STRATEGIC DEFAULT IN THE INTERNATIONAL COFFEE MARKET*

Arthur Blouin and Rocco Macchiavello

This article 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 sizable negative externalities on supplying farmers. *JEL* Codes: D22, L14, G32, O16.

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

*We are grateful to the editors, Pol Antràs and Andrei Shleifer, and to four anonymous referees. We thank 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 shared the data used in this article for many comments and suggestions. We 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, and Warwick. Matteo Benetton, Niclas Moneke, Juan Morales, and Caroline Nicol provided outstanding research assistance. We gratefully acknowledge IGC and IPA financial support.

The Quarterly Journal of Economics (2019), 895–951. doi:10.1093/qje/qjz004.
Advance Access publication on January 22, 2019.

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., Antràs 2015; Djankov et al. 2003; Fafchamps 2003).

This article develops a test to empirically identify strategic default and implements it in the context of the international coffee market. We build on 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; Hart 2009). The test identifies strategic default by studying how contractual defaults respond to large unanticipated changes in market conditions. Of course, large changes in market conditions could increase both revenues and costs, thereby affecting the likelihood of default through multiple channels. 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

1. See, for example, 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 (Holmström 1979; Grossman and Hart 1983). Note 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 affected differently by changes in the environment and therefore require different remedies. For instance, strategic default might require finding alternative partners to trade and will be affected by the market structure in ways that effort underprovision is not. Finally, they also have different legal implications.

market, not only for production efficiency but also for contract design, insurance, and credit availability.

We conduct our analysis on prefinancing 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 who import coffee. The lender 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 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 and audited financial accounts for the borrowing mills.

We are interested in two questions: (i) is there evidence of strategic default?, and (ii) 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] There are two types of forward sales contracts. In fixed-price 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 would not face this temptation.

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

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

A key challenge to test for strategic default is 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 3 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 made 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 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 because 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 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 nonperformance on the loan part of the contract and is standard in the literature on loans. However, it 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 implement our baseline test and the event study using these alternative indicators of contractual default as well as the union of all possible indicators. In 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, mills are less likely to receive a loan from the lender in the future. 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 that the relationship will 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, that is, when it was more likely to be strategic. Furthermore, consistent 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

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

contractual arrangement between the coffee mill, the buyer, and the lender. The possibility of strategic default induces a trade-off between insurance and counterparty 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, the risk of strategic default is large. Parties then sign differential contracts that eliminate temptations to default at the cost of forgoing 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 be unconstrained, credit constrained, insurance constrained, or both credit and insurance constrained. Perhaps counterintuitively, 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 on 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 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.

I.A. *Related Literature*

Our main contribution to the literature is to isolate a specific form of moral hazard and 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 article 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 article also relates to an emerging literature on contracting in environments with weak or nonexistent enforcement institutions. Finally, although we provide a test for strategic default in a particular market, the main idea can be fruitfully applied to strategic default in other contexts.

By studying prefinancing agreements, the article contributes to a body of empirical work on contracts (see Chiappori and Salanié 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 difference-in-differences 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 consideration of different aspects, most notably the importance of interfirm business relationships.

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.

Within the literature on contracting under imperfect enforcement, Antràs and Foley (2015) offer a notable contribution. They show that trade finance terms balance the risk that 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 Kenyan flower exporters exerted efforts to protect valuable relationships with foreign buyers during a negative supply shock. Unlike their article, we observe and test for strategic default and focus on how it influences contractual terms and efficiency.[7]

The analysis in this article relates to, 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, prefinancing 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 article can be adapted to isolate strategic default in other contexts. For instance, classic studies

7. Two recent articles 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; 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., Manova 2012) 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.

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 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 associated inefficiencies.[11]

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, and the 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.[12]

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 (~~forthcoming~~). 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 apply our test for strategic default on a sample of Costa Rican coffee mills and find similar results.

12. After the cherry skin is removed with a machine, beans are sorted by immersion in water and 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, sometimes with the help of machines. After drying, the hulling process removes the parchment skin before export.

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 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 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, prespecified date. The mill obtains working capital finance either in the form of an advance from the buyer or, as is the case in this article, as a working capital loan extended by a lender on the back of the forward sale contracts.[13] 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 prespecified 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

13. 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.

TABLE I
DESCRIPTIVE STATISTICS

Variable	Observations (1)	Mean (2)	Std. dev. (3)	Min (4)	Max (5)
Panel A: Loans					
Default (baseline definition)	967	0.082	0.238	0	1
Loan amount (US$)	967	473,012	553,040	8,500	4,500,000
Interest rate	967	9.8%	0.10%	8%	18%
New borrower	967	46.7%	49.9%	0	1
Length of loan (days)	967	247.75	64.3	42	365
Number of buyers providing collateral	967	1.93	1.35	1	11
Share loans backed by both fixed and differential contracts	967	0.48	0.499	0	1
Panel B: Contracts					
Fixed-price contract	967	0.448	0.498	0	1
Price on shipping date/ futures price shipping date (when contract was signed)	967	1.19	0.278	0.666	2.049
Contract matures during harvest season	1,228	0.22	0.41	0	1
Futures price shipping date (when contract was signed)	1,228	151.23	42.67	64.3	301.99
Price when contract matures	1,228	167.94	55.12	52.81	309.94
Year contract signed	1,228	2009	2.9	2000	2014
Year contract matured	1,228	2010	2.9	2000	2014
Panel C: Mills					
Number of loans from lender	272	3.6	2.86	1	12
Assets (1,000 US$)	113	2,035	2,954	9.24	17,894
Sales (1,000 US$)	106	3,713	5,278	28.6	39,677
Purchases (1,000 US$)	102	759	719	12.65	3,247
Sales/cherry purchases	102	3.77	1.08	2.26	11.604
Profit (1,000 US$)	106	56.4	30.78	34.9	260.9
Price paid to farmers (US$)	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

TABLE I

CONTINUED

Variable	Observations (1)	Mean (2)	Std. dev. (3)	Min (4)	Max (5)
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 are presented at four levels: the loan, the contract, the mill, and the buyer. There can be several contracts backing a single loan, because mills sign contracts with different buyers and sign contracts of different types (fixed/differential). There are 1,228 observations of this type. Sometimes 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, for example, the scheduled shipping date could be missing, resulting in fewer observations. While most analysis in the article requires shipping information, we also 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. These data are again available for the later part of the sample. Furthermore, the financial data is backward-looking and can only be matched to a loan-year when the mill receives another loan within the next three years. See Online Appendix B for more details. Within mills that we can match to financial statements, observations vary due to reporting inconsistencies.

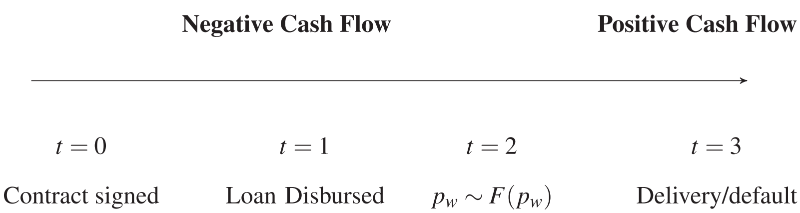

FIGURE I

Timing of Events

The figure illustrates the actual timing of events during a typical harvest season and 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 sell the cherries to another buyer at the prevailing spot price and strategically default ($t = 3$). Online 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, and then defaults only if it finds one. This alternative timing of events underpins our baseline definition of contractual nonperformance and the calibration in Section IV.

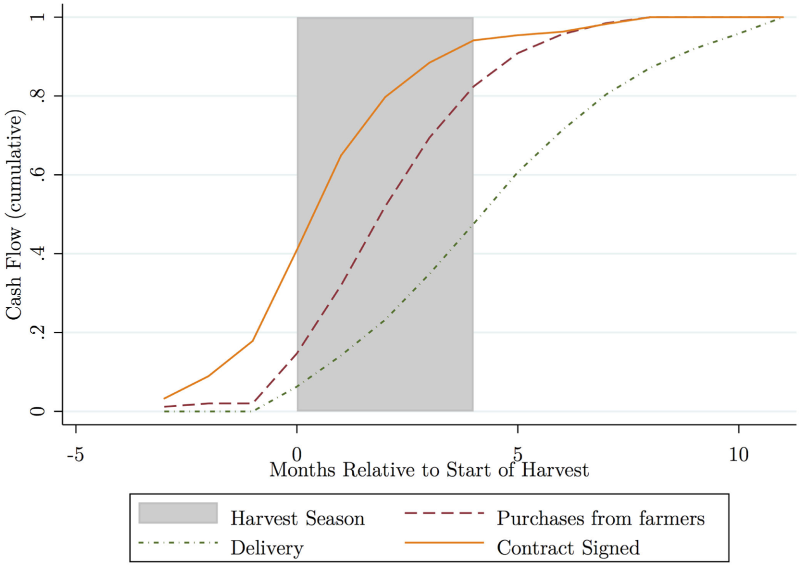

FIGURE II

Mills' Contracting and Cash Flow Profiles

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 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 timing of harvest is asynchronous across countries (see Online Appendix Figure A5). 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 article. 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 5 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-term 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 (postharvest) is crucial for our empirical strategy as detailed in Section III.

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 4 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 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 at which it decides whether to honor the contracts or default (postharvest) 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 12 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 (Online Appendix 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 directly repaid by the buyer once the mill delivers the coffee.[14] The lender's

14. 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

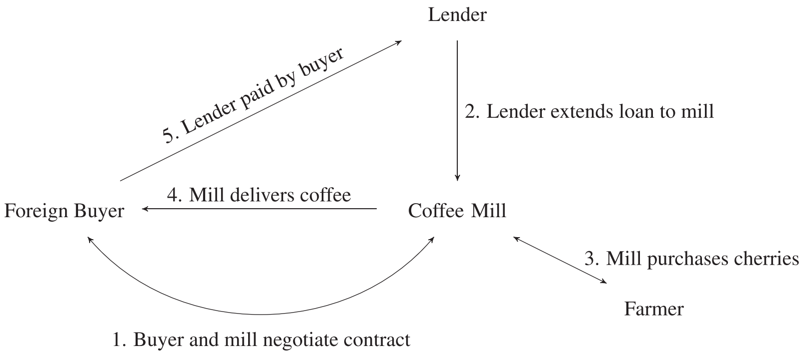

FIGURE III

Lending Model under Normal Circumstances

This figure shows the lending model 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 secures a loan from the lender. The loan amount is based on a formula that decides on a fraction of the value of the contract and 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-on quantity to the buyer. The buyer then repays loan to the lender directly.

contract terms and portfolio of clients appear to be broadly representative of the markets in which the lender operates (see Online Appendix Figures A1, A2, and A3).

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.[15]

the payment of principal plus interest. The interest rate is never contingent on coffee prices.

15. 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).

In fixed-price contracts the price is agreed on 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.[16] 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 prespecified differential. Theoretically, the basis price can be any published price in the industry. In practice, almost all differential contracts are signed against futures markets (i.e., Robusta coffee is traded against the London LIFFE Contract, while Arabica coffee, the object of this study, is traded against the New York ICE "C" Contract). Differential contracts remove counterparty risk but leave parties exposed to price fluctuations. A seller that has sold coffee on a differential basis will not be tempted to renege on the contract if prices suddenly increase, since the contracted price tracks spot market conditions.[17]

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 the lender's portfolio has remained fairly constant over time (see Online Appendix Figure A4). 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).[18]

16. 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.

17. 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.

18. Fair trade contracts specify a differential price above a fixed floor price and thus are classified as differential for the purpose of our analysis.

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 of 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 for 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, it then defaults on the sale contract. At this point, the mill decides whether to also default on the loan. If the mill decides to default, the loan is not repaid. This happens in 1% of cases. If 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 gray-shaded end nodes in Figure IV). The baseline definition has two advantages: (i) it is a directly observable measure of nonperformance 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

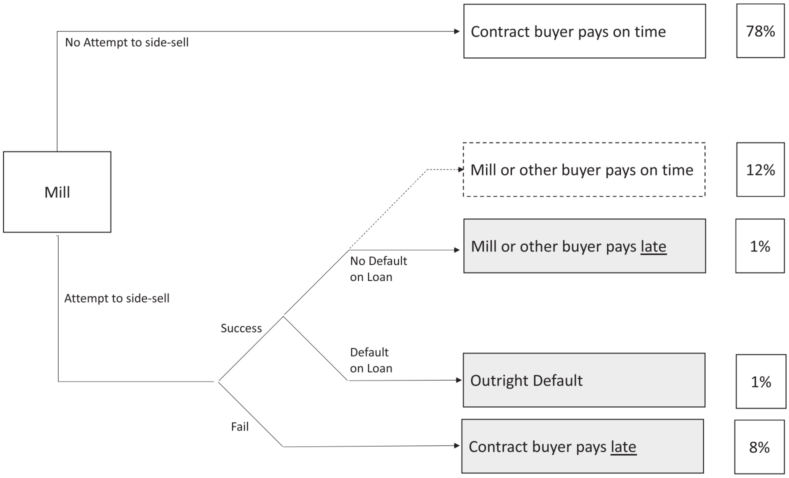

FIGURE IV

Contractual Nonperformance: Taxonomy and Baseline Definition

This figure illustrates and quantifies the different observable cases of contractual nonperformance. 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 Online 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 honor the sale contract and sell to the buyer or whether to search for 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, it 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 defaults on the sale contract. At this point, the mill decides whether to also default on the loan. 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 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 gray-shaded end nodes in the figure). The baseline definition has three advantages: (i) it rests on a directly observable measure of nonperformance on the loan part of the bundle; (ii) it is standard in the literature on loans; and (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, the baseline definition underestimates strategic default by omitting those cases in which the mill or an alternative buyer repays 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.

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.[19]

To gain a better understanding of contractual nonperformance, and to 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 pairing up with alternative buyers.[20]

Similarly, several buyers described instances in which, following increases in prices, suppliers attempted to renegotiate fixed-price contracts. Whether the buyer accepts the renegotiation is down to individual circumstances, including the extent to which the buyer would suffer from a canceled shipment. Attempted renegotiations are unambiguously perceived as a sign of an unreliable supplier.[21]

19. 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."

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

21. 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*

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 toward strategic default.[22]

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 counterparties for 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 the world price of coffee and other data. In addition, we conducted original field and internet searches to assess the extent of continued activity among all mills that 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. Online 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 data set 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 the 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 533 are differential. The remaining 258 contracts are mostly

---

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."

22. 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.

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.

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 prefinancing agreements to quantify the inefficiencies resulting from strategic default.

III.A. Conceptual Framework and Empirical Strategy

1. The Decision to Strategically Default. Consider a risk-neutral mill deciding whether to honor or default on a forward sales contract.[23] 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 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, that is, $p_c = p_w + \Delta_c$.

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:

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

Denote with $\mathbf{V} = \mathbf{U^R} - \mathbf{U^D}$ the value of the relationship.[24] The incentive compatibility constraint can be rewritten as

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

23. 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.

24. In our context, $\mathbf{V}$ bundles the relationships with both the buyer and the lender, as further discussed in Section IV.A.

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.

2. 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.

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 a measure of price surprise for each contract between mill m and buyer b signed at date t for delivery at date t':

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

3. 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 it is written off, restructured, or has no payments after 90 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

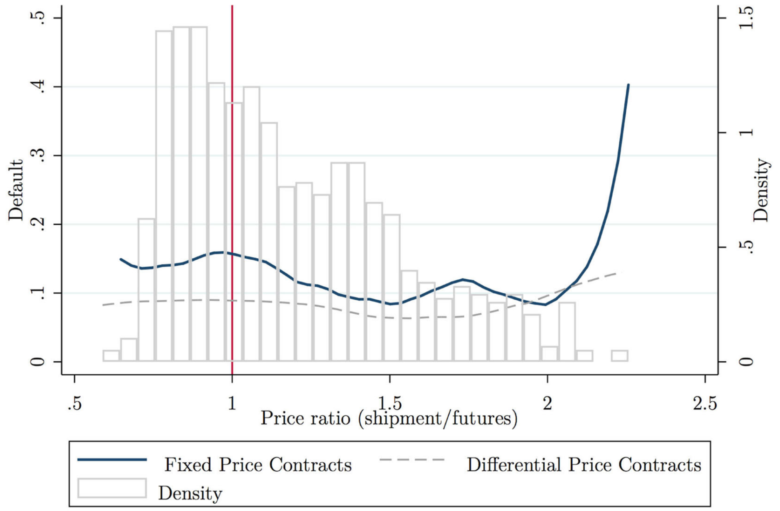

FIGURE V

Unexpected Price Increases and Contract Default

The figure illustrates the main evidence underpinning the test for strategic default. The gray 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 Online 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 default rates at price surprises in the interval 1.5 to 1.8 is sufficient to detect a differential effect across contract types.

default.[25] The evidence, however, could be driven by confounding factors. Online Appendix Figure A4 shows that the distribution

25. Endogeneity of price surprises to a mill's defaulting behavior isn't a concern. Mills are small relative to the global market. Furthermore, the asynchronous timing of harvest seasons across countries (Online Appendix Figure A5) implies that when, say, a Costa Rican mill is deciding whether to default, world coffee prices move because of development in other countries, for example, news of a frost just before the upcoming harvest season in Brazil. In any case, we report

of contract types has been relatively stable during the sample period despite significant swings in world coffee prices and volatility (Online Appendix Figure A6). Online Appendix Figure A7 shows that the distribution of price surprises across contract types are similar (and not statistically different from each other). 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

1. Baseline Specification and Results. We consider variations on the specification

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

where $D_{lmbt}^{t'}$ is a dummy equal to 1 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).[26] 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.[27] 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

below specifications that control for time fixed effects, which are identified by idiosyncratic variation induced by the timing and length of contracts signed.

26. 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.

27. Online Appendix 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.

TABLE II

STRATEGIC DEFAULT I: UNEXPECTED PRICE INCREASES AND DEFAULTS ON LOANS

Dependent variable	Default (baseline definition)							
	Contract level						Loan level	
Data	(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 × 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 that 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 because 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 = .027. Online Appendix E reports robustness checks on this table varying both the definition and the thresholds to assign loans to fixed contracts in columns (7) and (8). Standard errors are clustered at the loan level. ** denotes significance at 95%; * denotes significance at 90%.

(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 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).[28]

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)–(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 0. 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.[29]

2. 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 (Online Appendix 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, Online Appendix Table A5, column (3)

28. We discuss further robustness tests later.

29. 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.

explores a specification that controls for all independent sources of variation in Table II, columns (4) and (5) at 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. Online Appendix Table A5, columns (4) and (5) 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.

3. 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 nonlinear effects of price surprises. We define price surprises as a binary variable equal to 1 if the price surprise is above a certain threshold z and 0 otherwise. We reestimate 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. Online Appendix 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: only when removing 25% of the observations in the sample did the linear specification deliver a coefficient that, while positive, is not statistically significant at conventional levels. The table thus confirms that the results are not driven by a 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

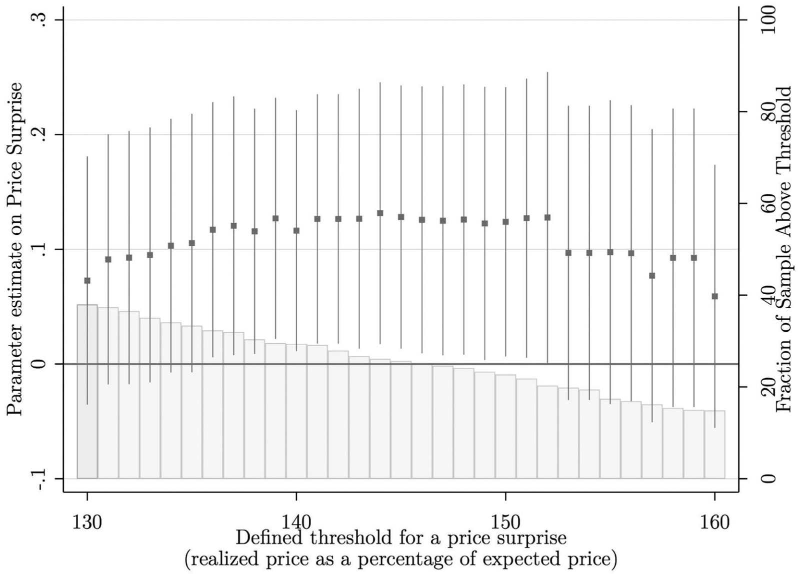

FIGURE VI

Baseline Test: Robustness to Nonlinearities and Outliers

The figure explores nonlinear 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 the value of 1 if the price surprise is above a certain threshold z and 0 otherwise. We reestimate 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 of 150%–180% of price surprises (see, again, Figure V). The figure also suggests that outliers are not driving results. We explore outliers directly in Online Appendix 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.

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.[30] To rule out this possibility, we take advantage of the

30. 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

stark separation between the timing of production and contract execution documented in Figure II. Using audited monthly cash flow data, the figure shows that (i) the vast majority (more than 90%) of payments to farmers occur during the harvest season, and (ii) the majority of forward sales contracts are executed after the end of the harvest season (but before the following harvest). That is, decisions to honor or default on the contract after harvest happen when 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 and just before and just after a sudden price increase.[31]

Table III implements the event study (see Figure VII for an illustration). We separate price increases into ones that happened in-season and out of 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

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, Online Appendix Table A7 shows that price surprises are not associated with lower seasonal profits for mills, suggesting that mills are not forced to default due to higher costs. Online Appendix 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, Online Appendix 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.

31. 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 countryside.

TABLE III

STRATEGIC DEFAULT II: UNEXPECTED OUT-OF-SEASON PRICE INCREASES AND DEFAULTS

Dependent variable	Default or 90+ days late on repayment				
In/out of harvest season	Out				In
	Fixed price			Differential price	Fixed price
	2 weeks	1 week	3 weeks	2 weeks	
Event window	(1)	(2)	(3)	(4)	(5)
Shipment scheduled after price jump	0.143*** (0.0132)	0.118*** (0.00352)	0.105*** (0.0387)	– 0.00479 (0.0584)	0.0438 (0.0856)
Control group mean of dependent variable	0.055	0.005	0.074	0.065	0.091
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 = .0134. Online Appendix E reports further robustness checks. Standard errors are clustered by event-day bins. *** denotes significance at 99%.

the analysis only at the contract level because of the small window around the more precisely relevant decision date.

Table III, columns (1)–(3) show the effect on default of out-of-harvest price increases, which can only be due to strategic default. The first column shows the difference using a two-week window and 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%.[32]

As expected, we see no analogous increase in defaults on differential-price contracts column (4). In column (5) we show

32. We also consider smaller increases of between 1% and 2.5% price increases. The resulting estimates are actually quite similar (Online Appendix 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 Online Appendix Table A12).

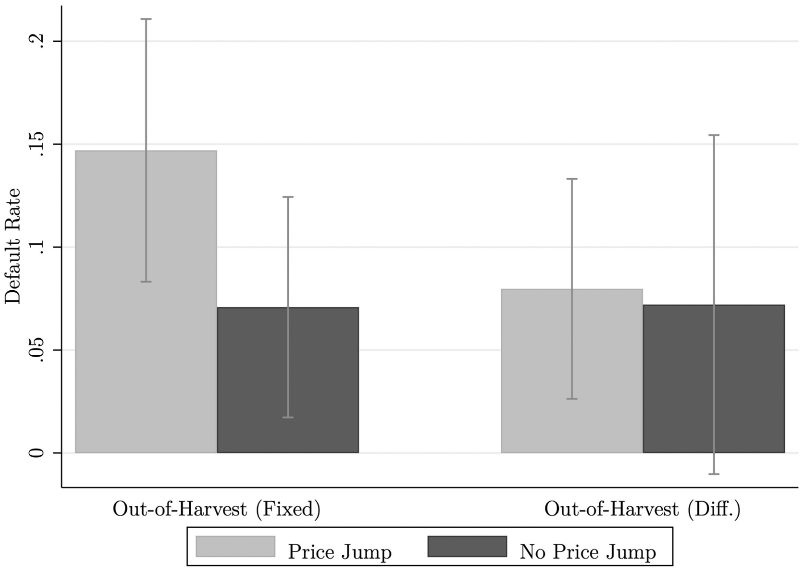

FIGURE VII

Event Study: Default before/after Price Jumps in/out of Harvest Season

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 Online 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 Online 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-of-season price increases only.

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.

III.D. Understanding Contractual Defaults

1. 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.[33] 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 behavior 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. Online Appendix 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

33. 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 nonperformance on the loan part of the contract bundle.

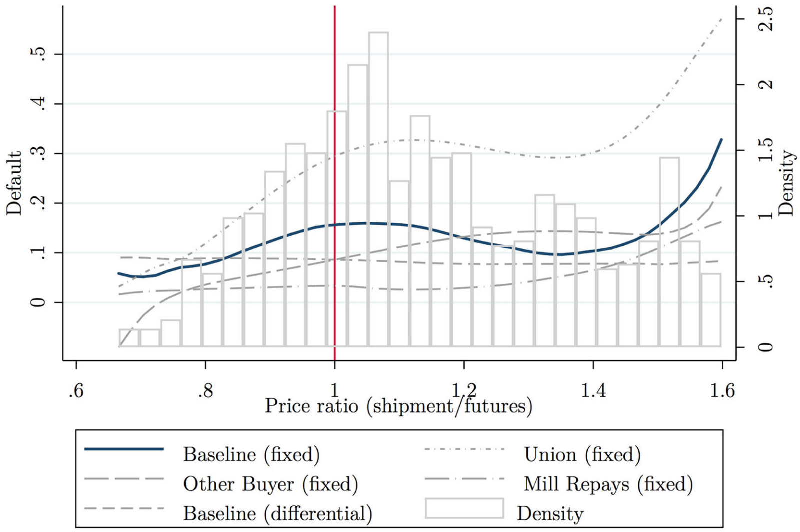

FIGURE VIII

Unexpected Price Increases and Contract Default (Alternate Definitions of Default)

The figure shows that the patterns in Figure V are robust to using alternative definitions of default. The gray 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).

of default on fixed contracts 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 behaviors 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.

2. Why Do Buyers on Fixed-Price Contracts Not Default When Prices Decrease? The logic of the test would suggest that unanticipated decreases 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.[34] 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 article. And so, there's that double incentive that when you buy article 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.[35]

III.E. *What Happens After a Default?*

1. Relationship Termination. Using 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 arbitration clauses are hardly ever enforced, it might 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).

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

35. 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.

TABLE IV

RELATIONSHIP TERMINATION FOLLOWING A DEFAULT

Dependent variable	Future loan (1)	Default (2)	Future loan (3)	Future loan (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^2	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 nonperformance. 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 a 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 (which are 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 columns (1) and (4) is consistent with strategic default being punished harshly. *** denotes significance at 99%; * denotes significance at 90%.

A possible interpretation is that the lender is less likely to supply loans following contractual nonperformance. This could be due to the buyer not offering contracts 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 surprises. Furthermore, this is only true for the fixed-price contracts, and not the differential-price contracts where we would not expect this pattern. This suggests the lender takes default after positive price surprises as a sign of strategic default rather than difficulties in repaying the loan.

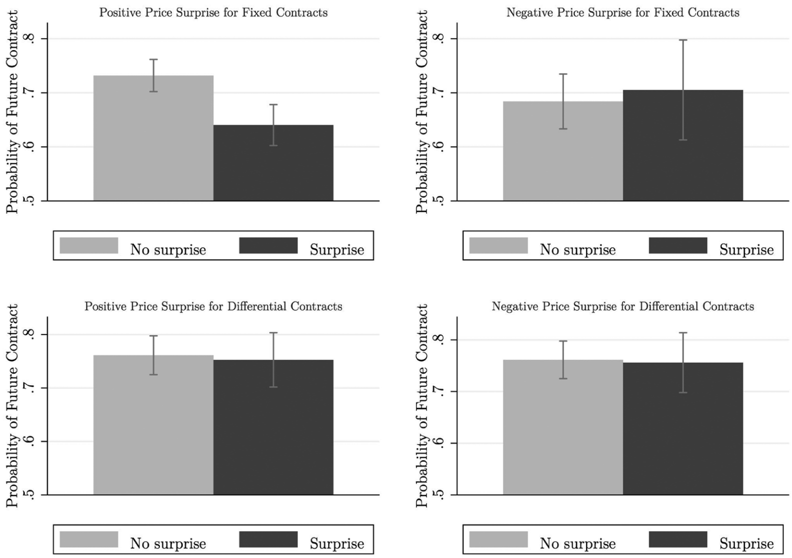

FIGURE IX

Relationship Termination following a Default: Positive versus Negative Price Surprises

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 nonperformance. 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.

The correlation in Table IV, column (1) could, however, also be driven by the mill's demand. For example, the mill could default when it anticipates 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 are being punished, but because they are no longer in operation). Building on Figure IX, we attempt to distinguish between the two hypotheses by instrumenting default with the price surprise. Column (2) shows a decent first stage (despite the small number of 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, loans are eventually repaid in essentially all cases of default. So the possibility that the mill needs fewer 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 that 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 a 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 (which are 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 Online Appendix C).

2. The (Surprising) Fate of Defaulting Mills. There is a completely different (albeit indirect) type of evidence supporting the idea that a large share of the documented defaults are indeed 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 still to be operating in the years following a default. We 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 have gone bust. Online Appendix Table A13 shows that almost none of the mills that defaulted have gone bust. The evidence is thus inconsistent with the hypothesis that defaulting mills simply suffered financial difficulties.[36]

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.[37]

The baseline estimate provides an upper bound as follows. Since we expect no default with a price surprise of less than 1, 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 with positive price surprises to the strategic motive. We observe a default rate of 14.5% when price surprises 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 $\frac{7.7\%}{13\%} = 59\,\%$.[38] We interpret this as an upper bound because we allow that some of this default is due to debt overhang 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 changes in the range of 1–3% (Online Appendix Table A11) and no effect outside the range. The estimates imply that 60–65% of

36. While the documented extent of mill survival might be surprising, fixed assets invested in the countryside 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 suboptimal decisions. Second, the decision might have been taken by the mill's manager in their 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.

37. Recall that by strategic default we specifically mean defaults caused by ex post moral hazard, rather than defaults caused by general willful 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 a case for exoneration).

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

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.[39] Note that because 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 need to explicitly embed the strategic default decision into an optimal contracting framework in which parties endogenously choose 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.

IV.A. Theoretical Framework

1. 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.

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

2. Production. One unit of coffee purchased from farmers produces $\frac{1}{a}$ 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 η, $\rho > 0$. The mill's cost of producing q units of output is given by $C(q) = q \times a \times \omega(q)$, that is, $C(q) = \gamma q^{1+\eta}$ with $\gamma = \rho \times a$.[40]

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

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

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

When the contract is on differential, the buyer's participation constraint collapses to $\Delta_c \leq 0$.[41]

For loan contracts, the mill borrows from the lender the working capital necessary for production. The mill is subject to limited liability, that is, 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 is given by $r_c = \frac{D}{L} - 1$. Assuming a risk-free interest rate equal to r, the lender's participation constraint is given by

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

40. 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.

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

4. Default and Enforcement. After p_w is realized, the mill decides between honoring the forward sale contract or selling the contracted coffee q_c to an alternative buyer at price p_w and defaulting.[42]

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, and with $\mathbf{U^D}$ 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 behavior. Learning about $\mathbf{V}$ is also necessary to perform counterfactuals.

5. 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 $u(\cdot)$, with $u' > 0$ and $u'' \leq 0$, and normalizing $\mathbf{U^D}$ to 0, expected utility is given by

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

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).[43]

6. 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

42. Online Appendix A considers a more elaborate side-selling process that underpins the baseline definition of default.

43. 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 because of the standard efficiency versus rent extraction trade-off. We also abstract from the mill's internal funds. Those would also not alter the qualitative predictions of the model and are taken into account in the calibration exercise.

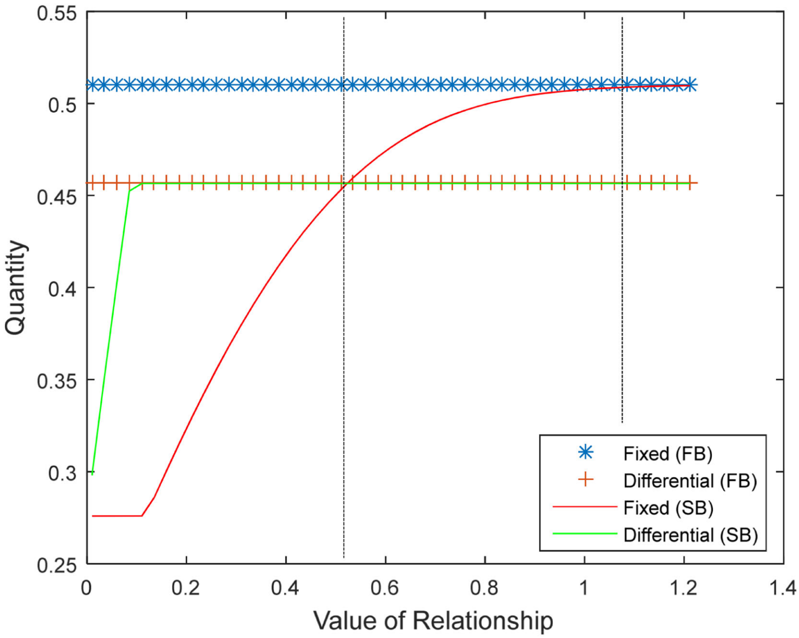

FIGURE X

Insurance versus Enforcement Trade-Off

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 pledgeable 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 forgoing price insurance and signing a contract on differential. This mitigates the strategic default motive and increases pledgeable income relative to a fixed-price contract. The model thus predicts that fixed-price contracts are more likely to be observed in relationships with high **V**. Evidence in Online Appendix E confirms this hypothesis.

which the mill can commit to repay the loan, that is, $\mathbf{I}[p_w] = 1$ 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 that is independent of the realized world price 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 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.

7. Strategic Default: Second Best. When contracts are not enforceable, the mill might decide to default. This decision trades off the short-run gains associated with side-selling and avoiding loan repayment against the loss in relationship value $\mathbf{V}$. Upon observing realized world prices p_w the mill defaults on the contract if[44]

$$\delta \mathbf{V} \leqslant u\big(\pi^D(p_w)\big) - u\big(\pi^R(p_w)\big). \tag{7}$$

8. Contract Choice. The possibility of strategic default introduces a trade-off between insurance and counterparty 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 pledgeable income and the amount of production: the mill is credit

44. The following incentive constraint adapts the incentive constraints in Section III to the case of a risk-averse mill, 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 $u(D) \leqslant \mathbf{V}$. Otherwise, high realizations of p_w make the mill less likely to default.

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.[45]

The model implies that relationships with higher value $\mathbf{V}$ sign fixed-price contracts that leave them exposed to strategic default. Because 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.[46]

9. Discussion of Modeling Choices. Before moving on to the calibration, it is worth pausing to discuss our main modeling choices. A first key assumption is that the model takes the value of the relationship $\mathbf{V}$ as exogenous. $\mathbf{V}$ could be microfounded in a variety of ways. For example, it could arise as part of a subgame perfect equilibrium in which strategic default is deterred, or in a model featuring both adverse selection over the mill's type and strategic default.[47] For example, in Antràs 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

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

46. Online 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 see why, note that under a fixed contract the likelihood of default is given by $P^F(\mathbf{V}) = 1 - F\left(\frac{u^{-1}(\mathbf{V}+u(p_c q_c - D))}{q_c}\right)$. 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.

47. 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 microfoundation for why a bad type on a fixed-price contract would (sometimes) default.

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 modeling strategy and abstract from adverse selection. We cannot and do not intend to rule out models with asymmetric information on types.[48]

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 **V**. An alternative modeling 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 we prefer our current modeling choice. Our conversations with the lender (and the evidence on relationship termination) suggest that although the loan is almost always repaid eventually, 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 question whether the relationship with the lender provides any

48. Because our data come from the lender, the relationship with the buyer is measured with significant error in our data. Because 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.

TABLE V
CALIBRATION: INPUTS

Panel A: Parameters						
Parameters	Values					Source
Price surprise $F(p_w)$	log-normal, $\mu = 0.0152$ and $\sigma = 0.225$					Data
Mills' risk aversion	$\alpha = 0.386$					Data (calibration)
Farmers' supply elasticity	$\eta = 0.6$					RDD
Panel B: Loan-specific values						
	No. 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 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 Online 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 the cost of raw materials $C_c(q_c) = \gamma_c \times q_c^{1+\eta}$ allows us to directly compute γ_c for all observations for which we have audited financial accounts. The value of the relationship $\mathbf{V}$ is then backed out for each loan. Specifically, given a set of parameters, we find the $\mathbf{V_c}$ that rationalizes a loan's key contractual outcomes: the interest rate r_c and whether the loan is backed by a fixed or a differential contract.

deterrent against default. The logic in Bulow and Rogoff (1989) holds if upon defaulting the borrower cannot be excluded from perfect insurance and saving markets. The evidence in this article 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 stealing the proceeds from not repaying the loan).[49]

IV.B. Calibration Strategy

The model is based on a limited set of parameters. Many are directly observed in the data or can be calibrated or estimated. The key parameter we want to learn about is $\mathbf{V}$, the value of the relationship(s) with the buyer and the lender. We pursue the following strategy. We "invert" the model and obtain an

49. Relational dynamics documented in Online Appendix B are consistent with slow build-up of relationships in which defaulting forgoes or delays future growth opportunities and is thus costly.

TABLE VI
CALIBRATION: RESULTS

Panel A: Baseline					
Variable	No. obs.	25th pctl.	50th pctl.	Mean	75th pctl.
Relationship value ($\mathbf{V_c}$)	307	34%	44%	158%	133%
Output loss $X^* = \left(1 - \frac{q_c}{q_c^F}\right)$	307	0%	19.7%	15.8%	19.9%
Output loss X_F^*	108	0%	0%	11.3%	32.8%
Output loss X_D^*	199	19.7%	19.8%	18%	20%
Output loss X_D^D	199	0%	0%	1%	0.1%
Output loss X_D^F	199	51.2%	55.2%	53.3%	55%
Wedge ($MPKV_c$ - r)	307	0%	0%	6%	4%
Wedge (if > 0)	112	4%%	8%	20%	16.7%
Panel B: Robustness to risk aversion (α)					
Moment	$\alpha = 0.286$	$\alpha = 0.336$	$\alpha =$ **0.386**	$\alpha = 0.436$	$\alpha = 0.486$
Output loss X^* (mean)	11.6%	11.5%	16%	16.4%	15.6%
Output loss X^* (Std. dev.)	10.5%	10.9%	12.2%	12.3%	13.4%
Panel C: Robustness to farmers' supply elasticity (η)					
Moment	$\eta = 0.50$	$\eta = 0.55$	$\eta =$ **0.60**	$\eta = 0.65$	$\eta = 0.70$
Output loss X^* (mean)	17.6%	16.2%	16%	13.8%	11.6%
Output loss X^* (Std. dev.)	13.6%	12.6%	12.2%	11.9%	12.4%

Notes. The table reports the results for the calibration exercise (see Online Appendix A for details). Panel A reports the baseline results with the parameters described in Table V. 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 Online 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. Panels 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.

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

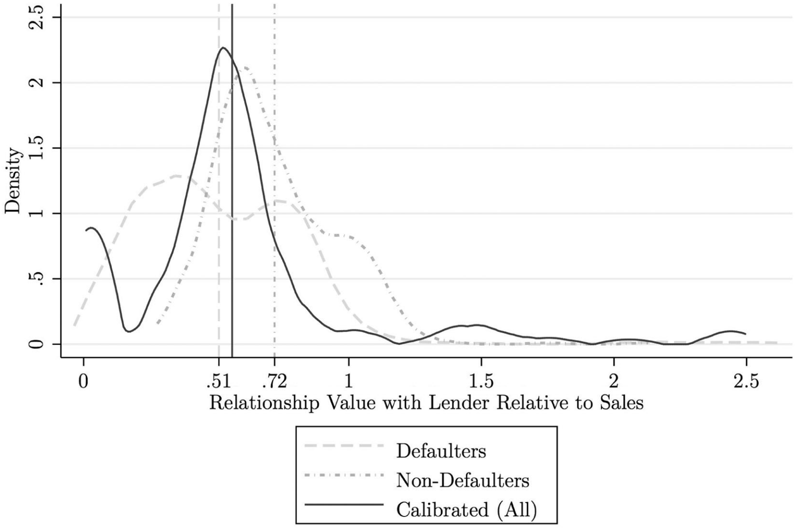

FIGURE XI

Upper Bound, Lower Bound, and Calibrated Relationship Values

The figure compares estimated relationship values **V** with those obtained from a revealed preference approach. Using the incentive compatibility constraint in Online Appendix A, it is possible to derive upper and lower bounds for the value of the relationship for defaulters and nondefaulters, 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 nonparametrically the distribution of relationship values **V**. Even a parametric approach would not 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. Despite using completely different sources of variation, the estimated relationship value **V** distributions display a significant overlap.

Figure X). This makes the identification of $\mathbf{V_i}$ particularly transparent.[50] They are also recorded without error in the data set.[51]

50. Conditional on γ_i, the interest rate and contract type are strongly correlated with each other (p-value of .00). The estimates match the correct contract type approximately 90% of the time.

51. See Online Appendix A for details.

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.[52] Finally, the slope of the coffee cherries supply curve, η, is estimated from the RDD analysis presented in Online 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 US$100,000. 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, 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). Although we do not have strong priors on the appropriate functional forms, the chosen parameterization 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 Online Appendix D for η). To assuage concerns, we report sensitivity checks on the calibrated values of 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 $\eta \in [0.5, 0.7]$.

52. 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.

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, nonparametrically, 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, Figure XI compares the estimated relationship value $\mathbf{V_i}$ 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

1. Estimated Relationship Values **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 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 nondefaulting (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?[53] We find that for the median (mean) observation, production would be 19.7% (15.8%) higher in the absence of strategic default.

The average effect masks substantial heterogeneity. The estimates suggest that 26% of mills produce at first best. That is, at the 25th percentile, the relationship value $\mathbf{V_i}$ is sufficiently large

53. 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.

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. The fourth, fifth, and sixth rows 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 (fourth row). The output gap relative to the optimal quantity conditional on a differential contract is minimal (fifth row). This implies that the majority of these mills (62%) are insurance constrained, that is, 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, the sixth row shows that these mills would produce 50% less output if they were forced to sell on a fixed contract. This is a large number that illustrates that mills signing differential contracts would be severely financially constrained if they had to rely on the collateral value of their relationships to insure against price risk.

Finally, in Table VI, the seventh and eighth rows look at the wedge between the physical marginal product of capital (MPK) and the risk-free interest rate. The MPK is the additional quantity the mill would produce if it was given an additional unit of capital at the loan interest rate. As in the second through sixth rows, 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 producing at first best (26%) or are insurance constrained (39%), the wedge is equal to 0: 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 MPK approximately equal to 30%. These results are in line with two pieces of evidence in the RDD analysis in Online Appendix 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.[54] Second, on the lower threshold (where we do find evidence of credit constraints) we estimate an average gap between MPK and r of about 7%, almost identical to the median estimated by calibrating the model.

Finally, it is worth noting that the lower output produced by the mills as a result of strategic default has implications for farmers welfare. In particular, we can bound farmers' welfare losses as follows. As an upper bound, we can interpret the cherries supply curve as the farmers' supply curve and infer a $\left(\frac{1}{0.84}\right)^{1+\eta} - 1 \approx 32\,\%$ higher welfare for farmers supplying the average mill in the absence of strategic default. As a lower bound, we can ignore any quantity response and simply use the increase in prices paid to farmers as a result of larger loans (Online Appendix Table A15, column (4)). These estimates suggest that at the average mill farmers' welfare would be $\left(\frac{15.8\%}{20.4\%}\right) \times 13.4\% \approx 10.4\%$ higher in the absence of strategic default, still a sizable effect.

2. Correlates of Estimated Relationship Values **V**. The calibration recovers estimates of relationship values. Online Appendix Table A18 projects the estimated relationship values on observable 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 favor of certain (causal) determinants of relationship values.

With respect to relationship-level variables, Online Appendix Table A18 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 match those in Online 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.

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

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 allows 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, 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 article develops a test to identify strategic default. The test builds on 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 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 article studies a common problem in a specific context. The article identifies strategic default and quantifies the costs generated by lack of contract enforcement. These costs appear 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 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 favor of a context-specific, one-size-does-not-fit-all approach to policy recommendations to be drawn from our analysis.

More specific policy implications can be drawn for developing countries aiming at improving exports, particularly in agricultural chains. Many developing countries heavily rely on export revenues generated in a 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 interfirm 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 economy standards, very large (see Hsieh and Olken 2014; Banerjee and Duflo 2014). Because 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

Supplementary Material

An Online Appendix for this article can be found at *The Quarterly Journal of Economics* online. Code replicating tables and figures in this article can be found in Blouin and Macchiavello (2019), in the Harvard Dataverse, doi:10.7910/DVN/CMBISS.

References

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

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

Antràs, 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," NBER Technical Report, 2008.

Baker, George, Robert Gibbons, and Kevin J. Murphy, "Subjective Performance Measures in Optimal Incentive Contracts," *Quarterly Journal of Economics*, 109 (1994), 1125–1156.

———, "Relational Contracts and the Theory of the Firm," *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.

Blouin, Arthur, and Rocco Macchiavello, "Replication Data for: 'Strategic Default in the International Coffee Market'," (2019), Harvard Dataverse, doi:10.7910/DVN/CMBISS.

Breza, Emily, and Andres Liberman, "Financial Contracting and Organizational Form: Evidence from the Regulation of Trade Credit," *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, Berkeley, 2016.

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*, 108 (2018), 3778–3813.

Chiappori, Pierre-André, and Bernard Salanié, "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.

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.

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," *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, 2014.

Ellingsen, Tore, and Eirik Gaard Kristiansen, "Financial Contracting under Imperfect Enforcement," *Quarterly Journal of Economics*, 126 (2011), 323–371.

Fafchamps, Marcel, *Market Institutions in Sub-Saharan Africa: Theory and Evidence* (Cambridge, MA: 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," *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," *Journal of Finance*, 68 (2013), 1473–1515.

Hart, Oliver, "Hold-Up, Asset Ownership, and Reference Points," *Quarterly Journal of Economics*, 124 (2009), 267–300.
Hart, Oliver, and John Moore, "Default and Renegotiation: A Dynamic Model of Debt," *Quarterly Journal of Economics*, 113 (1998), 1–41.
Holmström, Bengt, "Moral Hazard and Observability," *Bell Journal of Economics*, 10 (1979), 74–91.
Hsieh, Chang-Tai, and Peter J. Klenow, "The Life Cycle of Plants in India and Mexico," *Quarterly Journal of Economics*, 129 (2014), 1035–1084.
Hsieh, Chang-Tai, and Benjamin A. Olken, "The Missing 'Missing Middle'," *Journal of Economic Perspectives*, 28 (2014), 89–108.
Joskow, Paul L., "Price Adjustment in Long-Term Contracts: The Case of Coal," *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," *American Economic Review*, 98 (2008), 967–989.
Lacker, Jeffrey M., and John A. Weinberg, "Optimal Contracts under Costly State Falsification," *Journal of Political Economy*, 97 (1989), 1345–1363.
Little, Peter D., and Michael Watts, *Living Under Contract: Contract Farming and Agrarian Transformation in Sub-Saharan Africa* (Madison: University of Wisconsin Press, 1994).
Macchiavello, Rocco, and Josepa Miquel-Florensa, "Vertical Integration and Relational Contracts: Evidence from the Costa Rica Coffee Chain," Mimeo, LSE, 2018.
Macchiavello, Rocco, and Ameet Morjaria, "Competition and Relational Contracts: Evidence from Rwanda's Coffee Mills," Mimeo, 2015a.
———, "The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports," *American Economic Review*, 105 (2015b), 2911–2945.
MacLeod, Bentley W., "Reputations, Relationships, and Contract Enforcement," *Journal of Economic Literature*, 45 (2007), 595–628.
Manova, Kalina, "Credit Constraints, Heterogeneous Firms, and International Trade," *Review of Economic Studies*, 80 (2012), 711–744.
Mayer, Christopher, Edward Morrison, Tomasz Piskorski, and Arpit Gupta, "Mortgage Modification and Strategic Behavior: Evidence from a Legal Settlement with Countrywide," *American Economic Review*, 104 (2014), 2830–2857.
Oster, Emily, "Unobservable Selection and Coefficient Stability: Theory and Evidence," *Journal of Business & Economic Statistics*, (2017), 1–18.
Rampini, Adriano A., and S. Viswanathan, "Collateral, Risk Management, and the Distribution of Debt Capacity," *Journal of Finance*, 65 (2010), 2293–2322.
Roberts, Michael R., and Amir Sufi, "Financial Contracting: A Survey of Empirical Research and Future Directions," *Annual Review 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, Stanford, 2017.
Varangis, Panos, and Bryan Lewin, "Approaches to Managing Coffee Price Risks," in *Agricultural Commodity Markets and Trade: New Approaches to Analyzing Market Structure and Instability*, Alexander Sarris and David Hallam, eds. (Cheltenham: Edward Elgar, 2006), 282–303.