

Microequity and Mutuality: Experimental Evidence on Credit with Performance-Contingent Repayment*

Francesco Cordaro[†] Marcel Fafchamps[‡] Colin Mayer[§]
Muhammad Meki[¶] Simon Quinn^{||} Kate Roll^{**}

June 11, 2022

Abstract

We conduct the first field experiment of a performance-contingent microfinance contract. A large food multinational wishes to help micro-distributors in its supply chain with the financing of a lumpy productive asset. Working with the firm in Kenya, we compare asset financing under a traditional debt contract to three alternatives: (i) a novel equity-like financing contract, (ii) a hybrid debt-equity contract, and (iii) an index-insurance financing contract. Experimental results reveal large positive impacts from the contractual innovations. These findings demonstrate the economic appeal of microfinance contracts that leverage improved observability of performance to achieve a greater sharing of risk and reward.

*We thank Elizabeth Gatwiri, Elizabeth Githinji, Kevin Kisali, Lina Kwok, Michael Monari, Gerald Mwangi, David Ngigi, and Irene Wanjiru for invaluable assistance in this research. We are grateful to Anik Ashraf, Deniz Aydin, Haseeb Ashraf, Paola Bustos, Jing Cai, Paola Conconi, Zoë Cullen, Joshua Deutschmann, Christopher Eaglin, Seth Garz, Xavier Giné, Jonathan Greenacre, Selim Gulesci, Morgan Hardy, Simon Hess, Michal Hodor, Martin Kanz, Erin Kelley, Michael King, Cynthia Kinnan, Pramila Krishnan, Greg Lane, Mahreen Mahmud, Karol Mazur, Craig McIntosh, Timothy Ogden, Imran Rasul, Tavneet Suri, Alessandro Tarozzi, Christine Valente, Eric Verhoogen, Chris Woodruff, Yuanwei Xu and Jonathan Zinman for their very useful comments. We thank audiences at the Africa meeting of the Econometric Society, AFE at the University of Chicago, the CSAE Annual Conference, the CEPR/TCD Workshop in Development Economics, the IPA-GPRL Researcher Gathering, the Lahore School of Economics, the Munich Economics of Firms and Labor Conference, NEUDC, the NYUAD-PEDL-CEPR Workshop, Nuffield College Oxford, SEEDEC, and WEFIDEV. The research was funded by the Mutuality in Business project, based at the Oxford Saïd Business School. Ethical approval was granted by the University of Oxford Saïd Business School Ethics Committee (reference SSH_SBS_C1A_16_004). We registered our Pre-Analysis Plan with the AEA RCT Registry (AEARCTR-0004789, available at <https://www.socialscisceregistry.org/trials/4789>).

[†]francesco.cordaro@eom.org: Economics of Mutuality

[‡]fafchamp@stanford.edu: Freeman Spogli Institute, Stanford University

[§]colin.mayer@sbs.ox.ac.uk: Saïd Business School, University of Oxford

[¶]muhammad.meki@qeh.ox.ac.uk: Department of International Development and Centre for Islamic Studies, University of Oxford

^{||}simon.quinn@economics.ox.ac.uk: Department of Economics and Centre for the Study of African Economies, University of Oxford

^{**}k.roll@ucl.ac.uk: Institute for Innovation and Public Purpose, University College London

1 Introduction

Many large multinational firms operate ‘route-to-market’ programmes in developing countries. To get their products to the end customer – especially in harder-to-reach rural areas and informal urban settlements – multinationals often rely on a network of ‘micro-distributors’: self-employed individuals who move consumer products from the firm’s stock points to customers. This large group of micro-distributors is not formally employed by the firm, but often are highly reliant on the firm for their income. Such multinationals have access to high-quality administrative data on the performance of their distributors, and recent changes in technology and digital financial services is enabling new ways for them to manage, monitor and pay these ‘workers’ (Adams, Freedman, & Prassl, 2018; Higgins, 2019; Suri, 2017). While route-to-market activities are highly prevalent in developing countries, there is a surprising lack of supply chain finance or trade credit provided by multinationals to their distributors, and relatively little academic research on the topic.

In this paper, we run the first field experiment of a performance-contingent microfinance contract. To do so, we work within the supply chain of one of the largest manufacturers of food products in the world, leveraging novel data to support lumpy capital investments by micro-distributors. We refer to this corporation pseudonymously as ‘FoodCo’. FoodCo owns a large chewing gum producer in Kenya, and wishes to help its micro-distributors with the financing of a productive asset to increase their distribution activities. FoodCo uses two kinds of micro-distributor: (i) individuals who move confectionery products from stock-points to retailers and (ii) individuals who move products from stock-points for direct sale to consumers. Both forms of micro-distributor need to transport stock – and, traditionally, most do so on foot, without the help of a vehicle. In our experiment, we partner with a local microfinance institution in order to provide new bicycles for micro-distributors. Our purpose is to understand whether it is profitable for large corporations to finance the acquisition of productive assets for micro-distributors working within their supply chains through novel performance-contingent microfinance contracts that share risk and reward more effectively than standard debt contracts. We use a field experiment to investigate the efficacy of alternative contractual structures for financing the productive asset. In doing so, we link asset finance repayments to the performance of micro-distributors, using FoodCo’s administrative data on stock purchases, from which we calculate a profit measure that is not reliant upon self-reported data by microdistributors.

Specifically, we test four alternative contracts, randomly offered to individuals who were part of FoodCo’s distribution network and who had expressed an interest in purchasing a bicycle to expand their distribution operations. All contracts require a 10% deposit to finance the bike purchase; the contracts differ in the way that the repayments for the remaining 90% are structured. First, we test a traditional *debt* contract, in which micro-distributors agree to repay the value of the asset, plus 15%, spread evenly over 12 fixed monthly payments. This contract – as with other standard debt contracts – does not involve any linking of repayment obligations and performance. Second, we test an *equity-*

like contract – in which, over the course of 12 months, clients agree to pay half of the fixed monthly payment of the debt contract, plus a 10% share of their monthly profits from micro-distribution work. This contract represents the first innovation in our experiment: by working directly with FoodCo and having access to their administrative data on distributor stock purchases and the price at which gum is sold to the end customer, we offer clients the option of making their repayment obligations contingent upon their self-employment performance (but we do so in a way that does not rely upon self-reported profits). Third, we offer a *hybrid* contract; in this contract, monthly payments are calculated in the same way as the equity contract, but repayments end once the total amount paid reaches the level required under the debt contract (that is, the value of the asset plus 15%). Such hybrid contracts offer the repayment flexibility of the equity contract, with reduced risk for the MFI given that a fixed total amount must still be repaid, and are increasingly being used in high-income countries by digital payments firms to provide financing for small businesses (Rishabh & Schäublin, 2021).¹ Fourth, we offer an *index insurance* contract; this has the same basic structure as the equity contract, but the sharing payments were based on 10% of an index constructed from the profits of other micro-distributors in the region, utilising the fact that we have access to data for all 1,727 of FoodCo’s distributors across Kenya, not just the 161 who participated in our project. Finally, we have a pure control group, who are not offered any of the contracts to finance a bicycle but maintain full ‘business as usual’ access to the FoodCo microdistribution program.

We have two main results. The first relates to take-up. Since equity-like contracts directly tax micro-distributor profits, one may expect adverse selection into such contracts from the ‘best’ micro-distributors. We find little evidence of this. While overall take-up rates are lower for the equity contract (53%) than for the debt contract (69%), we find no evidence of systematically lower take-up for individuals who have higher business profits at baseline, nor for those with greater business management practices or cognitive ability. We find some evidence of greater take-up into equity contracts for individuals who at baseline were measured to be more risk averse and for those measured to be more loss averse, which is consistent with equity contracts providing implicit insurance and downside protection that is particularly valued by such individuals. Similarly, we do not find any evidence of systematically lower take-up of the hybrid contract by the top distributors, which has an overall take-up rate of 70%. The lowest take-up rate is for the insurance contract (45%); this is consistent with many previous studies of micro-insurance that have hypothesised – often without having the data to explicitly test – that their observed low take-up may relate to basis risk (low correlation between the individual-level performance and the movement of the index upon which payments is made). In our setting, we are

¹ Our hybrid contract is designed such that the duration of the contract extends if the distributor is experiencing low profits, with a maximum duration of two years, compared to a 12-month duration for all other contracts i.e. in the fixed-duration equity contract, the distributor may end up paying much less or much more than the debt contract in nominal terms if their profits are low or high during the 12-month period, whereas under the hybrid contract the distributor is always obligated to pay the same nominal amount as the debt contract. We provide further details of the calibration of these parameters in the paper.

able to measure basis risk precisely due to the detailed administrative data for all individuals inside our project and those outside of it, and we show that basis risk in such contracts is indeed quite large.

Our second main result concerns the downstream impact of the contracts on micro-distributor profits. Many studies have hypothesised that performance-contingent contracts are likely to induce moral hazard, either in terms of ex-ante moral hazard (lower effort in business activities) or ex-post moral hazard (greater likelihood of non-payment of contractual liabilities). We find no evidence of this; in fact, we find the opposite. The two performance-contingent contracts (equity and hybrid) actually *out-perform* the debt contract in a number of dimensions. First, we find no evidence that individuals assigned to performance-contingent payments move away from the activity in which they are being taxed; if anything, we actually find that they increase the proportion of their selling activities in chewing gum compared to individuals assigned to the debt contract (at baseline, the average micro-distributor has an approximate equal split in the portfolio of goods that they sell between chewing gum supplied by FoodCo and other non-FoodCo consumer goods). Further, our high-quality administrative data on profits reveals large positive impacts of the performance-contingent contracts on micro-distributor profits, relative to the control and all other contracts. Specifically, in intent-to-treat terms, we estimate an increase in monthly profits of US\$34 for individuals assigned to the hybrid contract, on a control mean of US\$11: a quadrupling of profit. We use a combination of daily administrative data and survey data to explore mechanisms. We find that the hybrid contract in particular caused distributors to exert more effort in their business and in their usage of the financed asset, and it led to increases in their own risk-taking (credit extension to their clients), as well as an improvement in their business management practices. We also find that here is a large and significant effect of the hybrid contract on the biggest category of household consumption expenditure, food (a 19% monthly increase relative to the control group), as well as a large positive impact of the hybrid contract on household expenditure on clothing. Further, we demonstrate that these results are robust to possible concerns about our main impacts being driven by ‘business stealing’ from respondents who were not offered a bicycle. Using data on the universe of entrepreneurs in the program, we find evidence of *positive* spillovers on participants outside of our experiment. Using survey questions about location of customers, as well as GPS data from trackers that were attached to bicycles, we show that this unexpected result is likely driven by the bicycles expanding respondents’ geographical market reach.

Our paper draws together two previously disparate strands of research: microfinance and supply chain finance. The first literature has identified limited impacts of the standard rigid microcredit contract on business performance and household outcomes (Banerjee, Karlan, & Zinman, 2015; Meager, 2018), notwithstanding some evidence of significant heterogeneity in business impacts (Banerjee, Breza, Duflo, & Kinnan, 2019) and positive general equilibrium effects (Breza & Kinnan, 2021). A related body of work has demonstrated the benefits of introducing more flexibility into the standard contract through ‘repayment grace periods’ following the seminal work of Field, Pande, Papp, and Rigol (2013), and contracts better tailored to the cash requirements of farmers (Barboni & Agarwal,

2018; Battaglia, Gulesci, & Madestam, 2017; Burke, Bergquist, & Miguel, 2019). Rather than adapting the structure of the classic microcredit contract to better align repayments with income streams, we push the frontier in this literature by investigating a more direct way to link repayments to income with what, to our knowledge, is the first field experiment of a performance-contingent contract for microenterprises.² Such contracts may be more appropriate for financing the investments of microenterprises with high expected but volatile returns, and especially for the most risk-averse microenterprise owners, but until now have only been tested in a laboratory setting or very small pilot studies.³

We also contribute to the supply chain finance literature, on which there is very little work in developing countries, despite the prevalence of large multinational route-to-market programmes and strong demand for trade credit and financing at various points in the supply chain (Breitbach, 2017; Maksimovic & Demirgüç-Kunt, 2001; Prahalad & Hammond, 2002; Sodhi & Tang, 2014). In an agricultural setting, Kremer, Jack, de Laat, and Suri (2019) work within a milk supply chain (where the output is also well observed, like in our setting) and find large benefits to financing a productive asset for farmers (a rainwater harvest tank). By conducting an experiment with a large multinational distributor, we shed light on the exciting potential for multinationals to help finance productive assets for micro-distributors, which has been described by CGAP as the 'next frontier in financial inclusion',⁵ and by leveraging innovations in technology and digital finance that improve the observability of microenterprise performance and mitigate problems arising from costly state verification (Townsend, 1979), which previously implied that performance-contingent contracts were not optimal for microenterprise financing.

2 Experimental design

2.1 Study context

In 2013, FoodCo developed a route-to-market micro-distribution program using self-employed micro-distributors in Kenya. The distribution system is built around small warehouses (called 'stockpoints'),

² There is, of course, a long tradition of research exploring in an agricultural setting the role of implicit insurance through share-cropping contracts (Burchardi, Gulesci, Lerva, & Sulaiman, 2019; Stiglitz, 1974)

³ Fischer (2013) uses a lab-in-field experiment to overlay profit sharing on top of joint liability credit arrangements, and finds that the inclusion of equity-like features can incentivise higher risk-return investments, especially for the most risk-averse individuals. De Mel, McKenzie, and Woodruff (2019) document results from a pilot with nine high-potential small firms, and they highlight the challenge of enforcing equity contracts with an actual ownership share in the whole microenterprise, which we avoid in our setting by focusing on the 'equity-like' component of performance-contingent payments based on an observable profit stream rather than explicit ownership in the whole firm.

⁴ Our work is also partially related to a literature on interlinked contracts in agricultural settings, such as those for coffee and cocoa farmers (Bandiera, Barankay, & Rasul, 2005; Casaburi & Macchiavello, 2014, 2019; Casaburi & Reed, 2020; Casaburi & Willis, 2018; Ghani & Reed, 2020; Kaur, Kremer, & Mullainathan, 2015; Macchiavello & Morjaria, 2015, 2021).

⁵ CGAP is a global partnership of over 30 leading development organisations focusing on financial inclusion and housed at the World Bank.

which are located in both rural and urban areas. Stockpoints receive deliveries of FoodCo chewing gum, which they sell alongside various non-FoodCo products.⁶ Micro-distributors purchase chewing gum (as well as other products) from stockpoints, before selling to customers (often on foot). They initially purchase the gum from the stockpoints with an up-front discount to the market price, which must be paid in full (i.e. no trade credit is provided to them). They additionally receive performance-related pay in the form of an end-of-month bonus via mobile money (M-Pesa) for every bag of gum sold. There is no obligation for distributors to sell gum exclusively, but selling FoodCo's product is relatively profitable, and they have a strong incentive to stay in the program. This setting is common to many route-to-market distribution programs run by multinational corporations around the world. Consumer goods and food and beverage firms including Hindustan Unilever, Danone, Hewlett-Packard and Coca-Cola have run 'last mile' route-to-market programs of this general form (Prahalad & Hammond, 2002). Such programs are used to build new markets amongst the poor, as well as to pursue social objectives through creating income-generating opportunities for participating micro-distributors (Simanis & Duke, 2014). Like many 'gig workers', these micro-distributors can often have an uncertain relationship with their host firm: they are not employees, and they have considerable autonomy – yet they also are directed by their managers and depend upon the firm. On the basis of feedback from FoodCo and from micro-distributors, we hypothesised that bicycle access could substantially improve distributors' sales. However, bicycles are often prohibitively expensive, costing approximately \$100 for a mid-market model.⁷

2.2 Contract variants

Our sample consists of micro-distributors who had been in the FoodCo program for at least three months and who had expressed an interest in purchasing a bicycle to expand their distribution activities. Interested micro-distributors were invited to a baseline workshop where they completed a number of surveys and were randomly offered one of four microfinance contracts (designed in collaboration with FoodCo and our partner MFI):

- (i). Debt. A contract requiring a total repayment amount equal to the asset financing amount plus a 15% mark-up, spread evenly over 12 fixed monthly payments;
- (ii). Equity. A 12-month contract that required clients to pay half of the fixed monthly payment of the

⁶ Appendix Figure A6 shows a simple illustration of the route-to-market flow of goods, and the two types of entrepreneur (one type that sells smaller packets of gum directly to end consumers, and another type that sells larger multi-packs to small retailers such as kiosks).

⁷ Henceforth, we use '\$' to refer to US\$, based on the actual Kenyan Shilling (KES) amounts and the baseline USD-KES exchange rate of approximately 102 at baseline.

⁸ The control group not offered the opportunity to finance a bicycle using any contract, but maintained full 'business as usual' access to the FoodCo microdistribution program; similarly, individuals who had rejected the contract that they drew a ball for were not offered any contract for bike financing. The experimental protocol is described in more detail in Section 2.4.

debt contract (calculated in the equivalent way), as well as paying a 10% share of their monthly profits (calculated from administrative data);

- (iii). Hybrid: A contract in which monthly payments were calculated in exactly the same manner as the equity contract, but with a flexible contract duration: repayments end when the cumulative payments reaches the level required under the debt contract plus the asset financing amount plus a 15% mark-up).
- (iv). Insurance An index insurance contract, which had a similar repayment structure to the equity contract, with the difference being how the profit-sharing was calculated: the 10% sharing payments were based on an index constructed from the profits of other micro-distributors in their region (again, calculated using administrative data).

All contracts required the micro-distributor to pay an initial deposit of 10%. The remaining 90% of the bicycle price was financed by the local MFI, who bore all the credit risk when financing bicycles for FoodCo's micro-distributors. Profit-sharing payments under the equity, hybrid and insurance contracts were calculated using administrative data on stock purchases and a profit margin given to us by FoodCo, based on the retail price at which the micro-distributors were required to sell their products. We designed the contracts to be similar in terms of expected net present value for the median entrepreneur, given the underlying distribution of entrepreneur incomes in the overall route-to-market program.⁹ Based on that pilot and baseline administrative data, we hypothesised that a contract with performance-contingent payments had the potential to benefit micro-distributors through the provision of implicit insurance, by automatically reducing repayment requirements when business conditions are challenging.¹⁰

The equity contract requires half of the fixed monthly repayments of the debt contract, plus 10% of the micro-distributor's monthly profit. Relative to the debt contract, the equity contract is particularly attractive for insuring downside risk: that is, if the micro-distributor has a bad month, the equity contract reduces the payments required. This aspect of the equity contract is likely to be especially valued if micro-distributors are quite risk-averse, or if they have reference-dependent preferences; in both cases, downside losses will have a particularly high marginal disutility, so the advantage of insuring such losses will be relatively high. Conversely, the disadvantage of the equity contract – from the

⁹ Specifically, we designed the contracts such that a micro-distributor with an income of \$40 per month would have the same repayments and contract duration under all three contracts; this represented the 56th percentile of the income distribution for micro-distributors from our pilot.

¹⁰ We confirmed the importance of such a mechanism by analysing in advance the variance of administrative data on incomes of all micro-distributors in the FoodCo program; this revealed the importance of both within-micro-distributor as well as between-micro-distributor variation. However, this benefit of performance-contingent payments could also in principle be achieved by an index insurance contract, which would not directly tax a micro-distributor's income. Index insurance may also be preferred to our 'equity-based' contracts, if common shocks are the primary concern for micro-distributors. Therefore, to provide a sharper test of the value of our 'equity-like' contracts, we added the 'insurance' contract to the set of contracts that we tested.

micro-distributors's perspective – is the implicit taxation of high performance. If such performance is concentrated among high-achieving micro-distributors – that is, if some micro-distributors are much more productive than others – then it is possible that this contract would suffer from adverse selection; in that case, the MFI would be left sharing profits only with the less productive clients. Moreover, the performance-contingent nature of payments under the equity contract highlights the potential – at least in theory – for problems of moral hazard; a sufficiently patient micro-distributor may prefer to take the bicycle but not to exert effort until after 12 months, when the contract is complete.

The hybrid contract involves performance-contingent repayments, but the contract terminates when cumulative repayments equal those that would be due under the debt contract. The hybrid contract thus provides the advantages of insuring against downside risk, but without the disadvantage of the taxation of high performance. We should, therefore, anticipate that the hybrid contract will prove more popular among micro-distributors than the equity contract. Further, if micro-distributors experience an endowment effect – such that they would prefer to bring forward the day on which they own the bicycle outright – then this contract directly incentivises effort. That is, with high profits the distributor is able to clear their total liability faster and end the contract early.

Finally, the index insurance contract takes a similar form to the equity contract – except that the payments are contingent on the average profits of other distributors in the region, rather than the individual's own profits.¹¹ This contract shares a similar advantage to the equity contract – namely, that it insures the micro-distributor against common shocks through the index. The index insurance contract is very similar in spirit to 'Area-based Yield Insurance' that is used to mitigate asymmetric information problems in agricultural settings by making payouts based on average yields over clearly defined geographic units (Carter, Galarza, & Boucher, 2007).¹² One advantage of this contract over standard index insurance contracts – for which take-up in other contexts has been extremely low (Cole et al., 2013) – is that the monthly calculations offer a greater opportunity for the 'experienced benefits' of insurance. This builds upon work in the literature that shows the benefits of experienced payouts on future take-up (see, for example, Karlan, Osei, Osei-Akoto, and Udry (2014)). As discussed by Cai, de Janvry, and Sadoulet (2020), with most insurance products it is difficult to learn about the value of insurance since a 'positive' experience with insurance only happens when there is a negative shock, which is rare. Our product allows more opportunity for experience to be built, given that calculations are done every month and so the benefits (lower required payments when average profits in the region are lower) are continuous rather than being rare. As with any index insurance contract, however, the contract has an important disadvantage: basis risk. That is, the micro-distributor may face a situation where others at the stock point perform well (perhaps, for example, aggregate demand in the local area is high), but her own performance is poor (perhaps, for example, she is ill in a particular month). Many

¹¹ To mitigate the risk of manipulation, we excluded from the index the profits of distributors in the individual's own stockpoint.

¹² For the index, 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, Mombasa).

previous microinsurance products have been characterised by very high basis risk, which can partially explain the very low take-up levels. (Carter, de Janvry, Sadoulet, Sarris, et al., 2014; Clarke, 2016; Cole et al., 2013). Figure A9 illustrates the relationship between micro-distributor performance and required payments under each contract. Panel A displays payments under the `deterministic' contracts, where payment amounts due are either completely unrelated to performance (debt contract) or related only to one's own performance (equity contract). In contrast, Panel B illustrates payments under the index contract, which are a realisation of a stochastic outcome (the sales of other micro-distributors in one's region).

2.3 Conceptual framework

To x ideas, we now discuss the trade-offs facing a stylised micro-distributor. We provide an intuitive discussion in this section, and we formalise the theoretical analysis in Appendix Section A2

Speci cally, we consider a stylised micro-distributor who is credit-constrained, and whose productivity will increase if she makes a lumpy purchase of a bicycle. The micro-distributor, faced with our menu of nancing contracts, needs to answer two questions. First, the **incentive compatibility** question: "under each available contract, how much effort shall I invest in sales for FoodCo (`on contract'), and how much effort shall I invest in other activities (`off contract')?" Second, the **individual rationality** question: "given a take-it-or-leave-it decision, which contracts should I accept?"

Risk plays two important roles in our conceptual framework – each of which, in our view, reflects important features of the actual experience of micro-distributors in our experiment. First – as our baseline behavioural games show – micro-distributors are **risk averse**¹³ This implies that, ceteris paribus, the micro-distributors value a contract that bundles some degree of risk-sharing. Second, micro-distributors operate in a risky environment – with the risk increasing along with the micro-distributors' use of the lumpy asset. This feature, too, is closely grounded in the real experience of our respondents. For example, a micro-distributor who cycles her bicycle further to serve new markets is also putting that bicycle at more risk of being stolen, or destroyed in an accident; similarly, new markets themselves are intrinsically likely to be more uncertain (Roll, Dolan, & Rajak, 2021).

Three implications follow from this conceptual framework. First, because micro-distributors are risk averse and risk exposure increases with effort, it follows that performance-contingent contracts can crowd in on-contract effort relative to the debt contract. It is tempting to think of performance-contingent contracts simply as `taxing entrepreneurial success' – and, therefore, creating a moral hazard problem, in which the client directs effort away from the `taxed' activity. Our framework shows that the opposite may be true – because bundled insurance can directly increase the marginal product of effort.

¹³Speci cally, in incentivised risk preference elicitation activities, we nd an average certainty equivalent clearly below the expected payoff.

Second, performance-contingent contracts should be particularly attractive to clients who are more risk averse. This follows straightforwardly from the insurance element that is implicitly bundled in such contracts (Burchardi et al., 2019; Stiglitz, 1975, 1989; Stiglitz & Weiss, 1981).

Finally, performance-contingent contracts may be profitable for the client, by facilitating capital investments – and, indeed, by encouraging additional effort – relative to the no-contract case. In this respect, our theoretical predictions follow the seminal insight of Boucher, Carter, and Guirkingner (2008) on ‘risk rationing’: namely, that when capital investment brings additional risks, an absence of bundled insurance implies that profitable investments often do not go ahead.

2.4 Descriptive statistics, contract assignment and take-up

Micro-distributors who had expressed an interest in purchasing a bicycle to expand their business were first invited to a baseline workshop, where they completed a household survey and a series of behavioural games. At the end of the session, each of the microfinance contracts was carefully explained to the respondents; this included several example scenarios and tests of understanding. Respondents were then introduced to a manager from our partner microfinance institution, who explained that they would be offering the financing contracts for bikes to a randomly-selected subset of participants. When communicating with participants, the words ‘debt’, ‘equity’ and ‘insurance’ were never used; contracts were explained using their cash-flow structure in the local language (Swahili), with each contract colour-coded for ease of remembering.

Between 2017 and 2019, 161 individual distributors participated in the study¹⁴. Their average age was 31, with 15% female and 70% married. 20% had a post-secondary education. On average, respondent households had three members. In the three months prior to the baseline survey, mean sales from all selling activities were \$995 (median \$418), and mean profits were \$133 (median \$107). Focusing just on profits from FoodCo products (taken from administrative data, rather than self-reported), average profits were \$53 (median \$34). Only 16% of distributors had employees, 26% also engaged in another income-generating activity (mostly casual labour), with average income of \$20 from that source (median \$0). Total household income was \$198 on average (median \$142), and total household expenditure was \$196 on average (median \$174).

The microfinance contracts were assigned using a public randomisation conducted with an opaque bag containing coloured balls. Micro-distributors had earlier specified their contract preferences to us, and those who drew a colour for which they had specified their acceptance were immediately directed to a representative from the MFI, who was present and proceeded with signing of contracts.¹⁵

¹⁴ During the same period, there was an average of 478 active micro-distributors per month (those with non-zero sales in a given month) in the wider FoodCo programme from which our participants were drawn. The total number of unique micro-distributors in the programme over that period was 1,727.

¹⁵ As mentioned, FoodCo provided the distributors that were part of its programme, and agreed to provide administrative data on profits, but the local MFI bore all the credit risk and was responsible for managing non-payments and defaults.

Individuals who drew a ball for the control group were not offered the opportunity to finance a bicycle using any contract, but they maintained full 'business as usual' access to the FoodCo microdistribution program; similarly, individuals who had rejected the contract that they drew a ball for were not given any contract. The choice of bicycle was made before contract decisions were elicited; most bikes were 'work friendly' bikes with a rear rack¹⁶. In total, 138 of the 161 participants were assigned to treatment (one of the four financing contracts), with the remainder assigned to control. Appendix Table A1 provides summary statistics, disaggregated by treatment assignment; the table also reports tests of randomisation balance. For all variables in Table A1, individual balance tests do not reject the null of no difference across treatment groups.

We begin by describing take-up, which is defined as whether the individual accepted the contract that they were offered and then followed through to providing the deposit and required documentation. The highest take-up rate was for the hybrid contract, at 70%, followed by 69% for the debt contract, 51% for equity and the lowest take-up rate for the insurance contract at 45%.¹⁷ Surprisingly, we found little evidence of adverse selection – the most profitable entrepreneurs (using baseline FoodCo income and a median split) are no less likely to select into the performance-contingent contracts compared to the debt contract. This is consistent with our simple theoretical framework that suggests that the performance-contingent contracts can crowd-in effort and allow greater risk-taking and higher expected earnings from using the bicycle, the value of which dominates the reduced appeal of contracts that directly tax profits. For supplementary analysis that provides further evidence against such adverse selection, Appendix Table A6 provides more detailed analysis of heterogeneous taking up using a series of pre-specified variables, which indicates some evidence of higher selection into the two contracts with performance-contingent repayments for individuals who at baseline were measured to be more risk averse and more loss averse (consistent with such individuals valuing the insurance-like characteristics of performance-contingent payments), but little evidence of lower selection into performance-contingent contracts for 'higher quality' entrepreneur (as proxied by business management practices score or cognitive test scores).

3 Treatment effects

We now analyse the consequences of the various contracts that we implemented. We had access to administrative data from FoodCo on stock purchases by all 1,727 unique distributors in their programme, regardless of whether they participated in our experiment. Distributors purchase gum from FoodCo-

¹⁶ The menu of bikes included 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.

¹⁷ A formal statistical test does not reject that take-up of the equity contract is the same as for that of debt or hybrid contracts. A formal test does reject that take-up of the insurance contract is the same as that of the hybrid contract and the equity contract (p-values of 0.034 and 0.061 respectively). See Appendix Table A5 for further details.

registered stockpoints, and FoodCo performs meticulous checks with field officers and stockpoints to verify the quality of data on purchases, based on which distributors are paid their end-of-month bonus. We create a panel of monthly profits by using the purchase data and the profit margin made by distributors for each of the six possible chewing gum products, based on the fixed price that FoodCo requires distributors to sell at. We aggregate the profits across the six products to create our primary outcome variable.¹⁸ We use the data for distributors outside of the experiment for the spillover analysis in the next section, but here we begin with the data for the 161 distributors who entered our experiment between 2017 and 2019. For all other variables, we use survey data collected quarterly for up to one year after treatment. Our data covers all available post-treatment months up until the COVID-19 lockdowns in March 2020.¹⁹ For each outcome, we use an intent-to-treat ANCOVA specification:

$$y_{it} = \alpha_0 + \sum_{k \in \{1, \dots, 4\}} \beta_k \text{Offered}_{ik} + y_{i0} + \epsilon_{it} \quad (1)$$

Here, Offered_{ik} is a dummy for whether individual had contract 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 data on profits). We cluster at the individual level. In the Appendix, we show that our results are robust to using randomisation inference. We also report LATE estimates in the appendix.

3.1 Impact of contracts on business activities

Our primary hypothesis, as specified in our pre-analysis plan, is that our treatments affected participants' profits. In column 1 of Table 1, the variable 'Performance-contingent' pools both the equity and hybrid contracts, while in the remaining columns the two performance-contingent contracts are separated. We find a large and significant positive effect of the hybrid and equity contract on profits. Column 1 of Table 1 shows this with the pooled variable, and column 2 displays the impacts separately for the equity and hybrid contracts. On average, micro-distributors assigned to the hybrid contract experienced a \$34 increase in their monthly profits (significant at the 5% level), which is very large in comparison to the control mean (at follow-up) of \$11. The estimated coefficient on the equity contract is \$20 (significant at the 10% level), while the coefficients on the debt and insurance contracts are

¹⁸ Technically speaking, our measure is a gross profit measure, or "net income before non-working-capital expenses", since it includes the main business expenditure of inventory, but excludes other expenditure such as travel and labour costs, which are negligible in this setting where most people are walking or cycling and most businesses are single-person (average business expenses on inventory are 40 times larger than the next biggest category of transportation costs).

¹⁹ We ended the project in March 2020, with approximately 85% of the planned survey follow-up data collected before the COVID-19 lockdown. As in many other settings, the lockdown presented a huge shock to operations; in our case this affected not only the operations of micro-distributors, but also structural changes in the way FoodCo managed the program, and an inability for the MFI to collect microfinance repayments. All of our analysis uses data up until and not including the lockdown. This means that we have an unbalanced panel, with an average of fifteen months of post-treatment administrative data. Note that our treatments were randomly assigned on a rolling basis (20 waves).

\$10 and \$12 (not statistically significant at conventional levels). In Appendix Table A5.1, we repeat the estimation using an Inverse Hyperbolic Sine transformation, with similar results. At baseline, by design, 100% of participants were engaging in micro-distribution work. Column 3 reveals that none of the contracts led to a reduction in the likelihood of doing some form of distribution work (which could either be selling FoodCo or non-FoodCo products). Turning from the extensive margin of engaging in any form of distribution work to the intensive margin of what proportion of their selling portfolio's profits is composed of FoodCo products, column 4 reveals the striking result that we do not find any evidence of moral hazard for the performance-contingent contracts. Individuals assigned to the equity and hybrid contracts do not decrease the proportion of FoodCo activities in their selling portfolio; in fact, individuals assigned to the hybrid contracts increase it significantly compared to individuals assigned to the debt contract. This lack of evidence for moral hazard in this particular context may have been due to the perceived benefit of the FoodCo micro-distribution program (a feature that may be shared with other route-to-market programs of multinational corporations), as well as the benefits of using the productive asset to sell more of the FoodCo product (which pays a relatively high margin and is non-perishable, allowing more time to sell to clients). These benefits may have outweighed any costs from performance-contingent contracts, which could explain why we do not find much evidence of adverse selection or moral hazard in our context. And the finding is consistent with our earlier conceptual framework that highlighted the potential for performance-contingent contracts to crowd in 'on-contract effort', in particular for individuals who are risk averse.

3.2 Spillovers and mechanisms

We now test whether our results are driven by 'business stealing' by treated respondents from control respondents. To test for spillovers, we exploit the fact that we have administrative data on the universe of micro-distributors in FoodCo's program, regardless of whether they participated in our project; distributors who were neither assigned to treatment nor to control. We test the consequence on these micro-distributors of random variation in the number of treated respondents at the stockpoint, conditional upon the number of experimental participants at the stockpoint. Denote by y_{ist} the profits of non-participant i , at stockpoint s , in period t . Denote by A_{st} the total number of participants who had been assigned to treatment at stockpoint s by period t , and by C_{st} the total number assigned to control. Denote by P_{st} the total number of participants assigned at stockpoint s by period t ; that is, $P_{st} = A_{st} + C_{st}$. We estimate:

$$y_{ist} = \alpha_0 + \alpha_1 A_{st} + f(P_{st}) + \epsilon_{ist}; \quad (2)$$

where $f(P_{st})$ denotes a flexible function of the number of participants (specifically, we include a different dummy variable for each different value of P_{st}), and where we cluster by stockpoint. In

this specification, α_1 tests for spillovers at the level of the stockpoint: if there are positive spillovers from providing bicycles, then $\alpha_1 > 0$, and if there are negative spillovers, $\alpha_1 < 0$. This identification strategy relies crucially upon the random assignment to treatment: holding fixed the total number of participants at each stockpoint, α_1 is determined randomly, so $E(\alpha_1 | A_{st} = 1) = 0$. (In Appendix A3, we illustrate the underlying variation that we exploit; there, we show the joint distributions of (i) stockpoint size, (ii) number of research participants, and (iii) number of treated participants.)

Table 2 displays the spillover results. Column 1 reports our primary specification. In column 2, we additionally control for time dummies. Columns 3 and 4 repeat the specifications in columns 1 and 2, but collapsing the analysis to the level of the stockpoint. In each case, we estimate significant positive spillovers from treatment. Specifically, we estimate an increase of about US\$4 for each non-participating entrepreneur at the stockpoint, for each additional person offered a contract. These combined findings are reassuring for the robustness of our main results: it indicates that, if anything, our main results are likely to be slight underestimates of the treatment effects of our contracts.

Our estimate of positive spillovers from treatment is consistent with the provision of bicycles having expanded the geographical reach of the stockpoint. Figure 1 shows several maps in support of this hypothesis, using GPS data from trackers that we installed on all bicycles. We can see that our bicycles were spread across all the most populous areas of Kenya, and – within a particular region – individuals are travelling across significant distances with the bicycles. We explore this further in Table 3 column 4, where we find that respondents assigned to all treatments are significantly more likely than control respondents to report selling to customers further than 1km away from their stock-point, and particularly so for the hybrid and insurance contracts.

Table 3 also explores other mechanisms. In columns 1 and 2, we use information from daily administrative data on stockpoint visits by distributors, which ranges from zero visits per month to every day of the month (with an average of 2.4 visits per month for the control group). We find that there are positive coefficients on all of the contracts, with the only significant coefficient being on the hybrid contract, indicating a large increase of 2.96 visits per month (value 0.056). Column 2 explores a similar question using a Herfindahl index that measures the extent to which a distributor's monthly profits are concentrated in a few days per month (which ranges from 0.03 at the 1st percentile to 1.0 at the 90th percentile, where 1.0 indicates a person who receives all of their monthly income in one day). Coefficients are negative on all of the contracts, indicating a greater smoothing of monthly profits over different days / stock-point visits, with the only significant coefficient again on the hybrid contract (indicating a reduction in concentration of -0.10 compared to a control mean of 0.55, value of 0.066). Column 3 explores whether the contracts led to a greater variety of products sold in the distributors' portfolio (which ranges from one product to six of FoodCo's chewing gum products). Although not significant at conventional levels, the coefficient on the hybrid contract indicates a relatively large amount of portfolio diversification, from an average of 1.3 products for the control group to over 2.0 products per month (a coefficient of 0.71, value 0.181). Column 5 of Table 3 provides evidence that

individuals assigned to the hybrid contract significantly increased their own risk taking, through an extension of trade credit to their customers, with a large increase of 5.7 percentage points likelihood of offering trade credit (significant at the 5% level) compared to a control mean of only 7 percentage points (which represents the percentage of their customers being offered trade credit).

Column 6 of Table 3 shows that individuals assigned to the hybrid contract and the insurance contract experienced positive impacts on business management practices, using an overall index of management practices calculated using the weighted sum of scores for four sub-categories: marketing practices, negotiation, cost and record-keeping practices, and sales targets. One plausible explanation for why the hybrid and insurance contracts had the greatest positive impacts are that these two contracts require the greatest amount of 'intellectual engagement' in calculating payments: the hybrid contract requires participants to pay a proportion of their monthly income, as well as 'carrying forward' the cumulative payments made to date (because when those cumulative payments reach a certain amount, the contract ends). The insurance contract similarly requires a greater cognitive exertion, since payments are linked to a proportion of an index that is calculated based on the regional sales average, excluding the sales of individuals in one's own stock point. In contrast, the equity contract is relatively more straightforward (calculated as a share of personal income for a fixed period of 12 months) and debt contract payments are the same payment every month. For further evidence that is consistent with this hypothesis, column 7 shows a positive treatment effect specifically on the record-keeping category of business management. Finally, columns 8 and 9 provide evidence that individuals assigned to the hybrid contract used their asset more intensively (especially compared to the debt contract, with the cross-coefficient test significant at the 1% level), both in terms of number of hours using the asset and likelihood of using it for business purposes.

3.3 Downstream outcomes

Table 4 presents treatment effects on three major components of household consumption expenditure. Column 1 indicates that there is a large and significant effect of the hybrid contract on the biggest category of household consumption expenditure, food: the coefficient of \$8 implies a 19% increase in monthly household expenditure on food compared to the control group (significant at the 10% level). Column 2 indicates a \$5 monthly increase in household expenditure on clothing for individuals assigned to the hybrid contract (significant at the 5% level, with effects not significant for any other contract). Column 3 investigates the effect on the third biggest category of household consumption expenditure, schooling, with the coefficient suggesting that individuals assigned to the debt contract actually experienced a decrease in monthly household spending on education (a decrease of \$5 per

²⁰ All indices are calculated as weighted sums, using the covariance matrix from the control group (at baseline), as in Anderson (2008). The business management practice questions are based on McKenzie and Woodruff (2015), amended to tailor it to micro-distribution businesses.

month (though not statistically significant on its own, it is significantly lower than the coefficient of positive \$3 on the hybrid contract, with a p -value for the cross-coefficient test of 0.029).

Finally, we consider health outcomes. The physical wellbeing of micro-distributors was one of the major motivations for providing bicycles – given that many of the respondents had previously been carrying very large bags on their back to deliver the stock. Column 4 of Table 4 shows treatment effects on whether health impedes work, and column 5 shows effects on whether work caused physical pain. In both columns, though not statistically significant at conventional levels, almost all coefficients are negative with large coefficient magnitudes (implying a lower likelihood of reporting that their physical health made work difficult or that their work caused them pain), and suggesting that being offered a bicycle led to positive effects on health.

4 Conclusion

Microfinance institutions have traditionally offered a very limited suite of products to their clients, mostly based on the canonical (inflexible) loan contracts pioneered by Grameen Bank. There is now a large literature testing the impact of that traditional form of microcredit contract; that literature generally indicates limited overall benefits, with promising results from contractual innovations that introduce greater repayment flexibility following the seminal work of Field, Pande, Papp, and Rigol (2013) (see Cai et al. (2021) for a recent summary). One benefit of debt-based contracts is that they are typically implemented with minimal direct monitoring of clients, and without the need for an MFI to have credible information on the performance of clients' enterprises.

In this paper, we run the first field experiment of a performance-contingent microfinance contract. We design this contract to help micro-distributors to finance the purchase of a lumpy business asset. We demonstrate significant benefits from such contracts. To do this, we worked in a setting in which performance-contingent contracts were already viable – given (i) the availability of detailed purchase data, and (ii) a clear mechanism for how the productive asset could be used to expand operations for microfinance clients. These two key features are already shared by a large variety of different self-employment contexts, in both low-income and high-income settings. First, the kind of micro-distributor program that we study is common to many route-to-market distribution programs, particularly for consumer goods and food and beverage firms. Second, and more generally, these characteristics are shared by many 'gig work' arrangements – where the host firm typically has a wealth of information about the quality and quantity of gig-worker performance. Thus, for example, Uber and other ride-sharing platforms could use contingent-repayment contracts to help their drivers to finance the purchase of their cars.²¹ Similarly, such contracts could readily apply to a very wide range of other

²¹ JUMO Drive is a novel vehicle finance product that allows Uber drivers in many African countries to buy their car by linking their Uber app to their loan contract: <https://www.jumo.world/press-release/uber-partners-with-jumo-to-provide-driver-partners-with-vehicle-finance/>

sub-contractors – for example, farmers who `nish' livestock animals for sale with equipment loans, or cut-and-trim manufacturers for their machinery, and so on. One could imagine such contracts being offered by the host firm – such as Uber – but one could also imagine third-party sharing agreements, more similar to the model adopted here by FoodCo – in which a specialized lender channels funds to a host firm for contingent lending to gig workers or sub-contractors, in return for clients agreeing to have the host firm share performance data with that lender.²²

Across a wide variety of contexts, rapid developments in financial technology – in particular, increasing adoption of mobile money and of electronic point-of-sale technologies – promise cheap access to credible information on the performance of microenterprises, gig workers and sub-contractors. The next generation of micro finance contracts can leverage these developments to expand the portfolio of products it offers clients – specifically, to include contracts with performance-contingent repayment obligations, offering better sharing of risk and reward. This opens several novel possibilities for contractual innovations that benefit both low-income entrepreneurs and large firms.

²² This is broadly analogous to factoring – in which a company sells its accounts receivable to a financial company, which then collects payment.

References

- Adams, A., Freedman, J., & Prassl, J. (2018). Rethinking Legal Taxonomies for the Gig Economy. *Oxford Review of Economic Policy*, 34(3), 475–494.
- Anderson, M. L. (2008). 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.
- Bandiera, O., Barankay, I., & Rasul, I. (2005). Social preferences and the response to incentives: Evidence from personnel data. *The Quarterly Journal of Economics*, 120(3), 917–962.
- Banerjee, A., Breza, E., Du o, E., & Kinnan, C. (2019). Can Micro nance Unlock a Poverty Trap for Some Entrepreneurs? Working paper
- Banerjee, A., Karlan, D., & Zinman, J. (2015). Six randomized evaluations of microcredit: Introduction and further steps. *American Economic Journal: Applied Economics*, 7(1), 1–21.
- Barboni, G., & Agarwal, P. (2018). Knowing what's good for you: Can a repayment exhibility option in micro nance contracts improve repayment rates and business outcomes? Working paper
- Battaglia, M., Gulesci, S., & Madestam, A. (2017). Contractual exhibility, rm growth, and information asymmetries in micro nance: Experimental evidence from bangladesh. Working paper
- Boucher, S. R., Carter, M. R., & Guirking, C. (2008). Risk rationing and wealth effects in credit markets: Theory and implications for agricultural development. *American Journal of Agricultural Economics*, 90(2), 409–423.
- Breitbach, T. W. (2017). Supply chain nancing in developing countries (Unpublished doctoral dissertation). Massachusetts Institute of Technology.
- Breza, E., & Kinnan, C. (2021). Measuring the equilibrium impacts of credit: Evidence from the indian micro nance crisis. *The Quarterly Journal of Economics*, 136(3), 1447–1497.
- Burchardi, K. B., Gulesci, S., Lerva, B., & Sulaiman, M. (2019). Moral hazard: Experimental evidence from tenancy contracts. *The Quarterly Journal of Economics*, 134(1), 281–347.
- Burke, M., Bergquist, L. F., & Miguel, E. (2019). Sell low and buy high: arbitrage and local price effects in kenyan markets. *The Quarterly Journal of Economics*, 134(2), 785–842.
- Cai, J., de Janvry, A., & Sadoulet, E. (2020). Subsidy policies and insurance de. *American Economic Review*, 110(8), 2422–53.
- Cai, J., Meki, M., Quinn, S., Field, E., Kinnan, C., Morduch, J., . . . Said, F. (2021). Micro nance. *VoxDevLit*, Issue.1
- Carter, M., de Janvry, A., Sadoulet, E., Sarris, A., et al. (2014). Index-based Weather Insurance for Developing Countries: A Review of Evidence and a Set of Propositions for Up-scaling. *Development Policies Working Paper* 11.
- Carter, M., Galarza, F., & Boucher, S. (2007). Underwriting area-based yield insurance to crowd-in credit supply and demand. *Savings and Development*, 33(3), 335–362.

-
- Casaburi, L., & Macchiavello, R. (2014). Saving constraints, reputation and market interlinkages: Evidence from the kenya dairy industry. *Unpublished. AQ, 2*.
- Casaburi, L., & Macchiavello, R. (2019). Demand and supply of infrequent payments as a commitment device: Evidence from kenya. *American Economic Review, 109*(2), 523–55.
- Casaburi, L., & Reed, T. (2020). Interlinked transactions and competition: experimental evidence from cocoa markets. *Working paper*.
- Casaburi, L., & Willis, J. (2018). Time versus state in insurance: Experimental evidence from contract farming in kenya. *American Economic Review, 108*(12), 3778–3813.
- Clarke, D. J. (2016). A Theory of Rational Demand for Index Insurance. *American Economic Journal: Microeconomics, 8*(1), 283–306.
- Cole, S., Giné, X., Tobacman, J., Topalova, P., Townsend, R., & Vickery, J. (2013). Barriers to household risk management: Evidence from india. *American Economic Journal: Applied Economics, 5*(1), 104–35.
- De Mel, S., McKenzie, D. J., & Woodruff, C. (2019). Micro-equity for microenterprises. *CEPR Discussion Paper No. DP13698*.
- Field, E., Pande, R., Papp, J., & Rigol, N. (2013). Does the classic microfinance model discourage entrepreneurship among the poor? experimental evidence from india. *American Economic Review, 103*(6), 2196–2226.
- Fischer, G. (2013). Contract structure, risk-sharing, and investment choice. *Econometrica, 81*(3), 883–939.
- Ghani, T., & Reed, T. (2020). Relationships on the rocks: Contract evolution in a market for ice. *Working paper*.
- Higgins, S. (2019). Financial technology adoption. *Northwestern University: Kellogg School of Management*.
- Karlan, D., Osei, R., Osei-Akoto, I., & Udry, C. (2014). Agricultural decisions after relaxing credit and risk constraints. *The Quarterly Journal of Economics, 129*(2), 597–652.
- Kaur, S., Kremer, M., & Mullainathan, S. (2015). Self-control at work. *Journal of Political Economy, 123*(6), 1227–1277.
- Kremer, M., Jack, W., de Laat, J., & Suri, T. (2019). *Borrowing requirements, credit access, and adverse selection: Evidence from kenya*.
- Macchiavello, R., & Morjaria, A. (2015). The value of relationships: evidence from a supply shock to kenyan rose exports. *American Economic Review, 105*(9), 2911–45.
- Macchiavello, R., & Morjaria, A. (2021). Competition and relational contracts in the rwanda coffee chain. *The Quarterly Journal of Economics, 136*(2), 1089–1143.
- Maksimovic, V., & Demirgüç-Kunt, A. (2001). Firms as financial intermediaries: Evidence from trade credit data. *Available at SSRN 632764*.
- McKenzie, D., & Woodruff, C. (2015). *Business practices in small firms in developing countries*. The World Bank.
- Meager, R. (2018). Understanding the average impact of microcredit expansions: A bayesian hierarchical analysis of seven randomized experiments. *American Economic Journal: Applied Economics*.
- Prahalad, C. K., & Hammond, A. (2002). Serving the world’s poor, profitably. *Harvard*

-
- business review*, 80(9), 48–59.
- Rishabh, K., & Schäublin, J. (2021). Payment Fintechs and Debt Enforcement. *Working paper*.
- Roll, K., Dolan, C., & Rajak, D. (2021). Remote (dis) engagement: Shifting corporate risk to the ‘bottom of the pyramid’. *Development and Change*, 52(4), 878–901.
- Simanis, E., & Duke, D. (2014). Profits at the bottom of the pyramid. *Harvard Business Review*, 92(10), 86–93.
- Sodhi, M. S., & Tang, C. S. (2014). Supply-chain research opportunities with the poor as suppliers or distributors in developing countries. *Production and operations management*, 23(9), 1483–1494.
- Stiglitz, J. (1974). Incentives and Risk Sharing in Sharecropping. *The Review of Economic Studies*, 41(2), 219–255.
- Stiglitz, J. (1975). The effects of income, wealth, and capital gains taxation on risk-taking. *Stochastic Optimization Models in Finance*, 291–311.
- Stiglitz, J. (1989). Financial markets and development. *Oxford Review of Economic Policy*, 5(4), 55–68.
- Stiglitz, J., & Weiss, A. (1981). Credit rationing in markets with imperfect information. *The American economic review*, 71(3), 393–410.
- Suri, T. (2017). Mobile money. *Annual Review of Economics*, 9, 497–520.
- Townsend, R. M. (1979). Optimal contracts and competitive markets with costly state verification. *Journal of Economic theory*, 21(2), 265–293.

Tables and figures

Table 1: **Business outcomes**

	(1)	(2)	(3)	(4)	(5)
	FoodCo profits	FoodCo profits	Activity: seller	FoodCo proportion	Other earnings
Debt	10.39 (11.535)	10.39 (11.520)	-0.05 (0.054)	-0.11** (0.046)	5.95 (15.253)
Performance-contingent	25.96** (10.786)				
Hybrid		34.43** (15.227)	0.03 (0.044)	0.03 (0.060)	-7.73 (13.347)
Equity		19.61* (11.742)	-0.03 (0.053)	-0.01 (0.046)	-1.68 (12.270)
Insurance	11.85 (10.312)	11.87 (10.269)	0.02 (0.040)	-0.06 (0.045)	3.07 (15.415)
Observations	2598	2598	468	468	468
Individuals	161	161	160	160	160
Control mean	11.32	11.32	0.93	0.48	70.67
Test: Hybrid = Debt		0.133	0.181	0.018	0.319
Test: Hybrid = Insurance		0.357	0.326	0.469	0.557
Test: Hybrid = Equity		0.472	0.741	0.023	0.541

Note: In this table we report *intent-to-treat* (ITT) estimates, obtained by least-squares estimation. Columns 1 and 2 use administrative data from FoodCo on business profits, for which there is an average of 15 months of post-treatment data (up to and excluding the start of Covid-19 lockdowns in March 2020). For all other columns, we use survey data collected by enumerators using quarterly follow-up surveys (again, up to and excluding Covid-19 lockdowns). Standard errors, clustered at the individual level, are reported in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: **Spillovers**

	(1)	(2)	(3)	(4)
Level of analysis:	Non-participating clients		Stockpoints	
Number treated at the stockpoint	3.96*** (1.343)	4.11*** (1.388)	4.07* (2.075)	4.03** (2.039)
Constant	11.04*** (1.298)	10.97*** (1.229)	11.64*** (1.002)	11.64*** (0.993)
Controls: Total participating at the stockpoint	yes	yes	yes	yes
Controls: Time	no	yes	no	yes
Observations	52948	52948	9737	9737

Notes: In this table, we use administrative data on micro-distributors who were not involved in our experiment, and test the consequence of random variation in the number of treated respondents at the stockpoint. We report intent-to-treat (ITT) estimates, and standard errors in parentheses (with clustering at the level of the stockpoint). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Mechanisms

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Stockpoint visits	Profit concentration	Product varieties	Sales expansion	Credit extension	Management practices	Record keeping	Bike use: business	Bike use: hours
Debt	1.28 (1.154)	-0.05 (0.048)	-0.02 (0.441)	0.10 (0.082)	0.01 (0.023)	0.00 (0.061)	-0.02 (0.072)	0.73*** (0.055)	22.32*** (2.142)
Hybrid	2.96* (1.539)	-0.10* (0.054)	0.71 (0.532)	0.19** (0.090)	0.05** (0.026)	0.10* (0.055)	0.14** (0.068)	0.90*** (0.037)	34.82*** (5.553)
Equity	1.29 (1.032)	-0.03 (0.044)	0.10 (0.468)	0.13 (0.087)	0.01 (0.020)	0.03 (0.055)	0.01 (0.067)	0.71*** (0.058)	24.90*** (2.067)
Insurance	0.27 (1.124)	0.01 (0.042)	0.07 (0.391)	0.22*** (0.076)	-0.00 (0.019)	0.11** (0.052)	0.11* (0.069)	0.79*** (0.068)	31.23*** (5.981)
Observations	2598	2598	2598	468	468	468	468	468	468
Individuals	161	161	161	160	160	160	160	160	160
Control mean	2.42	0.55	1.33	0.58	0.08	0.68	0.65	0.00	0.00
Test: Hybrid = Debt	0.307	0.311	0.140	0.228	0.157	0.089	0.014	0.008	0.036
Test: Hybrid = Insurance	0.241	0.137	0.231	0.486	0.104	0.161	0.036	0.006	0.094
Test: Hybrid = Equity	0.994	0.719	0.777	0.626	0.948	0.676	0.651	0.847	0.386

Note: In this table we report *intent-to-treat* (ITT) estimates. We use information from daily administrative data (columns 1 to 3), survey data from all participants (columns 4 to 7) and information on asset usage specifically from clients who took up the treatment (columns 8 and 9) to explore a number of variables that shed light on the mechanisms for our results from Table 1: how often distributors visit stock-points in a given month to purchase inventory (which ranges from 0 to 31), how concentrated their total monthly profit is over those visits (Herfindahl index), the number of FoodCo products they sell in their monthly portfolio (which ranges from 1 to 6), whether they sell to distant customers (greater than 1km from their stock-point), whether they extend credit to customers, their business management practices (an overall index and specifically record keeping) and finally whether they use the bicycle for business and the number of hours that they use it in a typical week. Standard errors, clustered at the individual level, are reported in brackets. * p < 0.10, ** p < 0.05, *** p < 0.01.

Table 4: Household consumption and health

	(1)	(2)	(3)	(4)	(5)
	Expenditure: food	Expenditure: clothing	Expenditure: schooling	Health impedes work	Work caused pain
Debt	8.99* (5.075)	0.25 (1.965)	-4.91 (3.420)	-0.09 (0.070)	-0.10 (0.062)
Hybrid	8.47* (5.117)	4.92** (2.372)	3.10 (4.360)	-0.06 (0.078)	-0.03 (0.073)
Equity	1.54 (4.152)	-0.16 (2.146)	-0.81 (3.649)	-0.07 (0.072)	-0.02 (0.067)
Insurance	8.18* (4.247)	-2.34 (1.974)	-0.44 (3.355)	-0.03 (0.079)	0.02 (0.078)
Observations	468	468	468	468	468
Individuals	160	160	160	160	160
Control mean	45.72	9.26	11.34	0.26	0.19
Test: Hybrid = Debt	0.927	0.032	0.029	0.644	0.204
Test: Hybrid = Insurance	0.155	0.030	0.312	0.792	0.883
Test: Hybrid = Equity	0.111	0.831	0.150	0.827	0.084

Note: In this table we report *intent-to-treat* (ITT) estimates of the impact of treatments on household consumption expenditures and participants' self-reported health outcomes. Standard errors, clustered at the individual level, are reported in brackets. * p < 0.10, ** p < 0.05, *** p < 0.01.

