose, consumer surplus was measured considering self-reported willingness to pay and assuming actuarially fair premiums. See Appendix A for more details.

study. Appendix B describes our pilot study and the methodology we employed to choose our four price combinations in Table 3.

4.3 Other SEWA products

Once households were assigned to an intervention, i.e., a price combination, they only had access to the 2 new hospi-cash policies chosen via the pilot studies. At the start of pilot 1, SEWA stopped marketing and sales of their of existing hospi-casjh policy (SSY). Contracts for pre-existing policies were honored, but the contracts were not renewed. By the time the baseline for the main study started, none of the 200 main study villages had ongoing SSY contracts.

Household had access throughout, however, to SEWA's other financial products, specifically their health care insurance product.

5 Design

5.1 Sample

The village and household samples for this study were assembled in two steps. *First*, we selected 214 villages in Ahmedabad and Gandhinagar districts in Gujarat state.[11] We worked in these two districts of Gujarat because SEWA, along with its insurance arm VimoSEWA, is headquartered in Ahmedabad. These are also two districts in which SEWA previously operated, but wanted to expand its insurance offerings.[12] Of these 214 villages, 5 were assigned to a pilot 1 that we used to determine which two hospi-cash policies we would study, 9 were assigned to the pilots 2 and 3 that we used to determine the price conditions for the experiment in the main study. The remainder were used for the main study.

Second, we selected households within each of these villages to survey. The universe was the set of households with SEWA members. While SEWA sold financial products to even non-members, we focused on members because SEWA had a census from which we could sample households in a manner that was representative of the universe. In total, there are 30,898 SEWA members in all 214 villages or 144 members per village on average. Some families have multiple SEWA members, so there are somewhat less than 144 SEWA-associated household per village.

[11] These villages are selected from a list of 292 Gujarat villages in which SEWA offered insurance products or was planning to do so. The 214 villages in our sample are selected because we have individual-level data for all SEWA members in these villages and because at least 18 SEWA members live in these villages. This 18 member cutoff was selected for 2 reasons. First, we wanted to makes sure there were enough SEWA members in a village to justify the fixed costs of surveying there. Second, we randomly select a subset of each village's members to sample. That random sampling is done by partitioning the population along terciles of age and (where available) income, which creates up to 9 partitions. 18 is a multiple of 9.

[12] Of the 214 villages, SEWA's earlier SSY wage insurance product was previously offered in 144 and was offered up until the study began in 106 villages. In the remaining villages, SEWA either had some operations but did not offer financial products, or had just recently entered that village.

In the next two sections we explain how we selected villages for the pilots and the main study. Then we explain how we selected households in the villages in which we operate.

5.1.1 Sample villages

Pilot 1 and 2. Out of the 214 total villages in our study, we selected 5 villages for each of pilots 1 and 2. The 5 villages for each pilot were selected to be representative of all the 214 villages in the following manner:

1. For each pilot, we decided to select 3 villages from Ahmedabad and 2 villages from Gandhinagar because the 214 villages had roughly 3 times as many villages in Ahmedabad as in Gandhinagar but we wanted more than 1 village from Gandhinagar.

2. In each district we had a few variables per household in each village. We wanted to create a distribution using these variables. However, the variables had different units and the variables we had differed across districts. Therefore, we standardized variables by calculating percentiles for the village means or sums of variables and use a different set of variables for this process in each district. In Ahmedabad district, we took village means of annual income of households, the ages of SEWA members, the number of years of education of SEWA members, and an indicator for whether the SEWA member works in agriculture, and we took village-level sums of the SEWA member population. In Gandhinagar district, we took village-level means of the age of SEWA members, an indicator for whether the member works in agriculture, and whether the member is literate, and we took village-level sums of the number of SEWA members per village.

3. Next we had to choose 2-3 villages per district that were representative from a 4+ dimensional distribution. We did this in three steps.

   (a) We chose what parts of the distribution to draw from. Because we chose 3 villages from Ahmedabad and 2 from Gandhinagar, we chose 3 percentile cutoffs (1/6, 1/2, and 5/6) for Ahmedabad and 2 cutoffs (1/4, 3/4) for Gandhinagar.[13]

   (b) We collapsed the 4+ dimensional distribution into a 1 dimensional problem. For each village j in district k, cutoff value γ_k in district k, and variable i observed in district k, we calculated $(z_{ijk} - \gamma_k)^2$, the squared difference between (a) the percentile z_{ijk} of the mean or sum of i in that village relative to all villages in the sample in that district and (b) the cutoff point. Then, for each cutoff γ_k in district k, we summed those squared differences across variables i measured in that district: $\sum_i (z_{ijk} - \gamma_k)^2$.

[13] E.g., for Ahmedabad, we divided the unit line into 1/3 segments and 1/6, 1/2, and 5/6 were the midpoints of each segment.

(c) We chose the two villages that were closest to each cutoff in each district. Specifically, for each cutoff γ_k in district k[14] we choose two villages j and j' in that district that minimized the sum of the squared distances from the cutoff value for each variable:

$$\min_{j \in k} \sum_i (z_{ijk} - \gamma_k)^2 \tag{1}$$

4. For each pair of villages that was closest to a cutoff in a district, we randomly assigned one village to pilot 1 and one to pilot 2. Since we had 2 villages per cutoff and 5 cutoffs, this yields 10 total villages, with 5 assigned to each pilot.

Pilot 3. In a similar fashion to how villages were selected for pilots 1 and 2, we selected 4 villages (3 in Ahmedabad, 1 in Gandhinagar) for a third pilot in October 2020. These villages were selected from the 204 villages left after villages were picked for Pilots 1 and 2. Because we only selected 1 village for Gandhinagar, the cutoff was 1/2 for that district. A single representative village with the lowest sum of squared differences from the means of the dimensions previously used is selected in Gandhinagar for pilot 3.

Main study. We assigned the remaining 200 villages to the main study.

5.1.2 Sample households

Household sample. Our inclusion criteria for households in the 3 pilot studies and the main study was that households have at least one SEWA member currently living there. (For example, the adult children of a SEWA member who no longer lives with them were ineligible to participate in the study.) While this criteria affects external validity, it provides a cohesive sample that already trusted SEWA and its financial services arm, VimoSEWA, and is thus likely to be cooperative with our study.

There were no exclusion criteria for pilot 1. For pilots 2 and 3 and the main study, however, we excluded members that were subscribed to VimoSEWA's existing hospi-cash policy. We excluded these members because a member who already has a hospi-cash policy is unlikely to purchase a second policy. Moreover, we are interested in demand for hospi-cash amongst those who do not already have it.[15] We also excluded members below the age of 18 and above the age of 54 from pilots 2 and 3 and the main study, as the hospi-cash policy is not available to individuals outside this age band.

[14] Here $\gamma_k \in \{1/6, 1/2, 5/6\}$ in Ahmedabad and $\gamma_k \in \{1/4, 3/4\}$ in Gandhinagar.

[15] This exclusion criteria did not bind by the main study because SEWA stopped selling their old hospi-cash policies in the 214 study villages. Households could keep their old policies if they were enrolled before the start of our pilots, but those policies would not be renewed. By the time the main study started, few if any household still had an old hospi-cash policy from SEWA.

We selected 250 households from the 5 villages assigned to pilot 1, 300 households from the 5 villages assigned to pilot 2, and 200 households from the 4 villages assigned to pilot 3. Our household selection processes for pilots 1, 2 and 3 are outlined in Appendices A, B and C, respectively. Here we describe how we selected 10,450 households for the main study.

In each village we assembled both a primary and backup sample of households. The primary households were the first households we visit for consent and surveying. If a household from the primary sample was unavailable, e.g., because no one was home or they refused consent, then we selected a household from the backup sample. If a backup sample household was unavailable, we select another backup sample household. Because we planned to sample households based on age and perhaps income-based partitions of the set of all SEWA households in a village, the backup sample for each primary sample household was drawn from the same partition as the primary sample.

In Ahmedabad district villages, SEWA has data on SEWA member income and age. We used both for selecting a representative sample. First, we partitioned SEWA households into terciles based on age. Within each tercile we partitioned SEWA households into terciles based on income. For each eligible SEWA household in one of the resulting 9 age $\times$ income partitions, we generated a uniform random number between 0 and 1 inclusive. Because the sample size we wanted (10,450) was $\sim$ 36.5% of all eligible SEWA households in the 200 main study villages, we assigned each household with a random number less than $\sim$0.365 to the primary sample. We assigned each household with a random number above that threshold to the backup sample.

In Gandhinagar villages, we only have data on age of SEWA members. Therefore, we partitioned SEWA households in a village into age terciles. Then we used the same random number assignment and threshold to assign $\sim$0.365 of households in each partition to the primary sample and the rest to the backup sample.[16]

5.2 Power calculations

Our experiment randomizes households (via the village to which they are assigned) into 4 conditions characterized by a pair of prices for the two hospi-cash products we offer for sale. We will estimate the fraction of the population that purchases each product at each price combination. However, these fractions are not our primary parameter of interest. Instead, following Tebaldi et al. (2019), we use these fractions to estimate bounds on a non-parametric demand curve for each of the two products. In short, we are not primarily interested in ensuring that we have enough power to estimate these fractions. Our challenge with sample size calculations was that there is no conventional approach to calculating sample size for set identification as this empirical strategy produces bounds, not confidence intervals.

[16]This random number generation process is not guaranteed to yield 10,450 households in the primary sample every time. Therefore, we ran the procedure until we obtain an iteration that yields exactly 10,450 households.

However, we will also examine the ITT effect of 4 price combinations on several secondary outcomes, e.g., uptake of health care insurance. Moreover, we will examine, e.g., the impact of uptake of hospi-cash policies on utilization of hospital care via a TOT estimator where assignment to 4 price conditions are instruments for the endogenous treatment variable uptake of hospi-cash insurance.

To address these secondary aims, we calculate a sample size assuming the primary outcome is hospitalization, there are 4 treatment arms and we use an ITT estimator. This is conservative because hospitalization is more rare than uptake of hospi-cash insurance or health care insurance. However, it is imperfect because we use usual one sample formula rather than a formula for IV estimation (e.g., Walker et al., 2017). Specifically, like Malani et al. (2021), we estimated that we would need 2,250 household per arm to detect a 25% change in hospitalization rate (off a baseline hospitalization rate of 10% per household per annum) with a 5% significance level and 80% power and allowing for 10% attrition. We did not use the intra-cluster correlation to adjust the sample size to account for the fact that we cluster randomized households, i.e., randomized villages, to price combinations. Instead we included an *ad hoc* design effect of 1.11. Thus, we chose a per arm sample size of 2,500 rather than 2,250, implying a total sample size of 10,000.

5.3 Treatment assignment

Each of the 200 villages are randomly assigned to one of four combinations of prices p in the set $P = (\{150, 500\}, \{250, 700\}, \{400, 950\}, \{550, 700\})$ for the two products $\theta \in$ {1+0 & INR 500, 2+3 & INR 600}. This is done in two steps.

1. Villages are blocked by taluka and distance from nearest hospital or community health center, whichever is closer. These distances are coded as 0-5 km, 5-10 km, and 10+ km.

2. Within each block, villages are randomly blocked into groups of four and are assigned to the four $p \in P$ without replacement.

3. Remainder villages, i.e., those that aren't evenly divided into groups of four, are randomly assigned $p \in P$ without replacement as well.

4. The algorithm assigning villages to price pairs stops when an equal number of villages are assigned across all price combinations.

5.4 Data collection

We conduct 6 rounds of data collection. Before we started our main study, we conducted 3 pilots, one to determine which 2 hospi-cash products to offer and two to determine the experimental prices for those two policies. Each pilot had an associated survey that is described in Appendices A and B.

For our main study, we conducted (1) a baseline survey just before SEWA begins marketing the 2 chosen hospi-cash policies to main study sample villages. The baseline will measure hospital utilization, financial status, and health. After the marketing begins, for a period of 6 to 12 months, we will gather (2) data from SEWA on uptake of under each hospi-cash policy and SEWA's other health insurance policies. Finally, (3) we will conduct an endline survey, 9 to 12 months after the marketing begins, to measure health insurance uptake, hospital utilization, claims made, financial status, and health.

5.5 Outcomes

During baseline we measure 2 categories of variables in order to estimate the degree of adverse selection into insurance:

1. Health status of SEWA members
2. Healthcare utilization of sample households.

Starting after baseline, we measure 4 primary outcomes measured at the household level for each household in the sample:

1. Uptake of each of the 2 selected hospi-cash products
2. Uptake of health insurance products (whether sold or distributed by SEWA or not)
3. Days of hospitalization, by each adult member and by any minor members of the household
4. Days of hospital treatment reimbursed under the hospi-cash policies.

In addition we measure the following secondary outcomes:

1. Uptake of other insurance and/or financial products
2. Number of days of work in last month and average daily wage on days worked
3. Monthly income
4. Asset index
5. Savings
6. Monthly consumption expenditure (net of medical expenses).
7. Measures of within village consumption inequality (e.g. variance of log consumption for households in a village)
8. Monthly medical expenditures
9. Battery of health and wellness measures.

5.6 Timeline

The following is a timeline of the project.

- **IRB.** We received IRB approval at IFMR on June 10, 2019. We received IRB approval at UChicago on June 18, 2019.
- **Pilot 1.** We conducted the first pilot in July 2019 to determine the features (family member eligibility and indemnity size) of the 2 hospi-cash policies the study would offer.
- **Pilot 2.** We conducted the second pilot in October-November 2019, to determine experimental prices for the 2 chosen hospi-cash policies.
- The COVID-19 pandemic hit in February 2020. We paused field operations.
- **Pilot 3.** We conducted the third pilot starting in late-December 2020. This was after India's first COVID-19 wave, but before its second wave in April 2021. We conducted Pilot 3 to revisit the experimental prices for the 2 chosen hospi-cash products, in case they had changed due to the pandemic.
- The second and third waves of the COVID-19 pandemic in India delayed the launching of baseline by about 12 months.
- **Baseline.** Baseline data collection launched in March 2022.
- **Treatment assignment**. We will reveal to SEWA the price combination to be offered in each study village for the 2 chosen hospi-cash policies.
- **Marketing.** Starting in mid-July 2022, SEWA will begin marketing the 2 chosen hospi-cash products we study at the prices assigned for each village. Marketing in a village will begin only after the baseline survey and backcheck are completed in that village.
- **Measurement of uptake and claims.** We measure purchase of insurance and hospi-cash claims from July 2022 to June 2023.
- **Endline.** We expect to conduct endline from December 2023 to April 2024.

6 Identification and Estimation

We will proceed in five ways to analyze the demand and cost impact of hospi-cash insurance and test hypotheses in Section 3.

- First, we will estimate nonparametric bounds on demand for hospi-cash insurance following the approach in Tebaldi et al. (2019), as well as parametric discrete choice models of hospi-cash demand (McFadden, 1973; Berry, 1994).

- Second, we will adapt results from Einav et al. (2010a,b); Geruso et al. (2019) to verify the presence of, and quantify the degree of adverse (or advantageous) selection into hospi-cash insurance (i.e., whether prior health status or health care utilization affects uptake conditional on price) and the presence of moral hazard (i.e., whether individuals consume more healthcare when insured, conditional on prior health status).

- Third, we will use intent-to-treat (ITT) estimators to measure the impact of different price combination on uptake of health care insurance.

- Fourth, we will use complier-average-treatment effect (CATE) or treatment-on-treated (TOT) estimators to determine the impact of enrollment on labor supply, financial status and health as well as to estimate heterogeneous treatment effects.

- We will follow Attanasio and Székely (2004) to study the effects of hospi-cash on risk-sharing by comparing changes in measures of consumption inequality across treatment arms (using ITT estimators).

We discuss each estimation strategy in turn.

6.1 Hospi-Cash Demand

The experiment assigns villages to 4 different price treatments (see Table 3), $p = P^1, P^2, P^3, P^4$. Since we collect information on effective purchases and on households who would be eligible for purchasing hospi-cash policies, after treatment we observe data that will provide, for every price vector p, a measurement of the share of individuals purchasing low-coverage (L), $s_L(p)$, the share of individuals purchasing high-coverage (H), $s_H(p)$, and the share of individuals not purchasing any policy, $s_0(p)$.

Identification of demand relies on the random assignment of prices across villages. To illustrate, consider an individual i, with observable characteristics $X_i = x$, and living in a village with the treatment assignment P^k, $k = 1, 2, 3, 4$. We let the indirect utility from purchasing hospi-cash option $j = 0, L, H$ be equal to

$$U_{ij} = V_{ij} - P_j^k. \tag{2}$$

Figure 1a shows the implication of this choice model. Every price vector partitions the space of V_L, V_H in three regions, which determine choice.

Using this model, Tebaldi et al. (2019) show how to compute sharp nonparametric bounds on the share of individuals purchasing any option at a given unobserved price vector $p^\star$, $s_j(p^\star)$ for $j = 0, L, H$. The method amounts to maximizing (minimizing) $s_j(p^\star)$ over all possible preferences such that $s_j(p)$ is consistent with the observed choice data across j and across price treatments p. We will complement this approach with parametric maximum likelihood methods (e.g., probit,

Figure 1: *Demand, Adverse Selection, and Moral Hazard*

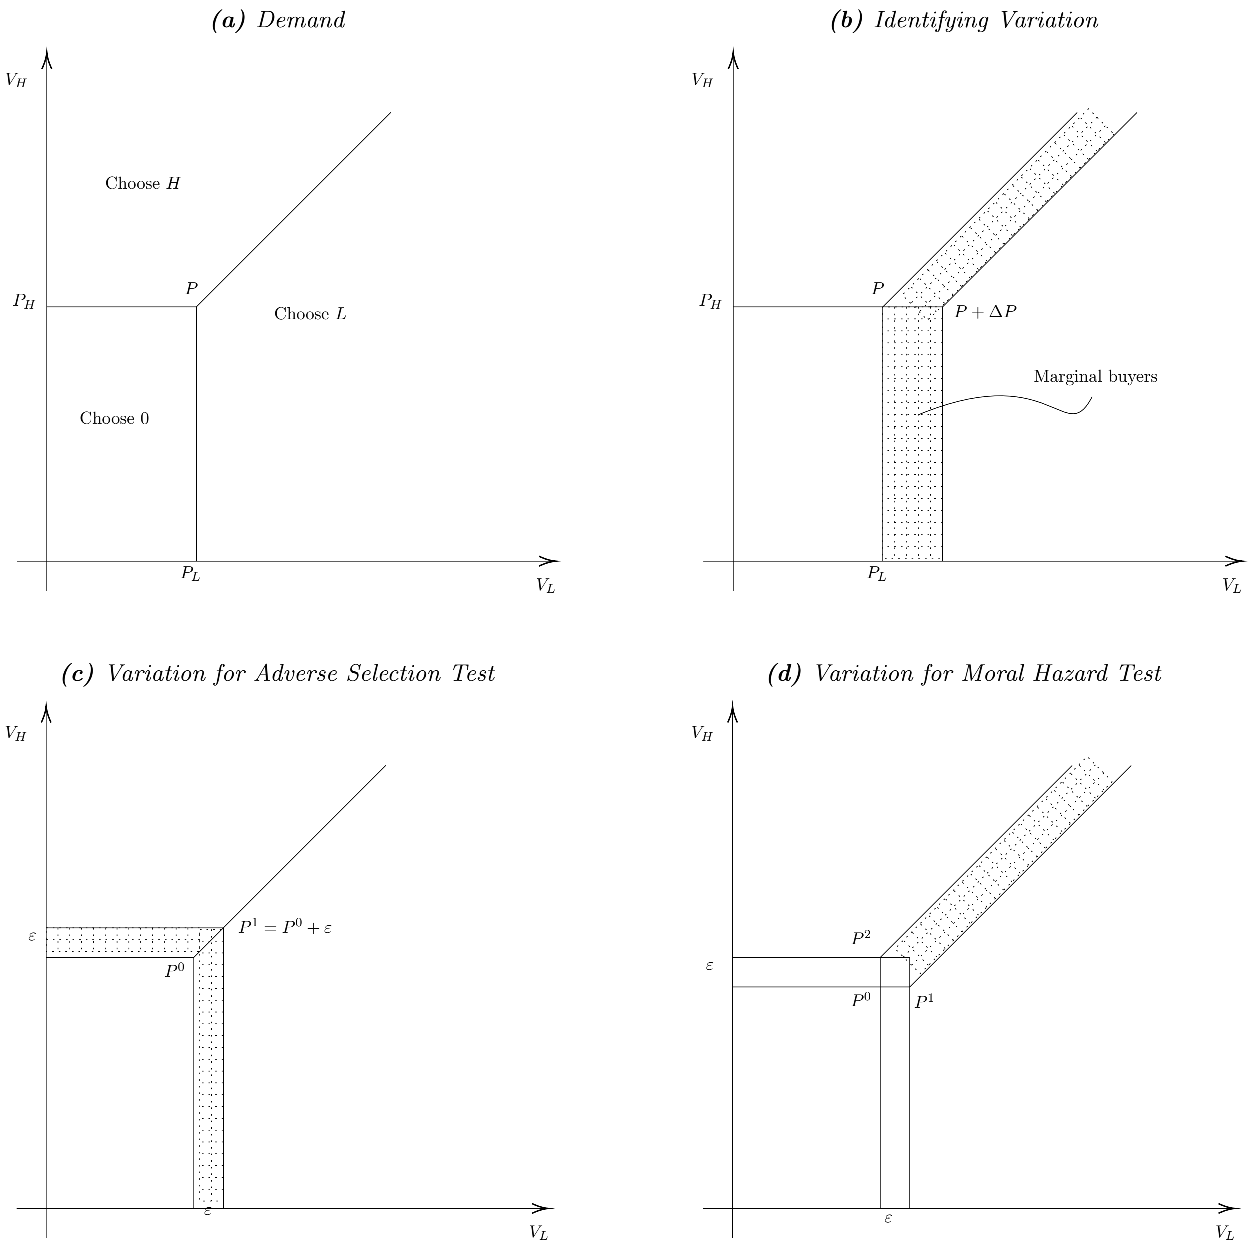

logit, mixed logit, see McFadden, 1973; Berry, 1994) to study the demand for hospi-cash policies over a range of prices.

Letting $f(V_i|X_i = x)$ be the (unknown) density of preferences conditional on $X_i = x$, the share of individuals purchasing j when assigned to treatment P^k is the integral of this density over the regions illustrated in Figure 1a:

$$s_j(P^k|X_i = x) = \int \mathbf{1}\left[\underbrace{V_{ij} - P^k_j \geq V_{i\ell} - P^k_\ell \quad \text{for all } \ell}_{\text{Choose } j}\right] f(V_i|X_i = x)\ dV_i. \tag{3}$$

If we assume that $V_{ij} = X_i\beta + \sigma\varepsilon_{ij}$ and that ε_{ij} is iid Type 1 extreme value (logit), we have

$$s_j(P^k|X_i = x) = \frac{e^{\left(x\beta - P^k_j\right)/\sigma}}{\sum_{\ell=1}^{3} e^{\left(x\beta - P^k_\ell\right)/\sigma}}. \tag{4}$$

Our experiment varies P^k randomly (see e.g. Figure 1b), and we will therefore maintain $f(V_i|X_i = x, P^k) = f(V_i|X_i = x)$ by assumption. This allows us to identify nonparametric bounds and confidence sets on $s_j(P^\star|X_i = x)$ for any $P^\star$. Moreover, imposing parametric assumptions as in (4), we can estimate β, σ, and obtain point estimates and confidence intervals on $s_j(P^\star|X_i = x)$ for any $P^\star$.

6.2 Adverse Selection and Moral Hazard

To test for and quantify adverse (advantageous) selection and moral hazard, we will adapt the approach outlined in Einav et al. (2010a,b), and recently revisited in Geruso et al. (2019).

In our context, individuals will have three options: no hospi-cash ($j = 0$), the "low-coverage" plan ($j = L$), and the more generous coverage plan ($j = H$). We can then completely define a household as a collection

$$\zeta_i = (V_{iL}, V_{iH}, C_{i0}, C_{iL}, C_{iH})\,; \tag{5}$$

V_{iL} and V_{iH} determine choice at any given price vector (see equation (2) and Figure 1a), C_{ij} describes the number of hospital days by each adult member and any minor member of the household when the household chooses j. Importantly, we only observe C_{ij} for the chosen j. Moreover, we expect the measurement of C_{i0} for households not purchasing hospi-cash coverage to be significantly noisier than the measurement of C_{iL} and C_{iH} among enrollees.

6.2.1 Adverse Selection

We will say that *the hospi-cash market features adverse selection if marginal enrollees are, on average, less risky than inframarginal enrollees.*[17] The following gives the spirit of our test and

[17] In a market with many products, like ours, the definition of adverse selection is less straightforward than in markets with two options (e.g. Einav et al., 2010a). In particular, one can define adverse selection between any two options (see Einav et al., 2010b; Geruso et al., 2019). Here we focus primarily on the extensive margin decision to enroll in any hospi-cash policy.

quantification of adverse selection for any given price vector P^0 (see Figure 1c), though we may revise somewhat the methodology as we refine it.

Let $TC(P)$ be the total number of hospital days among all hospi-cash enrollees in a village when the price vector is P, $Q_j(P)$ be the number of households enrolled in plan j, and $AC(P)$ be the average number of hospital days among hospi-cash enrollees at the same price vector:

$$AC(P) = \frac{TC(P)}{Q_L(P) + Q_H(P)}. \tag{6}$$

Consider a (small) price variation from P^0 to $P^1 = P^0 + \varepsilon$, $\varepsilon > 0$ (see Figure 1c). There is adverse selection at P^0 if

$$AC(P^1) > \frac{TC(P^0) - TC(P^1)}{Q_L(P^0) - Q_L(P^1) + Q_H(P^0) - Q_H(P^1)}. \tag{7}$$

In words, we can test for adverse selection if the average cost among inframarginal buyers is larger than the average cost among marginal buyers. Importantly, all quantities in (7) can be estimated directly using the exogenous variation across price treatments. Here we sketch the intuition for our procedure using a parametric approach, but we will endeavor to develop a nonparametric version extending the intuition developed in Tebaldi et al. (2019).

We observe household-level uptake and claims for every product $j \in \{0, L, H\}$ and for varying treatment assignments P^k, $k \in \{1, 2, 3, 4\}$. This enables us to estimate OLS regressions of the form:

$$D_{ij} = \alpha + \beta_{Hj} P^k_{Hi} + \beta_{Lj} P^k_{Li} + \epsilon_{ij} \tag{8}$$

$$C_{ij} = \gamma + \delta_{Hj} P^k_{Hi} + \delta_{Lj} P^k_{Li} + \upsilon_i \tag{9}$$

for each $j \in \{L, H\}$, where D_{ij} is equal to 1 if household i choose contract j and 0 otherwise, P^k_{ji} is price of policy j under price combination k, C_{ij} is the realized number of hospital days for household i with policy j, and the errors are clustered at the village level. (Note: we will use different measures of cost and utilization, but here we keep referring to number of hospital days.)

With the parameters from equations (8) and (9) we can write

$$Q_L(P) = \sum_i \alpha + \beta_{HL} P_H + \beta_{LL} P_L \tag{10}$$

$$Q_H(P) = \sum_i \alpha + \beta_{HH} P_H + \beta_{LH} P_L \tag{11}$$

$$TC(P) = Q_L(P) * (\gamma + \delta_{HL} P_H + \delta_{LL} P_L) + Q_H(P) * (\gamma + \delta_{HH} P_H + \delta_{LH} P_L). \tag{12}$$

We are then able to construct a test statistic and p-value for the null hypothesis of no selection:

$$H0: \;\; AC(P) = \frac{TC(P) - TC(P+\varepsilon)}{Q_L(P) - Q_L(P+\varepsilon) + Q_H(P) - Q_H(P+\varepsilon)}, \tag{13}$$

against one-sided alternatives that define adverse or advantageous selection, respectively.

6.2.2 Moral Hazard

We will say that *the hospi-cash market features moral hazard if more generous coverage corresponds to a larger number of hospital days.* The following defines a test and quantification of moral hazard (see Figure 1d).

Moral hazard is present if

$$C_{i0} < C_{iL} < C_{iH}. \tag{14}$$

We sketch a test of moral hazard using only hospitalizations among hospi-cash enrollees. Depending on the quality of the measurement of hospitalizations for those who select $j = 0$, the analysis can be extended accordingly.

Absent moral hazard, the distribution of hospitalizations among households who are at the margin between choosing H and L does not depend on their coverage choice. More precisely, the distribution of cost among those switching from L to H when the price of L increases by ε (see Figure 1d) is the same as the distribution of cost among those switching from H to L when the price of H increases by ε. If there is moral hazard, the cost among switchers in the former case will be higher than in the latter.

More precisely, consider a price vector $P^0 = (P_L^0, P_H^0)$, and two alternative price vectors $P^1 = (P_L^0 + \varepsilon, P_H^0)$, $P^2 = (P_L^0, P_H^0 + \varepsilon)$, see also Figure 1d. The average number of hospitalizations among the switchers from L to H when the price of L increases is

$$\frac{TC_H(P^1) - TC_H(P^0)}{Q_H(P^1) - Q_H(P^0)}. \tag{15}$$

The average number of hospitalizations among the (same) switchers from H to L when the price of H increases is

$$\frac{TC_L(P^2) - TC_L(P^0)}{Q_L(P^2) - Q_L(P^0)}. \tag{16}$$

We will detect moral hazard if

$$\frac{TC_H(P^1) - TC_H(P^0)}{Q_H(P^1) - Q_H(P^0)} > \frac{TC_L(P^2) - TC_L(P^0)}{Q_L(P^2) - Q_L(P^0)}, \tag{17}$$

since that is the case only if

$$\int \mathbf{1}\left[V_{iH} - P_H \approx V_{iL} - P_L > 0\right] C_{iH} f(V_i) dV_i > \int \mathbf{1}\left[V_{iH} - P_H \approx V_{iL} - P_L > 0\right] C_{iL} f(V_i) dV_i. \quad (18)$$

A parametrization like the one in (8) and (9) will allow us to construct a test statistic and p-value for the equality between (15) and (16) against (17).

6.3 Intent-to-treat (ITT) estimates

We use regression models of the form below to estimate the ITT effect of different price combinations on outcomes such as health insurance uptake:

$$y_i = \alpha + \beta_H P_{Hi} + \beta_L P_{Li} + \epsilon_i \quad (19)$$

where i indexes households and (P_{Hi}, P_{Li}) are the prices to which households are randomized. We weight households equally and the errors are clustered at the village level. The ITT effect of a change in the price of plan j is β_j.

6.4 Complier average treatment effects (CATE) or treatment on treated (TOT)

We use regression models of the form below to estimate CATE or TOT effects of each insurance plan on outcomes such as health, employment, financial status, and consumption:

$$y_i = \alpha + \gamma_L D_{Li} + \gamma_H D_{Hi} + \upsilon_i \quad (20)$$

where D_{ji} is an indicator for whether household i took up policy j and we instrument for these uptake indicators with the price vector (P_{Hi}, P_{Li}). We also extend this basic approach in the manner below to estimate heterogeneous treatment effects based on whether a household experienced a sickness episode:

$$y_i = \alpha + \gamma_L D_{Li} + \gamma_H D_{Hi} + \delta_L D_{Li} X_i + \delta_H D_{Hi} X_i + \phi X_i + \upsilon_i \quad (21)$$

Again these regressions weight each households equally and errors are clustered at the village level.

6.5 Risk Sharing

We will compare changes in village level consumption measures across villages in different arms treatment and control villages to assess whether the offer of the hospi-cash policy affects risk-sharing (as proxied by our consumption measures). Broadly speaking, this will follow the analysis outlined in Attanasio and Székely (2004) and following the general estimation approach in Section 6.3.

7 Funding

Funding from this project comes from the Tata Centre for Development (TCD) at the University of Chicago, a center in turn funded by the Tata Trusts and based in New Delhi. A research grant from the Weiss Fund for Research in Development Economics provides additional support for this project.

Appendix

A Pilot 1: Optimal selection of products

A.1 Overview

We wanted to offer 2 products, each at 4 prices, in the main study. In the first pilot, we selected which 2 products to offer. We did this by measuring preferences over policies that vary along two dimensions[18]:

1. The number of people covered $n \in \mathbf{N} = \{1 \text{ adult} + 0 \text{ children}, 2+0, 2+1, 2+2, 2+3\}$, and
2. The daily coverage rate $r \in \mathbf{R} = \{\text{INR } 200, 300, 400, 500, 600\}$.

Once we measured preferences, we used estimated consumer surplus and uptake estimates to choose two products, subject to certain constraints on the design of the main experiment. Here we describe the sample for this pilot and then describe our analytic methods.

A.2 Sample

We described in the main text how we select 5 villages for pilot 1. Here we describe how we selected 250 households from the universe of households with SEWA members in these villages. We chose 250 households not based on any power calculation but on the arbitrary target that we would sample 50 households per village, on average, throughout the study.

In each village we assembled both a primary and backup sample of households for the pilot, just as we did for the main study households. In Ahmedabad district villages, we used both age and income for selecting a representative sample. First, we partitioned SEWA households into quintiles[19] based on age. Within each quintile, we partitioned SEWA households into deciles based on income. We select one primary household from each of the 50 resulting age $\times$ income partitions. Then across each age partition, we randomly selected 8 to 14 households for the backup sample. In Gandhinagar villages, we only had data on age of SEWA members. Therefore, we partitioned SEWA households in a village into age quintiles. Then we randomly selected 10 households in each age partition for a primary sample and 11 to 14 households in each age partition for the backup sample.

[18] Other features of the products were determined by our partner, SEWA. SEWA decided that at least one adult beneficiary had to be a SEWA member between the ages of 18-54, that the total number of covered days across all household members would be 15 per year, and that coverage would include inpatient treatments and outpatient surgeries. We do not believe the 15 day cap is binding on many households. Singh and Ladusingh (2010) found that the 90th percentile length-of-stay for hospital treatment was 15 days.

[19] Our method of sampling for the pilot differed from our method for the main study because the smallest villages selected for the pilot have a larger number of SEWA members (over 50). We believed this would allow us to use a finer partition (quintiles) and still generate a reasonable number of primary and backup households per partition; however, such fine partitions resulted in logistical challenges in the field that induced us to change course in pilot 3, described below.

In total, this yielded 250 households in the main pilot 1 sample and 298 households in the backup pilot 1 sample. We were unable to assign exactly 14 backup households to each age quintile in each village in the pilot one sample, as several age quintiles did not have the requisite number of households (i.e., 24).[20] If all of the backup households were used in a particular age quintile-village pair during pilot 1 and our sample was still short of 250 households, surveyors proceeded to backup households in other age quintiles within the same village. If there weren't any, surveyors proceeded to survey backup households in other villages.

A.3 Survey design

The first step was to ask each household how much they are willing to pay for each of a set of hospi-cash policies. We do not want to ask each household how much they are willing to pay for each of the 25 possible hospi-cash policies we are considering as that would cause confusion. (There are 5 options for covered lives and 5 options for indemnity level, so 25 combinations.) Instead we want to ask each household about 5 policies where we hold covered lives constant and vary indemnity or vice versa. To accomplish this and get willingness to pay for 25 different policies we take the following steps:

1. Within each village, sample households are randomly blocked into groups of 10. Within each block, households are randomly blocked again into two subgroups of 5.
   (a) Within the first block of each subgroup, each household i_1 is randomly assigned a coverage number n_{i_1} from N, without replacement, and randomly assigned an indemnity rate $r_{i_1} \in \{200, 600\}$.
   (b) Within the second block of each subgroup, each household i_2 is randomly assigned an indemnity rate r_{i_2} from R, without replacement, and randomly assigned a coverage number $n_{i_2} \in \{1+0, 2+3\}$.
2. For each household i_1, we ask the following 5 questions:
   (a) What is the household's willingness-to-pay w_{i_1} for a hospi-cash policy that offers INR r_{i_1} in compensation per day in the hospital, up to a maximum of 15 days per annum, for up to n_{i_1} members of the household. (The respondent chooses from an array of prices ranging from 0 to INR 3000 in increments of INR 50.)
   (b) We ask the same question as above, but we increment up (down) the indemnity level by INR 50 if $r_{i_1} = 200$ ($r_{i_1} = 600$).
3. For each household i_2, we ask the following 5 questions:

[20] We initially wanted 40% more backup sample households than main sample households, just in case. This target was chosen arbitrarily.

(a) What is the household's willingness-to-pay w_{i_2} for a hospi-cash policy that offers INR r_{i_2} in compensation per day in the hospital, up to a maximum of 15 days per annum, for up to n_{i_2} members of the household. (The respondent chooses from an array of prices ranging from 0 to INR 3000 in increments of INR 50.)

(b) We ask the same question as above, but we increment up (down) the coverage number by 1 person in N if $n_{i_2} = 1 + 0$ ($n_{i_2} = 2 + 3$).

4. We ask each household approximately 40 questions about family composition, income, home quality, other assets, existing insurance subscriptions, and hospital use.

A.4 Data analysis.

After conducting the first pilot, we learned that VimoSEWA would only enroll adults between the ages of 18 and 54, and that the primary member on a plan must be female. Therefore, the first step of our analysis was to omit individuals greater than 54 from our analysis and all men who participated in the pilot 1 survey. From an original sample of 250 households, this left 153 households.

Second, we calculated the consumer surplus for each product. We began by obtaining an estimate of the predicted premiums that VimoSEWA, SEWA's insurance arm, would charge for a hospi-cash policy with features (r, n). These premiums are reported in the table below.

Table 4: *SEWA estimated premiums by hospi-cash policy feature*

Indemnity Rate	Coverage Rates				
	1 adult + 0 children	2 adult + 0 children	2 adult + 1 child	2 adult + 2 children	2 adult + 3 children
Rs. 200	Rs. 225	Rs. 300	Rs. 400	Rs. 400	Rs. 450
Rs. 300	Rs. 250	Rs. 350	Rs. 475	Rs. 475	Rs. 525
Rs. 400	Rs. 275	Rs. 375	Rs. 550	Rs. 550	Rs. 625
Rs. 500	Rs. 300	Rs. 425	Rs. 600	Rs. 600	Rs. 675
Rs. 600	Rs. 325	Rs. 450	Rs. 650	Rs. 650	Rs. 775

Third, we calculate the consumer surplus $c_i(r_j, n_k)$ for a given household i for a product with features (r_j, n_k) as the willingness to pay $w_i(r_j, n_k)$ obtained from the Pilot 1 survey minus the corresponding premium π_{r_j,n_k} from the table above. We thus obtain a set of five consumer surplus estimates $c_i(\mathbf{r}, \mathbf{n}) \in C = \{c_i(r_{j_1}, n_{k_1}), c_i(r_{j_2}, n_{k_2}), c_i(r_{j_3}, n_{k_3}), c_i(r_{j_4}, n_{k_4}), c_i(r_{j_5}, n_{k_5})\}$ for each household i. We then do pair-wise comparisons of the consumer surpluses in set C and create an indicator for insurance uptake by household i if $\max\{0, c_i(r_j, n_k), c_i(r_{j'}, n_{k'})\} > 0$ for each $j, j', k, k' \in 1, ..., 5$ offered to i. We also define the consumer surplus from each product pair for household i to be $s_i(r_j, n_k, r_{j'}, n_{k'}) = \max\{0, c_i(r_j, n_k), c_i(r_{j'}, n_{k'})\}$ for each $j, j', k, k' \in 1, ..., 5$ offered to i.

Fourth, for each product pair $\{(r_j, n_k), (r_{j'}, n_{k'})\}$, we calculate the proportion $\rho(r_j, n_k, r_{j'}, n_{k'})$ of households that would take up that product pair amongst the households offered that pair in Pilot 1. We also calculate for each product pair, the average consumer surplus $\bar{s}(r_j, n_k, r_{j'}, n_{k'})$ generated amongst households that were offered that pair.

Fifth, we wanted to use these to predict out of sample purchase shares and consumer surpluses for product combinations that were offered to no households in Pilot 1. Therefore, we ran the following regressions:

$$\begin{aligned}\rho(r_j, n_k, r_{j'}, n_{k'}) &= \alpha_0 + \alpha_{1.} r_j + \alpha_{2.} r_{j'} + \alpha_{11} r_j * n_j + \alpha_{22} r_{j'} * n_{j'} \\ &\quad + \alpha_{21} r_{j'} * n_j + \alpha_{12} r_j * n_{j'} + \epsilon_j \qquad (22a) \\ \bar{s}(r_j, n_k, r_{j'}, n_{k'}) &= \beta_0 + \beta_{1.} r_j + \beta_{2.} r_{j'} + \beta_{11} r_j * n_j + \beta_{22} r_{j'} * n_{j'} \\ &\quad + \beta_{21} r_{j'} * n_j + \beta_{12} r_j * n_{j'} + \upsilon_j \qquad (22b)\end{aligned}$$

where $\{(r_j, n_k), (r_{j'}, n_{k'})\}$ are indicators for each of the indemnity and coverage rates in R and N, respectively. We used estimates from this regression to predict $\rho(r_j, n_k, r_{j'}, n_{k'})$ and $\bar{s}(r_j, n_k, r_{j'}, n_{k'})$ out of sample.

Sixth, we selected two products by creating an index that weights the follow three objectives equally:

1. The absolute value of the difference between coverage under policy 1 and policy 2 in a pair.
2. Predicted consumer surplus under the product pair.
3. The square of the difference between predicted uptake under a policy pair and 0.67 (which is equivalent to 1/3 uptake of policy 1, 1/3 of policy 2 and 1/3 no purchase).

The first two goals were an objective of our project partner SEWA. They want to offer a wide range of products. They also wanted to maximize consumer surplus because the purpose of their organization is to improve the welfare of their members. The third goal generates equal market shares across policy 1, policy 2 and neither policy. Because we cannot assign people to policies, we want to do the next best thing: choose prices so that there are equal numbers of people that select each policy.

The three objective are measured by variables with different units; yet we want to weight each of the three variables equally. To do this, we create z-scores for each variables based on the distribution of that variable. We then take a simple sum of z-scores across the 3 variables for each product pair. We select the product pair with the highest summed "Z-score":

- Number covered: 1 adult + 0 children, indemnity: INR 500

- Number covered: 2+3, indemnity: INR 600

A.5 Cost estimate

Once we selected 2 products, we asked SEWA to estimate how much it would cost them to produce each product. The following table lists their estimate of the risk portion of their premium and their administrative costs. These are the results of their internal calculations and based on their prior experience selling (a different) hospi-cash policy.

Table 5: *Estimated loss and administrative cost to SEWA of each hospi-cash product in main study.*

	Premium components (INR)	
	Financial loss (risk)	Administration
Product 1 (1+0, INR 500)	137	163
Product 2 (2+3, INR 600)	530	245

B Pilots 2 (pre-COVID stoppage) and 3 (post-COVID stoppage): Optimal selection of prices

We selected prices for the hospi-cash policies selected in Pilot 1 in two steps. First, the Pilots 2 and 3 generated estimates of a distribution of willingness to pay (WTP) for the 2 products selected in the first pilot. We did not merely use stated WTP from Pilot 1 because we worried that individuals will overstate WTP unless they have money on the line. So, to elicit estimates of WTP, in Pilot 2 (pre-COVID stoppage) we conduct a Becker-DeGroot-Marshak (BDM) exercise. We conducted Pilot 3 despite having incentivizeded WTP maeasures from Pilot 2 because, in part, we wanted to see if WTP changed after the first 2 COVID waves and, in part, we wanted to adjust the design of the BDM to account for the fact that we will offer 2 products jointly.

Second, we designed an algorithm to choose optimal prices for our study given our plan to use the strategy in Tebaldi et al. (2019) to estimate bounds on the demand curve. In brief, the algorithm used as an input our prior on the demand curve and chose prices so as to minimize an increasing function of the width of bounds on the demand curve given a set of offer prices. Section B.1 describes Pilot 2 and 3, and our derivation of a prior on demand. Section B.2 describes our algorithm for mapping this prior into four optimal price combinations for our experiment.

B.1 Pilots to generate prior on choice shares at various prices

B.1.1 Sample (Pilot 2)

Whereas in Pilot 1 we selected 50 households per village to be in our sample, we selected 60 households per village to be in our sample for Pilot 2 to ensure that we have a high enough conversion rate (i.e., to ensure enough people enroll in insurance once they become eligible) to adequately calculate prices from this pilot. As with pilot 1, we assembled a primary and backup sample of households for pilot 2. Below we explain how we assembled these samples.

In Ahmedabad pilot 2 villages, we partitioned SEWA households in each village into age quintiles and then into income deciles within each quintile, resulting in 5 $\times$ 12 partitions in each village. We selected the household nearest each income decile within each age quintile, yielding 60 main sample households per village. We used the same process to select randomly an additional 7 to 10 households for the backup sample for each age quintile, which differed from our approach in assigning backup households in pilot 1. We constructed the backup sample in this manner to ensure our backup sample also spans the income distribution within each age quintile.[21] This left us with 19 to 22 households for each age quintile across both primary and backup samples. In Gandhinagar pilot 2 villages, we randomly chose between 15 and 28 households within each age quintile. We randomly assigned 12 from each age quintile to the main pilot 2 sample and 3 to 16 of them to the backup sample for each quintile.

In total, this yielded 300 households in the main pilot 2 sample and 225 households in the backup pilot 2 sample. We were unable to assign exactly nine backup households[22] to each age quintile in each village in the pilot 2 sample, as several age quintiles did not have the requisite number of households (i.e., 21). We therefore oversampled in age quintile-village pairs that have excess households and assigned these oversampled households to age quintile-village pairs that did not have nine backups. We selected households to survey from these two samples in the same manner we did in Pilot 1. If all of the backup households in a particular age quintile-village pair were used, surveyors proceeded to backup households in other age quintiles within the same village. If there were none, surveyors proceeded to the oversampled households assigned to each age quintile-village pair.

B.1.2 Sample (Pilot 3)

Our sampling procedure in pilot 3 is similar to our procedure in pilot 2; however, we selected 50 members per village instead of 60 to reduce costs and finish more quickly. Contrary to pilot 2, we divided selected Ahmedabad villages into age terciles and within each tercile, income terciles, leaving us with nine age $\times$ income partitions. In Gandhinagar villages, we divided SEWA households in a village into age terciles only. We used fewer partitions in pilot 3 than in the earlier pilots to avoid running out of sample in particular partitions, which held up surveyor progress in the earlier two pilots. Approximately 6 main sample members were randomly selected from each partition in each of the four villages. The program used to construct the sample had a stopping protocol when exactly 200 members were selected, meaning some partition-village pairs had fewer than six

[21] In pilot 3, we simply assigned all households not assigned to the main sample in a partition to the backup sample, as running out of backup households in pilot 2 due to this backup assignment method posed a challenge for surveyors.

[22] Our goal was to initially have 3/4 the number of main sample households as backup households in each partition. The 3/4 amount was arbitrarily chosen.

members selected, while others had slightly more. All households that weren't selected by the program were assigned to the backup sample. If all backup households in a particular partition were used, surveyors proceeded to backup households in other partitions within the same village. If there were none, surveyors proceeded to survey excess backup households in other villages.

B.1.3 Pilot 2 (pre-COVID stoppage): Eliciting willingness to pay via BDM

We offered each household one of the two selected products for sale, though the sale was conducted via the BDM random auction method. That means that a household bid an amount for the product; we randomly draw a price for the product; and, if the bid amount was higher than the price draw, then the household had to purchase the product at the bid price. The price schedules from which respondents selected their WTP covered the 10^{th} to 90^{th} percentiles of the WTP distributions elicited during pilot 1 for our two products of interest. The prices we randomly drew to complete the BDM exercise were uniformly distributed in increments of Rs. 50 between the 10^{th} and 90^{th} percentiles of the distributions of individuals' subjective WTP elicited during the first pilot for the 1+0/Rs. 500 and 2+3/Rs. 600 products.

The first step in implementing the pilot was to ensure households have information on each product. In the interval between the conclusion of Pilot 1 and prior to the implementation of the BDM exercise in Pilot 2, VimoSEWA conducted education training exercises for SEWA members in villages assigned to Pilot 2. This is a regular practice followed by VimoSEWA and done to ensure policy-holders understand the concept of insurance and the benefits that accrue to them.

Second, we randomly assigned each household to one of the two hospi-cash products selected in Pilot 1.

Third, for each household and its assigned product, we did the following:

1. Teach the household how the BDM mechanism works by playing it with a bar of soap.
2. Elicit the household's WTP for the product in the bin to which it is assigned via a BDM exercise.
3. Ask $\sim$ 40 questions about family composition, income, home quality, other assets, insurance subscriptions and hospital use.

Neither the study nor VimoSEWA offered any other product for sale to these households prior to Pilot 2, a fact that was made clear to households before conducting the BDM exercise. Members whose bid was above the price draw had until a fixed date (December 6, 2019 in Pilot 2) to pay for their hospi-cash policy. We gathered data on which purchase-eligible households actually paid for their policy. Households who did not pay do not obtain coverage were not permitted to buy the offered hospi-cash policy (or any other hospi-cash polict) from VimoSEWA.

B.1.4 Pilot 3 (post-COVID stoppage): Two-products BDM

Pilot 3 was designed to estimate the joint distribution of WTP for the 2 products selected in pilot 1. There were 2 reasons for conducting pilot 3 notwithstanding the completion of pilot 2. First, we needed to collect data following the 2020 COVID-19 lockdown, to account for the possibility that the pandemic affected (the distribution of) preference for hospi-cash products. Second, we wanted to fix the design of the pilot 2 survey to obtain a direct measure of the joint distribution of preferences for the 2 products, rather than information on preferences for only one product at a time in isolation.

Using the results from the BDM exercise in pilot 2, we set the WTP for the 1+0 person (L) coverage product to be bounded between $\underline{V}^L = 0$ and $\overline{V}^L = 3150$, and the WTP for the 2+3 person (H) coverage product likewise to be bounded between $\underline{V}^H = 0$ and $\overline{V}^H = 3150$.

Each individual was randomized in one of two groups, L or H, with equal probability. Individuals randomized in group L, were asked whether they would be willing to pay the midpoint $\frac{\overline{V}^L + \underline{V}^L}{2}$ (rounded to INR 50) for the L product. If they answered yes, they were asked whether they would be willing to pay $\frac{\overline{V}^L + \underline{V}^L}{2} + 50$. The process continued to add 50 as long as they said yes, or they hit the upper bound $\overline{V}^L$. If they said no, they were asked if they would be willing to pay $\frac{\overline{V}^L + \underline{V}^L}{2} - 50$. The process continued to subtract 50 until they said yes, or they hit the lower bound $\underline{V}^L$.

Once this process ended, individuals who were initially randomized in group L were assigned the elicited V_i^L. Analogously, the individuals who were initially randomized in group H were assigned the elicited V_i^H.

Individuals in group L (H) were then asked whether they would be willing to pay $V_i^L + 50$ ($V_i^H - 50$) for the opposite product H (L). If they said yes (no), they were asked if they would be willing to pay $V_i^L + 100$ ($V_i^H - 100$), and so on and so forth until they said no (yes), or they hit the upper (lower) bound.

At the end of the process, all individuals were assigned the elicited pair (V_i^L, V_i^H), with $V_i^L \leq V_i^H$. Importantly, this design did not impose statistical independence between V_i^L and V_i^H.

Lastly, for every individual i a bivariate chit p_i was drawn i.i.d. uniformly from the support

$$
\begin{aligned}
p_i^L &\in [\underline{V}^L, \overline{V}^L] \\
p_i^H &\in [\underline{V}^H, \overline{V}^H] \\
p_i^L &\leq p_i^H.
\end{aligned}
$$

If $p_i^L \leq V_i^L$ and $p_i^H \leq V_i^H$, the individual was allowed to purchase either L and H, paying the drawn chit price (e.g., p_i^L if selecting L). If $p_i^L > V_i^L$ and $p_i^H > V_i^H$, the individual was not allowed to purchase any product. If $p_i^L \leq V_i^L$ and $p_i^H > V_i^H$ ($p_i^L > V_i^L$ and $p_i^H \leq V_i^H$), the individual was allowed to purchase L (H), paying the price p_i^L (p_i^H).

The data from pilot 3 consisted of a list of elicited WTP pairs (V_i^L, V_i^H), one for each i. For every price pair $P = (P^L, P^H)$ in the grid $[0, 1600] \times [0, 1600]$ (in steps of 25), we computed the predicted prior shares at P as

$$\sigma^L(P) = \frac{\sum_i \mathbf{1}[V_i^L - P^L > V_i^H - P^H \text{ and } V_i^L - P^L \geq 0]}{\sum_i \mathbf{1}} \tag{23}$$

$$\sigma^H(P) = \frac{\sum_i \mathbf{1}[V_i^L - P^L \leq V_i^H - P^H \text{ and } V_i^H - P^H \geq 0]}{\sum_i \mathbf{1}} \tag{24}$$

$$\sigma^0(P) = 1 - \sigma^L(P) - \sigma^H(P). \tag{25}$$

B.2 Selecting Prices for the Main Experiment

In our experiment, we randomize villages over a set of 4 price vectors, $\boldsymbol{p} = \{P^1, P^2, P^3, P^4\}$, and measure the resulting choices. The data will then take the form $(\boldsymbol{p}, \boldsymbol{s})$, as in Tebaldi et al. (2019) (TTY): every price vector in $\boldsymbol{p}$ is associated to the corresponding choice shares, which we collect in $\boldsymbol{s}$.[23]

A key design decision was the determination of the four prices $\{P^1, P^2, P^3, P^4\}$ in Table 3. We selected these prices to satisfy the following criteria:

1. Each price vector P needed to satisfy the following constraints given by SEWA, where $\sigma^0(P)$ is the expected prior share of individuals not buying either product, defined in equation (25):

$$\begin{aligned} 1 - \sigma^0(P) &\geq 0.2 \\ P^L &\geq 125 \\ P^H &\geq 475 \\ P^L &< P^H \\ P^H/P^L &< 6. \end{aligned} \tag{26}$$

   This defined the set of feasible prices for the experiment, that here we denote $\mathcal{P}$. The elements of this set correspond to the black dots in Figure 2.

2. Prices were then chosen to minimize the (expected) average size of the identified set at 10 unobserved price vectors $\boldsymbol{p}^\star = \{P^{1\star}, P^{2\star}, \ldots, P^{10\star}\}$, where the expectation was based on the prior derived from pilot 3. We selected the price vectors $\boldsymbol{p}^\star$ as 9 equally distanced points (3-by-3 grid) in the interior of $\mathcal{P}$ (defined in (26)). The tenth price vector in $\boldsymbol{p}^\star$ was set to be

[23] TTY show how to estimate sharp nonparametric bounds on the share of individuals choosing a product j at the unobserved price vector $p^\star$ when the data consists of a collection $(\boldsymbol{p}, \boldsymbol{s}) \equiv \{p, s(p)\}_{p \in \boldsymbol{p}}$, where $p \in \mathbb{R}^J$ is a price vector and $s(p)$ is the vector of choice shares realized at p, one for each option in the set $\{0, 1, ..., J\}$. The approach introduced by TTY defines, for every product j, a pair of functions $\theta_j^{UB}(p^\star; \boldsymbol{p}, \boldsymbol{s})$ and $\theta_j^{LB}(p^\star; \boldsymbol{p}, \boldsymbol{s})$ such that the interval $[\theta_j^{LB}(p^\star; \boldsymbol{p}, \boldsymbol{s}), \theta_j^{UB}(p^\star; \boldsymbol{p}, \boldsymbol{s})]$ is the identified set for the probability that individuals choose j when facing the price vector $p^\star$.

Figure 2: *Selection of Price Vectors for Main Experiment*

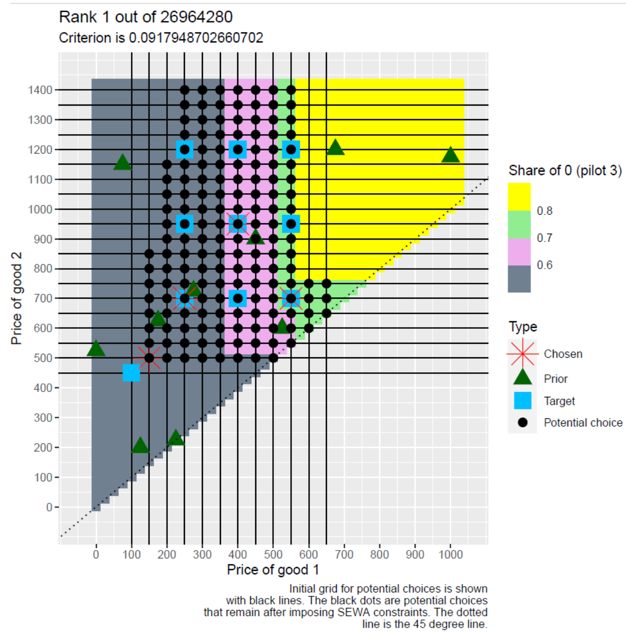

just outside the South-East vertex of $\mathcal{P}$. The 10 price vectors in $\boldsymbol{p}^\star$ correspond to light-blue squares in Figure 2.

B.2.1 Minimizing the Expected Size of the Identified Set

To select $\boldsymbol{p}$ we took as given the constraints in (26) defining the feasible set $\mathcal{P}$, and the 10 price vectors in $\boldsymbol{p}^\star$. We then proceeded as follows.

1. We selected 10 prices and shares from pilot 3, using K-means clustering of prices and prior shares. We collect these 10 pairs in $(\tilde{\boldsymbol{p}}, \tilde{\boldsymbol{s}})$. When selecting the prices for the main experiment we impose that preferences are consistent with $(\tilde{\boldsymbol{p}}, \tilde{\boldsymbol{s}})$.

   The 10 prices in the collection $\tilde{\boldsymbol{p}}$ correspond to the green triangles in Figure 2.

2. We defined, for every $\ell = 1, 2, ..., 10$, and every $j = 0, L, H$, the following function:

$$\lambda_j^\ell(\boldsymbol{p}, \boldsymbol{s}; \tilde{\boldsymbol{p}}, \tilde{\boldsymbol{s}}) \equiv \theta_j^{UB}(P^{\ell\star}; (\boldsymbol{p}, \tilde{\boldsymbol{p}}), (\boldsymbol{s}, \tilde{\boldsymbol{s}})) - \theta_j^{LB}(P^{\ell\star}; (\boldsymbol{p}, \tilde{\boldsymbol{p}}), (\boldsymbol{s}, \tilde{\boldsymbol{s}})), \tag{27}$$

where

$$\theta_j^{UB}(P^{\ell\star};(\boldsymbol{p},\tilde{\boldsymbol{p}}),(\boldsymbol{s},\tilde{\boldsymbol{s}})) = \max_{f\in\mathcal{F}} \sigma_j(P^{\ell\star};f) \tag{28}$$
$$s.t.$$
$$\sigma(p;f) = s(p) \ \ \forall\, p \in \boldsymbol{p}$$
$$\sigma(\tilde{p};f) = \tilde{s}(\tilde{p}) \ \ \forall\, p \in \tilde{\boldsymbol{p}},$$

and

$$\theta_j^{LB}(P^{\ell\star};(\boldsymbol{p},\tilde{\boldsymbol{p}}),(\boldsymbol{s},\tilde{\boldsymbol{s}})) = \min_{f\in\mathcal{F}} \sigma_j(P^{\ell\star};f) \tag{29}$$
$$s.t.$$
$$\sigma(p;f) = s(p) \ \ \forall\, p \in \boldsymbol{p}$$
$$\sigma(\tilde{p};f) = \tilde{s}(\tilde{p}) \ \ \forall\, p \in \tilde{\boldsymbol{p}}.$$

In the above, $\mathcal{F}$ collects all possible densities of willingness to pay for the two products L and H, and

$$\sigma_j(p;f) \equiv \int \mathbf{1}\left[V_j - p_j \geq V_k - p_k \text{ for all } k\right] f(V)dV$$

is the share of individuals purchasing j when facing the price vector p; $\sigma(p;f) = \{\sigma_j(p;f)\}_j$.

TTY proves the equivalence between the solutions of the problems in (28) and (29) and the solutions of (possibly large) finite, linear programs. This ensures computational tractability while preserving precision.

3. The value of $\lambda_j^{\ell}(\boldsymbol{p},\boldsymbol{s};\tilde{\boldsymbol{p}},\tilde{\boldsymbol{s}})$ is the width of the identified set for the share of individuals purchasing j when the price vector is $P^{\ell\star}$, if the data consists of $(\boldsymbol{p},\tilde{\boldsymbol{p}}),(\boldsymbol{s},\tilde{\boldsymbol{s}})$.

   We selected the 4 prices for the experiment to minimize the average width of the identified set, imposing the prior prices and shares, while following a "worse case scenario" approach for the shares that will be observed during the main experiment.

   Formally, the 4 prices in Table 3 solved

$$\min_{\boldsymbol{p}\in\mathcal{P}} \max_{\boldsymbol{s}} \sum_j \sum_{\ell} \lambda_j^{\ell}(\boldsymbol{p},\boldsymbol{s};\tilde{\boldsymbol{p}},\tilde{\boldsymbol{s}}). \tag{30}$$

   The final prices rank first out of 26,964,280 possible collections of four price vectors in the set $\mathcal{P}$. These are represented by the four red stars in Figure 2.

References

Attanasio, O. P. and Székely, M. (2004). Wage Shocks And Consumption Variability In Mexico During The 1990s. *Journal of Development Economics*, 73(1):1–25. 17, 22

Berman, P., Ahuja, R., and Bhandari, L. (2010). The impoverishing effect of healthcare payments in india: new methodology and findings. *Economic and Political Weekly*, pages 65–71. 4

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, pages 242–262. 16, 18

Dobkin, C., Finkelstein, A., Kluender, R., and Notowidigdo, M. J. (2018). The economic consequences of hospital admissions. *American Economic Review*, 108(2):308–52. 4

Einav, L., Finkelstein, A., and Cullen, M. R. (2010a). Estimating welfare in insurance markets using variation in prices. *The quarterly journal of economics*, 125(3):877–921. 6, 17, 19

Einav, L., Finkelstein, A., and Levin, J. (2010b). Beyond testing: Empirical models of insurance markets. *Annu. Rev. Econ*, 2:311–36. 17, 19

Geruso, M., Layton, T. J., McCormack, G., and Shepard, M. (2019). The two margin problem in insurance markets. *The Review of Economics and Statistics*, pages 1–46. 17, 19

Malani, A., Holtzman, P., Imai, K., Kinnan, C., Miller, M., Swaminathan, S., Voena, A., Woda, B., and Conti, G. (2021). Effect of health insurance in india: A randomized controlled trial. Working Paper 29576, National Bureau of Economic Research. 14

McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior. 16, 18

Shahrawat, R. and Rao, K. D. (2012). Insured yet vulnerable: out-of-pocket payments and india's poor. *Health policy and planning*, 27(3):213–221. 4

Singh, C. H. and Ladusingh, L. (2010). Inpatient length of stay: a finite mixture modeling analysis. *The European Journal of Health Economics*, 11(2):119–126. 24

Tebaldi, P., Torgovitsky, A., and Yang, H. (2019). Nonparametric estimates of demand in the california health insurance exchange. Technical report, National Bureau of Economic Research. 1, 6, 7, 9, 13, 16, 17, 20, 28, 32

Walker, V. M., Davies, N. M., Windmeijer, F., Burgess, S., and Martin, R. M. (2017). Power calculator for instrumental variable analysis in pharmacoepidemiology. *International journal of epidemiology*, 46(5):1627–1632. 28575313[pmid] PMC5837396[pmcid] 3858437[PII]. 14

World Health Organization (2019). Global health expenditure database. 4