Maua RCT
Pre-Analysis Plan

Francesco Cordaro[*], Marcel Fafchamps[†], Colin Mayer[‡],
Muhammad Meki[§], Simon Quinn[¶], and Kate Roll[‖]

Fieldwork location:	Kenya
Fieldwork dates:	From January 2018 (ongoing rolling baseline)
Date of Pre-Analysis Plan:	26 September 2019

1 Introduction

This document outlines our pre-analysis plan for an ongoing experiment that provides asset-based microfinance contracts to microentrepreneurs in Kenya. The document summarises our experiment, our data and our plan of regressions. We intend to submit this Pre-Analysis Plan to the AEA RCT Registry.

2 Sampling

2.1 Study context

We collaborate with one of the largest multinationals in the world, Mars Corporation (owners of the Wrigley Company), to offer microfinance contracts to fund the purchase of a bicycle for microentrepreneurs who work with Wrigley's micro-distribution programme in Kenya. Our sample consists of microdistributors, many of whom sell retail products (which include chewing gum) on

[*]Mars, Incorporated.
[†]Freeman Spogli Institute, Stanford University.
[‡]Saïd Business School, University of Oxford.
[§]Department of Economics, Pembroke College and CSAE, University of Oxford.
[¶]Department of Economics, St Antony's College and CSAE, University of Oxford.
[‖]UCL Institute for Innovation and Public Purpose.

foot, and who would like to finance the purchase of a bicycle to expand their micro-distribution activities. This setting is common to many 'route-to-market' distribution programs run by multinational corporations around the world. The project is part of Mars Corporation's 'Mutuality in Business' research program, in collaboration with the Said Business School in Oxford. We work with these microdistributors from Wrigley Kenya's route-to-market program, who purchase chewing gum from stock points (small warehouses). Stock points receive deliveries of Wrigley chewing gum, which they sell alongside various non-Wrigley products. Microdistributors purchase chewing gum as well as other products from stock points, before selling it to retailers, who are often small kiosks. Microdistributors initially purchase the gum at a discount per bag of gum (approximately US$ 0.05) to the market price (an up-front margin); and, for every bag of gums sold, they receive an end-of-month bonus of via mobile money (M-Pesa). There is no obligation for them to exclusively sell gum, but selling Wrigley's product is relatively profitable, and they have a strong incentive to stay in the program.[1]

Qualitative work with microdistributors revealed the two major constraints on their business to be 'materials and equipment', which were interpreted as working capital and assets. In this project, we finance the purchase of a productive asset; our sample consists of individuals who expressed an interest in expanding their business by purchasing a bicycle. We elicit preferences over different microfinance contracts to finance the cost of a bicycle, with contracts varying in the extent to which they feature 'mutuality' in the sharing of risk and reward, with one of the contracts being subsequently offered to a randomly selected subset of participants. We describe the details of the intervention in the following sections. We currently have 152 microdistributors enrolled in the project, and our target sample size is 250.

2.2 Baseline workshop structure

Microdistributors are invited to a workshop, where they complete a baseline survey, which includes questions asking about individual and household characteristics, household finances (incomes, expenditures, savings and loans), business performance (revenues, profits, expenditures and assets), business management practices, and a number of questions on social capital and health. Following the survey, all microdistributors participate in a set of detailed behavioural games, designed to

[1] Stock points also sell gum to people other than Wrigley microdistributors, but such sales are not eligible for an up-front discount or end-of-month bonus.

measure risk preferences, loss aversion, time preferences, and cognitive ability.[2] There are two risk preference elicitation activities; the first is a survey-based measure, using a series of questions that seek to gauge respondents' risk-taking in their occupation, financial matters, and faith in others.[3] The second measure of risk preferences is incentivised, where respondents are posed a series of 30 questions that require them to choose between a certain amount of money or an uncertain investment option, which has two possible outcomes: (i) a 'bad' outcome, with a payoff of zero; or (ii) a 'good' outcome, with a payoff of KES 1,000.[4] In the loss aversion activity, respondents are offered a series of binary-outcome investment choices that involve a large positive outcome or a (gradually increasing) negative outcome, which they can accept or reject. If they accept the investments and the loss aversion activity is chosen for payment at the end of the workshop,[5] then a realised loss is taken out of their guaranteed workshop participation fee; as such, this represents a potential real loss.[6] In the time-preference elicitation activity, individuals are offered a series of choices between an amount of money paid on the same day as the workshop or (gradually increasing) amounts of money one month from the workshop.[7] After the survey and behavioural games, participants have the different microfinance contracts explained to them, and then have their preferences over these elicited using an incentivised activity, described in the next section.

[2] All data collection takes place electronically, using tablet computers and uploading data to a server (SurveyCTO).

[3] Specifically, the questions are: *(i) "How would you rate your willingness to take risks in financial matters?"; (ii) "How would you rate your willingness to take risks in your occupation?"; (iii) "How would you rate your willingness to take risks when it comes to having faith in other people?"; and (iv) "How do you see yourself? Are you generally a person who is fully willing to take risks or do you try to avoid taking risks?"*. Responses are given on a scale of 1 to 10, with 0 representing 'risk-averse' and 10 representing 'fully prepared to take risks'. The questions were adapted from Dohmen, Falk, Huffman, Sunde, Schupp, and Wagner (2011), have been used by other researchers in a number of settings and have demonstrated a reasonably strong correlation with incentivised measures of risk preferences.

[4] We adapted the measures used by Barr and Packard (2002) and Vieider, Lefebvre, Bouchouicha, Chmura, Hakimov, Krawczyk, and Martinsson (2015).

[5] Before conducting all activities, participants were informed that, at the end of the behavioural games session, one of the incentivised activities would be selected for payment by physically drawing a ball from a bag. Within the selected activity, balls would be drawn to select the one final question that would be used for payment. As such, participants were required to answer all questions attentively, because any question could have been selected. This method also allowed the use of payment amounts that were relatively large, with the average payment being approximately three times as large as median daily business profits for microenterprises in the sample. From a methodological perspective, Charness, Gneezy, and Halladay (2016) show that paying for only a (randomly selected) subset of all activities is at least as effective as paying for all of them, and can actually be more effective in terms of helping to avoid wealth effects and hedging within the behavioural games session.

[6] We adapted the loss aversion measure used by Bartling, Fehr, and Herz (2014).

[7] The time-preference activity was also conducted using a 'far frame', where money was offered one month forward versus two months forward.

3 Contract structure, elicitation and assignment

After the survey and preference elicitation, participants are introduced to a manager from the microfinance organisation (Longitude Finance), who explains that they will be offering the financing contracts for bikes to a randomly selected subset of participants. There are five possible treatment contracts, which vary in the extent to which they contained 'debt-like', 'equity-like' or 'insurance-like' features:[8]

(i) **`Debt_1`**: A 12-month, fixed-repayment debt contract, based on a 15% mark-up of the financed amount, and with an option to pay more than one month's payment in any given month (and to subsequently reduce the contract duration by the number of voluntary payments);

(ii) **`Debt_2`**: A 12-month, fixed-repayment debt contract, based on a 15% mark-up of the financed amount, and with no early repayment option;

(iii) **`Equity_1`**: A flexible-duration contract that requires half of the fixed monthly payment of `Debt_2`, and additionally requires a 10% share of the microdistributor's monthly profits, with the contract ending as soon as cumulative payments reach the total payment due under `Debt_2`;

(iv) **`Equity_2`**: A 12-month contract that has half of the fixed monthly payment of `Debt_2`, and additionally requiring a 10% share of the microdistributor's monthly profits;

(v) **`Insurance`**: An index insurance contract, which has a similar repayment structure to `Equity_2`, with the difference being how the profit-sharing payments are calculated: the 10% payments are based on an index constructed from the profits of other microdistributors in their region.

Profits are calculated using administrative data from Mars on stock purchases, and the profit margin is based on the recommended retail price for each chewing gum product. Our hypothesis is that 'equity-like' contracts, through performance-contingent payments, have the potential to benefit microentrepreneurs through the provision of implicit insurance, by automatically reducing repayment requirements when business conditions are challenging. We confirmed the potential importance of

[8] When communicating with participants, the words 'debt', 'equity' and 'insurance' are never used; contracts are explained using their cash-flow structure in local language (Swahili), with each contract colour-coded for ease of remembering.

such a mechanism by analysis of variance of administrative data on incomes of all microdistributors in the Wrigley program, which revealed the importance of both within-microdistributor as well as between-microdistributor variation. However, this benefit of performance-contingent payments could also in principle be achieved by an index insurance contract, which would not directly tax an microdistributor's income. Index insurance may also be preferred to our 'equity-based' contracts, if common shocks are the primary concern for microdistributors. Therefore, to provide a sharper test of the value of our 'equity-like' contracts, we added the `Insurance` contract to the initial set of contracts that we tested in an earlier pilot (where we only implemented contracts `Debt_2`, `Equity_1` and `Equity_2`). For the `Insurance` contract, we explored different levels of aggregation, and decided on a region-based index (for the five major regions in Kenya: Nairobi, Central Kenya, Kisumu, Eastern Kenya, Mumbasa). We construct the index based on the average profits of every other microdistributor in the region for a particular month, excluding the microdistributor's own stock point, to mitigate the risk of manipulation. A historical analysis of the correlation between the constructed index and microdistributor profits reveals lower basis risk than in many other studies that have used index insurance products.[9] As we will explain below, the `Debt_1` contract is one for which we elicit a preference, but assign a very low probability of actual allocation. One possible benefit of the `Equity_1` contract, which has a flexible duration, is in the ability to end the contract early by making higher payments when performance is greater. This 'savings-like' element may be more important than the 'insurance-like' element of performance-contingent repayments. As such, to provide a sharper test of the value of the `Equity_1` contract relative to the fixed-repayment 12-month `Debt_2` contract, we added the `Debt_1` contract that was identical to the `Debt_2` contract except that microdistributors were free to make early repayments in multiples of the monthly amount. While it was important to investigate demand for this contract, in order to simplify the logistics of the experiment and the contracts being implemented in the field, we assigned a very low probability to this contract being drawn for microdistributors.

A strategy method is used to elicit the contract preferences for each microdistributor. Five enumerators wear coloured shirts, representing the different contracts. Each microdistributor is required to state whether they would accept the contract if they were to draw that coloured ball from the bag during the final randomisation, and they are permitted to accept or reject as many of the contracts as they desire; it is explained that their decisions do not influence their allocation of contracts.

[9] See Cole, Giné, Tobacman, Topalova, Townsend, and Vickery (2013); Carter, de Janvry, Sadoulet, Sarris, et al. (2014); Clarke (2016).

After recording decisions from every microdistributor for each of the five contracts, a public randomisation is conducted using a bag containing 101 balls. The bag contains 20 grey, blue, pink and yellow balls, as well as 20 green balls for a pure control group, who receive no contract offer. Only one ball is coloured white, representing contract `Debt_1`, to make it possible but unlikely to be drawn in practice.[10] Microdistributors who draw a colour for which they had pre-specified their acceptance are immediately directed to a representative from the MFI, who is present and proceeds with signing of contracts.[11]

All contracts require the microdistributor to pay an initial deposit of 10%. The remaining 90% of the bicycle price is financed by our MFI partner, Longitude Finance. For example, assuming a bicycle price of KES 9,275, the microdistributor is required to provide a 10% deposit (KES 928), with the remaining amount (KES 8,347) being financed. For the `Debt_2` contract, payments are calculated by simply increasing the financing amount by 15% (to KES 9,600), which would need to be paid in equal monthly instalments of KES 800 for 12 months. We did not use compound interest, which greatly facilitates explanation of total payments. The `Equity_2` contract also has a fixed-repayment component, which is half of that of the `Debt_2` contract (KES 400); in addition, microdistributors are required to pay 10% of their Wrigley income per month, which we proxy for using their observable stock purchases, which are measured by Wrigley for the purposes of computing microdistributors' monthly bonuses. The `Equity_1` contract has exactly the same repayment structure as the `Equity_2` contract, with the exception that it terminates once cumulative payments reach KES 9,600; as such, it can be considered as a hybrid between a 'debt-like' and an 'equity-like' contract. We designed the contracts such that an microdistributor with an income of KES 4,000 per month would have the same repayments and contract duration under all three contracts; this represents the 56th percentile of the income distribution for microdistributors from our pilot, and provides a relatively fair 'horse race' between the contracts, abstracting from risk.

[10] Microdistributors are not informed of the numbers of balls in the bag.

[11] The choice of bicycle is made before contract decisions are elicited; most bikes are 'work friendly' bikes with a rear rack. The menu of bikes includes one model that is of a higher quality and nearly twice as expensive, one that is female-friendly with a dipped bar, and one mountain bike.

4 Data

4.1 Construction of outcome variables

We define the following outcome variables:

DEFINITION	SOURCE (QUESTIONNAIRE CODE)	NEW NAME?
SALES (ADMINISTRATIVE DATA)		
Monthly sales	Taken from the administrative data	`admin_sales`
BUSINESS PERFORMANCE		
Whether their primary activity is as a seller / distributor	`occupation==1`	`biz_seller`
Whether their primary activity is as owner of another (non-distribution) business	`occupation==2`	`biz_other`
Current number of employees/apprentices in primary business (excluding self)	`biz_emps`	
Sales in the previous 30 days (distribution business, all income sources)	`biz_revL1`	
Profits in the previous 30 days (distribution business, all income sources)	`biz_profL1`	
Hours they spend selling (typical week)	`selling_days * selling_hours`	`selling`
HOUSEHOLD FINANCES		
Household income (previous 30 days)	`hinc_bizL1 + hinc_caslabL1 + hinc_fxdwagL1 + hinc_remittL1 + hinc_otherL10`	`hinc`
Household expenditures (previous 30 days)	`hexp_billsL1 hexp_foodL1 + hexp_clothesL1 + hexp_hhitemsL1 + hexp_schoolL1 + hexp_healthL1 + hexp_transportL1 + hexp_temptL1 + hexp_specialL1 + hexp_otherL1`	`hexp`

Household savings[12]	`sav_bank + sav_depos + sav_cash + sav_animal + sav_loan + sav_oth`	`hsav`
Household debt	`loan_bank_oweamt + loan_family_oweamt + loan_other_oweamt`	`hlon`
Value of household assets[13]	`hh_assets1_v + ... + hh_assets21_v`	`hast`

OTHER EARNINGS (RESPONDENT)

Proportion of microdistribution profits that was from Wrigley products (last 3 months)	`(maua_prop / 10 ) - 0.05`	`mauaprp`
Dummy: Has other source of income (in addition to distribution)	`wg_have`	
Number of hours per week in other work	`wagemp_hours` × `wagemp_days`	`wg_hrs`
Monthly wage income	`wg_amt`	

BUSINESS MANAGEMENT PRACTICES

Score for management practices (weighted using covariance matrix from the control group, as in Anderson (2008))	Weighted sum of all variables listed in the following four rows	`mp_all`
Score for marketing practices (weighted using covariance matrix from the control group, as in Anderson (2008))	Weighted sum of `bizmp_1`, `bizmp_2`, `bizmp_3` and `bizmp_11` (`bizmp_11` recoded so that 1, 2 or 3 is set to 1, and set to 0 otherwise)	`mp_mark`
Score for negotiation	`bizmp_4`	
Score for costing and record-keeping practices (weighted using covariance matrix from the control group, as in Anderson (2008))	Weighted sum of `biz_5` to `biz_8`	`mp_rcrd`
Keeping a sales target	`bizmp_9`	

SAVINGS / SPENDING PROBLEMS

Dummy: Has problems saving	`bhvrl_1=4 | bhvrl_1=5`	`s_svprb`

[12] This does not include any savings or debts from participation in ROSCAs/'committees'.

[13] This does not include any estimate for the value of land or property, and we only have explicit valuations from the 6m survey (before that, we have quantities of each asset but not value).

Dummy: Makes unecessary purchases	`bhvrl_2=1 | bhvrl_2=2 | bhvrl_2=3`	`s_unpur`
Dummy: Does things early to prevent forgetting	`bhvrl_7=1 | bhvrl_7=2 | bhvrl_7=3`	`s_timedo`
Dummy: Pays things early to prevent forgetting	`bhvrl_10=1 | bhvrl_10=2 | bhvrl_10=3`	`s_timepay`
Dummy: Families asks them for spare money	`bhvrl_3=4 | bhvrl_3=5`	`s_kintax`

LOCUS OF CONTROL, HAPPINESS AND TRUST

Dummy: believes that achievements are firstly a question of destiny and luck	`control_2 = 4 | control_2 = 5`	`locus1`
Dummy: doubts themselves when they encounter difficulties	`control_3 = 4 | control_3 = 5`	`locus2`
Trust and community relations (weighted using covariance matrix from the control group, as in Anderson (2008))	Weighted sum of `trust_1` to `trust_5`	`trust`
Satisfaction with income and ability to meet expenditure demands (weighted using covariance matrix from the control group, as in Anderson (2008))	Weighted sum of `wellbeing_1` and `wellbeing_2`	`happy1`
Dummy: Satisfied with materials and equipment at work	`wellbeing_3 = 1 | wellbeing_3 = 2`	`happy2`

HEALTH

Dummy: bad general physical health	`health_1 = 4 | health_1 = 5`	`health1`
Dummy: moderate or worse bodily pain (last 3 months)	`health_2 = 3 | health_2 = 4 | health_2 = 5`	`health2`
Dummy: physical pain caused or made worse by work	`health_3 = 1 | health_3 = 2`	`health3`
Dummy: physical health has made work difficult (last 3 months)	`health_5 = 1`	`health4`

4.2 Construction of other variables

For the purpose of testing balance and heterogeneity analysis, we define the following other variables, relating to individual and household characteristics:

DEFINITION	SOURCE (QUESTIONNAIRE CODE)	NEW NAME?
OTHER VARIABLES		
Dummy: Respondent is female	Dummy: (`gender` == 1)	`resp_gender`
Dummy: Respondent is married	Dummy: (`married` == 1)	`resp_married`
Respondent's age	`age`	
Respondent's highest completed level of education	`educ`	
Number of people in the household (including respondent)	`hh_size`	
Number of people in the household earning any form of income	`hh_earners`	
Index of incentivised risk preference elicitation activity	Unweighted sum of `rp1_25_0` to `rp1_25_10`, `rp1_50_0` to `rp1_50_10` and `rp1_75_0` to `rp1_75_10`.	`rp1`
Index of incentivised loss aversion elicitation activity	Unweighted sum of `loss_1` to `loss_10`	`loss`
Index of incentivised time preference elicitation activity	Unweighted sum of `tp1_1` to `tp1_10` and `tp2_1` to `tp2_10`	`tp`
Index of numeracy skills (weighted using covariance matrix from the control group, as in Anderson (2008))	Weighted sum of `math_1` to `math_8`, `math_9` and `math_10` (with each variable recoded so that a correct response to the mathematical question is coded as 1)	`math`

5 Analysis

5.1 Testing balance

We now index contracts by k, as follows:

TREATMENT	DESCRIPTION	k
`Control`	No contract	$k = 0$
`Debt_1`	Fixed-repayment debt with early repayment option	$k = 1$
`Debt_2`	Fixed-repayment debt without early repayment option	$k = 2$
`Equity_1`	Flexible-duration equity	$k = 3$
`Equity_2`	Fixed-duration equity	$k = 4$
`Insurance`	An index insurance contract	$k = 5$

We begin by testing for balance in contractual offers. As explained earlier, the `Debt_1` contract was offered with a very small probability (1/101); it is not meaningful to test for balance in respect of this contract. Therefore, we test for balance across control and the remaining four contracts. To do this, we will run the following regression:

$$y_{i0} = \beta_0 + \sum_{k \in \{2,\ldots,5\}} \beta_k \cdot \texttt{Offered}_{ik} + \varepsilon_{i0}. \tag{1}$$

Here, $\texttt{Offered}_{ik}$ is a dummy for whether individual i had contract k randomly drawn. We run this regression using baseline data, using heteroskedasticity-robust errors, and we will do this for all variables listed in section 4.1 and section 4.2. (For variable `admin_sales`, we will take the *average* monthly sales for the preceding three months, or part of those three months that the respondent was active at the stock-point; *e.g.* if someone was active in two months, we will average only over those two months.) For each variable, we will report (i) the overall sample mean, and (ii) the p-value from a joint test of the null hypothesis of balance: $H_0 : \beta_2 = \beta_3 = \beta_4 = \beta_5 = 0$. We interpret these p-values essentially as descriptive of imbalance; to test for imbalance, we will rely on an overall omnibus test that $\beta_2 = \beta_3 = \beta_4 = \beta_5 = 0$ jointly for all variables tested.

5.2 Contract take-up

We will begin by testing demand for each type of offered contract. Define a_{ic} as a dummy variable for whether individual i agreed to take up contract offer c. We will estimate average take-up rates by using a simple Linear Probability Model, in which we regress take-up on a series of dummy

variables for which contract was being offered in each respective choice:

$$a_{ic} = \sum_{k \in \{1,\ldots,5\}} \beta_k \cdot \mathcal{I}(k = c) + \varepsilon_{ic}, \tag{2}$$

where $\mathcal{I}$ refers to the indicator function.

We will cluster errors throughout at the level of the individual respondent, i. Our regression estimates, $\hat{\beta}_k$, will therefore show the average willingness to accept each type of contract offer. We will report a p-value from a test of the omnibus null hypothesis that these proportions are identical across contracts: $H_0 : \beta_1 = \beta_2 = \beta_3 = \beta_4 = \beta_5$.

We will then repeat this exercise for a series of median splits. Define X_i as a dummy for whether individual i is at or above the median in some given characteristic (to be discussed shortly). Then we will extend equation 2 by estimating the following:

$$a_{ic} = \sum_{k \in \{1,\ldots,5\}} \beta_k \cdot \mathcal{I}(k = c) + \sum_{k \in \{1,\ldots,5\}} \gamma_k \cdot \mathcal{I}(k = c) \cdot X_i + \varepsilon_{ic}, \tag{3}$$

Under this specification, we interpret $\hat{\beta}_k$ as the proportion of respondents having $X_i = 0$ who are willing to agree to contract k. We interpret $\hat{\gamma}_k$ as the difference between this proportion and the proportion of respondents having $X_i = 1$ who are willing to agree to contract k. Under this specification, we anticipate reporting a p-value from a test of the omnibus null hypothesis that these proportions are identical across contracts for respondents having $X_i = 0$ (*i.e.* $H_0 : \beta_1 = \beta_2 = \beta_3 = \beta_4 = \beta_5$), and a separate test of the omnibus null hypothesis that these proportions are identical across contracts for respondents having $X_i = 1$ (*i.e.* $H_0 : \beta_1 + \gamma_1 = \beta_2 + \gamma_2 = \beta_3 + \gamma_3 = \beta_4 + \gamma_4 = \beta_5 + \gamma_5$). Finally, we anticipate reporting a p-value for test of the omnibus null hypothesis that there is no heterogeneity in take-up between individuals having $X_i = 0$ and $X_i = 1$: *i.e.* $H_0 = \gamma_1 = \gamma_2 = \gamma_3 = \gamma_4 = \gamma_5$.

We will estimate equation 3 using six different definitions of $X_i = 1$; we will respectively define X_i as a dummy for whether individual i is at or above the sample median in terms of ...

(i). ...*average* monthly microdistribution income for the preceding three months, taken from administrative data (specifically, `admin_sales` as defined in sections 4.1 and 5.1);

(ii). ...incentivised risk aversion (specifically, `rp1` as defined in section 4.2);

(iii). ...incentivised loss aversion (specifically, `loss` as defined in section 4.2);

(iv). ...incentivised time preference (specifically, `tp` as defined in section 4.2);

(v). ...numeracy skills (specifically, `math` as defined in section 4.2); and

(vi). ...management practices (specifically, `mp_all` as defined in section 4.1).

To increase statistical power, we will also conduct the above analysis using an indicator variable that pools `Equity_1` and `Equity_2` (as described in Section 5.1, these two contracts have an identical month-to-month repayment structure, with the difference being in their duration).

5.3 Consequences of contracts

We will test the consequences of the various contracts. To do this, we will exploit the panel structure of the data: for variable `admin_sales`, we will construct a monthly panel, for the duration of the contracts; for all other outcome variables, we will use our face-to-face data, collected three months, six months, nine months, and twelve months after contract implementation. We also intend to conduct a longer-term follow-up either at the 18-month or 24-month stage, and are currently seeking funding for this. For each outcome, we will estimate two objects of interest.

First, we will test the consequences of being *offered* the various types of contract. To do this, we will use an intent-to-treat ANCOVA specification:

$$y_{it} = \beta_0 + \sum_{k \in \{2,\ldots,5\}} \beta_k \cdot \texttt{Offered}_{ik} + \gamma \cdot y_{i0} + \varepsilon_{it}. \tag{4}$$

Here, $\texttt{Offered}_{ik}$ is a dummy for whether individual i had contract k randomly drawn. In this specification, y_{i0} refers to the baseline value for outcome y (or the average prior outcome, in the case of administrative sales data). We will cluster at the individual level in all cases. Note that we estimate consequences of contracts 2 to 5 only; this is because, as noted earlier, contract 1 had a very small probability of being drawn. We will therefore interpret $\hat{\beta}_2$, $\hat{\beta}_3$, $\hat{\beta}_4$ and $\hat{\beta}_5$ as the estimated effect of having each contract type randomly drawn.

Second, we will test the consequences of *contracting* under each of the various types of potential agreement. This correlates with willingness to agree, which may itself proxy for a variety of unobservable factors. However, given our experimental design, we can fully control for such unobservable heterogeneity by controlling for the pattern of contractual agreement. Specifically, we define a vector of four dummy variables representing the respondent's agreement to contracts 2, 3, 4 and 5: $\boldsymbol{a}_i \equiv (a_{i2}, a_{i3}, a_{i4}, a_{i5})$. Conditional on $\boldsymbol{a}_i$, the realisation of contractual agreement is exogenous:[14]

$$y_{it} = \beta_0 + \sum_{k \in \{2,\ldots,5\}} \beta_k \cdot \texttt{Contracted}_{ik} + \gamma \cdot y_{i0} + \boldsymbol{\delta} \cdot \boldsymbol{a}_i + \varepsilon_{it}. \tag{5}$$

Under this specification, we will then interpret $\hat{\beta}_2$, $\hat{\beta}_3$, $\hat{\beta}_4$ and $\hat{\beta}_5$ as the estimated effect of having contracted under contracts 2, 3, 4 and 5. To increase statistical power, we will also conduct the above analysis using a dummy variable that pools the equity contracts (contracts 3 and 4).

5.3.1 Primary outcome: Administrative data on sales

Following Olken (2015), we begin by defining our primary outcome of interest. In testing the consequences of contracts, our key hypothesis is that our treatments affect participants' sales, as measured using the stock-points' administrative data. Our primary outcome, therefore, is the variable `admin_sales`, as defined in section 4.1. We will use the estimation strategy outlined earlier with `admin_sales` as the outcome variable, and will report the usual p-values from standard Wald tests.

5.3.2 Secondary outcomes

We have seven families of secondary outcomes — namely, each of the families defined in section 4.1 from 'BUSINESS PERFORMANCE' onwards. For each of the outcomes in these families, we will run the estimation and hypothesis tests outlined earlier. For each hypothesis test, we will report two values:

(i). The usual p-value from a Wald test; and

(ii). We will report False Discovery Rate q-values, taken across the family of outcomes (Benjamini et al., 2006).

[14] For maximal flexibility, as a robustness check we can also define $\boldsymbol{a}_i$ as a vector of the 16 possible combinations of a_{i2}, a_{i3}, a_{i4} and a_{i5}.

5.4 Further analysis

We anticipate conducting additional analysis to understand the mechanisms by which our contracts operate. This is likely to include

(i). GPS trackers : These have been installed on all bikes; however, we are currently unsure of the availability and quality of this data, due to having experienced a number of issues with batteries (many of which have completely discharged and have had to be replaced multiple times), particularly in remote areas of Kenya where the GPS signal is weaker;

(ii). Social network analysis: During the baseline workshops we collected (using pen and paper) details of the business-related social network of all microdistributors (specifically, their knowledge of and level of social and financial interaction with all other individuals at their respective stock point). This data may then be used to run a series of dyadic regressions to investigate externalities of the contracts (specifically, the relationship between outcomes under profit-sharing contracts and the outcomes of others in their social network).

References

M. L. Anderson. Multiple inference and gender differences in the effects of early intervention: A reevaluation of the abecedarian, perry preschool, and early training projects. *Journal of the American statistical Association*, 103(484):1481–1495, 2008.

A. Barr and T. G. Packard. Revealed preference and self-insurance: Can we learn from the self-employed in chile? *Working paper*, 2002.

B. Bartling, E. Fehr, and H. Herz. The intrinsic value of decision rights. *Econometrica*, 82(6): 2005–2039, 2014.

Y. Benjamini, A. M. Krieger, and D. Yekutieli. Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*, 93(3):491–507, 2006.

M. Carter, A. de Janvry, E. Sadoulet, A. Sarris, et al. Index-based weather insurance for developing countries: A review of evidence and a set of propositions for up-scaling. *Development Policies Working Paper*, 111, 2014.

G. Charness, U. Gneezy, and B. Halladay. Experimental methods: Pay one or pay all. *Journal of Economic Behavior & Organization*, 131:141–150, 2016.

D. J. Clarke. A theory of rational demand for index insurance. *American Economic Journal: Microeconomics*, 8(1):283–306, 2016.

S. Cole, X. Giné, J. Tobacman, P. Topalova, R. Townsend, and J. Vickery. Barriers to household risk management: Evidence from india. *American Economic Journal: Applied Economics*, 5 (1):104–35, 2013.

T. Dohmen, A. Falk, D. Huffman, U. Sunde, J. Schupp, and G. G. Wagner. Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association*, 9(3):522–550, 2011.

B. A. Olken. Promises and perils of pre-analysis plans. *Journal of Economic Perspectives*, 29(3): 61–80, 2015.

F. M. Vieider, M. Lefebvre, R. Bouchouicha, T. Chmura, R. Hakimov, M. Krawczyk, and P. Martinsson. Common components of risk and uncertainty attitudes across contexts and domains: Evidence from 30 countries. *Journal of the European Economic Association*, 13(3):421–452, 2015.