

Optimal Arrangements for Distribution in Developing Markets: Theory and Evidence*

William Fuchs[†]

Brett Green[‡]

David I. Levine[§]

October 19, 2017

Abstract

A large literature examines demand-side barriers to product adoption. In this paper, we examine frictions on the supply side. We model the relationship between a producer or distributor and its vendors, where credit constraints and limited contract enforcement present challenges for distribution. We show that providing vendors with an initial endowment and the option to buy additional units at a fixed price is an optimal way in which to overcome these frictions. The arrangement is straightforward to implement and is optimal both for profit-maximizing firms and for non-profit organizations with limited resources. We test the arrangement using a field experiment in rural Uganda. We find that the optimal arrangement increases sales by 3-4 times compared to a standard contract. However, the rate of sales growth was lower than predicted by the model due to (i) an unwillingness to extend credit to customers and (ii) the lack of a reliable savings technology.

JEL: I15, O1, O33, D02, D86

Keywords: Product adoption puzzle, credit constraints, limited enforcement, relational contracts

*We would like to thank Peter DeMarzo, Andy Skrzypacz, David Sraer, Steve Tadelis, and Rob Townsend for useful suggestions. We have also benefited from seminar participants at Kellogg School of Management, the Norwegian School of Economics, Stanford Graduate School of Business, Stockholm School of Economics, Universidad Carlos III de Madrid, Universidad Autónoma de Barcelona, and the Munich Relational Contracts Conference. We gratefully acknowledge financial support from the Clausen Center, the Gartwood Center for Corporate Innovation, the Hellman Family Faculty Fund, the Committee on Research at Berkeley, the European Research Council (Grant 681575), and the generous donors of our Indiegogo funding campaign. Vastinah Kemigisha, Moses Oundo, Aisha Nansamba, and Anne Kayiwa were instrumental in carrying out the field experiment. Norma Altshuler, Ross Green, Vincent Skiera, and Mark Whitmeyer provided excellent research assistance.

[†]UT Austin and UC3M, wfuchs@gmail.com

[‡]UC Berkeley, Haas School of Business, greenb@berkeley.edu

[§]UC Berkeley, Haas School of Business, levine@berkeley.edu

1 Introduction

Investment in a number of basic technologies (e.g., solar lights, efficient cook stoves, fertilizer, anti-malaria nets, water filters) appears to have large welfare benefits for many households in developing economies. To illustrate, consider the case of using solar lights and efficient cook stoves to replace kerosene lamps and cooking over an open flame. As of 2011 the International Energy Agency (IEA) reports that over 1.3 billion people lack access to electricity and 2.7 billion cook over an open flame. The negative health impacts are sobering. The United Nations Development Program and the World Health Organization (WHO) report that 1.6 million deaths per year in developing countries are caused by the indoor air pollution attributed to traditional fuels.¹ Efficiently designed cook-stoves can eliminate up to 94% of the smoke and 91% of the carbon monoxide emissions and have been demonstrated to lead to improved health outcomes. Solar lights can altogether eliminate the need for kerosene lamps, which are detrimental to indoor air quality and pose a serious fire risk.

In addition to the health benefits, these technologies appear to have substantial economic benefits. Consumers spend \$17 billion on kerosene each year to light their homes. The light cast from a kerosene lamp is poorly distributed, has a low intensity, and is expensive. The poor lighting quality from kerosene lamps makes it difficult for children to study, inhibiting literacy and education. In rural Uganda women are estimated to spend 2 hours a day gathering fuel for cooking, reducing the amount of time that can be devoted to income generating activities. Those in urban areas who purchase their fuel for cooking spend up to 30% of their income on it.² The WHO estimates that efficient cook stoves reduce fuel consumption by approximately 50%. The fuel savings alone would pay for the cost of the stoves in less than 3 months. The same is true for solar lights. Purchasing a \$25 solar light saves households \$1.50 per week in kerosene costs. By conservative estimates, the light has a useful lifespan of 2 years, which means that investing in a solar light has an internal rate of return of over 300%.

Despite these seemingly large economic and health benefits, adoption rates are low and markets for these technologies have been slow to develop. If these products are so valuable to households, why is their adoption rate so low? Why have private markets not developed for them? On the demand side, there are well-documented barriers impeding adoption. For instance, poor households face credit constraints (Cole et al., 2013; Tarozzi et al., 2014), lack information about the product benefits or durability (Feder and Slade, 1984; Conley and

¹See Fuel for life: household energy and health World Health Organization, 2006

²See the Worldwatch Institute report: Energy for Development: The Potential Role of Renewable Energy in Meeting the Millennium Development Goals 2004.

Udry, 2001; Giné and Yang, 2009), suffer from present bias (Duflo et al., 2011), or may be too risk adverse to experiment with a new technology (Foster and Rosenzweig, 1995; Conley and Udry, 2010; Bryan et al., 2014). Product quality can also be hard to assess; consumers may be concerned about receiving products of inferior quality (Bold et al., 2017).

To address these challenges, a variety of retail offers have been suggested. For example, Levine et al. (2013) proposed using a retail offer involving a free-trial period and installment payments. In two different randomized control trials (RCTs) they found that using a free-trial period and installment payments increased the adoption rate by more than 40 percentage points compared to the standard fixed-price retail contract (47% versus 5%). Moreover, repayments rates were extremely high (99%) and 31% of customers completed their installments early.

While this success is encouraging, to scale this approach to the firm or market level, supply-side frictions must also be addressed. In particular, to reach the final consumer, firms or organizations generally rely on a decentralized sales force (e.g., local vendors and shopkeepers), especially in rural areas.³ Local vendors face similar issues to the ones faced by consumers. Small shopkeepers lack the capital necessary to purchase, transport, and store inventory.⁴ Even with sufficient capital, most vendors have little experience with many of these products and may be uncertain, or even pessimistic, about the profitability of retailing new technologies. To further complicate matters, weak enforcement of contracts and limited liability renders many commonly used incentive schemes for vendors infeasible.

In this paper, we address several of these supply-side frictions. We model the problem of a firm that can employ a local vendor to distribute its goods at a lower per unit cost than the firm.⁵ The vendor is liquidity constrained and therefore cannot purchase inventory outright. Motivated by the limited enforceability of contracts in developing countries, we assume that the firm cannot prevent the local vendor from absconding with inventory. Therefore, the firm must use (credible) promises of future utility to provide incentives to the vendor. The question is how to do so in the most efficient manner.

We characterize the optimal self-enforcing arrangement and show it has an appealing and simple implementation. It entails an initial endowment of the good and a fixed price at which the vendor can buy additional units in the future. The initial endowment helps overcome the liquidity constraint and the opportunity to grow induces the vendor to continue

³In fact, we developed these ideas in discussions with one of the world's largest development NGOs, BRAC, in an effort to improve the distribution of development goods using their existing network of micro-entrepreneurs.

⁴See e.g., de Mel et al. (2008).

⁵The cost differential is meant to capture the idea that local vendors have superior information about the local market conditions and/or a greater ability to recover payments from the final consumer.

reinvesting. We show that this arrangement is optimal both for profit-maximizing firms as well as for non-profit organizations with a limited budget whose objective is to maximize product distribution.

An interesting feature of the optimal arrangement is that it involves “starting small.” That is, the initial endowment is below the vendor’s capacity and only over time does the vendor’s business grow to the efficient scale. We provide a closed form solution for the optimal size of the initial endowment and show that starting small is particularly important when profit margins are low and when the vendor is relatively impatient.

Using the model, we develop several hypotheses regarding the impact of different contractual arrangements. We then run a field experiment in order to test them and evaluate the performance of the optimal arrangement implied by the model. We recruited vendors in rural Uganda to sell solar lights and randomized over several features of the arrangement that we offered to each vendor.

Our results from the field suggest that vendor liquidity constraints are indeed an important barrier. Vendors that are provided an initial endowment had average sales 3-4 times higher than vendors who were not. In addition, consumer uncertainty about product quality appears to be another important factor. Vendors who were given a “loaner light” in order to provide potential customers with a free-trial period also had significantly higher sales. On the other hand, vendor uncertainty did not appear to be a limiting factor: providing the vendors with the right to return any unsold inventory had little effect on total sales.

Perhaps not surprisingly, vendor sales growth was lower and exhibited a different time-series pattern than predicted by the model. These results suggest that factors outside of the model limit vendors’ ability to grow their business. Exit surveys point to a general inability to save revenues from sales until the next delivery of lights (a period of only a couple weeks) as well as an unwillingness to offer credit (e.g., an installment plan) to customers.

Recent developments in “fintech” (i.e., financial technology) have facilitated ways to overcome these difficulties as well as the frictions that we address with our incentive scheme. For example, M-Kopa Solar sells a home solar system with a “kill switch” and using a contract that features a daily payments made via mobile money.⁶ Using mobile money ensures that vendors are not required to handle (and save) cash, while the kill switch makes the product worthless if the vendor absconds with it or if the customer defaults. M-Kopa has significantly higher production costs because each system must contain many of the capabilities of a mobile phone. Nevertheless, they have sold more than 500,000 units in East Africa. Their success suggests that the frictions we study in the model and identify in our

⁶See <https://www.bloomberg.com/features/2015-mkopa-solar-in-africa>. Date accessed: September 26, 2017.

experiment are of first-order importance for understanding the product adoption puzzle.

Part of our goal in this paper is to narrow the gap between the literature on dynamic/relational contracting and the development economics literature. Research in both areas has flourished in the last several decades. Yet, there has been relatively little work that applies the tools developed in dynamic contracting to problems in development economics.

A notable exception is the recent work of Townsend and co-authors.⁷ For example, Karaivanov and Townsend (2013) use data from the Townsend Thai surveys to evaluate which models best describe the patterns of investment and consumption.⁸ Consistent with our hypothesis of liquidity constrained vendors, they document that investment in rural areas is sensitive to cash flows and that a savings-only regime (i.e., no borrowing) provides the best fit with the data. Our approach is complementary in that we do not presume that the market outcome is constrained efficient. Instead, we ask if there are possible arrangements that might enhance welfare given the contractual constraints and experiment with these arrangements in a controlled study.

Our theoretical results are related to the literature on optimal contracting with investment.⁹ Starting with DeMarzo and Fishman (2007a) a number of papers in this literature have considered implementations where the agent has both a standard debt contract and a credit line. When cash flows realizations are low the agent draws on the credit line to make the coupon payments on the debt. When the cash flow realizations are high the agent pays back the credit line. After the credit line is repaid, the agent consumes the cash flows in excess of coupon payments. If the agent has reached the limit of the credit line and still cannot meet the coupon payments then the firm is liquidated. Although the dynamics are similar, their focus is on implementation through financial securities and the capital structure of the firm, whereas we focus on making the arrangement simple enough to be used in the field. Implementing the arrangement with an initial endowment and fixed price requires less book keeping for the principal. It also provides a transparent way to recognize that the lack of enforceability means that any liquidity provided to the agent must also be promised in future continuation value. Rampini and Viswanathan (2010, 2013) develop dynamic models of investment with both limited commitment and enforcement in which a fraction of the capital can be pledged as collateral. In our setting, the investment corresponds to working

⁷Another exception is Dubois et al. (2008), who study both formal and informal mechanisms for risk sharing. There is also an older literature that has focused more on the consumption smoothing problem at the household level. See for example Townsend (1994) and Ligon et al. (2002).

⁸See for example Townsend and Urzua (2009), Kaboski and Townsend (2011), Karaivanov and Townsend (2013), and Chiappori et al. (2014) among others.

⁹See, for example, Quadrini (2004), Albuquerque and Hopenhayn (2004), Clementi and Hopenhayn (2006), DeMarzo and Fishman (2007b), Demarzo et al. (2012).

capital and is not pledgeable.

Our model also builds on work in the relational contracting literature, most notably Thomas and Worrall (1994), Baker et al. (2002), and Levin (2003). While the motivation for our model is different from Thomas and Worrall (1994), it is mathematically quite similar. For instance, the solution in our model also involves increasing the continuation value of the agent at a rate proportional to her discount factor until the steady state is reached. Troya-Martinez (2015) studies trade credit in a relational contracting setting, which is implemented with trade credit being suspended (possibly permanently) when the agent fails to make a full repayment.¹⁰ In our implementation, the size of the credit line is fixed over time and the adjustment takes place via quantities.

2 Illustrative Example

In this section we present a simple example which illustrates the key ideas of the paper. We take as given, that an organization (the “NGO”) has identified a good with social benefits (e.g., bed nets for malaria prevention) and faces the problem of distributing the good throughout the economy without being able to write enforceable contracts with its distributors.

For simplicity, consider a single village in which there are a large number of households. The NGO has raised funds of B for the purpose of distributing bed nets throughout the village. The NGO can purchase bed nets from a producer at marginal cost c . Each household within the village is willing to pay up to p for a bednet.¹¹

We assume here that $p < c$, and so, in the absence of some form of subsidy, the market for bed nets would remain undeveloped in the village (see Dupas and Cohen, 2010). In order to reach households, the NGO must incur a distribution cost of d , for the transportation and time involved in delivering each unit. The objective of the NGO is to maximize the discounted sum of all bed nets distributed, $\sum_{t=0}^{\infty} \delta^t k_t$, subject to the constraint that the NGO has a limited amount of funds with which to purchase and distribute bed nets.¹² The question we seek to answer in this section is how the NGO should go about distributing the bed nets.

Although the NGO is not concerned with profits, it is constrained by its funding and thereby will find it advantageous to charge households their willingness to pay in order to

¹⁰See also Cunat and Cuñat (2007).

¹¹Note that p may represent household’s true value for the good, or the amount that they are able to pay, which might be less than their true value due to credit constraints.

¹²This is the same objective that would obtain if, for example, the NGO ascribes some sufficiently large social value, Δ_s , to each bednet distributed and has the objective of maximizing total social surplus.

distribute more bed nets throughout the community.¹³ By doing so, the NGO reduces their effective marginal cost to $c + d - p$. Thus, if the NGO decides to *procure and distribute* bed nets, it can afford to purchase and distribute a total quantity of bed nets equal to

$$K_{pd} = \frac{B}{c + d - p}.$$

Consider now the possibility of forming partnership with a local shopkeeper or vendor in order to assist with the distribution of bed nets. The natural advantage of the partnership is that local vendors can distribute bed nets at a cost of only $d_A < d$, based on their knowledge and retail experience within the local community. Thereby, the hope is that by forming this partnership, the NGO will be able to lower its costs and reach more households. As discussed earlier, one difficulty with this approach is that local vendors are financially constrained, and thus, do not have the capital to purchase inventory up front. This is further complicated by the inability to write enforceable contracts and the limited liability of local agents. Thus, to sustain a partnership, the NGO will find it necessary to provide local vendors with appropriate incentives. The question then is whether doing so can achieve a better outcome than the procure and distribute strategy described above.¹⁴

Maximal First-Period Distribution

One approach is for the NGO to purchase as many bed nets as feasible (B/c) and give them to the vendor to distribute. Provided the households' willingness to pay, p , exceeds the agents distribution cost, d_A , the agent will find it in her interest to distribute the bed nets to households, from which she derives a net profit of $(B/c)(p - d_A)$. If $d > p$, then the NGO will have achieved a larger distribution of bed nets using this approach relative to the procure and distribute strategy (otherwise, procure and distribute is preferable). However, because the NGO has used all of its resources, the vendor will have no incentive to continue the partnership. That is, the NGO will be unable to incentivize the agent to use some of her profits to reinvest and distribute more bed nets (recall that $p < c$, so without a subsidy, the vendor will not have an incentive to purchase additional units at their marginal cost and distribute them). Hence, the distribution process stops after the first period.¹⁵

¹³Charging households for the good may have the additional benefit of selecting those households with the higher willingness to pay, who are thus more inclined to use the product (Ashraf et al., 2010).

¹⁴Naturally, if the potential savings on distribution costs are sufficiently small, the NGO will not find overcoming the agency costs to be worthwhile.

¹⁵Since it has exhausted its budget, any promises made by the NGO suggesting otherwise are not credible.

Starting Small

Under certain conditions, the NGO can do better by starting small and enabling the vendor to grow her business over time. To provide incentives for the vendor to continue the partnership beyond the initial period, the NGO cannot exhaust all of its resources in the first period. Instead, the NGO procures an initial quantity of $k_0 < B/c$, and provides this to the vendor as “seed” capital. The promise of repeated business then provides the vendor incentives to reinvest. Such incentives can take different forms; one simple way is by offering to sell additional bed nets to the vendor at a subsidized price, p_A .

As before, provided $p > d_A$, the vendor will find it in her interest to sell the initial allocation to households leaving her with a net profit of $k_0(p - d_A)$. At this point, the vendor must decide whether to return to the NGO to purchase more bed nets or allocate this profit toward other uses. Assuming that bed nets take one month to sell and the vendor’s (monthly) discount factor is δ_A , the vendor will return to the NGO to buy more bed nets provided that

$$\delta_A \left(\frac{k_0(p - d_A)}{p_A} \times (p - d_A) \right) \geq k_0(p - d_A) \iff p_A \leq \delta_A(p - d_A). \quad (1)$$

That is, providing appropriate incentives to the vendor, amounts to charging a price low enough that she finds it in her interest to buy and distribute more bed nets from the NGO. Rewriting the equality in (1) as

$$1 \leq \frac{\delta_A(p - d_A)}{p_A},$$

yields a simple interpretation; for each dollar of revenue earned, the vendor must decide whether to consume it (the left-hand side), or reinvest it in the partnership (the right-hand side). By reinvesting in the partnership, the agent can purchase a quantity of $1/p_A$ bed nets, which can be sold over the next month at profit margin $p - d_A$, generating a total revenue of $\frac{p - d_A}{p_A}$ in the next period. It is straightforward to see that this condition ensures that the vendor will prefer to continue her relationship with the NGO until either the market is saturated or the NGO runs out of money and becomes unable to continue providing bed nets at the subsidized price. Let us now fix $p_A = \delta_A(p - d_A)$ —it can be shown that this is the optimal price for the NGO to charge the vendor—so that the quantity of bed nets will grow at a rate proportional to the vendor’s discount factor (i.e., $k_{t+1} = \delta_A^{-1}k_t$). Eventually, the NGO will exhaust its resources and the vendor will have distributed a total of $K^* = \frac{B - ck_0}{c + d_A - p}$ bed nets.

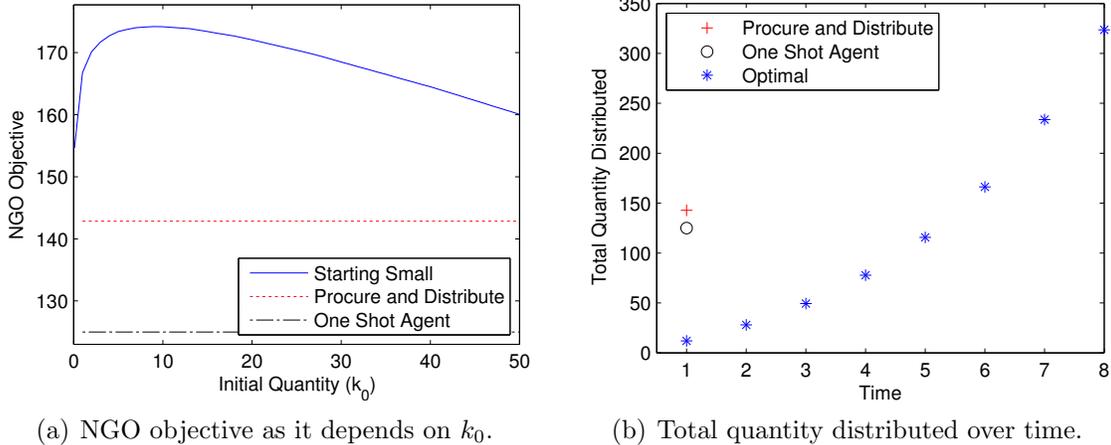


Figure 1: Illustrates the advantage of starting small in the relationship and building up over time. The figures use the parameters $p = 4, c = 8, d = 3, d_A = 0.5, B = 10^3, \delta = 0.99, \delta_A = 0.75$.

When considering the number of units to allocate initially to the agent, k_0 , the NGO faces a trade-off. A larger initial allocation increases the immediate distribution of the good, but it tightens the budget constraint, which reduces the subsidy the NGO can provide for units in the future. Given the vendor can always decide not to return and reinvest her profits, she must receive a continuation value in the relationship proportional to the number of units she is originally allocated. Figure 1(a) shows the value of the NGO’s objective as a function of k_0 . For this parametrization, it is optimal for the NGO to use only a small fraction of its total resources in the first period. Although this implies first-period distribution will be lower than with the other strategies (see panel (b)), as the vendor sells the bed nets she will return to the NGO to purchase additional units, which the NGO continues to subsidize. Intuitively, by starting small the NGO is able to recover a larger portion of its costs and therefore increase the total quantity of distribution (albeit at the cost of a delay in when the bed nets reach households).

Perhaps the most desirable feature of this arrangement is the simplicity with which it can be implemented; that is, the NGO provides an initial quantity to the vendor and charges a fixed (subsidized) price for all subsequent units. The initial quantity helps overcome the agent’s liquidity constraints and the subsidized price provides the necessary incentives for reinvestment. In the next section, we present a formal model and demonstrate the optimality of this arrangement.

3 The Model

We now relabel the NGO as simply the *principal*, which may also refer to a manufacturer or distributor seeking entry to a new market. As in the example, we endow the principal the technology to produce (or purchase) units of the good at a marginal cost, c . Similarly, we relabel the local vendor as simply the *agent*, who has a technological advantage in that her distribution cost per unit is lower than the principal's. We will focus on the case in which the agents distributional cost advantage is sufficiently large that the principal wants to use the agent to distribute its goods. To simplify notation and without loss of generality we normalize the agent's distribution costs to 0.

The principal and the agent can interact repeatedly over time $t = 0, 1, \dots, \infty$. The agent has the capacity to distribute up to \bar{k} units of the good per period. The goods be can sold on the local market to *households*. We assume there are arbitrarily many potential households in the economy.¹⁶ Each household has unit demand and is willing to pay \bar{p} for the good.

The agent has no capital and enjoys limited liability. The agent can also walk away from the arrangement in any period. The timing is as follows.

- At the beginning of period t , the principal gives the agent some amount, k_t , of the good for the agent to sell.
- The agent sells the goods and realizes a cash flow of $\bar{p}k_t$.¹⁷
- The agent then makes a transfer T_t to the principal and consumes the rest.
- Discounting occurs and then period $t + 1$ begins.

An *arrangement* is a relational contract between the principal and the agent consisting of a sequence of functions, $\{k_t, T_t\}_{t=0}^{\infty}$, mapping the relevant histories into quantities delivered by the principal and transfers made by the agent.

Both the principal and the agent are risk neutral and care about the expected present value of their per-period payoffs. We consider here a profit maximizing principal and discuss the connection to the NGO's objective of maximizing distribution in Section 3.4. The principal and agent have per period discount factors δ_P and δ_A respectively and we assume

¹⁶This assumption is convenient to preserve stationarity. However, it is not difficult to extend our results to a setting with a finite number of households.

¹⁷In Section 3.5, we extend our results to a setting in which sales revenue is risky and privately observed by the agent.

$\delta_P \geq \delta_A$. We use Π and U to denote these values:

$$\begin{aligned}\Pi &= E \left[\sum_{t=0}^{\infty} \delta_P^t (T_t - ck_t) \right] \\ U &= E \left[\sum_{t=0}^{\infty} \delta_A^t (\bar{p}k_t - T_t) \right]\end{aligned}$$

Though we do not explicitly incorporate termination of the arrangement, setting $k_t = T_t = 0$ for all $t \geq \tau$ is equivalent to terminating the arrangement at date τ .

We assume that the goods can be distributed and sold at an expected profit.

Assumption 1 (Profitability). $\delta_A \bar{p} - \delta_P c > 0$

This assumption guarantees the set of equilibria is non-trivial with a profit maximizing firm, but is not necessary if the principal is willing to spend resources in order to maximize distribution (see Section 3.4).

The history of the game observed by the principal at the start of period t is: $h_t^P \equiv \{k_s, T_s\}_{s=0}^{t-1}$. When choosing an action in period t , the history of the game for the agent is $h_t^A \equiv \{h_t^P, k_t\}$. A pure strategy for the principal is a sequence of functions $\{\sigma_t^P\}_{t=0}^{\infty}$ which determine the quantity of goods k_t to give to the agent as a function of h_t^P . A pure strategy for the agent is a sequence of functions $\{\sigma_t^A\}_{t=0}^{\infty}$ which for each period determine the agent's transfer T_t as a function of h_t^A . Mixed strategies are defined in the conventional way.

Because there is no external enforcement of contracts, the relationship will be governed by self-enforcing arrangements. An arrangement is said to be self-enforcing if the strategy pair that implements it is a Perfect Bayesian Equilibrium (PBE) of the game described above. While there are many PBE of the game, we will focus on *optimal* arrangements (i.e., the set of Pareto efficient PBE), which can be parameterized by the expected continuation utility of the agent (Abreu et al., 1990).

3.1 The Principal's Problem

Finding the optimal arrangement can be reduced to solving a dynamic program, which we undertake here. We use v to denote the continuation utility of the agent. The principal's maximization problem can be stated recursively as:

$$\Pi(v) = \sup_{K, T, W} \{T - cK + \delta_P \Pi(W)\}, \quad (\text{P})$$

subject to

$$T \in [0, \bar{p}K] \tag{2}$$

$$K \in [0, \bar{k}] \tag{3}$$

$$\delta_A W - T \geq V_{out} \equiv 0 \tag{4}$$

$$\bar{p}K - T + \delta_A W = v. \tag{5}$$

The liquidity constraint in (2) implies that neither agent nor principal has access to a borrowing technology. The only mechanism by which the consumption good is created is through selling to households.¹⁸ Equation (4) is the key incentive constraint, which can be interpreted as deriving from the principal’s inability to write an enforceable contract. This constraint ensures that the agent has incentive to make the transfer of T rather than consume it and forego her future continuation value.¹⁹ Finally, (5) is a standard promise-keeping constraint which requires the principal indeed delivers v in utility to the agent.²⁰

Due to the linearity of preferences, the solution to the dynamic program has a “bang-bang” feature. For low continuation values, the agent is compensated purely with continuation value and does not consume. For high continuation values, the agent consumes as much as possible.

Lemma 1. *There exists a solution to (P). The optimal policy is as follows:*

(i) For $v \in [0, \bar{v}]$, where $\bar{v} \equiv \bar{p}\bar{k}$,

$$K(v) = v/\bar{p}, \quad T(v) = \min\{\bar{p}K(v), \delta_A \bar{v}\}, \quad W(v) = \min\{\delta_A^{-1}v, \bar{v}\}.$$

(ii) For $v > \bar{v}$,

$$K(v) = \bar{k}, \quad T(v) = \max\{\bar{v}(1 + \delta_A) - v, 0\}, \quad W(v) = \max\{\bar{v}, \delta_A^{-1}(v - \bar{v})\}.$$

When the agent’s continuation value is low, the principal has to (inefficiently) restrict the amount of inventory in order to prevent the agent from stealing while respecting the

¹⁸This constraint eliminates the possibility of dynamic trading gains due to the agent’s relative impatience (Opp and Zhu, 2015).

¹⁹We have implicitly assumed that it is optimal to use a grim-trigger punishment if the agent deviates from the arrangement. This assumption is not restrictive (Mailath and Samuelson, 2006, Proposition 2.6.1).

²⁰To simplify the analysis we solve for the optimal arrangement assuming the agent does not have access to a private savings technology. This assumption is without loss of generality with a risk-neutral agent and thus our optimal arrangement is robust to introducing a private savings technology for the agent. See Kocherlakota (2004) for a discussion of how access to a private saving technology influences the optimal unemployment insurance in a setting with a risk averse agent.

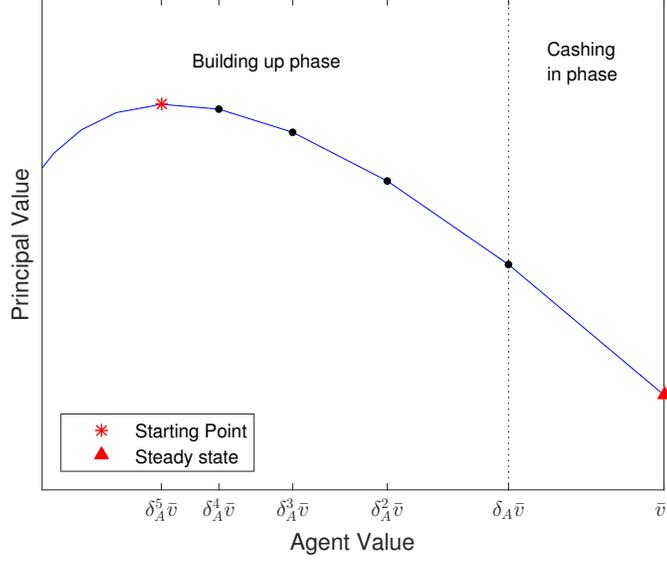


Figure 2: This figure illustrates the principal’s value function and the two phases of the arrangement for parameters such that the time-to-capacity is 5 periods (i.e., $N^* = 5$). The initial phase begins at the red asterisk and moves rightward along the black dots until the agent’s value reaches $\delta_A \bar{v}$, at which point investment is efficient and above which the agent begins to consume. The parameters used to generate the figure are $\delta_A = 0.78, \delta_P = 0.9, \mu = 0.25, \bar{k} = 1$.

promise-keeping constraint. The agent then transfers all of the proceeds back to the principal in exchange for a higher promised utility in the next period. When the agent’s continuation value is sufficiently high, this constraint stops binding and inventory reaches its efficient level. At this point, the agent only transfers a part of her revenues to the principal and consumes the rest.

3.2 Implementation

Our key theoretical insight is that the optimal arrangement can be implemented with a structure that is identical to the one used in Section 2.

Proposition 1. *The optimal arrangement can be implemented by a pair $(k_0^*, p_A^*) \in \mathbb{R}^2$. The principal provides the agent with a fixed initial endowment of the good, $k_0^* < \bar{k}$, and sets a fixed price, $p_A^* = \delta_A \bar{p}$ at which the agent can purchase all future units of the good.*

Intuitively, notice that while the agent’s liquidity constraint binds, the agent is induced to reinvest all her profits. Thus, she transfers an amount $T(v) = \bar{p}K(v) = v$ to the principal. In exchange, she obtains a continuation value of $\delta_A^{-1}v$, which means that in the next period she will receive a quantity of $K(\delta_A^{-1}v) = \frac{v}{\delta_A \bar{p}}$. Thus, the agent is effectively paying a price

of $p_A^* = \delta_A \bar{p}$ for each unit of inventory that she will receive in the next period. Once the liquidity constraint no longer binds (i.e. for $v \in (\delta_A \bar{v}, \bar{v}]$), the agent makes a transfer of $\delta_A \bar{v}$ in each period in exchange for \bar{k} units in the next period. Thus, again, the effective price is $\delta_A \bar{p}$.

Initially, the dynamics of the optimal arrangement resemble the example in Section 2. A profit-maximizing firm provides the agent with seed capital k_0^* and a subsidized price, which allows her to grow her business gradually. Once the agent has reached scale (i.e., $k_t = \bar{k}$), the agent begins to consume and enjoy the profits. The optimal arrangement thus involves two phases. The first is a *building up* phase in which the agent's business grows at a rate proportional to her discount factor. Although the agent does not consume during this period, her continuation value increases in this region as consumption nears. The second phase begins when the investment reaches its efficient level. At this point, the principal can no longer provide incentives to the agent by promising to grow the agent's business and thus the agent enters the *cashing in phase* and begins to consume. These dynamics are illustrated in Figure 2.

Remark. *The two distinct phases arise in part due to the linearity of the agent's preferences. With strictly concave utility, the optimal arrangement would have similar features, but the distinction between the two regions would be less dramatic as the agent would consume prior to reaching the efficient investment level. We maintain linear preferences so as to preserve the ease with which the arrangement can be implemented.*

3.3 Starting Small

Perhaps the most interesting feature of the optimal arrangement described above is that it involves “starting small.” That is, the initial endowment is below the agent's capacity and only over time does the agent's business grow to the efficient scale. In this section, we provide a closed-form solution for the optimal initial endowment and demonstrate several comparative statics.

In order to do so, let $N^* \equiv \min\{t : k_t = \bar{k}\}$, which denotes the number of period until the agent reaches capacity (henceforth, the “time-to-capacity”) under the optimal arrangement. Also, let $\gamma \equiv \frac{\delta_P}{\delta_A}$ denote the agent's impatience relative to the principal and $\mu \equiv \frac{\bar{p}-c}{\bar{p}}$ denote the profit margin of the good in the absence of any frictions.

Proposition 2. *The optimal initial endowment is $k_0^* = \delta_A^{N^*} \bar{k}$, where*

(i) If $\delta_A < \delta_P$ (i.e., $\gamma > 1$), then

$$N^* = \left\lceil \frac{\log\left(\frac{\mu}{1+\gamma(\mu-1)}\right)}{\log(\gamma)} \right\rceil,$$

where $\lceil x \rceil$ denotes the smallest integer weakly greater than x .

(ii) If $\delta_A = \delta_P$ (i.e., $\gamma = 1$) then $N^* = \left\lceil \frac{1-\mu}{\mu} \right\rceil$.

Notice that $N^* \geq 1$ (by Assumption 1) and therefore the initial endowment is always strictly less than \bar{k} . Intuitively, the endowment is designed to relax the liquidity constraints of the agent, but the principal never recovers the costs of these units. The revenues from first-period sales are the agent's rent.²¹ Hence, from the principal's perspective, there is no reason for the agent to operate at capacity in the first period. Provided $k_0 \geq \delta_A \bar{k}$, the revenues from first-period sales will be sufficient to purchase \bar{k} units for the next period.

Given Proposition 2, it is then straightforward to conduct comparative statics with respect to the two key parameters.

Corollary 1. *Under the optimal arrangement:*

(i) *The initial endowment is increasing in the profit margin (μ) and decreasing in the relative impatience of the agent (γ).*

(ii) *The time-to-capacity is decreasing in the profit margin (μ) and increasing in the relative impatience of the agent (γ).*

This corollary highlights that “starting small” is particularly important when profit margins are low and the agent is relatively impatient, both of which are likely to be important factors in our field experiment as well as in other applications of interest.

3.4 Relation to the NGO's Problem

When Assumption 1 holds the optimal arrangement for the NGO only differs from what a profit-maximizing firm would do in that the NGO would pick a higher initial k_0 . Indeed, the NGO would optimally set $k_0 = \bar{k} > k_0^*$ but would still charge p_A^* . Importantly, the budget constraint would not play a role because the operation is profitable.

When Assumption 1 fails, the NGO must take into account the available budget to finance the subsidized sale of the goods. Fortunately, as we show below, this can be handled

²¹Recall that the transfer in the first-period is for inventory to be received in the beginning of the second period.

by recasting the NGO's problem. Assume each unit adopted generates a social surplus of size Δ_S . Thus, the NGO objective of maximizing the total (discounted) social surplus can be written as:

$$\sum_{t=0}^{\infty} \delta_P^t (k_t \Delta_S).$$

The funding constraint requires that the present value of the cost of the operation cannot exceed the NGO's budget, denoted by B_0 ,

$$\sum_{t=0}^{\infty} \delta_P^t (ck_t - T_t) \leq B_0.$$

Instead of analyzing the problem of maximizing welfare for a given budget, consider the dual problem: minimize the total cost of the operation,

$$\min \sum_{t=0}^{\infty} \delta_P^t (ck_t - T_t), \tag{6}$$

subject to achieving a certain level of discounted social surplus, S ,

$$\Delta_S \sum_{t=0}^{\infty} \delta_P^t k_t \geq S.$$

Naturally, the objective in (6) can be rewritten as

$$\max \sum_{t=0}^{\infty} \delta_P^t (T_t - ck_t),$$

which is the same as the profit maximizing firm's objective.

Thus, the dual of the NGO problem (as formulated above) is the same as the profit-maximizing firm's problem we analyzed with the additional constraint of achieving a certain level of social surplus. One can then solve the NGO dual for different levels of social surplus, S , and choose the highest S for which the total cost satisfies the budget constraint. The dynamics of the relationship will continue to be characterized by a building-up phase where the agent grows her business followed by a cashing-in phase during which the agent consumes.

3.5 Risky Cash Flows and Private Information

Thus far, we have assumed that the agent can sell up to \bar{k} units each period at a fixed price of \bar{p} . In practice, there is likely to be uncertainty associated with both of these variables. Here, we extend our analysis by allowing the proceeds from the sale of the goods to be stochastic

and privately observed by the agent. We demonstrate two main findings.

First, the optimality of the arrangement in Proposition 1 is robust, though it may require the principal to provide the agent with access to a savings technology. Second, the dynamics of the relationship depend on cash flow realizations. As a result, growth is stochastic and the long-run outcome may be a termination of the relationship rather than operation at full capacity. In what follows, we elaborate further on these two findings.

To do so, let us denote the random variable representing the per unit proceeds or cash flow in period t by p_t , which is distributed according to the cumulative distribution function, F , with support $[p_{\min}, p_{\max}]$. F is i.i.d. over time with mean $\bar{p} \equiv \mathbb{E}[p_t]$.

In order for an arrangement to be an equilibrium, it must satisfy constraints similar to the benchmark case with the addition of an incentive compatibility constraint to ensure the agent will report the realized proceeds truthfully.

Given that the agent is risk neutral, we can interpret the promised value v directly as a favorable money balance the agent has with the principal. The principal gives this value v to the agent in two forms: 1) units to sell, denoted by $K(v)$, and 2) cash to be held in the agent's account during the period, denoted by $I(v)$. To fix ideas, we have normalized the intra-period gross return on this account to one. To satisfy the promise-keeping constraint, it must be that

$$\bar{p}K(v) + I(v) = v.$$

Thus, we can think of the units that the agents buys as having a beginning of period price of \bar{p} . At the end of the period t , the agent will have $p_t K(v) + I(v)$ dollars that she can choose to reinvest or consume. Note that for every marginal dollar the agent gives the principal she must be promised at least δ_A^{-1} dollars in the next period; otherwise the agent would prefer to consume than to transfer cash to the principal. Thus, to incentivize reinvestment, the principal can provide the agent with a savings technology with return δ_A^{-1} in addition to a beginning of period price per unit of capital, \bar{p} . Notice that with these two prices, the agent is indifferent as to how v is split into $K(v)$ and $I(v)$ and also indifferent as to how $p_t K(v)$ is split into consumption and transfers back to the principal. Therefore, such an arrangement is incentive compatible for any choice of K and transfers.

Proposition 3. *When cash flows are risky and privately observed by the agent, the optimal arrangement can be implemented with an initial endowment of units and cash, allowing the agent to buy units in period $t + 1$ for a price $p_A^* = \delta_A \bar{p}$ at the end of period t , and providing the agent with a savings technology that delivers an expected gross return of δ_A^{-1} .*

Now that we have determined the prices that satisfy incentive compatibility and the promise-keeping constraint, we think of the principal's problem as a single-person decision

problem with two components:

(1) A portofolio choice problem at the beginning of the period: how much to invest in the risky asset $K(v)$ (with the remainder, $I(v) = v - \bar{p}K(v)$, invested in the safe asset).

(2) A consumption-savings decision at the end of the period: how much to allow the agent to consume, $C(v, p)$, versus how much to save for the future, $T(v, p)$.

If the degree of asymmetric information is small, in particular, if $p_{\min} \geq \delta_A \bar{p}$ then the solution to both of these problems is the same as when output is deterministic. Regarding (1), when $p_{\min} \geq \delta_A \bar{p}$ (and therefore $p_{\min} \geq c$), the returns on the risky asset dominate the return on the safe asset. Thus $I(v) = 0$ provided the agent is not operating capacity. Regarding (2), there will be no need to use precautionary savings since if $p_{\min} \geq \delta_A \bar{p}$, upon reaching capacity, the agent will always have sufficient funds to repurchase the full stock in future periods even after a string of the worst possible realizations. Thus, if $\delta_A \bar{p} < p_{\min}$ the dynamics of the relationship are very similar to the ones for the deterministic case except that the growth is stochastic rather than deterministic. In the long run the agent will always operate at capacity.

When the degree of asymmetric information increases the solution to (1) and (2) may differ, resulting in different dynamics. First note that if $p_{\min} < c$ then the risky investment no longer dominates the safe investment and the solution to the portfolio choice problem may involve $I(v) > 0$ even for $K(v) < \bar{k}$. Even though the principal is risk neutral, once the agent's constraints are taken into account he is effectively risk-averse over the agent's end of period wealth. Hence, the principal may choose to allocate some wealth to the safe investment prior to reaching capacity.

Furthermore, if the difference in discount rates is small, the solution to the consumption-savings decision will delay consumption even beyond when full capacity is reached. This allows the agent to accumulate some precautionary savings that can be used to purchase inventory following a string of negative realizations. In this case, the dynamics of the relationship involve three phases (see Figure 3). There is the initial building up phase in which after high realizations of p the quantity allocated to the agent will increase and after very low ones it will decrease. Once the agent reaches full capacity then a new phase, the *precautionary-savings* phase begins. In this region, the agent does not consume nor does the quantity allocated grow. Instead, the agent deposits precautionary savings into a savings account, which translates into higher continuation values. Only when a sufficient level of precautionary savings has accumulated does the agent begin to consume (i.e., the *cashing-in* phase).

The amount of precautionary savings depends both on the relative impatience of the agent and on the distribution of risky cash flows. Regarding the long run outcome, there

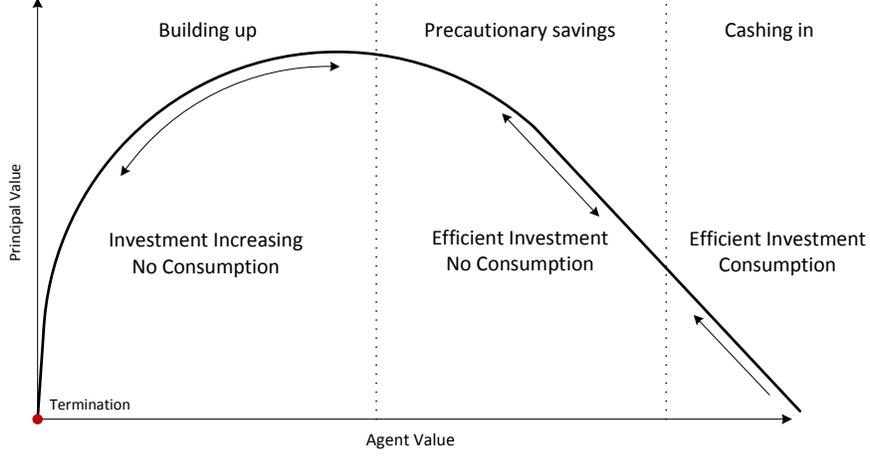


Figure 3: Illustration of the optimal arrangement with risky cash flows when it includes a precautionary savings phase (e.g., when δ_A is sufficiently close to δ_P).

are two cases: (i) If the savings buffer is sufficiently large (e.g., $\delta_A = \delta_P$) then the long run outcome is either full capacity or termination depending on the realization of prices; (ii) if the savings buffer is sufficiently small (i.e., $\delta_A \ll \delta_P$), then with probability one the relationship will eventually be terminated. The latter case is similar to the immiseration result of Thomas and Worrall (1990).

4 Testable Hypotheses

In this section, we revisit several assumptions of the model in order to develop hypothesis regarding factors that inhibit adoption and market development. We will later test these hypothesis in our field experiment.

4.1 Liquidity Constraints

In the model, we assumed that the agent had no wealth to invest in the project nor access to credit markets. Formally, this is represented by the constraint that $T \leq \bar{p}K$. As a result, a necessary feature of the optimal contract is to provide the agent with an initial endowment or “seed capital” (see Propositions 1 and 2). If instead, the agent did not face liquidity constraints then this feature would not be necessary. Rather, the optimal arrangement would involve an additional transfer to the principal at $t = 0$ in the amount \bar{v} and immediately moving to the steady state (i.e., where the agent continuation value is \bar{v}). In this case, the principal can extract the full surplus, and hence there is no need to start small. This

observation naturally leads us to the following hypothesis.

Hypothesis 1. *If the agent has sufficient initial wealth or access to credit markets then the performance of the arrangement should not depend on the size of the initial endowment.*

Thus, by varying the size of the initial endowment, we can evaluate the extent to which the liquidity constraints are relevant.

4.2 Risk Aversion or Pessimism

In the previous section, we argued that the optimal arrangement is robust to settings where cash flows are stochastic. There, we maintained the assumption that the agent is risk neutral and has the correct subjective beliefs about the distribution of cash flows.

If, on the other hand, the agent is risk averse or pessimistic about demand for the product then she may be unwilling to reinvest sales revenue to buy more units. One way to overcome such an aversion is to provide the agent with insurance in case the agent is unable to sell units at a profit. Thus, if demand turns out to be sufficiently low, then the agent can simply return the units without losing her investment.

Hypothesis 2. *If the agent is risk averse or pessimistic about the ability to sell the good for a profit then providing the “right to return” unsold units should improve the performance of the arrangement.*

4.3 Consumer Uncertainty and Learning

As discussed in the introduction, the literature documents a number of demand-side barriers. Several of these barriers pertain to the information available to the customer. For instance, customers are likely to be skeptical about the benefits and durability of the new technology (Feder and Slade, 1984; Conley and Udry, 2001; Giné and Yang, 2009). There is suggestive evidence that this problem has been caused in part by the proliferation of cheap and unreliable products in many sectors (e.g., counterfeit products). This general lack of information or skepticism may lead to an adverse selection problem between the agent and consumer, thereby reducing sales. Even absent an adverse selection problem, if households are uncertain about the quality of the good and risk averse, then they may be unwilling to invest in the new technology.

As suggested by Levine et al. (2013), providing the customer with a free-trial period may help to overcome these informational barriers. In our setting, the free-trial period also gives customers the chance to experience the financial benefits (i.e., reducing expenditures

on kerosene), thereby relaxing the liquidity constraints of households. For both of these reasons, we formulate a third testable hypothesis.

Hypothesis 3. *If consumers are uncertain about the quality of the good and/or are liquidity constrained then providing the agent with a “loaner” designated to provide customers with a free-trial period before their purchase should increase the performance of the arrangement.*

5 Field Experiment

In this section we describe our field experiment. The purpose of the experiment was (1) to evaluate the overall performance of the optimal arrangement implied by the model, and (2) to test the hypothesis formulated in the previous section.

5.1 Experimental Design

To conduct the experiment, we partnered with BRAC Uganda, a large non-profit organization. BRAC has a network of community health providers (CHPs) from which we recruited our “agents”. Effectively, there is one CHP per village; and prior to our intervention, these CHPs worked as vendors of health related consumable goods such as soap, sanitary pads, and malaria pills that they acquired from BRAC.

We visited 8 BRAC branches in rural Uganda. The trial in the first four branches ran from April 2013 to April 2014 and in the second wave of 4 branches from January to June 2014. Each branch was selected based on having low penetration of grid connections and limited distribution of low-cost solar lights.

A BRAC branch has a few dozen microfinance groups, each with 20 or so women. We divided each branch into 4 zones, each of which typically had 10 or more microfinance meetings. A BRAC credit officer escorted us to four microfinance group meetings per zone. The meetings were geographically dispersed so that each vendor would have a catchment area of 200 or so households. Our goal was that each catchment area would have enough residents to support one vendor.

At the microfinance meetings we presented the solar lights and explained we were recruiting vendors to sell these lights. We provided vendors with several different models of lights. One of these was the Firefly Mobile, produced and distributed by Barefoot Power. This solar light is bright enough for reading and studying, and it can also charge a mobile phone. We also introduced the basic Firefly, which is not quite as bright and cannot not charge a phone. The wholesale and suggested retail prices were \$20 and \$26 for the Firefly mobile and \$12 and \$16 for the basic.

We invited one woman per microfinance meeting to a recruitment meeting. If more than one woman at a recruitment meeting was interested, we gave preference to one who had access to SMS text messaging. In a few cases we asked the BRAC credit officer privately for a recommendation. If a recommendation from the credit officer was not possible, we selected the vendor who first expressed interest. During the recruitment meeting, we trained vendors on how to use the lights and informed them of their economic benefits.

Our training involved explaining features of the light, the operation of the light, and the terms of the retail sales offer (such as the one-year warranty). We showed vendors how the light can save customers money, where the savings on kerosene can quickly sum up to more than the cost of the light.

We anticipated that many of the vendors' customers would be liquidity constrained. Thus, we explained to the vendors the advantages of several sales offers that overcome liquidity constraints: layaway (which Guiteras et al. (2014) found worked well selling water filters in Bangladesh), installments (which Levine et al. (2013) found worked well selling cookstoves in urban and in rural Uganda), and selling via a rotating savings and credit association (ROSCA). In a ROSCA, a group of customers pool their funds each meeting to purchase one light. The group continues until all customers have purchased a light. ROSCAs are common in this part of Uganda.

After completing training, we gave each interested woman a solar light. She was then asked to pay for the light with mobile money over the next several weeks. The purpose of this exercise was two-fold: first, we hoped to familiarize potential vendors with the type of sales offer to use with their own customers and second, to partially screen out vendors that were unlikely to perform well. All vendors completed payments, though several of them took longer than was originally specified.

After the field staff made invitations and prior to the recruitment meeting, we randomized half of the zones to receive an arrangement that included a trade-credit line of up to 4 lights and the other half of the zones received an arrangement that did not include trade-credit line.²² We offered all vendors the right to purchase lights at the wholesale price during the first meeting or at any point in the future.²³ In order to test Hypothesis 2 and Hypothesis 3, we conducted two additional (and orthogonal) randomizations. First, we randomized over whether the agent was provided the right to return unsold inventory. Second, we randomly

²²In a setting without enforcement, storage costs, or default costs, a trade-credit line is equivalent to providing the agents with an endowment of the good. We framed the arrangement as a "credit line" in attempt to recover our costs in cases that the vendor was unable to sell lights. About 50% of the vendors who received trade credit and placed one order did not make a second order. Among these vendors, we recovered approximately 30% of the inventory.

²³Setting the vendor price equal to the wholesale price is roughly consistent with a distribution cost equal to 10% of sales and a monthly discount factor of $\delta_A = 0.85$.

selected a subset of the vendors to which we provided a loaner light. The light was clearly labeled as “Property of BRAC” and the vendor was instructed to use the loaner light in order to give potential customers a free-trial period.

We held a separate recruitment for all of the potential vendors who received the same vendor arrangement (that is, all the women in the same zone who had the same arrangement). At the recruitment meeting we discussed strategies for how to sell lights. We emphasized offering the customer a free-trial period and time payments. We then introduced the vendors to their designated arrangement. At the end of the recruitment meeting we took initial orders. Vendors not allocated an endowment paid cash for their initial orders, while vendors offered the endowment paid cash only if ordering more than 4 lights.

Vendors sold lights throughout the month. Once a month Barefoot sent a text message asking vendors to SMS back with their next order. Barefoot made a delivery a week after the text message requesting orders. Vendors met the delivery driver at the BRAC branch headquarters to make their payment and accept delivery of the lights they had ordered.

5.2 Summary of Findings

The 62 vendors with no credit line had average sales of 1.8 lights (SD = 6.1, median = 0).²⁴ The 67 vendors who received the “optimal” arrangement, which included a credit line sold an average of 6.6 lights (SD = 7.3, median = 4). The difference is highly statistically significant ($p < 0.01$) with both a two-sided t-test and with a nonparametric sign-rank test. The difference becomes even more significant when we omit a single outlier who sold 44 lights without a credit line, as we do in the remainder of our analysis. The results discussed above are summarized in Figure 4.

A regression of $\log(1 + \text{sales})$ on the characteristics of the arrangement can be found in Tables A.2 and A.3. Again, providing vendors with trade credit predicts significantly higher sales across all model specifications. For the t-test on sales and the regression on $\log(\text{sales}+1)$ we cluster standard errors by branch, the unit of stratification.

Qualitative evidence also supports the statistical result that providing trade credit to ease liquidity constraints is important. In interviews, vendors not offered trade credit stated that they would have ordered and sold more lights if they had been given credit for their initial orders. They said that it was challenging to get money in advance from customers to order more lights. Conversely, vendors with trade credit stated they would not have

²⁴The actual sale of a light from vendor to household was not observable to us. We therefore refer to “sales” as the total number of lights distributed to the vendor (presumably to sell to customers) during the life of the study.

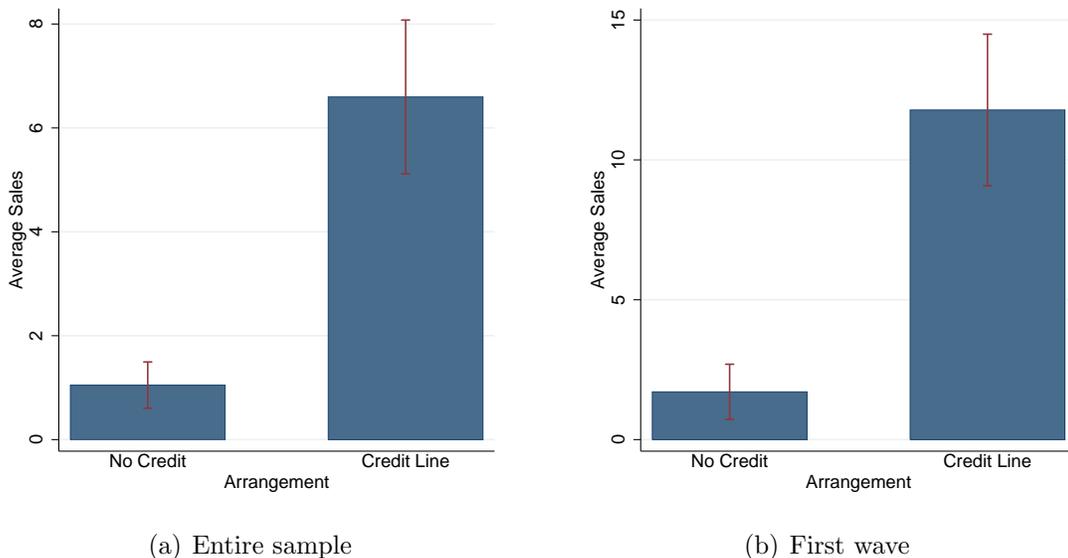


Figure 4: Credit Line versus No Credit: This graph shows the effect of including a line of trade credit in the arrangement for the entire sample (left-panel) and for the first-wave only (right panel). The blue bar is the average of total sales across all vendors who received an offer with the corresponding characteristic. The range between the red capped line indicates the 95% confidence interval.

sold as many lights without the initial credit. Thus, while the sample is small, our study provides strong evidence that liquidity constraints are important and providing trade credit can substantially increase sales. Although sales were much lower in the second wave of four branches, the credit line remains important.²⁵

To see if vendor uncertainty or pessimism was a relevant consideration, we gave a random sample of vendors the right to return unsold inventory for a full refund.²⁶ Perhaps surprisingly, there is little evidence to suggest that the right to return improves the performance of the arrangement (see Figure A.1). In the regression estimation the coefficient is (at best) only weakly positive and not statistically significant.

On the other hand, there is convincing evidence that consumer uncertainty is important factor. Within each of the study arms, we provided a randomly selected subset of the vendors

²⁵Our qualitative interviews partially explain why sales were so much lower in the second wave. Vendors suggested sales were low because competitors were selling low-quality lights for lower prices. As most potential customers could not detect the quality difference, demand was low for the Barefoot lights. In addition, two of the branches in the second wave were in areas with higher NGO penetration, so consumers may have become used to receiving free or deeply discounted goods thereby reducing their willingness to pay.

²⁶Vendors who received trade credit could return inventory rather than repay their debt, whereas vendors with no credit could return their inventory for cash.

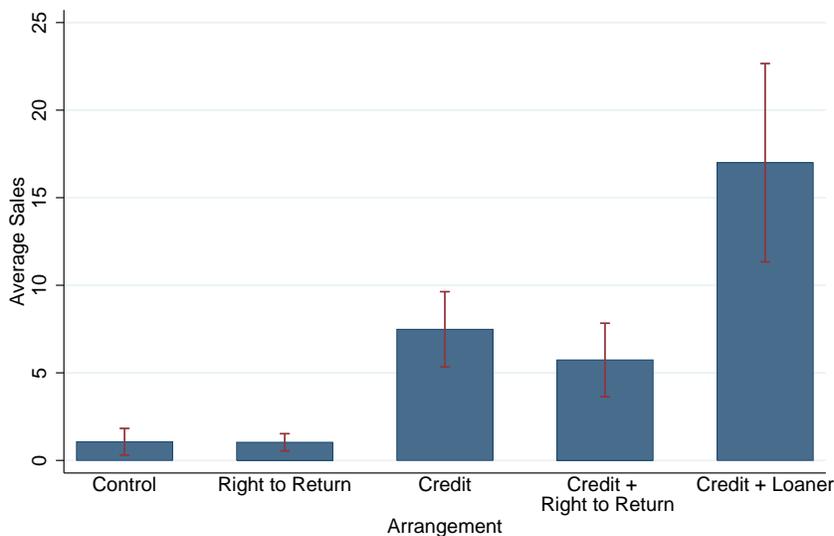


Figure 5: Average Sales by Arrangement. This graph shows the average total sales of vendors by the characteristics of the arrangement they were offered. The blue bar is the average of total sales across all vendors who received an offer with the corresponding characteristics. The range between the red capped line indicates the 95% confidence interval.

a loaner light to rotate among potential consumers. In interviews, vendors with free trials reported that they lent it to at least 5 households. The average vendor without a loaner light sold about 3 lights. The vendors allocated a loaner light sold more than 3 times as many (mean=11.2, median=10, SD=9.7). These results are illustrated in Figure A.2.

Figure 5 summarizes our main results from the field experiment. Providing agents with either a credit line or a loaner light had a significantly positive impact on sales, while the right to return did not. Perhaps most notable is the increase in sales observed via an arrangement that provided the agent with *both* a credit line and a loaner light.

5.3 Discussion

Our results suggest that vendor credit constraints and consumer uncertainty (or lack of information) are important barriers, while vendor uncertainty is not. One plausible explanation for the last finding is that vendors rarely purchased inventory and then sought customers. Instead, despite our encouragement to do otherwise, 70% of vendors reported waiting for a client to pay cash for the light prior to ordering it. Thus, while vendor uncertainty does not appear to be an important factor in our experiment, it may be more crucial in other settings where vendors use different sales techniques.

In the model, vendors retain earnings and increase their inventory and, on average, should grow sales over time at a rate proportional to δ_A^{-1} until reaching capacity. Unlike our model,

the initial burst of sales in our experiment only rarely led to a sustained growth. Even for vendors with trade credit, by four months after the first recruitment, the average sales was less than 1.5 lights per month. This could be explained by customer defaults, or by the vendors simply not being able to sell the lights at a profitable price. However, as we discuss below, we believe there were two other key reasons why vendors' sales growth was lower than predicted.²⁷

Factors that Inhibited Growth

Inability to Save. Many vendors noted that it was difficult to retain cash from sales until the next order, even if the next order was only a few weeks away. Vendors with cash in hand reported being subject to a lot of demands and found it difficult to avoid using it before the next delivery. Indeed, our model suggests that a savings technology may be necessary when cash flows are risky (see Proposition 3). The difficulty in saving might have been exacerbated by the fact that the vendors in our study were all women.²⁸ There are several RCTs that have studied ways in which to facilitate the commitment to savings.²⁹ In practical applications, we believe it will be important to find ways in which to build such mechanisms into the arrangement.

One way to avoid the savings problem is to have customers pay the principal directly using an electronic payment technology such as mobile money. With this technology, the vendor is not required to handle and save cash between delivery dates. Instead, the vendor could have an account with the producer whose balance increases whenever customers make payments. Moreover, the use of both mobile phones and mobile money is already widespread in rural Uganda. For instance, 98% of our vendors reported owning a mobile phone and 83% reported having prior experience with mobile money.

Failure of the Credit Chain. Existing literature has shown that poor households face credit

²⁷While there are certainly other factors that may have contributed to a lack of growth, we believe most of them can be overcome with proper screening and training. We intentionally did not engage in much screening so as to obtain cross-sectional variation. For instance, vendors reported they intended to work a median of 20 hours a week selling solar lights. At the same time, they reported working about 40 hours a week at their current jobs and having an average of 5 children at home. Thus, it seems unlikely that vendors would have anywhere near the 20 hours a week they forecast to market solar lights. Amplifying this concern, the median vendor reported taking 60 minutes to travel to the BRAC branch office. Thus, most vendors faced meaningful transaction costs. In addition, vendors had imperfect recall of the content of our product training. Almost all (98%) knew to keep the lamp out of the sun when charging the solar panel. A lower share (79%) knew the manufacturer's warranty was for one year, and even fewer (60%) recalled that the solar light should charge for 2 days prior to its first use. Inability to explain product features may also have reduced their sales effectiveness.

²⁸See Bobonis (2009), Bobonis et al. (2013) and references therein regarding the allocation of resources within households and domestic violence.

²⁹See Ashraf Nava and Yin (2006) and Basu (2014).

constraints, which is an important factor limiting adoption rates (Cole et al., 2013; Tarozzi et al., 2014). Therefore, as mentioned earlier, during their training we emphasized several techniques vendors should use to help alleviate these constraints and boost their sales (i.e., installments, lawaway, ROSCA). Despite this encouragement, only 30% of vendors actually employed these techniques according to surveys. Overall, vendors seemed generally unwilling to extend credit to their customers even when the vendors themselves were extended credit and vendors were explicitly encouraged to offer credit to their customers. Post-study interviews suggest vendors' unwillingness to extend credit was due to a fear of customers defaulting on their payments.

Is Technology the Answer?

When vendors are both unwilling to offer consumer credit and cannot easily save to increase their inventory, it is efficient if the producer or distributor can offer credit directly to consumers and receive payments from them. By doing so, the vendor's role is essentially reduced to that of a credit officer. At the same time, such an arrangement could lead to potential moral hazard problems of vendors selecting customers carelessly and/or not retrieving products when consumers default.

Some distributors and producers have combined payments via mobile money with a "kill" switch on their products to overcome these challenges as well as the frictions in our model. The advantage of the kill switch is that it makes the product worthless to the customers if they do not make payments and to vendors if they abscond with inventory, thereby increasing customer repayment rates and reducing vendors' ability to extract rents. For example, M-Kopa Solar sells a home solar system with a solar panel, three small lights, a phone charger, and a radio for an initial \$35 deposit, followed by 365 [daily] payments of 45 cents paid by mobile money.³⁰ A comparable system would retail for a cash price of roughly \$75.

M-Kopa has higher production costs because each system must contain many of the capabilities of a mobile phone. The benefits are apparently large; the present value of payments is about \$153 at a discount rate of 100% per year and \$185 at a discount rate of 20% per year. Thus, a consumer could purchase a similar system for about half the present value of payments if paying in cash. Nevertheless, M-Kopa has sold more than 500,000 systems in East Africa and is growing at a rate of 500 new systems per day.³¹ The success of M-Kopa and its peers suggests the severity of the credit constraints, savings constraints, and moral hazard emphasized in our model are of first-order importance.

It is important to note that many products are not obviously compatible with a kill

³⁰<http://www.m-kopa.com/products>

³¹<http://solar.m-kopa.com/about/company-overview>, date accessed: September 26, 2017.

switch (e.g., cook stoves, water filters, malaria nets). In this regard, solar lights are unique because they operate on electricity whereas many of the products that lack widespread adoption in developing countries do not. Thus, M-Kopa’s technological strategy is likely to be prohibitively expensive to incorporate into many of the other products to which our model applies.

6 Concluding Remarks

Markets for new technologies in emerging economies develop slowly due to a variety of economic frictions. Products that would enhance the welfare of many poor households are not adopted as fast as socially desirable. Most of the literature has focused on addressing demand-side barriers to product adoption. In this paper, we have developed a theory to address supply-side barriers. Two important issues to be overcome are the liquidity constraints and lack of enforceable contracts. The optimal arrangement involves providing the agent with a “small” amount of seed capital as well as the option to buy more units in the future at a fixed price. Interestingly, this arrangement is optimal for both profit maximizing firms and non-profit organizations with a limited budget.

We conducted a field experiment to test our theory. The evidence clearly indicates that liquidity constraints and consumer uncertainty about product quality are important factors. Providing the agent with credit increases sales by 3-4 times compared to a standard contract. The combination of both credit and a loaner light increases sales nearly tenfold. Growth, however, was lower than expected. As we learned from interviewing the agents, this lack of growth was largely due to vendors’ inability to save even for the short periods of time between orders and a failure of the “credit chain” (i.e., vendors were unwilling to offer credit to their customers). Therefore, we believe it is critical for researchers to incorporate reliable savings technologies in future experiments and encourage the flow of credit. Emerging technologies, such as those used by M-Kopa, have developed to help overcome these obstacles, but they come at a non-trivial increase in cost and are likely to be prohibitively expensive for many products. An interesting direction for future work is to explore the extent to which these obstacles can be partially overcome using appropriately designed incentive schemes rather than integrating expensive hardware.

References

- Abreu, Dilip, David Pearce, and Ennio Stacchetti**, “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 1990, *58*, 1041–1063.
- Abuquerque, Rui and Hugo A Hopenhayn**, “Optimal Lending Contracts and Firm Dynamics,” *The Review of Economic Studies*, 2004, *71*, 285–315.
- Ashraf Nava, Dean Karlan and Wesley Yin**, “Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines,” *The Quarterly Journal of Economics*, 2006, *121* (2), 635–672.
- Ashraf, Nava, James Berry, and Jesse M Shapiro**, “Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia,” *American Economic Review*, 2010, *100*, 2383–2413.
- Baker, George, Robert Gibbons, and Kevin J. Murphy**, “Relational Contracts and the Theory of the Firm,” *The Quarterly Journal of Economics*, 2002, *117* (1), 39–84.
- Basu, Karna**, “Commitment savings in informal banking markets,” *Journal of Development Economics*, 2014, *107*, 97–111.
- Bobonis, Gustavo**, “Is the Allocation of Resources within the Household Efficient? New Evidence from a Randomized Experiment,” *Journal of Political Economy*, 2009, *117* (3), 453–503.
- , **Melissa Gonzalez-Brenes, Roberto Castro, Gustavo Bobonis Melissa Gonzalez-Brenes, and Roberto Castro**, “Public Transfers and Domestic Violence: The Roles of Private Information and Spousal Control,” *American Economic Journal: Economic Policy*, 2013, *5* (1), 179–205.
- Bold, Tessa, Kayuki C. Kaizzi, Jakob Svensson, and David Yanagizawa-Drott**, “Lemon Technologies and Adoption: Measurement, Theory and Evidence from Agricultural Markets in Uganda*,” *The Quarterly Journal of Economics*, mar 2017, *103*, 76–84.
- Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak**, “Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh,” *Econometrica*, 2014, *82* (5), 1671–1748.
- Chiappori, Pierre-André, Krislert Samphantharak, Sam Schulhofer-Wohl, and Robert M. Townsend**, “Heterogeneity and risk sharing in village economies,” *Quantitative Economics*, 2014, *5* (1), 1–27.

- Clementi, Gian Luca and Hugo A. Hopenhayn**, “A Theory of Financing Constraints and Firm Dynamics,” *Quarterly Journal of Economics*, 2006, 121, 229–265.
- Cole, Shawn, Xavier Giné, Jeremy Tobacman, Petia B Topalova, Robert M Townsend, and James I Vickery**, “Barriers to household risk management: evidence from India,” *American Economic Journal: Applied Economics*, 2013, 5 (1), 104–135.
- Conley, Timothy and Christopher Udry**, “Social Learning through Networks: The Adoption of New Agricultural Technologies in Ghana,” *American Journal of Agricultural Economics*, 2001, 83 (3), 668–673.
- Conley, Timothy G and Christopher R Udry**, “Learning About a New Technology : Pineapple,” *American Economic Review*, 2010, 100 (1), 35–69.
- Cunat, Vincent and Vicente Cuñat**, “Trade credit: Suppliers as debt collectors and insurance providers,” *Review of Financial Studies*, 2007, 20 (2), 491–527.
- de Mel, Suresh, David McKenzie, and Christopher Woodruff**, “Returns to Capital in Microenterprises: Evidence from a Field Experiment,” *The Quarterly Journal of Economics*, 2008, 123 (4), 1329–72.
- DeMarzo, Peter and Michael Fishman**, “Agency and Optimal Investment Dynamics,” *Review of Financial Studies*, 2007, 20, 151–188.
- DeMarzo, Peter M. and Michael J. Fishman**, “Optimal long-term financial contracting,” *Review of Financial Studies*, 2007, 20, 2079–2128.
- Demarzo, Peter M., Michael J. Fishman, Zhiguo He, and Neng Wang**, “Dynamic Agency and the q Theory of Investment,” *Journal of Finance*, 2012, 67 (6), 2295–2340.
- Dubois, P, B Jullien, and T Magnac**, “Formal and informal risk sharing in LDCs: Theory and empirical evidence,” *Econometrica*, 2008, 76 (4), 679–725.
- Duflo, Esther, Michael Kremer, and Jonathan Robinson**, “Nudging farmers to use fertilizer: Theory and experimental evidence from Kenya,” *American Economic Review*, 2011, 101 (6), 2350–2390.
- Dupas, Pascaline and Jessica Cohen**, “Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment,” *The Quarterly Journal of Economics*, 2010, 125, 1–45.

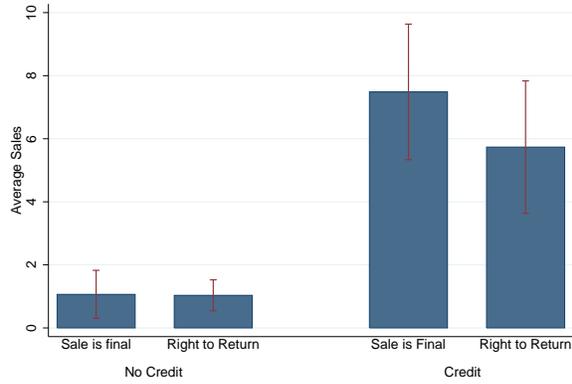
- Feder, Gershon and Roger Slade**, “The Acquisition of Information and the Adoption of New Technology,” *American Journal of Agricultural Economics*, 1984, 66 (3), 312–320.
- Foster, Andrew D. and Mark R. Rosenzweig**, “Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture,” *The Journal of Political Economy*, 1995, 103 (6), 1176–1209.
- Giné, Xavier and Dean Yang**, “Insurance, credit, and technology adoption: Field experimental evidence from Malawi,” *Journal of Development Economics*, 2009, 89 (1), 1–11.
- Guiteras, R, David I Levine, T Polley, and B Quistorff**, “Credit Constraints, Present Bias and Investment in Health: Evidence from Micropayments for Clean Water in Dhaka,” *Unpublished*, 2014.
- Kaboski, Joseph P. and Robert M. Townsend**, “A Structural Evaluation of a Large-Scale Quasi Experimental Microfinance Initiative,” *Econometrica*, 2011, 79 (5), 1357–1406.
- Karaivanov, Alexander and Robert Townsend**, “Dynamic Financial Constraints: Distinguishing Mechanism Design from Exogenously Incomplete Regimes,” *Econometrica*, 2013, 82 (3), 887–959.
- Kocherlakota, Naryana R**, “Figuring out the impact of hidden savings on optimal unemployment insurance,” *Review of Economic Dynamics*, 2004, pp. 541–554.
- Levin, Jonathan**, “Relational Incentive Contracts,” *American Economic Review*, 2003, 93, 835–857.
- Levine, David I, Theresa Beltramo, Garrick Blalock, and Carolyn Cotterman**, “What Impedes Efficient Adoption of Products? Evidence from Randomized Variation in Sales Offers for Improved Cookstoves in Uganda,” *IRLE working paper*, 2013.
- Ligon, Ethan, Jonathan P Thomas, and Tim Worrall**, “Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies,” *The Review of Economic Studies*, 2002, 69 (1), 209–244.
- Mailath, George J and Larry Samuelson**, *Repeated Games and Reputations*, Oxford University Press, 2006.
- Opp, Marcus M and John Y Zhu**, “Impatience Versus Incentives,” *Econometrica*, 2015, 83 (5), 1601–1617.

- Quadrini, Vincenzo**, “Investment and liquidation in renegotiation-proof contracts with moral hazard,” *Journal of Monetary Economics*, 2004, *51*, 713–751.
- Rampini, A. and S. Viswanathan**, “Collateral and capital structure,” *Journal of Financial Economics*, 2013, *109* (2), 466–492.
- Rampini, Adriano A. and S. Viswanathan**, “Collateral, risk management, and the distribution of debt capacity,” *Journal of Finance*, 2010, *65* (6), 2293–2322.
- Tarozzi, Alessandro, Aprajit Mahajan, Brian Blackburn, Dan Kopf, Lakshmi Krishnan, and Joanne Yoong**, “Micro-loans, insecticide-treated bednets, and malaria: Evidence from a randomized controlled trial in Orissa, India,” *American Economic Review*, 2014, *104* (7), 1909–1941.
- Thomas, Jonathan and Tim Worrall**, “Income fluctuation and asymmetric information: An example of a repeated principal-agent problem,” *Journal of Economic Theory*, 1990, *51* (2), 367–390.
- and —, “Foreign Direct Investment and the Risk of Expropriation,” *The Review of Economic Studies*, 1994, *61*(1), 81–108.
- Townsend, Robert M.**, “Risk and Insurance in Village India,” *Econometrica*, 1994, *62* (3), 539–591.
- Townsend, Robert M. and Sergio S. Urzua**, “Measuring the Impact of Financial Intermediation: Linking Contract Theory to Econometric Policy Evaluation,” *Macroeconomic Dynamics*, 2009, *13* (Suppl S2), 268–316.
- Troya-Martinez, Marta**, “Self-Enforcing Trade Credit,” *Working Paper*, 2015.

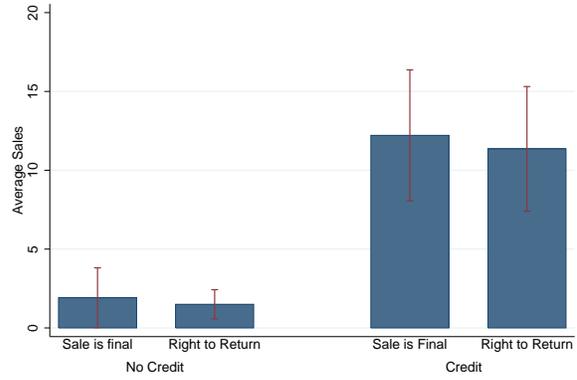
A Tables and Figures

Education	None 28%	Primary 60%	Secondary 12%	
Occupation	Retail 51%	Agriculture 40%	Livestock 13%	Other 40%
	Mean	Median	Std Dev	
Age	39.1	38	9.55	
Experience	9.4	6	10.2	
Work hours/day	8.46	8	3.98	
Children	4.91	5	2.74	
Residency	17.8	15	13.3	
Travel time to BRAC	58.3	60	37.7	
Kerosene expenditure				
Pre-solar	3,264	3,000	2,799	
Post-solar	305	0	826	
Female	100%			
Married	70%			
Own mobile phone	98%			
Use SMS	47%			
Use Mobile Money	83%			

Table A.1: Summary of vendor characteristics reported in surveys.



(a) Entire sample



(b) First wave

Figure A.1: Addressing Vendor Uncertainty with the Right to Return. This graph shows the effect of including the right to return unsold inventory in the arrangement for the entire sample (left-panel) and for the first-wave only (right panel). The blue bar is the average of total sales across all vendors who received an offer with the corresponding characteristics. The range between the red capped line indicates the 95% confidence interval.

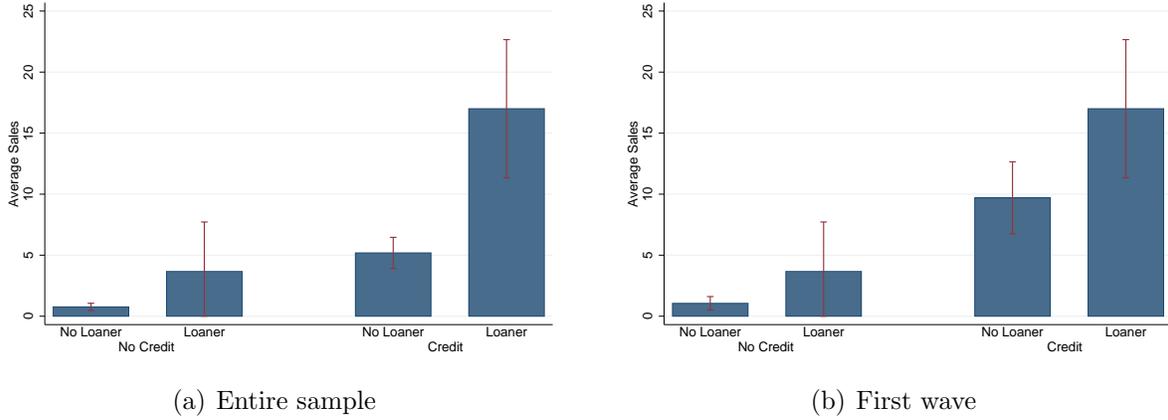


Figure A.2: Addressing Consumer Uncertainty with a Loaner Light. This graph shows the effect of including the right to return unsold inventory in the arrangement for the entire sample (left-panel) and for the first wave only (right panel). The bar denotes the average of total sales across all vendors who received an offer with the corresponding characteristics. The range between the red capped line indicates the 95% confidence interval.

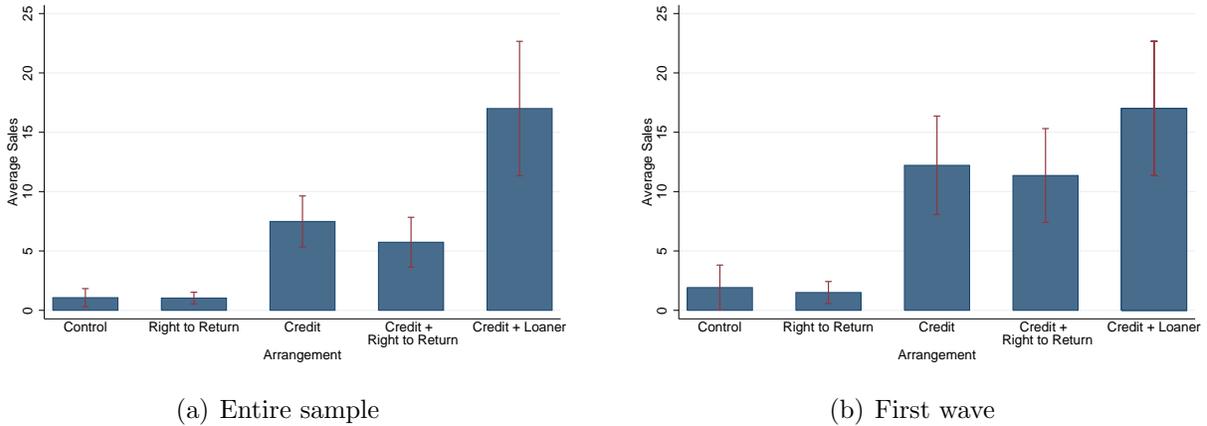


Figure A.3: Average Sales by Offer: This graph shows the average total sales of vendors by the characteristics of the arrangement they were offered for the entire sample (left-panel) and for the first-wave only (right panel). The blue bar is the average of total sales across all vendors who received an offer with the corresponding characteristics. The range between the red capped line indicates the 95% confidence interval.

Table A.2: Regression Estimates for Full Sample. This table gives the results of the estimation of the regression equation

$$\log(1 + \text{total sales}_i) = \alpha + \beta \text{Arrangement Characteristics}_i + \epsilon_i, \quad (7)$$

where the unit of observation (i) is the vendor, and $\text{Arrangement Characteristics}_i$ is a vector of attributes associated with the arrangement offered to that vendor. For instance, column (2) contains the estimates from the model in which $\text{Arrangement Characteristics}_i$ is a pair of dummy variables ($\text{Credit}_i, \text{Loaner}_i$), where the first dummy indicates whether vendor i 's arrangement included trade credit and the second dummy indicates whether vendor i 's arrangement included a loaner light. Standard errors are in parenthesis below the estimated coefficients and are clustered at the branch level. Standard errors are in parenthesis below the estimated coefficients and are clustered at the branch level. The data used for this estimation includes the full sample (i.e., both first and second wave).

	(1)	(2)	(3)	(4)	(5)
Credit	1.086*** (5.11)	1.078*** (5.03)	1.087*** (5.13)	1.078*** (5.05)	1.071*** (4.86)
Loaner		1.232** (2.77)		1.227** (2.75)	0.760 (1.57)
Right to Return			-0.122 (-1.45)	-0.100 (-1.15)	-0.115 (-1.30)
First wave					0.701** (2.90)
Constant	0.500*** (3.88)	0.361*** (3.90)	0.561*** (5.53)	0.411*** (6.33)	0.189 (1.24)
R ²	0.266	0.407	0.269	0.409	0.496
N	129	129	129	129	129

t-statistics in parenthesis

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.3: Regression Estimates for First Wave. This table gives the results of the estimation of the regression equation

$$\log(1 + \text{total sales}_i) = \alpha + \beta \text{Arrangement Characteristics}_i + \epsilon_i, \quad (8)$$

where the unit of observation (i) is the vendor, and $\text{Arrangement Characteristics}_i$ is a vector of attributes associated with the arrangement offered to that vendor. For instance, column (2) contains the estimates from the model in which $\text{Arrangement Characteristics}_i$ is a pair of dummy variables ($\text{Credit}_i, \text{Loaner}_i$), where the first dummy indicates whether vendor i 's arrangement included trade credit and the second dummy indicates whether vendor i 's arrangement included a loaner light. Standard errors are in parenthesis below the estimated coefficients and are clustered at the branch level. The data used for this estimation includes only the first wave.

	(1)	(2)	(3)	(4)
Credit	1.541*** (6.47)	1.537*** (6.73)	1.543*** (6.40)	1.539*** (6.66)
Loaner		0.763 (1.39)		0.770 (1.39)
Right to Return			0.0815 (0.88)	0.119 (1.92)
Constant	0.798** (3.77)	0.584** (3.27)	0.756** (3.50)	0.521** (3.28)
R ²	0.473	0.567	0.474	0.570
N	53	53	53	53

t-statistics in parenthesis

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B Proofs

Proof of Lemma 1. It is straightforward to verify that the stated policy satisfies (2)-(5). In particular, (2) binds from above for $v \leq \delta_A \bar{v}$, from below for $v \geq (1 + \delta_A)\bar{v}$ and is interior for $v \in (\delta_A \bar{v}, (1 + \delta_A)\bar{v})$; (3) is slack for $v \in (0, \bar{v})$ and binds for $v \geq \bar{v}$; (4) binds for $v \leq \bar{v}$ and is slack otherwise. The rest of the proof is by construction. We will first construct the principal's value function under the stated policy and then verify that it indeed solves (P).

Notice that $W(\bar{v}) = \bar{v}$, hence \bar{v} is the steady state. In the steady state, each period, $K(\bar{v}) = \bar{k}$, the agent consumes $(1 - \delta_A)\bar{v}$ and therefore transfers $\delta_A \bar{v}$ to the principal. The principal's value function at \bar{v} is therefore

$$\begin{aligned} \Pi(\bar{v}) &= \delta_A \bar{v} - c\bar{k} + \delta_P \Pi(\bar{v}) \\ &= \frac{(\delta_A \bar{p} - c)\bar{k}}{1 - \delta_P} \end{aligned}$$

Next, consider any $v \in (\delta_A \bar{v}, \bar{v})$, so that there is one period until the steady state is reached (i.e., $W(v) = \bar{v}$). Under the stated policy, we have that

$$\Pi_0(v) = \delta_A \bar{v} - c \frac{v}{\bar{p}} + \delta_P \Pi(\bar{v}) \quad (9)$$

Now, fix any integer $N \geq 0$, we claim that for $v \in (\delta_A^{N+1} \bar{v}, \delta_A^N \bar{v})$, the principal's value function under the stated policy is given by

$$\Pi_N(v) = v\mu \sum_{t=0}^{N-1} \left(\frac{\delta_P}{\delta_A} \right)^t + \delta_P^N \Pi_0 \left(\frac{v}{\delta_A^N} \right) \quad (10)$$

We have already demonstrated the base case (i.e., $N = 0$) in (9). Suppose that (10) holds for some $n \geq 1$ and consider any $v \in (\delta_A^{n+2} \bar{v}, \delta_A^{n+1} \bar{v})$, under the stated policy we have that $K(v) = v/\bar{p}$, $T(v) = v$ and $W(v) = \delta_A^{-1}v$, therefore the principal's value function is given by

$$\begin{aligned} \Pi_{n+1}(v) &= \left(v - \frac{cv}{\bar{p}} \right) + \delta_P \Pi_n(\delta_A^{-1}v) \\ &= \mu v + \delta_P \left(\frac{v}{\delta_A} \mu \sum_{t=0}^{n-1} \left(\frac{\delta_P}{\delta_A} \right)^t + \delta_P^n \Pi_0 \left(\frac{v}{\delta_A^{n+1}} \right) \right) \\ &= v\mu \sum_{t=0}^n \left(\frac{\delta_P}{\delta_A} \right)^t + \delta_P^{n+1} \Pi_0 \left(\frac{v}{\delta_A^{n+1}} \right), \end{aligned}$$

which is of the form in (10) and thus verifying the claim. Next, for $v \in (\bar{v}, (1 + \delta_A)\bar{v})$, we

have that

$$\Pi(v) = g_0(v) = \bar{v}(1 + \delta_A) - v - c\bar{k} + \delta_P \Pi(\bar{v}).$$

Using an induction argument similar to the one above, for any integer $k \geq 0$ and $v \in \left(\bar{v} \sum_{t=0}^k \delta_A^t, \bar{v} \sum_{t=0}^{k+1} \delta_A^t\right)$

$$\Pi(v) = g_k(v) \equiv -c\bar{k} \sum_{t=0}^k \delta_P^t + \delta_P^k g_0 \left(\delta_A^{-k} \left(v - \bar{v} \sum_{t=0}^k \delta_A^t \right) \right).$$

We have thus constructed the principal's value function under the stated policy for all $v > 0$.

Before verifying optimality of the policy, it is useful to observe several properties of the value function. First notice that Π is piecewise linear and concave in v . Next, notice for $\gamma \neq 1$ that (10) can be written as

$$\Pi_N(v) = v\mu \left(\frac{\gamma^N - 1}{\gamma - 1} \right) - \gamma^N \frac{c}{p} v + \delta_P^N (\delta_A \bar{v} + \delta_P \Pi(\bar{v})), \quad (11)$$

which is differentiable with respect to v for $v \in (\delta_A^{N+1} \bar{v}, \delta_A^N \bar{v})$, $N \geq 0$ with a slope given by

$$\Pi'_N(v) = \mu \left(\frac{\gamma^N - 1}{\gamma - 1} \right) - (1 - \mu) \gamma^N. \quad (12)$$

When $\gamma = 1$, (12) becomes $\Pi'_N(v) = \mu N - (1 - \mu)$. Finally, for $v \in \left(\bar{v} \sum_{t=0}^k \delta_A^t, \bar{v} \sum_{t=0}^{k+1} \delta_A^t\right)$, $k \geq 0$, the slope of the value function is γ^k .

To verify that the stated policy is indeed optimal, notice that by substituting the promise-keeping constraint into the objective and (4), the problem can be restated as:

$$\sup_{K, T} \left\{ T - cK + \delta_P \Pi \left(\frac{v + T - \bar{p}K}{\delta_A} \right) \right\}$$

subject to $T \in [0, \bar{p}K]$, $K \in [0, \min\{\bar{k}, v/\bar{p}\}]$. Since Π is concave, it is enough to check that local deviations are not profitable. The second constraint always binds at the top under the stated policy, so we only need to consider a reduction in K of ϵ . If (2) also binds at the top (i.e., $v \leq \delta_A \bar{v}$) then this deviation also requires a small reduction in T to satisfy the first constraint, which leaves the continuation value unchanged, and therefore reduces the objective by $(\bar{p} - c)\epsilon$. If (2) does not bind from above (i.e., $v > \delta_A \bar{v}$), then the marginal benefit of this local deviation is $c + \delta_P \Pi'(W(v)) \frac{\bar{p}}{\delta_A}$, which is negative provided that $\Pi'(W(v)) \leq -\frac{c}{\bar{p}} \frac{\delta_A}{\delta_P}$. Noting that the value function constructed above has a slope of $-c/\bar{p}$ for $v \in (\delta_A \bar{v}, \bar{v})$ and is concave verifies that such a deviation is not profitable.

Next, consider a deviation from the stated policy $T(v)$. Note that the objective is increasing in T if $\Pi'(W(v)^+) \geq -\delta_A/\delta_P$, which holds if and only if $v < \delta_A\bar{v}$, in which case an increase in T violates the first constraint. For $v \in [\delta_A\bar{v}, (1 + \delta_A)\bar{v}]$, $\Pi'(W(v)^+) = \Pi'(\bar{v}^+) = -1 \leq -\delta_A/\delta_P$ and $\Pi'(W(v)^-) = \Pi'(\bar{v}^-) = -c/p > -\delta_A/\delta_P$. Therefore, neither increasing nor decreasing $T(v)$ is profitable. Finally, for $v > (1 + \delta_A)\bar{v}$, $\Pi'(W(v)^+) \leq \Pi'(W(v)^-) \leq -1 \leq -\delta_A/\delta_P$. Hence, the principal does strictly worse by increasing $T(v)$ in this region, and a reduction in $T(v)$ violates the first constraint. Thus, we have shown that no profitable deviations from the stated policy exists, which completes the proof. \square

Proof of Proposition 1. Fix an arbitrary $k_0 \leq \bar{k}$. We first claim the the arrangement (k_0, p_A^*) implements the policy in Lemma 1 where the initial continuation value of the agent is $v_0 \equiv k_0\bar{p}$. To see this, notice that $K(v_0) = K(k_0\bar{p}) = k_0$ and that the agent's revenue in the first period equals $k_0\bar{p}$. Clearly, the agent will never optimally purchase more than \bar{k} units. If $v_0 > \delta_A\bar{v}$, then the agent will optimally choose to purchase exactly \bar{k} (i.e., $T(v_0) = \delta_A\bar{p}\bar{k}$) and consume the rest. If $v_0 < \delta_A\bar{v}$, then the agent will optimally choose to purchase k_0/δ_A units and consume nothing in the initial period. In either case, the number of units the agent will have in the next period is $k_1 = \min\{k_0/\delta_A, \bar{k}\} = K(W(v_0))$ when the transfer will be $t_1 = \min\{\bar{p}K(W(v_0)), \delta_A\bar{v}\} = T(W(v_0))$. Therefore, the policy in the next period is also replicated. Since k_0 was chosen arbitrarily, this completes the proof of the claim.

What remains is to prove that $k_0^* < \bar{k}$. For this, it suffices to show that $v_0^* \in \arg \max_0 \Pi(v) < \bar{v}$. This follows immediately from the fact that the principal's value function is strictly decreasing on $(\delta_A\bar{v}, \bar{v})$ (see (9)) and the concavity of the value function (see the proof of Lemma 1). Thus, the optimal initial endowment is strictly less \bar{k} . \square

Proof of Proposition 2. Recall that the principal's value function is concave and for $\gamma > 1$ has a slope given by (12) for $v \in (\delta_A^{N+1}\bar{v}, \delta_A^N\bar{v})$, $N \geq 0$. Setting the slope equal to zero and solving yields $N_1 = \frac{\log(\frac{\mu}{1+\gamma(\mu-1)})}{\log(\gamma)}$. If N_1 is an integer, then $N^* = N_1$ and the principal's value function has slope zero over the interval $v \in (\delta_A^{N_1+1}\bar{v}, \delta_A^{N_1}\bar{v})$, is upward sloping to the left and downward sloping to the right. Therefore, $\delta_A^{N^*}\bar{v} \in \arg \max_v \Pi(v)$. If N_1 is not an integer, then $\Pi'(\delta_A^{N^*}\bar{v}^-) > 0 > \Pi'(\delta_A^{N^*}\bar{v}^+)$. Therefore, $\delta_A^{N^*}\bar{v} = \arg \max_v \Pi(v)$. By Proposition 1, the principal can achieve this value by providing the agent with an initial endowment of $k_0^* = \delta_A^{N^*}\bar{v}/\bar{p} = \delta_A^{N^*}\bar{k}$, in which case it will take the agent N^* periods to reach capacity. For $\gamma = 1$, the result can be shown using the same technique by substituting $\Pi'_N(v) = \mu N - (1 - \mu)$ and hence $N_1 = \frac{1-\mu}{\mu}$. \square

Proof of Proposition 3. With risky cash flows that are privately observed by the agent, the

principal's (recursive) problem can be written as

$$\Pi(v) = \sup_{K,T,W} \mathbb{E} [T(v, p) - cK(v) + \delta_P \Pi(W(v, p))] \quad (13)$$

subject to for all (v, p) and \hat{p} such that $T(v, \hat{p}) < K(v)p$

$$T(v, p) \in [0, K(v)p] \quad (14)$$

$$K(v) \in [0, \bar{k}] \quad (15)$$

$$\delta_A W(v, p) \geq T(v, p) \quad (16)$$

$$\mathbb{E} [K(v)p + \delta_A W(v, p) - T(v, p)] = v \quad (17)$$

$$-T(v, p) + \delta_A W(v, p) \geq -T(v, \hat{p}) + \delta_A W(v, \hat{p}) \quad (18)$$

where (14) is the liquidity constraint, (15) is the agent's capacity constraint, (16) is the participation constraint, which ensures the agent will not abscond, (17) is the promise-keeping constraint, and (18) is the incentive compatibility constraint that ensures the agent will report cash flows (and make the associated transfers) truthfully.

Consider a candidate optimal solution for $K(v)$ and $T(v, p)$. Suppose first that the contract satisfies $\mathbb{E} [\delta_A W(v, p) - T(v, p)] = 0$. Since $\delta_A W(v, p) \geq T(v, p) \quad \forall p$ this implies $\delta_A W(v, p) = T(v, p) \quad \forall p$. Thus, if one interprets $T(v, p)$ as deposits the agent makes into an account, these deposits must deliver a rate of return $1/\delta_A$. Also, since $K(v) = \frac{v}{\bar{p}}$, one can think of \bar{p} as the price per unit of capital. Thus, if $\mathbb{E} [\delta_A W(v, p) - T(v, p)] = 0$, the optimal arrangement can be implemented with these two instruments.

Next suppose $\mathbb{E} [\delta_A W(v, p) - T(v, p)] > 0$. One can still think of \bar{p} as the beginning of period price of a unit of capital. But we now need to understand the endowment to mean either a cash amount of v or some capital $K(v) + \text{cash}$, where for each unit of capital, the cash is reduced at a rate \bar{p} . Let $I(v, p) = \delta_A W(v, p) - T(v, p)$ and $I(v) = \mathbb{E} [I(v, p)]$. Suppose that $I(v, p) = I(v)$ (i.e., it is independent of p). Then (18) must hold with equality for all (p, \hat{p}) and $(I(v) + T(v, p))/\delta_A = W(v, p)$. We can then interpret the arrangement as having an intra-period cash balance of $I(v)$, which also delivers the same inter-period return of $1/\delta_A$.

What remains to be shown is that $I(v, p) = I(v)$. To establish this, first note that $\Pi(v)$ must be weakly concave in v . Otherwise, the principal could always offer to mix over W in a way such that all the constraints are still satisfied.³² Now, take any $p' > p$ and suppose

³²To allow for this, let x be a random variable and let $W(v, p, x)$ denote the continuation value of the agent as it depends on the principal's randomization so that $W(v, p) = \mathbb{E}_x [W(v, p, x)]$. Because we allow for the continuation value to be stochastic, δ_A^{-1} must be interpreted as the expected return on the savings account.

$I(v, p') > I(v, p)$, since $W(v, p) = (I(v, p) + T(v, p)) / \delta_A$ letting $I(v, p)$ increase in p simply increases the variance of $W(v, p)$, which cannot lead to an improvement since $\Pi(\cdot)$ is weakly concave. Next, note that $I(v, p)$ cannot be strictly decreasing in p . Otherwise, the agent would have an incentive to under report and consume the difference (i.e., (18) would be violated). Hence, the principal can do no better than to set $I(v, p) = I(v)$ for all p . \square