

Working paper

Strategic default in the international coffee market

Arthur Blouin
Rocco Macchiavello

August 2017

When citing this paper, please
use the title and the following
reference number:
F-89305-CCP-1

IGC
International
Growth Centre



DIRECTED BY



FUNDED BY



from the British people

Strategic Default in the International Coffee Market

Arthur Blouin

Rocco Macchiavello

University of Toronto

London School of Economics

August 14, 2017*

This paper studies strategic default on coffee pre-financing agreements. In these common arrangements, mills finance coffee production through loans backed by forward-sales contracts with foreign buyers. We model how strategic default introduces a trade-off between insurance and counterparty risk: relative to indexed contracts, fixed-price contracts insure against price swings but create incentives to default when market conditions change. To test for strategic default, we construct contract-specific measures of unanticipated changes in market conditions by comparing spot prices at maturity with the relevant futures prices at the contracting date. Unanticipated rises in market prices increase defaults on fixed price contracts but not on price-indexed ones. We isolate strategic default by focusing on unanticipated rises at the time of delivery after production decisions are sunk. Estimates suggest that roughly half of the observed defaults are strategic. Strategic defaults are more likely in less valuable relationships which, in turn, tend to sign price-indexed contracts to limit strategic default. A model calibration suggests that strategic default causes 15.8% average losses in output, significant dispersion in the marginal product of capital and sizeable negative externalities on supplying farmers.

Keywords: Strategic Default, Contractual Forms, Counterparty Risk.

JEL Codes: D22, L14, G32, O16.

*We are grateful to David Atkin, Abhijit Banerjee, Sascha Becker, Jordi Blanes-i-Vidal, Lorenzo Casaburi, Marcel Fafchamps, Bob Gibbons, Amit Khandelwal, Claudio Michelacci, Dilip Mookherjee, Sharun Mukand, Andy Newman, Elias Papaioannou, Daniel Paravisini, Adriano Rampini, Debraj Ray, Antoinette Schoar, Catherine Thomas and numerous people working at the institution that has shared the data used in this paper for helpful comments. Matteo Benetton, Niclas Moneke and Juan Morales provided outstanding research assistance. We also thank seminar and workshop participants in ASSA, Bank of Italy, Berkeley, Berkeley ARE, BU, Columbia, Crest, Delhi, EIEF, ESSFM, Gremaq, INSEAD, IPA, Kellogg, Leicester, Manchester, MIT Sloan, Munich, NBER/BREAD, NEUDC, Oslo, Oxford, PSE, UCSD, USC, SSE, Stanford, UCLA, UPF, Warwick. We gratefully acknowledge IGC and IPA financial support.

1. INTRODUCTION

Contractual defaults occur either out of necessity or for strategic reasons. Well-documented examples of *strategic* default include medieval Maghribi agents (Greif (1993)); difficulties in sourcing at the East Indian Company (Kranton and Swamy (2008)) and in modern contract farming schemes (Little and Watts (1994)); and defaults on mortgages with negative equity (Guiso et al. (2013)). Indeed, the possibility of strategic default underpins many theoretical analyses of market frictions.¹ Empirically identifying it and quantifying its consequences, however, remains challenging. The main difficulty is distinguishing whether default occurs because the defaulting party cannot execute the contract, or does not want to. Nevertheless, understanding both the extent and drivers of strategic default could lead to better contract and policy design.² This is particularly so in the context of international transactions and in developing countries where formal contract enforcement is weak, or absent altogether (see, e.g., Antras (2015), Djankov et al. (2003) and Fafchamps (2003)).

This paper develops and implements a method to identify and assess the importance of strategic default empirically. We build upon a critical insight in the theoretical literature: strategic default occurs when market conditions change sufficiently to place a business relationship outside its self-enforcing range (see, Klein (1996), Baker et al. (2002) and Hart (2009)). The test identifies strategic default by studying how contractual defaults respond to large unanticipated changes in market conditions. Of course, large changes in market conditions could increase both revenues and costs, thereby affecting the likelihood of default through multiple channels. So to isolate the strategic motive, we focus on *contract specific* unanticipated changes in market conditions that *increase* revenues after all production costs are sunk. We quantify the importance of strategic default in the coffee industry, not only for production efficiency but also for contract design,

¹See, e.g., Lacker and Weinberg (1989), Hart and Moore (1998), Shleifer and Wolfenzon (2002), Hart (2009) and Ellingsen and Kristiansen (2011). Different contract terms cause changes in incentives to default strategically. Strategic default, then, is a form of moral hazard. It is, however distinct from standard moral hazard in which a costly action must be incentivized under conditions of uncertainty and limited observability (Hölmstrom (1979), Grossman and Hart (1983)).

²Distinguishing the two forms of moral hazard is potentially important. First, they have different welfare implications (strategic default is a transfer, while standard moral hazard reduces surplus directly). Second, they are differently affected by changes in the environment and, therefore, require different remedies. For instance, in the case of commercial transactions, strategic default might require finding alternative partners to trade and will, therefore, be affected by the market structure in ways that effort underprovision is not. Finally, they also have different legal implications.

insurance, and credit availability.

We conduct our analysis on pre-financing agreements in the international coffee market; a context that offers conveniences for empirical design but is also of intrinsic interest.³ In these common arrangements, coffee mills finance their operations through working capital loans backed by forward-sales contracts with foreign buyers. At the beginning of harvest, a lender advances funds to a coffee mill. During harvest, the mill uses the loan to source coffee from farmers and process it. The mill executes the forward sale contract by delivering coffee to the buyer after harvest. The buyer then repays the lender upon receiving the coffee. We use confidential data from a lender specialized in this type of loan. The data include detailed information on a portfolio of 967 loans, extended to 272 coffee mills in 24 countries.

We are interested in three questions: is there evidence of strategic default? How do parties adjust contracts to the possibility of strategic default? How large are the inefficiencies caused by strategic default? We first present a framework that captures the salient features of the contractual arrangement between the coffee mill, the buyer, and the lender.⁴ We have two goals: i) derive testable predictions to detect strategic default and explore its implications for contractual choice; ii) guide a calibration exercise. A risk-averse mill signs a forward-sale contract with a risk-neutral buyer-lender that advances funds to source coffee from farmers. At the time parties contract, spot market prices at the later delivery date are uncertain. If contracts are perfectly enforceable, the risk-averse mill signs a fixed-price forward contract and receives insurance from the buyer-lender. The mill receives funds required to produce the optimal quantity of coffee given that it is perfectly insured against price fluctuations.

When strategic default is a threat, however, a trade-off between insurance and counter-party risk emerges. In deciding whether to deliver the coffee or default, the mill trades-off financial gains against losses in the future relationship with the buyer-lender.⁵ If spot prices at delivery are much higher than anticipated at the time of contracting, the mill will be tempted to default and sell the coffee to a dif-

³Coffee, the most valuable agricultural export for several developing countries is the primary source of livelihood for approximately 25 million farmers worldwide.

⁴In these arrangements, the buyer and the lender are either the same firm or have a valued relationship. We bundle the two together in this section and provide details in the next. Henceforth, for simplicity, we refer to the buyer-lender.

⁵While it is standard practice in the industry to write formal contracts to obtain loans and accompany shipments and payments across borders, those are typically not enforced by courts or international arbitration in case of default. The losses can include moral costs and broader reputation costs associated with default.

ferent buyer for a higher price. In anticipation of this possibility, the buyer-lender is less willing to extend funds, and the mill is credit constrained. Alternatively, the parties could sign an indexed contract in which the price received by the mill tracks the spot market price at delivery (a so called, differential contract). Such a contract allows the mill to commit to delivery at the cost of foregoing insurance. The model yields three main predictions: 1) unexpected price surges *increase* the likelihood of default on fixed price contracts, but not on differential price contracts; 2) conditional on a fixed-price contract, the effect of price surges on default is lower for more valuable relationships; 3) more valuable relationships select into fixed-price contracts. As a result of strategic default, and depending on the value of the relationship with the buyer-lender, mills can either be unconstrained, credit constrained, insurance constrained or both credit and insurance constrained.

We test these three predictions in the data and find ample support. The key challenge to test Prediction 1 is to identify unexpected changes in market conditions. This requires controlling for contracting parties expectations about future prices at the later delivery date. The first advantage of our setting is that prices quoted in the futures markets reveal parties expectations about market conditions. We can, therefore, construct a contract-specific measure of unanticipated changes in market conditions by taking the ratio between the realized spot market price at the time of delivery and the corresponding futures price at the time the contract is signed. This contract level variation allows us to study the effect of unanticipated changes in market conditions on default controlling for mill fixed effects; buyer fixed effects; and season and country-specific seasonality effects.⁶ When the international price of coffee unexpectedly increases by 10% over the duration of the contract, default increases by almost three percentage points in fixed-price contracts but not in differential price contracts.

We show that these defaults are not driven by mill's inability to source coffee from farmers or increased incentives to divert the loan following a rise in sourcing costs. The second advantage of our setting is the stark separation between production and contract execution. Mills process coffee and incur costs during the harvest season. They deliver coffee and execute contracts well after the end of the harvest season. This timing allows us to isolate strategic default by conducting an event study that considers only price increases that occur after the end of harvest, i.e. when the mill's production decisions and sourcing costs are sunk. Defaults are

⁶In fact, two distinct sources of variation help identification: loans signed at the same time have different price surprises because of different lengths; loans with the same length but signed at different times will also have different price surprises. When we isolate each source of variation separately, we find they yield nearly identical estimates.

about 12 percentage points more likely in contracts when a shipment is scheduled to take place in the week after a price increase relative to the week before. A back of the envelope calculation combining estimates from both strategies suggests that 42%-59% of the defaults observed on fixed price contracts are strategic.

Mills that default are punished by the buyer-lender, with the punishment increasing in the severity of default. For example, being nine months late on delivery and repayment or defaulting outright reduces the likelihood of receiving another contract in the future by almost 30%. If a mill is only three months late, the probability of receiving another loan is only 12% lower. The cost of such punishment depends on how valuable the relationship with the buyer-lender is to the mill (Prediction 2). Using a variety of proxies for relationship value (including measures of history, network centrality and third party assessment), we find that mills that have less valuable relationships with both the lender and the buyer drive strategic defaults on fixed price contracts.

Finally, we show that mills with more valuable relationships with both the buyer and the lender are more likely to sign fixed price contracts (Prediction 3). This correlation holds controlling for mill fixed effects; buyer fixed effects; and season and country-specific seasonality effects. Perhaps counter-intuitively, then, strategic default can be detected on the more valuable relationships, the ones that sign fixed-price contracts. However strategic default also imposes larger indirect costs on the relationships that sign differential contracts and remain uninsured. A model calibration aims at quantifying both the direct and indirect costs.

The model predictions are, therefore, strongly supported by the data. The model derives those predictions from a relatively parsimonious set of parameters. Many of these parameters are directly observed in the data or can be calibrated/estimated. The key unobserved parameter is the value the mill places on keeping a good relationship with the buyer-lender. Knowing this parameter would allow us to assess the importance of informal enforcement and perform counterfactuals to quantify the inefficiency caused by strategic default. We take advantage of the model's relative simplicity to "invert" it, and obtain an estimate of the relationship value for each loan.⁷ We find that the value of the relationship is substantial: it is 44% (158%) of the value of the contract for the median (mean) observation in the sample. Furthermore, strategic default causes significant output distortions: the median (mean) mill production would be 19.7% (15.8%) higher if contracts were perfectly enforceable. The estimates suggest that 26% of mills

⁷Specifically, given a vector of loan-specific parameters, we find the relationship's value that best matches the observed interest rate and forward sales contract type for that particular loan.

are unconstrained; 39% of the mills are insurance constrained; and the remaining 35% of mills are credit constrained, many severely so. These distortions translate into a highly dispersed and skewed distribution of the marginal product of capital across mills. In the group of mills that are credit constrained, the marginal product of capital at the median (mean) firm is 8% (20%) higher than the interest rate (which is around 10%).⁸

Related Literature

Our main contribution to the literature is to isolate a specific form of moral hazard and to quantify the output losses that arise from imperfect enforcement, including their indirect effects through endogenous contract choice. This exercise contributes to a number of literatures. From a methodological point of view, the paper is most closely related to the empirical literature on contracts ([Chiappori and Salanié \(2003\)](#)). However, we study these inefficiencies in the context of exports from developing countries. As such the paper also relates to an emerging literature on contracting in environments with weak or non-existent enforcement institutions. Finally, although we provide a test for strategic default in a particular market, the main idea can be fruitfully applied to study strategic default in other contexts.

By studying pre-financing agreements, the paper contributes to a body of work on contractual imperfections and defaults (see, e.g., [Chiappori and Salanie \(2001\)](#) for a seminal contribution in insurance markets). The literature on credit markets has mostly focused on testing for, and distinguishing between, moral hazard and adverse selection. For example, [Karlan and Zinman \(2009\)](#) and [Adams et al. \(2009\)](#) offer experimental and structural analyses respectively that separate moral hazard from adverse selection in the consumer loan market. We focus on isolating strategic default as a specific source of moral hazard. Following the financial crisis, strategic default has been studied with different methodologies in the mortgage market (e.g., [Guiso et al. \(2013\)](#) survey people about strategic default; [Bajari et al. \(2008\)](#) use structural methods; and [Mayer et al. \(2014\)](#) use a diff-in-diff analysis of a mortgage modification program). In contrast to defaults on mortgages, which happen during economic downturns, we test for strategic default by looking at unexpected *increases* in prices that make the borrower better off. This difference greatly facilitates separating the strategic motive from other causes of default. Relative to the consumer credit and mortgage literatures, our focus on

⁸Furthermore, strategic default implies externalities along the supply chain: output losses at the mill level translate into lower prices and, therefore, lower welfare for farmers that supply the mills. Our estimates bound welfare losses for farmers supplying the average mill between 10% and 32%.

international working capital loans to large firms in developing countries requires considering different aspects, most notably the importance of inter-firm business relationships.^{9,10} More broadly, we also contribute to the empirical literature on financial contracting by highlighting the important role of endogenous contractual terms. For instance, in a review of the literature [Roberts and Sufi \(2009\)](#) argue that understanding how the expectation of renegotiation affects ex-ante contractual terms remains an important but underdeveloped area for research.

Within the literature on contracting under imperfect enforcement, [Antras and Foley \(2015\)](#) offer a notable contribution. They show that trade finance terms balance the risk the exporter does not deliver, and the importer does not pay (see also [Kranton and Swamy \(2008\)](#) for a theoretical treatment). As a result, trading relationships can endogenously become a source of capital and affect responses to shocks.¹¹ In a cross-section of contracts, [Banerjee and Duflo \(2000\)](#) also focus on contract selection, demonstrating that reputation plays a prominent role in the Indian software industry. [Macchiavello and Morjaria \(2015b\)](#) document how Kenya flower exporters exerted effort to protect valuable relationships with foreign buyers during a negative supply shock. Unlike their paper, we observe and test for strategic default and focus on how it influences contractual terms and efficiency (neither of which is studied in [Macchiavello and Morjaria \(2015b\)](#)). Two recent papers offer evidence that enforcement problems significantly impair economic output. [Bubb et al. \(2016\)](#) experimentally test for limited enforcement in water transactions between neighboring farmers in rural India and find that limited enforcement causes significant output losses. [Startz \(2017\)](#) finds that welfare in the Nigerian consumer goods import market would be nearly 30% higher in the absence of search and contracting problems.

The test developed in this paper can certainly be adapted to isolate strategic default in other contexts. For instance, classical studies by [Goldberg and Erickson \(1987\)](#) and [Joskow \(1988\)](#) document that price indexation is a common

⁹A related literature studies trade credit contracts in which suppliers extend credit to downstream buyers (see, e.g., [Klapper et al. \(2012\)](#) and [Breza and Liberman \(2017\)](#) for recent contributions and references). In [Burkart and Ellingsen \(2004\)](#) and [Giannetti et al. \(2011\)](#), for instance, trade credit is used to limit loan diversion, a different form of moral hazard. We provide a test to isolate strategic default and quantify the associated inefficiencies.

¹⁰We show that strategic default is large enough to generate credit constraints for a significant proportion of firms in the sample. These results complement [Banerjee and Duflo \(2014\)](#), to date the best direct evidence for credit constraints among (relatively) larger firms. We study firms that are significantly larger and identify a specific source of credit constraint.

¹¹A recent literature has studied trade finance and the effects of credit supply on exports (see, e.g., [Amiti and Weinstein \(2011\)](#), [Manova \(2012\)](#), [Paravisini et al. \(2014\)](#)) with a rather different focus.

feature of contracts in the petroleum coke and coal markets and also argue, without providing a direct test, that it is used to reduce opportunistic behavior. The trade-off between fixed-price and differential contracts is also related to [Rampini and Viswanathan \(2010\)](#). Their model show how collateral constraints introduce a trade-off between financing and risk management. In line with their predictions, we find that more constrained mills do not insure against price risk. [Rampini et al. \(2014\)](#) provide empirical evidence studying airline hedging decisions over fuel price risk and find evidence of limited risk management, particularly among financially constrained airlines.¹² We focus on the implications of the financing-insurance trade-off for strategic default, and both identify the implications for contractual terms, and quantify the associated inefficiencies. In the context of contract farming, strategic default might also alter the trade-off between insurance and credit provision, as suggested in a recent study by [Casaburi and Willis \(2016\)](#).¹³

2. BACKGROUND AND DATA

2.A. *Coffee Washing mills*

When coffee cherries change colour from green to red they are ripe for harvest. Most coffee-growing countries have only one harvest a year. The harvest season typically lasts for three to four months and its timing varies by country depending on latitude, altitude, soil and weather patterns. Coffee cherries must be processed immediately after harvest to obtain parchment coffee. There are essentially two processing methods: the dry method and the wet method. The dry method is directly performed by farmers. The wet method is performed by coffee washing mills, the objects of this study. Relative to the dry method, the wet method requires significant investment in specialized equipment but produces higher and more consistent quality.¹⁴

¹²In the same industry, [Benmelech and Bergman \(2008\)](#) find that airlines are better able to renegotiate their obligations when performance is poor. As noted above, our test identifies strategic defaults when exogenous changes in market conditions make the borrower better off.

¹³Like [Casaburi and Willis \(2016\)](#) we also contribute to the limited empirical literature on inter-linked transactions in developing countries. The literature focuses on different issues and settings with small traders and farmers (see, e.g., [Fafchamps \(2003\)](#)). Various papers study other aspects of the coffee sector. For instance, [Fafchamps and Hill \(2008\)](#) investigate international price transmission to farmers, [de Janvry et al. \(2010\)](#) and [Dragusanu and Nunn \(2014\)](#) look at fair trade, and [Macchiavello and Morjaria \(2015a\)](#) study how competition between mills affects relationships with farmers.

¹⁴After the cherry skin is removed with a machine, beans are sorted by immersion in water then left to ferment to remove the remaining skin. Once fermentation is complete, the coffee is washed in water tanks or in washing machines. The beans are then dried, sometime with the further help of machines. After drying the parchment skin is easily removed in the hulling

Despite having seasonal activities tied to the coffee harvest, coffee washing mills are large firms by developing country standards. In our sample, mills average over 3.5 million dollars a year in sales, hold about 2 million in total assets and receive average working capital loans of \$473,000 (see Table 1).¹⁵ The production function is relatively simple: the quantity of parchment coffee produced is a constant share, ranging between one fifth and one seventh, of the processed coffee cherries. Disbursements to purchase coffee cherries from farmers during harvest are, by far, the largest source of variable costs and account for 60%-70% of the overall costs. Other costs include labour, transport, electricity, marketing and, of course, costs of finance. Large volumes of working capital need to be mobilized over short periods of time in environments characterized by weak legal institutions and significant uncertainty.

2.B. Contractual Practices I: Loans

This paper studies pre-financing agreements, an extremely common source of working capital finance in the coffee sector and other agricultural commodity markets (see, e.g., [Varangis and Lewin \(2006\)](#)). In this type of agreement, working capital is provided before harvest either directly by buyers or by financial institutions (e.g., banks, specialized lenders). In both cases the working capital loan is backed by forward sales contracts stipulated before production takes place.¹⁶ We obtained access to the internal records of an international lender specialized in providing working capital loans to coffee washing mills. The data cover all loans ever disbursed by the lender over a period of twelve years for a total of 967 working capital loans. The mills are located in 24 countries, with Peru, Mexico, Nicaragua, Rwanda and Guatemala accounting for the majority of loans (Table A1).

The mills in the sample mostly supply the coffee specialty market. In this segment, coffee mills supply directly to foreign buyers (either roasters or traders). Figure 1 illustrates the timing of the lending cycle. First, before the harvest begins the buyer and the mill sign a forward sales contract, or a letter of intent, for the delivery of a certain amount of coffee of pre-specified quality at a later date. The lender provides funds to the mill to source coffee from farmers to fulfill the contract

process. All coffee is hulled before export.

¹⁵By comparison, the firms in the sample in [Banerjee and Duflo \(2014\)](#), probably the largest firms for which credit constraints have been rigorously documented, average \$140,000 in sales and \$17,000 in loan amount.

¹⁶Fixed assets invested in the mill are rarely, if ever, used as collateral for working capital loans. These assets are hard to liquidate: they are invested in rural areas and are highly specific. Repossessing collateral is also notoriously difficult in many developing countries.

with a foreign buyer. The mill sources coffee from farmers and then processes it. Once the coffee is delivered, the lender is repaid.¹⁷

The lender is aware of the difficulties involved in extending large working capital loans to seasonal businesses operating in rural areas of developing countries. The lending model is designed to cope with issues of adverse selection and moral hazard. The lender utilizes a comprehensive scoring system to rate loan applications and decide the size of the loan disbursed. Once the loan is approved and disbursed, the mill might divert the loan rather than use it to purchase cherries from farmers (ex-ante moral hazard). To ensure this does not happen, the lender disburses loans progressively through smaller instalments and actively monitors the sourcing of coffee through in-country loan officers. The mill could use the disbursed loan as intended, but still decide not to repay the lender if it is not convenient (ex-post moral hazard). This type of strategic default is the focus of our analysis.

To limit strategic default, the lender provides working capital loans based on sales contracts between the mills and buyers with whom the lender also has a business relationship. The lending model is illustrated in Figure 2. The lender advances funds up-front and is then directly repaid by the buyer once the mill delivers the coffee.

The system used by the lender is typical in the industry.¹⁸ Figure A1 shows that sales contracts are the most commonly used form of collateral for working capital loans. Figure A2 shows that the interest rates charged by the lender are broadly representative of those offered by other lenders in the markets in which the lender operates. Finally, the lender also attracts fairly typical clients. Indeed, as shown in figure A3, the mills that the lender deals with are similar in size, age, average sales price to mills that we have data on from projects in other countries. Furthermore, the loans themselves are representative of the industry.

¹⁷The loan covers a share of the funds required. In our case, the share depends on a scoring system. The scoring systems aggregates continuous sub-scores based on a large number of mill and loan characteristics into discrete categories (B, A, AA). The share extended to the mill varies between 40% and 70% of the value of the contract with the buyer depending on the score.

¹⁸This lending system is similar to invoice discounting in which the receivable is used as collateral for the loan. Bank loans secured by accounts receivable are the primary source of SME financing for working capital in the US and probably an even more important source of finance in countries with weak legal environments (Demirgüç-Kunt and Maksimovic (2001)). Invoice discounting is similar to factoring, a growing source of financing for SMEs around the world (Klapper (2006)). The key distinction is that in factoring a business *sells* its accounts receivable to a third party (called the factor).

2.C. Contractual Practices II: Sale Contracts

Given the lending arrangement, it is important to describe the incentives associated with sales contracts. Coffee is mostly traded through forward sales contracts: around the beginning of the harvest season buyers and sellers agree for coffee to be delivered at a future date, typically after the end of the harvest season. These forward sales contracts take a limited number of standard contractual forms. From the point of view of our research design, the key distinction is between trade at fixed (or outright) price versus trade at a differential (or price to be fixed (PTBF)). High coffee prices volatility (see Figure 3) implies that the two types of contracts have radically different implications.¹⁹

Fixed contracts were the only contractual form before active futures markets came into being. Fixed price contracts provide insurance against price fluctuations but leave parties exposed to counterparty risk. A seller that has sold coffee for a fixed price will be tempted to renege on the contract if spot prices at the time of delivery are much higher than anticipated at the time of contracting. The reverse is true for the buyer.²⁰

With the development of future markets, coffee has increasingly been exchanged on a differential basis. In this type of sale, the seller (buyer) commits to deliver (take) a certain amount of coffee for a price equal to a basis price plus/minus a pre-specified differential. Theoretically, the basis price can be any published price in the industry. In practice, almost all differential contracts are signed against futures markets (i.e., Robusta coffee is traded against the London LIFFE Contract while Arabica coffee, the object of this study, is traded against the New York ICE 'C' Contract). Differential contracts remove counterparty risk but leave parties exposed to price fluctuations. A seller that has sold coffee on a differential basis will not be tempted to renege on the contract if prices suddenly increase, since the contracted price tracks spot market conditions.²¹

¹⁹The two most frequently used contractual forms are those issued by the European Coffee Federation (ECF) and by the Green Coffee Association (GCA) in the United States. The basic conditions of sale are easily covered by stipulating the applicable standard form. Parties fill the standard form with the remaining important details of the individual transaction (quantity, quality, price).

²⁰Fixed price contracts do not completely remove price risk. An importer who buys coffee that has not already been sold (bought long) hopes that the price will stay the same or go up. Importers in rich countries however easily ensure against this risk through hedging. An exporter who sells coffee that has not already been sourced (sold short) hopes that the price stays the same or goes down. Exporters in developing countries lack access to hedging instruments. They limit this risk by entering stable arrangements with producers and timing production and sourcing decisions accordingly.

²¹Trading on a differential basis transforms outright price risk into differential price risk.

In the data, we observe a roughly equal split between fixed and differential contracts (45% and 55% of contracts respectively).²² Nearly 30% of loans are backed by a mix of fixed price and differential price contracts, typically signed with different buyers.²³ In nearly 80% of the contracts shipments are due after the end of the harvest season.

2.D. Bundling Relationships and Contract Default

A key advantage to studying loans is that default is *directly* observed: a loan is in default if it is not repaid (on time). It is standard practice in the industry to write formal contracts to obtain loans and accompany shipments and payments across borders. These contracts, however, are typically not enforced by courts or international arbitration in case of default. In practice, the loss in reputation and future business from the buyer and the lender is the sole deterrent towards strategic default.

The loan is directly repaid by the buyer to the lender after the mills delivers the coffee. Assuming the buyer and the lender acts as a single entity, then, the mill must default on the sale contract with the buyer to default on the loan - and vice-versa (see Figure A4 for an illustration). In other words, the system bundles the relationship with the lender to the relationship with the buyer. This enables high-risk exporters in developing countries to borrow by using their relationships with low-risk buyers in developed importing countries as collateral.

In principle, however, we might expect that the mill defaults on one party but not on the other. From the perspective of the lender (and therefore our data), the case in which the mill defaults on the lender but not on the buyer is indistinguishable from the case in which the mill defaults on both. The opposite case, when the mill defaults on the buyer but not on the lender, is potentially more interesting. The data show that this happens rarely. We have data on the financial transactions of the lender, and can therefore check which party repays

Although differential price risk is inherently lower, it is not zero. The fixed differential specified in the contract cannot perfectly track the evolution of actual market differentials. Differential price risk is therefore relatively stronger in the specialty coffee segment, where differentials are fine tuned to narrower origins with less liquid markets. Still, swings in differential are far rarer and smaller in magnitude than swings in prices.

²²Fair Trade contracts specify differential price above a fixed floor price. For our purpose, those contracts are therefore classified as differential.

²³We abstract away from the possibility that loans are backed by a mixture of fixed and differential contracts in the theoretical model (for simplicity). In the empirical analysis we account for loans backed by a mix of fixed and differential contracts by conducting our analysis both at the contract level (where contracts can only be fixed or differential) and at the loan level (where we examine robustness in the degree of mixing).

the loan. Appendix D shows that the vast majority of loans (around 90%) are directly repaid by the buyer with whom the sale contract is agreed with. This is as expected. For the average buyer, the lender finances 1.9 suppliers. The 5 largest buyers account for about a third of all loans disbursed. The lender is therefore reluctant to accept repayment from a mill that has defaulted on the buyer: to preserve one relationship with a delinquent mill the lender would risk jeopardizing a relationship with a buyer that helps to guarantee repayment from several of its suppliers. Appendix D provides further evidence on this.²⁴ To err on the side of caution, and for the sake of clarity, we also bundle the buyer and the lender into a single entity for both the theoretical and baseline empirical analysis.²⁵

2.E. Data Sources

The lending model, therefore, produces extremely detailed information on both the commercial and financial operations of hundreds of mills and their buyers in several countries. As described, the lending model and the lender loan portfolio are representative of common contractual practices in the industry. Working with the international lender allows us to analyze confidential data about the operations of many mills and many buyers across countries, which would otherwise be extremely hard to collect. The lender shared essentially all their operating data. We use loan application data (which include financial statements and all the information in the construction of the credit scores); actual financial transactions made by the lender (which includes timing, amounts and counterpart for both disbursements and repayments); the terms of all loans and text files of all sales contracts made between buyers and mills for the delivery of coffee. After substantial organization and cleaning we match the data to world price of coffee. Appendix B provides further details.

After putting each source of data together, we end up with a scheduled-shipment level dataset with 6,372 observations. Shipments are sometimes fixed price and sometimes differential price, even within the same loan. We therefore

²⁴Appendix D shows that the few loans that are repaid by the mill also happen at times with large price increases and are backed by fixed contracts signed with buyers that are not so important in the lender’s portfolio. While we do not directly observe default against the buyer, the evidence suggests that a significant share of these loans might be cases of strategic default against the buyer. If that is the case our baseline empirical analysis *underestimates* the extent of strategic default in the market.

²⁵A buyer on a fixed price contract has incentives to reject coffee delivery following an unexpected price decline and default. We find no evidence for this. This likely reflects the fact that most buyer-lender contracts are domestic (i.e., between companies located in the U.S.) and that, relative to the mills, the lender has much stronger relationships with a significantly more concentrated set of buyers.

define a contract to be a set of shipments within a loan with a common price-type. This leaves us with 1,228 contracts for 967 loans. There are some contracts where terms of the agreement are not specified. This is typically the case when a buyer and mill sign a promissory note instead of a contract. Therefore, of the 1,228 contracts, we have shipping information for 967; 434 of which are fixed price and 536 are differential. The remaining 258 contracts are mostly promissory notes, where the shipping details are unknown. Most of the analysis focuses on the contracts, but we also run the key specifications at the loan level as well.

3. THEORETICAL FRAMEWORK

We now present a theoretical framework to guide the empirical analysis. We have two goals. First, we derive a set of qualitative predictions to detect strategic default and explore its implications for contract choice. We test these predictions in Section 4. Second, we use the framework to guide the calibration exercise in Section 5.

3.A. Set Up

i) Players and Timing A risk-averse mill, a risk-neutral buyer and a risk-neutral lender contract for the delivery and financing of coffee. The timing is as in Figure 1; at time $t = 0$ parties contract. At $t = 1$ production takes place. At time $t = 2$ the world coffee price $p_w \in [0, \infty)$ is realized according to a cumulative distribution function $p_w \sim F(p_w)$ with finite expectation \bar{p}_w . Finally, at time $t = 3$ contracts are executed. Let $\mathbf{I}[p_w]$ be an indicator function denoting whether the mill delivers coffee to the buyer *and* repays the loan to the lender when the realized world price is p_w .²⁶

ii) Production One unit of coffee purchased from farmers produces $1/a$ units of output. Coffee purchased from farmers is the sole input. The aggregate supply of coffee to the mill is given by $\omega = \rho q^\eta$, with $\eta, \rho > 0$. The mill's cost of producing q units of output is given by $C(q) = q \times a \times \omega(q)$, i.e., $C(q) = \gamma q^{1+\eta}$ with $\gamma = \rho \times a$.²⁷

iii) Contracts Contracts consist of two parts: a sales contract and a loan.

²⁶Given the lending model described in Section 2 the two decisions are bundled together. Appendix D provides a discussion and supporting evidence.

²⁷For simplicity, we omit additional processing costs. An upward sloping supply captures mills market power in the rural areas which arises, inter alia, due to high transportation costs and the need to process coffee within hours of harvest.

Sales Contracts The sales contract specifies the delivery of q_c units of coffee at price p_c at date $t = 3$. There are two types of contracts: *fixed price* and *differential*. In a fixed price contract the buyer and the mill agree on a fixed unit price, p_c . In a differential contract the price to be paid by the buyer is equal to the price realized in the world market, p_w , plus a differential Δ_c , i.e., $p_c = p_w + \Delta_c$. At the delivery date, the buyer sells the coffee at the prevailing world market price, p_w . At the contracting stage, the participation constraint for the risk-neutral buyer is simply given by expected zero profits. The buyer is willing to accept the contract for q_c units at price p_c provided

$$(1) \quad \int_{p_w} \mathbf{I}[p_w] q_c (p_w - p_c) dF(p_w) \geq 0$$

When the contract is on differential the buyer's participation constraint collapses to $\Delta_c \leq 0$.²⁸

Loan Contracts The mill borrows from the lender the working capital necessary for production. The mill is subject to limited liability, i.e., at all dates and in all states of the world the payoff of the mill must be weakly positive. The mill signs a standard debt contract with the lender in which L denotes the amount borrowed and D the amount the mill commits to repay. The interest rate on the loan, then, is given by $r_c = (D/L) - 1$. Assuming a risk-free interest rate equal to r , the lender's participation constraint is given by

$$(2) \quad L(1 + r) \leq \int_{p_w} \mathbf{I}[p_w] \min\{p_c q_c, D\} dF(p_w)$$

iv) Default and Enforcement Given the lending scheme, to default on the loan, the mill has to default on the sales contract (and vice-versa). After p_w is realized the mill can decide to sell the contracted coffee q_c to an alternative buyer at price p_w and default.²⁹

The mill is in a relationship with both the buyer and the lender. To study the interaction between the informal enforcement in the relationship and the formal contract we collapse the dynamic relationship onto static parameters (see [MacLeod \(2007\)](#)). We denote with $\mathbf{U}^{\mathbf{R}}$ the discounted value of future expected profits when continuing the relationship with the buyer *and* the lender, \mathbf{U}^D the discounted

²⁸A delivery failure imposes no cost on the risk-neutral buyer. Relaxing the assumption does not alter the qualitative predictions.

²⁹In the Appendix we consider a more elaborate and realistic side-selling process that takes full advantage of the empirical definition of default.

value of future expected profits following a default. Let $\mathbf{V} = \mathbf{U}^R - \mathbf{U}^D$ denote the *value of the relationship*. \mathbf{V} is the key parameter in the analysis: it drives the testable predictions on contractual choice and mill's behaviour. Learning about \mathbf{V} is also necessary to perform counterfactuals.³⁰

v) Mill's Payoff The mill borrows $L = C(q)$. Given contracts, the mill's monetary payoff when the international price is equal to p_w is given by $\pi^R(p_w) = \max\{p_c q_c - D, 0\}$ if the mill repays the loan and by $\pi^D(p_w) = p_w q_c$ when the mill defaults and sells the coffee on the spot market at price p_w . Note that $\pi^R(p_w)$ depends on p_w if the contract is a differential one, in which case $p_c = p_w + \Delta_c$. Assuming the mill's utility function is given by $u(\cdot)$, with $u' > 0$ and $u'' \leq 0$, and normalizing \mathbf{U}^D to zero, expected utility is given by

$$(3) \quad \mathbf{E}[\Pi] = \int_{p_w} u(\mathbf{I}[p_w]\pi^R(p_w) + (1 - \mathbf{I}[p_w])\pi^D(p_w))dF(p_w) + \mathbf{I}[p_w]\mathbf{V}$$

A contract is then a N-tuple q_c, p_c, L_c, r_c . The agreed contract maximizes the mill's expected utility subject to the buyer and lender participation constraints 1 and 2.³¹

3.B. First Best

The contractual outcome is illustrated in Figure 4. The case in which contracts are perfectly enforceable is the first best. Formally, this corresponds to the situation in which the mill can commit to repay the loan, i.e., $\mathbf{I}[p_w] = 1$ for all spot price realizations p_w . Intuitively, with enforceable contracts the risk-averse mill receives insurance from the risk neutral buyer-lender. The mill is guaranteed a fixed payoff which is independent of the realized world prize p_w . This is achieved by signing a fixed price contract. The quantity financed and produced is then independent of the value of the relationship \mathbf{V} and is at the first best level, denoted $q_c = q_F^*$. The quantity produced is also larger than what the mill would

³⁰The future value of the relationship \mathbf{V} might depend on spot prices p_w at the time the mill decides whether to default. For notational simplicity we omit this since in the empirical analysis we control for spot prices p_w and test the model's predictions using unexpected deviations from \bar{p}_w .

³¹The assumption that the mill has all the bargaining power at the contracting stage does not affect the qualitative predictions of the model. The assumption allows us to isolate strategic default as the sole cause of output distortions. If the buyer/lender had bargaining power output distortions could arise due to the standard efficiency - rent extraction trade-off. We also abstract from mill's internal funds. Those would also not alter the qualitative predictions of the model and are taken into account in the calibration exercise.

produce under a differential contract, $q_c = q_D^*$. A differential contract would leave the mill exposed to uninsured price risk. By the standard logic, this lowers the mill's desired production.

3.C. Strategic Default: Second Best

When contracts are not enforceable the mill might decide to default. This decision trades-off the short-run gains associated with side-selling and avoiding loan repayment against the loss in relationship value \mathbf{V} . Upon observing realized world prices p_w the mill defaults on the contract if

$$(4) \quad \delta \mathbf{V} \leq u(\pi^D(p_w)) - u(\pi^R(p_w)).$$

Consider first the case of a fixed price contract. The mill defaults if

$$(5) \quad \delta \mathbf{V} \leq u(p_w q_c) - u(p_c q_c - D)$$

i.e., following (unexpectedly) *high* world price realizations p_w .

Under a differential contract, instead, the mill defaults if

$$(6) \quad \delta \mathbf{V} \leq u(p_w q_c) - u(\max\{(p_w + \Delta_c)q_c - D, 0\}).$$

Following high world price realizations the mill is - if anything - *less* likely to default. Substituting for the binding buyer's participation constraint (1) (i.e., $\Delta_c = 0$), the right-hand side of the constraint is bounded above by $u(D)$. If $u(D) \leq \mathbf{V}$ the mill never defaults strategically. Higher realizations of world prices, then, do not affect the likelihood of default under a differential contract. The contrasting response of default behaviour to unexpectedly high realizations of world prices across contractual forms gives our first testable prediction.

The second testable prediction relates relationship's value \mathbf{V} to the intensity of the response of default behaviour to unexpectedly high realizations of world prices. Under a fixed contract the likelihood of default is given by $P^F(\mathbf{V}) = 1 - F(u^{-1}(\mathbf{V} + u(p_c q_c - D))/q_c)$. Higher relationship value \mathbf{V} decreases the effect of unanticipated increases in the world price on the likelihood of default. To see why this is the case, note that $u' > 0$ and $F'' < 0$ in the right tail of the price distribution. In contrast, the effect of high realizations of world prices on the likelihood of default is zero regardless of relationship's value \mathbf{V} for differential price contracts. This gives our second prediction.

Finally, our third testable prediction relates the relationship's value \mathbf{V} to contract choice. The possibility of strategic default introduces a trade-off between the two contractual forms. A fixed contract protects the mill against price risk, but leaves the buyer and lender exposed to counterparty risk. A differential contract does not protect the mill against price risk, but allows the mill to commit to not strategically default. All else equal, then, the possibility of strategic default lowers the mill's pledgeable income under a fixed price contract relative to a differential contract.

The resulting trade-off is illustrated in Figure 4. For very large values of \mathbf{V} strategic default is very costly and, therefore, rare. A fixed price contract then is preferred as it offers insurance against price risk at relatively low costs. In the limit the mill receives the desired insurance and produces at first best levels $q_c = q_F^*$. For lower values of \mathbf{V} , however, the chances of a strategic default increase. This reduces the pledgeable income and the amount of production: the mill is credit constrained. For even lower values of \mathbf{V} the credit constraint becomes so severe that the mill prefers to switch to a differential contract and produce q_D^* . The mill is then insurance constrained, but not credit constrained.³² The model implies that relationships with higher value \mathbf{V} can afford signing fixed-price contracts that leaves them exposed to strategic default. This is our third testable prediction.

In sum, the model yields the three following testable predictions:

Testable Predictions

T1: *Unanticipated increases in spot prices increase the likelihood of default under fixed contract but not under differential contract;*

T2: *Conditional on a fixed contract, higher relationship value decreases the effect of unanticipated increases in the world price on the likelihood of default;*

T3: *More valuable relationships sign fixed price contracts.*

Section 4 tests these predictions. Note that because parties adjust the contractual form to the possibility of strategic default, strategic default can be detected only on fixed price contracts. The observed level of default, then, does not fully reveal the costs associated with imperfect enforcement. A possibly large share of the costs remains hidden under the lack of insurance and underinvestment of mills on differential contracts. To quantify these costs is the goal of the calibration exercise in Section 5.

³²For even lower values of \mathbf{V} the mill might be unable to fund the desired level of production even under a differential contract.

4. EMPIRICAL RESULTS

The empirical analysis is divided into two Sections. This Section tests the qualitative predictions from the model. We first exploit contract-specific unanticipated international coffee price movements to test for strategic default; we then explore heterogeneous responses depending on relationship strength; and finally, we look at selection into contract types. The evidence strongly supports the predictions of the model and suggests that strategic default might be responsible for either credit or insurance constraints in the market.³³ The next Section calibrates the model described in Section 3 to quantify the inefficiencies resulting from strategic default.

4.A. Prediction 1: Strategic Default

i) Baseline Test for Strategic Default: The key test for strategic default is that unanticipated increases in the world price of coffee p_w increase the likelihood of default on loans backed by fixed price contracts but not on loans backed by differential price contracts (Prediction 1). The key challenge to test this prediction is to control for parties expectations about prices at delivery \bar{p}_w . In futures markets the price quoted at the closing date for a future delivery date gives us parties expectations about market conditions at delivery. For each contract signed between mill m and buyer b at date t for deliveries at date t' we construct a measure of price surprise as:

$$(7) \quad P_{mbtt'} = \frac{p_w^{t'}}{\mathbf{E}[p_w^{t'}|t]}$$

in which $p_w^{t'}$ is the realized spot price at maturity, i.e., the random variable p_w in the model, and $\mathbf{E}[p_w^{t'}|t]$ is the futures price quoted at t for deliveries at t' , i.e., its expected value \bar{p}_w .

Figure 5 shows the relationship between increases in international coffee prices and loan defaults by contract type. The histogram shows the distribution of the ratio of New York ‘C’ Arabica coffee price at the scheduled shipment date, $p_w^{t'}$, divided by the futures price for the shipment date at the time the contract was signed, $\mathbf{E}[p_w^{t'}|t]$. The ratio gives a contract specific measure of price surprises.

³³Appendix F exploits discontinuous changes in loan size induced by the lender’s scoring model to implement a RDD to study the effects of larger loans. We find evidence that mills with relatively lower scores are indeed credit constrained. Some of the RDD estimates are also used to calibrate the model in Section 5.

Figure 5 also separates the defaults into fixed price contracts (blue) and differential price contracts (red). A loan is in default if it is written-off, restructured or has no payments after ninety days from its maturity dates.³⁴ Consistent with Prediction 1 all of the increase in defaults associated with unexpected surges in world coffee prices come from fixed price contracts. Conversely, we see no relationship between unexpected surges in world coffee prices and default when the contract is on differential.

The differential relationship between unexpected price surges and default across contract types is consistent with strategic default. The evidence, however, could be driven by confounding factors. For instance, mill or buyer characteristics might simultaneously affect contract choice and propensity to default at times of volatile prices. Different contracts could also be signed at times when defaults and price volatility co-move for unrelated reasons.

Table 2 provides an econometric investigation of the strategic default test that controls for these and other potential confounders. We control for time-invariant mill and buyer characteristics by taking advantage of the panel structure of the data including the relevant sets of fixed effects. We flexibly control for timing effects exploiting asynchronous timing of the harvest season across countries in the sample. Figure A5 shows seasonality patterns in the closing and maturity dates of loan contracts in the sample. The figure illustrates the bimodal distribution of both closing and maturity dates. The two peaks in each distribution are driven by asynchronous coffee harvest seasons across the two hemispheres.³⁵ The variation allows us to identify contract-specific changes in incentives to strategically default like in Figure 5 while controlling for time fixed effects in a flexible way in all the empirical specifications. Specifically, Table 2 reports results from the specification

$$(8) \quad D_{lmbt}^d = \alpha_0 + \alpha_1 \frac{p_{lmbt}}{\bar{p}_{lmbt}} + \alpha_2 \bar{p}_{lmt} + \alpha_3 p_{lmt} + \lambda_m + \mu_t + \gamma_b + \varepsilon_{lmbt}$$

where D_{lmbt}^d is a dummy taking value equal to one if mill m defaults on loan l closed at time t backed by buyer b . The main regressor of interest is $P_{mbtt'} = \frac{p_{lmbt}}{\bar{p}_{lmbt}}$, the ratio of the realized world price at the time maturity over the expected spot price for that date when the contract was signed defined in equation (7).³⁶ Furthermore,

³⁴Alternative definitions of default yields qualitatively similar results.

³⁵For example, most contracts in Peru (34% of the loans in the sample) are closed in the period May to June while in Nicaragua (11% of the loans in the sample) most contracts are closed in October to December.

³⁶The definition of maturity date depends on whether we run the specification at the loan level or at the contract level. In the first case we use the maturity date for the loan. In the second case, we use the date the shipment was scheduled.

we include buyer (γ_b) and mill (λ_m) fixed effects, as well as a set of year and month fixed effects (μ_t). Finally, ε_{lmbt} is an error term arbitrarily correlated across observations for the same mill m .

Table 2 reports results from variations of this specification. Column 1 presents OLS estimates for the fixed-price sample (on which we expect an effect). We see that a 10% increase in the world coffee price is associated with a three percentage point increase in the default rate. Column 6 shows the analogous estimate for the differential contracts, and in those contracts the same 10% increase is associated with only a 0.3 percentage point change in the default rate, an order of magnitude smaller and not statistically different from zero. To account for the asynchronous harvest seasons across countries, we allow the month fixed effects to vary by country in column 2 and the results are nearly identical. In column 3 we control for spot and futures prices, and again, the estimate is almost completely unaffected.³⁷ These estimates suggest that at most 59% of defaults are strategic.³⁸

Two distinct sources of variation identify our effect. First, loans that are signed at the same time might have different price surprises because they vary in length. Second, loans that have the same length are signed at different times. In the spirit of an over-identification test, we now isolate the two different sources of variation and verify that they produce nearly identical estimates. In column 4 we include a control for the length of the loan (in days) in order to estimate the effect solely based on the second source of variation in our treatment. There we find, again, a very similar estimate to the one in column 1. In column 5 we include year-month fixed effects to exploit only the variation that exists between loans that were signed in the same month. This strategy identifies the effect based solely on the first source of variation: loan length. This strategy again produces a very similar estimate to column 1.

We expect that this increase in defaults should only come on the fixed price contracts. In columns 6-8 we can see what happens to differential price contracts. Column 6 shows the estimate on the differential contracts with the baseline specification. This produces a statistically insignificant estimate about an order of

³⁷Results are robust to alternative definitions of default. See Table A2.

³⁸We obtain this upper bound as follows. Since we expect no default with a price surprise of less than one, default rates on fixed price contracts with those price surprises provide a baseline level of defaults due to other factors. We can then attribute the predicted difference in default associated to positive price surprises to the strategic motive. This produces an expected default rate of 14.5% when price surprise are greater than 1 and 6.8% otherwise. We compare this to the overall default rate on fixed price contracts (13%) and observe that $(14.5\% - 6.8\%)/13\% = 59\%$. Note that if we expect some of this default to be due to debt over-hang associated with increased costs from pass-through of world price increases to farmers, then this estimate should be considered an upper-bound (we offer an estimate of a lower bound in the next subsection).

magnitude smaller than the one estimated on fixed price contracts.³⁹ In columns 7 and 8 we show a difference-in-differences specification at the loan level rather than the contract level. A fixed price loan is now defined as a loan where the majority of the money is earned from fixed price contracts.⁴⁰ In this specification the $\frac{p_{lmbt}}{p_{lmbt}}$ variable represents the effect of a price surprise on differential contracts and again we estimate effects that are very close to zero, but if anything negative (consistent with theory). The effect on fixed-price contracts is confirmed in columns 7-8 when we look at the interaction between our price ratio $\frac{p_{lmbt}}{p_{lmbt}}$ and loans that are mostly backed by fixed price contracts. Column 7 shows the effect for the main specification, while column 8 allows for the observed asynchronous harvest seasons observed in the data, and both estimates are consistent with the estimates observed at the contract level.⁴¹

ii) Price Pass-through and Ex-post versus Ex-ante Moral Hazard: Evidence that the likelihood of default increases following unexpected surges in prices would normally also allow to distinguish strategic default from loan diversion as a source of moral hazard. Under loan diversion, unexpected surges in sales prices increase profits, reduce debt overhang, and provide the mill with stronger incentives to use the loan to source cherries. In our set-up, however, international coffee prices might be quickly transmitted to prices paid by mills to farmers. If this is the case, then unanticipated increases in international coffee prices during the harvest season increase the costs of sourcing raw material. If coffee is sold short, this reduces expected profits for mills with fixed price contracts but not for mills with differential price contracts. The reduction in expected profits, then, could induce the mill to divert the loan and trigger default via a standard debt-overhang effect. Because of this possibility, results in Figure 5 and Table 2 do not distinguish ex-ante moral hazard (the mill is unable to repay as the loan was not invested) from ex-post moral hazard (the mill is able to repay, but decides not to).

To isolate strategic default as a source of moral hazard we take advantage of the stark separation in time between production decisions (which occur during harvest) and contract execution (which happens after the harvest is over). Nearly 80% of contracts in the sample mature after the end of the harvest season (Ta-

³⁹Results for the corresponding specifications in Columns 2 to 5 are identical.

⁴⁰We use a 50% threshold. Results are robust to other thresholds (see Table A3 for a robustness exercise where we examine alternate thresholds of default.)

⁴¹The results also show that at the mean expected price surprise, fixed price contracts are less likely to default. This is also consistent with the model since fixed price contracts are positively selected (Prediction 3).

ble 1). To isolate strategic default motives, we distinguish out-of-harvest price surprises from in-harvest price surprises. After the end of harvest season, once cherries have been sourced and processed, price pass-through is no longer relevant and a price increase can only improve the profits of the mill. In this case defaults associated with unexpected price increases are unambiguously strategic. Meanwhile, in-season price increases could or not result in default either because of strategic default or because of ex-ante moral hazard, depending on the exact timing of coffee sourcing, price increases and transmission of prices to the country side.

Table 3 implements an event study approach (see figure 6 for an illustration). We separate price increases into ones that happened in-season and out-season. For each we consider a price increase to be a relevant event if it results in a weekly price increase of at least 3.0%.⁴² We then take small windows of between one and three weeks around the event and run a simple local-linear model to check whether shipments that were scheduled just before the price increase (and were therefore likely delivered before the realization of a price change) experience less default than shipments scheduled for just after a price increase. We run the analysis only at the contract level because of the small window around the more precisely relevant decision date.

Columns 1-3 of table 3 show the effect on default of out-of-harvest price increases, which can only be due to strategic default. The first column shows the difference using a two week window while the second shows a one week window and the third a three-week window. A two week window is our preferred specification given the trade-off between sample size and potential bias resulting from a big window. The results are consistent with table 2. In each case we find a large and statistically significant increase in the default rate of about 10-15%. We conclude from this that a large percentage of defaults under fixed price contracts are indeed strategic. In fact, to construct a back-of-the-envelope estimate we also report the control group means of the dependent variable. We see that for our main estimate default increased from 5.5% to nearly 20% as the result of the price jump, indicating that about 75% of defaults are strategic following a large sudden price increase. Of course, this comes from a 3% price increase which is very large: it is nearly at the 85th percentile of the distribution. We also consider smaller increases of between 1%-2.5% price increases. Those estimates though are actually

⁴²Other thresholds produce similar results (see table A4). The event study approach guarantees that in-season price increases are identical for contracts that mature just before and just after an out-of-season price increase. Results are robust to control for in-season price increases in the specification (see table A5).

quite similar (Table A4) and indicate that at least 42% of observed defaults on fixed price contracts are strategic.⁴³ In sum, combining the two empirical strategies leaves us quite confident that 42% to 59% of the defaults observed on fixed price contracts are strategic.

As expected, we see no analogous increase in defaults on differential price contracts (column 4). In column 5 we show the result for the in-season price increases. We find imprecise results. This could be because of the contrasting effects of the debt-overhang and pass-through mechanism described above, or simply because we have fewer in-season price increases since nearly 80% of contracts mature out-of-harvest. Regardless, the fact that the results are robust to considering only out-of-harvest price jumps suggest that strategic behaviour is an important source of defaults in this market. Taken together, the evidence strongly supports Prediction 1.⁴⁴ We now turn to Prediction 2.

4.B. Prediction 2: Heterogeneity

Using both the event-study methodology and an OLS-based approach we find that unexpected increases in the world price of coffee substantially increase the rate of default on fixed price, but not on differential price, contracts. However, given that coffee is primarily produced in countries with weak institutions and that arbitration clauses are hardly ever enforced, it might actually be surprising that *more* mills do not default when incentives to do so are strong. In the absence of formal contract enforcement, mills will weigh the short-run benefits of default against the long-run costs of jeopardizing valuable relationships with their partners. Conditional on a fixed-price contract, relationships with lower value are more likely to default in response to unanticipated price increases (Prediction 2). We test Prediction 2 by examining whether, conditional on a fixed price contract, unanticipated price surprises have a larger impact on the likelihood of default on less valuable relationships. Before testing the proposition, however, we provide evidence that indeed relationships deteriorate following a default. Specifically, we show that buyers and the lender both punish mills for defaulting by denying them

⁴³We obtain this lower bound as follows. Assume a constant effect of price surprises on default for week-over-week price change in the range 1-3% (Table A4) and no effect outside the range. The estimates imply that 60%-65% of default following price increases in the range is strategic. Approximately 64% of loans experience a week-over-week price change in the range. This suggests that about $64\% \times 65\% \approx 42\%$ of defaults are strategic. Note that since we expect that defaults may become more prevalent with larger price surprises, this estimate is a lower bound. It is also a lower bound if we expect that price surprises of less than 1% might induce strategic default.

⁴⁴Appendix E explores robustness of these results along several dimensions.

access to future loans.

i) Punishing Default and the Role of Building Relationships Table 4 shows OLS estimates that suggest that mills that default are less likely to be given a future loan by the lender, and that the severity of this punishment depends on the severity of the default. Column 1 presents the estimate of the least severe default, simply being three months late on repayment. When this occurs, mills are about 11.5% less likely to receive another loan from the lender. This estimate progressively increases, to the most severe definition of default, which is either outright non-payment or being more than 9 months late on repayment. Using this definition, the decrease in the likelihood of receiving a future loan more than doubles to almost 28%.

The evidence for buyer punishment is similar. Columns 5-8 suggest that when a mill is three months late on delivery, we are about 7.5% less likely to observe the buyer-mill pair again in the data, while if the mill is 9 months late we are about 12% less likely to observe the pair again. This evidence on the buyer's punishment is subject to the caveat that we do not observe every transaction made between the buyer and the mill. That is, we are unable to rule out that what we observe as buyer punishment may be driven by the fact that since the lender punishes the mill the buyer-mill pair would not be observed in the future (despite the fact that they may well be continuing their business relationship through another lender). In Appendix E.2 we investigate buyer punishment of the mill repaying the lender directly, which we argue is not (or is at least much less) subject to this caveat, and we get very similar estimates ($\approx -8\%$). Note that the evidence for lender punishment is not subject to the same concern since we observe all transactions between the mill and the lender. We interpret Table 4 as a whole as strong evidence that the mill faces punishments for default.

ii) Relationship Value and Strategic Default The effectiveness of this type of punishment depends on the value that the mill places on its relationship with the lender and the buyer (Prediction 2). We now test Prediction 2 by examining whether, conditional on a fixed price contract, unanticipated price surprises have a larger impact on the likelihood of default on less valuable relationships.

Table 5 explores the effect of price surprises, splitting the sample according to the strength of the relevant relationship. We consider measures of strength in the mill-lender, mill-buyer and lender-buyer relationships. Specifications are as in Table 2. In columns 1-3 we examine heterogeneity by the strength of the relation-

ship between the lender and the mill. Since we observe the entire history of the relationship between the mill and the lender, we can use the most straightforward proxy of relationship value: past history. We expect mills with more of a history (defined as the sum of past loans) to have a stronger relationship.⁴⁵ We find that both those with high and low relationship values are more likely to default given an unexpected price increase (columns 1-2), but those on fixed price contracts with relatively worse relationships with the lender are about five times as likely to default. We also find that even those with low relationship values on differential contracts are not more likely to default given an unexpected price increase (column 3).

In Columns 4-6 we analyze the buyer-mill relationship. This relationship is measured with more noise since transactions between the mill and the buyer that occurred before the mill receives loans from the lender are not observed in the dataset. However, one advantage of this data is that, unlike with the lender, we observe mills with many different buyers. We can therefore analyze the ‘fit’ between the buyer and the mill in a way that is much more difficult when looking at the lender’s relationships. To do this, we use a methodology similar to [Abowd et al. \(1999\)](#). We run a regression of relationship age on fixed effects for the year the contract was agreed to; the month the contract was agreed to; the year-quarter of the first observed buyer-mill transaction; and the buyer-mill pair. We use the buyer-mill pair fixed-effect as a proxy for relationship value (these are plotted in figure A6), as they capture unobserved heterogeneity in the buyer-mill pair that influences the length of their relationship. This accomplishes the same goal as using relationship history to proxy for relationship value, as we did with the lender, but also allows us to difference out buyer and time effects in a way that is impossible with the lender-relationships (since we observe only one lender). Using this measure of buyer-mill relationship, we find a very consistent pattern. We see again, that both those with high and low relationship values respond to positive price surprises by defaulting (columns 4-5), but the mills on fixed price contracts with the lowest relationship values are more than twice as likely to do so. Again, we do not observe similarly large default rates for those with low relationship values on the differential price contracts (column 6).

Finally, we examine buyer-lender relationships on the mill’s decision to default (column 7-9). The decision of the mill clearly depends on its own relationships

⁴⁵This measure is appropriate given the mill fixed-effects in each specification. Note that from the point of view of testing the prediction of the model it is not important to distinguish whether relationship’s age causes the relationship to be more valuable or whether relationship value simply correlates with relationship’s age due to, e.g., selection effects.

with both the lender and the buyer, since defaulting on one necessitates defaulting on the other (see Appendix D). However, the buyer-lender relationship may also matter for default. In particular, the mill may be more wary of defaulting on a buyer that is very important to the lender than they are of defaulting on a buyer that is less important to the lender.⁴⁶ Similar to the analysis with the mill's own relationships, we can test heterogeneity in default by the buyer-lender relationship. We take a measure of network centrality from the lender's network of buyers (see figure A9) with the idea that buyers connected to many different mills are more important to the lender's business. Defaults on fixed price contracts by high and low buyer-lender relationships are found in columns 7-8, and once again we find that relationships matter for mitigating default. Mills that have fixed price contracts with buyers that are unimportant to the lender are much more likely to default than mills with buyers that are of above average importance to the lender. Again, there is no similar rate of default among analogous differential price contracts (column 9). This is consistent with the mills perceiving some heterogeneity in the response of the lender to default. The lender, for example, may be more willing to accept repayment from a buyer not listed on the contract when the contracted buyer is not important to the lender. However when the buyer accounts for a large share of the lender's operations, the lender may be less willing to accept repayment following a default on the buyer. In this case, for the mill to default on the buyer it also must default on the lender, which may prove too costly as the mill then faces punishment by both.⁴⁷

4.C. Prediction 3: Relationship Value and Contract Selection

Finally, Table 6 tests whether fixed price contracts tend to be signed in more valuable relationships (Prediction 3). If the main punishment mechanism that buyers and the lender have is to end the relationship with the mill upon default, then we expect that when this relationship is strongest, the buyer and the lender feel most comfortable signing a fixed price contract. We test the model prediction running the following specification:

$$(9) \quad F_{cmbt} = \alpha_0 + \alpha_1 R_{mbt} + \alpha_2 X_{cmbt} + \lambda_m + \gamma_b + \mu_t + \varepsilon_{cmbt}$$

⁴⁶A buyer that has a weaker relationship with the lender might also have higher incentives to default, although not differentially across contract types at times of positive price surprises.

⁴⁷Appendix D provides further evidence on this.

where F_{cmbt} is a dummy taking value equal to one if contract c between mill m and buyer b signed at time t is fixed price. The main regressor of interest is R_{mbt} , a measure of relationship value between mill m and buyer b at the time they sign the contract. Furthermore, we include contract level controls X_{cmbt} as well as buyer (γ_b), mill (λ_m) and time fixed effects (μ_t). Finally, ε_{lmbt} is an error term arbitrarily correlated across observations for the same mill m .⁴⁸

We test the prediction using four measures of relationship value. First is our preferred measure, used in the previous table, and based on an [Abowd et al. \(1999\)](#) style regression (see a plot of the relevant fixed-effects in figure A6). Second, we can use a network-based proxy for buyer-mill relationship, similar to the one used to proxy for the buyer-lender relationship in table 5, columns 7-9. In the context of the buyer-mill relationship this variable is higher when more money used as collateral for the loan comes from fewer buyers - meaning that any one buyer becomes more important for the transaction. Third, we can use a measure analogous to the lender-mill relationship proxy used in the previous table: the history of business done between the firms. Finally, we use the lender's perceived quality of the mill-buyer relationship, represented by a score out of five in the credit application. This last measure seems to be ideal, except for the fact that it was only recently added to the credit applications, so well over half of the observations are missing.

Using all four measures, we find that the positive selection effect dominates. That is, buyers are more likely to sign fixed price contracts with mills with whom they have the best relationships, not the worst. The data suggests that buyers are wary of default and are careful about who they sign a fixed price contract with. In columns 1-4 we test the [Abowd et al. \(1999\)](#) style proxy, and find that a relationship expected to last 1,000 days longer is about 10% more likely to receive a fixed price contract. We test the robustness of the model to our main robustness exercises in Table 2, and the estimate is quite stable in each case.

Similarly in columns 5 we find that a higher concentration of collateral is associated with fixed price contracts (the units are more difficult to interpret); column 6 shows that an additional ten million dollars in past history with a buyer is associated with a 4% higher likelihood of receiving a fixed price contract; while contracts one point higher out of five on the lender's perception of the mill-buyer relationship are about 20% more likely to receive a fixed price contract.

Our analysis suggests that relationships are crucial to lending in this context.

⁴⁸Note that the inclusion of mill, buyer and time fixed effects implies that we cannot separately identify the strength of the relationships at the mill-lender and at the buyer-lender level.

They not only determine the terms of the agreements, but they heavily influence contract outcomes.

5. MODEL CALIBRATION

The evidence in the previous section lends strong support to the model in Section 3: a significant share of observed defaults is strategic (Proposition 1); relationship value is a key deterrent of strategic default (Proposition 2); parties structure contracts to minimize the adverse consequences of strategic default (Proposition 3). Because parties adjust the contractual form to the possibility of strategic default, the observed level of strategic default does not fully reveal its costs. A possibly large share of the costs remains hidden under the lack of insurance and the resulting underinvestment. Given the comprehensive empirical support for the model’s predictions, this Section calibrates the model to quantify the inefficiencies caused by strategic default.

5.A. Empirical Strategy

The model in Section 3 derives predictions from a limited sets of parameters. Many of these parameters are directly observed in the data or can be calibrated or estimated. The key parameter we want to learn about is \mathbf{V} , the value the mill places on keeping a good relationship with the buyer-lender.

We pursue the following strategy. We take advantage of the model’s relative simplicity to “invert” it and obtain an estimate of \mathbf{V}_i for each loan (see figure 7 for the distribution of \mathbf{V}_i). Specifically, given a set of parameters we find the \mathbf{V}_i that rationalizes a loan’s key contractual outcomes: the interest rate r_i and whether the loan is backed by a fixed or a differential contract. Although in principle we could estimate loan-specific \mathbf{V}_i matching additional outcomes, the interest rate and the contract type present two main advantages. They are both recorded without error in the dataset. Second, they are intimately connected with \mathbf{V}_i in the model (see Figure 4). This makes the identification of \mathbf{V}_i particularly transparent.⁴⁹

We distinguish two sets of parameters: those that are constant across loans; and those that vary. The former (denoted Z), captures the distribution of price surprises, the risk aversion of the mill and the slope of the farmers’ supply curve. The distribution of price surprises $F(p_w)$ is directly observed in the data and is well approximated by a log-normal distribution. We assume an utility function given by $u(x) = x^{1-\alpha}$. We calibrate α to match the average forward discount

⁴⁹See Appendix A for details.

in the data. That is, we assume that the risk averse mill is indifferent between a random draw from the price distribution $F(p_w)$ and a sure payoff equal to the current spot price.⁵⁰ Finally, the slope of the coffee cherries supply curve, η , is estimated from the RDD analysis presented in the Appendix. Table A10 shows the effects on the costs and average price paid to farmers of an (exogenous) increase in loan size of approximately 100,000 USD. These two estimates allow us to recover η . Finally, we let the cost parameter γ_i vary by loan. The cost parameter γ_i is directly reported in the financial accounts of the mill. The operating costs take the form $C(q_i) = \gamma_i \times q_i^{(1+\eta)}$. Operating costs $C(q_i)$ and production volumes q_i are directly observed in the financial records. Knowledge of η , then, allows us to assign γ_i to each loan for which financial accounts are available.⁵¹ The calibrated parameters are reported in Table 7.

To estimate the loan specific relationship value \mathbf{V}_i we feed the parameters Z and γ_i into the model and find the \mathbf{V}_i that best matches the observed interest rate r_i and contract type for that loan. Conditional on γ_i , interest rate and contract type are strongly correlated with each other (p-value of 0.00). The estimates match the correct contract type approximately 90% of the time. Appendix A provides further details.

5.B. Results

The main results, alongside counterfactuals and sensitivity checks on the calibrated values of α and η are reported in Table 8. The first row of the Table reports the estimated \mathbf{V}_i . We find that for the median (mean) observation in the sample, the value of the relationship amounts to 44% (158%) of the sales value on the contract. For loans backed by fixed contracts, these estimates can be directly compared with lower (upper) bounds for non-defaulting (defaulting) loans obtained from the incentive compatibility constraint. The estimated \mathbf{V}_i appear to be in the correct ball park (see Figure 7).⁵²

The second row quantifies inefficiencies by comparing the predicted production volume with the implied first best volume (which can be analytically computed).

⁵⁰The average price surprise is slightly above one (Table 1) reflecting forward discounts (see, e.g., Dana (1998)). Essentially we use the forward discount to calibrate α .

⁵¹Note that because we need the financial statement data to construct \mathbf{V}_i we can construct it only for the smaller sample of loans for which we have the financial data.

⁵²Estimated \mathbf{V}_i correlate relatively well with the proxies used for relationship. Note however that, by construction, estimated \mathbf{V}_i correlate with fixed contract type and can't therefore be used as a further test for Prediction 3. The fact \mathbf{V}_i can only be estimated on the smaller sample for which financial statements are available justifies why we do not use them to test Prediction 2, however our test for prediction 2 (table 5) is robust to using this estimated \mathbf{V}_i .

This comparison also yields our main counterfactual: by how much would production increase if we removed strategic default?⁵³ We find that for the median (mean) observation, production would be 19.7% (15.8%) higher in the absence of strategic default.

The average effect masks substantial heterogeneity. The estimates suggest that 26% of mills produce at first best. That is, at the 25th percentile, relationship value V_i is sufficiently large that there is no output loss due to strategic default. Looking at the third row, we see that 65% of the 108 mills predicted to be on fixed contracts produce at the first best level. The average mill on a fixed contract produces 11.3% less than the first best.

When the threat of strategic default is particularly severe, its consequences are mitigated by using differential contracts. Rows 4, 5 and 6 look at the remaining 199 mills that are predicted to be on differential contracts. These mills produce on average 18% less than at the first best (row 4). The output gap relative to the optimal quantity conditional on a differential contract is minimal (row 5). This implies that the vast majority of these mills (62%) are insurance constrained, i.e., they produce less than at the first best level due to exposure to price risk but, conditional on such exposure, they would not want to expand production. This group of mills account for 39% of the overall sample. Finally row 6 shows that these mills would produce 50% less output if they were forced to sell on a fixed contract. This is a very large number that illustrates that mills signing differential contracts would be very severely constrained if they had to rely on the collateral value of their relationships to insure against price risk.

Finally, Rows 7 and 8 look at the wedge between the physical marginal product of capital (MPK) and the risk free interest rate. The MPK is the additional quantity that the mill would produce if it was given an additional unit of capital at the loan interest rate. As done in rows 2 to 6 we therefore focus on quantity distortions and ignore uninsured risk (which generates a wedge between the expected marginal revenue and the interest rate for insurance constrained mills). For the majority of mills that are either producing at first best (26%) or are insurance constrained (39%), the wedge is equal to zero: these mills would not want to produce more if given additional capital. The remaining 35% of mills, however, are credit constrained, some severely so. These mills would take-up additional finance at the loan interest rate and use it to produce more. On this group of mills the

⁵³In practice, it is not going to be feasible to completely remove strategic default. Furthermore, even if it was, other incentive constraints that are currently not binding might become so. The counterfactual is useful to gauge the severity of strategic default, not to assess effects of any particular policy.

estimates suggest a median (mean) wedge of 8% (20%). This implies an average marginal product of capital approximately equal to 30%. These results are in line with two pieces of evidence in the RDD analysis in Tables A9 and A10. First, as predicted by the model and the calibration, we find evidence of credit constraints for some, but not all, borrowers. Specifically, borrowers around the lower of the two thresholds in the lender scoring system appear to be credit constrained. We do not find evidence of credit constraints for borrowers around the higher threshold.⁵⁴ Second, on the lower threshold (where we do find evidence of credit constraints) we estimate an average gap between MPK and r of about 7%, almost identical to the median estimated by calibrating the model.

Finally, it is worth noting that the lower output produced by the mills as a result of strategic default has implications for farmers welfare. In particular, we can bound farmers welfare losses as follows. As an upper bound, we can interpret the cherries supply curve as the farmers supply curve and infer a $(1/0.84)^{(1+\eta)} - 1 \approx 32\%$ higher welfare for farmers supplying the average mill in the absence of strategic default. As a lower bound, we can ignore any quantity response and simply use the increase in prices paid to farmers as a result of larger loans (Table A10, Column 4). These estimates suggest that at the average mill farmers welfare would be $(15.8\%/20.4\%) \times 13.4\% \approx 10.4\%$ higher in the absence of strategic default; still a sizeable effect.

6. CONCLUSION

Strategic default - the possibility that a party in a contractual agreement deliberately defaults even when successful performance is feasible - can severely hamper market functioning. Yet, empirically identifying strategic default and quantifying its consequences remains challenging. While we do observe defaults, we typically do not know if any particular default occurs because the defaulting party cannot execute the contract, or does not want to.

This paper developed and implemented a method to identify and assess the importance of strategic default empirically. The test builds upon a critical insight in the theoretical literature: strategic default occurs when market conditions change sufficiently to place a business relationship outside its self-enforcing range. We apply the test to a sample of pre-financing agreements involving coffee mills in several developing countries. We develop a theoretical model that clarifies how strategic default introduces a trade-off between insurance and counterparty risk.

⁵⁴Note that the lender assigns *higher* scores to loans with differential contracts.

Relative to contracts that index prices to market conditions, fixed-price contracts insure against price swings but create incentives to default when market conditions change. This observation has implications for the design of a test for strategic default as well as for sorting of parties into contractual types.

We find ample support for the predictions of the model. To test for the presence of strategic default, we construct contract specific measures of unanticipated changes in market conditions by comparing spot prices at contract maturity with the relevant futures prices at the contracting date. Controlling for mill; buyer; time; and country-specific seasonality effects, we show that unanticipated increases in market prices increase defaults on fixed price contracts but not on price-indexed ones. We isolate strategic default from competing causes by focusing on unanticipated rises at the time of contract execution after all production decisions are sunk. We show that strategic default appears to be more severe in less valuable relationships. In turn, those relationships tend to sign price-indexed contracts precisely to limit the negative consequences of strategic default.

A model calibration suggests that strategic default has severe consequences for the functioning of this market. Strategic default causes significant output distortions: the median (mean) mill production would be 19.7% (15.8%) higher if contracts were perfectly enforceable. The estimates suggest that 26% of mills are unconstrained; 39% of the mills are insurance constrained; and the remaining 35% of mills that are credit constrained, many severely so. These distortions translate into a highly skewed distribution of the marginal product of capital across mills. Furthermore, strategic default implies externalities along the supply chain: output losses at the mill level translate into lower demand, and lower prices paid, for coffee delivered from farmers. Our estimates bound welfare losses for farmers supplying the average mill between 10% and 32%.

This paper studies a common problem in a specific context. The results have policy implications, particularly so for developing countries aiming at improving their business environments and exports. For instance, many developing countries heavily rely on export revenues generated in few, highly volatile, mineral and agricultural markets. However, access to risk-management tools is limited. The paper shows that fostering contract enforcement and strengthening interfirm relationships along the supply chain can yield significant degrees of insurance and expand output. Our results also suggest that a combination of counterparty risk and limited liability might significantly reduce both the supply and the demand for hedging tools, even among relatively large exporters. Finally, our results also suggest the existence of externalities along the domestic value chain: supplying

farmers might suffer significant welfare losses as a result of governance failures at the processing stage. Strengthening the governance of large exporters might yield large payoffs upstream.

What is perhaps most striking about our results is that the possibility of strategic default appears to severely hamper the working of firms that are, by developing countries standards, very large (see, [Hsieh and Olken \(2014\)](#), [Banerjee and Duflo \(2014\)](#)). Understanding barriers to the operation of large firms has significant implications given that even in developing countries those account for a large share of capital invested and, perhaps more importantly, there is limited evidence that small firms can bootstrap their growth and generate much needed high-quality jobs ([Hsieh and Klenow \(2014\)](#)). The results however also call for caution and context specific approaches: we have documented that the relevant binding constraints can be different across firms even within a narrowly defined sector. Further research to establish the form and extent through which contractual frictions hamper the operation of large firms in other contexts should be an important area for future research.

REFERENCES

- Abowd, John M, Francis Kramarz, and David N Margolis**, “High wage workers and high wage firms,” *Econometrica*, 1999, *67* (2), 251–333.
- Adams, William, Liran Einav, and Jonathan Levin**, “Liquidity Constraints and Imperfect Information in Subprime Lending,” *American Economic Review*, *99* (1),” *The American Economic Review*, March 2009, *49*, 84.
- Amiti, Mary and David E Weinstein**, “Exports and financial shocks,” *The Quarterly Journal of Economics*, 2011, *126* (4), 1841–1877.
- Antras, Pol**, *Global Production: Firms, Contracts, and Trade Structure*, Princeton University Press, 2015.
- and **C Fritz Foley**, “Poultry in Motion: A Study of International Trade Finance Practices,” *Journal of Political Economy*, 2015, *123* (4).
- Bajari, Patrick, Chenghuan Sean Chu, and Minjung Park**, “An empirical model of subprime mortgage default from 2000 to 2007,” Technical Report, National Bureau of Economic Research 2008.
- Baker, George, Robert Gibbons, and Kevin J Murphy**, “Relational Contracts and the Theory of the Firm,” *The Quarterly Journal of Economics*, 2002, *117* (1), 39–84.
- Banerjee, Abhijit and Esther Duflo**, “Reputation effects and the limits of contracting: A study of the Indian software industry,” *The Quarterly Journal of Economics*, 2000, *115* (3), 989–1017.
- and —, “Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program,” *Review of Economic Studies*, 2014, *81* (2), 572–607.
- Benmelech, Efraim and Nittai K Bergman**, “Liquidation values and the credibility of financial contract renegotiation: Evidence from US airlines,” *The Quarterly Journal of Economics*, 2008, *123* (4), 1635–1677.
- Breza, Emily and Andres Liberman**, “Financial contracting and organizational form: Evidence from the regulation of trade credit,” *The Journal of Finance*, 2017, *72* (1), 291–324.
- Bubb, Ryan, Supreet Kaur, and Sendhil Mullainathan**, “Barriers to Contracting in Village Economies: A Test for Enforcement Constraints,” 2016.

- Burkart, Mike and Tore Ellingsen**, “In-kind finance: A theory of trade credit,” *American Economic Review*, 2004, pp. 569–590.
- Casaburi, Lorenzo and Jack Willis**, “Time vs. State in Insurance: Experimental Evidence from Contract Farming in Kenya,” 2016.
- Chiappori, Pierre-André and Bernard Salanie**, “Testing for asymmetric information in insurance markets,” *Journal of Political Economy*, 2001, 108 (1), 56–78.
- Chiappori, Pierre-Andre and Bernard Salanié**, “Testing contract theory: A survey of some recent work,” in L. Hansen M. Dewatripont and P. Turnovsky, eds., *Advances in Economics and Econometrics - Theory and Applications, Eighth World Congress*, Econometric Society Monographs, Cambridge University Press, Cambridge 2003, pp. 115–149.
- Dana, James D**, “Advance-purchase discounts and price discrimination in competitive markets,” *Journal of Political Economy*, 1998, 106 (2), 395–422.
- de Janvry, Alain, Craig McIntosh, and Elisabeth Sadoulet**, “Fair Trade and free entry: Generating benefits in a disequilibrium market,” *Unpublished Mimeo*, 2010.
- Demirgüç-Kunt, Asli and Vojislav Maksimovic**, *Firms as financial intermediaries: Evidence from trade credit data*, Vol. 2696, Citeseer, 2001.
- Djankov, Simeon, Rafael La Porta, Florencio Lopez de Silanes, and Andrei Shleifer**, “Courts,” *The Quarterly Journal of Economics*, 2003, 118 (2), 453–517.
- Dragusanu, Raluca and Nathan Nunn**, “The Impacts of Fair Trade Certification: Evidence From Coffee Producers in Costa Rica (Preliminary and Incomplete),” *mimeo - Harvard University*, 2014.
- Ellingsen, Tore and Eirik Gaard Kristiansen**, “Financial Contracting Under Imperfect Enforcement,” *The Quarterly Journal of Economics*, 2011, 126 (1), 323–371.
- Fafchamps, Marcel**, *Market institutions in sub-Saharan Africa: Theory and evidence*, MIT press, 2003.

- **and Ruth Vargas Hill**, “Price transmission and trader entry in domestic commodity markets,” *Economic Development and cultural change*, 2008, 56 (4), 729–766.
- Giannetti, Mariassunta, Mike Burkart, and Tore Ellingsen**, “What you sell is what you lend? Explaining trade credit contracts,” *The Review of Financial Studies*, 2011, 24 (4), 1261–1298.
- Goldberg, Victor P and John R Erickson**, “Quantity and price adjustment in long-term contracts: A case study of petroleum coke,” *The Journal of Law and Economics*, 1987, 30 (2), 369–398.
- Greif, Avner**, “Contract enforceability and economic institutions in early trade: The Maghribi traders’ coalition,” *The American economic review*, 1993, pp. 525–548.
- Grossman, Sanford J and Oliver D Hart**, “An analysis of the principal-agent problem,” *Econometrica*, 1983, pp. 7–45.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales**, “The determinants of attitudes toward strategic default on mortgages,” *The Journal of Finance*, 2013, 68 (4), 1473–1515.
- Hart, Oliver**, “Hold-up, asset ownership, and reference points,” *The Quarterly Journal of Economics*, 2009, 124 (1), 267–300.
- **and John Moore**, “Default and renegotiation: A dynamic model of debt,” *The Quarterly Journal of Economics*, 1998, 113 (1), 1–41.
- Hölmstrom, Bengt**, “Moral hazard and observability,” *The Bell journal of economics*, 1979, pp. 74–91.
- Hsieh, Chang-Tai and Benjamin A Olken**, “The missing “missing middle”,” *The Journal of Economic Perspectives*, 2014, 28 (3), 89–108.
- **and Peter J Klenow**, “The life cycle of plants in India and Mexico,” *The Quarterly Journal of Economics*, 2014, 129 (3), 1035–1084.
- Joskow, Paul L**, “Price adjustment in long-term contracts: the case of coal,” *The Journal of Law and Economics*, 1988, 31 (1), 47–83.

- Karlan, Dean and Jonathan Zinman**, “Observing unobservables: Identifying information asymmetries with a consumer credit field experiment,” *Econometrica*, 2009, 77 (6), 1993–2008.
- Klapper, Leora**, “The role of factoring for financing small and medium enterprises,” *Journal of Banking & Finance*, 2006, 30 (11), 3111–3130.
- , **Luc Laeven, and Raghuram Rajan**, “Trade credit contracts,” *Review of Financial Studies*, 2012, 25 (3), 838–867.
- Klein, Benjamin**, “Why hold-ups occur: the self-enforcing range of contractual relationships,” *Economic inquiry*, 1996, 34 (3), 444–463.
- Kranton, Rachel and Anand V Swamy**, “Contracts, hold-up, and exports: textiles and opium in colonial India,” *The American Economic Review*, 2008, 98 (3), 967–989.
- Lacker, Jeffrey M and John A Weinberg**, “Optimal contracts under costly state falsification,” *The Journal of Political Economy*, 1989, pp. 1345–1363.
- Little, Peter D and Michael Watts**, *Living under contract: contract farming and agrarian transformation in sub-Saharan Africa*, Univ of Wisconsin Press, 1994.
- Macchiavello, Rocco and Ameet Morjaria**, “Competition and Relational Contracts: Evidence from Rwanda’s Coffee Mills,” *mimeo*, 2015.
- and – , “The value of relationships: evidence from a supply shock to Kenyan rose exports,” *The American Economic Review*, 2015, 105 (9), 2911–2945.
- MacLeod, Bentley W**, “Reputations, relationships, and contract enforcement,” *Journal of economic literature*, 2007, 45 (3), 595–628.
- Manova, Kalina**, “Credit constraints, heterogeneous firms, and international trade,” *Review of Economic Studies*, 2012, 80 (2), 711–744.
- Mayer, Christopher, Edward Morrison, Tomasz Piskorski, and Arpit Gupta**, “Mortgage modification and strategic behavior: evidence from a legal settlement with Countrywide,” *The American Economic Review*, 2014, 104 (9), 2830–2857.

- McCrary, Justin**, “Manipulation of the running variable in the regression discontinuity design: A density test,” *Journal of Econometrics*, 2008, *142* (2), 698–714.
- Paravisini, Daniel, Veronica Rappoport, Philipp Schnabl, and Daniel Wolfenzon**, “Dissecting the effect of credit supply on trade: Evidence from matched credit-export data,” *The Review of Economic Studies*, 2014, *82* (1), 333–359.
- Rampini, Adriano A and S Viswanathan**, “Collateral, risk management, and the distribution of debt capacity,” *The Journal of Finance*, 2010, *65* (6), 2293–2322.
- Rampini, Adriano, Amir Sufi, and S Viswanathan**, “Dynamic risk management,” *Journal of Financial Economics*, 2014, *111*, 271–526.
- Roberts, Michael R and Amir Sufi**, “Financial contracting: A survey of empirical research and future directions,” *Annu. Rev. Financ. Econ.*, 2009, *1* (1), 207–226.
- Shleifer, Andrei and Daniel Wolfenzon**, “Investor protection and equity markets,” *Journal of Financial Economics*, 2002, *66* (1), 3–27.
- Startz, Meredith**, “The value of face-to-face: Search and contracting problems in Nigerian trade,” *Mimeo*, 2017.
- Varangis, Panos and Bryan Lewin**, “Approaches to managing coffee price risks,” *Agricultural Commodity Markets And Trade: New Approaches to Analyzing Market Structure And Instability*, 2006, p. 282.

MAIN FIGURES

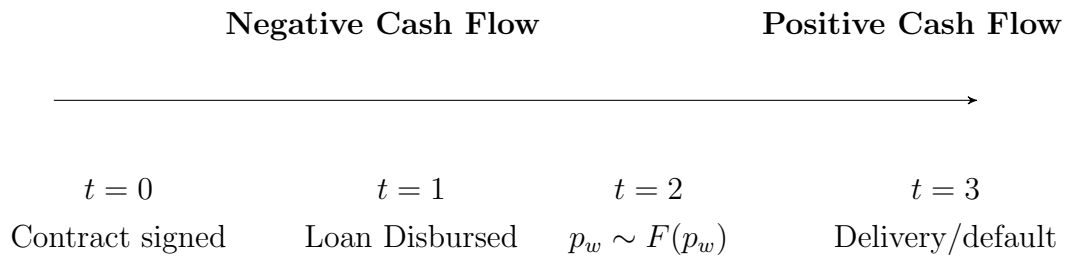


Figure 1: Timing of Events

Notes: The contract between the buyer and mill is signed at the beginning of the harvest season ($t = 0$). This contract is used to secure a loan from the lender, and the loan money is disbursed as needed throughout the harvest season for the mill to purchase cherries ($t = 1$). It is during harvest season, then, that the mill could potentially divert the loan (ex-ante moral hazard). After purchasing cherries it is possible that the world price of coffee changes. Price changes *after* the end of the harvest season are not passed through to farmers. The relevant spot market price p_w for the delivery date is drawn from the distribution $F(p_w)$ ($t = 2$). Once mills know the realized spot market price p_w , they decide whether to follow through with the contract they signed at $t = 0$ or to sell the cherries to another buyer at the prevailing spot price and strategically default ($t = 3$).

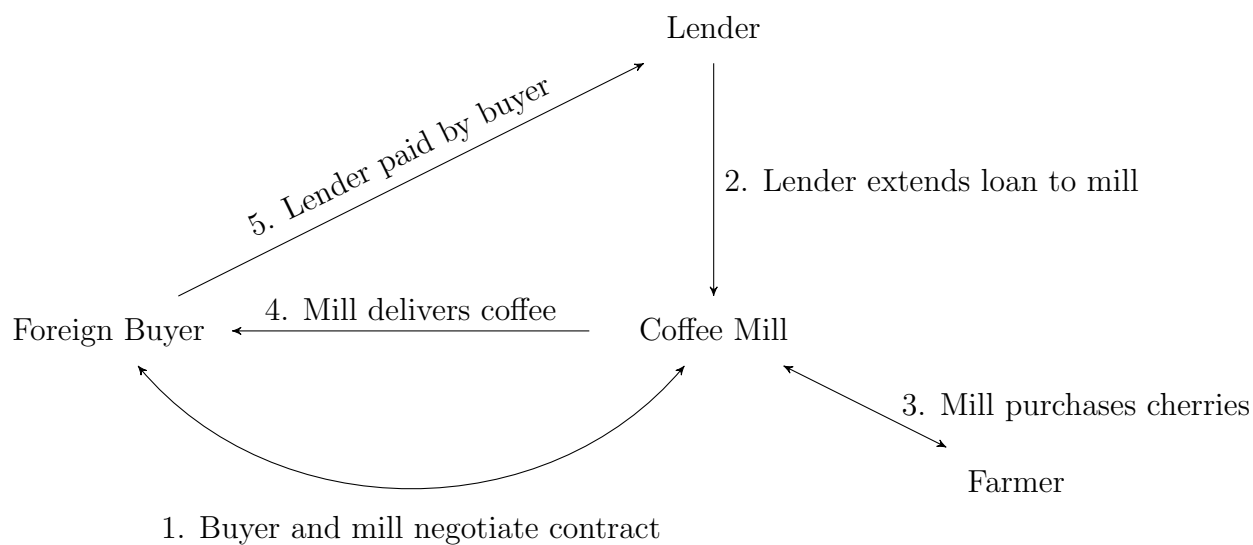


Figure 2: Lending Model Under Normal Circumstances

Notes: This figure shows the lending model that the lender uses when everything goes as planned (see the case of when a mill defaults in figure A4). Each step is numbered based on the sequence in which events typically occur. In this case the mill and the buyer agree on a contract at the beginning of the harvest season which sets a price and quantity of coffee to be delivered by the mill at a specific future date. Using this contract as collateral, the mill then secures a loan from the lender. The loan amount is based on a formula which decides on a fraction of the value of the contract, and which varies based on a credit score received by the mill during the application process. The mill uses the loan money to purchase coffee cherries from farmers, they process the cherries and deliver the agreed upon quantity to the buyer. The buyer then repays loan to the lender directly.



Figure 3: Time Series Graph of World Coffee Prices

Note: The graph plots the time series of coffee prices over time. To provide context relative to a harvest season, shaded in grey is one harvest period for Honduras. We chose this harvest season because it is of typical length but also because it has experienced one of the largest price increases over a harvest period in the sample (nearly 50%).




Figure 4: Insurance vs. Enforcement Trade-Off

Note: The Figure illustrates how the Value of the Relationship \mathbf{V} alters the solution of the model. The x-axis reports the Value of the Relationship \mathbf{V} , the y-axis reports the quantity produced by the mill under different contracts and scenarios. In the first best, there is no strategic default. When this is the case the quantity produced by the mill does not depend on the value of the relationship. By providing price insurance, a fixed-price contract induces the mill to produce a higher quantity than a differential contract. In the second best, however, there is strategic default. When this is the case, fixed price contracts leave the buyer-lender exposed to the risk of strategic default. This lowers the mill's pledgeable income, the amount the mill can borrow and, consequently, the quantity produced. A higher relationship value \mathbf{V} reduces the likelihood of strategic default and allows the mill to borrow more. Eventually, for very high values of \mathbf{V} the solution approaches the first best. For lower values of \mathbf{V} , however, the mill is better off foregoing price insurance and signing a contract on differential. This mitigates the strategic default motive and increases pledgeable income relative to a fixed price contracts. The model is numerically solved assuming the functional forms and parameters described in Section 5.

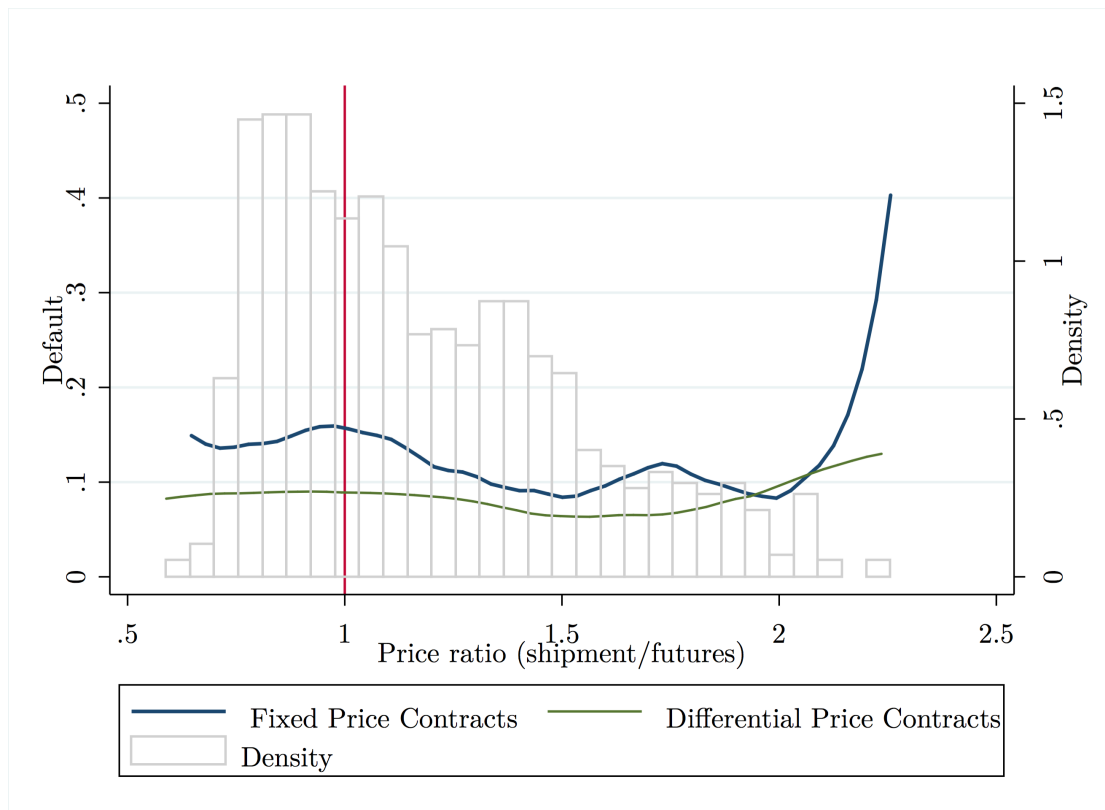


Figure 5: Unexpected price increase and contract default

Notes: The grey bars indicate the frequency of a given price surprise (x-axis), which is defined as the price at the time of loan maturity divided by the futures price for that date, at the time that the loan was signed. The figure is at the loan level and uses the 50% threshold for fixed price loans described in table 2. For a given price surprise, the darker blue line plots defaults on fixed price contracts while the lighter green line plots the same for differential price contracts. Defaults are defined as any outright default, or any loan that is completely unpaid as of 90 days past-due.

The figure shows that default is driven by large price surprises and fixed price contracts. The highest default rates are among those that experienced world prices that were much larger than the price at closing, as they were due to ship the coffee. This size of price increase, however, was a fairly rare occurrence. Nevertheless, we see a roughly two-fold increase in defaults among those experiencing the highest rate of price increases. We observe that for large jumps in the world price of coffee that defaults are entirely driven by fixed price contracts. We see little to no increase in defaults among contracts that tie prices to the world price.

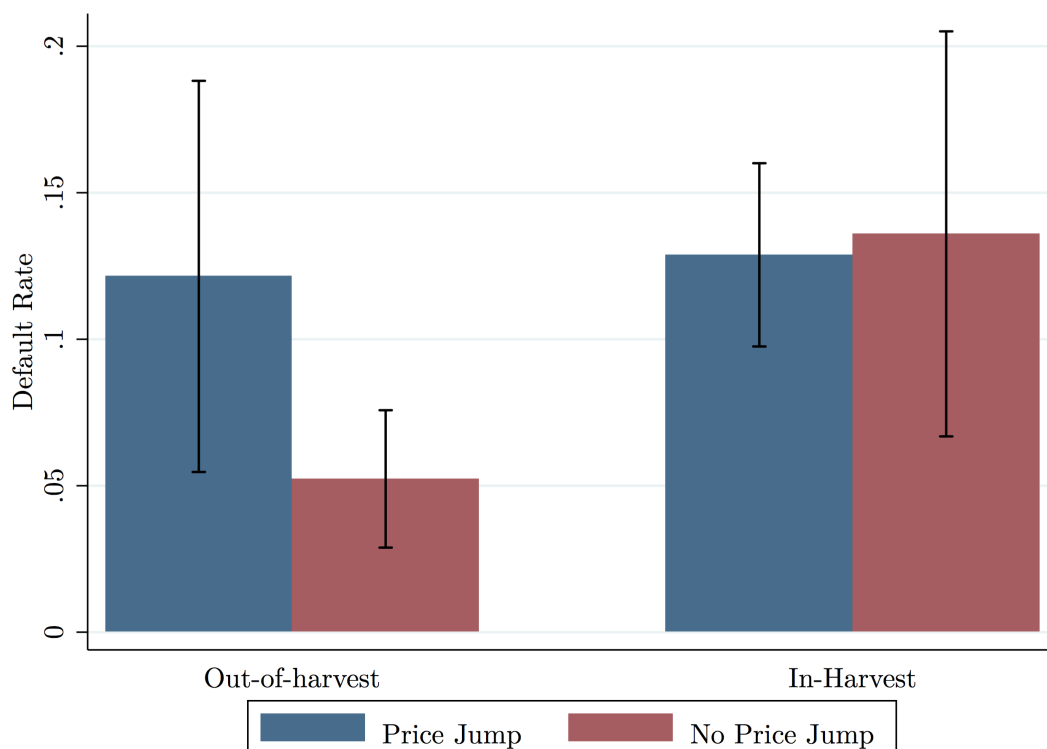


Figure 6: Default before/after price jumps in / out of harvest season

Notes: The figure examines defaults on fixed price contracts before and after a large price jump, that occur in-harvest and out-of-harvest. Contracts are defined as being in default if the first payment occurs at least 90 days past due. The ‘No price jump’ bars represent default rates when a shipment was scheduled within a two week window before a large price jump. The ‘price jump bars’ represent default rates when a shipment was scheduled in the two week window after a price jump. We define a ‘price jump’ here as any weekly price increase of at least 3% (see the appendix for robustness on this dimension). The figure shows that after an unexpected price jump, the defaults among the fixed price contracts rise for out-season price increases only.

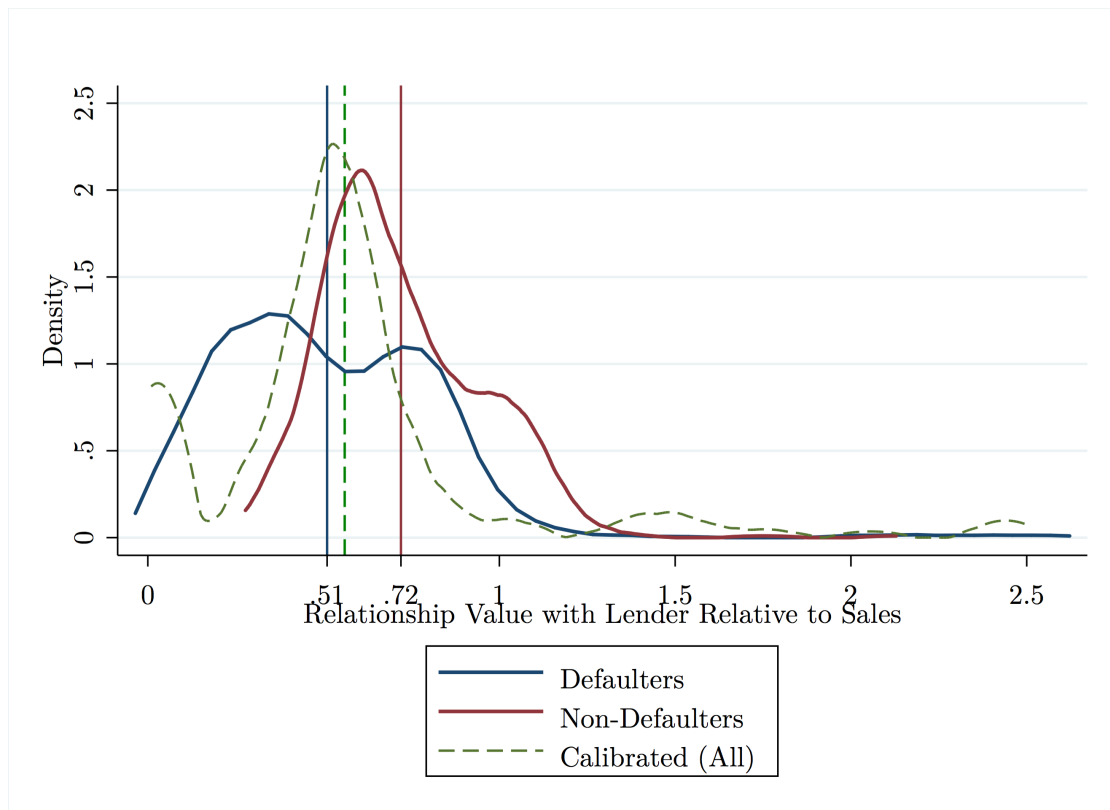


Figure 7: Upper Bound, Lower Bound and Calibrated Relationship Values

Notes: This figure shows upper and lower bounds for the relationship value derived using a revealed preference approach. We plot the mean relationship value for defaulters as a lower bound and the relationship value for non-defaulters as an upper bound. This is constructed using the probability of repeat interactions minus the estimated likelihood of punishment in the case of default plus the potential premium for defaulting. The idea is that those that defaulted did take advantage of the premium, thus for defaulters we have an upper bound of their relationship value (they must value the relationship by less than the premium). For non-defaulters we have a lower-bound (they did not default so they must value the relationship more than the premium). We also plot the estimated relationship value using the calibration exercise which is derived using a completely unrelated mechanism. See appendix A for details. We note that the calibrated relationship value that we construct is right within the bounds suggested by the revealed preference approach.

TABLES

Table 1: Descriptive Statistics

Variable	Observations	Mean	SD	Min	Max
	(1)	(2)	(3)	(4)	(5)
Panel A: Loans					
Default	967	0.082	0.238	0	1
Loan amount (USD)	967	473,012	553,040	8,500	4,500,000
Interest Rate	967	9.8%	0.10%	8%	18%
New Borrower	967	46.7%	49.9%	0	1
Length of Loan (days)	967	247.75	64.3	42	365
Number of Buyers Providing Collateral	967	1.93	1.35	1	11
Share loans backed by both fixed and differential contracts	967	0.48	0.499	0	1
Panel B: Contracts					
Fixed Price Contract	967	0.448	0.498	0	1
Price on Shipping Date / Futures Price Shipping Date (when contract was signed)	967	1.19	0.278	0.666	2.049
Contract Matures During Harvest Season	1,228	0.22	0.41	0	1
Futures Price Shipping Date (when contract was signed)	1,228	151.23	42.67	64.3	301.99
Price when contract matures	1,228	167.94	55.12	52.81	309.94
Year Contract Signed	1,228	2009	2.9	2000	2014
Year Contract Matured	1,228	2010	2.9	2000	2014
Panel C: Firms					
Number of Loans from Lender	272	3.6	2.86	1	12
Assets (1,000 USD)	113	2,035	2,954	9.24	17,894
Sales (1,000 USD)	106	3,713	5,278	28.6	39,677
Purchases (1,000 USD)	102	759	719	12.65	3,247
Sales / Cherry Purchases	102	3.77	1.08	2.26	11.604
Profit (1,000 USD)	106	56.4	30.78	34.9	260.9
Price paid to farmers (USD)	92	56.50	13.9	38.85	73.85
Growers Supplying Coffee	126	1,114	1,817	1	12,455
Share of Purchases Financed by Lender	102	57%	29%	5%	100%
Panel D: Buyers					
Number of Clients	102	1.86	2.07	1	11
Number of Loans	102	7.15	17.1	1	145
Dollars Guaranteed (\$1,000)	102	162	504	4	5,030
Share of Loan Guaranteed	102	51%	26%	4%	100%

Note: Data is presented at four levels: the loan, the contract, the mill and the buyer. There can be several contracts backing a single loan, because mill's sign contracts with different buyers, and sign contracts of different types (fixed / differential). There are 1,228 observations of this type. Sometime the contract information is missing. This typically happens when the buyer and the mill have only signed a promissory note or a letter of intent. In these cases, e.g., the scheduled shipping date could be missing resulting in fewer observations. While most analysis in the paper requires shipping information, we also do perform our main tests at the loan level using the loan maturity date, which is never missing. At the loan level we have 967 observations. Unfortunately, detailed scorecards for loan applications were introduced by the lender only later in the sample. As a result, we have fewer loans that have a credit score (previously the lender used a letter system only). The detailed scorecards are also our main source of information for mill level characteristics, since they include financial audits and statements submitted by the mill during the application process. This data is again available for the later part of the sample. Furthermore, the financial data is backwards-looking and can only be matched to a loan-year when the mill receives another loan within the next 3 years. See the data appendix for more details. Within mill's for whom we can match to financial statements, observations vary due to reporting inconsistencies.

Table 2: Strategic Default I: Unexpected price increases and defaults on loans

Dependent Variable:	Default or 90+ days late on repayment							
	Contract Level						Loan Level	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Price Surprise	0.304** (0.121)	0.343** (0.154)	0.305** (0.139)	0.279** (0.122)	0.369** (0.171)	0.0360 (0.0679)	-0.0661 (0.0767)	-0.0253 (0.0875)
Fixed							-0.253** (0.111)	-0.288** (0.125)
Fixed x Price Surprise							0.196** (0.0907)	0.201* (0.103)
Sample	Fixed	Fixed	Fixed	Fixed	Fixed	Differential	All	All
Mill Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Buyer Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes
Month Fixed Effects	Yes	No	Yes	Yes	No	Yes	Yes	No
Country-Month Fixed Effects	No	Yes	No	No	No	No	No	Yes
Year-Month Fixed Effects	No	No	No	No	Yes	No	No	No
Length of Loan Control	No	No	No	Yes	No	No	No	No
Spot and Future Price	No	No	Yes	No	No	No	No	No
Number of observations	434	434	434	434	434	533	967	967
R^2	0.495	0.621	0.499	0.502	0.664	0.427	0.387	0.479

Notes: Regressions are at the contract level or the loan level. At the loan level we sometimes have loans with both fixed price and differential price shipments, so we define a loan to be a ‘fixed price loan’ if more than half of the sales (in dollars) come from fixed price shipments. In all cases our dependent variable is default or severely late payments, where lateness is defined as being at least 90 days past due. Price surprise is defined as being the price at the time of shipment is due divided by the futures price for that time at the time the agreement was made. At the loan level we use the maturity date instead of the shipment date to determine the price surprise since there are typically several shipments financed by a loan. Appendix E reports robustness checks on this Table varying both the definition and the thresholds to assign loans to fixed contracts in Columns 7-8. Standard errors are clustered at the mill level. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

Table 3: Strategic Default II: Unexpected out of season price increases and defaults

Dependent Variable:	Default or 90+ days late on repayment				
	Fixed Price			Differential Price	Fixed Price
In / Out of Harvest Season	Out			In	
Event Window:	2-weeks	1-week	3-weeks	2-weeks	
	(1)	(2)	(3)	(4)	(5)
Shipment Scheduled After Price Jump	0.143*** (0.0132)	0.118*** (0.00352)	0.105*** (0.0387)	-0.00479 (0.0584)	0.0438 (0.0856)
Control Group Mean of Dependent Variable	0.055	0.005	0.074	0.065	0.091
Observations	123	70	154	150	72
R-squared	0.026	0.044	0.015	0.000	0.002

Notes: Local linear regressions are executed at the contract level. In all cases our dependent variable is default or severely late payments, where lateness is defined as being at least 90 days past due. All regressions use an event study methodology, where an event is defined as a weekly price increase of at least 3%. Appendix E reports further robustness checks. Standard errors are clustered by event-day bins. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

Table 4: Punishment: How does the lender respond to strategic default?

Dependent Variable:	Indicator for observed future business between the mill and the lender or the buyer							
	Punishment by the Lender				Punishment by the Buyer			
Definition of Default (months late):	3+ months (1)	5+ months (2)	7+ months (3)	9+ months (4)	3+ months (5)	5+ months (6)	7+ months (7)	9+ months (8)
Default	-0.117** (0.0519)	-0.222*** (0.0630)	-0.244*** (0.0673)	-0.277*** (0.0687)	-0.0754* (0.0446)	-0.0950* (0.0492)	-0.114** (0.0529)	-0.119** (0.0527)
Mill Fixed Effects	No	No	No	No	No	No	No	No
Buyer Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	434	434	434	434	434	434	434	434
R ²	0.373	0.386	0.387	0.392	0.245	0.245	0.246	0.247

Notes: The unit of observation is a contract. We focus on fixed price contracts for which we have documented strategic default (unreported results show that default on differential contracts is less severely punished). The dependent variable in each regression is a binary indicator for whether the lender ever lends again to the mill, or whether the buyer-mill pair appear again within the time-period covered by the data. Mill fixed effects are not included from this model because defaults are most common among borrowers that obtain their first loan from the lender (consistently with the punishment we are trying to detect). Those observations would not contribute to the estimate if mill fixed effects were included, but are nevertheless of considerable interest for the purpose of the exercise. All regressions are estimated using a linear probability model. Standard errors are clustered at the loan level. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

Table 5: Heterogeneity: Strategic Default by Relationship Strength

Dependent Variable: Relationship Type: (Proxy): Sample:	Default or 90+ days late on repayment											
	Lender - Mill (Lending History)			Buyer - Mill (Relationship fit (AKM))			Buyer - Lender (Buyer network centrality)					
	Fixed price & Low Value	High Value	Differential Low Value	Fixed price & Low Value	High Value	Differential Low Value	Fixed price & Low Value	High Value	Differential Low Value	Fixed price & High Value	Low Value	Differential Low Value
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)				
Price Surprise	0.667* (0.402)	0.133 (0.128)	-0.0414 (0.129)	0.645** (0.279)	0.236 (0.164)	0.0697 (0.0986)	0.484*** (0.184)	-0.0889 (0.213)	0.107 (0.0743)			
Mill Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Buyer Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	204	230	229	221	213	276	300	134	366			
R-squared	0.599	0.535	0.699	0.614	0.732	0.649	0.545	0.590	0.572			

Notes: The unit of observation in the table is a contract. The table examines heterogeneity by relationship value. For each dimension examined, a binary variable is constructed to split the sample into high and low (high is always any value greater than the median), and a separate regression is run on each of those samples. For the relationship between the mill and the lender, we use the cumulative dollars loaned to date between the lender and the mill as a proxy for the relationship value. For the relationship between the mill and the buyer, we use a function of the age of the observed mill-buyer relationship. For the relationship between the lender and the buyer we use the eigencentality of the buyer in the mill network of the lender. This captures the idea that if the lender defaults on the buyer then the buyer can get financing from another firm, and will take either many or few mills with them. See figure A9 for a visual depiction of the lender's network of mills. Standard errors are clustered at the loan level. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

Table 6: Contract Selection: Selection of fixed versus differential contracts

Dependent Variable:	Fixed price contract						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Age of Relationship (1000 days) (AKM)	0.136** (0.0610)	0.145** (0.0693)	0.137** (0.0609)	0.137** (0.0615)			
Centralization of loan collateralization (more money by fewer buyers)					0.237*** (0.0833)		
Cumulative history with buyer (\$1,000,000)						0.00372*** (0.000560)	
Lender's Perception of Mill-Buyer relationship (from credit score)							0.209** (0.102)
Mill Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Buyer Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	No	Yes	Yes	No	No	No
Country-Month Fixed Effects	No	Yes	No	No	Yes	Yes	Yes
Length of Loan Control	No	No	No	Yes	No	No	No
Spot and Future Price	No	No	Yes	No	No	No	No
Observations	967	967	967	967	967	967	331
R-squared	0.258	0.291	0.260	0.261	0.297	0.283	0.353

Notes: The unit of observation in these regressions is a contract. Notably the sample in column 7 is different. The lender score of the buyer-mill relationship comes from an addition to the credit application process that was added part-way through the sample, and it is only observed for more recent loans. The main proxy for relationship value is the age of the relationship, defined in the same way we define relationship value in table 5. Standard errors are clustered at the mill level. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

Table 7: Calibration: Inputs

Panel A: Parameters						
Parameters	Values					Source
Price Surprise $F()$	Log-Normal, $\mu = 0.0152$ and $\sigma = 0.225$					Data
Mill's Risk Aversion	$\alpha = 0.386$					Data (Calibration)
Farmers' Supply Elasticity	$\eta = 0.6$					RDD

Panel B: Loan Specific Values						
	N. Obs.	p25	Median	Mean	p75	Source
Input: Cost (γ_c)	307	3.12	4.64	6.20	8.12	Data (Calibration)
Target: Interest Rate (r_c)	307	9%	9.5%	9.6%	10%	Data

Notes: The Table reports the inputs for the calibration exercise in Section 5. Panel A reports the parameters that are common to all observations. The distribution of price surprises $F(p_w)$ is directly observed in the data and is well approximated by a Log-Normal with mean μ and standard deviation σ . The mill's utility function is assumed to be $u(x) = x^{1-\alpha}$. The parameter α is calibrated from the data to match the average advance purchase discount implied by the distribution of price surprises $F(p_w)$. Specifically, α is chosen so that a mill would be indifferent between an uninsured random draw from $F(p_w)$ (with expected price $\bar{p}_w > 1$) and a fixed price contract with price 1. The farmers' supply elasticity η is estimated from the RDD estimates in Appendix F. Specifically, we estimate the effect of a larger loan on the amount of cherries purchased and the unit prices paid to farmers. The two effects combined identify the slope of the supply curve. Panel B focuses on the loan-specific parameters and target. The loan-specific cost parameter γ_c is directly inferred from the audited financial accounts. Knowledge of η , production volumes q_c and cost of row material $C_c(q_c) = \gamma_c \times q_c^{1+\eta}$ allows us to directly compute γ_c for all observations for which we have audited financial accounts. The value of the relationship \mathbf{V} is then backed out for each loan. Specifically, given a set of parameters we find the \mathbf{V}_c that rationalizes a loan's key contractual outcomes: the interest rate r_c and whether the loan is backed by a fixed or a differential contract.

Table 8: Calibration: Results

Panel A: Baseline					
Variable	N. Obs.	25th pctl.	50th pctl.	Mean	75th pctl.
Relationship Value (\mathbf{V}_c)	307	34%	44%	158%	133%
Output Loss $X^* = (1 - q_c/q_c^F)$	307	0%	19.7%	15.8%	19.9%
Output Loss X_F^*	108	0%	0%	11.3%	32.8%
Output Loss X_D^*	199	19.7%	19.8%	18%	20%
Output Loss X_D^D	199	0%	0%	1%	0.1%
Output Loss X_D^F	199	51.2%	55.2%	53.3%	55%
Wedge ($MPKV_c - r$)	307	0%	0%	6%	4%
Wedge (if > 0)	112	4%%	8%	20%	16.7%

Panel B: Robustness to Risk Aversion (α)					
Moment	$\alpha = 0.286$	$\alpha = 0.336$	$\alpha = \mathbf{0.386}$	$\alpha = 0.436$	$\alpha = 0.486$
Output Loss X^* (Mean)	11.6%	11.5%	16%	16.4%	15.6%
Output Loss X^* (St. Dev.)	10.5%	10.9%	12.2%	12.3%	13.4%

Panel C: Robustness to Farmers Supply Elasticity (η)					
Moment	$\eta = 0.50$	$\eta = 0.55$	$\eta = \mathbf{0.60}$	$\eta = 0.65$	$\eta = 0.70$
Output Loss X^* (Mean)	17.6%	16.2%	16%	13.8%	11.6%
Output Loss X^* (St. Dev.)	13.6%	12.6%	12.2%	11.9%	12.4%

Notes: The Table reports the results for the calibration exercise in Section 5 (see Appendix A for details). Panel A reports the baseline results with the parameters described in the previous Table. The value of the relationship \mathbf{V}_c is backed out for each loan by solving the model matching the observed interest rate and contract type in the data. The result is then scaled down by a factor of 1.64 in accordance with the market liquidity τ and punishment parameter λ as described in Appendix A. The output loss X^* computes the percentage deviation between the predicted production at \mathbf{V}_c and the first best quantity q_c^F . Output loss X_F^* is for loans predicted to be on fixed price contracts only. In this case if there is an output loss, it arises due to credit constraints. X_D^* is the output loss for loans predicted to be on differential contracts. This output loss can be decomposed into two: the gap relative to the optimal quantity conditional on a differential contract (X_D^D) and the predicted gap if that relationship had a fixed price contract instead (X_D^F). Wedge refers to the difference between the lender risk free interest rate (set at $r = 0.08$, the lowest interest rate contracted by the lender over the relevant sample period) and the predicted physical marginal product of capital (MPK). This is obtained by solving for the model in a counterfactual scenario in which the mill has all parameters fixed and is endowed with a small amount of liquidity. Panel B and C explore the robustness of the results to changes in risk aversion α and coffee cherries supply slope η . The Table focuses on those two parameters as those are either calibrated (α) or estimated (η). The other key parameters are directly observed in the data.

A.1. *Remarks on the Model*

Proposition *Under perfect contract enforcement the mill offers a fixed price contract with price $p_c^* = \bar{p}_w$ and produces quantity $q_F^* = (\bar{p}_w/(1 + \eta)\gamma)^{1/\eta}$. This quantity is larger than the one the mill would optimally produce with a differential contract q_D^* .*

Proof: Consider a fixed price contract. Standard arguments imply that both the buyer's and lender's participation constraints bind. Expected profits of the mill are equal to $\mathbf{E}[\Pi] = \int_{p_w} u(\bar{p}_w q_c - C(q_c)) dF(p_w)$. Taking first order condition, $\bar{p}_w = C'(q_c^*)$ establishes the result. We now show that a differential contract does (weakly) worse. This is intuitive given *i*) the mill is risk averse, and *ii*) lender's and buyer's participation constraints bind. But we also prove that a differential contract - if chosen - leads to a quantity always lower than the optimal one.⁵⁵ The buyer's participation constraint is now $\Delta_c = 0$. Defining $\tilde{p} = \frac{D}{q}$ and setting up the Lagrangian we obtain

$$\max_{q,D} \int_{p_w} u(p_w q - D) dF(p_w) + \zeta \left(\int_0^{\tilde{p}} (p_w q_c - D) dF(p_w) + D - C(q) \right).$$

The two first order conditions are given by

$$\begin{aligned} \int_{\tilde{p}} u'(\cdot) p_w dF(p_w) + \zeta \left(\int_0^{\tilde{p}} p_w dF(p_w) - C'(q) \right) &= 0 \\ - \int_{\tilde{p}} u'(\cdot) dF(p_w) + \zeta (1 - F(\tilde{p})) &= 0 \end{aligned}$$

Substituting for ζ into the first condition we obtain

$$\int_{\tilde{p}} u'(\cdot) p_w dF(p_w) + \frac{\int_{\tilde{p}} u'(\cdot) dF(p_w)}{(1 - F(\tilde{p}))} \left(\int_0^{\tilde{p}} p_w dF(p_w) - C'(q) \right) = 0.$$

To establish the result it suffices to show that at $C'(q) = \bar{p}_w$ the expression

⁵⁵This is useful in the calibration to set the extreme points over which the algorithm solves for the optimum.

above is *always* negative.⁵⁶ The expression above is smaller than

$$\frac{\int_{\tilde{p}} u'(\cdot) p_w dF(p_w)}{(1 - F(\tilde{p}))} - \frac{\int_{\tilde{p}} u'(\cdot) dF(p_w)}{(1 - F(\tilde{p}))} \left(\frac{\int_{\tilde{p}} p_w dF(p_w)}{(1 - F(\tilde{p}))} \right)$$

and this is negative since the covariance of $u'(\cdot)$ and p_w is negative for $u'' < 0$. ■

A.2. Incentive Compatibility for the Calibration

In calibrating the model we actually use a slightly more elaborate incentive compatibility constraint that allows us to take full advantage of our definition of default and observed gradient in punishment (Table 4). Specifically, we assume that in trying to side-sell coffee, the mill searches for an alternative buyer willing to pay price p_w and finds one with probability τ . If the mill does not find the buyer, it repays the loan late. This allows us to distinguish between three continuation values: \mathbf{U}^R , when the mill repays; \mathbf{U}^D when the mill defaults and \mathbf{U}^L when the mill is late. We assume that the only punishment available to the buyer and the lender is to discontinue the relationship. In line with the evidence in Table 4, we assume $\mathbf{U}^D = \mathbf{U}$ and $\mathbf{U}^L = \lambda \mathbf{U}^R + (1 - \lambda) \mathbf{U}^D$. That is, the punishment that follows a late payment is in between the punishment that follows a default and the continuation value following repayment. The assumption implies that *i*) the continuation value following a late payment only depends on whether the loan is renewed or not and *ii*) conditional on no future loan, the continuation value does not depend on loan default.

With this set up, the mill will look for an alternative buyer willing to buy at spot price p_w (and default if such a buyer is found) if

$$u(\max\{p_c q_c - D, 0\}) + \delta \mathbf{U}^R \leq \tau (u(p_w q_c) + \delta \mathbf{U}^D) + (1 - \tau) (u(\max\{p_c q_c - D, 0\}) + \delta \mathbf{U}^L)$$

which, using our notation $\mathbf{V} = \mathbf{U}^R - \mathbf{U}^D$, can be rewritten as

$$\delta \mathbf{V} \leq \frac{\tau}{(1 - (1 - \tau)\lambda)} (u(p_w q_c) - u(\max\{p_c q_c - D, 0\})).$$

This constraint is identical to the one in the main text, except for the fact that the value of the relationship \mathbf{V} is scaled by the factor $\frac{1 - (1 - \tau)\lambda}{\tau}$. In the calibration exercise we apply this scaling to the estimated \mathbf{V} .

⁵⁶From the expression above it is easy to verify that a risk neutral mill is indifferent between the two contracts at the optimum.

A.3. Details of Calibration

The parameters to be calibrated are: price surprise distribution $F(p_w)$; risk aversion α ; slope of supply curve η ; loan specific cost parameter γ_i ; market thickness τ and differential punishment λ . For expositional simplicity, the last two parameters only appear in the Appendix incentive constraint. The parameters are calibrated as follows:

$F(p_w)$: The price surprise distribution $F(p_w)$ is directly observed in the data. Given the unit of observation for the calibration is a loan, we take price surprises over the period in between the loan closing date and the loan maturity date. Price surprises are defined as the ratio between the spot price at maturity and the futures price for the maturity date at closing. We consider loans of regular length, between six and nine months. The empirical distribution has mean 1.05 and variance equal to 0.107. The residual variance drops to 0.047 after controlling for month, year and country of closing fixed effects. The empirical distribution is well approximated by a log-normal. We therefore assume the price surprise distribution $F(p_w)$ to be lognormal and estimate its shape and scale parameters $\mu = 0.0152$ and $\sigma = 0.225$ respectively.

α : We assume a utility function given by $u(x) = x^{1-\alpha}$. We calibrate α to match the average forward discount in the data. That is, we assume that the risk averse mill is indifferent between a random draw from the price distribution $F(p_w)$ and a sure payoff equal to the current spot price. Suppose the mill sells Q units of coffee. Then α solves $\int_{p_w} (p_w Q)^{1-\alpha} dF(p_w) = (1 \times Q)^{1-\alpha}$ where the normalization of the current spot price to 1 comes directly from the definition of price surprise p_w . The advantage of the chosen functional form is the parameter α does not depend on the quantity sold Q . We recover α by solving this equation with the command `fsolve` in Matlab using the calibrated parameters μ and σ for the distribution $F(p_w)$ and an initial search value 0.111.

η : The slope of coffee cherries supply curve is recovered from the RDD estimates in Appendix F. For those mills that are credit constrained, the RDD identifies the effect of additional \$100,000 loan L . Total costs of purchasing cherries from farmers is $C(q) = \gamma_i q^{1+\eta}$. Taking logs, the RDD estimates gives $\frac{\partial \ln C(q)}{\partial L} = (1 + \eta) \times \frac{\partial \ln q}{\partial L}$. The estimate is $\frac{\partial \ln C(q)}{\partial L} = 0.34$ (Table A10, Column

2). The unit price paid to farmers is $\omega = \rho q^\eta$. Taking logs, the RDD estimates gives $\frac{\partial \ln \omega}{\partial L} = \eta \times \frac{\partial \ln q}{\partial L}$. The estimate is $\frac{\partial \ln \omega}{\partial L} = 0.13$ (Table A10, Column 4). Combining the two we obtain $\eta = 0.13 / (0.34 - 0.13) \approx 0.6$.

γ_i : We let the cost parameter γ_i vary by loan. The parameter is directly reported in the financial accounts of the mill. The operating costs take the form $C(q) = \gamma_i q^{1+\eta}$. The operating costs and production volumes q_i are directly observed in the financial records. Knowledge of η allows us to assign γ_i to each loan for which financial accounts are available. Note that because we need the financial statement data to construct we can construct it only for the small sample of loans for which we have the financial data.

τ : The market thickness parameter τ is given by the probability that the mill defaults conditional on either being late or defaulting. This is directly observed in the data and is equal to 0.39. We approximate τ at 0.4.

λ : The punishment parameter λ is identified by looking at the differential punishment depending on the severity of the default. In particular, denote with U^R the value of not defaulting and normalize to zero the value of discontinuing the relationship. The estimate in column 1 of table 4 gives $U^L = (1 - 0.117)U^R$. Column 4 in the same Table implies outright default is punished more harshly: $U^D = (1 - 0.28)U^R$. Combining these two estimates with the definition of $U^L = \lambda U^R + (1 - \lambda)U^D$ we obtain $\lambda = 0.57$. The estimated V_i are therefore scaled by a factor $\frac{1-(1-\tau)\lambda}{\tau} = 1.64$.

Given a set of (loan specific) parameters and a candidate \mathbf{V}_i , the model predicts an interest rate $\hat{r}_i(\mathbf{V}_i)$ and a probability that a fixed contract is chosen, $\hat{\phi}_i(\mathbf{V}_i)$. Specifically, denoting with $\mathbf{EU}^C(V_i)$ the predicted expected utility under contract type $C \in \{F, D\}$ we let $\hat{\phi}_i(\mathbf{V}_i) = \frac{\mathbf{EU}^F(V_i)}{\sum_{C \in \{F, D\}} \mathbf{EU}^C(V_i)}$. For each loan i the value \mathbf{V}_i is estimated as

$$\mathbf{V}_i \in \arg \min \frac{(\hat{r}_i(\mathbf{V}_i) - r_i)^2}{\sigma_r^2} + \frac{(\hat{\phi}_i(\mathbf{V}_i) - C_i)^2}{\sigma_C^2}$$

where σ_r^2 and σ_C^2 are the population variances of the loan interest rate and contract types.

APPENDIX B. DATA SOURCES (FOR ONLINE PUBLICATION)

B.1. Contract Data at the Shipment Level

Besides detailed information on each loan (borrower, size of the loan, contracting date, maturity date, collateral, interest rate, final repayment status, etc.) the lender provided us with the contracts made between the buyer and the mill. These contracts are boiler-plate, and typically include the buyer's name, the mill's name, and each promised shipment from the mill to the buyer. For each shipment the contracts list the date of delivery, quantity, price, and price-type. In most cases if the price is fixed for one shipment on the contract it is fixed for all, but there are some contracts that are mixed: where some shipments are at a fixed rate and others use a differential rate. We use these contracts to construct a shipment level dataset.

These files came in PDFs so they had to be coded as well. This was done using a similar process as the one outlined above. We wrote a text-analysis script in Python to scrape every contract that we had, to construct a dataset with all of the information we were interested in. We then had a research assistant manually check 20% of the sample randomly for errors. In this case though, because of the consistency of how the contracts are written, there were almost no errors found by the manual check so we decided not to enter any of the contract information by hand.

B.2. Transaction Level Data

We also received from the lender a file that outlined every transaction they made over the sample period. This file included a loan ID, a dollar amount either sent out or repaid, the identity of both the sender and receiver of the transfer, and the date of the transfer. From this we can infer default. Due to the nature of the agreements, overwhelmingly the buyer repays the lender (see figure A7). From the financial transactions we can also see that buyer typically repays the lender on the delivery date.

The transaction data provides us with the identity of the party repaying the loan, as well as the date and the amount of repayment. This is helpful because it could be that the mill is able to sell their coffee on the world market to a different buyer on the day of the original shipment, and tries to repay the lender directly, going around the buyer. We are able to match repayments by the buyer on the loan, and identify whether and when each specific buyer on a contract repaid

their portion of the loan. Sometimes the buyer never repays the loan, but instead the mill repays the lender directly. This happens very infrequently, as the lender does not want to facilitate default on the buyer, especially since there are very few large buyers (figure A9 shows the network map between buyers and mills). However, this does occasionally occur. When the mill does this on within 90 days of the scheduled shipment we do not observe this as default, which may introduce measurement error into our default measure by underestimating the true default rate. However, our results are robust to using direct repayments by the mill as a definition of default, under the presumption that in this case the mill side-sold the coffee at a better price, risking the relationship with the buyer, but not wanting to risk the relationship with the lender (See Appendix D).

B.3. Application Data and Financial Data

The lender's files made available to us include information from all applications, including income statements and balance sheets, both of which are typically audited by the lender. This financial data comes at the mill-year level. We only have a subset of mills with this data. The detailed scorecards were kept in organized soft copies only for the later part of the sample. Furthermore, because the information is collected at the time of application, and is therefore backward-looking. Financial statements are typically collected for the three previous years, so whenever a mill signs another loan in the three years after receiving a loan, we can match the financials for the year of the loan to the loan.

The application data also includes all the information from the scoring model that is used to determine the size of the loan. We received spreadsheets that provide scores on a number of elements, such as liquidity risk, history with the lender, relationship with the buyer, environmental practices, etc. All of these sub-scores are aggregated by the lender into an overall score based on a weighting scheme. The overall score is aggregated further into a letter grade. The spreadsheets provide all of the sub-scores, scores and formulas used in aggregation. In addition to this, they include the terms of the loan given to the mill, and often include general background information about the mill itself, such as location, number of employees, management history, etc.

These spreadsheets were used to construct a mill-year panel that includes relevant expenses, sales, existing loans, credit scores and sub-scores. This involved a three-step process. First, a Python-script was written to scan and pull all of the relevant information from the spreadsheets. Second, a research assistant pulled

20% of the sample at random to check for systematic errors. Through this process we found that the Python code was not accurately capturing some of the financial statement data (but did capture loan and credit score data very well) due to inconsistencies in the way it was entered into the spreadsheet. So as a third step, the financial statements were manually coded. Still, there were some minor inconsistencies. For example, in a given spreadsheet, statements are typically provided for the past three years, and we often see the same firm apply in back to back years, meaning that there are two years of overlap in data. In a few cases these data did not agree. In these cases we first prioritized data that had been audited, and if there were still disagreements, we used the more recent file.

B.4. World Price Data

Finally, all of this data is matched to world coffee prices. We collect data on spot prices as well as futures prices for the closing date on the contract, the shipment date and the maturity date of the loan. Futures prices for the date of shipment at the closing date are used to control for expected price changes in each regression that relies on world price changes. One issue that should be noted is that while we have dates for the closing dates and maturity dates of the contracts (which come from the lenders spreadsheets) for every loan, we only have shipping dates (which come from the buyer-seller contracts) for a subset of loans. This means that we can only construct the price at shipping for about 70% of the sample. Our approach is to use this more limited sample whenever the analysis requires the shipping price information, but to otherwise rely on the full sample.

APPENDIX C. ADDITIONAL DESCRIPTIVE TABLES AND FIGURES (FOR ONLINE PUBLICATION)

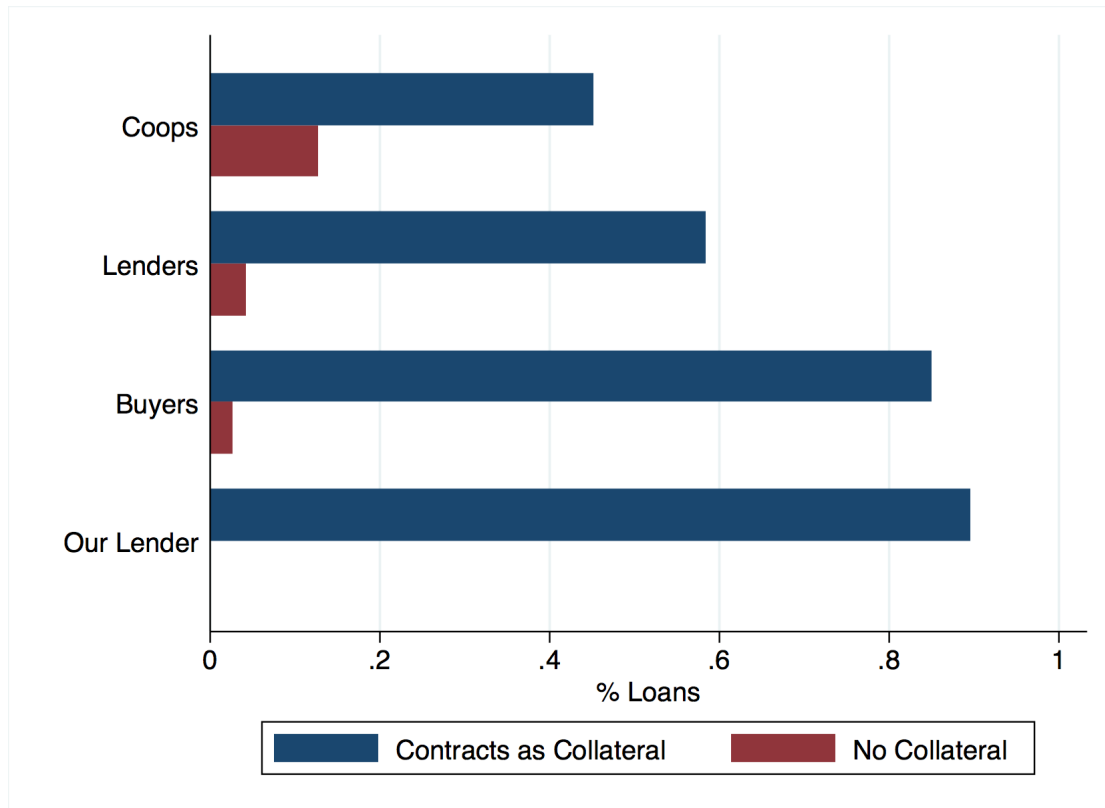


Figure A1: Use of Collateral

Notes: This Figure describes the use of collateral by different types of lenders. The Figure plots the fraction of loans that use collateral, and the fraction that use forward sales contracts as collateral. The Figure distinguishes loans from our lender as well as other sources of loans: upper-tier cooperatives (e.g., federations of cooperatives); other financial institutions (labelled as lenders) and buyers. The Figure confirms that forward sales contracts are the dominant form of collateral for working capital loans for all types of lenders in the industry.




Figure A2: Representativeness of Lender’s Interest Rates

Notes: This Figure describes interest rates. For many borrowers, the data include information about working capital loans extended by other lenders. The Figure reports the distribution of interest rates on working capital loans. On average, our lender charges interest rates that are nearly identical to those charged by other lenders.

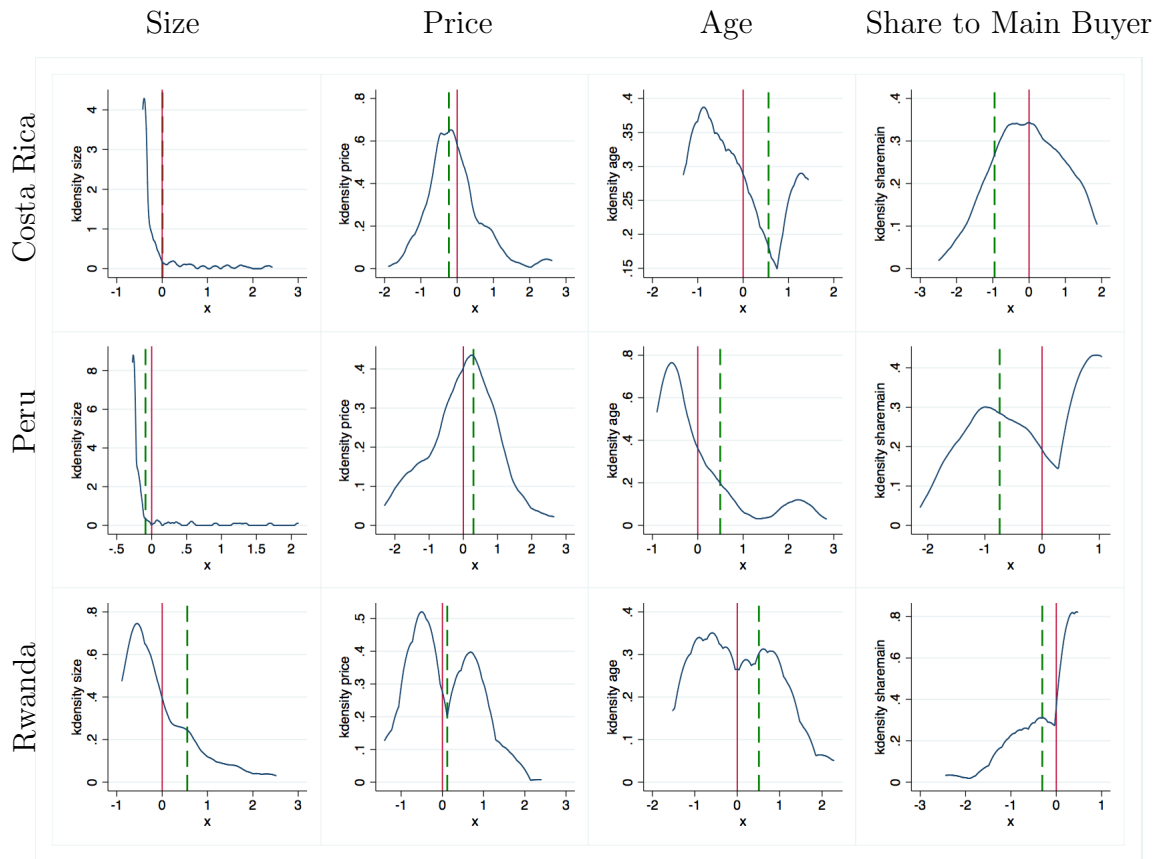


Figure A3: Representativeness of Lender's Portfolio

Notes: The lender extends working capital loans across the world in several countries. For three countries (Costa Rica, Peru and Rwanda) we have data on the universe of coffee washing stations and/or exporters from other projects. In the three countries we can compare the mills in the lender's portfolio to the rest of the industry. We focus on four variables: size, price, age and share sold to the main buyer. For each variable and country the figure reports the standardized distribution (centered at the mean) and the mean for clients in the lender's portfolio (represented by the dashed-vertical line). The figure shows that the clients in our lender's portfolio are broadly representative of the industries in these three countries. Consistently with the loans from our lender relaxing the mill dependency on buyers for finance, in each country the clients in our lender's portfolio sell a lower share of their produce to their main buyer.

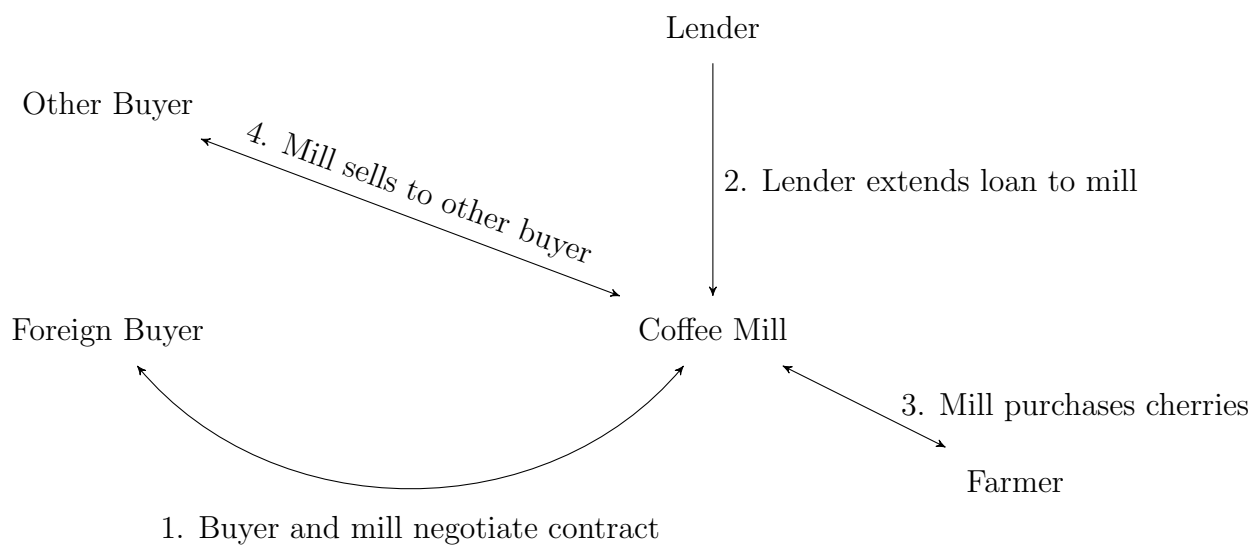


Figure A4: Lending Model With Default

Notes: This figure shows the lending model that the lender uses when the mill defaults (see ‘normal’ case in figure 2). In this case the mill and the buyer agree on a contract at the beginning of the harvest season which sets a price and quantity of coffee to be delivered by the mill at a specific future date. Using this contract as collateral, the mill then secures a loan from the lender. The loan amount is based on a formula which sets a fraction of the value of the contract, which varies based on a credit score received by the mill during the application process. The mill uses the loan money to purchase coffee cherries from farmers, they process the cherries and deliver the agreed upon quantity to a different buyer who presumably pays more than the agreed upon price. In this case the lender typically refuses side payments since they value their relationship with the buyers and do not want to facilitate default on them.



Figure A5: Lender’s Cash Flow Profile

Note: The graph plots the density of closing months and maturity months for the lender’s portfolio. The graph shows two peaks in both closing and maturity, reflecting the asynchronous harvest seasons across countries located in different emispheres.




Figure A6: Histogram of Unobserved Relationship Effects from AKM-style Regression

Note: This figure plots the distribution of fixed-effects estimated from a regression of relationship age on fixed effects for the year the contract was agreed to; the month the contract was agreed to; the year-quarter of the first observed buyer-mill transaction; and the buyer-mill pair. We use buyer-mill pair fixed-effect as a proxy for relationship value, as it captures any conditional unobserved heterogeneity in the buyer-mill pair that influences the length of their relationship (a pair-specific measure of fit).

Table A1: Summary Statistics by Country

	% Lender Portfolio (1)	Year of Entry (2)	N. of Mills (3)	N. of Loans (4)	% Fixed (5)	% Default (6)
Panel A: Americas and Carribbean						
Nicaragua	34.46829	2002	25	106	0.331897	0.067953
Peru	22.02435	2002	58	301	0.308516	0.099759
Mexico	12.1392	2000	44	167	0.364185	0.0963391
Honduras	10.15018	2005	22	64	0.326754	0.0254777
Guatemala	6.74886	2000	24	91	0.335821	0.0509259
CostaRica	4.02112	2000	10	27	0.596244	0.0046296
Colombia	2.05848	2005	13	33	0.336842	0.1633663
Bolivia	1.52627	2004	9	27	0.461538	0.1027397
Ecuador	0.82804	2005	4	13	0.71134	0.1
ElSalvador	0.19217	2006	1	2	0.8	0.2
Haiti	0.08989	2010	2	4	1	0.125
Brazil	0.0471	2006	1	2	1	0
DominicanRepublic	0.00195	2012	1	1	0	0
Panel B: Africa						
Uganda	2.15768	2005	8	21	0.529915	0.2892562
Rwanda	1.55691	2004	34	77	0.753623	0.112782
Zambia	0.70918	2009	1	4	1	0
Tanzania,UnitedRepublicOf	0.67792	2008	4	9	0.969697	0.0606061
Congo,TheDemocraticRepublicOfThe	0.08	2013	3	3	0.083333	0
Ethiopia	0.07807	2005	2	3	1	0.9444444
Malawi	0.06168	2009	1	2	0.708333	0.0416667
SierraLeone	0.01569	2012	1	1	.	0
Kenya	0.01457	2005	1	2	.	0
Panel C: Southeast Asia						
Indonesia	0.34459	2006	2	6	1	0
EastTimor	0.00781	2006	1	1	.	0
Total	100	.	272	967	.	.

Note: This table shows the breakdown of loans by country. For each country we show the percentage of loans coming from that country; the year that the lender first agreed to make loans in the country; the number of mills they have ever lent to in the country; the number of loans they have ever made in the country; the fraction of contracts ever made that are fixed price; and finally the default rate, by country. The table is sorted by the importance of the country to the lender's portfolio, within each geographic region.

APPENDIX D. UNTANGLING THE LENDER FROM THE BUYERS (FOR ONLINE PUBLICATION)

The lending model prescribes that the loan is directly repaid by the buyer to the lender after the mills delivers the coffee. Assuming the buyer and the lender acts as a single entity, then, the mill must default on the sale contract with the buyer to default on the loan - and vice-versa. Because of this, we have bundled the buyer and the lender into a single entity for both the theoretical and baseline empirical analysis.

This Appendix explores the extent to which this assumption is warranted and its likely consequences for our analysis. We are interested in three questions: 1) is indeed the case that the majority of loans are directly repaid by the contracted buyer?; 2) how do the loans that are not directly repaid by the contracted buyer compare with those that are?; 3) how do the buyers that do not repay directly compare with those who do? The Appendix shows that our baseline assumption is justified and that, if anything, it might *underestimate* the extent of strategic default in this market.

We start by looking at who repays the loan to the lender. The lender's lending model specifies that the loan should be disbursed to the mill, who then delivers cherries to the buyer, who then repays the lender. In Figure A7 we see that indeed the vast majority of loans are repaid by the buyer initially on the contract. Nearly 90% of the time repayment is made by one of the buyers on the contract (rather than directly by the mill). In just over 10% of loans some portion of the loan is repaid either by the mill directly, or by another buyer that was not initially on the contract. Even in those cases, it is typically a very small portion of the loan that is repaid by an unexpected party - that is, is very rare for more than 50% of the loan to be repaid by an unexpected party.

This brings us to the second question: are loans directly repaid by the mill indicative of side-selling against the buyer? Figure A8 provides evidence that supports this conjecture. The vast majority of cases in which the mill directly repays the lender are i) on fixed price contracts, and ii) occur at times with large positive price surprise. That is, mills directly repay the loans *precisely* in those circumstances in which we expect them to have incentives to strategically default against the buyer. Note that a direct consequence of this observation is that our baseline analysis *underestimates* the extent of strategic default in the market.

Finally, this brings to our last question: if these loans are indeed cases of strategic default against the buyer, why does the lender accepts default on some

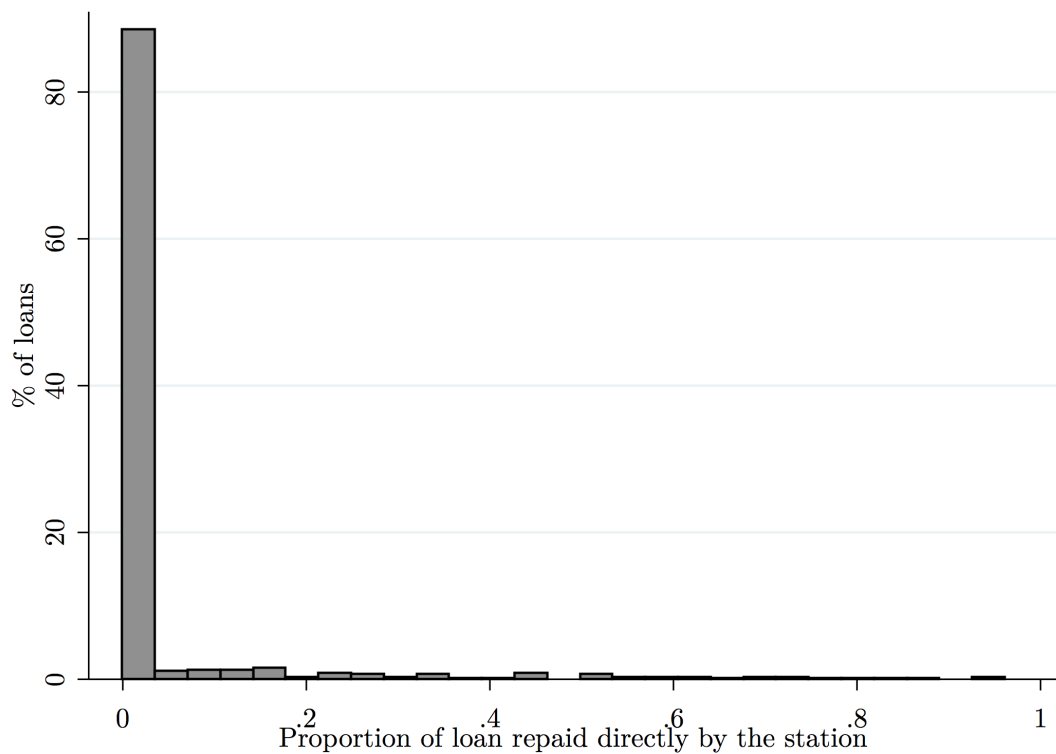


Figure A7: Histogram of Repayment Source

Notes: The figure simply shows the histogram of direct repayment by the mill. Direct repayment is defined as the repayment transaction listing the payee as the mill rather than the expected buyer. It shows that over 80% of the time none of the loan is repaid by the mill. This is consistent with the lenders description of the loans. Since the buyer always guarantees the loan, they repay it directly to the lender upon receiving delivery, and additional profits over and above the amount due to the lender is sent to the mill. The lender very rarely allows repayment directly from the mill since that would indicate a high likelihood of default on the buyer, who is integral to the operations of the lender.

buyers and not others? Indeed, the lender faces a decision as well. They can facilitate default on the buyer, but may then be punished by the buyer. In this case we might expect that the lender may be willing to recoup the loan and facilitate default on the buyer only if the buyer and lender do not have a strong relationship. We proxy for the lender-buyer relationship using a measure of network centrality, using the same measure as in table 5. Figure A10 shows that indeed the lender is most likely to accept direct payment from the mill whenever the buyer listed on the contract is less central in the lender's network of mills (see also Figure A9 for an illustration). This is at least suggestive evidence that actually these instances in which the buyer and lender might not act as a single entity could actually be strategic defaults (even though we do not observe them as such). Omitting these observations from our main default analysis provides a lower-bound on the prevalence of strategic default in the market (Prediction 1) and also works against

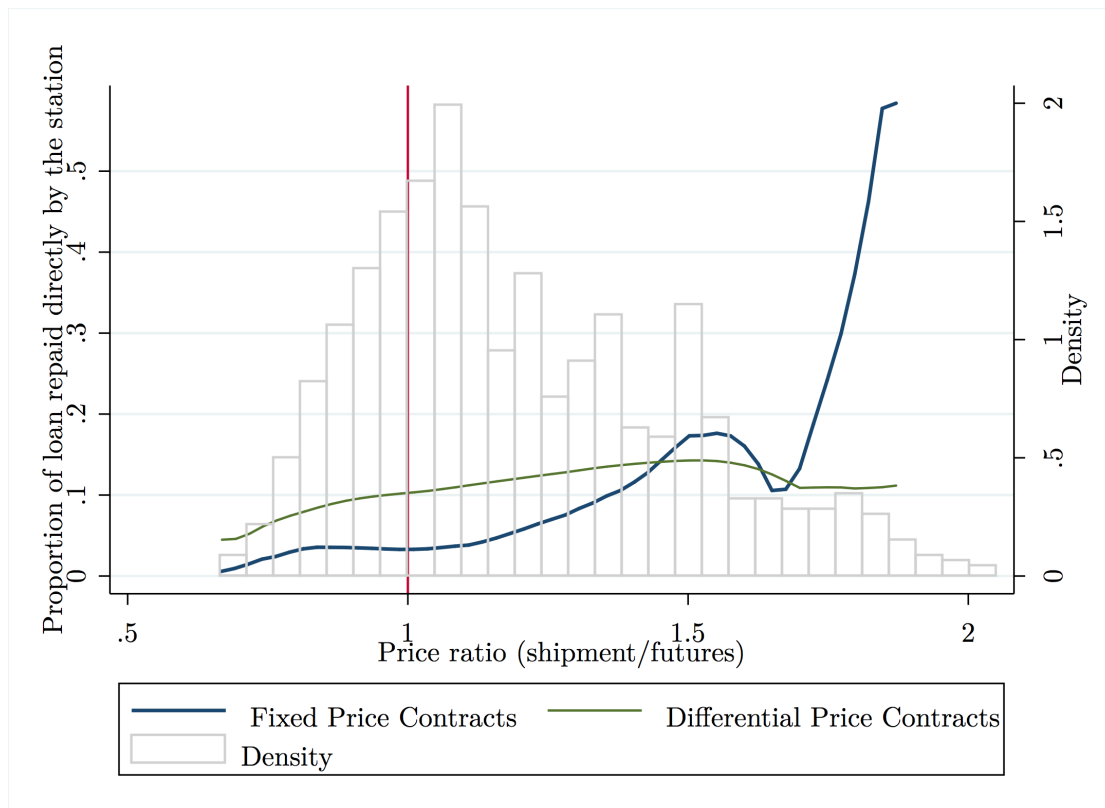


Figure A8: Repayment directly by mill

Notes: The grey bars indicate the frequency of a given price surprise (x-axis), which is defined as the price at the time of the scheduled shipment divided by the futures price for that date, at the time that the loan was signed. The blue line plots direct repayments on fixed price contracts for a given price surprise, while the green line plots the same thing for differential price contracts. Direct repayment is defined as the repayment transaction listing the payee as the mill rather than the expected buyer.

Typically the lender provides the loan to the mill, who uses the money to purchase cherries which are sold to the buyer, who then repays the loan to the lender directly and giving the difference to the mill. In the case of strategic default, the mill does not provide the buyer with the coffee, leaving the buyer unable to repay the lender. In some cases we see that the lender does accept direct payments from the mill, which would imply that the mill sold the coffee to a different buyer. We see that this is much more likely to occur after large price increases.

us finding heterogeneity in strategic default by relationship value (Prediction 2).

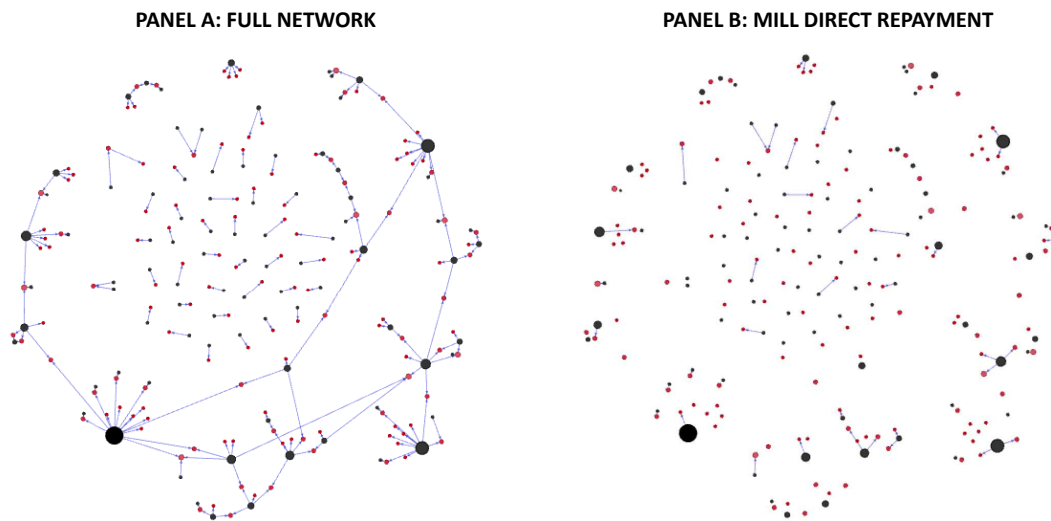


Figure A9: Network of buyers and mills

Notes: Panel A in the Figure describes the network of links between buyers and mills in the lender's portfolio. The darker dots (with radius proportional to the number of links) represent buyers whose contracts are used as collateral by the lender (the mill might have other buyers). The lighter dots represent mills. The Figure shows substantial heterogeneity in buyer's importance from the lender's perspective: certain buyers are central to the network of relationships and bring several suppliers into the lender's portfolio. Panel B only plots the buyer-mill links for which direct repayment from the mill is observed. The comparison between Panel A and Panel B reveals that direct repayment is predominantly observed for buyers that are not central in the lender's portfolio. The Figure shows why the lender does not want to accept payment directly from the mill. Direct payment from the mill facilitates default on the buyer. Buyers are extremely important to the business of the lender: relatively few of them recruit mills that require financing for the lender. The lender does not want to risk upsetting the buyer, because losing business from a buyer could mean losing multiple clients. Hence, it is costly for the lender to collude with the mill against the buyer. If the lender upsets the mill, e.g., by not allowing direct/partial repayment following the mill's default against the buyer, the lender will at most lose a single account from an unreliable client.

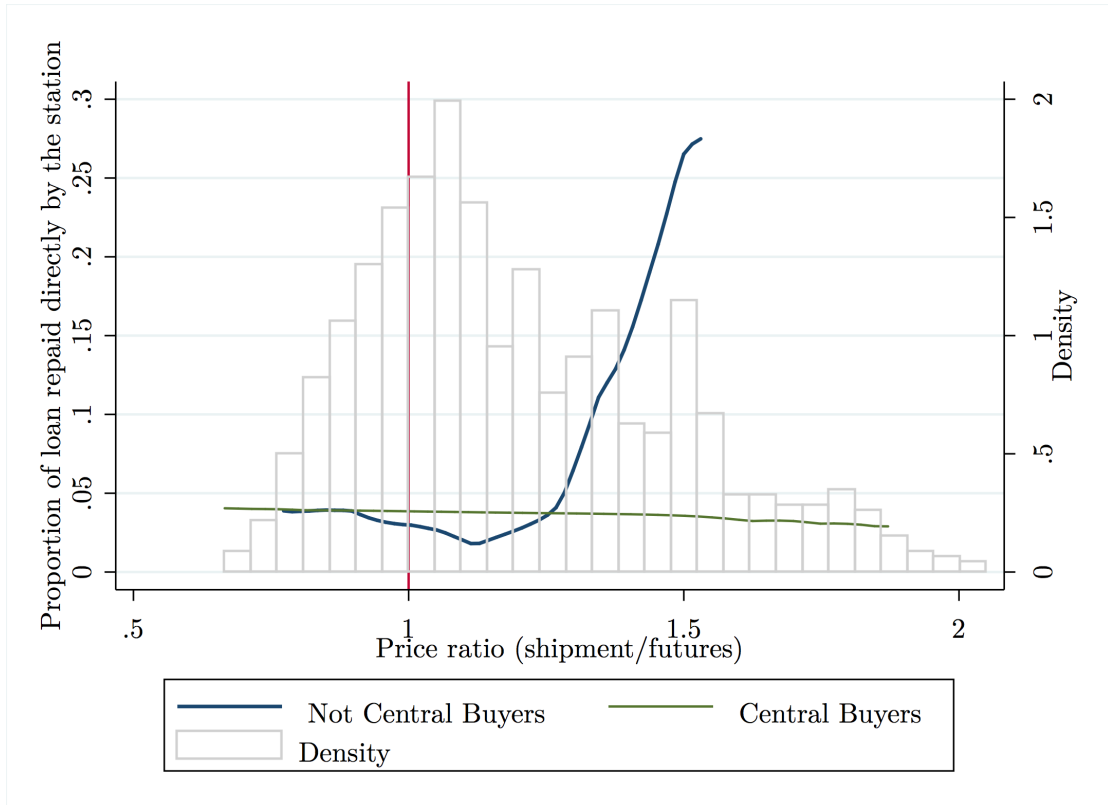


Figure A10: Repayment directly by mill

Notes: The grey bars indicate the frequency of a given price surprise (x-axis), which is defined as the price at the time of the scheduled shipment divided by the futures price for that date, at the time that the loan was signed. The blue line plots direct repayments for a given price surprise for buyers that are important to the lender, while the green line plots the same thing for buyers that are not. Importance of the buyer is defined by above or below median eigencentrality in the lender's network of buyers and mills. Direct repayment is defined as the repayment transaction listing the payee as the mill rather than the expected buyer.

APPENDIX E. ROBUSTNESS (FOR ONLINE PUBLICATION)

E.1. Robustness of the main result

This Appendix explores robustness of the main results to alternative definitions and specifications. We begin with results that test Prediction 1, then turn to results that test Prediction 3. Table 2 tests for Prediction 1. The table uses two main thresholds in the various definitions required. First, we define a default to be any loan more than 90 days late; second, we define a loan as a fixed price loan if more than 50% of the loan is fixed.

We start with the first threshold and look at a range of lateness around the 90 day mark. We show that the result is robust to a fairly wide range of definitions. If we use too lenient a definition, we may get many defaults that are simply due to regular delays, and that introduces measurement error into our default measure, while if we use too strict a default definition, we may run into a problem with very little variation within mills. However, table A2 shows that actually, a fairly large and plausible range of default definitions works. We look at 2, 3, and 4 months late.⁵⁷ Using all three definitions we see positive and significant estimates, consistent with table 2.

Table A2: Robustness of Table 2 to different definitions of ‘default’

	Threshold lateness for default	2 months	3 months	4 months
	(1)	(2)	(3)	
Price Surprise	0.268*	0.301**	0.176**	
	(0.140)	(0.121)	(0.0837)	
Mill Fixed Effects	Yes	Yes	Yes	
Buyer Fixed Effects	Yes	Yes	Yes	
Year Fixed Effects	Yes	Yes	Yes	
Month Fixed Effects	Yes	Yes	Yes	
N	434	434	434	
R^2	0.579	0.497	0.478	

Note: Regressions are at the contract level. In all cases our dependent variable is default or severely late payments, where lateness is defined as being at least 60, 90 and 120 days past due for columns 1, 2, and 3 respectively. Price surprise is defined as being the price at the time of shipment is due divided by the futures price for that time at the time the agreement was made. Standard errors are clustered at the loan level. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

The second dimension of robustness that we can examine is the definition of fixed price loan that we use in table 2, columns 7 and 8. At the contract level, every

⁵⁷We base a month here on a 31 day month, so that 2 months late is 62 days, 3 months is 93 days and 4 months is 124 days. Note that in the main estimate we used 90 days, which accounts for any difference between column (2) in here and column 1 in table 1

contract is either a fixed price contract or a differential price contract. However loans are typically comprised of several contracts, and we can therefore have a loan that is partly backed by fixed price contracts and partly backed by differential priced ones. In the analysis in Table 2, Columns 7, 8 we run specifications at the loan level, and define a loan as being *fixed* if the share of the value of fixed price contracts that are used as collateral exceeds 50%. Here we test robustness to that definition.

We show again the main specification (Table A3 column 2), and also show a range around that specification. We show that using a range from 45% to 60% we find very similar results. As with the robustness on the previous dimension, there are trade-offs. If we go too far towards zero in our definition of fixed price contract, we should not expect to find any effect. This is because in fact most contracts are differential price contracts, and we have already shown (and indeed the theory predicts we should not expect) that there is no reaction to a price surprise for differential price contracts. On the other hand, if we go too far towards one, we expect only the firms with exceptional relationships receiving *exclusively* fixed price contracts since relationship value is a strong predictor of contract type (see table 6), and we have also shown in the main analysis that it is the firms with the worst relationships, among those with fixed price contracts, that are most likely to strategically default. In any event, the robustness exercise suggests that the results are not very sensitive to changes in the definition we use for fixed price loans, with a range from 45% to 60% producing very similar results.

We can also examine robustness of the event study (table 3) as well. The event study is important to separate strategic default from loan diversion as an alternative source of default. The main decision we made in that table was regarding how to define a price surprise. We decided to examine very large price surprises, and chose a definition for an ‘event’ as being near the upper end of weekly price shocks such that we would still have sufficient variation in our main specification (i.e. with a two week window). In fact, there are only 123 observations using this definition. This is not too surprising. We use a 3% weekly threshold which is quite a demanding definition of an event. If this level of weekly price increase were sustained for a year, it would result in a 365% increase in the price of coffee ($100 \cdot (1.03)^{52} = 465.09$). Of course with so few observations we should also check that the results are not driven by only a few observations. One way we do this is by increasing the event-study window to three weeks, and results there are consistent (Table 3 column 3). Here we increase the number of observations on another dimension: by reducing the threshold of price jump to be considered an event.

Table A3: Robustness of Table 2 to different definitions of ‘fixed price loan’

Threshold for fixed price loan (% loan fixed)	45% fixed (1)	50% fixed (2)	55% fixed (3)	60% fixed (4)
Fixed Price Loan x Price Surprise	0.182** (0.0907)	0.196** (0.0907)	0.168* (0.0913)	0.173* (0.0922)
Fixed Price Loan	-0.233** (0.111)	-0.253** (0.111)	-0.217* (0.111)	-0.221** (0.112)
Price Surprise	-0.0618 (0.0766)	-0.0661 (0.0767)	-0.0517 (0.0766)	-0.0534 (0.0757)
Mill Fixed Effects	Yes	Yes	Yes	Yes
Buyer Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes
<i>N</i>	967	967	967	967
<i>R</i> ²	0.386	0.387	0.386	0.386

Note: Regressions are at the loan level. At the loan level we sometimes have loans with both fixed price and differential price shipments, so we define a loan to be a ‘fixed price loan’ if more some threshold of the sales (in dollars) come from fixed price shipments. The threshold used varies as specified in the table. In all cases our dependent variable is default or severely late payments, where lateness is defined as being at least 90 days past due. Price surprise is defined as being the price at the time of maturity divided by the futures price for that time at the time the agreement was made. Standard errors are clustered at the loan level. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

We look at 4 different definitions of an event in table A4, looking at 0.5% intervals below our main estimate at 3%. So we check 1% (column 1); 1.5% (column 2); 2% (column 3) and 2.5% (column 4). This gives us more than our baseline 123 observations in each case, with a range from 124 to 223. The results are extremely consistent across the board. With the 1% definition the estimate is slightly larger, and with a 2% definition the estimate is slightly smaller, but all are well within the confidence intervals of each other, and all are positive and significant, as expected.

Table A4: Robustness of Table 3 to different definitions of ‘event’

Threshold price increase for event	1% jump (1)	1.5% jump (2)	2% jump (3)	2.5% jump (4)
Shipment Scheduled After Price Jump	0.173** (0.0787)	0.153** (0.0727)	0.116* (0.0677)	0.131*** (0.0189)
Control Group Mean of Dependent Variable	0.097	0.094	0.079	0.071
Observations	223	177	145	124
R-squared	0.025	0.020	0.015	0.023

Note: Local linear regressions are executed at the contract level. In all cases our dependent variable is default or severely late payments, where lateness is defined as being at least 90 days past due. All regressions use an event study methodology, where an event is defined as a weekly price increase of varying amounts. The definition of a price-jump event is as listed at the top of the table (i.e. ranging from 1%-2.5% price increases). Standard errors are clustered by event-day bins. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

Finally, we examine the robustness of the event study to controlling for in-season price increases. One concern could be that in-season price increases are highly correlated with out-season price increases and that defaults are driven by the in-season and not the out-season jump. This concern should be mitigated somewhat by the narrow window surrounding an event that we can use (our narrowest window of one week should contain very few in-season days) however it is prudent and straightforward to check that our results are robust to this control as well.

Table A5 checks this, running the main result (for the out-season and fixed-price sample), including a control for any in-season price swings. The results are nearly identical, consistent with the fact that the event study window is sufficiently narrow to control for this possibility. The estimate in the main specification falls from 0.143 (table 3, column 1) to 0.136 (table A5 column 1), with the estimates using alternate event windows being similarly close (in fact more so).

Table A5: Robustness of Table 3 to in-season price control (main sample)

Dependent Variable:	Default or 90+ days late		
	2-weeks	1-week	3-weeks
Event Window:	(1)	(2)	(3)
Shipment Scheduled After Price Jump	0.136*** (0.0190)	0.120*** (0.00645)	0.0945** (0.0422)
In-season price change control	Yes	Yes	Yes
N	123	70	154
R^2	0.042	0.045	0.036

Notes: local linear regressions are executed at the contract level. In all cases our dependent variable is default or severely late payments, where lateness is defined as being at least 90 days. All regressions use an event study methodology, where an event is defined as a weekly price increase of at least 3%. Standard errors are clustered by event-day bins. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

We conclude from our battery of robustness checks (both presented in the main tables as well as the supplementary robustness checks here) that the results on prediction 1 are very stable, and we are quite confident that in fact strategic default is quite prevalent in this context.

E.2. Punishment

In the body of the text we discuss both the buyer and the lender seem to punish default, however given the nature of the data it is difficult to identify buyer punishment given that the lender also punishes default. This is because when we fail to observe a buyer-mill pair later in the data it could be because the mill never shows up in the data again because they have been punished by the lender, or it could be because the buyer has punished the mill which may or may not lead to the mill re-appearing in the data. In any event, the theory simply requires that the mill is punished by the lender-buyer pair, and that is identified, it is simply difficult to disentangle whether punishment is coming from only the lender, or from both.

In order to identify buyer punishment we need a scenario where the lender would not punish the mill, but the buyer would. We can investigate this issue by looking at cases where the mill repays the lender directly. We show in Appendix D that this tends to occur when the mill has strategically defaulted on the buyer, but when the buyer is not particularly important to the lender. In those cases, the lender is willing to accept direct payment from the mill, risking its relationship

with the buyer. In these cases, because the mill repaid the loan to the lender, we might not expect the lender to punish the mill, but since the buyer has been defaulted on, they may punish the mill.⁵⁸

For these cases to be able to identify buyer punishment distinctly from lender punishment we need to show that while direct payment by the mill does not impact the likelihood that the mill reappears in the data, it does affect the likelihood that the buyer-mill pair reappears in the data. Note that another advantage of this context is that it makes some sense to include mill fixed-effects in this specification. In the main specification we do not include mill fixed effects because a large portion of the punishments occur on first interactions, and because these firms are punished they disappear from the data. With mill fixed effects, any observation that appears once in the data does not contribute to identification, although we are still interested in this source of variation. In this case though, a default does not result in the observation being removed from the data as long as the lender does not punish direct repayment. That is, it may still be true that direct repayment is more common on first interactions between a buyer and a mill, and it might be true that this buyer-mill pair are never seen again, but since we expect the mill to still re-appear in the data in these cases, we can still identify punishment from them.

Table A6: Identifying Buyer Punishment Using Direct Repayments by the Mill

Dependent Variable	Lender Punishment (1)	Buyer Punishment (2)
Direct Repayment by the Mill to the Lender	-0.0458 (0.0420)	-0.0797* (0.0449)
Mill FE	Yes	Yes
Buyer FE	Yes	Yes
Year FE	Yes	Yes
Month FE	Yes	Yes
Observations	434	434
R-squared	0.775	0.607

Note: Regressions are at the contract level and only include fixed price contracts. In all cases our dependent variable is whether the mill appears again in the data with either the lender or the buyer. Direct Repayment by the Mill to the Lender denotes whether the paying agent on the loan repayment transaction matches the name of the mill. Standard errors are clustered at the loan level. Standard errors are clustered at the loan level. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

Table A6 shows the results from using direct repayment by the mill instead of lateness or outright default as the object of punishment. Here we see a much

⁵⁸One thing to note is that if the buyer does not have a strong relationship with the lender they may be inconsequential to the mill as well. This may make punishment less effective and therefore less prevalent in practice.

smaller effect of lender punishment, as expected, and it is now indistinguishable from zero. On the other hand, the buyer effect is similar to what it was in the main estimates - indicating that in about 8% of defaults the buyer will punish the mill by not doing business with them again. This effect is almost twice as large as the lender effect, indicating that the identification problem inherent in the main results is much less of a concern here. Based on these estimates we are a bit more comfortable with the interpretation that both the lender and the buyer punish default by the mill.

E.3. Contract Length

We now discuss the role of contract length in the test of Prediction 3. A longer forward contract exposes parties to a higher variance in the price surprise. Fixed price contracts might then be associated with short forward contracts where there is little chance for price movement, and these short contracts are precisely made between parties with a strong relationships. Although the alternative scenario in which parties with strong relationships can afford to sign longer forward contracts is equally plausible, it is worth exploring the robustness of the results on contract selection to the joint consideration of contract length as a contractual outcome. To examine this we jointly estimate these two dimensions of the terms of the contract, since both may be endogenous to relationship value, but could be related.

Effectively, the joint estimation runs the same specification on each of the outcomes (contract type, loan length) allowing for correlation between the residuals of the two models. So while the estimates themselves will not change, they may get much less precise when we allow for some correlation between the errors (which we should since the terms of the contract are most likely jointly determined). Table A7 tests this type of specification using our main empirical model. We find that very little changes, and each of the specifications remain precise at conventional levels when we are more systematic in our handling of the errors.

Perhaps more relevant is how loan length influences the estimates themselves. We can run placebo estimates to determine whether the length of the loan (on a number of dimensions) influences the type of contract signed. Of particular concern is whether once mills and buyers know that a contract will be small, or short, that they are then ok with a fixed price contract.

We therefore take three measures to capture this general concept, and look at their conditional correlation with contract type. First we look at the number of shipments on a loan, with the concern in mind that more shipments might indicate

Table A7: Joint Estimation of Effect of Relationship Value on Loan Length and Contract type

Outcome:	Fixed (1)	Loan Length (2)	Fixed (3)	Loan Length (4)	Fixed (5)	Loan Length (6)	Fixed (7)	Loan Length (8)
Buyer-Mill Fit (AKM)	0.115** (0.0533)	0.563 (0.992)						
Centralization of loan collateralization (more money by fewer buyers)			0.232*** (0.0631)	-5.239*** (0.935)				
Cumulative history with buyer (1,000,000)					3.585*** (0.454)	0.357 (18.10)		
Lender's Score for Buyer-Mill Relationship							0.0377* (0.0226)	-1.891** (0.891)
Mill Fixed Effects		Yes		Yes		Yes		Yes
Buyer Fixed Effects		Yes		Yes		Yes		Yes
Year Fixed Effects		Yes		Yes		Yes		Yes
Month Fixed Effects		Yes		Yes		Yes		Yes
Number of observations		995		1,228		1,228		437

Notes: Regressions are at the contract level. In all cases our dependent variable is default or severely late payments or length of the loan in days. Lateness is defined as being at least 90 days past due. Loan length is defined as the number of days between the closing date and the maturity date. The main proxy for relationship value is the age of the relationship, defined in the same way we define relationship value in table 5. Standard errors are clustered at the loan level. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

Table A8: Placebo specification: do buyers simply use fixed contracts for short contracts?

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Number of Shipments	-0.00652 (0.00597)	-0.00654 (0.00665)	-0.00645 (0.00600)						
Time between Closing and Last Shipment				-0.00157 (0.00137)	-0.00135 (0.00155)	-0.00136 (0.00139)			
Quantity Delivered (500 tons)							0.000965*** (0.000148)	0.00100*** (0.000156)	0.000992*** (0.000149)
Client Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Buyer Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	No	Yes	Yes	No	Yes	Yes	No	Yes
Country-Month Fixed Effects	No	Yes	No	No	Yes	No	No	Yes	No
Spot and Futures Prices	No	No	Yes	No	No	Yes	No	No	Yes
Observations	967	967	967	967	967	967	967	967	967
R-squared	0.251	0.283	0.253	0.250	0.282	0.252	0.251	0.283	0.253

Notes: The unit of observation throughout the table is a contract. The number of shipments variable refers to a count of the shipments that take place as collateral backing a particular loan. The time between closing and last shipment is the number of days between the date that the loan was approved and the date that the last shipment on the contract was scheduled. The quantity delivered is the scheduled amount of coffee to be delivered on a particular contract, in 500 tons. The dependent variable in each case is a binary variable for whether the contract between the buyer and the mill was tied to the world price of coffee (1 if not, 0 if it was). Standard errors are clustered by loan. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

a longer time between the contract being signed and any potential price surprise. However, we find that the number of shipments does not explain the contract type. Second, we directly look at the number of days between closing and shipment, and again find no correlation with contract type. Finally we look at whether larger or smaller contracts are more or less likely to be fixed price. It would be worrisome if only the smaller contracts received fixed price contracts - indicating that a buyer may be understand that default is possible but can not be bothered to impose its desire for a differential price for such an insignificant transaction. We actually find that the more substantial contracts are more likely to come with fixed prices. This likely reflects the fact that the mills with the best relationships are the ones that receive fixed price contracts, and are also the ones that receive large contracts. In sum, the test for Prediction 3 is robust to considering the joint determination of contract length and contract type.⁵⁹

⁵⁹Additional placebos show that price surprise do not correlate with contract type.

APPENDIX F. CREDIT CONSTRAINTS (FOR ONLINE PUBLICATION)

F.1. Main Evidence on Credit Constraints

We find that buyers and the lender are aware of the risks associated with strategic default, and try to balance filling the demand for credit and insurance while also using access to these instruments to discipline mills and mitigate default. The model predicts that (some) mills will therefore be credit constrained. Due to lending process used by the lender, we can test for credit constraints among mills, following a methodology related to [Banerjee and Duflo \(2014\)](#). We exploit the fact that the lender assigns loans based on credit scores that are rounded up or down. This allows us to compare ostensibly identical mills - one just barely being rounded up, another just barely being rounded down - who received different loan amounts from the lender. It is worth noting that in so doing we are able to test for credit constraints employing a RDD methodology on a sample of very large firms.⁶⁰

The idea behind the test for credit constraint is simple. A mill is credit constrained if the marginal product of capital is larger than the interest rate paid on the marginal dollar borrowed $MPK > r$.⁶¹ As clarified in [Banerjee and Duflo \(2014\)](#), then, the test for credit constraints is as follows: *a mill receiving a larger loan is credit constrained if i) takes-up the loan to expand production and ii) does not substitute any existing loan.* Note that the test is valid if three conditions are satisfied. First, the additional loan should now change the (marginal) interest rate. Second, the larger loan must affect operations at the margin, i.e., its availability shouldn't fund fixed costs or determine whether the mill operates at all or not. Third, the mill must be able to use the loan to pay down other loans. The first assumption can directly be tested for in the RDD analysis. Furthermore, our lender interest rates are representative of the industry (see figure A3). The lender

⁶⁰To implement the RDD we need information on the terms of the loan given by the lender as well as information on other loans that the mills have. This information is collected by the lender as part of the application process. The information, however, is available only for mills that apply for a loan in a subsequent year. If the following application came within three years of the former, then we observe an overlap in the financial records. That is, for the first loan, we can see what the mill did with the money. Because the lender only tracks historical financial health of the mill at the time of a loan application, and does not follow up with audits of their books after the loan is received and paid back, we can only run this analysis for a selection of mills. Indeed, based on our previous analysis, this suggests that we are least likely to observe the mills that defaulted on loans previously. As such, we should interpret estimates as evidence that *some* level of credit constraints exist in the market, and not as an indication of representative averages.

⁶¹A mill with infra-marginal loans at lower interest rate is credit rationed, in the sense that would like to borrow more at cheaper rates, but not necessarily credit constrained.

will then be the marginal one for some mills, but not for others. This doesn't affect the test, however. If our lender is cheaper than the marginal source of finance, and the mills are not credit constrained, they will simply substitute away from other loans that are available to them at a higher interest rate. In order for this substitution to still lead to an increase in production, it would have to be the case that the mill substitutes for *all* other sources of working capital. This is clearly not the case in the data. The second and third requirements are also satisfied in our context. Our analysis focuses on working capital loans in an industry with a clear seasonality pattern. Mills entry/exit is therefore not the relevant margin and, since our loans are signed at the beginning of the season, the mill can always reduce working capital borrowed from other sources.⁶²

We test for credit constraints implementing an RDD design on the size of the loan. The lender determines the size of the working capital loan based on a letter score. The letter score is calculated rounding an underlying continuous numerical score which weights a large number of mill and loan characteristics (see details below). We have two thresholds where rounding takes place. Mills are classified as having a C, B, A or AA credit score, based on a continuous numerical score between one and five given by an auditor.⁶³ We have follow-up data on 523 loans. We conduct the analysis on the B-A and A-AA thresholds separately.⁶⁴ In general, the lender provides 40%, 60% and 70% of the required funds to fulfil the sale contracts on loan applications with a B, A and AA scores respectively.

Table A9 reports the first set of results. Columns 1-2 show that on both the B-A and A-A thresholds firms received substantially larger loans if they got rounded up instead of rounded down (also see figure A12). The estimates are consistent with what reported by the lender. The estimates also show that the RDD approximates a very large increase in loan size, significantly larger than the one studied in [Banerjee and Duflo \(2014\)](#). Furthermore, in neither case was there an associated decline in the interest rate (figure A13 and Table A9, columns 3-4). However while we find no evidence of substitution away from other loans for the B-A threshold (figure A14 and table A9, column 5), we find basically one to one substitution out of other loans for the A-AA threshold (table A9, column 6). This suggests that the B-A mills could be credit constrained while the A-AA mills are

⁶²A mill could also have $MPK > r$ if it is insurance constrained. In this case however, the mill would simply use the additional loan to substitute for other more expensive sources of finance.

⁶³For example, mills receiving between 2.7 and 3.35 are classified as receiving a B score and mills receiving scores between 3.36 and 4.1 receive an A. Anyone over 4.1 receives an AA.

⁶⁴We can not run the analysis on the C to B threshold since very few mills receive a C and among those even fewer apply again in the future, leaving us with no follow-up information for them.

certainly not.

Table A9: Credit Constraints I: Is there exogenous variation in (just) credit?

Dependent Variable:	Loan Amount		Interest Rate		Other Loans	
	B-A	A-AA	B-A	A-AA	B-A	A-AA
	(1)	(2)	(3)	(4)	(5)	(6)
RD Estimate: Optimal Bandwidth	112,268 (47,078)**	237,093 (113,322)**	0.0001 (0.001)	-0.0003 (0.0003)	48,045 (71,967)	-194,920 (154,741)
RD Estimate: 75% Bandwidth	132,705 (39,877)***	299,342 (87,801)***	-0.0002 (0.0009)	-0.0005 (0.0004)	105,241 (67,215)	-360,359 (172,908)**
RD Estimate: 125% Bandwidth	118,612 (38,253)***	211,233 (122,506)*	0.0005 (0.0012)	-0.0004 (0.0003)	-5,793 (83,036)	-114,031 (166,904)
Observations	523	523	523	523	523	523

Notes: An observation in these regressions is a loan. We have 387 loans (out of 499) that have a numerical score. The missing observations are earlier loans from before the current scoring system was put in place in 2008. While the entire sample contains loans from 2003 to 2014 the first loan with a numerical score closes late in 2008. In the RDD regressions observations are clustered by numerical score bins of size 0.05. In each case the RDD is estimated using the kernel density method, with an optimal bandwidth chosen using the Imbens and Kalyanaraman (2011). We report the number of observations used prior to the bandwidth computation. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

The second part of the test for credit constraints entails examining how the additional money is utilized. Given that the A-AA mills substitute away from other loans we should not expect them to also expand their operations. Indeed, we find no evidence that they expand operations or that they earned a higher profit (neither result is reported). We focus instead on the mills at the B-A threshold. Ex-ante we might expect the B-A mills to expand operations since they received loan money which did not induce a substitution effect away from other loans.⁶⁵

We find that for the B to A mills purchases do expand by about \$130,000 (figure A15 and table A10 column 1). This estimate quite similar to the estimate of the additional loan amount they received from the lender (about \$115,000). The increase in cherries purchased reflects about a 33% increase in production (column 2). Furthermore, sales do seem to increase as well, by about 45% (figure A16 and table A10 column 3), reflecting an MPK of about 12%. The average interest rate is about 8%, which suggests that some firms are very likely to be credit constrained, but this relatively small difference suggests that many are not. In fact this is remarkably consistent with our calibration exercise, which produces an MPK -

⁶⁵Mills may also undertake unproductive activities with the additional loan money (ex-ante moral hazard). This should however increase the likelihood of default on the loan. We see no such effect in the data, providing further evidence that loan diversion is unlikely to be a source of credit constraint in the lender portfolio.

$r = 6\%$ on average (Table 8, Panel A, row 7). Combined, these two pieces of evidence imply that many firms at the B-A threshold are credit constrained.

Furthermore, and important for the calibration results (i.e. this estimate is what we use for η in the calibration), we find some evidence of a pass-through effect down the supply chain. The prices paid to farmers increases by about 13% (table A10 column 4).

Table A10: Credit Constraints II: Did firms use additional credit to expand? (B-A threshold only)

Dependent Variable:	Cherries Purchased	log(Cherries Purchased)	log(Sales to other buyers)	log(Prices Paid to Farmers)
	(1)	(2)	(3)	(4)
RD Estimate: Optimal Bandwidth	133,115** (56,161)	0.338** (0.148)	0.452** (0.222)	0.134*** (0.0290)
RD Estimate: 75% Bandwidth	122,259** (49,252)	0.283** (0.125)	0.403* (0.227)	0.138*** (0.0421)
RD Estimate: 150% Bandwidth	152,646** (64,615)	0.376** (0.152)	0.385** (0.187)	0.189*** (0.0581)
Observations	212	212	206	166

Notes: An observation in these regressions is a loan. We have 387 loans (out of 499) that have a numerical score. While the entire sample contains loans from 2003 to 2014 the first loan with a numerical score closes late in 2008. The missing observations are earlier loans from before the current scoring system was put in place in 2008. Of those loans that have a score, we have data on the performance of the mill after receiving the loan for only 212 mills. This is because the data on financial performance comes from pre-loan audits, and so we only see ‘future performance’ if a mill applies for another loan from the lender. For only 166 mills we have both the value and quantity of cherries purchased and we use this data to construct price paid to farmers. When purchases are reported directly we use that value, and in cases where purchases are not directly reported, we use COGS (Cost of Goods Sold) as a measure of purchase value, if the only item they sell is coffee. For these firms we therefore have information on how much the mill spent on purchases but we cannot compute a price for the COGS because we do not have the necessary quantity information. In the RDD regressions observations are clustered by numerical score bins of size 0.05. In each case the RDD is estimated using the kernel density method, with an optimal bandwidth chosen using the Imbens and Kalyanaraman (2011). We report the number of observations used prior to the bandwidth computation. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

F.2. Robustness of Credit Constraints Evidence

We take a number of measures to ensure the validity of the regression discontinuity design. For instance, one major concern is that the lender is able to influence the scores of the audit, and either nudge preferred mills that are near the threshold required to get a better loan up, or conversely move mills that they have less confidence in down. As a first check of this we examine the McCrary density test, which looks at the density of contracts just above and just below the discontinuity threshold (McCrary (2008)). This can be seen in figure A11, which plots the density of loans against their scores. The B-A threshold is indicated with a red line, and we find no reason at all to be concerned about systematically nudging

up or down of different loans.

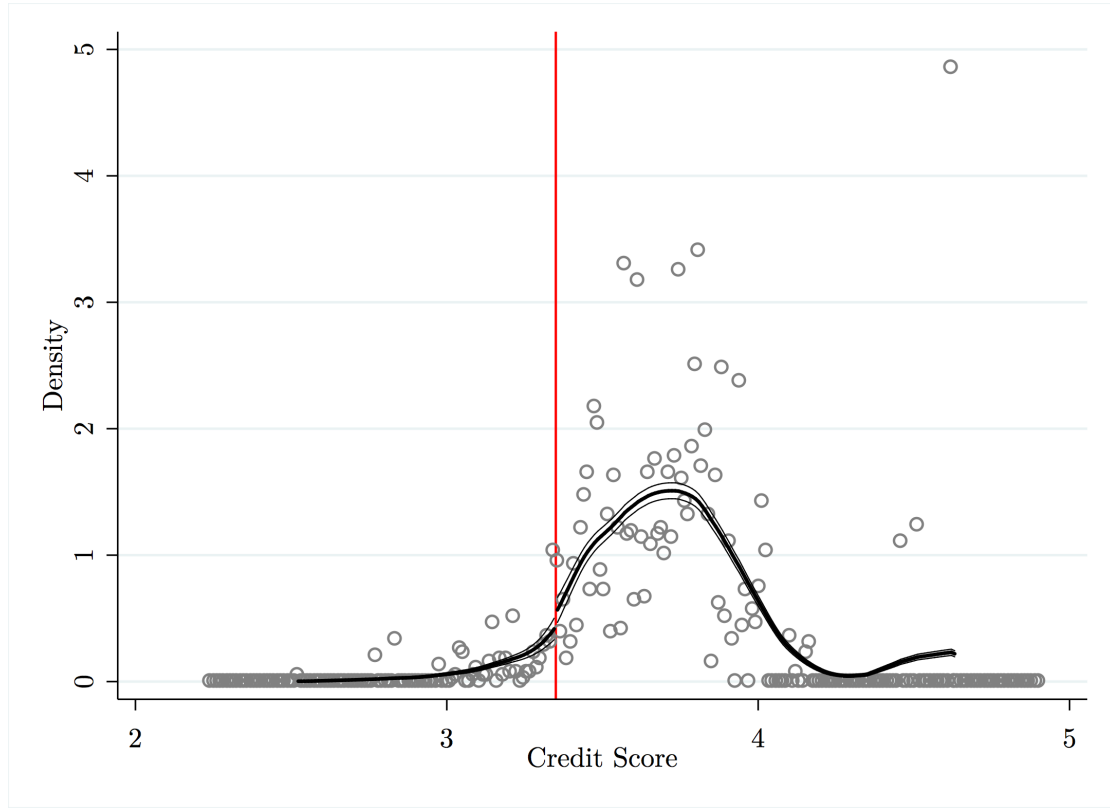


Figure A11: Sorting around the threshold 1: McCrary density test

Notes: This graph shows the density of credit scores on either side of the B-A credit score threshold (thick line), and plots 95% confidence intervals (thin lines) to illustrate that there is no difference in the density of loans on either side of the threshold. One concern is that the lender manipulates the algorithm to provide some high quality firms with more loan money. This would violate one of the assumptions of the RDD. We test for this by examining the density of contracts just above or just below the threshold. If we believe that agents nudge their ‘favorite’ clients to above the threshold then we should see a higher density just above relative to just below the cut-off. In fact we do not see that at all. We conclude that either the lender does not manipulate the scores or that they nudge up and nudge down with fairly equal regularity.

Of course if the lender nudges mills up and down at the same rate, the McCrary test would be satisfied, but the regression discontinuity would be invalidated. Accordingly we take additional measure to ensure that audits were not manipulated by the lender. The aggregated score is based on a host of sub-scores, so the only way to manipulate the score would be to manipulate the sub-score. However, there are 64 sub-scores, so it would be unlikely for any one of them to be systematically contributing to moving the aggregate over the required threshold. We can therefore check to see if firms that end up just above the aggregate threshold are also significantly higher for any of the sub-scores. Of course, all should mechanically be slightly higher, since they do contribute to the aggregate score, but we should

not expect significant differences. We find, of the 64 sub-scores that only 5 are significantly different on either side of the threshold, well within what we would expect based on random chance (Table A11). Furthermore, the sub-scores that are different are not the ones we would expect to be manipulated because they carry very little weight in the make-up of the overall score. Only a handful of contracts are actually so close to the threshold that their letter score could be influenced by one of these sub-scores. Therefore, aside from the fact that the lender manipulating the audits would undermine having any audits in the first place, we are quite confident based on the two validity tests that we ran, that there is no manipulation of the scores around the threshold.

Table A11: Sorting around the threshold 2: Sub-scores

Sub-score considered	Dependent Variable: Credit Score			
	B-A Threshold		A-AA Threshold	
	Estimate (1)	p-value (2)	Estimate (3)	p-value (4)
Quality of Accounting	-3.21	0.204	0.04	0.519
Liquidity Risk	-0.073	0.797	0.262	0.516
Profitability	0.260	0.121	0.289	0.370
Credit History	0.293	0.179	0.544	0.233
Asset Quality	0.192	0.181	0.481	0.290
Agriculture	-0.195	0.164	0.559	0.204
Productivity and Yields	-0.150	0.366	0.693	0.122
Quality of Materials	-0.263	0.113	0.465	0.315
Processing	0.272	0.056	0.269	0.491
Appropriate Technology	0.136	0.328	0.239	0.577
Staff Capacity	0.123	0.347	0.342	0.435
Supplier Reliability	0.361	0.057	0.433	0.351
Crop Security	0.228	0.142	0.375	0.410
Managerial Quality	0.051	0.744	0.173	0.475
Finance and Accounting	0.208	0.183	0.177	0.737
Operations	-0.193	0.136	0.107	0.802
Internal Controls	-0.062	0.370	0.485	0.247
Marketing and Sales	-0.177	0.199	0.397	0.390
Staff Retention	0.384	0.109	0.493	0.268
Report Quality	0.099	0.482	0.190	0.454
E-mail Promptness	0.129	0.487	0.439	0.336
E-mail Quality	0.190	0.248	0.278	0.499
Visit Quality	-0.319	0.097	0.148	0.326
Relationship with Buyer	0.168	0.319	0.381	0.408
Strength of Buyer	0.0016	0.991	0.333	0.456
Appropriateness of Buyer	-0.177	0.466	0.477	0.256
Agreement Type	-0.179	0.307	-0.198	0.191
Weather	-0.068	0.537	0.868	0.026
Pests	-0.192	0.145	0.522	0.170
Country Stability	-0.039	0.842	0.297	0.442
Product Demand	-0.216	0.122	0.354	0.409
Sales Price Volatility	-0.369	0.069	0.486	0.189
Regulations	0.023	0.868	0.400	0.350
Perishability	0.167	0.253	0.406	0.335
Total p-values < 0.1	.	4	.	1
Total p-values < 0.05	.	0	.	1
Total p-values < 0.01	.	0	.	0

Note: This table reports p-values and estimates from local linear RDD regressions of numerical score on sub-score. There are two separate regressions on each line, one for the B-A threshold and another for the A-AA threshold. In total there are 68 regressions. We find in total 5 estimates significant at the 10% level or lower; 1 estimate at the 5% level or lower; and 0 estimates at the 1% level or lower. By random chance we would expect 6-7; 3-4; 0-1 respectively.

F.3. Supplementary RDD Figures




Figure A12: RDD 1: Firms above the threshold and access to credit

Notes: We exploit a 'rounding RDD' in the credit score given to mills to show that firms are credit constrained. Firms receive an 'A' score if they receive a numerical score above 3.81 and they receive a 'B' score if the numerical score is below 3.81. The red vertical line shows illustrates the 3.81 cut-off. We present means from equidistant x-axis bins on either side of the threshold. Bins are 0.03 in size in each case. The figure shows that firms below that threshold receive much smaller loans than firms above the threshold. If these firms do not also receive lower interest rates, and they use the money to expand operations instead of substituting away from more costly loans then we can conclude that they are credit constrained.




Figure A13: RDD 2: Firms above the threshold and interest rates

Notes: We exploit a 'rounding RDD' in the credit score given to mills to show that firms are credit constrained. Firms receive an 'A' score if they receive a numerical score above 3.81 and they receive a 'B' score if the numerical score is below 3.81. The red vertical line shows illustrates the 3.81 cut-off. We present means from equidistant x-axis bins on either side of the threshold. Bins are 0.03 in size in each case. The figure shows that firms above and below this threshold receive similar interest rates. This is important because if they received lower interest rates it would be hard to know whether the larger loans or lower interest rates were responsible for increases in purchases, sales, etc. Indeed, because of the similar rates shown here we can conclude that we have reasonable exogenous variation in credit availability, as long as firms do not substitute away from more costly loans as they get access to additional credit.

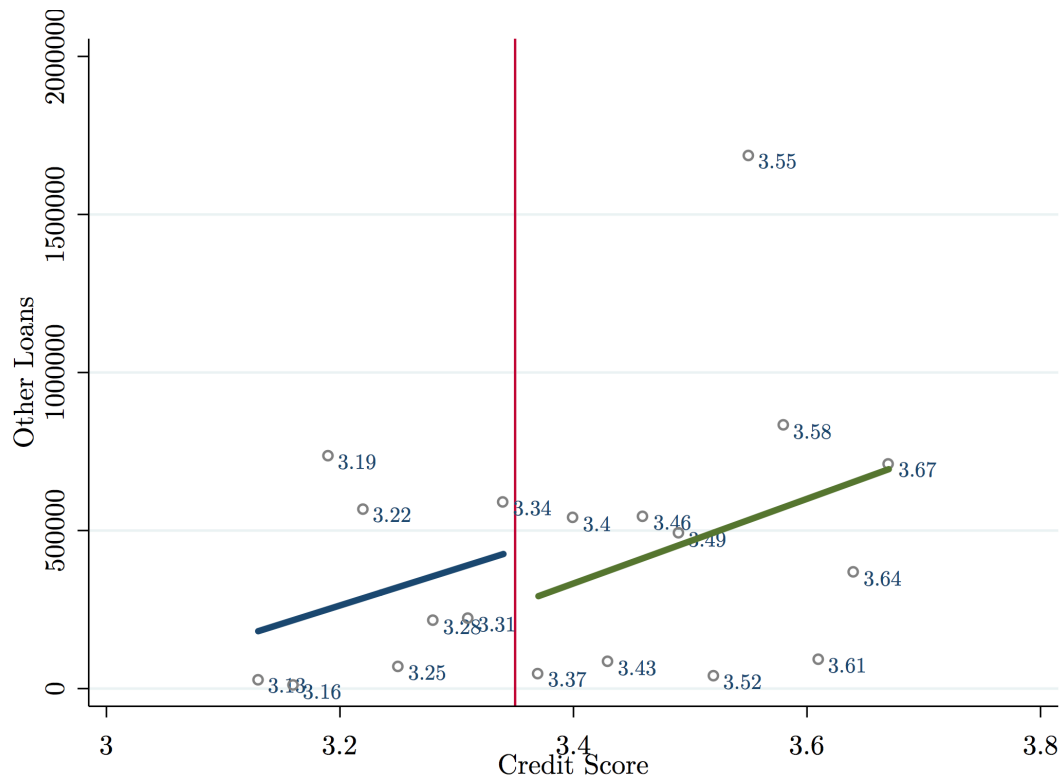


Figure A14: RDD 3: Firms above the threshold and other loans

Notes: We exploit a 'rounding RDD' in the credit score given to mills to show that firms are credit constrained. Firms receive an 'A' score if they receive a numerical score above 3.81 and they receive a 'B' score if the numerical score is below 3.81. The red vertical line shows illustrates the 3.81 cut-off. We present means from equidistant x-axis bins on either side of the threshold. Bins are 0.03 in size in each case. The figure shows that firms above and below this threshold receive amounts of loans from other sources. This is important because if they substituted away from costlier loans it would suggest that they were not credit constrained, and we should not expect any impact on purchases or sales.




Figure A15: RDD 4: Firms above the threshold and purchases

Notes: We exploit a 'rounding RDD' in the credit score given to mills to show that firms are credit constrained. Firms receive an 'A' score if they receive a numerical score above 3.81 and they receive a 'B' score if the numerical score is below 3.81. The red vertical line shows illustrates the 3.81 cut-off. We present means from equidistant x-axis bins on either side of the threshold. Bins are 0.04 in size in each case. The figure shows that firms above the 3.81 threshold purchased more cherries from farmers.




Figure A16: RDD 5: Firms above the threshold and sales

Notes: We exploit a 'rounding RDD' in the credit score given to mills to show that firms are credit constrained. Firms receive an 'A' score if they receive a numerical score above 3.81 and they receive a 'B' score if the numerical score is below 3.81. The red vertical line shows illustrates the 3.81 cut-off. We present means from equidistant x-axis bins on either side of the threshold. Bins are 0.04 in size in each case. The figure shows that firms above the 3.81 threshold earn more revenue.

The International Growth Centre (IGC) aims to promote sustainable growth in developing countries by providing demand-led policy advice based on frontier research.

Find out more about our work on our website www.theigc.org

For media or communications enquiries, please contact mail@theigc.org

Subscribe to our newsletter and topic updates www.theigc.org/newsletter

Follow us on Twitter [@the_igc](https://twitter.com/the_igc)

Contact us
International Growth Centre,
London School of Economic
and Political Science,
Houghton Street,
London WC2A 2AE

IGC

**International
Growth Centre**

DIRECTED BY



FUNDED BY



Designed by soapbox.co.uk