

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

William Fuchs[†]

Brett Green[‡]

David I. Levine[§]

June 23, 2018

Abstract

A large literature examines demand-side barriers to product adoption. In this paper, we examine supply-side barriers in a setting with limited contract enforcement. We model the relationship between a distributor and its credit-constrained vendors. We show that the optimal self-enforcing arrangement can be implemented by providing vendors with a line of credit and the option to buy additional units at a fixed price. Moreover, the structure of this arrangement 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 model-implied optimal arrangement increased distribution significantly compared to a standard contract. However, growth was lower than predicted by the model due to (i) Vendors being unwilling to extend credit to customers and (ii) the lack of a reliable savings technology. We discuss recent technological innovations that help to overcome both of these challenges.

*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 Chicago Booth, 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 (Rehfuess, 2006). 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 (Flavin and Aeck, 2004). 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 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.¹ Local vendors face similar issues to the ones faced by consumers. Small shopkeepers lack the capital necessary to purchase, transport, and store inventory.² 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.³ 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 a line of trade credit and an option to buy inventory in the future at a fixed price. The credit line helps overcome the liquidity constraint and the opportunity to grow induces the vendor to reinvest. Moreover, the structure of this arrangement is optimal for both profit-maximizing firms and non-profit organizations with a limited budget whose objective is to maximize product distribution.

¹In 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.

²See e.g., de Mel et al. (2008).

³The 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.

An interesting feature of the arrangement is that it involves “starting small.” That is, the limit on the credit line 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 credit line and demonstrate that it increases with profitability and decreases with the relative impatience of vendors. These findings suggest that starting small is particularly relevant for distribution in the developing world.

Our model-implied optimal arrangement is similar to the “micro-consignment model” (MCM) in the development literature (Van Kirk, 2010; Smith, 2010). The MCM attempts to alleviate credit constraints of local entrepreneurs while simultaneously providing them with access to a distributor. This model has been used for products ranging from eye-glasses to solar lights. For example, Solar Sister is a non-profit organization that uses micro-consignment (Misra, 2011). They provide women with a start-up kit of inventory of portable solar lamps and mobile phone chargers, training and marketing materials, and ledgers to keep sales records. Solar Sister entrepreneurs sell on consignment: the women do not have to pay for the inventory until they sell it and they earn a commission on each sale. Two common themes among organizations that adopt the MCM are that (i) the amount of start-up inventory is relatively small, and (ii) the arrangement is designed so that successful entrepreneurs grow over time until they reach the efficient scale.

Micro-franchising is a closely related approach. Living Goods is a non-profit organization that has adopted the micro-franchising model (Keeton, 2017). They recruit local entrepreneurs and provide them with training, access to a distributor and a loan for working capital. Perhaps the key distinction between the two approaches is that micro-consignment gives entrepreneurs the right to return unsold inventory, whereas micro-franchisors are the residual claimants on unsold inventory and are obliged to repay their loan regardless of sales performance. To the extent that the entrepreneur faces default costs, micro-consignment provides downside protection to an entrepreneur who is unable to sell inventory, whereas micro-franchising provides higher powered incentives.

We calibrate the model to identify the key parameters of the model-implied optimal arrangement and develop several hypotheses regarding the impact of different contractual features. We then design and run a field experiment in order to test our hypotheses and evaluate the performance of the model-implied optimal arrangement. Through a partnership with a large non-profit organization, 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 credit constraints are indeed an important barrier. Vendors that are provided with a credit line had significantly higher sales 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. Most notably, combining the model-implied optimal arrangement with a loaner light led to roughly a ten-fold increase in sales relative to a standard fixed price contract. On the other hand, the distinction between micro-consignment and micro-franchising was not particularly important: providing the vendors with the right to return unsold inventory had little effect on overall performance.

Perhaps not surprisingly, 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 difficulties in saving sales revenue between re-supply meetings (a period of several weeks) as well as a general unwillingness to extend credit to customers.

Interestingly, there are several recent developments in financial technology have facilitated ways to overcome these difficulties as well as the frictions that we address with our incentive scheme. Most notably, PAYGO solar products, which have an embedded lockout technology (i.e., a “kill switch”) and use a contract that features frequent payments made via mobile money directly from the customer to the producer. Using mobile money ensures that vendors are not required to handle (and save) cash, while the lockout technology switch makes the product worthless if the vendor absconds with it or if the customer stops making payments. PAYGO providers have higher production and servicing costs because each system requires additional hardware and software in order to embed the lockout technology. Nevertheless, PAYGO solar products have emerged as the fastest growing segment in the off-grid solar industry.⁴ Their success suggests that the frictions we study in the model and identify in our experiment are of first-order importance for understanding the product adoption puzzle.

Related Literature

One of our goals in this paper is to bridge the gap between the development economics literature and the literature on dynamic contracting. 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.⁵ For example, Karaivanov and Townsend (2013) use survey data to evaluate which models best describe

⁴The 2018 Global Off-Grid Solar Market Trends Report. Date Accessed: May 9, 2018.

⁵Another 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).

the patterns of investment and consumption.⁶ 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 experiment.

Our theoretical results are related to the literature on optimal contracting with investment.⁷ 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. 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 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), Ray (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. With no risk and equal discount rates, Ray (2002) provides a unifying general analysis demonstrating convergence to the efficient outcome obtains in a large class of related problems. 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.⁸ In our implementation, the size of the credit line is fixed over time and the adjustment takes place via quantities.

⁶See also Townsend and Urzua (2009), Kaboski and Townsend (2011), Karaivanov and Townsend (2013), and Chiappori et al. (2014) among others.

⁷See, for example, Quadrini (2004), Albuquerque and Hopenhayn (2004), Clementi and Hopenhayn (2006), DeMarzo and Fishman (2007b), Demarzo et al. (2012).

⁸See also Cuñat (2007).

Our work relates to a large literature on vertical relations in industrial organization. This literature predominantly focuses on contractual relations between two parties at successive stages in the vertical chain of production, where issues of double marginalization and externalities across multiple downstream firms arise (see Katz, 1989 for an overview). We abstract from these issues and focus instead on addressing liquidity constraints of the downstream vendor and the moral hazard problem that arises from providing credit.

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.⁹

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. This 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. 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 distribute more bed nets throughout the community.¹⁰ By charging household’s a price p ,

⁹Note that p may represent household’s value for the good, or the amount that they are able to pay, which might be less than their value due to credit constraints. Importantly, p need not correspond to the true social value of the good.

¹⁰Charging 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). On the other hand, doing so reduces demand and may have negative distributional consequences Dupas and Cohen (2010). We intentionally abstract from these issues by assuming all households have the same value for the good,

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.¹¹

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.¹²

and instead focus on how to best provide incentives to vendors in order to maximize distribution subject to a budget constraint.

¹¹Naturally, if the potential savings on distribution costs are sufficiently small, the NGO will not find overcoming the agency costs to be worthwhile.

¹²Since 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 a form of trade credit. The promise of repeated business then provides the vendor incentives to repay over time. 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, the vendor 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.

When considering the number of units to allocate initially to the agent, k_0 , the NGO

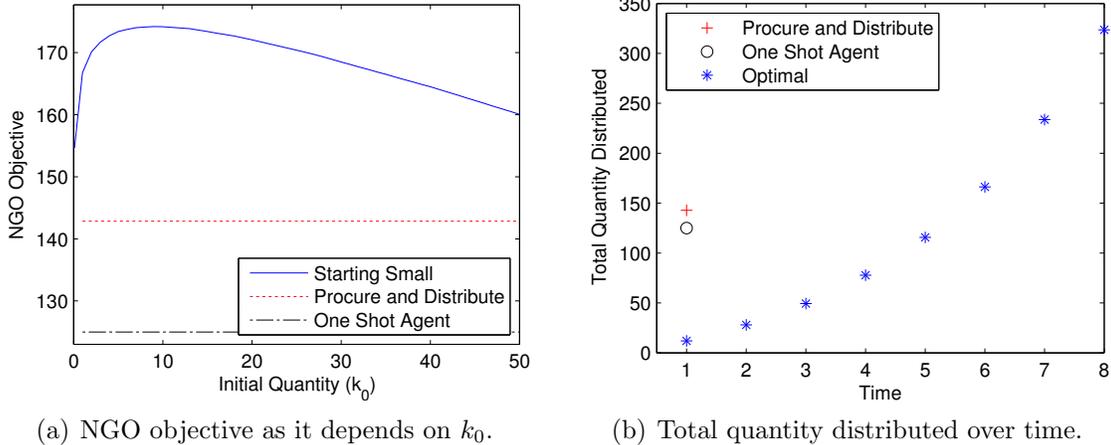


Figure 1: Illustrates the advantage of starting small in the relationship and building up over time. The figures use the parameters $p = 4, c = 8, d = 3, d_A = 0.5, B = 10^3, \delta = 0.99, \delta_A = 0.75$.

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. In other words, the NGO must decide whether to distribute to fewer households more quickly or to a larger share of the population but more slowly. 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 distribution in the first several periods 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.

3 The Model

In this section, we present a formal model and demonstrate the optimality of the arrangement described above. We relabel the NGO as 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 agent’s 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 zero.

The principal and the agent can interact repeatedly over an infinite horizon, $t = 0, 1, 2, \dots$. The agent has the capacity to distribute up to \bar{k} units of the good per period. The goods can be sold in the local market to households. We assume there are arbitrarily many potential households in the economy.¹³ Each household has unit demand and is willing to pay \bar{p} for the good.¹⁴

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 $\bar{p}k_t$.¹⁵
- 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

¹³This assumption is convenient to conserve on the number of state variables required in our recursive formulation of the problem. However, it is not difficult to extend our results to a setting with a finite number of households.

¹⁴Our model implicitly assumes that vendors are local monopolies, which seems reasonable given our observation that distribution channels for these goods do not exist and our hypothesis that it is due to a variety of economic frictions.

¹⁵In 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.¹⁶ 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.¹⁷ Finally, (5) is a standard promise-keeping constraint which requires the principal indeed delivers v in utility to the agent.¹⁸

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

¹⁶This constraint eliminates the possibility of dynamic trading gains due to the agent’s relative impatience (Opp and Zhu, 2015).

¹⁷We 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).

¹⁸We solve for the optimal arrangement assuming the agent does not have access to a private savings technology. However, since the agent is risk neutral, there is no incentive to privately save and thus our optimal arrangement is robust to introducing a private savings technology. 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.

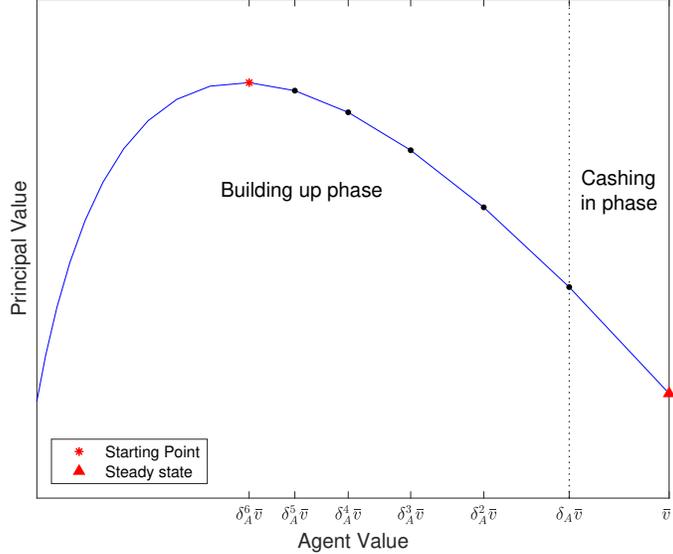


Figure 2: This figure illustrates the principal’s value function and the two phases of the arrangement for parameters such that the time-to-capacity is 6 periods (i.e., $N^* = 6$). The initial phase begins at the red asterisk and moves rightward along the black dots until the agent’s value reaches $\delta_A \bar{v}$, at which point investment is efficient and above which the agent begins to consume. The parameters used to generate this figure can be found in Table 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 trade credit line with limit $k_0^* < \bar{k}$, and sets a fixed price, $p_A^* = \delta_A \bar{p}$ for both repayment and purchasing additional units.*

Notice that while the agent’s liquidity constraint binds, the solution in Lemma 1 specifies that the agent transfer all of her revenue back to the principal. For example, in the first period, she transfers an amount $\bar{p}k_0^*$ to the principal. Of this, a fraction δ_A goes to repaying the credit line and a fraction $(1 - \delta_A)$ goes toward purchasing additional units. In the next

period, the agent will use the credit line to again borrow k_0^* units and will also receive the pre-purchased $(\delta_A^{-1} - 1)k_0^*$. Thus, in the second period the agent will have a total of $\delta_A^{-1}k_0^*$ units of inventory. The growth of inventory and continuation value persists at a rate of $(\delta_A^{-1} - 1)$ until the liquidity constraint no longer binds (i.e. for $v \in (\delta_A \bar{v}, \bar{v}]$) after which the agent makes a transfer of $\delta_A \bar{v}$ in each period. Here again, one can decompose this payment into repayment of the credit line, $\delta_A \bar{p} k_0^*$, and payment for units in the next period $\delta_A(\bar{v} - \bar{p} k_0^*)$. In the next period, the vendor will have access to a total of $k_0^* + \frac{\delta_A(\bar{v} - \bar{p} k_0^*)}{\delta_A \bar{p}} = \bar{k}$ units and thus will be operating at capacity.

Initially, the dynamics of the optimal arrangement resemble the example in Section 2. A profit-maximizing firm provides the agent with a small allocation of inventory 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.

It is worthwhile to note that 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 limit on the credit line 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 credit line limit and demonstrate several comparative statics.

In order to do so, let $N^* \equiv \min\{t : k_t = \bar{k}\}$, which denotes the number of periods 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 limit on the credit line 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 credit limit is always strictly less than \bar{k} . Intuitively, the credit line is designed to relax the liquidity constraints of the agent, but the revenues from first-period sales also correspond to the agent's rent. 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 credit limit 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. This can be handled by recasting the NGO's problem in

terms of its dual. 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 a credit line, 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 portfolio 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 at 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

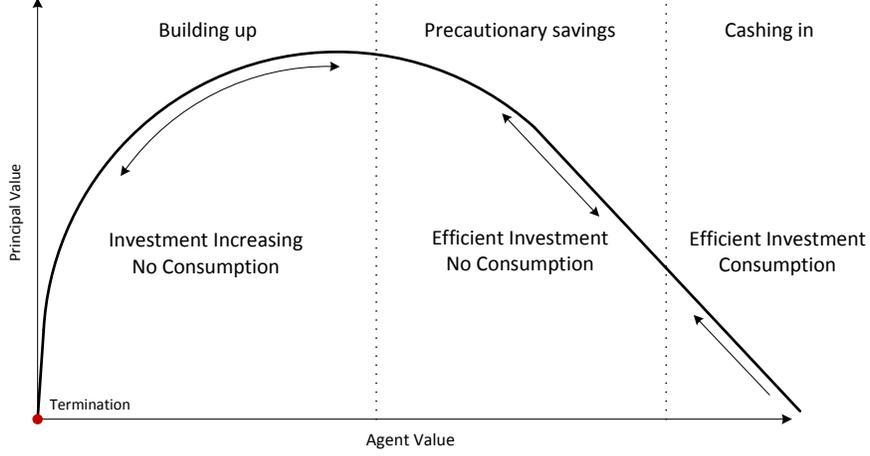


Figure 3: Illustration of the optimal arrangement with risky cash flows when it includes a precautionary savings phase (e.g., when δ_A is sufficiently close to δ_P).

agent and on the distribution of risky cash flows. Regarding the long run outcome, there 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 a credit line (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 credit line limit.*

Thus, by varying the limit on the credit line, we can evaluate the extent to which the liquidity constraints are relevant.

4.2 Risk Aversion and the Right to Return

In the previous section, we argued that the model-implied 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 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.*

As discussed in Section 1, the primary difference between two commonly used arrangements in practice, micro-consignment and micro-franchising, is that the former endows the right to return unsold inventory whereas the latter does not.

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, which took place in rural Uganda. The purpose of the experiment was (1) to evaluate the overall performance of the model-implied optimal arrangement, and (2) to test the hypothesis formulated in the previous section.

5.1 The Setting

The experiment took place in rural villages of central Uganda located within several hours of Kampala. Rural Uganda is very poor. At the time the experiment was conducted, the average monthly household income in rural central Uganda was roughly \$250. The main sources of earnings were subsistence farming (42.4%), wage employment (24.2%), and non-agricultural enterprise (23.9%).¹⁹ Rural Ugandans also generally do not have access to formal financial institutions. For instance, only 9% of women and 16% of men use a bank account.²⁰

We used solar lights as the product with which to conduct the experiment. We selected solar lights because (1) they are relatively small and easy to transport, thereby easing the logistical burden of transporting inventory, (2) they have an immediate financial benefit to households due to the reduction in kerosene expenditures and mobile charging costs, and (3) the penetration of solar lamps in rural Uganda was very low at the time of the experiment. In fact, most of the women in the recruitment meetings as well as onlookers from the villages had never seen a solar lamp prior to our demonstration.

¹⁹Source: *Uganda National Household Survey 2012/13*. Uganda Bureau of Statistics. <http://ubos.org/unda/index.php/catalog/2>. Date Accessed: May 9, 2018.

²⁰Source: *Uganda Demographic and Health Survey 2016*. Uganda Bureau of Statistics. <https://www.dhsprogram.com/pubs/pdf/FR333/FR333.pdf>. Date Accessed: May 9, 2018.

5.2 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. The most popular model was the Firefly Mobile, produced and distributed by Barefoot Power. The FireFly Mobile easily provides enough light for reading and studying as well as completing basic household chores. It is also capable of charging mobile phones.²¹ The FireFly Mobile was packaged in a small box which included a solar panel, a lamp with a built-in battery, necessary cables, and multiple adapters to enable charging for the most popular mobile phones in the region.

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 the 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

²¹Mobile charging capability was a highly attractive feature as many households spent a significant amount of time and resources traveling to village centers in order to charge their mobile devices.



Figure 4: A picture of the FireFly Mobile produced by Barefoot Power. The package included a solar panel, a lamp with a built-in battery, and adapters for mobile charging.

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 line of trade credit and the other half of the zones received an arrangement that did not include trade-credit line. We offered

all vendors the right to purchase lights at a fixed price during the first meeting or at any point in the future. 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 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. We also emphasized to vendors that the arrangement could continue as long as they were successful and that by reinvesting they could grow their business eventually generate substantial income. At the end of the recruitment meeting we took initial orders. Vendors who were not allocated credit were required to pay cash for their initial orders upon delivery, while vendors treated with a credit line paid cash only for units in excess of the credit limit.

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.

5.3 Model-Implied Inputs

We parameterize the model in order to guide the key parameters of the arrangement offered to vendors. The model parametrization is specified in Table 1.

| Parameter | Value |
|--|-------|
| Retail price (p) | \$26 |
| Wholesale price (c) | \$19 |
| Vendor distribution cost per unit (d_A) | \$1.8 |
| Monthly vendor capacity (\bar{k}) | 10 |
| Monthly vendor discount factor (δ_A) | 0.855 |
| Monthly principal discount factor (δ_P) | 0.983 |

Table 1: Model parameters

Retail and wholesale costs were obtained from Barefoot. The vendor distribution cost

per light is consistent with an hourly wage of \$0.30 and the assumption that it would take vendors an average of 6 hours to complete the sale of each light (which includes sales and marketing, placing and collecting orders, transportation time, and payment collection).²² We chose $\bar{k} = 10$ based on the assumption that vendors would spend about one third of their working hours on solar product sales and distribution. By selling 10 lights per month, we estimated that vendors could increase their monthly income by approximately \$40 (80%). The monthly discount factor for the agent is consistent with the median monthly discount rate found in Guiteras et al. (2014) of 17%. The discount factor chosen for the principal corresponds to an annual interest rate of 23%, which is in line with the lending interest rate in Uganda at the time of the experiment.²³

| Parameter | Value |
|-------------------------------|------------|
| Credit line limit (k_0^*) | 3.9 lights |
| Vendor price (p_A^*) | \$20.7 |
| Time to Capacity (N^*) | 6 months |

Table 2: Model-implied optimal arrangement

As illustrated in Proposition 1, the two key parameters of the optimal arrangement are the limit on the trade credit line (k_0^*) and the price charged to the vendor for additional units (p_A^*). The model-implied optimal values for these parameters, based on the inputs in Table 1, are displayed in Table 2. Correspondingly, the vendors that were treated with an arrangement that included the credit line had a limit of 4 lights, which was to be repaid at a rate of 50,000 Ugandan shillings (UGX) per light (roughly \$20). All recruited vendors were allowed to purchase any additional lights at a price of 50,000 UGX.

5.4 Summary of Findings

The 67 vendors who were treated with the model-implied optimal arrangement, which included a credit line sold an average of 6.6 lights (SD = 7.3, median = 4).²⁴ The 62 vendors not treated with a credit line had average sales of 1.8 lights (SD = 6.1, median = 0). The difference is highly statistically significant ($p < 0.01$) with both a two-sided t-test and with

²²The wage estimate derives from the median monthly earnings for females in paid employment according to the Uganda National Panel Survey 2011-12 and assuming a 40 hour work week.

²³The lending interest rate in Uganda is the rate charged by banks on loans to its prime customers. See https://www.bou.or.ug/bou/rates_statistics/statistics.html.

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.

a nonparametric sign-rank test. The difference becomes even more significant when we omit a single outlier in the control group who sold 44 lights, as we do in the remainder of our analysis. The results discussed above are summarized in Figure 5.

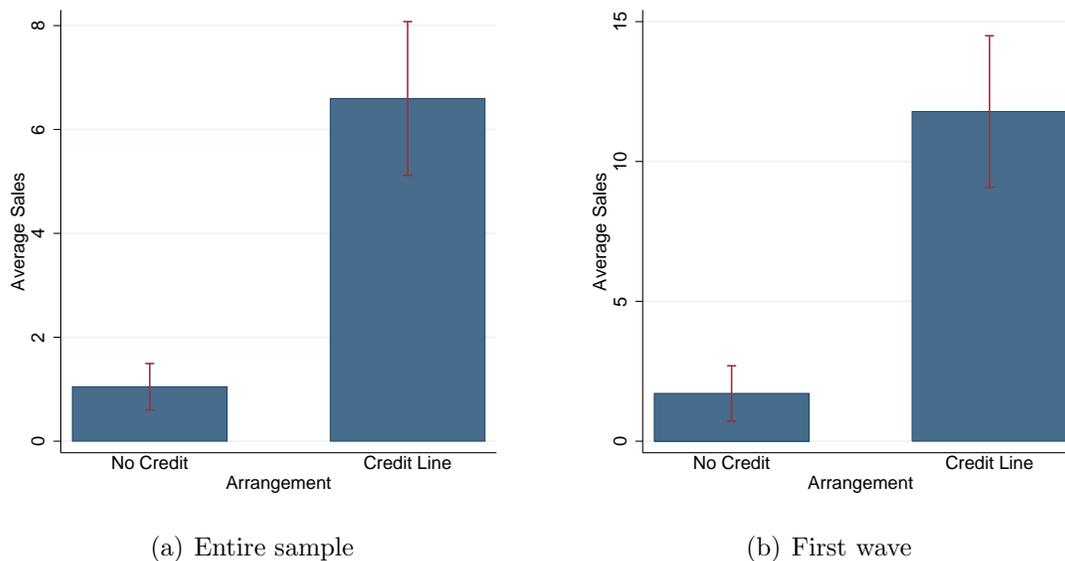


Figure 5: Credit Line versus No Credit: This graph shows the effect of including a line of trade credit in the arrangement for the entire sample (left-panel) and for the first-wave only (right panel). The blue bar is the average of total sales across all vendors who received an offer with the corresponding characteristic. The range between the red capped line indicates the 95% confidence interval.

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 $\log(1 + \text{sales})$ 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 in the treatment group stated they would not have sold as many lights without the line of credit. Thus, while the sample is small, our study provides strong evidence that liquidity constraints are important and providing trade credit can increase sales. Although sales were much lower in the second wave of four branches, the credit line remains important.²⁵

²⁵Our qualitative interviews partially explain why sales were so much lower in the second wave. Vendors

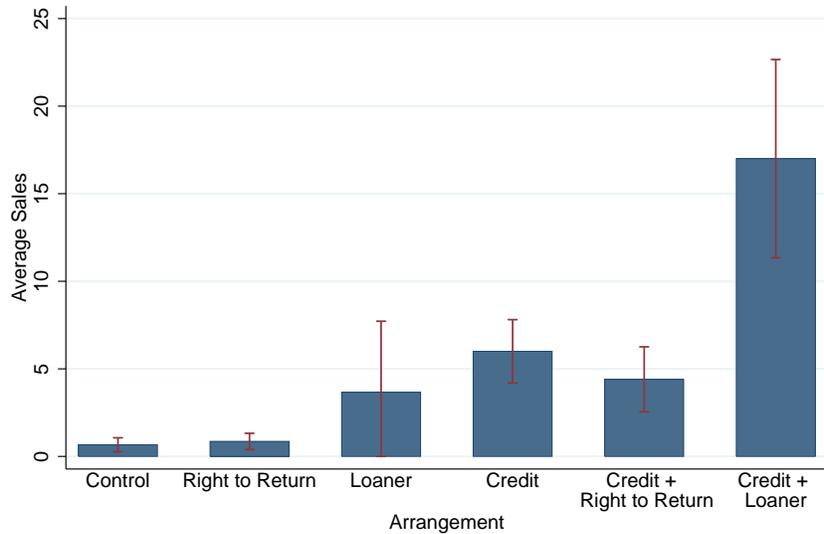


Figure 6: Average Sales by Arrangement. This graph shows the average total sales of vendors by the characteristics of the arrangement they were offered. The blue bar is the average of total sales across all vendors who received an offer with the corresponding characteristics. The range between the red capped line indicates the 95% confidence interval.

To see if vendor risk aversion 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 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.

Another possibility is that the loaner light helped to relax liquidity constraints of households. That is, during the trial period, household’s kerosene expenditures are lower which makes it easier for them to save money to purchase the light. Yet, in surveys, vendors re-

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.

ported lending the loaner for a period of one to three days, which would limit the scope for this channel.

Figure 6 summarizes our main results from the field experiment. Providing agents with either a credit line or a loaner light had a 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. Our regression estimates suggest that the interaction between the credit line and the loaner light leads to an additional 20% increase in sales (See Column (6) of Table A.2), but the coefficient is not statistically significant.

5.5 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 the inverse of the agent's discount factor 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, four months after initial recruitment the average sales was less than 1.5 lights per month. One possibility is market saturation. That is, perhaps vendors exhausted the set of households with high valuations and had a difficult time selling additional lights for a profit. While there is surely some truth to this hypothesis, we believe this explanation is incomplete because the catchment area of each vendor had at least several hundred households without access to the grid and with little to no penetration of solar products. As we discuss below, we believe there are several other key factors that explain why vendors' sales growth was lower than predicted by the model.

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.²⁷ There are several RCTs that have studied ways in which to facilitate the commitment to savings.²⁸ 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 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.

Other Factors Contributing to a Lack of Growth. While 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

²⁷See Bobonis (2009), Bobonis et al. (2013) and references therein regarding the allocation of resources within households and domestic violence.

²⁸See Ashraf et al. (2006) and Basu (2014).

(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.

Is Technology the Answer?

When vendors are both unwilling to offer consumer credit and cannot easily save, it may be more 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 not repossessing products when households default.

Recent technological innovations have allowed producers to combine payments via mobile money with an embedded lockout technology on their solar products. The lockout technology makes the product worthless to the customers if they do not make payments and worthless to vendors if they abscond with inventory, thereby increasing customer repayment rates and reducing vendors’ ability to extract rents. A typical retail contract involves a downpayment of roughly 10-20% with daily or weekly payments made over a 1-2 year term. When a customer makes a payment, the system will be “unlocked” for a fixed amount of time until another payment is due at which point the system locks (i.e., becomes unusable) unless the payment is made. Once the household has made all of the specified payments, the system is permanently unlocked. This emerging product and loan combination is known as “pay-as-you-go” (PAYGO) solar and has become one of the fastest growing sectors of solar home systems (SHS) in the developing world.²⁹

M-Kopa is one of the largest PAYGO solar providers in East Africa. They sell a SHS that comes with a solar panel, three small lights, a phone charger, and a radio. The PAYGO contract involves a downpayment of \$35 deposit, followed by 365 [daily] payments of 45 cents paid by mobile money. A comparable system would retail for a cash price of roughly \$75. M-Kopa’s systems are more expensive to produce and service because each system contains a GSM chip and 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 600,000 systems in East Africa

²⁹See <https://www.lightingglobal.org/2018-global-off-grid-solar-market-trends-report/>. Date Accessed: April 20, 2018.

and is growing at a rate of 500 new systems per day.³⁰ More generally, the recent success of PAYGO solar products suggests the severity of the frictions we study in this paper are of first order importance.

Despite the innovative promise of PAYGO-like contracts, most household investments are not obviously compatible with a lockout technology (e.g., fertilizer, school fees, efficient cook stoves) because installation is either infeasible or prohibitively expensive. In this regard, mobile and solar products are unique because they operate on electricity and therefore can be easily and remotely disconnected, whereas most household investments that lack widespread adoption are not lockout compatible. One therefore might be tempted to conclude that the broader scope of the innovation is somewhat limited.

On the other hand, several firms operating in PAYGO solar have started experimenting with re-collateralizing SHS (and thus re-using the lockout technology). For instance, Fenix International recently started offering follow-up loans to its PAYGO solar customers who had exhibited a strong repayment history. The loan effectively allows its customer to borrow against their equity in their SHS, which is paid back over time and where the SHS will once again lock if payments are not made. While re-collateralizing SHS is not yet a common practice in the industry, there is clearly potential for this innovation to “unlock” households’ access to credit.

6 Conclusion

Most of the literature has focused on addressing demand-side barriers to product adoption. In this paper, we developed a theory to address supply-side barriers. We model the relationship between a producer and its vendors, which features credit constraints and limited contract enforcement. The optimal self-enforcing arrangement can be implemented by providing vendors with a line of credit and the option to buy additional units at a fixed price. Moreover, the structure of the arrangement is optimal both for profit-maximizing firms and for non-profit organizations.

We calibrate the model and conduct a field experiment to test the model-implied optimal arrangement. The evidence from the field clearly indicates that vendor liquidity constraints and consumer uncertainty about product quality are important factors. The combination of a credit line and a loaner light increases sales nearly tenfold over a standard contract. Sustained growth, however, was lower than predicted by the model. Survey evidence suggests that the lack of sustained growth was largely due to vendors’ inability to save even for the

³⁰<http://solar.m-kopa.com/about/company-overview>, date accessed: April 20, 2018.

short periods of time between orders and a failure of the “credit chain” (i.e., vendors were unwilling to offer credit to their customers).

The development of PAYGO solar products, which combines frequent time-payments via mobile money with a lockout technology, has recently emerged to help overcome these obstacles, but it comes at a non-trivial increase in cost and is prohibitively expensive for most products that lack widespread adoption in the developing world. Several promising directions for future work are (1) to investigate the extent to which these obstacles can be overcome using incentive schemes rather than integrating expensive hardware, and (2) to explore re-collateralizing systems with an embedded lockout technology in order to promote households’ access to credit.

References

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

- Clementi, G. L. and H. A. Hopenhayn (2006). A Theory of Financing Constraints and Firm Dynamics. *Quarterly Journal of Economics* 121, 229–265.
- Cole, S., X. Giné, J. Tobacman, P. B. Topalova, R. M. Townsend, and J. I. Vickery (2013). Barriers to Household Risk Management: Evidence from India. *American Economic Journal: Applied Economics* 5(1), 104–135.
- Conley, T. and C. Udry (2001). Social Learning through Networks: The Adoption of New Agricultural Technologies in Ghana. *American Journal of Agricultural Economics* 83(3), 668–673.
- Conley, T. G. and C. R. Udry (2010). Learning About a New Technology : Pineapple. *American Economic Review* 100(1), 35–69.
- Cuñat, V. (2007). Trade credit: Suppliers as debt collectors and insurance providers. *Review of Financial Studies* 20(2), 491–527.
- de Mel, S., D. McKenzie, and C. Woodruff (2008). Returns to Capital in Microenterprises: Evidence from a Field Experiment. *The Quarterly Journal of Economics* 123(4), 1329–1372.
- DeMarzo, P. and M. Fishman (2007a). Agency and Optimal Investment Dynamics. *Review of Financial Studies* 20, 151–188.
- DeMarzo, P. M. and M. J. Fishman (2007b). Optimal long-term financial contracting. *Review of Financial Studies* 20, 2079–2128.
- Demarzo, P. M., M. J. Fishman, Z. He, and N. Wang (2012). Dynamic Agency and the q Theory of Investment. *Journal of Finance* 67(6), 2295–2340.
- Dubois, P., B. Jullien, and T. Magnac (2008). Formal and informal risk sharing in LDCs: Theory and empirical evidence. *Econometrica* 76(4), 679–725.
- Duflo, E., M. Kremer, and J. Robinson (2011). Nudging farmers to use fertilizer: Theory and experimental evidence from Kenya. *American Economic Review* 101(6), 2350–2390.
- Dupas, P. and J. Cohen (2010). Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment. *The Quarterly Journal of Economics* 125, 1–45.
- Feder, G. and R. Slade (1984). The Acquisition of Information and the Adoption of New Technology. *American Journal of Agricultural Economics* 66(3), 312–320.

- Flavin, C. and M. H. Aeck (2004). Energy for Development: The Potential Role of Renewable Energy in Meeting the Millenium Development Goals. Technical report, The Worldwatch Institute.
- Foster, A. D. and M. R. Rosenzweig (1995). Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture. *The Journal of Political Economy* 103(6), 1176–1209.
- Giné, X. and D. Yang (2009). Insurance, credit, and technology adoption: Field experimental evidence from Malawi. *Journal of Development Economics* 89(1), 1–11.
- Guiteras, R., D. I. Levine, T. Polley, and B. Quistorff (2014). Credit Constraints, Present Bias and Investment in Health: Evidence from Micropayments for Clean Water in Dhaka. *Working Paper*.
- Kaboski, J. P. and R. M. Townsend (2011). A Structural Evaluation of a Large-Scale Quasi Experimental Microfinance Initiative. *Econometrica* 79(5), 1357–1406.
- Karaivanov, A. and R. Townsend (2013). Dynamic Financial Constraints: Distinguishing Mechanism Design from Exogenously Incomplete Regimes. *Econometrica* 82(3), 887–959.
- Katz, M. L. (1989). Chapter 11 Vertical contractual relations.
- Keeton, C. (2017). Social innovation for health-care delivery in Africa. *Bulletin of the World Health Organization* 95(4), 241–312.
- Kocherlakota, N. R. (2004). Figuring out the impact of hidden savings on optimal unemployment insurance. *Review of Economic Dynamics*, 541–554.
- Levin, J. (2003). Relational Incentive Contracts. *American Economic Review* 93, 835–857.
- Levine, D. I., T. Beltramo, G. Blalock, and C. Cotterman (2013). What Impedes Efficient Adoption of Products? Evidence from Randomized Variation in Sales Offers for Improved Cookstoves in Uganda. *IRLE working paper*.
- Ligon, E., J. P. Thomas, and T. Worrall (2002). Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies. *The Review of Economic Studies* 69(1), 209–244.
- Mailath, G. J. and L. Samuelson (2006). *Repeated Games and Reputations*. Oxford University Press.

- Misra, N. (2011). Solar Sister: Bringing a market based, gender inclusive, bottom up clean energy revolution to Africa.
- Opp, M. M. and J. Y. Zhu (2015). Impatience Versus Incentives. *Econometrica* 83(5), 1601–1617.
- Quadrini, V. (2004). Investment and liquidation in renegotiation-proof contracts with moral hazard. *Journal of Monetary Economics* 51, 713–751.
- Rampini, A. and S. Viswanathan (2013). Collateral and capital structure. *Journal of Financial Economics* 109(2), 466–492.
- Rampini, A. A. and S. Viswanathan (2010). Collateral, risk management, and the distribution of debt capacity. *Journal of Finance* 65(6), 2293–2322.
- Ray, D. (2002). The time structure of self-enforcing agreements. *Econometrica* 70(2), 547–582.
- Rehfuess, E. (2006). *Fuel for Life: Household Energy and Health*. World Health Organization.
- Smith, B. R. (2010). The MicroConsignment Model Reconsidered: Invention-Led Development by Shifting Risk, Transferring Knowledge, and Scaling Capabilities: Innovations Case Discussion: The MicroConsignment Model. *Innovations* 5(1), 129–135.
- Tarozzi, A., A. Mahajan, B. Blackburn, D. Kopf, L. Krishnan, and J. Yoong (2014). Microloans, insecticide-treated bednets, and malaria: Evidence from a randomized controlled trial in Orissa, India. *American Economic Review* 104(7), 1909–1941.
- Thomas, J. and T. Worrall (1990). Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *Journal of Economic Theory* 51(2), 367–390.
- Thomas, J. and T. Worrall (1994). Foreign Direct Investment and the Risk of Expropriation. *The Review of Economic Studies* 61(1), 81–108.
- Townsend, R. M. (1994). Risk and Insurance in Village India. *Econometrica* 62(3), 539–591.
- Townsend, R. M. and S. S. Urzua (2009). Measuring the Impact of Financial Intermediation: Linking Contract Theory to Econometric Policy Evaluation. *Macroeconomic Dynamics* 13(Suppl S2), 268–316.

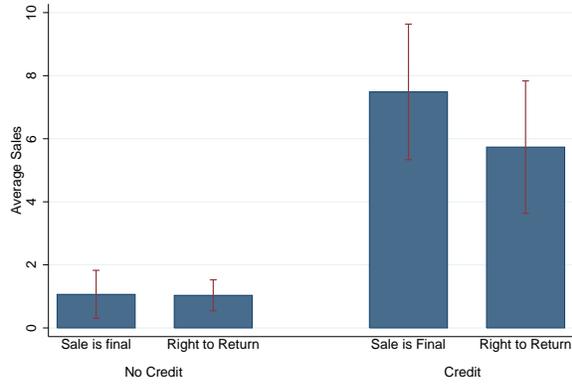
Troya-Martinez, M. (2015). Self-Enforcing Trade Credit. *Working Paper*.

Van Kirk, G. (2010). The MicroConsignment Model: Bridging the Last Mile of Access to Products and Services for the Rural Poor. *Innovations* 5(1), 101–127.

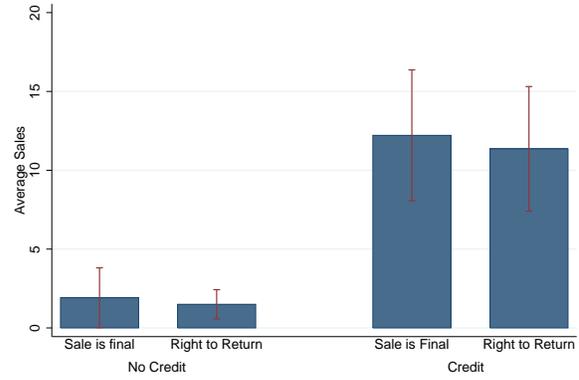
A Tables and Figures

| | | | | |
|----------------------|---------------|--------------------|------------------|--------------|
| 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.

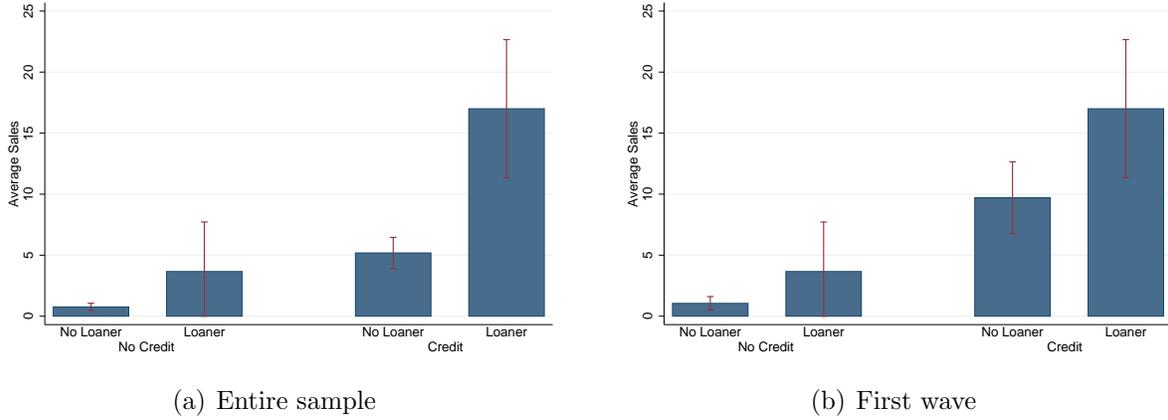


Figure A.2: Addressing Consumer Uncertainty with a Loaner Light. This graph shows the effect of including the right to return unsold inventory in the arrangement for the entire sample (left-panel) and for the first wave only (right panel). The bar denotes the average of total sales across all vendors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

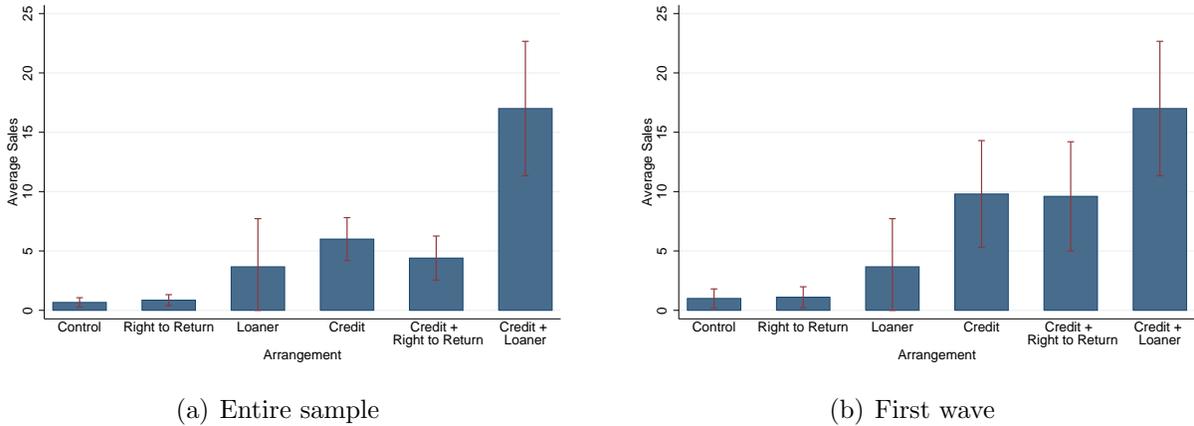


Figure A.3: Average Sales by Offer: This graph shows the average total sales of vendors by the characteristics of the arrangement they were offered for the entire sample (left-panel) and for the first-wave only (right panel). The blue bar is the average of total sales across all vendors who received 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

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

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) | (6) |
|-----------------|--------------------|--------------------|--------------------|--------------------|--------------------|--------------------|
| Credit | 1.086*** (5.11) | 1.078*** (5.03) | 1.087*** (5.13) | 1.078*** (5.05) | 1.054*** (4.90) | 1.046*** (4.77) |
| Loaner | | 1.232** (2.77) | | 1.227** (2.75) | 1.116 (1.58) | 0.643 (0.83) |
| Right to Return | | | -0.122 (-1.45) | -0.100 (-1.15) | -0.102 (-1.16) | -0.117 (-1.32) |
| Loaner x Credit | | | | | 0.210 (0.35) | 0.219 (0.37) |
| First Wave | | | | | | 0.701** (2.88) |
| Constant | 0.500*** (3.88) | 0.361*** (3.90) | 0.561*** (5.53) | 0.411*** (6.33) | 0.425*** (5.29) | 0.203 (1.41) |
| R ² | 0.266 | 0.407 | 0.269 | 0.409 | 0.410 | 0.497 |
| N | 129 | 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

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

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) | (5) |
|-----------------|--------------------|--------------------|--------------------|--------------------|---------------------|
| Credit | 1.541*** (6.47) | 1.537*** (6.73) | 1.543*** (6.40) | 1.539*** (6.66) | 1.655*** (15.83) |
| Loaner | | 0.763 (1.39) | | 0.770 (1.39) | 0.987 (1.14) |
| Right to Return | | | 0.0815 (0.88) | 0.119 (1.92) | 0.129* (2.57) |
| Loaner x Credit | | | | | -0.407 (-0.68) |
| Constant | 0.798** (3.77) | 0.584** (3.27) | 0.756** (3.50) | 0.521** (3.28) | 0.454 (2.03) |
| R ² | 0.473 | 0.567 | 0.474 | 0.570 | 0.577 |
| N | 53 | 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

$$\Pi_0(v) = \delta_A \bar{v} - c \frac{v}{\bar{p}} + \delta_P \Pi(\bar{v}) \quad (7)$$

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

$$\Pi_N(v) = v\mu \sum_{t=0}^{N-1} \left(\frac{\delta_P}{\delta_A} \right)^t + \delta_P^N \Pi_0 \left(\frac{v}{\delta_A^N} \right) \quad (8)$$

We have already demonstrated the base case (i.e., $N = 0$) in (7). Suppose that (8) 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 (8) 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 (8) can be written as

$$\Pi_N(v) = v\mu \left(\frac{\gamma^N - 1}{\gamma - 1} \right) - \gamma^N \frac{c}{p} v + \delta_P^N (\delta_A \bar{v} + \delta_P \Pi(\bar{v})), \quad (9)$$

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

$$\Pi'_N(v) = \mu \left(\frac{\gamma^N - 1}{\gamma - 1} \right) - (1 - \mu) \gamma^N. \quad (10)$$

When $\gamma = 1$, (10) 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 (7)) and the concavity of the value function (see the proof of Lemma 1). Thus, the optimal credit limit 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 (10) 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 a trade credit line with limit $k_0^* = \delta_A^{N^*}\bar{v}/\bar{p} = \delta_A^{N^*}\bar{k}$, 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

$$\Pi(v) = \sup_{K,T,W} \mathbb{E} [T(v, p) - cK(v) + \delta_P \Pi(W(v, p))] \quad (11)$$

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

$$T(v, p) \in [0, K(v)p] \quad (12)$$

$$K(v) \in [0, \bar{k}] \quad (13)$$

$$\delta_A W(v, p) \geq T(v, p) \quad (14)$$

$$\mathbb{E} [K(v)p + \delta_A W(v, p) - T(v, p)] = v \quad (15)$$

$$-T(v, p) + \delta_A W(v, p) \geq -T(v, \hat{p}) + \delta_A W(v, \hat{p}) \quad (16)$$

where (12) is the liquidity constraint, (13) is the agent's capacity constraint, (14) is the participation constraint, which ensures the agent will not abscond, (15) is the promise-keeping constraint, and (16) 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 $\delta_A W(v, p) \geq T(v, p) \forall p$ this implies $\delta_A W(v, p) = T(v, p) \forall p$. 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 credit line 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 (16) 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.³¹ Now, take any $p' > p$ and suppose

³¹To 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., (16) would be violated). Hence, the principal can do no better than to set $I(v, p) = I(v)$ for all p . \square