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

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

November 2024

Abstract

We collaborate with a large multinational that wishes to increase the productivity of micro-distributors in its supply chain by financing a business asset. We conduct a field experiment comparing financing under traditional debt contracts with three alternatives that offer a greater sharing of risk and reward. We find the largest impacts from a novel hybrid contract, which combines debt-like features with performance-contingent payments. Our findings suggest substantial mutual benefits for the multinational, its micro-distributors, and stockpoints in its supply chain. These results highlight the potential of financial contracts that leverage improved performance data observability in low- and middle-income countries.

*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 audiences at the Africa meeting of the Econometric Society, University of Chicago AFE, CEPR TCD Workshop, WEFIDEV, Yale University NEUDC, Northwestern University, GPRL-IPA, World Bank, Georgetown University, Bocconi University LEAP, NYUAD-PEDL-CEPR workshop, Cambridge Centre for Alternative Finance, DIW Berlin, Munich Economics of Firms and Labor Conference, Nordic Conference in Development Economics, NOVAFRICA, SEEDEC, Society for Experimental Finance, LUMS, CAFRAL, CSAE and Nuffield College Oxford. We are also grateful to Anik Ashraf, Haseeb Ashraf, Deniz Aydin, Abhijit Banerjee, Peter Bossaerts, Paola Bustos, Jing Cai, Stefano Caria, Kevin Carney, Lorenzo Casaburi, Paola Conconi, Kim Fe Cramer, Zoë Cullen, Joshua Deutschmann, Christopher Eaglin, Seth Garz, Xavier Giné, Jonathan Greenacre, Selim Gulesci, Morgan Hardy, Sean Higgins, Simon Hess, Michal Hodor, Martin Kanz, Erin Kelley, Michael King, Cynthia Kinnan, Pramila Krishnan, Nirupama Kulkarni, Greg Lane, Friederike Lenel, Nicola Limodio, Rocco Macchiavello, Karen Macours, Mahreen Mahmud, Karol Mazur, Craig McIntosh, David McKenzie, Mushfiq Mobarak, Timothy Ogden, Imran Rasul, Raghavendra Rau, Jānis Skrastinš, Tavneet Suri, Alessandro Tarozzi, Christine Valente, Eric Verhoogen, Jack Willis, Chris Woodruff, Yuanwei Xu, Erina Ytsma and Jonathan Zinman, for their helpful comments. 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).

†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, University of Oxford
‖simon.quinn@imperial.ac.uk: Department of Economics and Public Policy, Imperial College London
**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 low- and middle-income 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 stockpoints to customers. Although not formally employed by the multinational, these distributors often depend heavily on it for income and face significant risks due to the nature of their 'last-mile distribution'. Multinationals increasingly discuss 'corporate social responsibility' and 'improved stakeholder engagement' – but often these discussions lack specific innovations. In this paper, we test one specific approach: *can a multinational facilitate profitable investments within its supply chain, to the advantage of itself and its distributors?*

To answer this question, we conduct a field experiment in collaboration with one of the largest manufacturers of food products in the world. We refer to this corporation pseudonymously as 'FoodCo'. FoodCo owns a large chewing gum producer in Kenya, and wishes to increase the productivity of micro-distributors in its supply chain by financing a business asset. FoodCo's distribution system is built around small micro-retailers, called 'stockpoints,' which receive deliveries from FoodCo. Such micro-retailers are central to the economic fabric of many low- and middle-income countries, serving as the primary channel through which millions of households purchase fast-moving consumer goods (Kruijff, Sawhney, & Wright, 2024). Micro-distributors in our setting purchase the chewing gum from these stockpoints before selling it to their customers. Traditionally, most distributors travel on foot. Qualitative work highlighted that many distributors could significantly increase their productivity with a transportation asset (such as a bicycle), but most are credit-constrained. An additional feature of this setting is that distributors require up-front financing for inventory before making any sales, and do not receive trade credit from their stockpoints. These constraints make it difficult for distributors to save for lumpy investments. FoodCo has historically not engaged in financing activities as it views its role as a supplier of food products rather than capital, avoiding the regulatory and enforcement complexities that such financing would entail. To enable the investment opportunity given the organizational constraints, we encouraged a collaboration between the multinational and a local lender that provides the capital for the bicycles. Through this collaboration, the lender gains access to FoodCo's administrative data, facilitating the development of innovative financial contracts that allow repayments to be linked to distributor performance.

These contracts use a gross profit measure calculated from stock purchases and known final selling prices, bypassing the need for self-reported profits. Traditionally, in data-poor settings, debt has been considered the optimal financing contract due to issues of costly state verification (Townsend, 1979). Our aim in this paper is to test whether, in data-rich environments, performance-contingent contracts can more effectively finance lumpy business investments by offering a better sharing of risk and re-

ward. Multinationals often have access to high-quality administrative data on the performance of their distributors. Recent technological advances offer new potential for multinationals to use this detailed data to manage, monitor, make payments, and predict creditworthiness (Annan, Cheung, & Giné, 2024; Berg, Burg, Gombović, & Puri, 2020; Fuster, Goldsmith-Pinkham, Ramadorai, & Walther, 2022; Higgins, 2024; Suri, 2017). However, there is relatively little economic research on the effects of different financial contracts for facilitating investments within retail distribution networks and supply chains (Jack, Kremer, de Laat, & Suri, 2023).[1]

First, we examine the overall impact of offering asset financing. Among a sample of distributors screened for interest in expanding their operations through the purchase of a bicycle, we find a 57.6% overall take-up rate when offered any contract. This results in a substantial increase in business profits from selling FoodCo's products, with no evidence of effects on other sources of income. Specifically, in intent-to-treat (ITT) terms, we estimate an 88% increase in monthly business profits for those assigned to a financing contract compared to the control group, in the three years following the experiment (standard error: 43%). The corresponding Local Average Treatment Effect (LATE) estimate indicates a 132% increase in monthly profits (standard error: 63%). This is consistent with increased business effort, as those who take up the contract nearly double their visits to stockpoints each month to purchase inventory. They also significantly expand the geographical area in which they sell, which is consistent with our GPS tracker data indicating that distributors used the bicycles to cover vast distances.

Second, we address whether novel financial contracts can support such investments more effectively than standard debt contracts. To do this, we test four contract types: a standard debt contract, a revenue-sharing contract, an index insurance contract, and a hybrid contract that combines debt-like payments with revenue sharing. Specifically, the debt contract, the revenue-sharing contract and the index insurance contract each have a fixed duration of 12 months. The hybrid contract operates on a revenue-sharing basis, but ends as soon as total repayments match the repayments that would be owed under the debt contract; in this way, the hybrid contract provides implicit insurance against economic shocks, but caps the distributor's upside sharing. We find that performance-contingent features lead to improved business performance – with the hybrid contract the most successful of the four. We find that the hybrid contract outperforms the standard debt contract on several fronts, despite similar take-up rates (69.2% and 67.7%, respectively). In ITT terms, the hybrid contract leads to a 170% increase in monthly profits compared to the control group (standard error: 69%), while the LATE estimate indicates a 219% increase (standard error: 91%). For the debt contract, the ITT estimate shows a 59% increase in monthly profits (standard error: 48%), with a LATE estimate of 77% (standard error: 63%). Cross-coefficient tests confirm that the hybrid contract outperforms the debt contract in terms of business profit impacts, across multiple specifications.

Exploring mechanisms, we find that the hybrid contract also outperforms the standard debt con-

[1] For related literature in other fields, see (for example) Kolk, Rivera-Santos, and Rufín (2014) and Singh, Bakshi, and Mishra (2015).

tract across several measures of business effort — including whether the financed asset is used for business purposes, the intensity of asset usage, geographical sales expansion, business management practices, and risk-taking through credit extension to distributors' own customers. The hybrid contract also leads to higher repayments to the lender compared to the debt contract. Given the almost identical take-up rates between the hybrid and debt contracts, we conclude that the difference in profits is driven by effort on the intensive margin (that is, an increase in profits conditional on adopting the contract) rather than by differences on the extensive margin; we show robustness of this conclusion using a Lee (2009) bounding exercise.

To interpret these results, we develop a dynamic stochastic model – in which a risk-averse distributor decides how much effort to exert on selling activities, and whether to accept or reject various financing contracts. The model formalizes the intuition that distributors are exposed to greater risk as they use the fixed asset, notwithstanding its high expected return. The model illustrates that the hybrid contract breaks the traditional trade-off between implicit insurance and reduced effort by offering repayment flexibility (and, with it, implicit insurance that helps mitigate liquidity risk) while incentivizing additional effort to clear the debt sooner. This framework aligns with our empirical finding of greater effort and profits under the hybrid contract.

Finally, we conduct a comprehensive cost-benefit analysis, which incorporates the combined impacts of the intervention across all of the relevant actors: the distributors, the multinational, the stockpoints, and the lender. We find large mutual benefits along the supply chain, and remarkably high benefit-cost ratios across all contracts – particularly for the hybrid contract. This is the case even when assuming minimal persistence of treatment effects beyond the three-year mark of the project. For instance, for the pooled estimate, we find a benefit-cost ratio of 6.3 when assuming zero years of treatment effect persistence, with a confidence interval of 1.5 to 11.1, corresponding to an IRR of 203%. This increases to 7.8 when assuming five years of persistence (confidence interval: 1.9 to 7.8), with an IRR of 210%. For the hybrid contract, we find a benefit-cost ratio of 10.8 with zero years of persistence (confidence interval: 3.2 to 18.4), corresponding to an IRR of 356%. This increases to 13.4 when assuming five years of persistence (confidence interval: 4.0 to 22.8), with an IRR of 360%.

Our paper draws together two 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), notwithstanding evidence of significant heterogeneity in business impacts (Banerjee, Breza, Duflo, & Kinnan, 2019; Bryan, Karlan, & Osman, 2024) and positive general equilibrium effects (Breza & Kinnan, 2021; Cai & Szeidl, 2022). A related body of work has shown the benefits of introducing more flexibility into standard contracts through 'repayment grace periods' (Barboni & Agarwal, 2023; Battaglia, Gulesci, & Madestam, 2024; de Haas, Crepon, Pariente, & Devoto, 2022; Field, Pande, Papp, & Rigol, 2013), though sometimes at the cost of higher default rates (Brune, Giné, & Karlan, 2022; Field et al., 2013). This is unsurprising, as greater risk-taking by financed businesses exposes lenders to more downside

risk, while lenders are often constrained in how much they can raise interest rates to capture the upside from more profitable investments. We advance this literature by improving the alignment of incentives between capital providers and businesses – exploring a more direct method of linking repayments to business profits. Performance-contingent contracts may be more appropriate than traditional debt for financing investments of small firms, particularly those with high underlying returns to capital (De Mel, McKenzie, & Woodruff, 2012; Udry & Anagol, 2006), but, until now, have only been tested in laboratory settings or small pilot studies.[2] One valuable feature of our setting is the homogeneity of our sample of business owners, which – along with an understanding of the supply chain, administrative data, and modeling of the distributor production function – provides an ideal opportunity to explore the mechanisms through which contractual terms affect investment behavior and effort.

The second literature – on supply chain finance – has involved relatively little empirical work in developing countries. Nonetheless, there is increasing prevalence of large multinational route-to-market programs, and strong demand for financing at various points in the supply chain (Casaburi & Willis, 2024; Macchiavello, 2022). In an agricultural setting, Jack, Kremer, de Laat, and Suri (2023) work within a milk supply chain (where output is also well observed, as in our context) and find large benefits to financing a productive asset for farmers (a rainwater harvest tank). Other literature in this space emphasizes strong theoretical justifications for suppliers acting as financial intermediaries – due to their comparative advantage in assessing the performance and creditworthiness of customers, and their ability to use relational contracting and informal means to enhance repayment likelihood (Biais & Gollier, 1997; Breza & Liberman, 2017; Burkart & Ellingsen, 2004; Klapper, Laeven, & Rajan, 2012; Macchiavello & Morjaria, 2021; McMillan & Woodruff, 1999; Petersen & Rajan, 1997).

Further, by conducting an experiment within a multinational's supply chain, we shed light on the exciting potential for large firms to help finance productive assets for their 'dependent contractors'. In doing so, we contribute to a growing literature on the benefits of improved market access for small firms (J-PAL, 2024), including the advantages of integrating small firms into multinational supply chains (Alfaro-Urena, Manelici, & Vasquez, 2022; Rodrik & Sandhu, 2024). In this way, we view this paper as an important proof of concept for a new class of financing contracts for small firms operating within a supply chain that permits observability of their performance.

The paper proceeds as follows. Section 2 explains the context, design, and implementation of the experiment, and outlines our conceptual framework. Section 3 reports our treatment effects. We provide a cost-benefit analysis in section 4, and conclude in section 5.

[2] See Fischer (2013), De Mel, McKenzie, and Woodruff (2019), and Meki (2024). Separately, there is a long tradition of research in agricultural settings exploring the role of implicit insurance through share-cropping and other insurance-like arrangements bundled into loans to farmers (Burchardi, Gulesci, Lerva, & Sulaiman, 2019; Giné & Yang, 2009; Karlan, Kutsoati, McMillan, & Udry, 2011; Stiglitz, 1974). In particular, the benefits of revenue-sharing contracts for the risk averse was central to Udry's (1994) analysis of informal state-contingent loans in Nigeria). Our paper contributes to the limited non-agricultural literature, with a focus on more formal actors within a multinational supply chain. There is also a literature on equity-like arrangements for financing human capital investments (Herbst & Hendren, 2024; Mueller & Yannelis, 2022).

2 Experimental design

2.1 Study context

In 2013, FoodCo developed its route-to-market program around small warehouses ('stockpoints') in rural and urban areas. These stockpoints receive and sell FoodCo products (six types of chewing gum) alongside non-FoodCo products. Micro-distributors purchase gum from stockpoints and sell to customers; in doing so, distributors wear distinctive FoodCo-branded shirts, but do not carry company IDs.[3] There is an agreement between FoodCo and the distributors, but this functions as a code of conduct rather than as a formal employment contract. Like many gig workers, distributors exist on the 'ill-defined periphery of the firm' (Barratt, Goods, & Veen, 2020; Hickson, 2024), taking on the risks of last-mile distribution. This setting is common to many route-to-market distribution programs run by multinational corporations around the world (Prahalad & Hammond, 2002).

Distributors initially purchase the gum from the stockpoints. When purchasing gum, distributors receive a discount to the market price, but are required to pay up-front; stockpoints do not provide trade credit. Additionally, distributors receive performance-related pay for every bag of gum sold; this is paid as an end-of-month bonus via mobile money (M-Pesa). There is no obligation for distributors to sell gum exclusively, but selling FoodCo's product is relatively profitable, and distributors have strong incentives to stay in FoodCo's program.

Historically, FoodCo has not engaged in financing activities – avoiding the regulatory and enforcement complexities that such financing would entail. To facilitate investment for our distributors, we encouraged a collaboration between the multinational and a local microfinance institution (MFI). The MFI used FoodCo's administrative performance data for initial screening at loan origination and to calculate payments due, and provided the capital to finance bicycles for distributors who accepted their contract offer. It also assumed full responsibility for the risk of collecting payments on the contracts.

2.2 Contract variants

Each bicycle cost approximately three times the average monthly profit from sales of FoodCo products. We tested four alternative contracts to finance this purchase. All contracts required the distributor to pay an initial deposit of 10%, with the remaining 90% of the bicycle price financed by the local MFI; the MFI bore all of the credit risk, and maintained ownership of the bicycles until completion of each contract. All financing was digital – no cash changed hands, either in disbursing funds to clients (payments were sent via mobile money to procure bikes), or in repayments (which were again made by mobile money).

The contracts were as follows:

[3] In Appendix Figure A4, we provide a graphical illustration of the route-to-market product flowchart.

(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). *RevShare*: A 12-month contract that required clients to pay half of the fixed monthly payment of Debt (calculated in the equivalent way), as well as paying a 10% share of their monthly profits (calculated from FoodCo administrative data, and described in further detail below). Relative to Debt, RevShare is particularly attractive for insuring downside risk: if the distributor has a bad month, RevShare reduces the payments required. Conversely, it is possible for the distributor to owe substantially more under RevShare than under Debt if monthly profits are high.

(iii). *Hybrid*: The monthly payment under Hybrid is the same as under RevShare (that is, half of the fixed monthly payment of Debt, plus 10% of monthly FoodCo profits). However, this contract terminates once the cumulative payments reach the level required under Debt (that is, the asset financing amount plus at 15% mark-up). Therefore, the maximum possible duration for Hybrid is 24 months (in the hypothetical case of a distributor with zero profits every month); the minimum possible duration is just one month (in the unlikely event that profits are so exceptionally high that the initial monthly payment matches the Debt liability immediately). Hybrid thus provides the advantage of mitigating liquidity risk, while avoiding the adverse incentive effects of total wealth being exposed to unlimited upside sharing. Further, if distributors experience an endowment effect (Carney, Kremer, Lin, & Rao, 2022), such that they would prefer to bring forward the day on which they own the bicycle outright, or debt aversion (Azmat & Macdonald, 2020; Martínez-Marquina & Shi, 2024; Paaso, Pursiainen, & Torstila, 2020), then the contract directly incentivizes effort.

(iv). *IndexShare*: This is an index insurance contract. Here, monthly payments are calculated in the same manner as under RevShare – but the 10% sharing payments are based on an index constructed from the profits of other distributors in their region (again, calculated using FoodCo administrative data). This contract shares a similar advantage to the RevShare contract – namely, that it insures the distributor against common shocks – but it does not penalize high effort as RevShare does. This contract is similar in spirit to index insurance contracts in agriculture. These are commonly used to mitigate asymmetric information and adverse incentive issues by basing crop insurance payouts on average yields over a clearly defined area, rather than on their own reported yield (Carter, Galarza, & Boucher, 2007).[4]

Finally, respondents in the *control* group were not offered the opportunity to finance a bicycle using any contract, but maintained full 'business as usual' access to the FoodCo micro-distribution program.

[4] In our study, the index is aggregated at the regional level, encompassing Nairobi, Central Kenya, Kisumu, Eastern Kenya, and Mombasa. The index that we calculate for each treated individual excludes their profits and the profits of other distributors at their stockpoint.

Our administrative measure of profits from selling FoodCo products – which forms the basis for payments under RevShare, Hybrid, and IndexShare, and constitutes the main variable for our empirical analysis – aligns with standard accounting definitions of 'gross profits': the value of sales minus the cost of goods sold. We know the exact cost of goods sold because distributors purchase their gum directly from FoodCo stockpoints. FoodCo performs meticulous checks with field officers and stockpoints to verify the quality of data on purchases, based on which distributors are paid their monthly bonuses (described in Section 2.1). The value of sales is also known because distributors adhere to the recommended retail price set by FoodCo. There are six FoodCo products, and each has a specific profit margin; we aggregate across the six products to form our primary measure of gross profits.[5]

The repayment terms of the different contracts were calibrated to be similar in terms of expected net present value for the median distributor, given (i) the baseline distribution of distributor profits in the broader route-to-market program and (ii) estimates, based on qualitative interviews, of the expected impact of the bicycles on profits.

2.3 Descriptive statistics and contract assignment

We advertised within FoodCo's network for distributors who had been in the FoodCo program for at least three months and who were interested in acquiring a bicycle to expand their business operations. Interested distributors were invited to a workshop, where they completed a survey. Distributors were given the opportunity to inspect several kinds of bicycles on offer; most bicycles were 'work friendly' models with a rear rack.[6]

After the surveys, each of the four possible financing contracts was carefully explained to the respondents in a group activity; this included several example scenarios and tests of understanding. When communicating with participants, the expressions 'Debt', 'RevShare', 'Hybrid', and 'IndexShare' were never used; contracts were explained using their cash-flow structure in the local language (Swahili). Respondents were then introduced to a manager from the partner microfinance institution, who explained that they would be offering the financing contracts for bikes to a randomly selected subset of participants. Each respondent then made a 'take-it-or-leave-it' decision on each of the contracts.

[5] Note that our measure of gross profits does not include other costs commonly referred to in accounting as 'selling, general, and administrative' in standard financial reporting. However, we are confident that the cost of goods sold represents the largest cost category, which we confirmed through in-person surveys where distributors were asked about all selling activities (both FoodCo and non-FoodCo products). Distributors reported that the cost of raw materials was by far the largest expense, representing 85% of total operating costs. The next largest category was transportation, comprising 7% of total costs. In the cost-benefit analysis in Section 4, we incorporate estimates from survey data to calculate a profit measure net of selling, general and administrative costs.

[6] The menu of bicycles included one model that is of a higher quality and nearly twice as expensive, and a 'female-friendly' bicycle with a dipped bar. See Fiala, Garcia-Hernandez, Narula, and Prakash (2022) for evidence of the significant benefits of bicycles for young women, in a setting geographically similar to ours and involving a similar Kenyan bicycle manufacturer. See also Van Doornik, Gomes, Schoenherr, and Skrastins (2024) for evidence of large returns to another transportation asset for women, motorcycles, in Brazil, using a novel asset financing contract.

Contracts were assigned using public randomization. The experiment was implemented across 19 waves, each corresponding to one workshop; in 18 of the 19 workshops, the assignment probability for each of the five treatment arms (control and four financing contracts) was equal.[7] Randomization was carried out using an opaque bag containing 100 colored balls, with 20 balls assigned to each of the five treatment arms, drawn with replacement. Respondents who drew a contract for which they had specified their acceptance were immediately directed to a representative from the MFI, to proceed to sign the contract.[8] 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 micro-distribution program; similarly, individuals who had rejected the contract for which they drew a ball were also not given any contract.

The experiment comprised 161 individual distributors, and was designed (considering the sample size) to be sufficiently powered to detect effects when considering the large hypothesised treatment effect, homogeneous sample, and high-frequency profit data that results in nearly 3,000 data points for administrative data regressions. Our approach bears some resemblance to that of Bloom, Eifert, Mahajan, McKenzie, and Roberts (2013), who conducted an experiment involving only 17 firms (28 plants) and relied on an intervention with a large hypothesised impact, homogeneity of the sample (from a similar business sector), and the availability of a long time series complemented by high-frequency data. These features are also present in our study.[9]

In total, 138 of the 161 participants were assigned to treatment (one of the four financing contracts), with the remainder assigned to control. In the appendix, we provide summary statistics, disaggregated by treatment assignment; the table also reports tests of randomization balance. An omnibus balance test, assessing the equality of coefficients for each treatment across all variables, comfortably passes ($p = 0.971$).[10] Respondents' 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 profits from all selling activities were Ks 13,329 (median Ks

[7] The exception was the first workshop, in which IndexShare was not offered. In the first workshop, the assignment probability for Debt, Hybrid, RevShare, and control was set to be equal, at one quarter. This leads to a slight overall under-sampling of IndexShare relative to the other treatments, given the relatively large size of the first workshop. Nonetheless, this poses no issues for our balance tests or follow-up regression analysis (McKenzie, 2015).

[8] Had the take-it-or-leave-it decisions been binding on both acceptance and rejection, it would have been possible to conduct tests for selection and moral hazard similar to those by Karlan and Zinman (2009) and Jack et al. (2023). Ethically, this was not possible in our context: some individuals who agreed in principle to accept the contract were unable to produce the 10% deposit, and it would not be appropriate to insist upon this.

[9] We anticipated a substantial treatment effect in our experiment because we were providing a highly targeted asset to specifically address a major constraint faced by distributors seeking to expand their businesses: transportation. This hypothesis was strongly supported by qualitative work during the experimental design phase; this consistently highlighted transportation as a key impediment to business growth.

[10] For robustness, we also estimate a multinomial logit specification to test balance between each treatment and control across all the variables using randomization inference, following recent recommendations (Kerwin, Rostom, & Sterck, 2024). This test also passes comfortably ($p = 0.844$).

10,666), and Ks 2,874 (median Ks 2,261) from just FoodCo products (for which we use administrative data).[11] Only 16% of distributors had employees; 26% also engaged in another income-generating activity (mostly casual labor), with average income of Ks 2,000 from that source (and a median of zero).

Several variables indicate that distributors faced liquidity and credit constraints — consistent with our qualitative findings that, despite believing that it would lead to significant profit increases, they struggled to save for the lumpy asset. First, the median household had total monthly consumption expenditure of Ks 17,375 compared to total household income from all sources of Ks 14,225. Second, more than half of distributors report that none of their FoodCo purchases are received on credit. Further, the median distributor only extends trade credit for 5% of their sales. Even where trade credit is provided, the duration is extremely short — for those who receive trade credit from their stockpoint, the average number of days of credit is 2.9 (median of 1), and for those who extend trade credit, average days to repay is 2.2 (median of 1).

To situate our experiment in the broader context, we run a comparison exercise using data from a 2016 general survey of all active distributors (conducted independently by FoodCo).[12] This survey includes 55 distributors who later joined our experimental sample in 2017 and 2018. Appendix A4 compares the characteristics of those 55 distributors with the characteristics of the general population of active distributors. We find no significant differences between the experimental sub-sample and the broader population of distributors in terms of age, ethnicity, religious affiliation, marital status, education level, monthly business profits from all sources, annual household income, and a household asset index. Our experimental sample — which focused on distributors interested in acquiring a bicycle — did have a lower proportion of females compared to the broader sample (27%).

2.4 Conceptual framework

To guide intuition, we now discuss the trade-offs facing a stylized micro-distributor. Specifically, we consider a distributor who is credit-constrained, and whose productivity will increase if she acquires a bicycle. The distributor, faced with our menu of financing 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?". 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 reflects important features of the actual experience of distributors in our experiment. First — using incentivised baseline

[11] We use Ks throughout to refer to Kenyan Shillings (KES). The USD-KES exchange rate at baseline was approximately equal to 102.

[12] This survey had also provided motivation for our project, as it highlighted that many distributors expressed dissatisfaction with the materials and equipment available to help them reach their business potential.

behavioral measures — we find that distributors are risk averse.[13] This implies that, *ceteris paribus*, distributors value a contract that bundles some degree of risk-sharing. Second, distributors operate in a risky environment — with the risk increasing along with the distributors' use of the lumpy asset. This feature, too, is closely grounded in the real experience of our respondents. For example, a distributor who cycles her bicycle further to serve new markets may increase and diversify her sales — but 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).

To formalise these ideas, we assume that the distributor has an exponential utility function with r being the coefficient of absolute risk aversion. We model the distributor's net income as the sum of income that does not depend upon the bicycle at all ('sure thing' sales: a constant, π_0), and income that depends upon a Cobb-Douglas form in on-contract labour effort (e) and capital (where $k \equiv 1$ for no bicycle, and $k >> 1$ for a bicycle): $\pi_1(e, k, \eta_t) = \eta_t \cdot e \cdot k$. Further, for simplicity, we assume that the distributor has neither any credit nor any savings technology; this accords approximately with the empirical reality (in which the respondent distributors were unable to accumulate sufficient funds to purchase the bicycle without a financing contract), and allows the model to focus clearly upon the consequences of the experimental contracts. Finally, we allow that a distributor who is debt-free enjoys a per-period gain of $\phi \geq 0$; this could describe a psychic gain (for example, due to debt aversion (Azmat & Macdonald, 2020; Martínez-Marquina & Shi, 2024; Paaso et al., 2020), or the endowment effect (Carney et al., 2022)), or reflect the compliance costs of needing to meet a loan officer on a regular basis.

Suppose that, each month, the distributor must pay a fixed sum F and then a proportion $(1-\omega)$ of her total net income. We assume that, each month (in advance of the realisation of η_t), the distributor chooses her effort. We allow a quadratic effort cost (in currency-equivalent terms); this reflects both the psychic cost of effort and the opportunity cost of the distributor's time on other projects (including, in particular, sales off-contract). This is a stationary problem; with monthly discount factor β, the distributor's infinite-horizon value is:

$$V(k, F, \omega) = \frac{1}{1-\beta} \cdot \max_{e \geq 0} \; \mathbb{E}_\eta \left(u \left\{ \omega \cdot \left[\pi_0 + \pi_1(e, k, \eta) \right] - 0.5e^2 + F \right\} \right). \tag{1}$$

Equation 1 can be used to describe four important cases. First, the value of refusing any financing contract; in this case, the distributor has no bicycle ($k = 1$), keeps all of her own income, and enjoys being debt-free ($\omega = 1$; $F = \phi$). The value can therefore be written as $V(1, \phi, 1)$. Second, the value of completing a financing contract; this resembles the contract refusal case, but the distributor

[13] For example, using incentivised risk preference elicitation activities, we find that — for a binary outcome lottery with expected payment of Ks 500 — the average certainty equivalent was Ks 374; for a lottery with expected payment of Ks 750, the average certainty equivalent was Ks 478. A simple structural estimation of $u(x; \alpha) = x^\alpha$ using all the data from our incentivised games returns $\hat{\alpha} = 0.69$.

has a bicycle: $V(k, \phi, 1)$. Third, the debt contract involves 12 months of fixed repayments of F_d, after which the client owns the bicycle; the initial value of taking that contract is therefore $\left(1 - \beta^{12}\right) \cdot V(k, -F_d, 1) + \beta^{12} \cdot V(k, \phi, 1)$. Fourth, by analogous logic, the initial value of taking the revenue-sharing contract is $\left(1 - \beta^{12}\right) \cdot V(k, -0.5F_d, \omega) + \beta^{12} \cdot V(k, \phi, 1)$.

Three insights flow immediately from this setup for the distributor's effort with the bicycle. First, since the marginal product of effort is increasing in k, the model predicts greater effort with the bicycle than without. Second, the model predicts similar effort for a distributor on Debt as for a distributor who owns the bicycle outright (because, in both cases, the distributor receives all of the income — and the fixed repayments F_d do not affect the marginal return to effort under constant absolute risk aversion).[14] Third, for reasonable values of risk aversion, RevShare reduces effort by 'taxing' the distributor's returns; as Angrist, Caldwell, and Hall (2021) elegantly put it, output sharing 'inserts a wedge between effort and income'.[15] Together, these results indicate that performance-contingent repayment structures require a trade-off: they are valued for their implicit insurance, but this comes at the cost of reduced effort. Indeed, this is a familiar trade-off from many studies of performance-contingent remuneration (Holmström, 1979; Lazear, 2000), including in the famous case of sharecropping: Burchardi et al. (2019); Stiglitz (1975); Stiglitz and Weiss (1981).

IndexShare offers one potential mechanism for breaking this trade-off — by providing implicit insurance based on the income shocks of other distributors, while eliminating the taxation of individual effort. Specifically, effort under IndexShare is expected to approximate effort under Debt (where, again, the distributor enjoys all of the gains from her own effort).[16] However, demand for IndexShare will depend upon how well the index correlates with the shocks that the client faces: if the index does not correlate closely, IndexShare exposes clients to substantial basis risk (see, for example, Karlan, Osei, Osei-Akoto, & Udry, 2014).

[14] Formally, because of the assumption of constant absolute risk aversion, the model predicts identical effort. This follows straightforwardly from equation 1, which can be rewritten as follows (emphasizing that the optimal choice of effort is invariant to the level of fixed repayments):

$$V(k, F, \omega) = \frac{1}{1-\beta} \cdot \exp\left(r \cdot F\right) \cdot \max_{e \geq 0} \; \mathbb{E}_{\eta} \left\{ -\exp\left(-r\omega \cdot \Big[\pi_0 + \pi_1(e, k, \eta)\Big] + 0.5re^2\right) \right\}.$$

If this assumption were relaxed, the effort may then differ slightly — but our model emphasizes that this difference would be due solely to relatively small wealth effects, and not likely to be large.

[15] In the appendix, we use a second-order approximation to argue that this result need not be universal; a distributor who is extremely risk averse — indeed, probably implausibly risk averse, in this context — might increase effort as ω decreases, because of an implicit-insurance channel.

[16] The model framework presented here can be extended to think about IndexShare; if we represent by θ the payments owing based upon the index, equation 1 can be modified as follows:

$$V(k, F, \omega) = \frac{1}{1-\beta} \cdot \max_{e \geq 0} \; \mathbb{E}_{(\eta, \theta)} \left(u\left\{ \omega \cdot \Big[\pi_0 + \pi_1(e, k, \eta)\Big] - 0.5e^2 - 0.5F_d - \theta \right\} \right).$$

The Hybrid contract provides an alternative way of breaking this trade-off: Hybrid offers repayment flexibility (and, thus, some implicit insurance) and, in the case $\phi > 0$, incentivises additional effort in order to clear the distributor's debt earlier. Under Hybrid, increased on-contract effort in any given month can bring forward the date at which the total contract is repaid — and, therefore, change the path of future repayments. Hybrid can, therefore, be understood as a dynamic optimisation problem with the outstanding debt (D_t) being the state variable. For a distributor entering a given period with outstanding debt D_t, the value of Hybrid can be written as follows:

$$V^h(D_t) = \max_{e \geq 0} \mathbb{E}_\eta \left[u \left(\max \left\{ \underbrace{\omega \cdot \Big[\pi_0 + \pi_1(e, \eta; k)\Big] - 0.5F_d}_{\text{contract ongoing}}, \underbrace{\pi_0 + \pi_1(e, \eta; k) - D_t}_{\text{contract ending/ended}} \right\} - 0.5e^2 \right) + \beta \cdot V^h(D_{t+1}) \right], \tag{2}$$

where the law of motion for D_t is:

$$D_{t+1} = \max \left\{ \underbrace{D_t - 0.5F_d - (1-\omega) \cdot \Big[\pi_0 + \pi_1(e, \eta; k)\Big]}_{\text{contract ongoing}}, \underbrace{0}_{\text{contract ended}} \right\}. \tag{3}$$

We can solve $V^h(D_t)$ numerically by backward induction in D_t; we show the solution graphically in the appendix.[17] Three features of the solution deserve discussion. First, where $\phi = 0$, average effort under Hybrid lies between effort on Debt and effort on RevShare. (As $\beta \to 1$, the average effort under Hybrid approximates the average effort under Debt — since the total repayment under Hybrid matches that of Debt and, by experimental design, the expected monthly payment is approximately equal to that of Debt.) Second, even under this case where $\phi = 0$, Hybrid can be preferred to Debt, owing to its flexibility. Third, if respondents have a strong desire to clear their debt ($\phi >> 0$), Hybrid additionally incentivises effort to achieve this — and, in doing so, is particularly valued by distributors.

As we noted in section 2.1, distributors require up-front financing for inventory purchases — and distributors do not receive trade credit from stockpoints. Our model could be extended to incorporate liquidity constraints and inventory financing requirements by introducing cash-on-hand as an additional state variable. Introducing this channel would not change qualitatively the model predictions — but would further emphasize the potential advantages to performance-contingent finance. In particular, under liquidity constraints, the fixed repayments required by Debt would be more onerous relative to the flexible repayments allowed by performance-contingent contracts.[18]

[17] First, note that the terminal value is known: once the debt is repaid, the distributor owns the bicycle outright, so $V^h(0) \equiv V(k, \phi, 1)$. Second, the state dynamics are monotonic: the total debt always decreases until it is repaid (if $D_t > 0$, it follows that $D_{t+1} < D_t$).

[18] We thank an anonymous referee for this insight.

3 Treatment effects

To measure the impacts of our treatments, we use a combination of administrative data (available directly from FoodCo stockpoints) and face-to-face surveys (which we collected each quarter, in person, for up to a year after treatment). Our data covers all available post-treatment months until the COVID-19 lockdowns in March 2020.[19] For each outcome, we use an intent-to-treat ANCOVA specification:

$$y_{it} = \beta_0 + \sum_{k \in \{1,\ldots,4\}} \beta_k \cdot \texttt{Offered}_{ik} + \gamma \cdot y_{i0} + \varepsilon_{it}. \tag{4}$$

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 data on profits). We cluster standard errors at the individual level and include month fixed effects in the regressions with administrative data. Given the skewed nature of the profits variable, we winsorize at 90% and check for robustness at several alternative levels of winsorization (99%, 97.5%, 95%, 92.5%). We also test for robustness using Poisson regressions and randomization inference.

3.1 First-stage: Take-up, bicycle ownership, and household finances

Table 1 presents results for take-up of the financing contracts offered, and the impacts of treatments on debt levels and asset ownership. Panel A presents results pooling all financing contracts, to explore the overall impact of being offered any form of asset-based financing. Panel B presents regressions with separate dummies for each contract.

We begin by describing take-up — by which we mean that a respondent had agreed to an offered contract, provided the requisite 10% deposit and supporting documentation, and received the bicycle. Results are presented in column 1. The overall take-up rate of any financing contract was 57.6%. The take-up rates for Hybrid and Debt were similar, at 69.2% and 67.7%, respectively. Take-up rates for RevShare and IndexShare were lower, and similar, at 48.8% and 46.9%, respectively. Formal statistical tests indicate that take-up of Hybrid is significantly greater than take-up of RevShare (p = 0.087) and IndexShare (p = 0.077). A formal test also indicates that take-up of Debt is significantly greater than take-up of RevShare (p = 0.092) and IndexShare (p = 0.082).

[19] We concluded the project in March 2020, having collected approximately 85% of the planned follow-up survey data before the lockdown – for 161 enrolled participants, below our original target sample size of 250 as documented in our pre-analysis plan. As in many other contexts, the lockdown caused significant disruption; in our case, it affected not only the operations of distributors but also led to structural changes in how FoodCo managed the program and hindered the MFI's ability to collect microfinance repayments. All of our analysis uses data up to, but not including, the lockdown period. For the survey data, attrition is very low, with an overall attrition rate of approximately 4%, which is uncorrelated with the treatments. There is no attrition in the administrative data—a zero represents an actual zero, indicating a month in which the distributor did not purchase any stock.

To support these empirical results, we collected qualitative data on reasons for rejecting different contracts. For Debt, the most common reason — cited in 35% of rejection cases — was the desire to be able to end the contract early. (Only 4% of those refusing Hybrid provided this explanation.) This is consistent with our conceptual framework — in which the parameter ϕ reflects a desire to complete the contract early.[20]

Column 2 does not reveal any notable patterns in the amounts financed across contracts; the average financed amount (Ks 8,698) is large relative to the average household debt level at baseline (Ks 2,498). Columns 3 and 4 explore whether our treatment changed the overall household stock of debt (excluding the amount outstanding under our financing product); we do not indicate any significant treatment effect on overall household debt. Columns 5 and 6 reassuringly show that, in the year following our intervention, the treatment group is significantly more likely to own a bicycle than the control group (for whom the mean ownership rate during the follow-up period is 7.8%). This suggests no significant sales of the financed asset; the ownership rates reflect the pattern of take-up rates from column 1.

3.2 Impact on business performance

Table 2 presents results for the main outcome: business profits. Panel A again presents pooled results; Panel B presents results by contract (including both ITT and LATE estimates).

Our primary hypothesis, as specified in our pre-analysis plan, was that our treatments affected participants' business profits. Column 1 pools all follow-up data (up to three years after treatment); it shows that assignment to treatment nearly doubled monthly business profits from the sale of FoodCo products, as measured from administrative data. Specifically, we estimate an ITT of Ks 792 per month (SE: 386); this compares to a control mean of Ks 897. Given the skewed nature of the outcome variable, in column 2 we show robustness to using a Poisson specification, and reach a similar conclusion: a coefficient of 0.65 (SE: 0.308) represents an increase in profits of approximately 92% (that is, $\exp(0.65) - 1$). In column 3, we report the LATE: we estimate a 132% increase in monthly business profits from taking up a contract (Ks 1,181; SE: 562). Compared to the average asset price of Ks 9,658, these large treatment effects suggest very favorable benefit-cost ratios, which we explore more systematically in Section 4.

Columns 4, 5, and 6 explore the treatment effects over time, using LATE estimates. Specifically, we restrict the time period to (i) months one to six after delivery of the assets; (ii) months seven to 12, and (iii) months 13 to 24. First, it is clear that large treatment effects appear soon after the

[20] The relatively lower take-up of IndexShare is unsurprising, given the potential role of basis risk (Carter, de Janvry, Sadoulet, Sarris, et al., 2014; Clarke, 2016; Cole et al., 2013). In the appendix, we highlight the relationship between distributor performance and required payments under each contract, confirming the role of basis risk. Indeed, our qualitative data supports this: several distributors, in refusing the IndexShare contract, expressed concerns about the index being tied to others' performance, and the risk of owing large payments unrelated to their own sales.

asset is disbursed. Second, the magnitude of the coefficient remains relatively consistent over time, indicating that the effects do not dissipate quickly and exhibit some level of persistence over the three-year period. Finally, in column 7, we find little evidence that the increase in profits from FoodCo products, as observed in the previous columns, crowded out other sources of income (including profits from selling non-FoodCo products and wage income). The estimated LATE coefficient is 613, though the standard error is large (SE: 1,690).

In Panel B, we disaggregate by financing contract type. The top-performing contract is Hybrid, and the following discussion focuses on the outperformance of Hybrid compared to the more standard Debt contract. We find that Hybrid has a large and statistically significant impact across all time periods, and in nearly all specifications, this effect is significantly greater than that of Debt. Specifically, in column 1, the estimated ITT for Hybrid (using the full three years of data) is Ks 1,529 (SE: 609), compared to an estimate for Debt of Ks 530 (SE: 435), with a p-value of 0.091 for the cross-coefficient test. The Poisson regression in column 2 shows an even stronger statistical difference between the two, with a coefficient of 1.10 (SE: 0.346) for Hybrid and 0.23 (SE: 0.428) for Debt, and a p-value of 0.021 for the cross-coefficient test. Columns 3 to 6 show similar results for the LATE estimates over the full three-year period, and also explore dynamic treatment effects. Specifically, Hybrid significantly outperforms Debt in (i) months one to six, (ii) months seven to 12, and (iii) months 13 to 24, with p-values of 0.076, 0.022, and 0.096, respectively. In Appendix A3.6, we demonstrate that these conclusions remain robust when using a Poisson specification; we again find large and stable treatment effects for Hybrid over time, with cross-coefficient tests confirming that Hybrid consistently outperforms Debt. In Appendix A3.7, we also show that the results are robust to the use of randomization inference. In Appendix A3.8, we demonstrate robustness at several alternative levels of winsorization (99%, 97.5%, 95%, 92.5%).

Could these differences between Hybrid and Debt be driven by differences in the composition of respondents accepting offers? We can answer this question in two complementary ways: (i) by allowing for heterogeneity on observables and (ii) by allowing for heterogeneity on unobservables. First, to allow for heterogeneity on observables, we repeat in Table 3 the analysis from Table 2, now incorporating controls for demeaned baseline values of total profits, risk aversion, and loss aversion, as well as the interactions between these demeaned variables and each treatment indicator. If — for example — our earlier results were driven by heterogeneity in take-up along these dimensions, this exercise would generate very different results to the original regressions. However, to the contrary, Table 3 shows that all of the previous results remain robust (indeed, the precision of many estimates increases).

Second, in the alternative, we allow for heterogeneity on unobservables by imposing a standard Lee (2009) monotonicity assumption (implying, in this context, that any respondent who takes Debt would also take Hybrid). Recall that, in Table 1, we found almost identical take-up rates between Hybrid and Debt (69.2% and 67.7% respectively, and not significantly different). This fact alone

makes it highly implausible — given the monotonicity assumption — that the difference in profits between Hybrid and Debt would be driven by differences on the extensive margin (that is, heterogeneity in contractual take-up). We conclude that the difference in profits is likely driven by differences in effort on the intensive margin (that is, an increase in profits conditional on adopting the contract). In Appendix A5, we show robustness of this conclusion using a formal Lee (2009) bounding exercise.

The previous discussion has emphasized particularly the performance of Hybrid compared to the more standard Debt contract. It is worth noting that the coefficients on RevShare — the other performance-contingent contract — are also large in several of the specifications in Table 2 (and, in several cases, are significantly different from the control). In particular, the estimated treatment effect of RevShare is quite stable across time periods: 981 (SE: 709) in months one to six, 1194 (SE: 722) in months seven to 12, and 1100 (SE: 804) in months 13 to 24. Finally, the coefficients on IndexShare are consistently small in magnitude, suggesting that liquidity risk (mitigated under Hybrid and RevShare but not under IndexShare) may play a more significant role in this context than the adverse incentive effects of effort taxation, which IndexShare addresses.

3.3 Contract repayments

We next analyze repayments to the capital provider under each of the financing contracts, for those individuals who took up the contract. Figure 2 plots the average repayment amount over time for each of the four financing contracts. It is evident — from as early as the first quarter — that repayment under Hybrid begins to increase above that of other contracts. This is consistent with our previous finding of higher treatment effects of Hybrid on business effort and performance. Note that this is not a mechanical result — contract payments were *not* deducted automatically at source by the multinational.[21] Rather, payments are manually made by distributors via M-Pesa. This is reassuring, and consistent with our empirical results that individuals under Hybrid had greater *ability* to pay due to greater impacts on profits.

By month six, average payments under RevShare also begin to diverge from payments under Debt and IndexShare. Once again, this finding is reassuring and aligns with our previous results, which suggested that the second-largest treatment effect in terms of business profits was observed under RevShare. Cumulative payments under RevShare converge with those under Hybrid by month nine. This outcome is as expected, given that Hybrid limits upside sharing to the total due amount of an equivalent Debt contract. In contrast, RevShare allows some individuals to pay significantly larger amounts, surpassing the initial capital plus 15% interest. A clear dichotomy becomes evident in month 10, where cumulative payments under Hybrid and RevShare converge — while diverging from cumulative payments under Debt and IndexShare. By month 15, payments under Hybrid and RevShare

[21] One may want to implement such a model when scaling up such contracts, but we were not able to implement that change in the payment system in time for this project.

are nearly equal and notably higher than cumulative payments under Debt and IndexShare, which are also almost identical in absolute terms.

Figure 3 illustrates the average total repayment under each of the contracts after 15 months. Recall that the contractual duration of Debt, RevShare, and IndexShare is 12 months, but the duration of Hybrid may be as little as one month, or as many as 24 months, depending on performance. None of the contracts generated 100% repayment. This is partly due to the Covid-19 shock; although the majority of repayment delinquency began before Covid-19, the latter part of our project coincided with the pandemic — during which the Kenyan Central Bank asked banks to provide relief to borrowers. The MFI was therefore not able to apply its standard enforcement procedures — which, ordinarily, would have resulted in significantly higher collection of outstanding amounts due. More generally, this is consistent with evidence that default rates under digital credit are higher than traditional cash-based lending (Carlson, 2017; Suri, Bharadwaj, & Jack, 2021).[22] That said, the magnitude of the MFI's loss is very small compared to the positive treatment effects for distributors outlined in section 3.2. Specifically, the average default amount is approximately Ks 3,000, compared to a treatment effect on *monthly* profits of Ks 1,182 (Panel A, column 3 of Table 2), indicating highly favorable benefit-cost ratios over time, which we explore further in section 4. Figure 3 reveals that repayment under Hybrid is 78% of the total capital disbursed on average, and repayment under RevShare is 81% on average. This is significantly higher than average repayment under Debt (59%) and IndexShare (58%) (p=0.055 for a formal test that the repayment rate under Hybrid and RevShare is the same as that under Debt and IndexShare).

3.4 Mechanisms: business practices and effort

To shed light on the mechanisms underlying (i) the positive overall treatment effects observed in Table 2, and (ii) the particularly large effects under Hybrid, Table 4 explores these mechanisms through the channels of business practices and effort. In columns 1 and 2, we address a key initial question: Did the distributors who received our bicycle actually use it for their business? In general, the answer is yes: the pooled estimates in Panel A indicate that 78% of distributors who took the bicycle mainly used it for business purposes, and the average number of hours that the bicycle was used per week was 27.3 (SE: 2.0). Turning to cross-contract comparisons in Panel B, we again find that Hybrid is the standout performing contract, with greater asset utilization for business purposes and increased effort. Specifically, 90% (SE: 4.3%) of individuals under Hybrid used their bicycle primarily for business

[22] There is evidence from several African settings of higher default rates for digital credit. For example, Brailovskaya, Dupas, and Robinson (2021) report from a digital credit experiment in Malawi that 11% of loans were never repaid, 4% were partially repaid, 47% were fully repaid but late, and only 38% were fully repaid on time. Kruijff et al. (2024) report from a nationally representative survey in Côte d'Ivoire that 78% of digital borrowers repaid their loans late. The authors also cite similar surveys in Kenya and Tanzania, which found late repayment rates of 47% and 56% respectively.

purposes, compared to 73% (SE: 5.3%) under Debt.[23] A formal test confirms that the difference is statistically significant, with a p-value of 0.014. Average weekly bicycle usage was 33.3 hours (SE: 5.2) under Hybrid, significantly larger than the 21.8 hours (SE: 2.0) under Debt ($p = 0.037$ for a cross-contract test).

Column 3 explores an administrative measure of business effort: how often distributors visit stockpoints in a given month to purchase inventory. The pooled estimate in panel A indicates an increase of 2.3 visits per month (SE: 1.3), compared to a control mean of 2.6, implying an 89% increase in visits for those who took up the treatment. In panel B, the highest coefficient is again on Hybrid, with a value of 3.7 (SE: 2.0); this compares with a coefficient on Debt of 1.9 (SE: 1.5), although we are not able to formally reject equality of the coefficients.

Column 4 explores a different measure of effort, captured by survey data asking distributors about the percentage of their selling portfolio that comes from customers greater than 1 km from their stockpoint. Results indicate that the treatment led to a large geographical expansion of customers, with a coefficient of 28 percentage points (SE: 12 percentage points) on the pooled estimate, which implies that distributors who took up any contract now generate 85% of their profits from customers greater than 1 km from their own stockpoint, compared to a control mean of 57%. There is no significant difference in the estimates across contracts; the individual estimates are large and individually significant versus control for all of the contracts except Debt. These results, indicating a substantial expansion of distributors' sales networks, align with the data captured from GPS trackers installed on all bikes. As illustrated in Figure 4, the GPS data from various project implementation sites across the country reveal that distributors covered vast geographical areas with their bicycles, highlighting the extensive reach of their operations.

In columns 5 to 7, we explore the impact of the contracts on business management practices. In column 5, we use an index of overall business management practices, comprising questions on marketing, negotiation, cost, record-keeping, and sales targeting. The questions are based on McKenzie and Woodruff (2015), amended for a micro-distribution business. Results suggest that individuals assigned to Hybrid and IndexShare experienced the greatest positive impacts on overall business management practices, with coefficients of 0.13 (SE: 0.07) and 0.21 (SE: 0.11) standard deviations.[24] One plausible explanation for why we see impacts on these contracts in particular is that they are the two contracts that require the greatest amount of 'mental engagement' in calculating payments: Hybrid requires clients to pay a proportion of their monthly income and to carry forward the 'state variable' (as modelled in our conceptual framework) of cumulative payments made to date and the re-adjusted notional debt outstanding, and IndexShare provides sharing based on the average sales of all other distributors

[23] The vast majority of individuals who took up the bicycle report that they primarily used it themselves — only 7% report that someone outside of their household used it for any period of time.

[24] Each index ranges from zero to one, indicating the proportion of questions that receive a positive response regarding whether a specific business management practice is undertaken.

in one's region. Column 6 provides evidence consistent with this hypothesis; there, we use a specific sub-category of questions that measure record-keeping, and we again find positive and statistically significant effects only on Hybrid and IndexShare, with coefficients of 0.20 (SE: 0.09) and 0.23 (SE: 0.14) standard deviations, respectively.

In column 7, we analyze one particular business practice that relates to distributors' risk-taking: the extent to which they offer credit to their own customers. Financial contracts that provide a greater extent of risk-sharing may themselves allow business owners to take more risk (Karlan et al., 2014). We find evidence that is consistent with this for Hybrid, which is again the contract with the greatest coefficient magnitude. The increase for Hybrid is a relatively large 6.0 percentage points (SE: 3.3 percentage points). The control group mean of 9.0% indicates that distributors typically extend very little credit to their customers, and that Hybrid led to greater risk-taking through credit extension.

Finally, we again allow for heterogeneity on observables. We repeat in Table 5 the analysis from Table 4, now incorporating the controls for demeaned baseline values of total profits, risk aversion, and loss aversion, as well as the interactions between these demeaned variables and each treatment indicator. All of the previous results remain robust, and again the precision of most estimates increases.

3.5 Downstream outcomes: Consumption

Table 6 presents the treatment effects on three major components of household consumption expenditure. Starting with food expenditure, column 1 shows coefficients that are large in magnitude relative to the control mean of Ks 4,626 per month, though accompanied by large standard errors. The largest coefficient is for IndexShare (Ks 1,545; SE: 985), followed by Debt (Ks 1,230; SE: 691) and Hybrid (Ks 707; SE: 726), with cross-coefficient tests failing to reject equality of coefficients.

In column 2, there is a notably large treatment effect on monthly household expenditure on clothing specifically for Hybrid (Ks 666; SE: 306), compared to a control mean of Ks 909. The cross-coefficient test rejects equality with the coefficient on Debt of Ks 82 (SE: 266), with a p-value of 0.043.

Column 3 examines the effect on household expenditure on schooling. Relative to a control mean of Ks 1,113, the coefficient for Hybrid is substantial at Ks 535 (SE: 565), and significantly larger than the coefficient for Debt of -Ks 425 (SE: 516), with a p-value of 0.065 from the cross-coefficient test.

In the appendix, we present results on health, happiness, and trust, as specified in our pre-analysis plan. While the standard errors are large, some coefficient magnitudes are substantial. For instance, regarding health, one key motivation for providing bicycles was respondents' concerns about carrying large bags on their backs, as identified in our qualitative work — this concern was a central reason for introducing a transportation asset in this experiment. In columns 1 and 2 of Appendix A9, we examine the impact of treatments on binary indicators for whether distributors report that their health impedes their work, and whether work caused physical pain. Compared to the control means of 26% and

19%, respectively, the estimated coefficients for most contracts are meaningfully large and negative, indicating that respondents are *less* likely to report health problems. Similarly, in columns 3, 4, and 5, we observe large positive coefficients representing the effects of treatments on happiness related to (i) income, (ii) ability to meet expenditure demands, and (iii) adequacy of work materials and tools for conducting sales work. However, these standard errors are again large. Nonetheless, the overall pattern of the coefficients offers some suggestive evidence of improvements in health and happiness, though there is no evidence of an increase in trust in others.

3.6 Robustness: Spillovers and GPS trackers

Next, we test for spillovers. There are two plausible mechanisms by which such spillovers might operate. First, treated respondents might take business from control participants — or, indeed, from distributors outside of the experiment. (We would expect this to bias upward our main estimates.) Alternatively, distributors on performance-contingent contracts might engage in 'side-selling' — purchasing their gum through peers to avoid increasing their contract payments. (We would expect this to bias downward our estimates of the impact of Hybrid and RevShare.)

To test for spillovers, we use administrative data on distributors who were in FoodCo's program but not in our experiment. Between 2017 and 2019, there were 1,727 unique distributors in the FoodCo program; we have daily data on all of their purchases of all gum products from FoodCo, and we can use this to test directly for spillovers. To do this, we exploit detailed baseline data in which 100 of our participants answered a series of dyadic questions about the extent of their relationship (if any) with other distributors at the stockpoint. In Appendix A2, we discuss this analysis and show results. In that appendix, we conclude that there are no meaningful spillover effects. (Nonetheless, as a further robustness exercise, we allow for varying degrees of spillovers as part of the benefit-cost analysis in section 4. There, we find high benefit-cost ratios and high internal rates of return even under very conservative assumptions about spillovers.)

The lack of meaningful spillover effects is consistent with the provision of bicycles having expanded the geographical reach of the distributors, as indicated in the results from column 4 of Table 4 (which showed, using survey data, a large increase in the likelihood of treated distributors selling to customers further than 1km away from their stockpoint). Figure 4 provides further evidence of this expansion using GPS data from trackers installed on all bicycles (with clients' consent). The maps illustrate that the bicycles were widely dispersed across Kenya's most populous areas, with distributors traveling considerable distances within their regions. In particular, Panel B plots all unique latitude and longitude points collected from the GPS data for distributor location, after filtering outliers, and compares them to the locations of all stockpoints, including those not involved in the experiment. The analysis reveals that distributors frequently ventured far from stockpoints, supporting the hypothesis that they accessed new markets rather than merely competing with non-treatment distributors (who are

more likely to travel on foot and service customers closer to stockpoint locations).

4 Cost-benefit analysis

4.1 Estimating the total return along the multinational supply chain

We now estimate cost-benefit ratios and internal rates of return (IRR), building upon the methodology of Banerjee, Duflo, et al. (2015), Bandiera et al. (2017), Alfonsi et al. (2020), and Bari, Malik, Meki, and Quinn (2024). Four distinct market participants contribute to the returns factoring into our cost-benefit calculations:

(i). The micro-distributors;

(ii). The multinational that produces the product, FoodCo;

(iii). The stockpoints purchasing gum from FoodCo and selling it to the distributors, earning their own margin from each sale;

(iv). The external capital provider responsible for providing the financing contracts, and bearing all of the contract repayment risk.

We begin by considering the returns for the three key stakeholders in the multinational supply chain; in the cost-benefit calculations in 4.2, we incorporate the returns for the capital provider.

For distributors, we adjust the numbers from our previous analysis, which represented the administrative data measure of gross profits: the sale price for each product, minus the cost of goods sold. The gross profit therefore does not include other costs that are commonly referred to as "selling, general, and administrative (SG&A)" in standard financial reporting. We now use survey data to approximate such costs, to calculate income after all operating expenses, but before distributors pay themselves. In our in-person surveys, we asked distributors about all of their selling activities (FoodCo and non-FoodCo). Distributors reported that the cost of raw materials represented by far the largest cost category, at 85% of total operating costs on average. The next biggest category was transportation, comprising 7% of total costs on average. The remaining categories consisted of: phone airtime and data costs, payments to any employees they have (only 9% of distributors hire anyone else), selling permit fees, and any other bills or expenses. We use these estimates of other costs to convert our previously used administrative measure of gross profits, which already included inventory costs, into a measure of operating income that we can compare with the operating income of FoodCo and stockpoints.

For FoodCo, we use a straightforward method to approximate the operating income they earn for every dollar of operating income earned by distributors, once again drawing from standard accounting practices and measures of costs and profits. For each of the six possible products sold by distributors, we calculate the value of sales generated by FoodCo based on information provided to us about the

price paid per product by stockpoints (from who distributors purchase their products). To get from the value of sales to an estimate of operating income, we use data from the last three years of publicly available annual financial reports for the company. The 3-year average gross profit margin (revenue minus cost of goods sold, as a percent of sales) was 52.9%. The annual report documented expenditure by the company of 10.7% of sales on advertising, 5.4% on merchandising and promotions, 10.4% on selling and marketing costs, and 8.5% on general and administrative costs, leading to a final operating income of 18% of sales. We apply that ratio to our administrative measure of the value of sales for FoodCo and find that, for every 1 Shilling of operating income earned by distributors, FoodCo earns 3.25 Shillings. In Appendix A10, we show that this is not affected by our treatments.

For stockpoints, we have information on their gross profits from administrative data. We do not have any information on their costs. We therefore use the cost ratios documented in FoodCo's annual report, excluding advertising and merchandising costs, which are not relevant for the stockpoints in consideration. Using an assumed gross profit margin of 52.9%, and 18.9% in all other operating costs, leaves us with a 34% operating income margin that we apply to the value of sales we back out for stockpoints from the administrative data. We find that, for every 1 Shilling of operating income earned by distributors, stockpoints earn 0.95 Shillings. In Appendix A10, we again show that this is not affected by our treatments.

Table 7 illustrates the returns to the three key stakeholders in the multinational supply chain. Columns 1 and 2 show the treatment effects on our new measure of operating income for distributors, using ITT and LATE specifications, respectively. Although the operating income is lower in magnitude than our previously used gross profits (due to the subtraction of further estimated costs — for example, the control mean goes from 897 in Table 2 to 521 in Table 7), the treatment effects remain large: the ITT coefficient is Ks 471 (SE: 221), and the LATE coefficient is Ks 708 (SE: 325).

Columns 3 and 4 present the equivalent estimates for FoodCo's operating income. As noted, our calculations suggest that FoodCo earns just over three times each shilling earned by distributors, which is reflected in the coefficient estimates: the ITT coefficient is Ks 1,534 (SE: 718), and the LATE coefficient is Ks 2,305 (SE: 1,057). Finally, in columns 5 and 6, we observe a large return for the stockpoints, mirroring the return of distributors: the ITT estimate is Ks 449 (SE: 210), and the LATE estimate is Ks 674 (SE: 309).

Columns 7 and 8 aggregate the returns across the three actors in the supply chain, showing that — particularly relative to the bike cost of Ks 9,000 on average — the total monthly return generated was substantial, with an ITT estimate of Ks 2,454 (SE: 1,148) and a LATE estimate of Ks 3,687 (SE: 1,691).

4.2 Benefit-cost ratios and IRR

We now combine these results to conduct a more detailed analysis of the benefit-cost ratio under different assumptions about the persistence of treatment effects after the three-year period of the project.

Beginning with costs, these comprise: (i) the capital disbursed for the initial asset purchases for take-up clients, subtracted from the total recovered capital (factoring in the small overall loss to the MFI, as discussed in Section 3.3); (ii) staff salaries; and (iii) other implementation expenses like venue rentals for workshops. The total costs are then compounded up to the two-year mark using a conservative 10% social discount rate.[25] We divide the total costs by the number of take-up clients in each contract and then incorporate the benefits from each contract. We provide further details of all our assumptions in Appendix A12.

For benefits, we first take the coefficient estimates from our LATE estimations in Table 7 as the total benefits from the intervention for the first three years after implementation. We then estimate the net present value of future benefits with various assumptions regarding the persistence of effects beyond the three-year period, ranging from zero to 10 years, consistent with the sensitivity analysis in comparable cost-benefit literature.

Figure 5 presents the results. We find large mutual benefits along the supply chain, and remarkably high benefit-cost ratios across all contracts — particularly for Hybrid, even when assuming minimal persistence of treatment effects. For instance, for the pooled estimate, we find a benefit-cost ratio of 6.3 when assuming zero years of treatment effect persistence (confidence interval: 1.5 to 11.1), corresponding to an IRR of 203%. This increases to 7.8 when assuming five years of persistence (confidence interval: 1.9 to 7.8), with an IRR of 210%. For Hybrid, we find a benefit-cost ratio of 10.8 with zero years of persistence (confidence interval: 3.2 to 18.4), corresponding to an IRR of 356%. This increases to 13.4 when assuming five years of persistence (confidence interval: 4.0 to 22.8), with an IRR of 360%. Given these remarkably high treatment effects, in the conclusion, we discuss the constraints to FoodCo realizing these high returns.

As a further robustness exercise, we allow for varying degrees of spillovers, which we simulate by reducing the magnitude of the treatment effect from that which we estimated in the previous analysis. In Appendix A14, we show that, even when reducing the treatment effect by 25%, we find a benefit-cost ratio for the pooled estimate of 4.7 when assuming zero years of treatment effect persistence, corresponding to an IRR of 147%, increasing to 5.9 when assuming five years of persistence (IRR: 157%). In Appendix A15, we show that even a 50% reduction in treatment effects leads to a benefit-cost ratio for the pooled estimate of 3.2 when assuming zero years of treatment effect persistence (IRR: 90%), and 3.9 when assuming five years of persistence (IRR: 104%). We conclude that our finding of very high benefit-cost ratios and internal rates of return are robust to even very conservative assumptions about spillovers reducing effects sizes.

[25] This rate falls within the range recommended by the World Bank (Lopez, 2008).

5 Conclusion

We ran a field experiment within one of the world's largest food manufacturers, to test whether the firm could facilitate productive asset investment for its credit-constrained distributors. Our results reveal substantial positive impacts along the supply chain — benefiting distributors, stockpoints, and the multinational — with highly favorable benefit-cost ratios and internal rates of return. We find particularly large gains from a novel hybrid contract that combines debt-like features with performance-contingent payments. To interpret these findings, we use a dynamic stochastic model to understand how the hybrid contract can break the traditional trade-off between implicit insurance and reduced effort: the contract provides repayment flexibility and implicit insurance while also incentivizing additional effort to clear the debt.

Our setting was an ideal one in which to test the effectiveness of performance-contingent contracts for productive asset financing — given, in particular, (i) a relatively homogeneous sample of distributors operating the same type of business; (ii) the availability of detailed administrative data on purchases; and (iii) a clear mechanism by which the productive asset could be used to expand operations for microfinance clients. These three key features are already shared by a large variety of 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 programs and retail distribution networks, particularly for consumer goods and food and beverage firms. Second, and more generally, these characteristics are shared by many 'gig work' and 'dependent contractor' arrangements — where host firms typically have extensive information about the quality and quantity of worker performance.

Indeed, as consumer markets expand in low- and middle-income countries, and as route-to-market programs grow, large companies are likely to place increasing reliance on 'dependent contractors' — many of whom are risk averse, economically precarious, and lack the fixed capital necessary to operate effectively. Our paper provides a proof of concept for a new class of financing contract, and our results show that such contracts may be particularly useful for such workers. Across a wide variety of contexts, rapid developments in financial technology — in particular, increasing adoption of mobile money and point-of-sale technologies — promise cheap access to credible information on the performance of microenterprises, gig workers, and sub-contractors, to improve screening and enforcement (Aggarwal, Brailovskaya, & Robinson, 2020; Annan et al., 2024; Berg et al., 2020; Higgins, 2024; Riley, 2018; Russel, Shi, & Clarke, 2023; Suri et al., 2023). The next generation of financial contracts can leverage these developments to expand the portfolio of products available to small firm owners – specifically, to include contracts with performance-contingent repayment obligations, offering better sharing of risk and reward. Our approach of offering flexible financing has some similarities to the long-standing German *Mittelstandsfinanzierung* model — a relationship-based approach to lending in which small and medium enterprises are allowed to roll over unpaid amounts to future periods in cases of financial difficulty. More recently, one promising innovation that facilitates asset financing is 'lock-

out technology', which has been shown to reduce moral hazard and improve credit risk management and repayment rates (Gertler, Green, & Wolfram, 2021).

Our partner MFI could not have achieved these returns without access to the multinational's administrative data. Theoretical work highlights the comparative advantage of suppliers as financial intermediaries over traditional lenders, given their superior ability to assess creditworthiness and enforce repayments by withholding future supplies.[26] Given the theoretical advantages and the remarkably high returns observed in our intervention, a natural question arises: *why wasn't the multinational already offering this financing?* What frictions might prevent multinationals from realizing such high potential returns within their supply chains?

Part of the answer may lie in a general lack of innovation by large firms in their supply chains. However, deeper frictions likely play a role. To explore this further, we conducted a qualitative discussion with a senior representative of the Kenyan financial services sector. That discussion highlighted significant legal and regulatory hurdles for a manufacturing firm to undertake lending activities. Compliance with central bank regulations requires extensive legal, financial, and operational documentation, along with thorough background checks on company directors. Additionally, the need for detailed Anti-Money Laundering (AML) and Know Your Customer (KYC) policies, complex IT systems, and risk management protocols may be beyond the capacity of a food manufacturer. Further, the central bank would likely require the establishment of a separate legal entity for lending, given the regulatory complexities and differences from the manufacturer's core operations. The practical challenges of lending to this demographic — such as collecting repayments and dealing with the complications of repossessing and auctioning assets — also pose significant risks, particularly to a brand's reputation. One alternative could be to partner with a financial institution and delegate borrower screening and debt collection, as was done in this collaboration. However, this approach would still involve navigating several regulatory hurdles, which may be too onerous for a food manufacturer.

Nonetheless, further research is needed to understand how some of these frictions can be overcome to unlock the high potential returns from such investments. Ride-sharing and delivery platforms, for example, could use contingent-repayment contracts to facilitate vehicle financing for drivers. Similarly, these contracts could apply to a broad range of sub-contractors – for example, farmers who 'finish' livestock animals for sale with equipment loans, or cut-and-trim manufacturers for their machinery (Casaburi & Willis, 2024). While host firms could offer these contracts, one could also envision third-party sharing agreements – similar to the model adopted here by FoodCo – where a specialized lender provides funds to a host firm for contingent lending to gig workers or sub-contractors, with the firm sharing performance data with the lender.[27] Such models open opportunities for contractual innovations that benefit both low-income microentrepreneurs and large firms.

[26] Further, in low-income countries, traditional lenders face distinct challenges in providing riskier, longer-term financing due to substantial liquidity risks from unstable funding sources and volatile deposits (Choudhary & Limodio, 2022).

[27] This is akin to 'factoring,' where a company sells accounts receivable to a financial company for collection.

References

Aggarwal, S., Brailovskaya, V., & Robinson, J. (2020). Cashing in (and out): Experimental evidence on the effects of mobile money in malawi. In *Aea papers and proceedings* (Vol. 110, pp. 599–604).

Alfaro-Urena, A., Manelici, I., & Vasquez, J. P. (2022). The effects of joining multinational supply chains: New evidence from firm-to-firm linkages. *The Quarterly Journal of Economics*, *137*(3), 1495–1552.

Alfonsi, L., Bandiera, O., Bassi, V., Burgess, R., Rasul, I., Sulaiman, M., & Vitali, A. (2020). Tackling youth unemployment: Evidence from a labor market experiment in Uganda. *Econometrica*, *88*(6), 2369–2414.

Angrist, J. D., Caldwell, S., & Hall, J. V. (2021). Uber versus Taxi: A driver's eye view. *American Economic Journal: Applied Economics*, *13*(3), 272–308.

Annan, F., Cheung, C., & Giné, X. (2024). Digital payments. *forthcoming, Oxford Review of Economic Policy*.

Azmat, S., & Macdonald, I. H. (2020). The psychological cost of debt: Evidence from housing mortgages in Pakistan. *Working Paper*.

Bandiera, O., Burgess, R., Das, N., Gulesci, S., Rasul, I., & Sulaiman, M. (2017). Labor Markets and Poverty in Village Economies. *The Quarterly Journal of Economics*, *132*(2), 811–870.

Banerjee, A., Breza, E., Duflo, E., & Kinnan, C. (2019). Can Microfinance Unlock a Poverty Trap for Some Entrepreneurs? *Working paper*.

Banerjee, A., Duflo, E., Goldberg, N., Karlan, D., Osei, R., Parienté, W., ... Udry, C. (2015). A Multifaceted Program Causes Lasting Progress for the Very Poor: Evidence from Six Countries. *Science*, *348*(6236), 1260799.

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. (2023). How do flexible microfinance contracts improve repayment rates and business outcomes? experimental evidence from india. *Working paper*.

Bari, F., Malik, K., Meki, M., & Quinn, S. (2024). Asset-based microfinance for microenterprises: Evidence from pakistan. *American Economic Review*, *114*(2), 534–574.

Barratt, T., Goods, C., & Veen, A. (2020). 'i'm my own boss...': Active intermediation and 'entrepreneurial'worker agency in the australian gig-economy. *Environment and Planning A: Economy and Space*, *52*(8), 1643–1661.

Battaglia, M., Gulesci, S., & Madestam, A. (2024). Repayment flexibility and risk taking: Experimental evidence from credit contracts. *Review of Economic Studies*, *91*(5), 2635–2675.

Berg, T., Burg, V., Gombović, A., & Puri, M. (2020). On the rise of fintechs: Credit scoring using digital footprints. *The Review of Financial Studies*, *33*(7), 2845–2897.

Biais, B., & Gollier, C. (1997). Trade credit and credit rationing. *The Review of Financial Studies*, *10*(4), 903–937.

Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., & Roberts, J. (2013). Does management

matter? evidence from india. *The Quarterly Journal of Economics*, *128*(1), 1–51.
Brailovskaya, V., Dupas, P., & Robinson, J. (2021). *Is digital credit filling a hole or digging a hole? Evidence from Malawi* (Tech. Rep.). National Bureau of Economic Research.
Breza, E., & Kinnan, C. (2021). Measuring the equilibrium impacts of credit: Evidence from the indian microfinance crisis. *The Quarterly Journal of Economics*, *136*(3), 1447–1497.
Breza, E., & Liberman, A. (2017). Financial contracting and organizational form: Evidence from the regulation of trade credit. *The Journal of Finance*, *72*(1), 291–324.
Brune, L., Giné, X., & Karlan, D. (2022). *Give me a pass: Flexible credit for entrepreneurs in colombia* (Tech. Rep.). National Bureau of Economic Research.
Bryan, G., Karlan, D., & Osman, A. (2024). Big loans to small businesses: Predicting winners and losers in an entrepreneurial lending experiment. *American Economic Review*, *114*(9), 2825–2860.
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.
Burkart, M., & Ellingsen, T. (2004). In-kind finance: A theory of trade credit. *American economic review*, *94*(3), 569–590.
Cai, J., & Szeidl, A. (2022). *Indirect effects of access to finance* (Tech. Rep.). National Bureau of Economic Research.
Carlson, S. (2017). Dynamic incentives in credit markets: An exploration of repayment decisions on digital credit in africa. *Department of Economics*.
Carney, K., Kremer, M., Lin, X., & Rao, G. (2022). *The Endowment Effect and Collateralized Loans* (Tech. Rep.). Retrieved from https://www.nber.org/papers/w30073
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*, *111*.
Carter, M., Galarza, F., & Boucher, S. (2007). Underwriting area-based yield insurance to crowd-in credit supply and demand. *Savings and Development*, 335–362.
Casaburi, L., & Willis, J. (2024). Value chain microfinance. *Oxford Review of Economic Policy*, *40*(1).
Choudhary, M. A., & Limodio, N. (2022). Liquidity risk and long-term finance: Evidence from a natural experiment. *The Review of Economic Studies*, *89*(3), 1278–1313.
Clarke, D. J. (2016). A Theory of Rational Demand for Index Insurance. *American Economic Journal: Microeconomics*, *8*(1), 283–306.
Cohen, A., & Einav, L. (2007). Estimating risk preferences from deductible choice. *American Economic Review*, *97*(3), 745–788.
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 Haas, R., Crepon, B., Pariente, W., & Devoto, F. (2022). Microcredit made to measure. experimental evidence from rural morocco. *Journal Of Development Economics*.

De Mel, S., McKenzie, D., & Woodruff, C. (2012). One-time Transfers of Cash or Capital have Long-lasting Effects on Microenterprises in Sri Lanka. *Science*, *335*(6071), 962–966.

De Mel, S., McKenzie, D. J., & Woodruff, C. (2019). Micro-equity for microenterprises. *CEPR Discussion Paper No. DP13698*.

Fiala, N., Garcia-Hernandez, A., Narula, K., & Prakash, N. (2022). Wheels of change: Transforming girls' lives with bicycles. *IZA Discussion Paper*.

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.

Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., & Walther, A. (2022). Predictably unequal? the effects of machine learning on credit markets. *The Journal of Finance*, *77*(1), 5–47.

Gertler, P., Green, B., & Wolfram, C. (2021). *Digital collateral* (Tech. Rep.).

Giné, X., & Yang, D. (2009). Insurance, credit, and technology adoption: Field experimental evidencefrom malawi. *Journal of development Economics*, *89*(1), 1–11.

Herbst, D., & Hendren, N. (2024). Opportunity unraveled: Private information and the missing markets for financing human capital. *American Economic Review*, *114*(7), 2024–2072.

Hickson, J. (2024). Freedom, domination and the gig economy. *New Political Economy*, *29*(2), 321–336.

Higgins, S. (2024). Financial technology adoption: Network externalities of cashless payments in mexico. *American Economic Review*, *114*(11), 3469–3512.

Holmström, B. (1979). Moral hazard and observability. *The Bell Journal of Economics*, 74–91.

Jack, W., Kremer, M., de Laat, J., & Suri, T. (2023). Credit Access, Selection, and Incentives in a Market for Asset-Collateralized Loans: Evidence From Kenya. *The Review of Economic Studies*, *90*(6), 3153-3185.

J-PAL. (2024). *Market access: Connecting firms and entrepreneurs to markets to spur business and job growth* (Tech. Rep.). Abdul Latif Jameel Poverty Action Lab.

Karlan, D., Kutsoati, E., McMillan, M., & Udry, C. (2011). Crop price indemnified loans for farmers: A pilot experiment in rural ghana. *Journal of Risk and Insurance*, *78*(1), 37–55.

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.

Karlan, D., & Zinman, J. (2009). Observing unobservables: Identifying information asymmetries with a consumer credit field experiment. *Econometrica*, *77*(6), 1993–2008.

Kerwin, J., Rostom, N., & Sterck, O. (2024). *Striking the right balance: Why standard balance tests over-reject the null, and how to fix it* (Tech. Rep.). Institute of Labor Economics (IZA).

Klapper, L., Laeven, L., & Rajan, R. (2012). Trade credit contracts. *The Review of Financial*

Studies, *25*(3), 838–867.

Kolk, A., Rivera-Santos, M., & Rufín, C. (2014). Reviewing a decade of research on the "base/bottom of the pyramid"(bop) concept. *Business & Society*, *53*(3), 338–377.

Kruijff, D., Sawhney, S., & Wright, R. L. (2024). *Empowering small giants: Inclusive embedded finance for micro-retailers.* Retrieved from https://www.cgap.org/research/empowering-small-giants-inclusive-embedded-finance-for-micro-retailers (Focus Note 2024. Washington, D.C.: CGAP)

Lazear, E. P. (2000). Performance pay and productivity. *American Economic Review*, *90*(5), 1346–1361.

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

Lopez, H. (2008). The social discount rate: Estimates for nine Latin American countries. *World Bank Policy Research Working Paper*(4639).

Macchiavello, R. (2022). Relational contracts and development. *Annual Review of Economics*, *14*(1), 337–362.

Macchiavello, R., & Morjaria, A. (2021). Competition and relational contracts in the rwanda coffee chain. *The Quarterly Journal of Economics*, *136*(2), 1089–1143.

Martínez-Marquina, A., & Shi, M. (2024). The opportunity cost of debt aversion. *American Economic Review*, *114*(4), 1140–1172.

McKenzie, D. (2015). *Tools of the trade: a joint test of orthogonality when testing for balance.* Retrieved from https://blogs.worldbank.org/impactevaluations/tools-trade-joint-test-orthogonality-when-testing-balance (Published on Development Impact, The World Bank)

McKenzie, D., & Woodruff, C. (2015). *Business practices in small firms in developing countries*. The World Bank.

McMillan, J., & Woodruff, C. (1999). Interfirm relationships and informal credit in vietnam. *The Quarterly Journal of Economics*, *114*(4), 1285–1320.

Meki, M. (2024). Small Firm Investment under Uncertainty: The Role of Equity Finance. *Oxford Department of International Development (ODID) Working Paper Series, University of Oxford*.

Mueller, H., & Yannelis, C. (2022). Increasing enrollment in income-driven student loan repayment plans: Evidence from the navient field experiment. *The Journal of Finance*, *77*(1), 367–402.

Paaso, M., Pursiainen, V., & Torstila, S. (2020). Entrepreneur debt aversion and financing decisions: Evidence from covid-19 support programs.

Petersen, M. A., & Rajan, R. G. (1997). Trade credit: theories and evidence. *The Review of Financial Studies*, *10*(3), 661–691.

Prahalad, C. K., & Hammond, A. (2002). Serving the world's poor, profitably. *Harvard Business Review*, *80*(9), 48–59.

Riley, E. (2018). Mobile money and risk sharing against village shocks. *Journal of Development Economics*, *135*, 43–58.

Rodrik, D., & Sandhu, R. (2024). *Servicing development: productive upgrading of labor-*

absorbing services in developing economies (Tech. Rep.). National Bureau of Economic Research.

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.

Russel, D., Shi, C., & Clarke, R. P. (2023). Revenue-based financing. *Available at SSRN*.

Singh, R., Bakshi, M., & Mishra, P. (2015). Corporate social responsibility: linking bottom of the pyramid to market development? *Journal of Business Ethics*, *131*(2), 361–373.

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., & 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.

Suri, T., Aker, J., Batista, C., Callen, M., Ghani, T., Jack, W., ... Sukhtankar, S. (2023). Mobile money. *VoxDevLit*, *2*(2), 3.

Suri, T., Bharadwaj, P., & Jack, W. (2021). Fintech and household resilience to shocks: Evidence from digital loans in kenya. *Journal of Development Economics*, *153*, 102697.

Townsend, R. M. (1979). Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory*, *21*(2), 265–293.

Udry, C. (1994). Risk and insurance in a rural credit market: An empirical investigation in northern Nigeria. *The Review of Economic Studies*, *61*(3), 495–526.

Udry, C., & Anagol, S. (2006). The return to capital in ghana. *American Economic Review*, *96*(2), 388–393.

Van Doornik, B., Gomes, A., Schoenherr, D., & Skrastins, J. (2024). Financial access and labor market outcomes: evidence from credit lotteries. *American Economic Review*, *114*(6), 1854–1881.

Tables and figures

Table 1: FIRST STAGE: TAKE-UP, BICYCLE OWNERSHIP, AND HOUSEHOLD FINANCES

	(1) Take-up	(2) Asset financing	(3) Household debt	(4) Household debt	(5) Owns a bicycle	(6) Owns a bicycle
Panel A: Pooled treatment						
Any contract	0.576***	8698.03***	-391.78	-651.31	0.52***	0.88***
	(0.043)	(105.915)	(655.929)	(1085.561)	(0.055)	(0.057)
Panel B: By contract						
Debt	0.677***	8953.04***	130.27	184.15	0.65***	0.93***
	(0.080)	(130.706)	(876.349)	(1232.579)	(0.085)	(0.054)
Hybrid	0.692***	8510.00***	-456.64	-617.48	0.66***	0.92***
	(0.091)	(238.456)	(732.582)	(981.379)	(0.092)	(0.049)
RevShare	0.488***	8415.00***	-850.27	-1616.30	0.46***	0.87***
	(0.078)	(229.676)	(673.032)	(1292.523)	(0.085)	(0.069)
IndexShare	0.469***	8910.00***	-309.69	-636.30	0.36***	0.77***
	(0.088)	(229.260)	(831.122)	(1722.240)	(1722.240)	(0.084)
Data source	Admin	Admin	Survey	Survey	Survey	Survey
Estimation	Take-up	LATE	ITT	LATE	ITT	LATE
Timeframe	Baseline	Baseline	1m-12m	1m-12m	1m-12m	1m-12m
Observations	161	161	496	496	496	496
Individuals	161	161	161	161	161	161
Control mean	0.00	0.00	2498.43	2498.43	0.07	0.07
Test: Debt = Hybrid	0.896	0.103	0.439	0.451	0.875	0.800
Test: Debt = RevShare	0.092	0.042	0.166	0.119	0.093	0.190
Test: Hybrid = RevShare	0.087	0.774	0.445	0.279	0.078	0.222

Note: In Panel A, "Any contract" pools all financing contracts, while Panel B presents regressions with separate dummies for each contract. We present both intent-to-treat (ITT) and local average treatment effect (LATE) estimations (instrumenting take-up with assignment), given the differential take-up seen in column 1. Column 2: amount of financing taken under one of our treatment contracts (equal to zero for the control group, by construction). Columns 3 and 4: total household debt levels, excluding our asset financing from column (2). Columns 5 and 6: whether the distributor owns a bicycle. The bottom three rows of the table display p-values for the three main cross-coefficient tests of interest: the difference in treatment effects between Hybrid, Debt, and RevShare. Standard errors, clustered at the individual level, are reported in brackets. * p < 0.10, ** p < 0.05, *** p < 0.01. All currency amounts are in Kenyan Shillings. The USD-KES exchange rate at baseline was approximately equal to 102.

Table 2: IMPACTS ON BUSINESS PROFITS

	(1) Profits: Foodco	(2) Profits: Foodco	(3) Profits: Foodco	(4) Profits: Foodco	(5) Profits: Foodco	(6) Profits: Foodco	(7) Other earnings
Panel A: Pooled treatment							
Any contract	791.63**	0.65**	1181.93**	1125.54**	1143.08*	865.53	612.85
	(385.626)	(0.308)	(562.082)	(537.083)	(595.220)	(631.056)	(1690.369)
Panel B: By contract							
Debt	530.38	0.23	692.24	1053.38**	458.68	104.65	1302.56
	(434.809)	(0.428)	(562.699)	(483.775)	(513.308)	(724.360)	(1928.553)
Hybrid	1528.51**	1.10***	1967.80**	2234.63***	2230.76***	1636.95*	-194.96
	(609.176)	(0.346)	(818.753)	(741.920)	(854.621)	(881.756)	(1607.719)
RevShare	781.65*	0.67*	1305.30*	981.15	1193.63*	1099.64	30.23
	(450.003)	(0.364)	(701.126)	(709.523)	(721.622)	(803.888)	(2020.717)
IndexShare	172.65	0.37	300.60	130.37	661.76	-11.78	1255.25
	(444.100)	(0.354)	(810.654)	(820.417)	(992.092)	(832.634)	(3006.037)
Data source	Admin	Admin	Admin	Admin	Admin	Admin	Survey
Estimation	ITT	ITT-Poisson	LATE	LATE	LATE	LATE	LATE
Observations	2888	2888	2888	785	817	910	496
Individuals	161	161	161	160	145	119	161
Timeframe	1m-36m	1m-36m	1m-36m	1m-6m	7m-12m	13m-24m	1m-12m
Control mean	897.45	897.45	897.45	1388.67	939.52	805.70	6528.46
Test: Debt = Hybrid	0.091	0.021	0.108	0.076	0.022	0.096	0.396
Test: Debt = RevShare	0.561	0.263	0.320	0.892	0.199	0.207	0.478
Test: Hybrid = RevShare	0.209	0.144	0.425	0.109	0.209	0.562	0.886

Note: In Panel A, "Any contract" pools all financing contracts, while Panel B presents a regression with separate dummies for each contract. ITT refers to Intent-to-Treat regressions, while LATE refers to Local Average Treatment Effect estimations (instrumenting take-up with assignment). Columns 1 to 6 use administrative data on profits from selling FoodCo products. Column 7 uses survey data that measures profits from all other sources (including profits from selling non-FoodCo products, as well as wage income). Standard errors, clustered at the individual level, are reported in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All currency amounts are in Kenyan Shillings. The USD-KES exchange rate at baseline was approximately equal to 102.

Table 3: IMPACTS ON BUSINESS PROFITS: CONTROLLING FOR BASELINE HETEROGENEITY

	(1) Profits: Foodco	(2) Profits: Foodco	(3) Profits: Foodco	(4) Profits: Foodco	(5) Profits: Foodco	(6) Profits: Foodco	(7) Other earnings
Panel A: Pooled treatment							
Any contract	938.32**	1.22***	1361.07***	1381.28**	1569.49***	1139.20**	647.98
	(370.186)	(0.399)	(525.879)	(556.328)	(538.977)	(564.319)	(1994.940)
Panel B: By contract							
Debt	605.09	0.53	789.65	1288.82**	777.31*	304.84	1339.21
	(374.726)	(0.505)	(487.980)	(519.817)	(464.468)	(537.859)	(2173.250)
Hybrid	1481.16***	1.67***	1869.66**	2231.57***	2300.34***	1500.52**	-165.76
	(564.037)	(0.458)	(737.126)	(712.227)	(757.774)	(699.507)	(1757.993)
RevShare	1098.50**	1.42***	1815.23***	1463.65*	1892.36***	1867.32**	535.96
	(457.052)	(0.466)	(696.630)	(748.026)	(705.841)	(831.617)	(2656.846)
IndexShare	382.15	1.00**	625.13	465.82	1281.18	488.54	827.02
	(456.454)	(0.456)	(749.186)	(791.678)	(908.041)	(805.908)	(2730.512)
Estimation	ITT	ITT-Poisson	LATE	LATE	LATE	LATE	LATE
Observations	2888	2888	2888	785	817	910	496
Individuals	161	161	161	160	145	119	161
Timeframe	1m-36m	1m-36m	1m-36m	1m-6m	7m-12m	13m-24m	1m-12m
Control mean	897.45	897.45	897.45	1388.67	939.52	805.70	6528.46
Test: Debt = Hybrid	0.094	0.002	0.118	0.131	0.033	0.072	0.375
Test: Debt = RevShare	0.236	0.020	0.084	0.738	0.070	0.041	0.702
Test: Hybrid = RevShare	0.499	0.415	0.944	0.307	0.607	0.663	0.694

Note: We repeat the analysis from Table 2, now incorporating controls for de-meaned baseline values of total profits, risk aversion, and loss aversion, as well as the interactions between these de-meaned variables and each treatment indicator. In Panel A, "Any contract" pools all financing contracts, while Panel B presents a regression with separate dummies for each contract. ITT refers to Intent-to-Treat regressions, while LATE refers to Local Average Treatment Effect estimations (instrumenting take-up with assignment). Columns 1 to 6 use administrative data on profits from selling FoodCo products. Column 7 uses survey data that measures profits from all other sources (including profits from selling non-FoodCo products, as well as wage income). Standard errors, clustered at the individual level, are reported in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All currency amounts are in Kenyan Shillings. The USD-KES exchange rate at baseline was approximately equal to 102.

Table 4: MECHANISMS: BUSINESS PRACTICES AND EFFORT

	(1) Bike use: business	(2) Bike use: hours	(3) Stockpoint visits	(4) Sales expansion	(5) Management practices	(6) Record keeping	(7) Credit extension
Panel A: Pooled treatment							
Any contract	0.78***	27.34***	2.28*	0.28**	0.10	0.10	0.03
	(0.030)	(2.046)	(1.345)	(0.119)	(0.077)	(0.098)	(0.028)
Panel B: By contract							
Debt	0.73***	21.79***	1.85	0.15	0.01	-0.01	0.02
	(0.053)	(2.009)	(1.466)	(0.118)	(0.085)	(0.104)	(0.032)
Hybrid	0.90***	33.33***	3.72*	0.26**	0.13*	0.20**	0.06*
	(0.043)	(5.163)	(1.955)	(0.121)	(0.071)	(0.092)	(0.033)
RevShare	0.71***	24.83***	2.21	0.27*	0.07	0.05	0.03
	(0.060)	(2.270)	(1.677)	(0.161)	(0.100)	(0.125)	(0.037)
IndexShare	0.79***	31.23***	0.58	0.48***	0.21**	0.23	-0.01
	(0.068)	(5.981)	(2.167)	(0.181)	(0.105)	(0.140)	(0.038)
Data source	Survey	Survey	Admin	Survey	Survey	Survey	Survey
Estimation	LATE	LATE	LATE	LATE	LATE	LATE	LATE
Observations	496	496	2888	496	496	496	496
Individuals	161	161	161	161	161	161	161
Timeframe	1m-12m	1m-12m	1m-36m	1m-12m	1m-12m	1m-12m	1m-12m
Control mean	0.00	0.00	2.57	0.57	0.69	0.66	0.09
Test: Debt = Hybrid	0.014	0.037	0.358	0.298	0.105	0.015	0.246
Test: Debt = RevShare	0.838	0.315	0.822	0.362	0.516	0.574	0.896
Test: Hybrid = RevShare	0.013	0.132	0.413	0.988	0.435	0.113	0.302

Note: We explore the impact of treatment on business effort and practices. Column 1: a binary variable indicating whether the bicycle financed through the intervention was used for business purposes. Column 2: number of hours that they use the project-financed bicycle in a typical week. (Columns 1 and 2 are coded as zero for individuals without a bicycle financed through the intervention.) Column 3: how often distributors visit stockpoints in a given month to purchase inventory. Column 4: proportion of the distributor's sales that comes from selling to customers that are greater than 1km from their stockpoint. Column 5: an index of business management practices, based on a set of questions developed by McKenzie and Woodruff (2015), and amended for a micro-distribution business. (Each index ranges from zero to one, indicating the proportion of questions that receive a positive response regarding whether a specific business management practice is undertaken.) Column 6: a specific sub-category of that index that relates to record-keeping. Column 7 is a proxy for distributors' risk-taking: the extent to which they offer credit to their own customers. Standard errors, clustered at the individual level, are reported in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: MECHANISMS, CONTROLLING FOR HETEROGENEITY

	(1) Bike use: business	(2) Bike use: hours	(3) Stockpoint visits	(4) Sales expansion	(5) Management practices	(6) Record keeping	(7) Credit extension
Any contract	0.82***	30.00***	2.79**	0.24**	0.19**	0.18	0.06**
	(0.028)	(2.096)	(1.313)	(0.109)	(0.089)	(0.110)	(0.026)
Debt	0.75***	23.83***	1.94	0.10	0.09	0.05	0.05
	(0.062)	(1.882)	(1.370)	(0.108)	(0.103)	(0.118)	(0.034)
Hybrid	0.94***	36.74***	4.17**	0.25**	0.21***	0.26**	0.09***
	(0.025)	(4.941)	(1.744)	(0.113)	(0.081)	(0.101)	(0.032)
RevShare	0.74***	25.23***	3.75**	0.26*	0.19	0.15	0.06
	(0.062)	(2.329)	(1.837)	(0.155)	(0.120)	(0.140)	(0.040)
IndexShare	0.83***	34.32***	0.65	0.34***	0.29***	0.26**	0.04
	(0.053)	(5.068)	(2.009)	(0.130)	(0.105)	(0.130)	(0.030)
Estimation	LATE	LATE	LATE	LATE	LATE	LATE	LATE
Observations	496	496	2888	496	496	496	496
Individuals	161	161	161	161	161	161	161
Timeframe	1m-12m	1m-12m	1m-36m	1m-12m	1m-12m	1m-12m	1m-12m
Control mean	0.00	0.00	2.57	0.57	0.69	0.66	0.09
Test: Debt = Hybrid	0.001	0.022	0.207	0.159	0.151	0.021	0.311
Test: Debt = RevShare	0.941	0.200	0.285	0.196	0.319	0.323	0.779
Test: Hybrid = RevShare	0.003	0.114	0.813	0.906	0.812	0.255	0.498

Note: We repeat the analysis from Table 4, now incorporating controls for de-meaned baseline values of total profits, risk aversion, and loss aversion, as well as the interactions between these de-meaned variables and each treatment indicator. Column 1: a binary variable indicating whether the bicycle financed through the intervention was used for business purposes. Column 2: number of hours that they use the project-financed bicycle in a typical week. (Columns 1 and 2 are coded as zero for individuals without a bicycle financed in our project.) Column 3: how often distributors visit stockpoints in a given month to purchase inventory. Column 4: proportion of the distributor's sales that comes from selling to customers that are greater than 1km from their stockpoint. Column 5: an index of business management practices, based on a set of questions developed by McKenzie and Woodruff (2015), and amended for a micro-distribution business. (Each index ranges from zero to one, indicating the proportion of questions that receive a positive response regarding whether a specific business management practice is undertaken.) Column 6: a specific sub-category of that index that relates to record-keeping. Column 7 is a proxy for distributors' risk-taking: the extent to which they offer credit to their own customers. Standard errors, clustered at the individual level, are reported in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: HOUSEHOLD CONSUMPTION EXPENDITURE

	(1) Expenditure: Food	(2) Expenditure: Clothing	(3) Expenditure: Schooling
Debt	1230.10*	81.68	-425.07
	(690.67)	(265.59)	(516.16)
Hybrid	706.74	665.89**	535.13
	(725.82)	(306.28)	(564.82)
RevShare	117.43	103.54	283.39
	(826.19)	(386.14)	(725.32)
Index	1545.52	-368.21	71.13
	(985.33)	(392.39)	(674.24)
Data source	Survey	Survey	Survey
Estimation	LATE	LATE	LATE
Observations	496	496	496
Individuals	161	161	161
Timeframe	1m-12m	1m-12m	1m-12m
Control mean	4626.37	908.79	1113.74
Test: Hybrid = Debt	0.471	0.043	0.065
Test: Hybrid = RevShare	0.442	0.121	0.697
Test: RevShare = Debt	0.112	0.945	0.234

Note: We explore treatment effects on downstream household outcomes, focusing on the largest categories of household consumption expenditure. Standard errors, clustered at the individual level, are reported in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All amounts are in Kenyan Shillings. The USD-KES exchange rate at baseline was approximately equal to 102.

Table 7: TOTAL RETURN ANALYSIS

	(1) Distributors	(2) Distributors	(3) FoodCo	(4) FoodCo	(5) Stockpoints	(6) Stockpoints	(7) Total Return	(8) Total Return
Assignment (ITT)	471**		1534**		449**		2454**	
	(221)		(718)		(210)		(1148)	
Take-up (LATE)		708**		2305**		674**		3687**
		(325)		(1057)		(309)		(1691)
Observations	2888	2888	2888	2888	2888	2888	2888	2888
Individuals	161	161	161	161	161	161	161	161
Timeframe	1m-36m	1m-36m	1m-36m	1m-36m	1m-36m	1m-36m	1m-36m	1m-36m
Control mean	521	521	1693	1693	495	495	2709	2709

Note: We display the returns from the intervention to each of the three participants in FoodCo's supply chain. For distributors, we adjusted our administrative measure of 'gross profits' (sales minus cost of goods sold) that we use in our previous analysis to better approximate net profits by incorporating survey-based estimates of additional operating expenses, such as transportation and raw materials. This adjustment allowed us to convert gross profits into 'operating income,' aligning with standard accounting measures. We then applied the same process to FoodCo and stockpoints for comparability. For FoodCo, we estimated operating income by applying gross profit and operating cost ratios from their publicly available financial reports. These ratios were applied to the value of sales generated by FoodCo based on stockpoint purchases, providing an estimate of their operating income. For stockpoints, we estimated operating income assuming the same gross profit and cost ratios as FoodCo, excluding costs not applicable to stockpoints (e.g., advertising). This provided a comparable measure of operating income for Stockpoints. ITT refers to Intent-to-Treat regressions, while LATE refers to Local Average Treatment Effect estimations (instrumenting take-up with assignment). Standard errors, clustered at the individual level, are reported in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All amounts are in Kenyan Shillings. The USD-KES exchange rate at baseline was approximately equal to 102.

Figure 1: **Model predictions**

PANEL A: EFFORT AS A FUNCTION OF RISK AVERSION

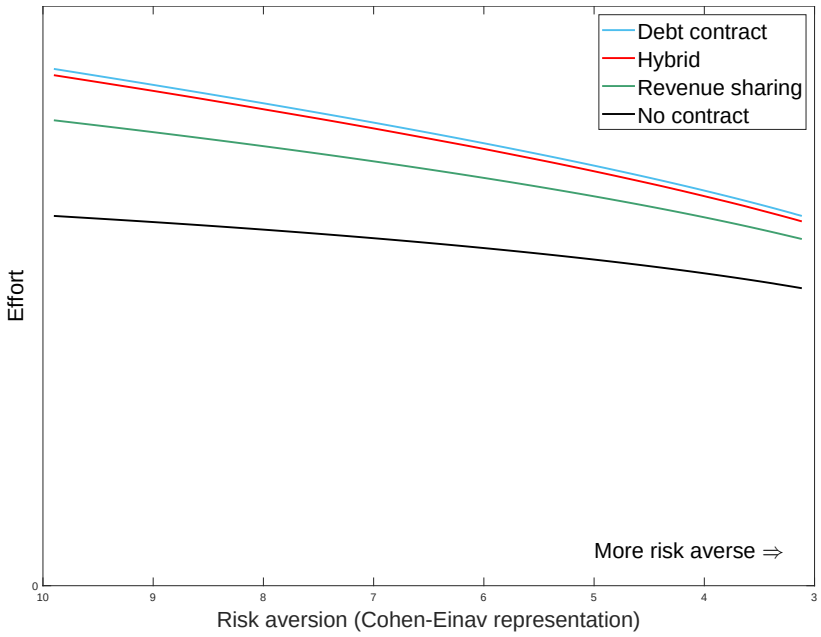

PANEL B: TAKE-UP AS A FUNCTION OF RISK AVERSION

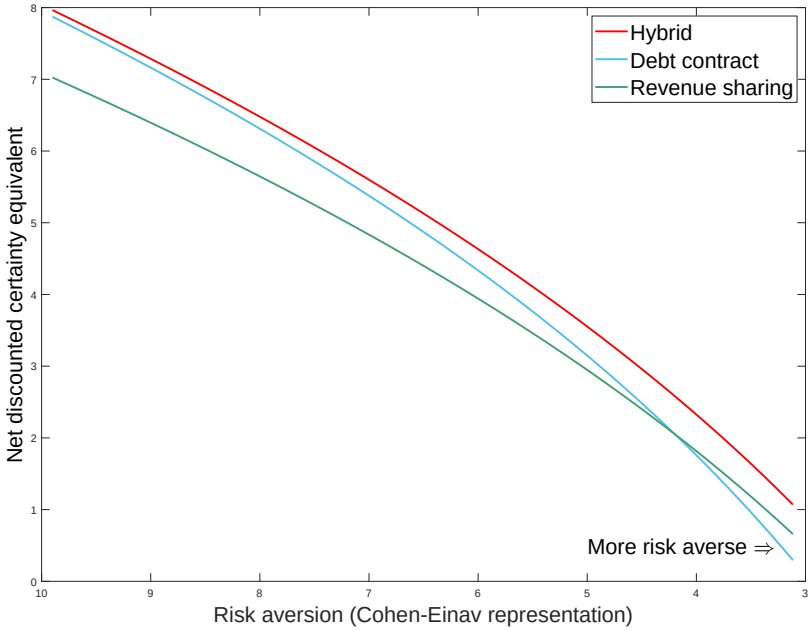

Note: *We illustrate numerically the theoretical predictions as to effort and take-up under no contract, the Debt contract, the RevShare contract and the Hybrid contract. For ease of interpretation, we use the representation of Cohen and Einav (2007); we imagine a 50-50 gamble where the gain is \$10 and the loss is x. For each given coefficient of absolute risk aversion, we solve for x so that the respondent is indifferent between taking the gamble and not; this is given by* $x \equiv \log\left[2 - \exp(-10r)\right]/r$. *The net discounted certainty equivalent is calculated as* $\left[-\log(1-\beta) - \log(-V)\right]/r$.

Figure 2: **Contract payments over time**

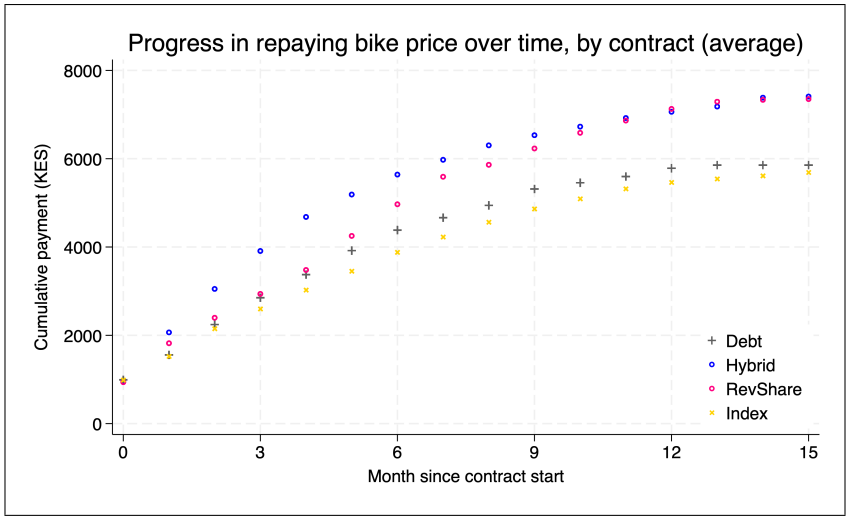

Note: *Each line represents the average repayment amount over time under each of the four contracts. All amounts are in Kenyan Shillings.*

Figure 3: **Total percentage paid**

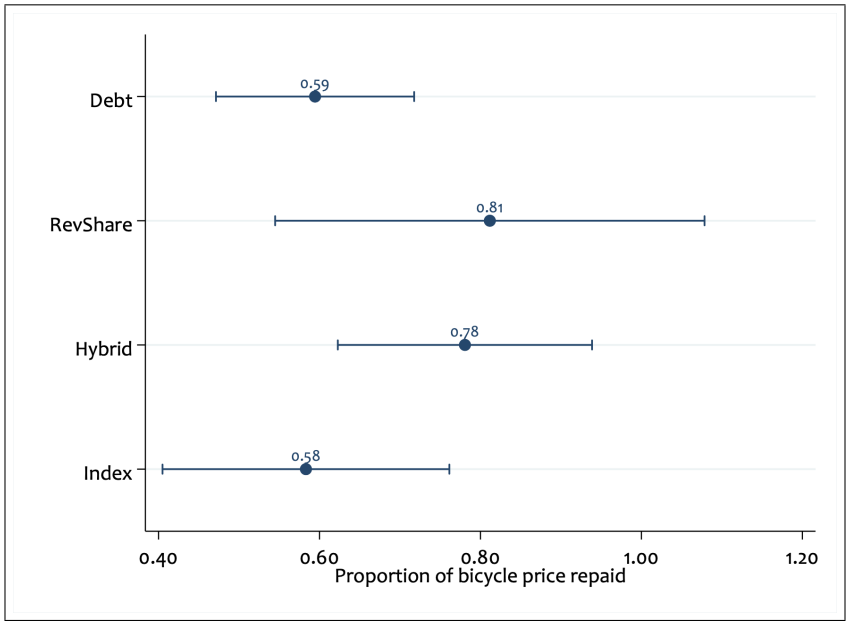

Note: *This figure illustrates the average repayment amount under each contract, as a percentage of the capital amount disbursed.*

Figure 4: **Bicycle GPS data**

PANEL A: OVERALL KENYA HEATMAP

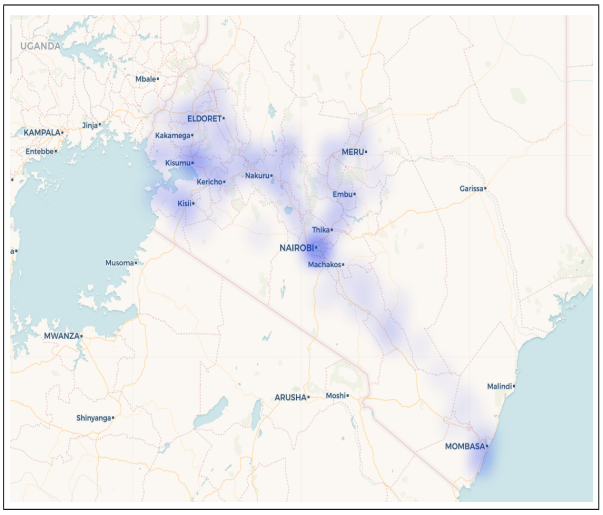

PANEL B: DISTRIBUTORS & STOCKPOINTS

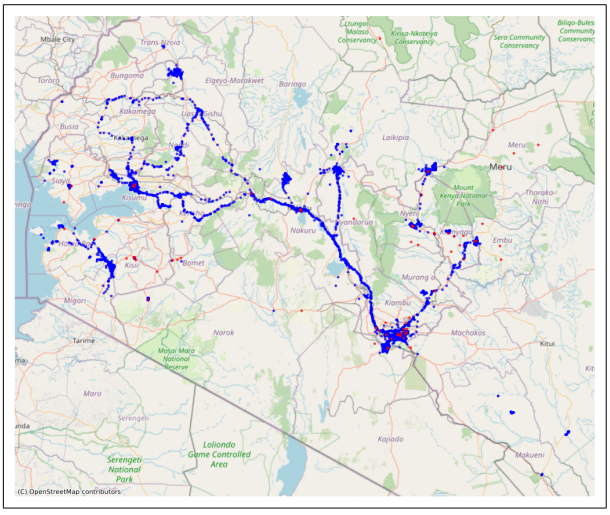

PANEL C: NAIROBI

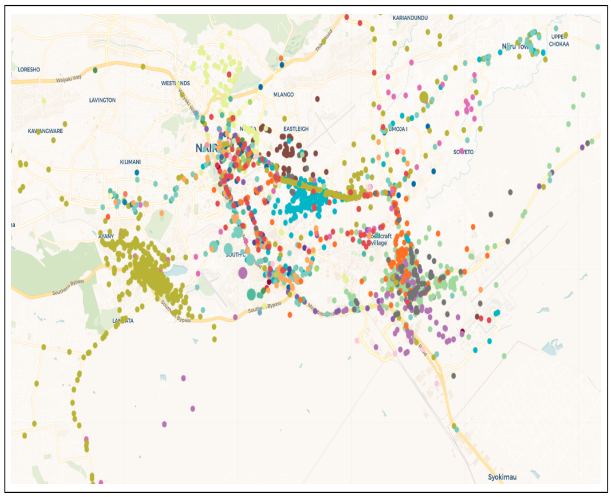

PANEL D: WESTERN KENYA

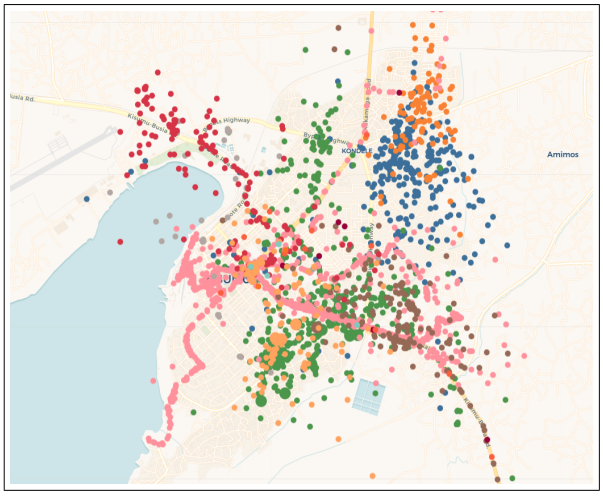

Note: *This figure utilises data from GPS trackers that were attached to each bike, between 2018 and 2020. The heat map in Panel A represents the density of visits in each location. The GPS data has been processed to build 20,683 areas with a resolution of approximately 5x5 meters and counting the frequency of trackers present in that area throughout the period of analysis. The colour intensity is proportional to the frequency of visits ranging from 1 visit (lighter blue) up to 1,954 visits (darker blue, corresponding to 69 visit per month on average). The picture shows the existence of clusters centred around the most populated areas (Nairobi, Western Kenya, Mombasa) and displacements between them. In Panel B, the blue dots represent all unique latitude and longitude pairs from data collected on distributor location via GPS trackers, after removing outliers (defined as the top 10% most extreme distances from a central reference point for each individual). This filtering allows for a more focused analysis of regular travel patterns for distributors and localized market activity. The red dots indicate the locations of all stockpoints, (including those not involved in the experiment). In Panels C and D, each colour represents data points for a distinct individual, highlighting the trip across the regions around Nairobi and Western Kenya (the two most populous regions in Kenya). On average each individual travelled 4.8 km per day (with a standard error of 0.4 km per day, and a median of 4.0 km per day).*

Figure 5: **Benefit-cost ratios and internal rates of return**

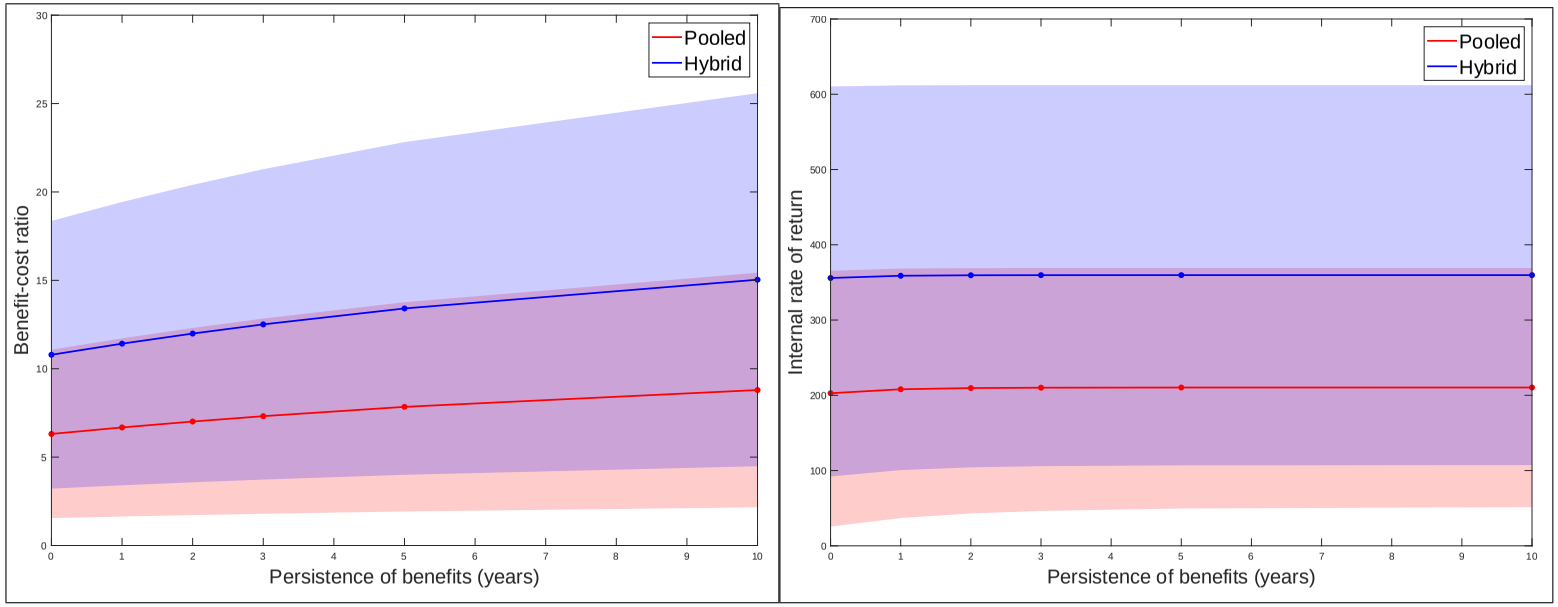

Note: *We present benefit-cost ratios and internal rates of return, both pooled across our four contracts and for the highest-performing contract, Hybrid. For benefits during the project period, we sum the treatment effects on business profits for the three key actors in the FoodCo's supply chain: the distributors, FoodCo, and stockpoints, as outlined in the total return analysis in Table 7. The estimated treatment effects from our LATE regressions are used, which we extrapolate when assuming future benefits. Our analysis is conducted under varying assumptions regarding the persistence of benefits beyond the three-year project period. For costs, we include: (i) the capital disbursed for the initial asset purchases for take-up clients, minus the total recovered capital (accounting for non-payment of contractual obligations); (ii) staff salaries; and (iii) other implementation expenses, such as venue rentals for workshops. The total costs are compounded to the two-year mark using a conservative 10% social discount rate (Lopez, 2008).*