STRATEGIC DEFAULT IN THE INTERNATIONAL COFFEE MARKET*

Arthur Blouin   Rocco Macchiavello

December 14, 2018

This paper studies strategic default on forward sale contracts in the international coffee market. 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 and suppliers have been paid. Estimates suggest that roughly half of the observed defaults are strategic. 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. A model calibration suggests that the possibility of strategic default causes 15.8% average losses in output, significant dispersion in the marginal product of capital and sizeable negative externalities on supplying farmers. **JEL Codes**: D22, L14, G32, O16.

Keywords: Strategic Default, Contractual Forms, Counterparty Risk.

*We are grateful to the Editors, Pol Antras and Andrei Shleifer, and to four anonymous referees. Also 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 many comments and suggestions. 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. Matteo Benetton, Niclas Moneke, Juan Morales and Caroline Nicol provided outstanding research assistance. We gratefully acknowledge IGC and IPA financial support. Correspondence Address: Rocco Macchiavello, Department of Management, London School of Economics and Political Science, Houghton Street, London, WC2A 2AE United Kingdom, r.macchiavello@lse.ac.uk

I. 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, Sapienza and Zingales 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; Fafchamps 2003).

This paper develops a test to empirically identify strategic default and implements it in the context of the international coffee market. 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, Gibbons and Murphy 2002 and Hart 2009). The test identifies strategic

[1]See, e.g., Lacker and Weinberg (1989); Hart and Moore (1998); Shleifer and Wolfenzon (2002); Hart (2009); Ellingsen and Kristiansen (2011). Different contract terms cause changes in incentives to default. 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). Note also that we refer to strategic default in the narrow sense of ex-post moral hazard (as opposed to the standard, ex-ante, one). Both forms of moral hazard are strategic, in the broader sense of being willful acts.

[2]Distinguishing the two forms of moral hazard is important. First, they have different welfare implications (strategic default is a transfer, while the standard moral hazard reduces surplus directly). Second, they are differently affected by changes in the environment and, therefore, require different remedies. For instance, 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.

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 decisions and payments to suppliers have been made. We quantify the importance of strategic default in the coffee market, not only for production efficiency but also for contract design, insurance, and credit availability.

We conduct our analysis on pre-financing agreements in the international coffee market. The context offers conveniences for empirical design but is also of intrinsic interest.[3] We use confidential data from a lender specialized in this common type of working capital loan. The structure of the contract is as follows. Before harvest begins, coffee mills sign forward sale contracts with foreign buyers that import coffee. The lender then advances funds to the coffee mill, backed by forward sale contracts. 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 directly repays the lender upon receiving the coffee. We obtained detailed information on the universe of 967 loans extended by the lender to 272 coffee mills in 24 countries. In addition to customary information on the loan contract terms, the files include detailed information on the forward sales contracts used as collateral as well as audited financial accounts for the borrowing mills.

We are interested in two questions: **1)** is there evidence of strategic default? and **2)** how large are the inefficiencies caused by strategic default? To answer the first question we lay out a framework of the mill's decision to strategically default on the forward sale contract. In deciding whether to deliver the coffee or default, the mill trades-off financial gains against losses in the future relationship with the buyer and the lender.[4]

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

[4]While it is standard practice in the industry to write formal contracts to obtain loans and accompany

There are two types of forward sales contracts. In **fixed prices** contracts the price of coffee is fixed at the time parties sign the agreement. In **differential price** contracts the price eventually received by the mill tracks the world spot market price at delivery. If spot prices at delivery are much higher than anticipated at the time of contracting, a mill on a fixed price contract will be tempted to default and sell the coffee to a different buyer for a higher price. A mill on a differential contract, instead, would not face such temptation.

A key challenge to test for strategic default is thus to identify unexpected changes in market conditions. This requires controlling for contracting parties expectations about the market prices prevailing at the later delivery date. 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 while controlling for confounding factors. When the international price of coffee unexpectedly increases by 10% over the duration of the contract, contractual default increases by almost three percentage points in fixed-price contracts but not in differential price contracts.

A second challenge is that unexpected increases in world coffee prices could be passed-through to farmers, thus raising costs and forcing the mill to default. To rule out this possibility and isolate strategic default we conduct an event-study that considers only price increases that occur after the end of harvest. The event-study takes advantage of the fact that the decision to default on the coffee delivery is taken when all payments to farmers have been made. Using audited monthly cash flow data we show that mills exhaust all payments to farmers during the harvest season (which typically lasts four

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.

to five months). In contrast, the majority of forward sales contracts are executed after the end of the harvest season (but well in advance of the following harvest). Among contracts that are due for delivery after harvest, the event study compares those that are due just before and just after a sudden price increase. Defaults are about 12 percentage points more likely on fixed-price contracts when a shipment is scheduled to take place in the week after a price increase relative to the week before.

A final challenge is that, since our data come from the lender, we do not directly observe delivery failures on the forward sale contracts. In the spirit of forensic economics we use alternative observable indicators of (attempted) default on the forward sale contract. Our baseline definition uses both failure to pay altogether and severely late repayments of the loan as indicators for default on the forward sale contract.[5] These combine for less than 10% of scheduled deliveries. The baseline definition is a directly observable measure of non-performance on the loan part of the contract and is standard in the literature on loans. The baseline definition, however, misses instances in which, the mill breaks the agreement, but still manages to repay the lender on time, either themselves or through another buyer. Combined, those involve 12% of scheduled deliveries. We thus implement both our baseline test and the event study using these alternative indicators of contractual default as well as the union of all possible indicators. In both the baseline and event-study specifications, results are remarkably robust: regardless of the set of indicators used, unexpected price increases are always associated with a higher likelihood of default on fixed contracts and never on differential ones. Using the baseline (and most conservative) definition of default we combine estimates from both empirical strategies to bound the prevalence of strategic default. The estimates imply that between 42% and 59% of the defaults observed on fixed price contracts are strategic.

We investigate the fate of mills involved in contractual defaults. Following a default,

[5]We consider default to include any significant breach of contract, including both outright failure to deliver, as well as very long delivery delays. We therefore use the terms 'default' and 'breach' interchangeably.

mills are less likely to receive a loan in the future from the lender. This correlation could either reflect that the lender and/or the buyer punish the defaulting mill or that mills simply default on contracts when they anticipate the relationship to end for unrelated reasons. To distinguish between these two hypotheses we show that, conditional on a default, the likelihood of relationship termination is higher when the default happened following an unexpected increase in world coffee prices, i.e., when it was more likely to be strategic. Furthermore, consistently with the strategic nature of many defaults, original field and internet searches to track all mills that have defaulted over the sample period reveal that essentially all defaulting mills are still in operation (although not dealing with our lender) years after the default occurred.

Strategic default thus appears to be fairly pervasive in this market. How costly is it? To answer this question we first need to understand how contracting parties adjust contractual terms to take into account the possibility of strategic default. We develop a parsimonious framework that captures the salient features of the contractual arrangement between the coffee mill, the buyer, and the lender. The possibility of strategic default induces a trade-off between insurance and counter-party risk. If contracts are perfectly enforceable, a risk-averse mill would sign a fixed-price forward contract in which the buyer, who has access to hedging instruments, provides insurance against price risk. When contracts are not perfectly enforceable the value of business relationships, and thus the degree of informal contract enforcement, matter. In more valuable relationships the risk of strategic default is limited. Parties still sign fixed-price contracts, possibly for lower amounts than optimal (i.e., the mill is credit constrained). In less valuable relationships, however, the risk of strategic default is large. Parties then sign differential contracts that eliminate temptations to default at the cost of foregoing price insurance (i.e., the mill is insurance constrained). As a result of strategic default, and depending on the value of the relationships with the buyer and the lender, mills can either be unconstrained, credit constrained, insurance constrained or both credit and insurance constrained. Perhaps counter-intuitively, then, strategic default can be detected on relatively more valuable relationships, the ones that sign fixed-price contracts. However the risk of strategic default imposes larger costs to those mills that are perceived by buyers to be at greater risk of strategically defaulting and therefore end up signing differential contracts that are not strategically defaulted against.

We calibrate the model to quantify both the direct and indirect costs associated with strategic default. Most of the parameters in the model are either 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 and the lender. We take advantage of the model's relative simplicity to "invert" it, and obtain an estimate of the relationship value for each contract. 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 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) is 8% (20%) higher than the interest rate (which is around 10%). We discuss policy implications of our findings in the concluding section.

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. 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 empirical work on contracts (see 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.[6] 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, Sapienza and Zingales (2013) survey people about strategic default; Bajari, Chu and Park (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 working capital loans to large firms in developing countries requires considering different aspects, most notably the importance of inter-firm business relationships.

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. As a result, trading relationships can endogenously become a source of capital and affect responses to shocks. 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

[6]For example, Karlan and Zinman (2009) and Adams, Einav and Levin (2009) offer experimental and structural analyses respectively that separate moral hazard from adverse selection in the consumer loan market.

how it influences contractual terms and efficiency.[7]

The analysis in this paper relates, and can be applied, to the study of other commonly observed financial arrangements between firms.[8] We contribute to this literature by testing for, and isolating, strategic default as a different source of moral hazard and by quantifying the associated efficiency losses.[9] Although we focus on the coffee market, pre-financing agreements are an extremely common source of working capital finance in other agricultural commodity markets (see, e.g., Varangis and Lewin 2006).[10] More broadly, we contribute to the empirical literature on financial contracting by highlighting the important role of endogenous contractual terms (Roberts and Sufi 2009).

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

[7] Two recent papers offer evidence that enforcement problems significantly impair economic output. Bubb, Kaur and Mullainathan (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.

[8] The large literature on trade credit mostly studies contracts in which suppliers extend credit to downstream buyers (see, Klapper, Laeven and Rajan (2012) and Breza and Liberman (2017) for recent contributions and references). Burkart and Ellingsen (2004) and Giannetti, Burkart and Ellingsen (2011), for instance, posit that trade credit is used to limit loan diversion, a different form of moral hazard. A large literature (see, e.g., **?**) studies the effects of credit supply on exports.

[9] We 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.

[10] They are also reminiscent of invoice discounting, factoring and other arrangements in which account receivables are used as collateral.

associated inefficiencies.[11,12]

II. Background and Data

II.A. Coffee Washing Mills

When coffee cherries change color from green to red they are ripe for harvest. Most coffee-growing countries have only one harvest a year whose timing varies by country depending on latitude, altitude 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 performed directly by farmers. The wet method is performed by coffee washing mills, the object of this study. Relative to the dry method, the wet method requires significant investment in specialized equipment but produces higher and more consistent quality.[13]

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, 27 employees, hold about 2 million in total assets and receive average working capital loans of $473,000 (see Table I). The production function is relatively simple: the quantity of parchment coffee produced is a constant proportion 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 labor, transport, electricity, marketing and, of course, costs of finance. The mills in the sample mostly supply the coffee

[11] 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 (2199).

[12] Various papers study other aspects of the coffee sector. For instance, de Janvry, McIntosh and Sadoulet (2015) and Dragusanu and Nunn (2014) look at fair trade, Macchiavello and Morjaria (2015a) study how competition between mills affects relationships with farmers, and Macchiavello and Miquel-Florensa (2018) study vertical integration between exporters and processors. Macchiavello and Miquel-Florensa (2018) also applies our test for strategic default on a sample of Costa Rican coffee mills, and find similar results.

[13] After 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 help of machines. After drying the hulling process removes the parchment skin before export.

specialty market. In this segment, coffee mills supply directly to foreign buyers.

II.B. Timing of Events and Mills' Cash Flow Profiles

To follow our empirical strategy it is crucial to have a clear understanding of the timing of events and mills' cash flow profiles. Figure I illustrates the timing of the harvest season. First, before the harvest begins, mills sign forward sales contracts with foreign buyers. These contracts specify the delivery of a certain amount of coffee of a certain quality at a later pre-specified date. The mill then obtains working capital finance either in the form of an advance from the buyer or, as is the case in this paper, as a working capital loan extended by a lender on the back of the forward sale contracts.[14] During the harvest season, which typically lasts between four and five months, mills source coffee from farmers and process it. Coffee is then delivered to the buyer in bulk shipments of pre-specified volume and quality. Finally, the working capital loan is repaid to the lender (or the buyer deducts the corresponding amount from the price).

The resulting cash flow profile is illustrated in Figure II. The Figure documents the stark separation in the timing of production and contract execution that underpins our empirical strategy. We collect monthly cash flow data from the audited financial records of the mills in our sample. The horizontal axis is the number of months from the beginning of the harvest season (at zero). The Figure reports the cumulative values of contracts signed, disbursements to farmers and scheduled coffee deliveries, averaged over seasons and mills in the sample. By the time harvest begins, mills have already signed forward contracts worth approximately 50% of overall production (contract line). Two key facts stand out. First, the vast majority (almost 85%) of payments to farmers occur during the harvest season (i.e., by month four after the beginning of harvest). This reflects the fact (confirmed by interviewed loan officers and mills) that mills rarely source coffee from farmers on credit or make payments to farmers after the

[14]Fixed 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.

end of the season. By the end of the harvest season, then, mills' financial obligations with farmers are essentially fully executed. Second, sales and deliveries are executed later. While processing takes only about two weeks, delivery takes longer as shipments are in bulk and, in some cases, mills might wait to have the right types and volumes of coffee to mix. By the time harvest ends, about 50% of contracted deliveries have occurred (delivery line). That is, about half of the contracted sales take place at a time in which production decisions are sunk and payments to farmers have been executed. The separation between the time at which the mill purchases coffee from farmers (harvest) and the time in which it decides whether to honor the contracts or default (post-harvest) is crucial for our test to detect strategic default in Section III.

II.C. Contractual Practices I: Loans (and Lender)

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).

Our lender provides working capital loans backed by forward contracts between the mill and buyers with whom the lender also has a business relationship. The lending model is illustrated in Figure III. The lender advances funds up-front and is then directly repaid by the buyer once the mill delivers the coffee.[15] The lender's contract terms and portfolio of clients appear to be broadly representative of the markets in which the lender operates (see Figures A1, A2 and A3 in the Appendix).

[15]Loan amounts vary between 40% and 70% of the value of the sales contract used as collateral depending on a comprehensive scoring system. The lender disburses loans progressively through smaller instalments and monitors sourcing of coffee from farmers to limit the possibility of loan diversion. The loan requires the payment of principal plus interest. The interest rate is never contingent on coffee prices.

II.D. Contractual Practices II: Sale Contracts

The working capital loan is disbursed on the back of one or more forward sales contracts. 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 **fixed** price and **differential** price (or price to be fixed (PTBF)) contracts.[16]

In fixed price contracts the price is agreed upon at the time the contract is signed. Fixed price contracts provide insurance against price fluctuations but leave parties exposed to counterparty risk. For example, 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.[17] In differential price contracts, instead, 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 counterpart 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.[18]

In the data, we observe a roughly equal split between fixed and differential contracts (45% and 55% of contracts respectively). The relative share of contract types in

[16] The 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).

[17] The reverse would be true for buyers. This, however, doesn't happen in our data. We discuss why that is the case in Section III.D.

[18] Fixed price contracts do not completely remove price risk, as international prices might be passed-through to farmers. Exporters for the most part lack access to hedging instruments and thus limit the remaining price risk by timing production and sourcing decisions accordingly. Conversely, differential contracts transform outright price risk into differential price risk. Although differential price risk is inherently lower, it is not zero, since the fixed differential specified in the contract cannot perfectly track the evolution of actual market differentials for coffee of specific origin/quality.

the lender's portfolio has remained fairly constant over time (see Figure A4 in the Appendix). Nearly 30% of loans are backed by a mix of fixed price and differential price contracts, typically signed with different buyers. 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).[19]

II.E. Contract Default

The next section develops and implements a test to detect strategic default. To implement the test we need a definition of contractual breach (in short, default). The test relies on the idea that following unanticipated rises in the world price a mill on a fixed price contract (but not on a differential contract) will be tempted to default on the forward sale contract. Because the data come from the lender, however, we do not directly observe breaches to the forward sale contract. In the spirit of forensic economics we use a number of different indicators of (attempted) default on the forward sale contract.

Figure IV illustrates and quantifies the different observable cases of contractual nonperformance. Consider the following timing of events (also used in the model in Section IV). At the time of scheduled delivery, the mill observes market conditions. It first decides whether to honor the sale contract and sell to the buyer or whether to search an alternative buyer and attempt to default. In the first case, the original buyer on the contract repays the loan on time. This happens in 78% of the cases. If, however, the mill searches for an alternative buyer, the mill might not find one. In this case the original buyer on the contract still repays the loan, but late. This happens in 8% of the cases. If the mill searches for, and finds, an alternative buyer the mill then defaults on the sale contract. At this point, the mill decides whether to also default on the loan or not. If the mill decides to default, the loan is not repaid. This happens in 1% of cases. If,

[19] Fair Trade contracts specify a differential price above a fixed floor price and thus are classified as differential for the purpose of our analysis.

instead, the mill does not default, the loan is repaid late by either the mill directly or by a different buyer. This happens in 13% of the cases (12% on time and 1% late).

Our baseline measure of default includes outright defaults as well as being late in repaying the loan (the grey-shaded end nodes in Figure IV). The baseline definition has two advantages: **i)** it is a directly observable measure of non-performance on the loan part of the bundle and is standard in the literature on loans; **ii)** under the timing of events described above, Figure IV clarifies that the baseline definition captures the majority of instances in which the mill attempted to default against the buyer.

The baseline definition makes two types of errors. First, defaults and late repayments can occur for reasons other than opportunism. This is of course unavoidable and the empirical tests in Section III bound the likelihood of these occurrences. Second, default on the forward sale might not take time, in which case the loan is repaid on time directly by the mill or by an alternative buyer (the dash-boxed node in the Figure). By not including those instances, the baseline measure underestimates strategic default. Section III thus reports results in which our empirical tests are conducted using all these alternative indicators of default as well as their union.[20]

To gain a better understanding of contractual non-performance, and make sure that our default classification resonates with the actual experience of practitioners, we conducted qualitative interviews with several loan officers at the partner institution and with both buyers and mills. Conversations with loan officers reveal that direct repayment from the mill or having a new buyer repaying the loan is likely indicative of default on the forward sale contract with the original buyer. Although the lender might exert some pressure on the mill to honor the original forward sale contract, when that is not possible the lender will try to get the loan repaid in a number of different ways, including by pairing up with alternative buyers.[21]

[20]Note that when all indicators of potential default against the buyer are considered, in the vast majority of cases the loan is eventually repaid. For simplicity, we refer to all considered instances as "default".

[21]A loan officer stated that *"This is one thing we try to make people understand: if the buyer doesn't make the payment it doesn't mean you're off the hook for the money. So generally they could instruct other buyers to send the money, but it has to be tied to a contract, so they're not going to say to some*

Similarly, several buyers described instances in which, following increases in prices, suppliers attempted to renegotiate fixed price contracts. Whether the buyer accepts the renegotiation or not is down to individual circumstances, including the extent to which the buyer would suffer from a cancelled shipment. Attempted renegotiations are unambiguously perceived as a sign of a non-reliable supplier.[22]

More broadly, what all conversations have confirmed is that, although it is standard practice in the industry to write formal contracts to obtain loans and accompany shipments and payments across borders, these contracts are next to impossible to enforce in court or in international arbitration in case of default. In practice, the loss in reputation and future business from the buyer and the lender is the main deterrent towards strategic default.[23]

II.F. Data Sources

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 and other data. Additionally, we conducted original field and internet searches

buyer 'hey buyer, can you send some money to [the lender]'...but yes, sometimes they might replace a contract...".

[22] For example, one buyer told us: *"yes we do renegotiate, we don't like to do it - and it depends on who and how they come to you because coffee enterprises are so dramatically different. So, you'll have some where you have no professional management at all [...] and then there's places that have a bunch of people with MBAs [...]. So it really depends on how they approach you and who they are to renegotiate pricing when the price does go up. But they definitely do it, and some do it more than others. And we do accept it in certain circumstances. Because, you can either try to enforce the contract, which is almost impossible to do; or you can say I don't want your coffee and stop buying coffee from them but that's not always a good choice; or you can accept it and so I would say the majority of the time we accept it and sometimes we say no, I'm sorry."*

[23] The buyers are located in western countries and therefore the lender could take a buyer that accepted coffee from the mill but did not repay the loan to the lender to court in a country with strong institutions. The buyer defaulting on the lender is thus not a concern.

to assess the extent of continued activity among all mills that have defaulted in the sample. We also collected detailed options data from Bloomberg for all start/end dates of all contracts to calculate implied coffee price volatilities for each contract. Appendix B provides further details on data construction, cleaning, matching and these other data sources.

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 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.

III. Detecting Strategic Default

This Section exploits contract-specific unanticipated international coffee price movements to detect strategic default. The next Section calibrates a model of pre-financing agreements to quantify the inefficiencies resulting from strategic default.

III.A. Conceptual Framework and Empirical Strategy

The Decision to Strategically Default Consider a risk-neutral mill deciding whether to honor or default on a forward sales contract.[24] At a certain date t the mill has agreed to deliver to the buyer q_c units of coffee at price p_c at a later date $t' > t$. There are

[24]The framework calibrated in Section IV embeds this decision into an optimal contracting model with a risk-averse mill signing both a sale and a loan contract.

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 spot market price in t', p_w, plus a differential Δ_c, i.e., $p_c = p_w + \Delta_c$.

At the time the mill takes the decision, all payments to farmers have been made (Figure II) and thus only revenues and future continuation values matter for the decision. Let $\mathbf{U^R}$ and $\mathbf{U}^D$ be the discounted value of future payoffs following a delivery and a default respectively. The incentive compatibility constraint is thus given by:

$$q_c \times p_c + \delta \mathbf{U^R} \leq q_c \times p_w + \delta \mathbf{U^D}. \tag{1}$$

Denote with $\mathbf{V} = \mathbf{U^R} - \mathbf{U^D}$ the **value of the relationship**.[25] The incentive compatibility constraint can be rewritten as

$$\delta \mathbf{V} < \begin{cases} (p_w - p_c)q_c & \text{if contract is } \textbf{fixed price} \\ \Delta_c q_c & \text{if contract is } \textbf{differential} \end{cases}$$

The key idea behind the test for strategic default is thus that (unanticipated) increases in the world price of coffee p_w increase the likelihood of default on fixed price contracts but not on differential contracts.

Empirical Strategy: A key challenge to implement the test is to obtain contract-specific exogenous variation in the temptation to default, $(p_w - p_c)q_c$. First, contract terms p_c and q_c are endogenously chosen by contracting parties and are thus potentially correlated with other factors that affect the likelihood of default. Furthermore, the spot market price at delivery, p_w, only varies over time. To implement the test we would instead like to control for time-varying factors that might also affect the likelihood of default.

[25] In our context **V** bundles the relationships with both the buyer and the lender, as further discussed in Section IV.A

Our strategy relies on building contract-specific measures of unanticipated increases in the world price of coffee p_w. In futures markets the price quoted at the closing date t for a future delivery date t' gives us parties' expectations about market conditions at delivery, $\mathbf{E}\left[p_w^{t'}|t\right]$. Using the New York 'C' Arabica coffee price at the scheduled shipment date, $p_w^{t'}$, we construct for each contract between mill m and buyer b signed at date t for delivery at date t' a measure of price surprise:

$$(2) \qquad P_{mbtt'} = \frac{p_w^{t'}}{\mathbf{E}\left[p_w^{t'}|t\right]}.$$

Price Surprise, Contract Default and Contract Type: A First Look Figure V shows the relationship between contract specific price surprises, $P_{mbtt'}$, and loan defaults. The histogram shows the distribution of $P_{mbtt'}$ in the sample. The Figure separates defaults into fixed price contracts (solid) and differential price contracts (dashed). We use the baseline definition of default (i.e., a contract is in default if there is written-off, restructured or has no payments after ninety days from its maturity date). The Figure shows that fixed price contracts display sharp increases in defaults associated with unexpected surges in world coffee prices. In contrast, there is no relationship between unexpected surges in world coffee prices and default when the contract is on differential.

The differential relationship between price surprises, $P_{mbtt'}$, and default across contract types is consistent with strategic default.[26] The evidence, however, could be driven by confounding factors. Figure A4 shows that the distribution of contract types has been relatively stable during the sample period despite significant swings in world coffee prices and volatility (Figure A6). Figure A7 shows that the distribution of price surprises across contract types are similar (and not statistically different from each other).

[26]Endogeneity of price surprises to mill's defaulting behaviour isn't a concern. Mills are small relative to the global market. Furthermore, the asynchronous timing of harvest seasons across countries (Figure A5 in the Appendix) implies that when, say, a Costa Rica mill is deciding whether to default, world coffee prices move because of development in other countries, e.g., news of a frost just before the upcoming harvest season in Brazil. In any case, we report below specifications that control for time fixed effects thus identifying off idiosyncratic variation induced by the timing and length of contracts signed.

We now proceed to an econometric investigation of the strategic default test that controls for additional potential confounders.

III.B. Baseline Test for Strategic Default:

Baseline Specification and Results We consider variations on the specification

$$D^{t'}_{lmbt} = \alpha_0 + \alpha_1 P_{mbtt'} + \lambda_m + \gamma_b + \mu_t + \varepsilon_{lmbtt'} \tag{3}$$

where $D^{t'}_{lmbt}$ is a dummy taking value equal to one if mill m is in default on loan l backed by buyer b closed at t and maturing at t'. The main regressor of interest is the price surprise $P_{mbtt'}$ defined in equation (2).[27] We control for time-invariant mill and buyer characteristics by including the relevant sets of fixed effects, λ_m and γ_b. We control for time effects, μ_t, and exploit asynchronous timing of the harvest season across countries in the sample to also control for country-specific season and seasonality effects.[28] Finally, $\varepsilon_{lmbtt'}$ is an error term arbitrarily correlated across observations for the same mill m.

Table II reports results from variations of this specification. Column 1 presents OLS estimates for the fixed-price sample (on which we expect an effect). A 10% increase in the world coffee price is associated with a three percentage point increase in the default rate. Columns 2-5 explore alternative specifications. To account for the asynchronous harvest seasons across countries, we allow the month and year 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. Column 4 controls for the length of the loan (in days) which could correlate with the magnitude of

[27] The 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.

[28] 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. For 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.

the price surprise. Column 5 includes year-month fixed effects and thus only exploits variation that exists between loans that were signed in the same month (i.e., identifying the effect mostly from variation in loan length). All specifications produce very similar estimates to Column 1.[29]

We expect that positive price surprises are associated with defaults only on the fixed price contracts. Columns 6-8 explore 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 magnitude smaller than the one estimated on fixed price contracts. Results for the corresponding specifications in Columns 2 to 5 are similar. Columns 7 and 8 explore a difference-in-differences specification at the loan level rather than the contract level. A fixed price loan is defined as a loan backed by a majority of sales contracts that are fixed price. In this specification the $P_{mbtt'}$ variable represents the effect of a price surprise on differential contracts. Again we estimate effects that are very close to zero. The effect on fixed-price contracts is confirmed in columns 7-8 when we look at the interaction between the price ratio $P_{mbtt'}$ 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.[30]

Additional Controls and Specifications These baseline results are robust to a variety of robustness tests reported in Online Appendix C. Specifically, results are robust to alternative definitions of late repayment (Table A2), to alternative thresholds to define fixed-price loans (Table A3), to alternative clustering strategies (Table A4) and to additional controls (Table A5). In particular, Column 3 in Table A5 explores a specification that controls for all independent sources of variation in Columns 4 and 5 in Table II at

[29] We discuss below further robustness tests.

[30] The results show that at the mean expected price surprise, fixed price contracts are less likely to default. This is consistent with the model in the next Section.

once (i.e., month-year fixed effects and contract length). Although we might be concerned that this specification identifies the effect from unreasonably few observations, we show that results are robust. Columns 4 and 5 in Table A5 explore an option-view of the decision to strategically default by recovering, and then controlling for, the implied volatility in coffee prices at the time the contract is signed. Results are, again, robust.

Functional Forms and Outliers Figure V might give the impression that the results in Table II are driven by a handful of extreme events with particularly large price hikes. While this does not necessarily undermine the analysis, it may be inconsistent with the linear specifications in Table II. Figure VI explores non-linear effects of price surprises. We define price surprises as a binary variable taking value equal to 1 if the price surprise is above a certain threshold z and 0 otherwise. We re-estimate the baseline regressions using definitions of z ranging from 130% (i.e., approximately 40% of the sample witnessed a price surprise) to 160% (only 17% did). The Figure reports estimated coefficients and shows that the results are remarkably stable in the definition of price surprise. This is because a lot of the identification in the data comes from the increasing portion of the default curve over the range 150% to 180% of price surprises (see, again, Figure V.)

We also explore robustness to outliers. Table A6 reverts to the linear baseline specification. The four columns in the Table report results removing observations in the top 1%, 5%, 10% and 25% of observed price surprises from the sample. Once again, results are remarkably robust: it is only when removing 25% of the observations in the sample that the linear specification delivers a coefficient that, while positive, is not statistically significant at conventional levels. The Table thus confirms that results are not driven by few outliers.

III.C. Ruling Out Input Price Pass-Through: an Event-Study Approach

The comparison between fixed price and differential contracts strengthens the case that default is a consequence of unexpected price surprises. The comparison, however, is not sufficient to establish the strategic nature of the observed defaults. A potential challenge in interpreting the results as evidence of strategic default is that unexpected increases in world coffee prices could be passed-through to farmers, raising input costs for mills and forcing mills on fixed-price contracts (for which revenues do not increase), but not those on differential ones (for which they do), to default.[31] To rule out this possibility, we take advantage of the stark separation between the timing of production and contract execution documented in Figure II. Using audited monthly cash flow data, the Figure shows that **1)** the vast majority (more than 90%) of payments to farmers occur during the harvest season; and **2)** the majority of forward sales contracts are executed after the end of the harvest season (but before the following harvest). That is, decisions to honour or default on the contract after harvest happen at a time in which all payments to farmers have been made. We thus isolate strategic default by conducting an event study that considers only price increases that occur after the end of harvest. After the end of harvest season, once cherries have been sourced, processed and paid for, 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. On the sample of contracts that are due for delivery after the end of harvest, the event study compares contracts that are due for delivery after harvest, just before

[31]The monthly cash-flow data from the audited financial accounts report the amount disbursed to farmers but not the volume sourced. It is thus not possible to compute prices paid to farmers and directly verify the extent of in-season price pass-through. Our understanding, however, is that in most countries the international price is passed-through to farmers almost fully and quickly. That being said, Table A7 in the Appendix shows that price surprises are not associated with lower mill's seasonal profits suggesting that mills are not forced to default due to higher costs. Table A11 shows that out of season price jumps do not correlate with the average price paid by mills to farmers, thus lending support to the identifying assumption underpinning the event study. Finally, Table A12 shows that price increases do not correlate with prices paid to farmers and with mill profits in the following season. This rules out the possibility that the mill strategically defaults in anticipation of future losses following a price surprise.

and just after a sudden price increase.[32]

Table III implements the event study (see Figure VII for an illustration). We separate price increases into ones that happened in-season and out-season. An event is defined as a weekly price increase of at least 3.0%. 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 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 II. In each case we find a large and statistically significant increase in the default rate of about 10-15%.[33]

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 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 suggests that strategic behavior is an important source of defaults in this market.

[32]Meanwhile, in-season price increases could 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.

[33]We also consider smaller increases of between 1%-2.5% price increases. The resulting estimates are actually quite similar (Table A11). 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 controls for in-season price increases in the specification (see Table A12).

III.D. Understanding Contractual Defaults

Untangling the Lender from the Buyers The evidence presented so far makes particular sense if mills were defaulting on buyers. Given that our data come from the lender, however, we do not have a direct measure of contractual breach on the forward sale contract. In the style of forensic economics, we use late repayment (and much rarer, defaults) on the loan to infer (attempted) default against the buyer.[34] Mills, however, might default against the buyer on the contract but still repay the loan on time, either directly or through an alternative buyer. Figure IV suggests that these cases are as frequent as all instances of defaults according to the baseline definition. Omitting to consider that such behaviour might also be indicative of strategic default would underestimate the extent of strategic default.

We thus explore robustness of our results to alternative definitions of default. Figure VIII shows that the main pattern in Figure V is confirmed in the raw data when considering alternative definitions of default. Table A9 reports both the baseline specification in Table II (Panel A) and the event study specification in Table III (Panel B) using alternative measures of default separately as well as their combination. Columns (1) and (5) consider loans backed by fixed and differential contracts respectively and define default to be the case in which any party (original buyer, a different buyer, or the mill) repays the loan late or the loan is defaulted against (baseline definition). Positive price surprises are associated with a higher likelihood of default on loans backed by fixed contracts but not on loans backed by differential contracts.

The next two sets of columns directly consider behavior consistent with the mill defaulting against the buyer, but not the lender. Columns (2) and (6) define default as whenever a different buyer from the one originally on the contract repays the loan. Results again show that price surprises increase the likelihood of default on fixed contracts

[34]Under the timing assumptions of the model calibrated in Section IV, the baseline definition of default captures all instances in which the mill attempted to default against the buyer (see Figure IV). The baseline definition is standard in the literature on loans and is a direct measure of contractual non-performance on the loan part of the contract bundle.

but not on differential contracts. Columns (3) and (7) define default as whenever the mill directly repays the loan. Once again, positive price surprises increase the likelihood of default on fixed contracts but not on differential contracts. Finally, columns (4) and (8) define a contractual default as whenever any of the behaviours separately considered in the three previous sets of regressions is observed. Positive price surprises are associated with higher likelihood of default on fixed price contracts but not on differential contracts.

In sum, from this analysis we conclude that, if anything, by focusing on directly observed defaults against the lender our baseline definition underestimates the extent of strategic default in the data.

Why Do Buyers on Fixed-Price Contracts Not Default when Prices Decrease? The logic of the test would suggest that unanticipated **de**creases in the world price of coffee would increase the buyer's incentive to default on fixed price contracts. Figure V shows that this is not the case.[35] The main reason for this is that foreign buyers hedge against movements in the price of coffee. Once that is done, defaulting on a shipment from a supplier is more costly. According to an interviewed Director of Purchasing and Production "*If you default on your own contract, if you outright say - like - 'I'm not buying that,' you're losing money because you've already invested in your book in paper. And so, there's that double incentive that when you buy paper against your physicals, it shrinks your range of options.*" Once the buyer has hedged, defaulting on a delivery increases the risk of defaulting on another (enforceable) obligation.[36]

[35] Buyers could, of course, strategically default under other circumstances. We just do not have a test for that.

[36] Another reason, more specific to our context, is that the lender might be financing multiple suppliers of the buyer. Defaulting on a supplier might jeopardize the buyer's ability to source from other suppliers as well.

III.E. What Happens After a Default?

Relationship Termination 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 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 trade-off the short-run benefits of default against the long-run costs of jeopardizing valuable relationships with their partners. We now explore what happens to the mill following a default.

The probability of getting a new loan from the lender is lower following a default or a late repayment (see Table IV, Column 1). A possible interpretation is that the lender is less likely to supply loans following contractual non-performance. This could either be due to the buyer not offering contracts and/or to the lender updating beliefs about the borrower's reliability by observing her decision to default on the buyer. If loan supply drives the correlation, we would expect the mill to be punished more harshly if the late repayment is due to strategic default. Figure IX shows that, conditional on a default, the mill is indeed less likely to receive a loan following a default that happened at the time of a positive price surprise as opposed to defaults happening at times of no or negative price surprise. Furthermore, this is only true for the fixed-price contracts, and not the differential price contracts where we would not expect this same patten. This suggests the lender takes default after positive price surprises a sign of strategic default rather than difficulties in repaying the loan.

The correlation in Column 1 of Table IV could, however, also be driven by the mill's demand. For example, the mill could default when it anticipates either not getting or not needing a loan from the lender and/or a contract from the buyer in the future (e.g. they fail to receive a future loan not because they they are being punished, but because they are no longer in operation). Building on Figure IX we attempt to distinguish between the two hypothesis by instrument default with the price surprise. Column 2 shows a decent first stage (despite the few defaults) and Column 3 shows that the IV estimate is significantly larger than the OLS estimate in Column 1. The larger IV estimate is consistent with loan supply drying up after a strategic default. If future loan demand was lower following a default, instead, we would expect the OLS to be biased in the opposite direction.

The IV exercise should be interpreted cautiously. First, the IV does not separate the supply and demand channel if the demand channel is correlated with the price surprise. However, as noted above loans are eventually repaid in essentially all instances of default. So the possibility that the mill needs less loans in the future because it kept the defaulted balance is unlikely to be relevant. Second, the approach does not really identify any parameter of interest. For example, we do not know if the relationship termination is part of a repeated-game equilibrium in which the mill is being punished by the lender or is simply the result of the lender updating beliefs about the mill's reliability as a borrower. Column 4 explores an alternative approach which simply adds the price surprise to the baseline specification in Column 1. Consistent with Figure IX, we find that positive price surprises are associated with lower likelihood of getting a new loan. Furthermore, the estimate for default is lowered and is now statistically insignificant at conventional levels. This is consistent with the hypothesis that defaults that occur at times of positive price surprises (and thus, more likely to be strategic) are more likely to lead to relationship termination. A simple exercise in the spirit of Oster (2017) reveals that the magnitude of the change is consistent with strategic default being punished harshly (For details see Appendix C).

The (Surprising) Fate of Defaulting Mills There is a completely different, albeit indirect, type of evidence supporting the idea that indeed a large share of the documented defaults are strategic. If positive price surprises simply induce mills to default or go bust, we would expect at least some of the defaulting mills to actually go bust (unless **every single** instance of observed default was marginal, in the sense that it kept the mill alive). In contrast, if many defaults are strategic, we expect mills to be still operating in the years following a default. We thus conduct original field and internet searches for all instances of observed defaults in the sample to check whether the defaulting mills are still active or, instead, have gone bust. Table A13 in the Appendix shows that almost none of the mills that have defaulted have gone bust. The evidence is thus inconsistent with the hypothesis that defaulting mills simply suffered financial difficulties.[37]

III.F. How many defaults are strategic?

Our preferred estimate is that a share between 42% and 59% of the defaults observed on fixed price contracts are strategic.[38]

The baseline estimates provides an 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. We observe a default rate of 14.5% when price surprise are greater than 1 and 6.8% otherwise, so the difference, 7.7%, we attribute to the strategic dimension we can isolate. We compare this to the overall default rate on fixed price contracts (13%). Doing so provides the upper bound of $7.7\%/13\% = 59\%$.[39] We

[37]While the documented extent of mill survival might be surprising, fixed assets invested in the country-side are specific to the coffee business and hard to liquidate. This might facilitate mills' staying in business. Note that we do not intend to imply that the decision to strategically default is necessarily optimal from the mill's viewpoint. First, it is quite possible for someone to strategically breach a contract to enjoy a short-term gain as a result of greed, envy, or other emotions that lead them to make sub-optimal decisions. Second, the decision might have been taken by the mill's manager in his/her interest, not in the interest of stakeholders. In this case, we would expect the defaulting management to be kicked-out by the mill's shareholders. Unfortunately we were not able to gather systematic evidence on this.

[38]Recall that with by strategic default we specifically mean defaults caused by ex-post moral hazard, rather than defaults caused by general wilful acts. Under this broader meaning of the word strategic, we identify a clear-cut case of strategic default; other defaults may be strategic or not (with a court potentially still having case for exoneration).

[39]If we construct the same upper bound in the same way using the union of all default definitions from Table A9, the comparable value is 57%.

interpret this as an upper bound because we allow that some of this default to be due to debt over-hang associated with increased costs from pass-through of world price increases to farmers.

The event study provides the lower bound. Assume a constant effect of price surprises on default for week-over-week price change in the range 1-3% (Table A11) 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.[40] 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 occasionally induce strategic default.

IV. Quantifying the Costs of Strategic Default

The reduced-form estimates in the previous Section suggest that a large share (about half) of observed defaults are likely to be strategic. How large are the efficiency losses caused by (the possibility of) strategic default? To answer this question we need to understand how parties structure their contractual arrangements in anticipation of the possibility of strategic default. To make progress we thus need to explicitly embed the strategic default decision into an optimal contracting framework in which parties endogenously chose contractual arrangements based on their circumstances. After presenting such a framework and discussing its main assumptions, we calibrate it and present results that quantify the costs associated with strategic default and the value of informal enforcement in this market.

[40] If we construct the same lower bound in the same way using the union of all default definitions from Table A9, the comparable value is 36%.

IV.A. Theoretical Framework

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 I. 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 $\overline{p}_w$. 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.

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 $\gamma = \rho \times a$.[41]

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

Sales Contracts The sales contract specifies the delivery of q_c units of coffee at date $t = 3$. The price is p_c in **fixed price** and $p_c = p_w + \Delta_c$ in **differential** contracts. 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

$$\int_{p_w} \mathbf{I}[p_w] q_c (p_w - p_c) dF(p_w) \geq 0. \tag{4}$$

When the contract is on differential the buyer's participation constraint collapses to $\Delta_c \leq 0$.[42]

[41] For 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.

[42] A delivery failure imposes no cost on the risk-neutral buyer. Relaxing the assumption does not alter

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

$$L(1+r) \leq \int_{p_w} \mathbf{I}[p_w] \min\{p_c q_c, D\} dF(p_w) \tag{5}$$

Default and Enforcement After p_w is realized the mill decides between honoring the forward sale contract or sell the contracted coffee q_c to an alternative buyer at price p_w and default.[43]

The mill is in a relationship with both the buyer and the lender. We bundle the two together. We collapse the continuation of the relationship onto static parameters (see MacLeod 2007). We denote with $\mathbf{U^R}$ the discounted value of future expected profits when continuing the relationship with the buyer **and** the lender, $\mathbf{U}^D$ the discounted 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 the mill's behaviour. Learning about $\mathbf{V}$ is also necessary to perform counterfactuals.

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 $\pi^D(p_w) = p_w q_c$ if the mill defaults and sells the coffee on the spot market at price p_w. Assuming the mill's utility function is given by $u(\cdot)$,

the qualitative predictions.

[43] Appendix A considers a more elaborate side-selling process that underpins the baseline definition of default.

with $u' > 0$ and $u'' \leq 0$, and normalizing $\mathbf{U}^D$ to zero, expected utility is given by

$$\mathbf{E}[\Pi] = \int_{p_w} u(\mathbf{I}[p_w]\pi^R(p_w) + (1 - \mathbf{I}[p_w])\pi^D(p_w))dF(p_w) + \mathbf{I}[p_w]\mathbf{V} \tag{6}$$

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 (4) and (5).[44]

First Best The contractual outcome is illustrated in Figure X. 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 produce under a differential contract, $q_c = q_D^*$. A differential contract leaves the mill exposed to uninsured price risk. To the extent that risk aversion reduces investment, this lowers the mill's desired production and coffee purchases.

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

[44]The 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.

observing realized world prices p_w the mill defaults on the contract if[45]

$$\delta \mathbf{V} \leq u\left(\pi^D(p_w)\right) - u\left(\pi^R(p_w)\right). \tag{7}$$

Contract Choice The possibility of strategic default introduces a trade-off between insurance and counter-party risk. 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. The resulting trade-off is illustrated in Figure X. 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 pledge-able 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.[46]

The model implies that relationships with higher value $\mathbf{V}$ sign fixed-price contracts that leave them exposed to strategic default. Since parties adjust the contractual form accordingly, strategic default can be detected only on fixed price contracts, and 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.[47]

[45] The following incentive constraint adapts the incentive constraints in Section III to the case of a risk-averse mill and using the notation distinguishing the loan from the sales contract. As in the simpler framework, higher realizations of world prices p_w increase the likelihood of default on fixed-price contracts. Higher realizations of p_w do not affect the likelihood of default under a differential contract if $u(D) \leq \mathbf{V}$. Otherwise, high realizations of p_w make the mill less likely to default.

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

[47] Appendix B provides empirical support for two predictions of the model. First, the model predicts that more valuable relationships (higher $\mathbf{V}$) sign fixed price contracts. Second, higher relationship value $\mathbf{V}$ decreases the effect of unanticipated increases in the world price on the likelihood of default. To

Discussion of Modelling Choices Before moving on to the calibration it is worth pausing to discuss our main modelling choices. A first key assumption is that the model takes the value of the relationship $\mathbf{V}$ as exogenous. $\mathbf{V}$ could be micro-founded in a variety of ways. For example, it could arise as part of a sub-game perfect equilibrium in which strategic default is deterred; or in a model featuring both adverse selection over the mill's type and strategic default.[48] For example, in Antras and Foley (2015) when an importer of the bad type is hit by a shock to her discount factor she can get away without paying the exporter precisely because contracts are not enforceable. The good type never defaults because she is always patient enough: continuation values are such that temptations to cheat are never too strong. With the primary goal of our framework being to guide the quantitative exercise we stick to the simpler modelling strategy and abstract from adverse selection. We cannot, and do not intend to rule out models with asymmetric information on types.[49]

A second assumption is that we bundle the mill's default decision against the buyer and the lender and only consider a combined relationship value $\mathbf{V}$. An alternative modelling strategy would be to endow the mill with two distinct decisions and relationship values: a relationship value with the buyer (meant to deter side-selling), and a relationship value with the lender (meant to deter default). There are two reasons why we prefer our current modelling choice. Our conversations with the lender (and the evidence on relationship termination) suggest that although the loan is eventually almost

see why note that under a fixed contract the likelihood of default is given by $P^F(\mathbf{V}) = 1 - F(u^{-1}(\mathbf{V} + u(p_c q_c - D))/q_c)$. The second prediction then follows from $u' > 0$ and $F'' < 0$ in the right tail of the price distribution. We proxy the value of the relationship $\mathbf{V}$ with measures of relationship history (past volumes of transactions, age) between the mill and both the lender and the buyer. We show that estimated relationship values $\mathbf{V}$ correlate with these measures of relationship history.

[48] For a pure adverse selection model to be consistent with the evidence, it would have to be the case that bad types default with a certain probability that depends positively on spot prices but only if the contract is a fixed-price contract. Furthermore, such a model would need additional assumptions to explain why a bad type occasionally defaults if on a fixed price contract if the temptation is large enough. Imperfect contract enforcement/strategic default offers the most natural micro-foundation for why a bad type on a fixed-price contract would (sometimes) default.

[49] Since our data come from the lender, the relationship with the buyer is measured with significant error in our data. As the panel dimension is critical to distinguish a model with types from alternative models generating relationship dynamics (Macchiavello and Morjaria 2015b) our data are not suitable to conduct such an exercise.

always repaid, the two relationships are not independent. The lender learns about a mill's reliability by observing how they behave with the buyer. Moreover, if we model the two decisions separately we would also need to estimate the costs of late repayment for the lender and failed delivery for the buyer. We prefer to stay away from introducing so many unobservables into the framework.

A third assumption is that defaulting mills can walk away with the loan. In equilibrium the loan is (almost) always repaid. One might thus question our assumption. Obviously, if the mill cannot walk away with the loan, relationship values will be overestimated. The extreme case in which the debt contract is fully enforceable provides a lower bound. A "lower-bound-to-the-lower-bound" is then simply given by subtracting from the estimated values the value of the loan. Although estimates would change, the qualitative conclusions from Table VI would not. A related issue is that, as noted in Bulow and Rogoff (1989), lending cannot be based on dynamic incentives alone: if the seller does not repay the loan, the seller no longer needs to borrow. One might thus question whether the relationship with the lender provides any deterrent against default. The logic in Bulow and Rogoff (1989), however, holds if upon defaulting the borrower cannot be excluded from perfect insurance and saving markets. The evidence in this paper suggests that these firms do not have access to well-functioning insurance markets. Weak governance might also prevent the firm from holding on to the money (e.g., a defaulting cooperative might have to leave some rents to the manager to prevent the manager from steeling the proceeding of not repaying the loan).[50]

IV.B. Calibration Strategy

The model is based on a limited set 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 **V**, the value of the relationship(s) with the buyer and the lender.

[50]Relational dynamics documented in Appendix B are consistent with slow build-up of relationships in which defaulting foregoes, or delays, future growth opportunities and is thus costly.

We pursue the following strategy. We "invert" the model and obtain an estimate of $\mathbf{V_i}$ for each loan (see figure XI for the distribution of $\mathbf{V_i}$). Specifically, given a set of parameters we find the $\mathbf{V_i}$ that rationalizes the observed 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 are intimately connected with $\mathbf{V_i}$ in the model (see Figure X). This makes the identification of $\mathbf{V_i}$ particularly transparent.[51] They are also recorded without error in the dataset.[52]

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 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.[53] Finally, the slope of the coffee cherries supply curve, η, is estimated from the RDD analysis presented in Appendix D. Table A15 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.

The calibrated parameters are reported in Table V. Two remarks are in order. First,

[51] 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.

[52] See Appendix A for details.

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

we need the financial statement data to construct γ_i (and estimate η). The calibration exercise can therefore only be performed on the smaller sample of loans for which we have the financial data. Second, while the distribution of price surprises and input costs are directly observed in the data, the mills' utility and cost functions are not. We parametrize both functions with simple functional forms that depend on one parameter only (α and η respectively). While we do not have strong priors on the appropriate functional forms, the chosen parametrization has the benefit that both parameters can be transparently recovered from relevant empirical moments (observed forward discounts for α and the RDD on loan size in Appendix D for η). To assuage concerns, we report sensitivity checks on the calibrated values of both α and η. We estimate $\alpha = 0.386$ and report results spanning the interval $\alpha \in [0.286, 0.486]$. We estimate $\eta = 0.6$ and report results spanning the interval $\eta \in [0.5, 0.7]$.

It is worth noting that we do not use actual defaults to calibrate the model (and, thus, the results are totally unaffected by extreme events). The calibration recovers relationship values $\mathbf{V_i}$ from interest rates (which reflect, *inter alia*, the likelihood of defaults) rather than from actual defaults. The continuous variation in interest rates allows us to recover, non-parametrically, the distribution of relationship values $\mathbf{V_i}$. Even a parametric approach would not perform well if we were to recover the relationship values $\mathbf{V_i}$ from the relatively few observed defaults. As a sanity check, however, Figure XI compares the estimated the relationship values $\mathbf{V_i}$ with bounds inferred from the actual observed decision to default or not. Despite using completely different sources of variation, the estimated relationship value $\mathbf{V_i}$ distributions display a significant overlap.

IV.C. Results

Estimated Relationship Values $\mathbf{V}$ *and Counterfactuals* The main results, alongside counterfactuals and sensitivity checks on the calibrated values of α and η are reported in Table VI. 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 XI).

The second row quantifies inefficiencies by comparing the predicted production volume with the implied first best volume (which can be analytically computed). This comparison also yields our main counterfactual: by how much would production increase if we removed strategic default?[54] 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, the relationship value $\mathbf{V_i}$ 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 accounts 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 severely financially constrained if they had

[54]In 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.

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 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 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 A14 and A15. 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.[55] 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 farmer's 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

[55]Note that the lender assigns higher scores to loans with differential contracts.

paid to farmers as a result of larger loans (Table A15, Column 4). These estimates suggests 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.

Correlates of Estimated Relationship Values **V** The calibration recovers estimates of relationship values. Appendix Table A18 projects the estimated relationship values on observables characteristics. There are three sets of observable characteristics that we would like to explore: **i**) relationship-level variables; **ii**) market-level variables; and **iii**) country-level institutional variables. Before describing the results it is important to stress that these can only be interpreted as suggestive correlations rather than as establishing evidence in favour of certain (causal) determinants of relationship values.

With respect to relationship-level variables, columns 1 and 2 in Table A18 show that the estimated relationship values are positively correlated with measures of the past amount of business between the mill and the buyer and between the mill and the lender. These results match those in Appendix Tables A16 and A17 documenting that these measures of relationship histories also correlate with responses to incentives to strategically default and with contract choices as predicted by the model.

With respect to market-level variables, it would be interesting to explore correlations between estimated relationship values and the availability of alternative buyers and lenders in the market. One might hypothesize that availability of many alternative trading partners increases outside options and reduces relationship values (e.g. Macchiavello and Morjaria 2015a). On the other hand, a good reputation might be more valuable in a market with many partners that allow for further growth. Unreported specifications fail to find any robust correlation between the number of alternative buyers and lenders in the market and the estimated relationship values. Because our data come from one lender only, however, we only observe other lenders and buyers dealing with our lender's clients, rather than the entire market. We are thus unable to construct precise measures of the number of alternative partners and interpret the lack of correlation in the data.

Finally, Columns 3 and 4 explore country-level institutional variables. The Table documents that the quality of **debt** contract enforcement (from Djankov et al. 2008) negatively correlates with estimated relationship values. Our preferred explanation is that in countries in which it is harder to enforce debt contracts, lenders will be more reluctant to lend. A relationship with any lender, including our lender, is then likely to be more valuable. On the other hand, we fail to detect a correlation between the standard Doing Business measure of **commercial** contract enforcement in the country and the estimated relationship values. This confirms what interviewed buyers told us: the quality of contract enforcement in the mill's host country has little to do with the ability of the buyer to enforce the international transaction.

V. 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 develops a test to identify strategic default. 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 forward sale contracts in the international coffee market. 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. We isolate the strategic motive by focusing on unanticipated changes in market prices at the time of contract execution, after production decisions are sunk and suppliers have been paid.

Our preferred estimates suggest that a large share (around 50%) of observed contractual breaches are likely due to strategic motives. 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. Strategic default introduces a trade-off between insurance and counterparty risk. 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. The relevant missing market then varies across firms. 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 paper identifies strategic default and quantifies the costs generated by lack of contract enforcement. These costs appears to be significant. Perfect contract enforcement, however, is not achievable in practice and is thus not the correct policy benchmark. As pointed out in the theoretical literature (see, e.g., Baker, Gibbons and Murphy 1994) partial improvement in contract enforcement might either increase or decrease efficiency. So, while the general spirit of our results is that reducing contracting frictions **could** yield large payoffs, we would like to advocate in favour of a context-specific, one-does-not-fit-all, approach to policy recommendations to be drawn from our analysis.

More specific policy implications, however, can be drawn for developing countries aiming at improving exports, particularly in agricultural chains. 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. Our analysis suggests that imperfect contract enforcement reduces both the supply and the demand for hedging tools, even among relatively large exporters. Fostering contract enforcement and strengthening inter-firm relationships along supply chains can yield significant degrees of insurance and expand output. Furthermore, the existence of externalities along the domestic value chain suggests that strengthening contract enforcement for large exporters downstream might yield large payoffs upstream.

At a broader level, a striking aspect of 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). As there is limited evidence that small firms can bootstrap their growth (Hsieh and Klenow 2014) it is important to understand barriers to the operation of large firms. Further research to establish the form and extent through which contractual frictions hamper the operation of these firms in other contexts remains an important area for future research.

University of Toronto

London School of Economics

REFERENCES

Adams, William, Liran Einav, and Jonathan Levin, "Liquidity Constraints and Imperfect Information in Subprime Lending," The American Economic Review, 99 (2009), 49–84.

Antras, Pol, *Global Production: Firms, Contracts, and Trade Structure*: Princeton University Press (2015).

Antras, Pol and C Fritz Foley, "Poultry in Motion: A Study of International Trade Finance Practices," Journal of Political Economy, 123 (2015), 853–901.

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, "Subjective performance measures in optimal incentive contracts," The Quarterly Journal of Economics, 109 (1994), 1125–1156.

——— "Relational Contracts and the Theory of the Firm," The Quarterly Journal of Economics, 117 (2002), 39–84.

Banerjee, Abhijit and Esther Duflo, "Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program," Review of Economic Studies, 81 (2014), 572–607.

Breza, Emily and Andres Liberman, "Financial contracting and organizational form: Evidence from the regulation of trade credit," The Journal of Finance, 72 (2017), 291–324.

Bubb, Ryan, Supreet Kaur, and Sendhil Mullainathan, "Barriers to Contracting in Village Economies: A Test for Enforcement Constraints," mimeo.

Bulow, Jeremy and Kenneth Rogoff, “A constant recontracting model of sovereign debt,” Journal of Political Economy, 97 (1989), 155–178.

Burkart, Mike and Tore Ellingsen, “In-kind finance: A theory of trade credit,” American Economic Review, 94 (2004), 569–590.

Casaburi, Lorenzo and Jack Willis, “Time vs. State in Insurance: Experimental Evidence from Contract Farming in Kenya,” American Economic Review, forthcoming.

Chiappori, Pierre-André and Bernard Salanie, “Testing for asymmetric information in insurance markets,” Journal of Political Economy, 108 (2001), 56–78.

Dana, James D, “Advance-purchase discounts and price discrimination in competitive markets,” Journal of Political Economy, 106 (1998), 395–422.

Djankov, Simeon, Oliver Hart, Caralee McLiesh, and Andrei Shleifer, “Debt enforcement around the world,” Journal of Political Economy, 116 (2008), 1105–1149.

Djankov, Simeon, Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer, “Courts,” The Quarterly Journal of Economics, 118 (2003), 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.

Ellingsen, Tore and Eirik Gaard Kristiansen, “Financial Contracting Under Imperfect Enforcement,” The Quarterly Journal of Economics, 126 (2011), 323–371.

Fafchamps, Marcel, *Market institutions in sub-Saharan Africa: Theory and evidence*: MIT press (2003).

Giannetti, Mariassunta, Mike Burkart, and Tore Ellingsen, “What you sell is what you lend? Explaining trade credit contracts,” The Review of Financial Studies, 24 (2011), 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, 30 (1987), 369–398.

Greif, Avner, “Contract enforceability and economic institutions in early trade: The Maghribi traders’ coalition,” American Economic Review, 83 (1993), 525–548.

Grossman, Sanford J and Oliver D Hart, “An analysis of the principal-agent problem,” Econometrica, 51 (1983), 7–45.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, “The determinants of attitudes toward strategic default on mortgages,” The Journal of Finance, 68 (2013), 1473–1515.

Hart, Oliver, “Hold-up, asset ownership, and reference points,” The Quarterly Journal of Economics, 124 (2009), 267–300.

Hart, Oliver and John Moore, “Default and renegotiation: A dynamic model of debt,” The Quarterly Journal of Economics, 113 (1998), 1–41.

Hölmstrom, Bengt, “Moral hazard and observability,” The Bell Journal of Economics, 10 (1979), 74–91.

Hsieh, Chang-Tai and Peter J Klenow, “The life cycle of plants in India and Mexico,” The Quarterly Journal of Economics, 129 (2014), 1035–1084.

Hsieh, Chang-Tai and Benjamin A Olken, “The missing “missing middle”,” The Journal of Economic Perspectives, 28 (2014), 89–108.

Imbens, Guido and Karthik Kalyanaraman, “Optimal bandwidth choice for the regression discontinuity estimator,” The Review of Economic Studies, 79 (2012), 933–959.

de Janvry, Alain, Craig McIntosh, and Elisabeth Sadoulet, “Fair Trade and free entry: Generating benefits in a disequilibrium market,” Review of Economics and Statistics, 97 (2015), 567–573.

Joskow, Paul L, "Price adjustment in long-term contracts: the case of coal," The Journal of Law and Economics, 31 (1988), 47–83.

Karlan, Dean and Jonathan Zinman, "Observing unobservables: Identifying information asymmetries with a consumer credit field experiment," Econometrica, 77 (2009), 1993–2008.

Klapper, Leora, Luc Laeven, and Raghuram Rajan, "Trade credit contracts," Review of Financial Studies, 25 (2012), 838–867.

Klein, Benjamin, "Why hold-ups occur: the self-enforcing range of contractual relationships," Economic inquiry, 34 (1996), 444–463.

Kranton, Rachel and Anand V Swamy, "Contracts, hold-up, and exports: textiles and opium in colonial India," The American Economic Review, 98 (2008), 967–989.

Lacker, Jeffrey M and John A Weinberg, "Optimal contracts under costly state falsification," The Journal of Political Economy, 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 Josepa Miquel-Florensa, "Vertical integration and relational contracts: Evidence from the Costa Rica coffee chain," mimeo.

Macchiavello, Rocco and Ameet Morjaria, "Competition and Relational Contracts: Evidence from Rwanda's Coffee Mills," mimeo.

——— "The value of relationships: evidence from a supply shock to Kenyan rose exports," The American Economic Review, 105 (2015b), 2911–2945.

MacLeod, Bentley W, "Reputations, relationships, and contract enforcement," Journal of economic literature, 45 (2007), 595–628.

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, 104 (2014), 2830–2857.

Oster, Emily, "Unobservable selection and coefficient stability: Theory and evidence," Journal of Business &; Economic Statistics, 0 (2017), 1–18.

Rampini, Adriano A and S Viswanathan, "Collateral, risk management, and the distribution of debt capacity," The Journal of Finance, 65 (2010), 2293–2322.

Roberts, Michael R and Amir Sufi, "Financial contracting: A survey of empirical research and future directions," Annual Revue of Financial Econonomics, 1 (2009), 207–226.

Shleifer, Andrei and Daniel Wolfenzon, "Investor protection and equity markets," Journal of Financial Economics, 66 (2002), 3–27.

Startz, Meredith, "The value of face-to-face: Search and contracting problems in Nigerian trade," Mimeo.

Varangis, Panos and Bryan Lewin, "Approaches to managing coffee price risks," Agricultural Commodity Markets And Trade: New Approaches to Analyzing Market Structure And Instability, 282.

Main Figures

Negative Cash Flow **Positive Cash Flow**

$t = 0$	$t = 1$	$t = 2$	$t = 3$
Contract signed	Loan Disbursed	$p_w \sim F(p_w)$	Delivery/default

Figure I: Timing of Events

Notes: The Figure illustrates both the actual timing of events during a typical harvest season as well as the timing of events in the model in Section IV. 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 ($t = 3$). Appendix A lays out a different timing of events in the decision to default. The mill first decides whether to search for an alternative buyer or not, and then defaults only if it finds one. This alternative timing of events underpins our baseline definition of contractual non-performance as well as the calibration in Section IV.

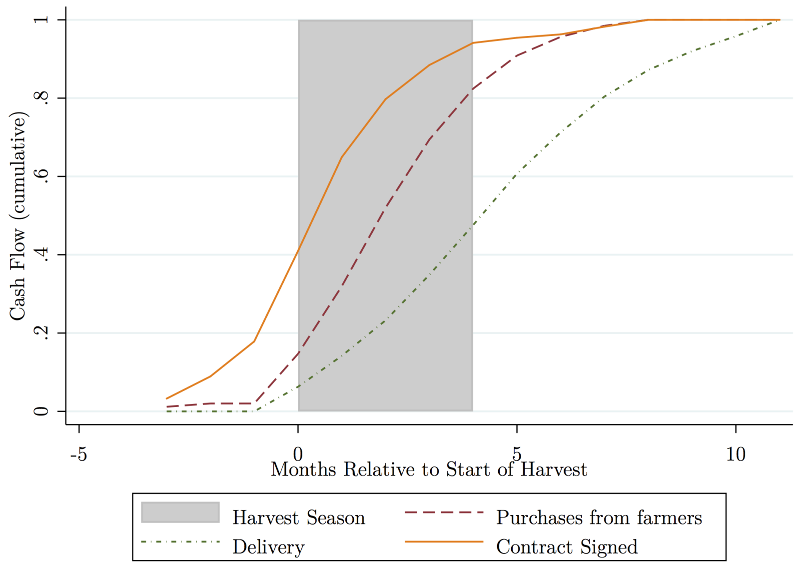

Figure II: Mills' Contracting and Cash Flow Profiles

Note: The Figure illustrates the timing of contracting and the cash flow profiles of coffee mills. The Figure documents the stark separation in the timing of production and contract execution that underpins our empirical strategy. We collect monthly cash-flows data from the audited financial records of the mills in our sample. The horizontal axis is the number of months from the beginning of the harvest season (at zero). The timing of harvest is asynchronous across countries (see Figure A5 in the Appendix). The Figure then reports the cumulative values of contracts signed, disbursements to farmers and scheduled coffee deliveries, averaged over seasons and mills in the sample. By the time harvest begins, mills have already signed (forward) contracts worth approximately 50% of overall production (contract line). These are the (type of) contracts used as collateral and that we study in this paper. Two key facts stand out. First, the vast majority (more than 90%) of payments to farmers occur during the harvest season (i.e., by month five after the beginning of harvest). This reflects the fact (confirmed by loan officers and mills surveys in Rwanda) that very rarely mills source coffee from farmers on long-credit or share profits (so called second-payments) at the end of the season. By the end of the harvest season, then, mills financial obligations with farmers are essentially fully executed. Second, sales and loans contracts are executed later. By the time harvest ends, only 50% of contracted deliveries have occurred (contract line). That is, about half of the contracted sales take place at a time in which all production decisions are sunk and contractual payments to suppliers have been executed. This separation between the time at which the mill purchases inputs from farmers (harvest) and the time in which it executes contracts (post-harvest) is crucial for our empirical strategy as further detailed in Section III.

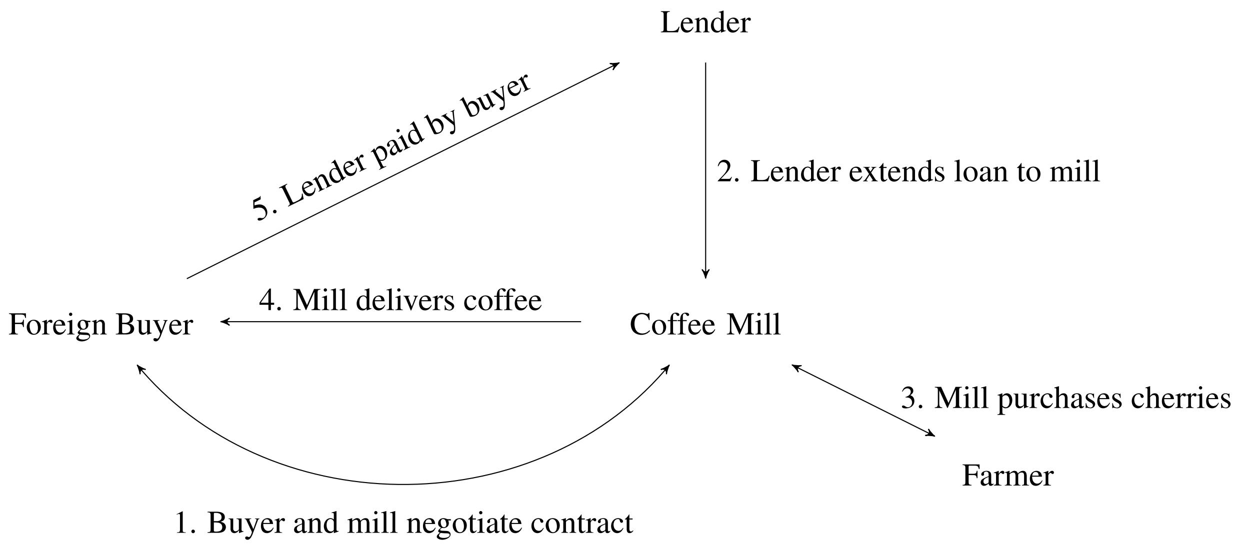

Figure III: Lending Model Under Normal Circumstances

Notes: This figure shows the lending model that the lender uses when everything goes as planned. Each step is numbered based on the sequence in which events 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, then process the cherries and deliver the agreed upon quantity to the buyer. The buyer then repays loan to the lender directly.

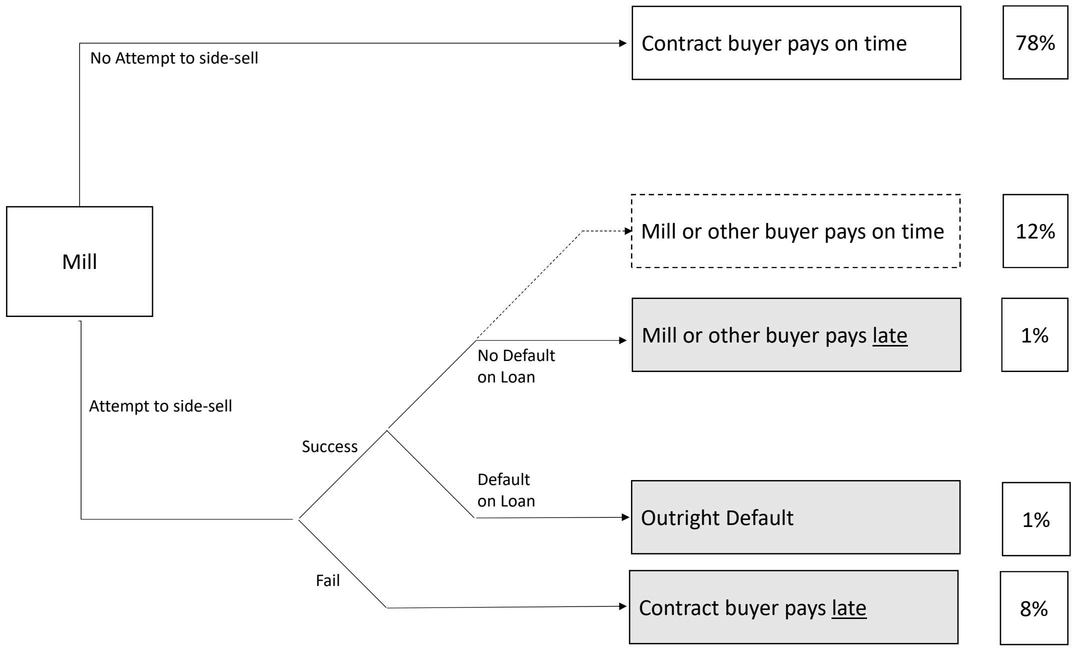

Figure IV: Contractual Non-Performance: Taxonomy and Baseline Definition

Notes: Figure IV illustrates and quantifies the different observable cases of contractual non-performance. We distinguish cases depending on **i)** whether the loan is defaulted, repaid late or fully repaid on time, and **ii)** who repays the loan (buyers on the original contract, other buyers, or the mill). To fix ideas, consider the timing of events used to calibrate the model in Section IV (see Appendix A for details). At the time the mill is supposed to execute the contract, the mill observes market conditions. It first decides whether to honour the sale contract and sell to the buyer or whether to search an alternative buyer and attempt to default. In the first case, the original buyer on the contract repays the loan on time. This happens in 78% of the cases. If, however, the mill searches for an alternative buyer, the mill might not find one. In this case the original buyer on the contract still repays the loan, but late. This happens in 8% of the cases. If the mill searches for, and finds, an alternative buyer the mill then defaults on the sale contract. At this point, the mill decides whether to also default on the loan or not. If the mill decides to default, the loan is not repaid. This happens in 1% of cases. If, instead, the mill doesn't default, the loan is repaid late by either the mill directly or by a buyer not originally on the contract. This happens in 13% of the cases (12% on time and 1% late). Our baseline measure of default includes outright defaults as well as being late in repaying the loan (the grey-shaded end nodes in the Figure). The baseline definition has three advantages: **i)** it rests on a directly observable measure of non-performance on the loan part of the bundle; **ii)** it is standard in the literature on loans; *iii)* under the timing of events described above, the baseline definition captures all instances in which the mill attempted to default on the forward sale contract. If defaulting on a forward sale contract does not take time, however, the baseline definition underestimates strategic default by omitting those cases in which the mill or an alternative buyer repay the loan on time (the dash-boxed node in the Figure). Section III reports results in which the empirical tests are conducted using all alternative indicators of default as well as their union.

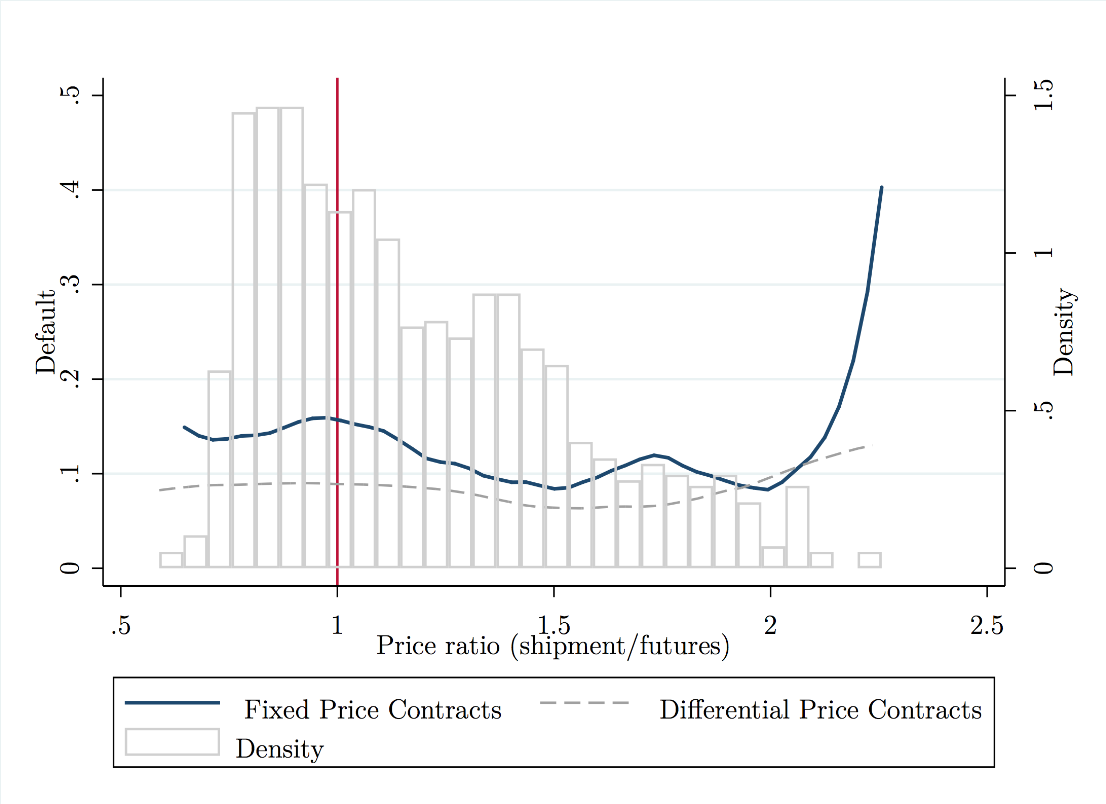

Figure V: Unexpected Price Increases and Contract Default

Notes: The Figure illustrates the main evidence underpinning the test for strategic default. 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 II. For a given price surprise, the solid line plots defaults (baseline definition: outright default and cases in which the loan is not yet fully repaid 90 days past-due) on fixed price contracts while the dashed line plots the same for differential price contracts. Analysis in Section III (see Figure VIII) and in Appendix C show that the patterns are robust to alternative measures of default. The figure shows that default is largely driven by fixed price contracts that experience large price surprises. There is no increase in defaults associated with price surprises among differential contracts. The highest default rates are among those that experienced world prices that were much larger than the price at closing. Further analysis in Section III shows that the patterns are not driven by outliers experiencing huge price surprises. Once controls are included, the bump in defaults rate at price surprises in the interval 1.5 to 1.8 is sufficient to detect a differential effect across contract types.

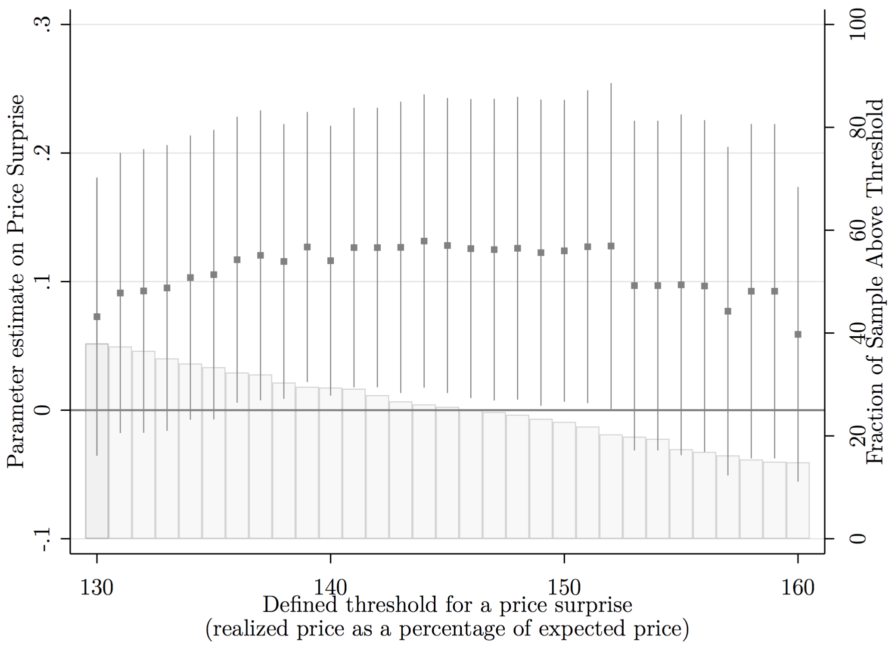

Figure VI: Baseline Test: Robustness to Non-Linearities and Outliers

Note: The Figure explores non-linear effects of price surprises and robustness to outliers. Figure V might give the impression that the results in Table II are driven by a handful of extreme events with particularly large price hikes. We thus define price surprises as a binary variable taking value equal to 1 if the price surprise is above a certain threshold z and 0 otherwise. We re-estimate the baseline regressions using definitions of z ranging from 130% (i.e., approximately 40% of the sample witnessed a price surprise) to 160% (only 17% did). The Figure reports estimated coefficients and shows that the results are remarkably stable in the definition of price surprise. This is because a lot of the identification in the data comes from the increasing portion of the default curve over the range 150% to 180% of price surprises (see, again, Figure V.) The Figure thus also suggests that outliers are not driving results. We explore outliers directly in Table A6. Reverting to the linear specification in Table II we show that results are robust when removing the top 1%, 5%, 10% and 25% of observed price surprises from the sample.

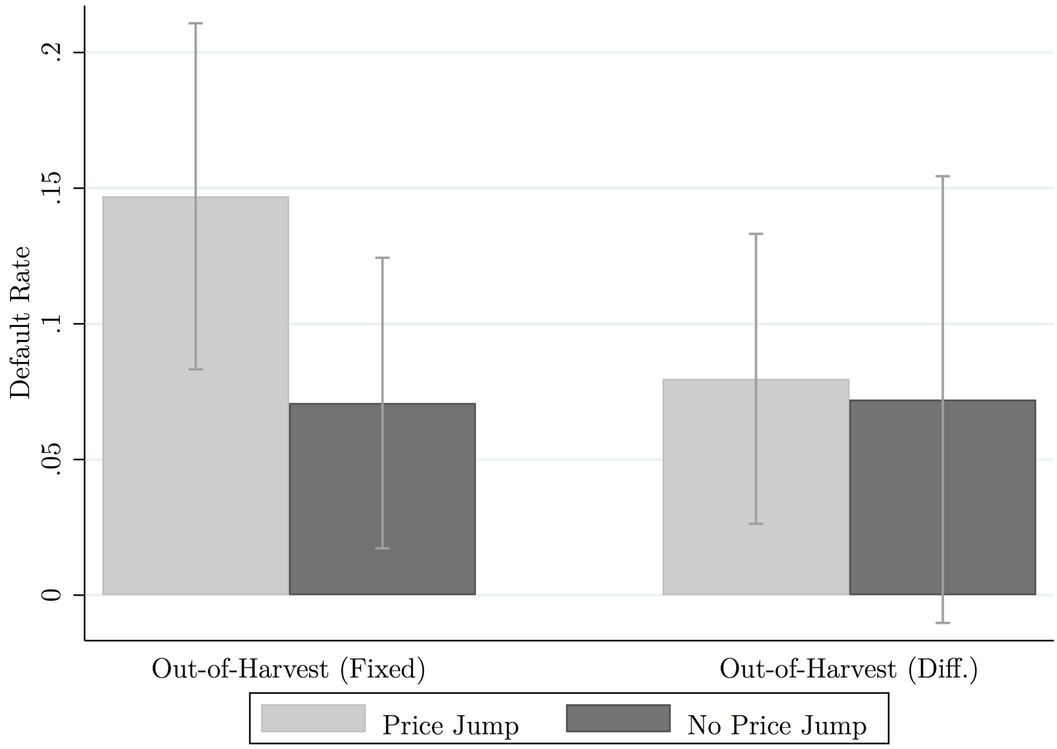

Figure VII: Event Study: Default before/after price jumps in/out of harvest season

Notes: The figure examines defaults on fixed price contracts before and after a large price jump, that occur in-harvest and out-of-harvest. A potential challenge in interpreting Figure V as evidence of strategic default is that unexpected increases in world coffee prices could be passed-through to farmers, thus raising input costs for mills and forcing them to default. To rule out this possibility, we take advantage of the stark separation in the timing of production and contract execution documented in Figure II. We exploit this timing and isolate strategic default by conducting an event study that considers only price increases that occur after the end of harvest, once price pass-through is no longer relevant and a price increase can only improve the mill's profits. The event study compares contracts that are due for delivery after harvest just before and just after a sudden price increase. The Figure uses the baseline definition of contractual default. Further analysis in Section III and in Appendix C shows that results are robust using alternative definitions of default. 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%. Further analysis in Appendix C shows a similar pattern when using different thresholds to define a 'price jump'. The figure shows that after an unexpected price jump, the defaults among the fixed price contracts rise for out-season price increases only.

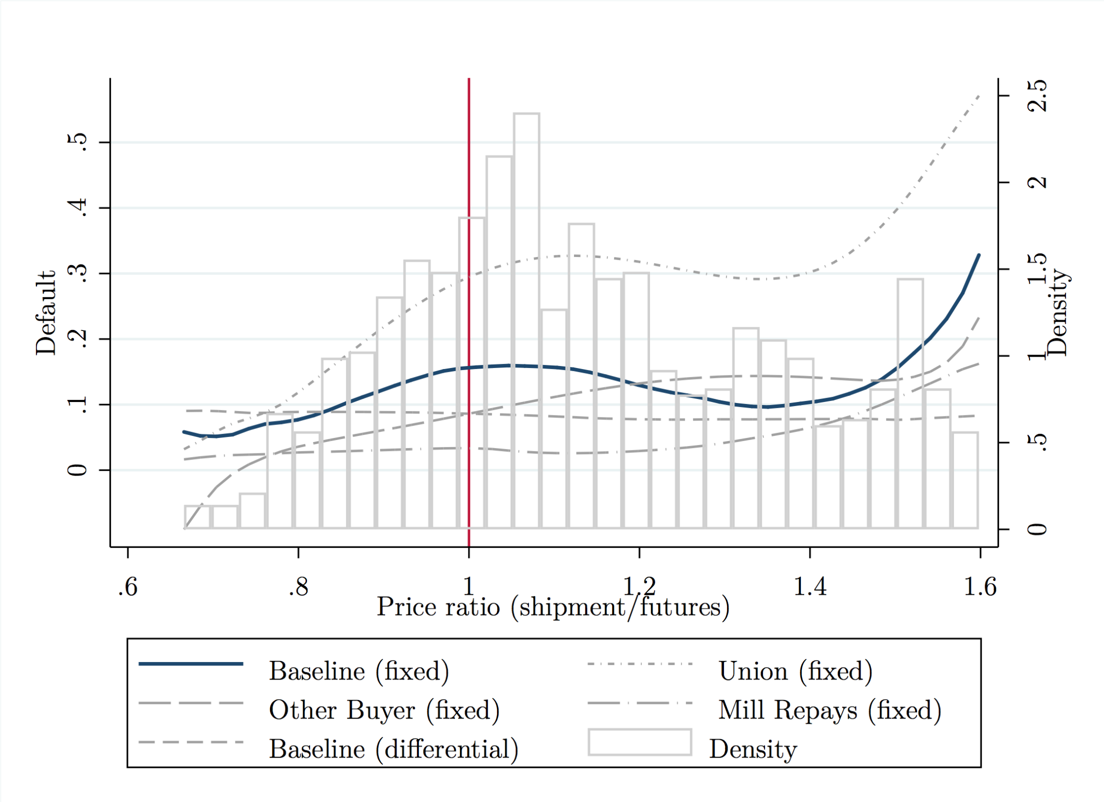

Figure VIII: Unexpected Price Increases and Contract Default (alternate definitions of default)

Notes: The Figure shows that the patterns in Figure V are robust to using alternative definitions of default. 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 II. In the raw data there are a few outliers that distort the scale of the graph (see Figure V). We therefore remove the top end of the distribution of price surprises. For a given price surprise, the solid line plots defaults (baseline definition: outright default and cases in which the loan is not yet fully repaid 90 days past-due) on fixed price contracts while the dashed line plots the same for differential price contracts. The Figure also includes, in addition to our baseline definition (solid), the default definitions that rely on direct repayment by the mill (dash-dot); other buyers (long-dash); and the union of the baseline and two alternate repayment measures (short-dash).

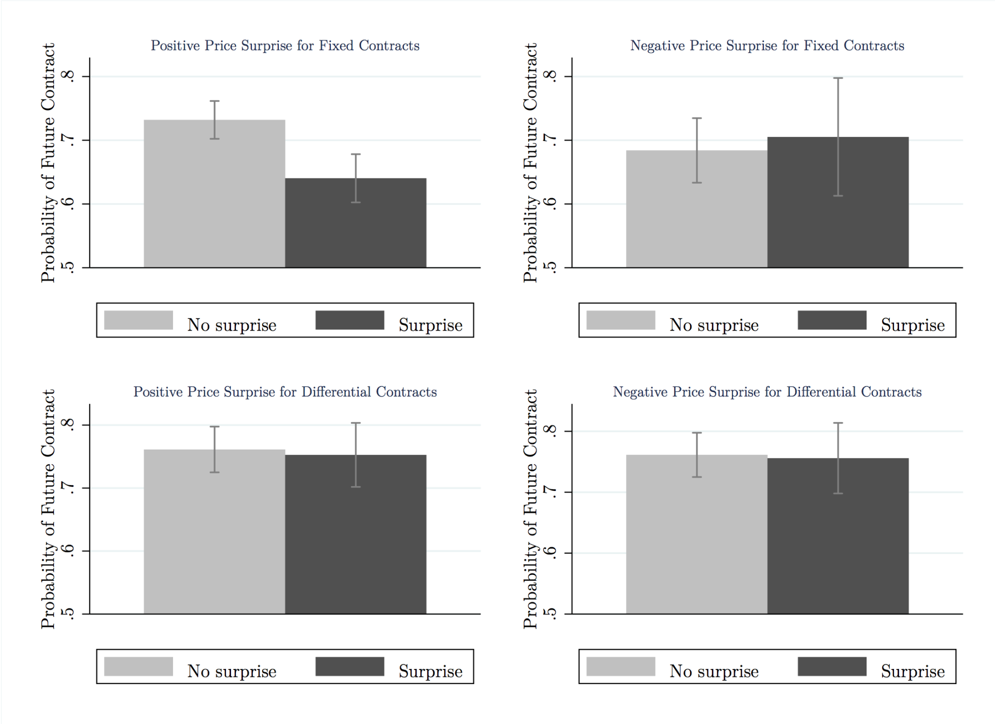

Figure IX: Relationship Termination following a Default: Positive vs Negative Price Surprises

Note: The Figure explores the likelihood that a defaulting mill receives a new loan in the season following a default. The probability of getting a new loan from the lender is lower following a default or a late repayment (see Table IV, Column 1). A possible interpretation is that the lender is less likely to supply loans following contractual non-performance. The mill should be punished more harshly if the late repayment is due to strategic default. The Figure shows that, conditional on a default, the mill is indeed less likely to receive a loan following a default that happened at the time of a positive price surprise as opposed to defaults happening at times of negative or no price surprise. Moreover, the effect exists only for fixed price contracts. A positive price surprise is defined as being in the top 25% of the price change distribution and a negative price surprise as anything in the bottom 25% of the price change distribution. All bars are conditioned on a default having taken place.

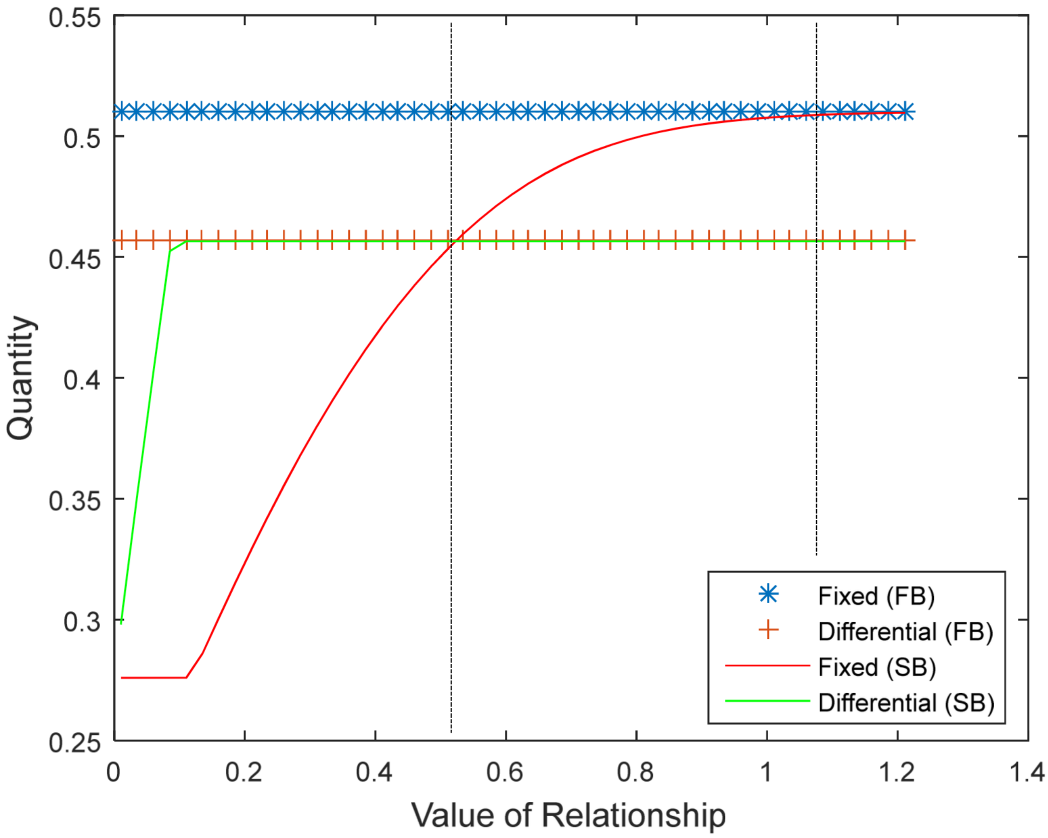

Figure X: Insurance vs. Enforcement Trade-Off

Note: The Figure illustrates how the Value of the Relationship **V** alters the solution of the model. The model is numerically solved assuming the functional forms and parameters described in Section IV. The x-axis reports the Value of the Relationship **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 **V**. 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 pledge-able income, the amount the mill can borrow and, consequently, the quantity produced. A higher relationship value **V** reduces the likelihood of strategic default and allows the mill to borrow more. Eventually, for very high values of **V** the solution approaches the first best. For lower values of **V**, however, the mill is better off foregoing price insurance and signing a contract on differential. This mitigates the strategic default motive and increases pledge-able income relative to a fixed price contracts. The model thus predicts that fixed price contracts are more likely to be observed in relationships with high **V**. Evidence in Appendix E confirms this hypothesis.

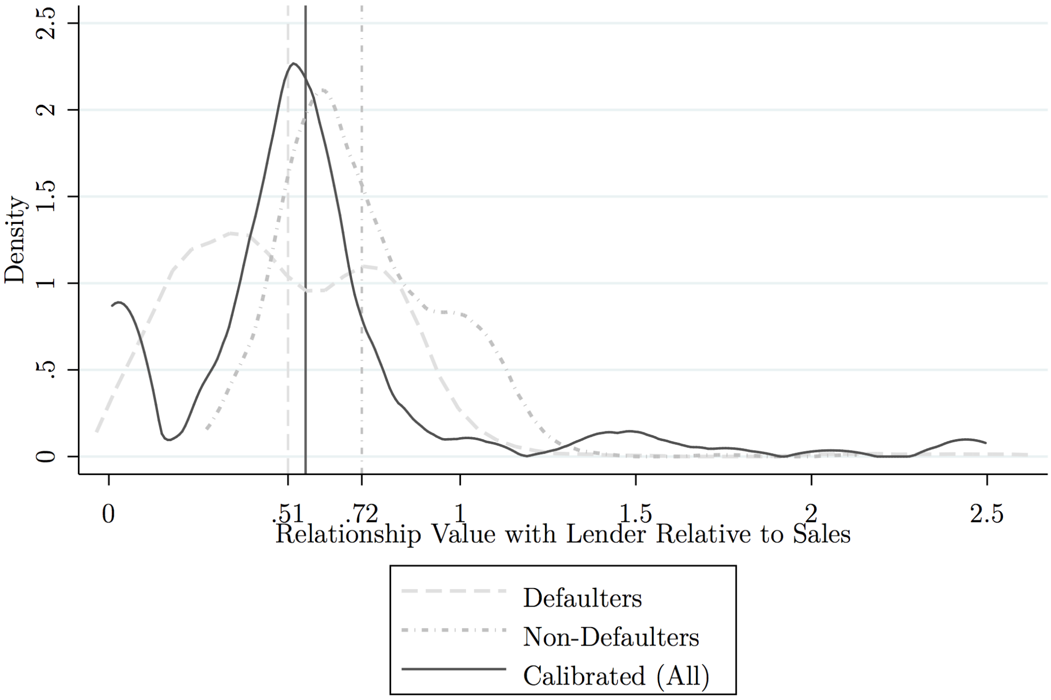

Figure XI: Upper Bound, Lower Bound and Calibrated Relationship Values

Notes: The Figure compares estimated relationship values **V** with those obtained from a revealed preference approach. Using the incentive compatibility constraint in Appendix A it is possible to derive upper and lower bounds to the value of the relationship for defaulters and non-defaulters respectively. The calibrated relationship values are within the bounds obtained from the revealed preference approach. It is worth noting that the calibration exercise and the revealed preference approach rely on completely different sources of identification. The revealed preference approach relies on observed defaults. The calibration does not. Actual defaults are not used to calibrate the model. The calibration recovers relationship values **V** from interest rates (which reflect, *inter alia*, the likelihood of defaults) rather than from actual defaults. The continuous variation in interest rates allows us to recover non-parametrically the distribution of relationship values **V**. Even a parametric approach wouldn't perform well if we were to recover relationship values **V** from the relatively few observed defaults. As a sanity check, however, the Figure compares the estimated relationship values **V** with bounds inferred from the actual observed decision to default or not. Despite using completely different sources of variation, the estimated relationship value **V** distributions display a significant overlap.

TABLES

Table I: Descriptive Statistics

Variable	Observations	Mean	SD	Min	Max
	(1)	(2)	(3)	(4)	(5)
		Panel A: Loans			
Default (Baseline Definition)	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: Mills			
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%
Number of Full Time Employees	48	10.4	7.4	1	32
Number of Seasonal Employees	43	17.5	35.5	2	196
		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%

Notes:

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 Appendix B for more details. Within mills for whom we can match to financial statements, observations vary due to reporting inconsistencies.

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

Dependent Variable:	**Default (Baseline Definition)**							
Data:	**Contract Level**						**Loan Level**	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Price Surprise	0.304**	0.343**	0.305**	0.279**	0.369**	0.0360	-0.0661	-0.0253
	(0.121)	(0.154)	(0.139)	(0.122)	(0.171)	(0.0679)	(0.0767)	(0.0875)
Fixed							-0.253**	-0.288**
							(0.111)	(0.125)
Fixed x Price Surprise							0.196**	0.201*
							(0.0907)	(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	−−	Yes	Yes	Yes
Month Fixed Effects	Yes	−−	Yes	Yes	−−	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
Mean of the dependent variable	0.127	0.127	0.127	0.127	0.127	0.080	0.139	0.139
N	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:

The Table reports results for the baseline test for strategic default. Main contract level results use a sample of 434 loans backed by fixed price contracts, which come from 180 mills, while the loan level analysis uses 967 loans backed by either fixed or differential contracts, from 272 mills. 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. We also test for the equality of coefficients between columns (1) and (6) and are able to reject equality, with a p-value = 0.027. 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 loan level. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

Table III: 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***	0.118***	0.105***	-0.00479	0.0438
	(0.0132)	(0.00352)	(0.0387)	(0.0584)	(0.0856)
Control Group Mean of Dependent Variable	0.055	0.005	0.074	0.065	0.091
N	123	70	154	150	72
R^2	0.026	0.044	0.015	0.000	0.002

Notes: The Table reports results for the event study test for strategic default. 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%. We also test for the equality of coefficients between columns (1) and (4) and are able to reject equality, with a p-value = 0.0134. 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 IV: Relationship Termination Following a Default

Dependent Variable:	Future loan	Default	Future loan	Future loan
	(1)	(2)	(3)	(4)
Price Surprise		0.328***		-0.153*
		(0.118)		(0.0920)
More than 90 Days Late	-0.0753*		-0.532*	-0.0657
	(0.0392)		(0.299)	(0.0399)
Year Fixed Effects	Yes	Yes	Yes	Yes
Maturity Month Fixed Effects	Yes	Yes	Yes	Yes
Buyer Fixed Effects	Yes	Yes	Yes	Yes
Mill Fixed Effects	Yes	Yes	Yes	Yes
Observations	434	434	434	434
R-squared	0.776	0.617	0.691	0.778

Notes:

The Table explores the likelihood that a mill receives a new loan from the lender following a default. The unit of observation is a contract. We focus on fixed price contracts only (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. All regressions are estimated using a linear probability model. Default is measured according to the baseline definition. Column 1 shows a negative correlation between a default and the likelihood the mill receives a loan in the future. A possible interpretation is that the lender is less likely to supply a loan following contractual non-performance. The correlation could, of course, also be driven by the mill's demand. Columns 2 and 3 explore an IV strategy in which default is instrumented with price surprises. Column 2 shows a decent first stage (despite the few defaults) and Column 3 shows that the IV estimate is significantly larger than the OLS estimate in Column 1. The larger IV estimate is consistent with the mill being punished following a (strategic) default. Column 4 reverts to an OLS specification. Consistent with Figure IX, we find that positive price surprises are associated with lower likelihood of getting a new loan. Furthermore, the OLS estimate for default is lowered and is now statistically insignificant at conventional levels. This is consistent with the hypothesis that defaults that occur at times of positive price surprises (and thus, more likely to be strategic) are more likely to lead to relationship termination. A bounding exercise along the lines of Oster 2017 reveals that the magnitude of the change in coefficients between Columns1 and 4 is consistent with strategic default being punished harshly. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

Table V: 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 IV. 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 $\overline{p}_w > 1$) and a fixed price contract with price 1. The farmers' supply elasticity η is estimated from the RDD estimates in Appendix D. 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 **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 VI: 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	$\alpha = 0.286$	$\alpha = 0.336$	α **= 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	η = 0.50	η = 0.55	η **= 0.60**	η = 0.65	η = 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 (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.

STRATEGIC DEFAULT IN THE INTERNATIONAL COFFEE MARKET

ONLINE APPENDIX

Arthur Blouin Rocco Macchiavello

Table of Contents

A1. Theory Appendix

A1.A. Remarks on the Model

Proposition *Under perfect contract enforcement the mill offers a fixed price contract with price* $p_c^* = \overline{p}_w$*and produces quantity* $q_F^* = (\overline{p}_w/(1+\eta)\gamma)^{1/\eta}$. *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 $\mathbf{E}[\Pi] = \int_{p_w} u(\overline{p}_w q_c - C(q_c))dF(p_w)$. Taking first order condition, $\overline{p}_w = C'(q_c^*)$ 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.[56] The buyer's participation constraint is now $\Delta_c = 0$. Defining $\widetilde{p} = \frac{D}{q}$and setting up the Lagrangian we obtain

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

The two first order conditions are given by

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

Substituing for ζ into the first condition we obtain

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

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

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

always negative.[57] The expression above is smaller than

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

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

A1.B. 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 IV). 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 IV, 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

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

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

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

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

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 $\mathbf{V}$.

A1.C. 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 D. 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 $\frac{\partial lnC(q)}{\partial L} = (1+\eta) \times \frac{\partial lnq}{\partial L}$. The estimate is $\frac{\partial lnC(q)}{\partial L} = 0.34$ (Table A10, Column 2). The unit price paid to farmers is $\omega = \rho q^{\eta}$. Taking logs, the RDD estimates gives $\frac{\partial ln\omega}{\partial L} = \eta \times \frac{\partial lnq}{\partial L}$. The estimate is $\frac{\partial ln\omega}{\partial L} = 0.13$ (Table A15, 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 IV 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 $\widehat{r}_i(\mathbf{V}_i)$ and a probability that a fixed contract is chosen, $\widehat{\phi}_i(\mathbf{V}_i)$. Specifically, denoting with $\mathbf{EU}^C(V_i)$ the predicted expected utility under contract type $C \in \{F, D\}$ we let $\widehat{\phi}_i(\mathbf{V}_i) = \frac{\mathbf{EU}^F(V_i)}{\sum_{C\in\{F,D\}} \mathbf{EU}^C(V_i)}$. For each loan i the value $\mathbf{V}_i$ is estimated as

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

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

A2. Data Sources and Additional Descriptive

A2.A. 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 boilerplate, 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.

A2.B. 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. 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, however, it does occasionally occur. When the mill does this 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 C).

A2.C. 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, 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.

A2.D. World Price Data

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.

A2.E. Internet and Field Data Collection

To assess whether firms who defaulted on the lender went out of business following the loan, we conducted a search for all firms in the sample that had defaulted. There were two main elements of this. First, we went back to the lender and asked whether any of the firms on the list had emailed, or made unsuccessful applications following the default at anytime. This happened in November 2017, so we have up to date contact of any form with the lender as of that date (which note: is outside the study range).

We also have data on management of most firms, as part of the application process. For the subset of firms that defaulted we systematically pulled this data from the applications. We conducted internet searches of firm websites, news articles and social media for last known activity of firms themselves, and if we could not find firm-level activity we searched for the management team to try and figure out if they are still active in the coffee industry.

For each defaulting firm, between these sources of information, we coded a last-known activity date. This captures either the last known activity of the firm, or the last known coffee-related activity of the management team. We ignored non-coffee-related activity of the management team. When possible, we also collect a location of activity for the management team with the idea that if the same management group is active in the same village under a different name, then if we cannot find evidence of firm activity, it may not be that the firm went bankrupt, but instead rebranded.

A2.F. Interviews of Industry Practitioners

We conducted 90 minute interviews with the key actors in the industry. This included an interview of a member of the senior management team of the lender, interviews with the lender loan officers who are the point of contact with the mills, and also with two different senior managers for buyers. All interviews were conducted in English, except for one with a loan officer was conducted in Spanish.

In each we asked a similar set of questions, from the perspective of the practitioner. The interviews were in a journalistic style of a mix between pre-set questions and on the fly follow-ups when something interesting was said.

A2.G. Bloomberg Options Data

We collected options data in order to include implied volatility of coffee as a control in the main regressions. The data comes from Bloomberg and the Intercontinental Exchange (ICE). Trade transaction data for both Futures and Options were downloaded from the ICE platform, following the link http://data.theice.com/, last accessed October, 2017. This options trading data was only available after 2008.

Reference daily implied volatilities were downloaded from a Bloomberg Terminal under a University of Toronto institutional subscription, in October, 2017.[58]

Risk-free rates are proxied with 3-Month Government T-bills from FRED, as the expiration interval for Coffee "C" Futures is roughly every three months.[59]

A2.H. Cash Flow Data

We manually collected cash flow statement data for every firm for which it was available. These files are not a standard part of the application process, but are a standard part of the audit process. Sometimes the information is volunteered as part of the appli-

[58] Download executed with the following code: =BDH($E$1,$F$2:$G$2,"27/10/1999","27/11/2017", "Dir=V","Dts=S","Sort=A","Quote=C","QtTyp=Y","Days=T","Per=cd","DtFmt=D", "UseDPDF=Y","CshAdjNormal=N","CshAdjAbnormal=N","CapChg=N","cols=3;rows=4384")

[59] Link: https://fred.stlouisfed.org/series/TB3MS, last accessed October, 2017

cation process, but not in a consistent manner. Similarly, the format of the information collected during the audit process is not in a standard format, is not always in English, and is often found in multiple different files. We therefore had a very difficult time automating the process using text analysis routines, so we decided to collect the data manually.

We hired five MA economics students at the University of Toronto, who divided up the firms and began the process of systematically organizing the data. The process began in December 2017, and the slowest RA had finished their share of loans on March 27th, 2018. We had randomly double assigned some cases to two RAs to check quality and consistency, since RAs had to in some cases make a judgement call on data quality (e.g. this often involved checking the same numbers against several financial statements, and pulling the ones with the most standardized reporting practices), and had to make considerations for differences in reporting practices between countries. In many cases RAs did not agree on the correct numbers.

The RA team met as a group to discuss conflicts. Enough conflicts existed that we decided to review the entire file again. This time one RA re-checked every financial statement we have on file for every firm. She either confirmed the numbers were correct in the case of RA agreement or in the case of not having been originally double checked, or played arbitrator in the case of disagreement. This ‘second round’ of data collection was finished August 23rd, 2018. The final dataset includes all cash-inflows and cash-outflows by month for all firms that we are confident used reporting practices that are consistent and standard.

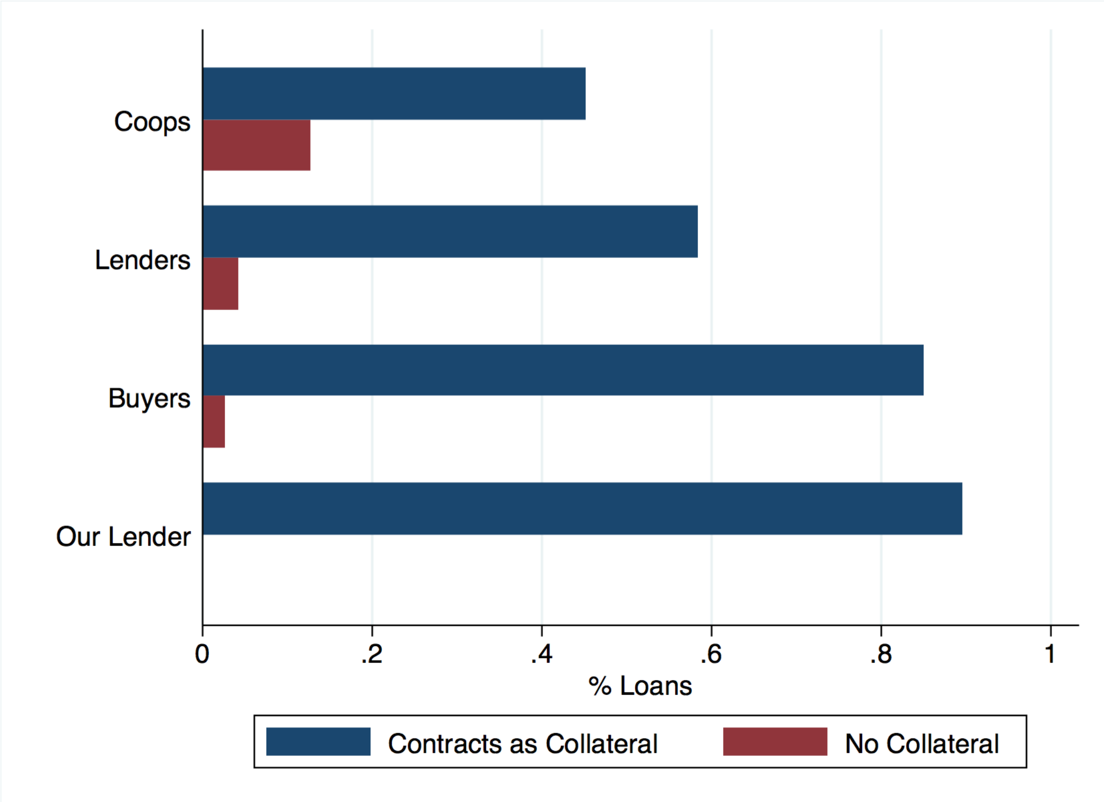

Figure A1: Use of Collateral

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.

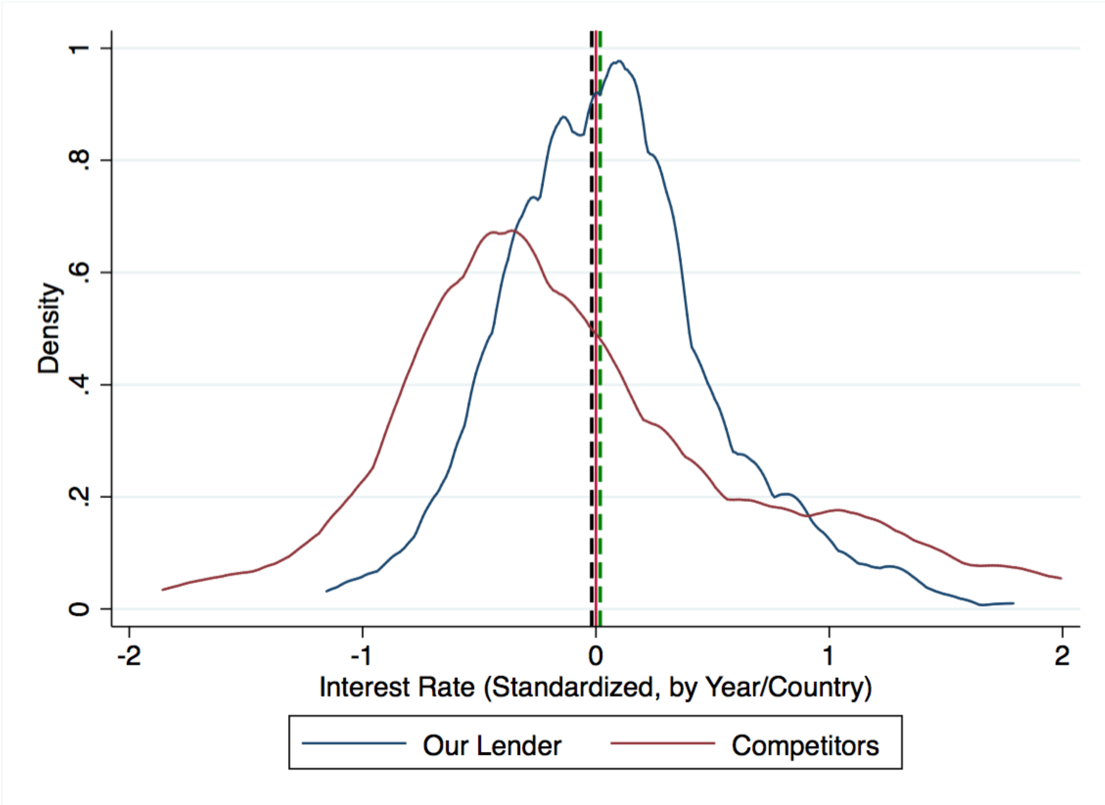

Figure A2: Representativeness of Lender's Interest Rates

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.

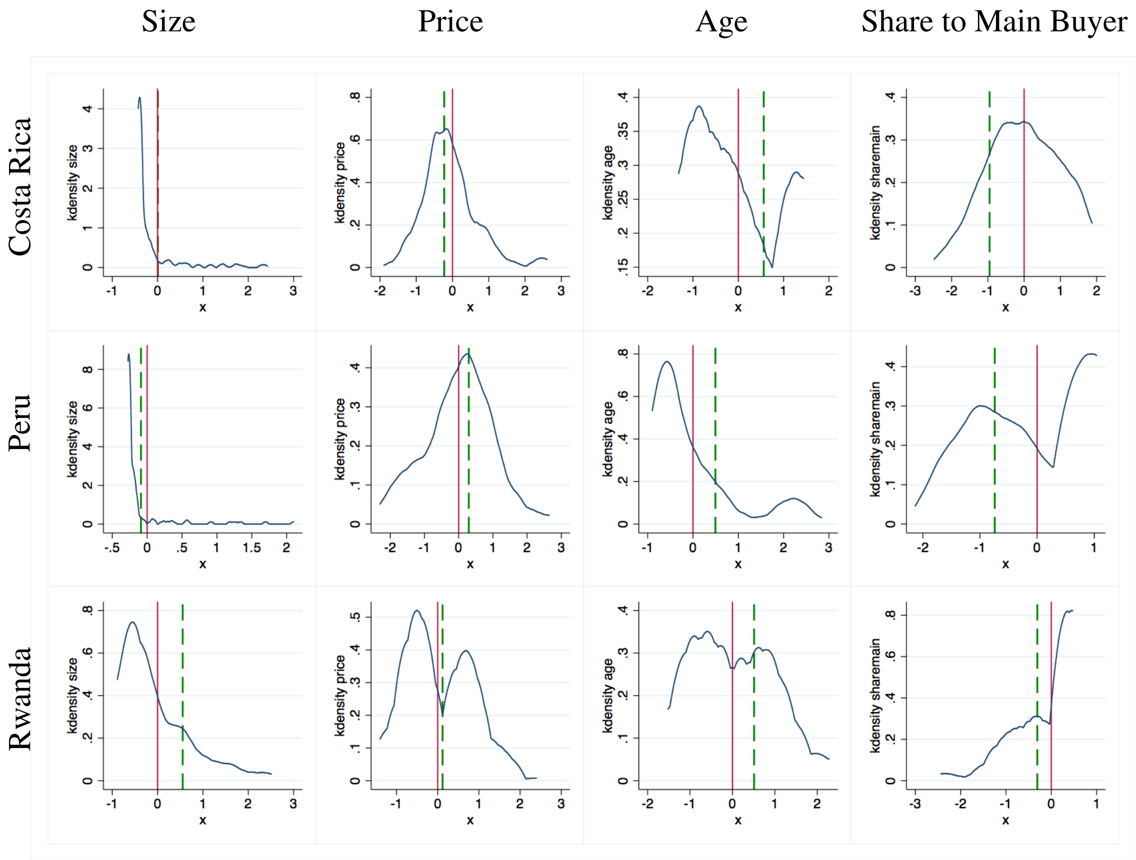

Figure A3: Representativeness of Lender's Portfolio

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.

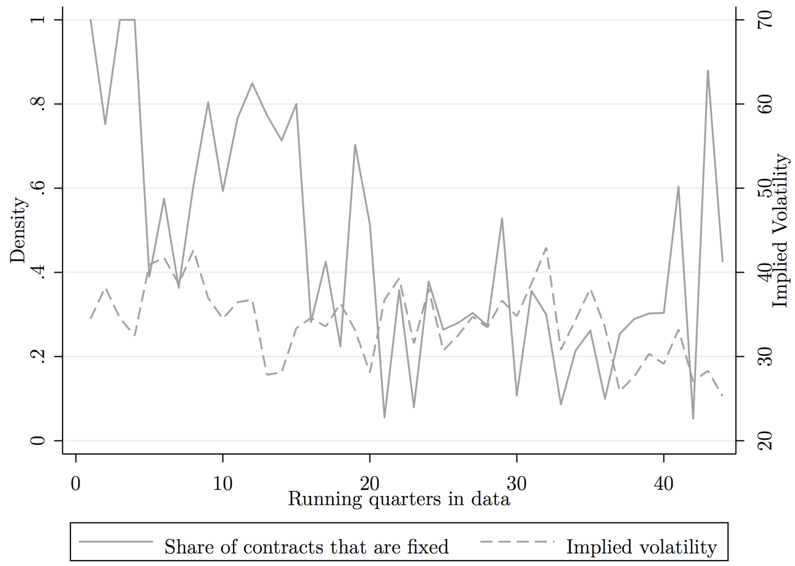

Figure A4: Evolution of Contract Types and Coffee Price Volatility

Note: The graph plots the density of price surprises by contract types.

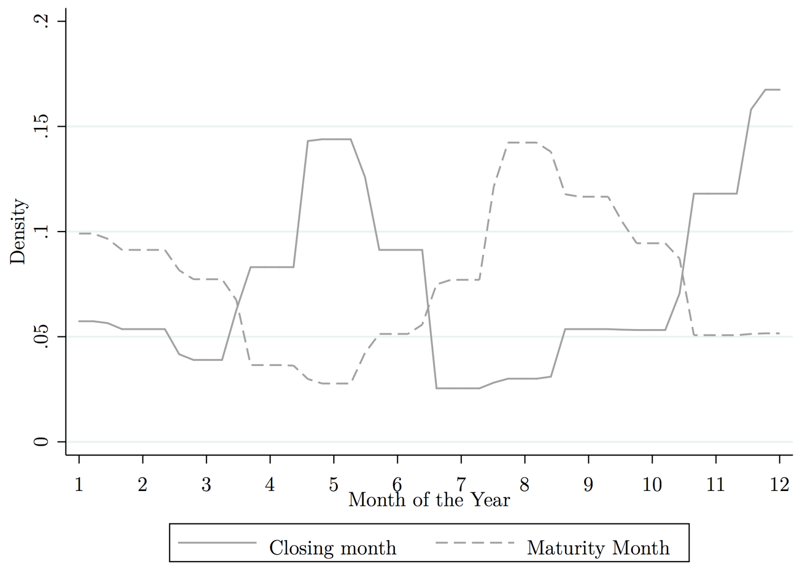

Figure A5: Asynchronous Harvest Timing Across Countries

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 hemispheres.

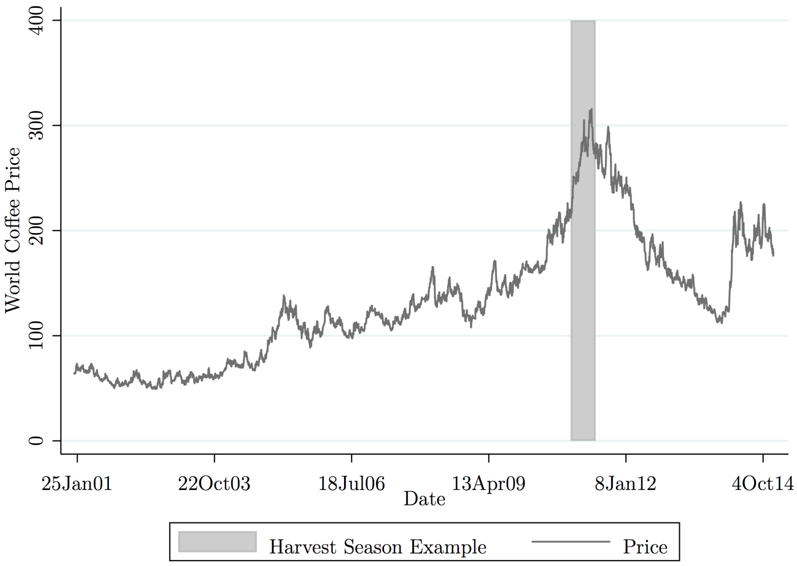

Figure A6: Time Series Graph of World Coffee Prices

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%).

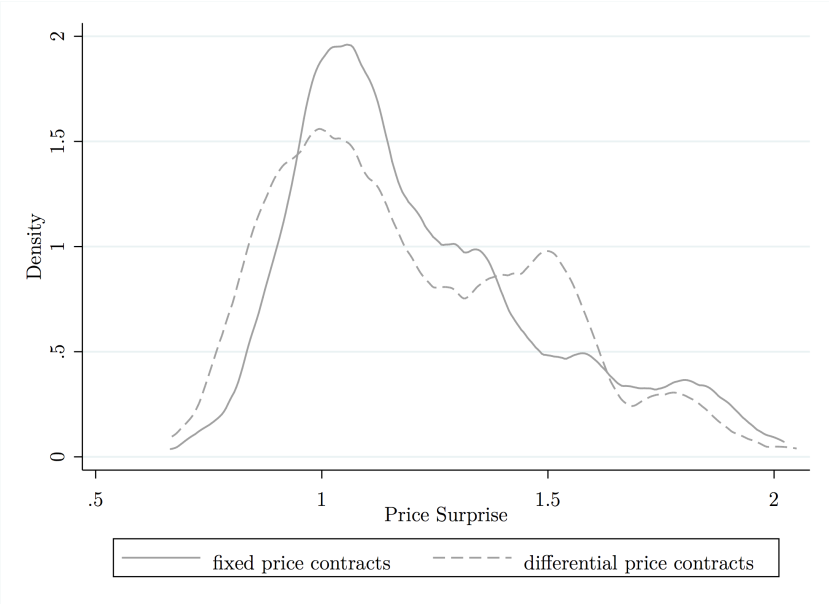

Figure A7: Price Surprises by Contract Types

Note: The graph plots the density of price surprises by contract types. We ran the associated Kolmogorov-Smirnov test on the distribution of price surprises (conditional on time/seasonality variables) and report that the difference is 0.07. Based on this test, we cannot reject the equality of the distributions at conventional levels of significance.

Table A1: Summary Statistics by Country

	% Lender Portfolio	Year of Entry	N. of Mills	N. of Loans	% Fixed	% Default
	(1)	(2)	(3)	(4)	(5)	(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.

A3. Robustness to Section *III.*

A3.A. Robustness to Section III.B

Table A2: Robustness of Table II: 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: The Table reports results of Column 1 in Table II using alternative thresholds to define late repayment. The baseline definition of default includes outright loan defaults and loans that are not been fully repaid after three months. Results are robust to a wide range of definitions. Note that a too lenient definition may classify as defaults instances that are due to regular delays. A too strict threshold may leave too few cases of default. The Table reports results considering thresholds of 2, 3 (baseline) and 4 months. Regressions are at the contract level. 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%.

Table A3: Robustness of Table II: Definitions of 'Fixed-Price Loan'

Threshold for fixed price loan (% loan fixed)	**45% fixed**	**50% fixed**	**55% fixed**	**60% fixed**
	(1)	(2)	(3)	(4)
Fixed Price Loan x Price Surprise	0.182**	0.196**	0.168*	0.173*
	(0.0907)	(0.0907)	(0.0913)	(0.0922)
Fixed Price Loan	-0.233**	-0.253**	-0.217*	-0.221**
	(0.111)	(0.111)	(0.111)	(0.112)
Price Surprise	-0.0618	-0.0661	-0.0517	-0.0534
	(0.0766)	(0.0767)	(0.0766)	(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
N	967	967	967	967
R^2	0.386	0.387	0.386	0.386

Note: The Table reports results for Column 7 in Table II using alternative thresholds to define a fixed-price loan. At the contract level, every 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 II, 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 replicate the specification showing a range from 45% to 60% and find 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 Appendix E). In all cases our dependent variable is the baseline definition of default. 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%.

Table A4: Robustness of Table II: Alternative Clustering Strategies

Dependent Variable:	Default or 90+ days late in repayment					
	(1)	(2)	(3)	(4)	(5)	(6)
Price Surprise	0.304***	0.304***	0.304***	0.304**	0.304**	0.304***
	(0.114)	(0.108)	(0.0900)	(0.121)	(0.133)	(0.0892)
Controls	baseline	baseline	baseline	baseline	baseline	baseline
Standard Errors	year-month	country-month	year	loan	client	two-way: client, year-month
Observations	434	434	434	434	434	434
R-squared	0.495	0.495	0.495	0.495	0.495	0.495

Note: The Table reports results of Column 1 in Table II using alternative clustering strategies. The Table considers clustering along the time dimension (i.e., at the year-month level), at the mill level, and two-way clustering at both mill level and country-time level. Note that the limited number of countries (around 20) suggests to avoid clustering at the country level. The unreported specification with standard error clustered at the country level yield, for the coefficient of interest, a standard error of 0.1137 with associated p-value 0.015. Standard errors and, more broadly, the interpretation of our results are unaffected by these alternative strategies.

Table A5: Robustness of Table II: Alternative Controls

Dependent Variable	Writeoff, Reneg. or 90 days late				
	(1)	(2)	(3)	(4)	(5)
Price Surprise	0.369***	0.325***	0.419***	0.258***	0.208**
	(0.0972)	(0.0967)	(0.111)	(0.0998)	(0.101)
Month-Year FE	Yes	Yes	Yes	No	No
Time to Maturity	No	Yes	Yes	Yes	Yes
Today's Price and Today's Expected Price at Maturity (Future price)	No	No	Yes	Yes	Yes
Expected Volatility	No	No	No	Bloomberg	Constructed
Mill FE	Yes	Yes	Yes	Yes	Yes
Buyer FE	Yes	Yes	Yes	Yes	Yes
Observations	434	434	434	422	230
R-squared	0.664	0.671	0.800	0.528	0.734

Notes: The Table reports results of Column 1 in Table II using alternative sets of controls. Table II reports a number of different specifications that include varying sets of controls. Columns 1, 2, and 3 in this Table reports results from the same specification adding progressively *all* controls at once. Results are virtually unchanged. Given there are relatively few defaults in the data, the main caveat against the specification in Column 3 is that it risks overfitting the data, i.e., identifying the effect off unreasonably few observations. Columns 4 and 5 explore an option view of the decision to strategically default. The possibility of defaulting has an option value for the mill. In the case of the fixed price contract, the default option looks like a deeply out-of-the-money call (as default must be unlikely, otherwise the lender and the buyer would never agree to the contract in the first place). Specifications in Column 4 and 5 directly control for predetermined objects that correlate with the option value: today's price; time to maturity and expected volatility. (The strike price, i.e., the minimum price at which it is worth to default, is unobservable). We use two measures of expected volatility: an off-the-shelf measure of implied volatility (Column 4) and a "first-principles" construction (Column 5). The first-principles approach gives the proper way to construct volatility but the necessary data are not available in Bloomberg going all the way back in our sample. The off-the-shelf measure takes the average volatility on a particular date over all possible contract lengths rather than the implied volatility for the actual contract start/end dates. The two measures turn out to be very highly correlated.

The exact measure constructs expected volatility for the specific start and end dates of each contract. The implied volatility σ is calculated by equalizing actual market price with the theoretical price for the relevant option using the standard Black-Scholes model. We consider the case in which the mill can exercise the "defaulting" option at the loan maturity date, i.e., a European option. The price of the associated call option is given by $C = SN(d_1) - Ke^{-r(T-t)}N(d_2)$ where $d_1 = \frac{ln\frac{S}{K}+(r+\frac{\sigma^2}{2})(T-t)}{\sigma\sqrt{T-t}}$ and $d_2 = d_1 - \sigma\sqrt{T-t}$ with S the future's price, K the strike price, T is the option's expiration time (i.e., the loan maturity), t is the contracting date, r is the risk-free interest rate, and σ the market's volatility expectations. The implied volatility is computed by finding the σ such that the difference between the calculated price and actual market price is minimized. On any given date thousands of options are traded with different strike prices, expiration dates and "moneyness". As commonly done, we compute the implied σ as the solution with the highest moneyness and then average between the implied volatility calculated from the call option and the put option to recover the implied volatility for that future on that trading date (see, e.g., Hull, J. *Options, Futures, and Other Derivatives*. 8th Edition, 2012: P224-226, 309-315.) In both cases, results are virtually unchanged.

Table A6: Importance of Extreme Events

Dependent Variable:	Default (Baseline Definition)			
% of Upper tail of Price Surprise Dropped	1%	5%	10%	25%
	(1)	(2)	(3)	(4)
Price Surprise	0.324**	0.336**	0.399**	0.183
	(0.129)	(0.138)	(0.181)	(0.229)
Year Fixed Effects	Yes	Yes	Yes	Yes
Maturity Month Fixed Effects	Yes	Yes	Yes	Yes
Buyer Fixed Effects	Yes	Yes	Yes	Yes
Mill Fixed Effects	Yes	Yes	Yes	Yes
N	430	420	395	332
R^2	0.495	0.523	0.524	0.561

Notes: The four columns in the Table report results removing observations in the top 1%, 5%, 10% and 25% of observed price surprises from the sample. Specifications are like in Column 1, Table II. Only when removing 25% of the observations in the sample that the linear specification delivers a positive, but not statistically significant at conventional levels, coefficient. The Table thus confirms that results in Table II are not driven by outliers.

A3.B. Robustness to Section III.C

Table A7: Background to Table III: Price Pass-Through

	Fixed Contracts		Differential Contracts	
Dependent Variable:	log(Price to Farmers)	log(1+Profit)	log(Price to Farmers)	log(1+Profit)
	(1)	(2)	(3)	(4)
Price Surprise	0.594	0.0294	0.536	0.203***
	(0.824)	(0.0918)	(0.550)	(0.0597)
Year Fixed Effects	Yes	Yes	Yes	Yes
Maturity Month Fixed Effects	Yes	Yes	Yes	Yes
Buyer Fixed Effects	Yes	Yes	Yes	Yes
Mill Fixed Effects	Yes	Yes	Yes	Yes
Price at Closing	Yes	Yes	Yes	Yes
Futures Price at maturity quoted at closing	Yes	Yes	Yes	Yes
Observations	148	220	217	303
R-squared	0.888	0.951	0.598	0.929

Note: One concern with the results in Table II could be that on the fixed price contracts the price surprise may increase the input prices due to pass-through of the price shock along the supply chain, but the mill does not see profits, because sales prices were already agreed upon. For differential contracts on the other hand, a price increase is likely to increase both their input prices and their sales prices. This could lead to fixed price contracts differentially defaulting - not strategically - but due to debt overhang. Indeed, while we do not see decreased profits for mills on fixed price contracts following a positive price surprise, and we actually see increased profits for differential price contracts, we are unable to rule out substantial pass-through of the world prices to farmers, as we observe fairly large (though imprecisely estimates) effects of price increases on the mill's input costs for those both on fixed and differential price contracts.

Table A8: Placebo to Table III: No Price Pass-Through of Out-Season Price Jumps

Dependent Variable:	**Prices paid to farmers by mills**				
	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	-17.62	-23.08	-19.65	-8.188	-5.767
	(10.69)	(14.30)	(13.12)	(13.05)	(11.96)
Observations	53	26	64	74	27
R-squared	0.053	0.094	0.080	0.013	0.022

Note: Note: One concern with the results in Table III could be that on the fixed price contracts the price surprise may increase the input prices due to pass-through of the price shock along the supply chain, but the mill does not see profits, because sales prices were already agreed upon. This could lead to fixed price contracts differentially defaulting - not strategically - but due to debt overhang. We therefore run an event study out-of-harvest season to ensure this pass-through is not driving results, since essentially all purchases have already been made by this time. Indeed, the table confirms that there is no price pass-through for any contract-type when the price increases after harvest.

Table A9: Robustness to Various Definitions of Default

Dependent Variable:	**Subsample: Fixed Price Contracts**				**Subsample: Differential Price Contracts**			
	Very late outright default	'Off-contract' buyer repays	Mill repays directly	Union (1)-(3)	Very late outright default	'Off-contract' buyer repays	Mill repays directly	Union (4)-(7)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	PANEL A: Price Surprise At Any Time							
Price Surprise	0.343**	0.192†	0.0455†	0.452***	0.00173	0.0693	0.0309	0.169
	(0.133)	(0.128)	(0.0311)	(0.166)	(0.0663)	(0.0816)	(0.0372)	(0.137)
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	Yes	Yes	Yes	Yes
Country-Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	434	434	434	434	533	533	533	533
R^2	0.634	0.637	0.864	0.621	0.590	0.498	0.520	0.522
	PANEL B: Event Study - Price Surprise Out of Harvest							
Price Surprise	0.143***	0.175***	0.0682***	0.375***	-0.0253	0.155	0.0117	0.132
	(0.0132)	(0.0516)	(0.0198)	(0.0837)	(0.0558)	(0.106)	(0.0783)	(0.102)
N	123	123	123	123	150	150	150	150
R^2	0.026	0.057	0.026	0.075	0.003	0.023	0.000	0.004

Notes: The Table explores the robustness of results in the baseline specification in Table II (Panel A) and in the event-study specification in Table III (Panel B) to alternative definitions of contractual default. Columns (1) and (5) consider loan backed by fixed and differential contracts respectively and define default to be the case in which any party (original buyer, a different buyer, or the mill) repays the loan late or the loan is defaulted against (baseline definition). Columns (2) and (6) define default to happen whenever a different buyer from the one originally on the contract repays the loan. Columns (3) and (7) define default to happen whenever the mill directly repays the loan. Columns (4) and (8) define a contractual default whenever any of the behaviours separately considered in the three previous sets of regressions is observed. Regardless of the definition of default used, positive price surprises are associated with higher likelihood of default on fixed price contracts but not on differential contracts. We also test for the equality of coefficients between columns (4) and (8) and are able to reject equality, with a p-value = 0.0606. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%, † denotes significance at 85%.

Table A10: Placebo to Table III: No Long Run Effects

	Fixed Prices		Differential Prices	
	Price to Farmer	Profit	Price to Farmer	Profit
	(1)	(2)	(3)	(4)
Last Year Price Surprise	50.10	19,363	56.52	5,867
	(111.4)	(30,343)	(41.68)	(16,424)
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
N	78	119	92	134
R^2	0.925	0.957	0.963	0.954

Note: One concern could be that a price surprise this year, even out-of-harvest, might reduce the profitability of the mill in the next year because of an increase in next year's input prices. In this case the mill might default this year due to expected future debt overhang. To test this we examine long-run effects on input prices and profits. As expected we do see an increase in prices paid to farmers (a price surprise last year is lasting), but if anything we see that profits are *higher* in the year following a price increase, ruling out the possibility that mills default because the net present value of future profits is no longer sufficient to cover long-term financial obligations.

Table A11: Robustness of Table III: 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: The Table reports specifications as those in Table III using alternative definitions of price-jump. The Table considers 4 definitions of an event at 0.5% intervals below the baseline estimate at 3%: 1% (column 1); 1.5% (column 2); 2% (column 3) and 2.5% (column 4). These alternative definitions expand the sample of contracts included in the sample from the 123 in the baseline 123 to 223. Results are consistent. Local linear regressions are executed at the contract level. In all cases our dependent variable is the baseline definition of default. 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%.

Table A12: Robustness of Table III: In-Season Price Controls

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

Notes: The Table examines the robustness of the event-study to controlling for in-season price increases. A concern with the event study is 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. While this concern is mitigated by the narrow window used in the event-study it is prudent and straightforward to check that our results are robust to this control as well. The Table runs 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 III, column 1) to 0.136 with the estimates using alternate event windows being similarly close. Local linear regressions are executed at the contract level. In all cases our dependent variable is the baseline definition of default. 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%.

A3.C. Additional Results for Section III.D

Table A13: What Happens to Defaulting Mills?

	Panel A: All defaults (N=147)	
	N	% of Sample
Firm definitely still exists	117	80%
Firm may still exist (unclear)	28	19%
Firm definitely does not exist	2	1%
	Panel B: Unclear cases (N=28)	
	Mean	Standard Deviation
last known activity year-default year | Firm may exist	4.4	2.9
1(last observed year $<$ 1+ default year | Firm may exist)	14%	0.35

Note: This Table reports the results from field and internet searches of all observed instances of default ($N = 147$) in the data. If positive price surprises simply induce mills to default or go bust, we would expect at least some of the defaulting mills to actually go bust (unless *every single* instance of observed default was marginal, in the sense that it kept the mill alive). In contrast, if many defaults are strategic, we expect mills to be still operating in the years following a default. We thus conduct original field and internet searches for all instances of observed defaults in the sample to check whether the defaulting mills are still active or, instead, have gone bust. For each defaulting mill, we try to establish if the mill is still in operation and, if not, when was the last time the mill was operating. Panel A shows that in 80% of the cases the defaulting mill is still operating in 2018; in 19% of the cases the mill is still possibly operating in 2018; in only 1% of cases the defaulting mill is no longer operating in 2018. Panel B takes a closer look at the 19% of cases ($N = 28$) that are not clear. Among those, the average mill was still operating 4.4 years after the default and in only 14% of the cases we couldn't verify mills operation in any year after the default. Examples of the research assistant search results for the unclear cases are: "Default probably caused a change in management but the ______(sic) mill is still viable and it is under the umbrella of ______(sic)""; "I could only find a website that shows a seller contract with ______(sic) from 2009" [note that this is the default year]. We also followed up with the lender, and they received an (unsuccessful) application for a loan in the following year (2010)"; "Based on the identity of the management group and location of the mill, I suspect that the mill rebranded after the default, and if so they were active on social media several years after defaulting."

A3.D. Details of the Calculation Suggesting that Default is Punished Harshly

The change in magnitude in the estimated coefficients between column 1 and 4 might appear small at first glance. A simple exercise in the spirit of Oster 2017, however, reveals that the magnitude of the change is consistent with strategic default being punished harshly. To see why, consider a short regression $F_{it} = \beta_0 + \mathring{\beta} Default_{it} + X'_{it}\Gamma + \varepsilon_{it}$ and a long regression $F_{it} = \beta_0 + \tilde{\beta} Default_{it} + \beta_1 PriceSurprise_{it} + X'_{it}\Gamma + \varepsilon_{it}$ in which F_{it} is a dummy taking value one if mill i receives a loan after year t. Let β^S and β be the coefficient on relationship termination following a strategic default (that we can identify) or another type of default, respectively. Using an Oster 2017 inspired technique[60] to solve for β and plugging in the estimates from Table IV we get $\beta = 0.933$. The bias we are interested in is then $\beta^S = 0.5(-0.0657 - 0.933) \approx -0.5$. This method gives us almost the same answer as the IV technique, and similarly implies that strategic default results in future loans with much less regularity than other types of default. While the estimates have to be interpreted very cautiously,[61] they do appear to be indicate a significant punishment following strategic defaults.[62]

[60]Our set-up is different from the one considered in Oster 2017 in the sense that our $\tilde{\beta}$ parameter is a superset of the 'confounding' β^s. To adjust for this, we consider a δ of 0.5 rather than her baseline of 1. In our case a $\delta = 1$ would assume that the correct benchmark was that only strategic defaults mattered. Our choice of $\delta = 0.5$ is quite conservative since it seems unlikely that strategic defaults matter *less* than other types of default with respect to punishment, and increasing δ will lead to the result that punishment of strategic default was even more prevalent. In other words, $\delta = 0.5$ stacks the deck against our claim as much as reasonably possible, yet still allows us to say that strategic default is punished very harshly relative to other types of default.

[61]Especially in light of ad hoc assumptions about δ.

[62]The evidence using alternative definitions of default in Table A9 is similar. When a mill is late on loan repayment, we are also less likely to observe the buyer-mill pair again in the data. This evidence is subject to the same caveat above as well as to the additional one that we do not observe transaction between the buyer and the mill when they do not deal with our lender.

A4. Credit Constraints

Table A14: 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	237,093	0.0001	-0.0003	48,045	-194,920
	(47,078)**	(113,322)**	(0.001)	(0.0003)	(71,967)	(154,741)
RD Estimate: 75% Bandwidth	132,705	299,342	-0.0002	-0.0005	105,241	-360,359
	(39,877)***	(87,801)***	(0.0009)	(0.0004)	(67,215)	(172,908)**
RD Estimate: 125% Bandwidth	118,612	211,233	0.0005	-0.0004	-5,793	-114,031
	(38,253)***	(122,506)*	(0.0012)	(0.0003)	(83,036)	(166,904)
Observations	523	523	523	523	523	523

Notes:

Tables A14 and A15 test for credit constraints using a methodology related to Banerjee and Duflo 2014. We exploit rounding up in the lender's credit score system to implement a RDD design. A mill is credit constrained if the marginal product of capital is larger than the interest rate paid on the marginal dollar borrowed $MPK > r$. The test for credit constraints is then 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. The test is valid if three conditions are satisfied: *a*) the additional loan should not change the (marginal) interest rate; *b*) the larger loan affects operations at the margin (i.e., how much to produce) and not whether the mill operates at all or not; *c*) the mill must be able to use the loan to pay down other loans. Condition *a*) can be tested directly. Our lender will 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. Conditions *b*) and *c*) are satisfied since we look at working capital laons signed at the beginning of the harvest season.

The lender determines the size of the working capital loan based on a letter score which is calculated rounding an underlying continuous numerical score which weights a large number of mill and loan characteristics (see details below). Mills are classified as having a C, B, A or AA credit score. We conduct the analysis on the B-A and A-AA thresholds separately (very few mills receive a C). 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 A14 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. The estimates are consistent with what reported by the lender. Furthermore, in neither case was there an associated decline in the interest rate (columns 3-4). However while we find no evidence of substitution away from other loans for the B-A threshold (column 5), we find basically one to one substitution out of other loans for the A-AA threshold (column 6). This suggests that the B-A mills could be credit constrained while the A-AA mills are not.

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 2012. We report the number of observations used prior to the bandwidth computation. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

Table A15: 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**	0.338**	0.452**	0.134***
	(56,161)	(0.148)	(0.222)	(0.0290)
RD Estimate: 75% Bandwidth	122,259**	0.283**	0.403*	0.138***
	(49,252)	(0.125)	(0.227)	(0.0421)
RD Estimate: 150% Bandwidth	152,646**	0.376**	0.385**	0.189***
	(64,615)	(0.152)	(0.187)	(0.0581)
Observations	212	212	206	166

Notes: The Table examines how the additional money is utilized focusing on the B-A threshold. 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 find that coffee purchases expand by about $130,000 (column 1), an amount similar to the additional loan received (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% (column 3), reflecting an MPK of about 12%. The average interest rate is about 8%, which suggests that a fair share of mills are credit constrained. Furthermore, we find that prices paid to farmers increases by about 13% (column 4). This estimate is used to calibrate the elasticity of farmers' supply η. We find no evidence (unreported) that larger loans increase the likelihood of default, which suggests that the incentive constraint to prevent loan diversion is not binding.

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 2012. We report the number of observations used prior to the bandwidth computation. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

Additional results available upon request ensure the validity of the RDD. We perform the McCrary density test and find no evidence for manipulation. If the lender nudges mills up and down the threshold at similar rates, the McCrary test would be satisfied but the RDD would be invalidated. We therefore check if any of the 64 sub-scores is significantly higher/lower around the threshold. Of the 64 sub-scores we find only 5 to be significantly different on either side of the threshold, well within what we would expect based on random chance. This lends support to the validity of the RDD.

A5. TESTING ADDITIONAL PREDICTIONS

Table A16: Heterogeneity: Relationship History and Strategic Default

Dependent Variable:	**Default or 90+ days late on repayment**					
Relationship Type:	**Lender - Mill**			**Buyer - Mill**		
(Proxy):	**(Lending History)**			**(Sales History)**		
Sample:	**Fixed price &**		**Differential**	**Fixed price &**		**Differential**
	Low Value	**High Value**	**Low Value**	**Low Value**	**High Value**	**Low Value**
	(1)	(2)	(3)	(4)	(5)	(6)
Price Surprise	0.258*	0.165	0.0489	1.553*	0.171	0.145
	(0.139)	(0.336)	(0.0986)	(0.834)	(0.142)	(0.256)
Mill Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Buyer Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
N	200	234	262	221	213	237
R^2	0.572	0.643	0.498	0.846	0.693	0.864

Notes: The Table explores the effect of price surprises splitting the sample according to the value of the relationship. We use the past history in a relationship to proxy for relationship value. Table A18 below confirms that estimated relationship values **V** correlate with relationship history between the mill and the lender and between the mill and the buyer. Specifications are as in Table II. Columns 1-3 examine heterogeneity by the strength of the relationship between the lender and the mill. We find that mills on fixed price contracts with relatively worse relationships with the lender are much more likely to default. We find no such effects on mills with low relationship values on differential contracts (Column 3). In Columns 4-6 examine buyer-mill relationship. This relationship is measured with more noise since transactions between the mill and the buyer at times not covered by the audited accounts used by the lender are not observed in the data. We find again that mills on fixed price contracts with low relationship values are much more likely to default. We find no such effects on mills with low relationship values on the differential price contracts (column 6). Unreported results show that mills on fixed price contracts are much more likely to default against buyers that are not central in the lender's network. Standard errors are clustered at the loan level. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

Table A17: Contract Selection: Selection of fixed versus differential contracts

Dependent Variable:	**Fixed price contract**						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Relationship History (buyer)	0.399***	0.426***	0.400***	0.386***	0.426***		
	(0.0915)	(0.0839)	(0.0927)	(0.0921)	(0.0839)		
Relationship History (lender)						0.160***	
						(0.000441)	
Lender's Perception of Mill-Buyer relationship							0.209**
(from credit score)							(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
N	967	967	967	967	967	967	331
R^2	0.258	0.291	0.260	0.261	0.297	0.283	0.353

Notes: The model predicts that fixed price contracts tend to be signed in more valuable relationships. The Table reports results from the specification:

$$F_{cmbt} = \alpha_0 + \alpha_1 R_{mbt} + \alpha_2 X_{cmbt} + \lambda_m + \gamma_b + \mu_t + \varepsilon_{cmbt} \tag{8}$$

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. We use the past history in a relationship to proxy for relationship value. Table A18 below confirms that estimated relationship values **V** correlate with relationship history between the mill and the lender and between the mill and the buyer. The results in the Table confirm the prediction of the model. Unreported specifications show that these results are robust to the estimation of a model in which contract length and contract type are jointly estimated allowing for correlation between the residuals of the two models. Additional estimates show that the shipments on a loan, the number of days between closing and shipment, and contract size do not affect the estimates. We find that more valuable relationships sign both larger and fixed price contracts. Additional placebos show that price surprise do not correlate with contract type. Standard errors are clustered at the mill level. *** denotes significance at 99%; ** denotes significance at 95%; * denotes significance at 90%.

Table A18: Correlates of Estimated Relationship Values **V**

Dependent Variable	Calibrated Value of Relationship			
	(1)	(2)	(3)	(4)
Relationship History with Buyer (% cumulative dollars)	1.353**			
	(0.678)			
Relationship History with Lender (% cumulative dollars)		1.427*		
		(0.731)		
Debt Contract Enforcement (z-score)			-0.645*	
			(0.379)	
Commercial Contract Enforcement (z-score)				0.317
				(0.332)
Buyer Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes
Country Fixed Effects	Yes	Yes	No	No
Observations	412	412	349	412
R-squared	0.174	0.169	0.214	0.165

Notes:

The Table projects estimated relationship values **V** on observables characteristics. The reported correlations are suggestive but cannot be interpreted as establishing evidence in favour of certain (causal) determinants of relationship value. There are potential three sets of relevant observable characteristics: *i*) relationship-level variables; *ii*) market-level variables; and *iii*) country-level institutional variables. Columns 1 and 2 show that the estimated relationship values are positively correlated with measures of the past amount of business between the mill and the buyer and between the mill and the lender. These results nicely match those in Appendix Tables A16 and A17 documenting that measures of relationship histories also correlate with responses to incentives to strategically default and with contract choices as implied by the model. Unreported results show no strong correlation pattern between the number (or concentration) of alternative trading partners in the country and estimated relationship values. While theoretical predictions are ambiguous, we note that our data come from one lender only and thus we are able to construct only very noisy proxies for alternative partners. The resulting measurement error makes it difficult for us to interpret the lack of correlation in the data and thus we do not report the results. Columns 3 and 4 explore country-level institutional variables. Debt contract enforcement is minus ”Time to resolve debt insolvency in the country” from Djankov et al. 2008. Note that the loss of observations in column (3) is due to the following countries that are in our dataset but not in the Djankov et al. 2008 dataset: Bolivia (N=11); Nicaragua (N=39); Rwanda (N=7); Tanzania (N=3); Uganda (N=3). Quality of *debt* contract enforcement negatively correlates with estimated relationship values. Our preferred explanation is that in countries in which it is harder to enforce debt contracts, lenders will be more reluctant to lend. A relationship with any given lender, including our lender, is then likely to be more valuable. Commercial contract enforcement is minus ”Time to enforce commercial contracts in country” from the World Bank Doing Business Database. Quality of *commercial* contract enforcement in the country does not correlate with estimated relationship values. The quality of contract enforcement in the mill’s host country has little to do with the ability of the buyer to enforce contracts on an international transaction. All buyers we interviewed confirmed that courts in the mill’s host country are irrelevant from the point of view of enforcing the forward sale contracts.