Journal of Development Economics 179 (2026) 103611

Contents lists available at ScienceDirect

Journal of Development Economics

journal homepage: www.elsevier.com/locate/devec

Regular article

Unreliable firms: Evidence from Rwanda ☆

Vishan Gandhi Nigam [a], Brandon Joel Tan [b,*]

[a] Massachusetts Institute of Technology, United States of America
[b] Harvard University, United States of America

ARTICLE INFO

JEL classification:
F14
O12
O14

Keywords:
Trade
Firms and development
Technological choice

ABSTRACT

This paper measures *reliability* – whether firms execute transactions on-schedule – for the universe of Rwandan formal firms using transaction timing data and describes the characteristics of reliable firms. Reliable firms have larger interfirm sales, export more, supply exporters and multinationals, and transact with other reliable firms. Reliable firms are less sensitive to supply chain disruptions. Supplying an MNC increases seller reliability even when servicing non-MNC buyers.

1. Introduction

A common complaint about firms in developing countries is their *unreliability*: products do not arrive according to plan. Aimed at combating this challenge, policymakers often enact policies targeting '24/7 electricity', provide expedited border customs processing to selected firms, and construct special economic zones (SEZs) where manufacturers can operate free of interruption. Given limited budgets, these efforts can have large tradeoffs: 24/7 electricity, for instance, can mean slower expansion of power to new regions. Yet despite considerable work in the economics literature on specific constraints to reliability – including poor transport infrastructure, electricity outages, contracting challenges, and customs delays – we still have a limited understanding of why reliable input supply matters to buyers, why suppliers might choose to be more reliable, and what exogenous factors influence that choice.

The goal of this paper is to measure one notion of reliability – whether firms execute transactions on-schedule – and describe the characteristics of reliable firms.

We begin by offering a framework for how reliability matters for firm performance – in other words, how it might enter the firm production function. In our theoretical framework, firms receive a premium for reliable output, which is complementary in two inputs: the reliability of the suppliers that provide upstream materials, and the firm's exogenous management quality. We examine this production function because, in the spirit of an O-Ring model (Kremer, 1993), such complementarities have important development implications. In particular, reliability depends critically on whether input suppliers can match the final good providers' choices. Thus, entire chains of firms need to reliably supply each other to achieve successful production – particularly when the customers are ex-ante productive firms – and the presence of unreliable firms upstream can inhibit production downstream.

In order to provide empirical evidence for such complementarities, we first generate estimates of firm-level reliability that can apply to an entire developing country. Most existing work on this 'second moment' of firm production follows one of two approaches. The first is to measure variability at some aggregate level. For instance, Bloom (2009) uses a structural model to study an aggregate uncertainty shock, and Asker et al. (2014) use production surveys to document industry-level volatility in productivity. In contrast, we focus on firm-level reliability, which differs from economy- or industry-wide shocks in that unreliability is endogenous (optimizing firms choose their reliability level) and other firms in the economy have the option of substituting to more or less reliable firms.

A second approach implicitly infers a firm's reliability from its resilience to a specific kind of shock. These shocks include both extreme events, such as civil wars (Macchiavello and Morjaria, 2015), earthquakes (Carvalho et al., 2021; Boehm et al., 2019) and weather shocks (Barrot and Sauvagnat, 2016); and recurring disruptions, such as cost overruns in the IT industry (Banerjee and Duflo, 2000), power outages

☆ We thank Pol Antras, David Atkin, Vittorio Bassi, Lauren Bergquist, Dave Donaldson, Esther Duflo, Nathan Hendren, Edward Glaeser, Michael Kremer, Marc Melitz, Anna Twum, seminar participants at NEUDC, and workshop participants at MIT and Harvard for helpful comments. We thank the Rwanda Revenue Authority and Rwanda Development Board for providing the data and the international growth centre for facilitating access. All errors are our own.

* Corresponding author.
E-mail addresses: vishan@mit.edu (V.G. Nigam), btan@g.harvard.edu (B.J. Tan).

https://doi.org/10.1016/j.jdeveco.2025.103611
Received 10 September 2024; Received in revised form 26 June 2025; Accepted 12 August 2025
Available online 16 September 2025

(Allcott et al., 2016), freight congestion (Firth, 2017), and delayed ice deliveries (Ghani and Reed, 2022). While we share the firm-level focus of these papers, our contribution here is to develop a metric of ‘day-to-day’ supplier reliability that can be measured for every formal firm in the economy, regardless of industry, and used to examine longer-run differences in output.

Our metric of unreliability depends on the *timing* of a firm’s sales: based on the intuition that reliable firms should transact at more regular intervals with the same customer, we compute the coefficient of variation (CV) of the time between consecutive transactions and average across all of a firm’s relationships. A firm that sells every Monday will have a CV of 0, while deviations due to missed or delayed transactions will increase the CV. This calculation is possible only because our dataset – the universe of electronic VAT records from Rwanda for 2017 – is unique in recording the exact date of each transaction between any two formal firms.[1] Since the time between transactions can also vary due to buyer demand, we leverage the network structure of the VAT data and exclusively compare firms selling to the *same buyer*, therefore ensuring all conclusions are about a firm’s reliability as a seller. We therefore directly observe the average output reliability level of each firm, as well as that of its suppliers and buyers; in contrast, most economy-wide work in trade treats non-price characteristics of a firm’s output (i.e. its quality) as an unobserved residual (Khandelwal, 2010).

Since the CV of time between transactions is technically a measure of transaction *irregularity*, before moving to analysis we state and validate the assumptions necessary to interpret the CV as an *unreliability* metric. First, we assume buyers weakly prefer suppliers that follow the same schedule in consecutive time periods, which means we can learn transaction schedules from relationship-specific patterns. Second, we assume that schedules consist of deliveries at *evenly-spaced intervals*, allowing us to summarize deviations from schedules via a single coefficient of variation. Intuitively, evenly-spaced deliveries would be desirable if the buyer has to plan coordinate the arrival of many inputs for their own production or faces continuous demand. To ensure we observe relationships in which this is likely to hold, and for which we can estimate mean time between deliveries with limited noise, in all specifications we subset to relationships with sufficient repeated interaction – at least 10 transactions per year or 5 per quarter – and in which both the seller and the buyer are VAT-paying firms. Since this metric is somewhat restrictive, the Appendix also includes a series of case studies in which we show that our metric coincides with other intuitive notions of reliability. We first show that in low-CV relationships, transactions are more concentrated on fewer unique weekdays or days of the month, where concentration is measured via a Herfindahl–Hirschman index (HHI). We also show that, among both food wholesalers and distributors of Rwanda’s leading beer (Bralirwa), low-CV firms also appear to exhibit smaller swings in inventories because they more closely time purchases and sales. In other words, selling less reliably may correlate with poor inventory management behavior that can lead to stockouts, as in Kremer et al. (2013).

Armed with our measures of firm-level reliability, we document several patterns that are consistent with the production function for reliability laid out in our theoretical framework. First, we find that reliable firms sell more to other firms, so a significant reliability premium does exist in interfirm transactions. Second, we document that exporters, who are typically more productive, are also more reliable (Melitz, 2003). This result suggests that reliable output depends on internal management quality, for instance to coordinate production. Third, looking upstream, typically productive firms (MNCs and exporters) choose significantly more reliable suppliers, suggesting that productive are willing to pay for reliable supply or a close correlate of it. The fact that ex-ante productive firms buy reliable inputs and sell reliable output strongly indicates that reliability choices are complementary across firms. In contrast, if a firm could simply choose high reliability on its own – that is, by investing in inventories or hiring additional units of labor – exporters and MNCs would not need reliable suppliers. We also go on to estimate the mean unreliability of every supplier, and find evidence of assortative matching: across all Rwandan formal firms, reliable firms trade with each other. This result is consistent with one of two closely related stories: first, that input reliability *mechanically* creates output reliability because delayed inputs make smooth production impossible; and second, that conditional on choosing an output reliability level, reliable firms *endogenously* choose reliable suppliers due to higher willingness to pay. In either story, the implication is that there are a limited number of reliable suppliers from which to choose.

The next part of our paper augments the cross-sectional evidence for complementarities with two event study exercises.

First, we examine how reliable and unreliable firms respond to a disruption of trade with Kenya, Rwanda’s second-largest source of imports after China, due to a disputed presidential election. Using a differences-in-differences design, with exposure based on the share of imports coming from Kenya in a pre-period, we document that more reliable firms (below median CV) experience smaller declines in total imports during the disruption. The natural experiment provides a separate piece of evidence that firms who supply reliable output also invest in secure input supplies. More importantly, it suggests that reliability levels are partially due to a fixed characteristic of the firm, as in our theoretical framework. If firms could adjust their chosen reliability each period, then the input flows of ex-ante reliable firms would be equally sensitive to input disruptions rather than less so.

Second, MNC buyers might causally change supplier reliability, either due to unpriced productivity spillovers or because their demand leads to reliability upgrading (Atkin et al., 2017; Verhoogen, 2008). To test for such causal effects, we estimate event studies of the impact of selling to MNCs on supplier reliability following methods in Alfaro-Urena et al. (2020). We find that, in the months after their first link to an MNC, Rwandan firms become more reliable (lower-CV) in their existing relationships with other Rwandan firms; this suggests that MNC entry can force improvements that spill over to other trading relationships, providing a potential justification for developing-country policies that seek to attract foreign direct investment. Together, our results suggest that firm reliability matters for performance, and is strongly linked with the reliability of trading partners.

This paper builds on several literatures. First, by developing a firm-level measure of reliability that applies to an entire economy, we build on existing work in macroeconomics on industry- or economy-wide uncertainty (e.g., Bloom, 2009), as well as microeconomic analyses of firm-level responses to specific supply chain disruptions (e.g., Macchiavello and Morjaria, 2015).

Second, while our focus on reliability is novel, this paper is closely related to existing work on product quality in international trade. In particular, our theoretical framework is heavily based on the O-Ring model of development (Kremer, 1993) and similar models of interlocking quality choices (Kugler and Verhoogen, 2011; Demir et al., 2021; Fieler et al., 2018), and our measurement of reliability follows a long literature on measuring quality across firms (see Schott, 2004; Hallak and Schott, 2011; Khandelwal, 2010). In most of this literature (Atkin et al., 2017 being a notable exception) quality is typically not observed directly, so it inferred either from unit prices of material and labor inputs (i.e. wages) or as a structural residual that rationalizes observed choices in a demand model (as in Khandelwal, 2010). Our approach is closer to the latter in that we observe the reliability of a firm’s practices directly under the assumption that buyers demand transactions at (weakly) more evenly-spaced intervals.

Finally, because our metric of firm reliability is based on patterns of repeated interaction, our approach relates to the existing development literature on relational contracts, as summarized in Macchiavello

[1] VAT is only chargeable upon delivery.

(2021). Much of this work focuses on relationship length: for instance McMillan and Woodruff (1999) show that older trade credit relationships have better repayment behavior, and more recently Houeix and Wiles (2025) conduct an intervention that increases imports of high-quality garments to Senegal by enabling long-term relationships. In contrast, our metric leverages heterogeneity in behavior within long-term relationships. Specifically, we rely on the time between subsequent transactions within a specified time period (a quarter or year), not on how many time periods elapse before the relationship ends. This focus on timing is enabled by the precise transaction timestamps in the Rwandan electronic VAT system; in contrast, most countries' VAT systems rely on monthly or quarterly firm filings, and thus lack precise transaction timing information.

The remainder of this paper proceeds as follows. Section 2 details our theoretical framework in which firms jointly choose the reliability and quantity of their output. Section 3 describes the Rwandan administrative data and our methodology for calculating firm-level reliability. Section 4 contains cross-sectional evidence on the characteristics of reliable firms. Section 5 presents our estimates of firm import responses to the 2017 Kenya election. Section 6 presents event-study evidence of how MNC linkages impact supplier reliability. Section 7 concludes.

2. Theoretical framework

To formalize intuition about why reliability can matter for firm performance, in this section we describe a production function in which firms choose both quantity produced and reliability upfront. In our model, the reliability of a firm's output is complementary in a firm's management quality, or TFP, and in the reliability of a firm's inputs. Looking ahead, the model predicts that in the long-run exporters and other high-TFP firms (Melitz, 2003) will have reliable output and choose more reliable input suppliers. In the short-run, since productive firms will endogenously choose more robust supply chains, they should reduce inputs by less in response to shocks. We proceed to verify these predictions in Sections 3 and 4.

Formally, we follow Kugler and Verhoogen (2012) in representing reliability as a form of quality. Let r_y be the (observed) reliability of a firm's output and r_s be the reliability of its materials supplier. Intuitively, r might be the probability that a good arrives at its destination as desired. A is the firm's management quality or 'type'. Producing reliable output is then Leontief:

$$r_y = \min(A^\alpha, r_s) \tag{1}$$

for some $\alpha < 1$. This production function embeds two assumptions about reliable production. First, reliable production requires reliable suppliers, because even one missing input can stop production and short-run substitution is impossible for many intermediates. This assumption plausibly holds for electricity Fried and Lagakos, 2020; Burgess et al., 2020 as well as for US importers of intermediate goods (Boehm et al., 2019). Second, supplier reliability is complementary to a firm's own management quality, reflecting the idea that certain firm capabilities – in inventory software, 'just-in-time' production methods (Bloom and Van Reenen, 2007), vertical integration of some inputs, or relationship formation with new suppliers (Grossman et al., 2021) – enable them to take advantage of reliable input supply. For simplicity we represent A^α and r_s as perfect complements.

The firm receives (pays) a premium p_r (c_r) for producing output (sourcing inputs) with an additional unit of reliability. We assume the firm is a price-taker in both input and output goods markets, which is plausible for a small open economy like Rwanda.[2] It then chooses quantity q_s and reliability r_s of its inputs to maximize profit, subject to reliability prices p_r and c_r, production function (1), quality-adjusted input price c_q per unit of output, and quality-adjusted output price normalized to 1:

$$\max_{q_s, r_s} \quad p_r \min(A^\alpha, r_s)\ Aq_s^\beta - (c_r r_s)(c_q q_s) \tag{2}$$

where $\beta < 1$. Thanks to Leontief production of reliability, if the firm produces it will always set:

$$r_y = r_s = A^\alpha \tag{3}$$

Note that as long as r_y is complementary in r_s and A, we would still have that higher-type firms (high A) are more reliable (high r_y) and have reliable input supply (high r_s). The firm's problem is then:

$$\max_{q_s} \quad p_r A^\alpha\ Aq_s^\beta - (c_r A^\alpha)(c_q q_s) \tag{4}$$

Taking the FOC with respect to q_s and solving, we obtain:

$$q_s^* = \left(\frac{A\beta p_r}{c_r c_q}\right)^{\frac{1}{1-\beta}} \tag{5}$$

which means we will obtain total sales of:

$$Sales^* = p_r r_y^* A(q_s^*)^\beta \ = p_r A^\alpha A\left(\frac{p_r A\beta}{c_r c_q}\right)^{\frac{1}{1-\beta}} = p_r^{1+\frac{1}{1-\beta}} A^{\alpha+1+\frac{1}{1-\beta}} \beta^{\frac{1}{1-\beta}} (c_r c_q)^{\frac{-1}{1-\beta}} \tag{6}$$

and input use of:

$$Inputs^* \ = c_r A c_q q_s^* \ = c_r c_q A\left(\frac{A\beta p_r}{c_r c_q}\right)^{\frac{1}{1-\beta}} \tag{7}$$

Eqs. (3) and (6) give rise to several testable predictions that we should to observe in a large sample of firms. Assuming firms are price-takers in partial equilibrium, we should expect the following:

Prediction 1: $Cov(r_y, Sales) > 0$. Firms with larger sales should be more reliable. Since $Sales = p_r r_y A q_s^\beta$, this correlation is partly mechanical: the prices used to calculate sales will reflect any premium p_r firms face for reliability. However it also may reflect that, incentivized by high p_r, firms (particularly those with underlying high productivity A) will endogenously sell more. In both cases, a positive correlation implies that reliability is a 'good' with a positive premium p_r. This test is necessary because, while residual measures of quality are 'goods' by construction (Khandelwal, 2010), the same is not true of measures of r_y that are directly observed in the data.

Prediction 2: $Cov(r_y, A) > 0$. Productive firms are more reliable. If we see productivity directly, or proxies for it such as direct or indirect export status (Melitz, 2003; Alfaro-Urena et al., 2020), we can avoid the mechanical correlation in Prediction 1 and isolate whether productive firms choose higher reliability due to scale.

Prediction 3: $Cov(r_s, A) > 0$ and $Cov(r_s, p_r) > 0$. If buyers are physically productive and/or have customers willing to pay for reliability, their suppliers will also be more reliable. This prediction reflects that management quality and supplier reliability are complements. Without sufficiently reliable inputs, it is impossible to produce at level A^α despite having type A. This prediction would not arise, for instance, if reliable output just required additional labor-hours per unit produced, or if firms could ensure reliable output by stocking infinite inventories.[3]

Note that Prediction 3 has potentially important aggregate implications. First, by virtue of the complementarity with management, productive (high-A) firms have higher returns to reliable inputs r_s, and thus input shortages that affect high-A firms will be particularly bad for output. Second, if reliability choices are interlinked, an exogenous lack of reliable inputs upstream in the production chain can inhibit reliable production elsewhere in the network even if superior management quality A exists, as in Kremer (1993) or Liu (2019). We discuss these possibilities more fully in Section 6.

[2] In our data, over 30% of all firms are direct importers and the majority of firms make a purchase from an importer.

[3] Kugler and Verhoogen (2011) generate similar predictions for correlations between productivity and input *price*, which is their proxy for quality.

Since Eqs. (3) and (6) holds in the long-run, we test for Predictions 1–3 using a single cross-section of all Rwandan VAT-paying firms. Looking ahead, we find patterns consistent with our model and as well as multiple general equilibrium models that, while not explicitly about reliability, would generate similar patterns. For instance, Kugler and Verhoogen (2011) embed unobserved quality in a Melitz model using the production function in (1), but with r_s obtained through additional units of labor rather than an upstream firm, and show that it generates positive correlations between firm sales, input prices, and output prices. O-Ring-style models (Demir et al., 2021; Kremer, 1993) also predict that productive firms will both hire high-skilled workers and source from other firms that do the same. In contrast, in models where output reliability or quality does not depend on suppliers, we should not observe correlations in reliability across the supply chain.

Finally, to consider how the firm responds to supply chain disruptions, we briefly discuss what happens to output when c_r increases so that an uninterrupted input supply becomes temporarily difficult to obtain. We focus on input disruptions to stay close to the quasi-experiment in Section 5, in which we examine how more- and less-reliable firms respond to a border closure that temporarily inhibits Rwandan firms' access to imported inputs.

Prediction 4: In response to an increase in c_r, average input use fall by more for low-A firms than for high-A firms. Let $\bar{A}(c_r)$ be the productivity level at which $Sales^* - Inputs^* = 0$. In response to an increase to $c_r' > c_r$, firms $A > \bar{A}(c_r')$ will continue to set $r_y = A$ but reduce input use according to Eq. (7). Firms with $\bar{A}(c_r) < A < \bar{A}(c_r')$ will temporarily shut down and set $Inputs = 0$.

Intuitively, no firm changes its production process in the short run because it is pinned down by A. However, high-A firms will keep operating because they produce with enough output reliability to make operation during the cost shock profitable, while low-A producers will not. In Section 5, we test for Prediction 4 by measuring reliability of all Rwandan importers in the cross-section and then estimating differential responses across importers to a specific trade disruption: the 2017 Kenyan presidential election. We find, consistent with the idea that reliability levels are determined in advance, that firms with high *output* reliability in the cross-section reduce total imports by a smaller amount in response to an equally-large shock to their *input* supply.

3. Measuring reliability

In this section we describe our measure of firm reliability, which is derived from interfirm transactions recorded as part of Rwanda's national value-added tax (VAT) system.

3.1. Data

We use several administrative datasets from the Rwanda Revenue Authority (RRA) that cover the universe of Rwandan taxpaying firms. First, to construct our reliability measure, we use transaction-level data from electronic billing machines (EBMs) in Rwanda. EBMs are cash register-like machines that all Rwandan VAT-paying firms are legally required to possess. When a VAT-registered firm in Rwanda makes a sale, it is required to issue an EBM receipt to the customer. Upon doing so, details on the seller tax ID, buyer tax ID, date, time, and size (in Rwandan francs) of the transaction are automatically transmitted to the Rwandan Revenue Authority, which then taxes firms on the basis of these receipts along with monthly filings. EBM machines were gradually rolled out in Rwanda from 2013–2015 (Eissa et al., 2014; Mascagni et al., 2019). We use transactions from January 1 to December 31, 2017 for our main analysis.

The unique feature of Rwandan EBM data is that transactions are auto-recorded at the time of sale – that is, when a receipt is generated by the seller's EBM machine. Since buyers need to collect their receipt in order to claim input tax credits, the time stamp is when goods are delivered. We therefore observe transaction-specific *dates* within each unique seller-buyer pair for the entire formal economy. In contrast, under most national VAT systems, buyers and sellers self-report transactions at the end of a month or quarter, so the date of each transaction is not known. Starting with the universe of recorded transactions, we subset to transactions in which both the seller and buyer are registered companies who paid corporate tax in the previous year. While some households also appear as sellers, subsetting to registered companies (hereby 'formal firms') allows us to observe firm characteristics, such as industry, age, and import/export history, that may vary with reliability. In addition, subsetting to registered buyers addresses a missing data problem: consumers do not have unique identifiers, and even if they did, VAT sales to consumers are typically under-reported in low-enforcement environments (Pomeranz, 2015). In contrast, reporting of interfirm sales is self-enforcing because the buyer is incentivized to report purchases to reduce her tax liability.

We merge the EBM data by firm ID to several other datasets. First, we use yearly corporate tax returns for 2016 to obtain firm industry and an independent measure of sales. Second, we merge to each firm's full history of export transactions, as well as to a list of multinational corporations (MNCs) provided by the Rwanda Development Board (RDB), allowing us to use MNC status and export status as additional firm characteristics. Third and most importantly, we use transaction-level import data to construct a *daily* panel of import transactions at the importing firm-export country-product (HS6) level from January 1st to December 31, 2017. These data enable us to construct firm-level exposure to a country-level shock – the August 2017 Kenyan election and accompanying threat of violence – based on each firm's pre-period share of imports coming from Kenya. We do so and estimate the subsequent effects on firm-level imports in Section 5. Note that we observe an import for over 30% of registered firms in the EBM sample.

Summary statistics by supplier are shown in Table 1. There are 6070 unique suppliers. The average supplier has 265 million RWF (approximately 250,000 USD) in yearly interfirm sales and links with 22 unique buyers. In all analyses, we subset to links with at least 10 transactions per year or at least 5 transactions when using a six-month period.[4]

Much of our analysis below examines whether certain groups of firms (multinationals, exporters, and their suppliers) exhibit higher reliability. Table 2 shows why these tests are possible with our data: exporters and MNCs source from and sell to many other firms, so the share of *relationships* involving each group exceeds the share of *sellers* that belong to each group. In particular, while only 2.3% of sellers are MNCs as shown in 1, Table 2 shows that MNCs are involved as sellers (buyers) in 7.1% (5.9%) of domestic links. Similarly, while 15.3% of sellers are exporters, exporters are involved as sellers (buyers) in 39.9% (21.8%) of links.

Finally, note that the EBM data are recorded when VAT receipts are exchanged, *not* necessarily when goods are delivered or payment is made, so we cannot observe transactions that fell through after the receipt was issued.

3.2. Calculation of reliability metric

This section describes our definition of reliability and how we construct a firm-level reliability measure for every supplier in Rwanda. While there is no single definition of reliability in the economics literature, in general we interpret reliability as 'delivering when required'. For instance, a reliable electricity grid is one where power is available 24/7, a reliable pharmacist is one that has medicines available when required, and a reliable corner store owner is one that never fails to make sales due to stockouts. The challenge with this definition is

[4] The CV measure in the next section requires at least 3 transactions and is volatile when fewer than 5 transactions are used.

Table 1
Summary statistics at supplier level.

	Mean	SD	Min	Max	N
Supplier ID					6070
Number of buyers	22.467	52.020	1.000	1456.000	6070
Average link value (bn Rwf)	0.035	0.241	0.000	8.147	6070
CV of days between shipments, mean across links	0.810	0.293	0.000	3.173	6070
CV of weeks between shipments, mean across links	0.666	0.243	0.000	2.080	5956
CV of days between shipments, mean across link-quarters	0.736	0.264	0.000	3.173	5852
Interfirm sales (bn Rwf)	0.265	1.081	0.000	25.345	6070
Other sales (bn Rwf)	2.108	5.918	0.000	68.675	6070
=1 if MNC	0.023	0.150	0.000	1.000	6070
=1 if exporter	0.153	0.360	0.000	1.000	6070
Share of sales to MNCs	0.027	0.092	0.000	1.000	6070
Share of sales to exporters	0.075	0.143	0.000	1.000	6070

Notes: Table gives summary statistics at the supplier level. Sample consists of all Rwandan firms that paid corporate tax in 2017 (formal firms). A link is a supplier-buyer pair. CV measures are computed separately for each supplier-buyer with at least 3 transactions in 2017 and then averaged across buyers. Interfirm sales are sales where the buyer is a formal firm, and other sales are to firms with other IDs or no ID. All data are from the Rwanda Revenue Authority (RRA).

Table 2
Link-level exposure to trade.

	N	Share
All links	146 185	1.000
Only buyer exports	18 524	0.127
Only seller exports	45 011	0.308
Both export	13 254	0.091
Neither export	69 396	0.475
Only buyer is MNC	7746	0.053
Only seller is MNC	9466	0.065
Both are MNCs	911	0.006
Neither is MNC	128 062	0.876

Notes: Table gives summary statistics at the link level for our preferred sample year (calendar year 2017). A link is a supplier-buyer pair. Sample consists of all links between suppliers and buyers that paid corporate tax (formal firms) and transacted in 2017. All links shown are domestic Rwandan transactions from electronic billing machines (EBMs). A firm exports if any export transaction is observed in Rwandan Customs records, and is an MNC if is fully foreign-owned or part of a joint venture. All data are from the Rwanda Revenue Authority (RRA).

that while we observe realized transactions in the RRA data, details of planned transactions, such as what the buyer hoped to purchase that day or what was agreed in contract for a customized good, are unobserved. In absence of such data, we must place assumptions on what is demanded:

Assumption 1. in relationships with frequent transactions, buyers weakly prefer their suppliers to be consistent – that is, to follow the same schedule week after week or month after month. Demand for consistency is plausible in interfirm transactions, especially in wholesale and manufacturing. Even if transaction *volumes* fluctuate, intermediate goods producers may prefer regular schedules, with little deviation, in order to best plan production and avoid locking up scarce working capital in inventory. However, this assumption may be violated when buyers prefer 'on-demand' goods and services, such as repairs, in patterns deviating from past histories.

Assumption 2. in relationships with frequent transactions, buyers weakly prefer transactions with the same supplier at equally-spaced intervals. While stronger than Assumption 1, the benefit of this assumption is that a schedule is summarized by a single parameter: the time between consecutive shipments. For instance, a buyer might prefer a weekly transaction schedule (every Monday) from a particular supplier. But if transactions are missing in certain weeks, or occur in 'clumps' of days before running out of stock, under Assumption 2 we can conclude that a supplier follows its schedule less reliably.

We now formalize this intuition. Consider a sequence of n dates where transactions occur between the same seller i and buyer j:

$$T^{ij} = \{d_1^{ij}, d_2^{ij}, \dots, d_n^{ij}\} \tag{8}$$

The number of days between consecutive transactions is:

$$Diff^{ij} = \{(d_2^{ij} - d_1^{ij}), (d_3^{ij} - d_2^{ij}), \dots (d_n^{ij} - d_{n-1}^{ij})\} \tag{9}$$

so we can summarize deviations from equally-spaced transactions using the coefficient of variation of $Diff$:

$$CV^{ij} = \frac{\sqrt{Var(Diff^{ij})}}{\mathbb{E}[Diff^{ij}]} \tag{10}$$

CV_{ij} (henceforth 'the CV') is our metric of unreliability. Note that CV_{ij} is always positive or zero. The CV is 0 if transactions in a relationship occur every x days (regardless of the frequency of transactions $\frac{1}{x}$, since we normalize by the mean time between shipments) and increases otherwise.

As an illustrative example, Fig. 1 plots two months of transactions between the two largest dairy manufacturers in Rwanda and a common buyer. Both these suppliers make regular deliveries to the buyer. However, the top supplier's sales are at irregular intervals, for instance around August 1 or 15. Since the time between consecutive transactions varies from 1 to 14 days, the red supplier has a high CV of 0.78. In contrast, the bottom supplier's transactions are always between 2 and 9 days apart, and are visibly more evenly dispersed, leading to a low CV of 0.463. If these patterns repeat themselves across all of the suppliers' shared buyers, we conclude that the bottom supplier is more *reliable*.

To generate our reliability metric, we first compute the CV as in Eq. (10) for each pair with at least 10 transactions.[5] recorded in 2017 and then average across all pairs involving the supplier. While the CV is itself unitless, to interpret magnitudes note that across suppliers, the mean CV has a mean of 0.81 with a standard deviation of 0.29, as shown in Table 1 We also compute the CV using weeks instead of days to account for idiosyncrasies of the week – for instance it is mechanically impossible to deliver every 3.5 days, so a twice-a-week schedule would have $CV > 0$ – as well as separately for each link-by-quarter, to remove the possibility of long gaps due to a relationship ending and restarting (see Martin et al., 2020). These alternative measures are similarly distributed, have correlations of 0.8 and 0.6 respectively with our primary measure, and do not qualitatively affect any of our empirical results in the following sections.

While the assumption of evenly-spaced transactions is restrictive, our CV metric is correlated with alternative definitions of a regular

[5] The minimum transaction count is to reduce measurement error in the pair-specific CV estimates. We need at least 3 transactions to compute a coefficient of variation.

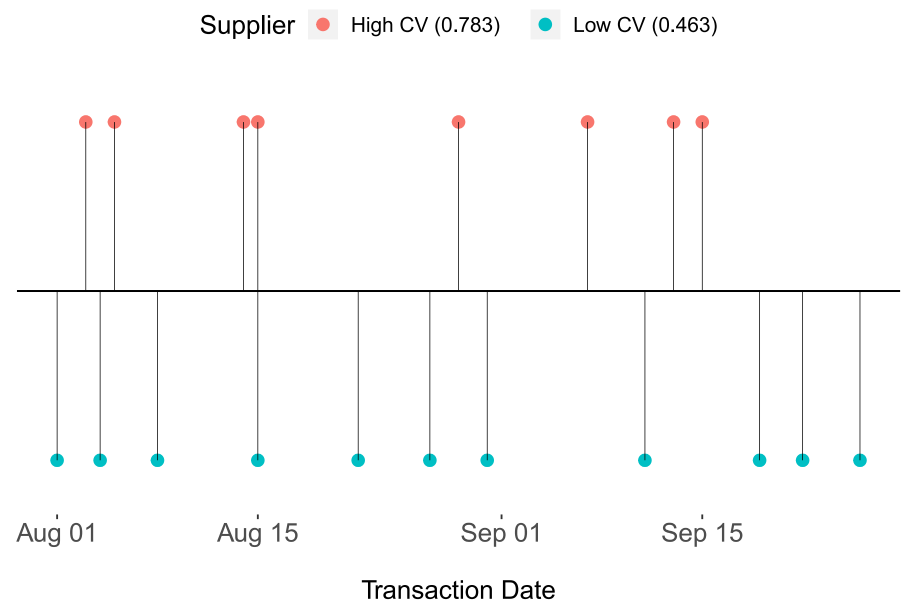

Fig. 1. Reliability calculation example using two dairy manufacturers.
Notes: Figure shows transaction dates for two suppliers of dairy products from August 1 to September 30, 2017 to a common buyer. The CV shown is the coefficient of variation of the time, in days between consecutive shipments. Transaction dates are obtained from VAT electronic billing machine (EBM) records of the Rwanda Revenue Authority (RRA)

schedule. For each relationship in our sample, we compute the share of transactions that occur on each weekday or day of the month and then compute the Herfindahl–Hirschman concentration index of this vector. The HHI will be closer to 1 if transactions happen on the same days of the week (ex: Monday and Wednesday) or month (ex: 1st, 5th, and 15th) even if these days are not evenly spaced apart. In Appendix Figure A1 we show that these concentration indices are negatively associated with our CV metric, with a stronger negative correlation when we compute the HHI using the number of transactions (which should capture actual interactions) rather than the share of value transacted. In other words, the two definitions of a schedule coincide in practice: low-CV firms concentrate their transactions on particular days of the week or month.

As an additional validation check, we can zoom in on a particular narrow industry and examine if low-CV firms engage in practices consistent with reliable supply, and in particular whether they smoothly manage their stocks of inventories. To do so, we narrow in on the 15 wholesalers of Bralirwa products in our sample: these 15 sellers sell the same set of products – either Rwanda's national beer or a Coca-Cola product – and are exclusively supplied by the same firm, the Bralirwa brewing company. For each wholesaler, we plot total purchases from the manufacturer, and total sales summed across all customers, together by date in the VAT data. Appendix Figure A2 shows the plots for a wholesaler with a medium-to-high CV of 0.864 and one with a relatively low CV of 0.642. The two exhibit differing sales patterns: the low-CV supplier makes a sale almost every day, suggesting that it has smooth demand *across* its buyers, while the high-CV supplier has many days with no sales. In addition, in the high-CV firm variation in purchases is much larger than variation in sales. These patterns (which are similar when we visually inspect all 15 suppliers) suggest that low-CV firms have less variable inventories, consistent with the idea that – at least in the fast moving consumer goods sector – low CV arises at least partially from avoiding stockouts

Last, our CV metric is highly correlated with shipment cancellations using shipment data from Karnataka, India. In India, every registered business is required to submit an electronic document (known as an e-way bill) to the government for any planned movement (or shipment) of goods valued above the threshold of Rs. 50,000 ($700 USD).[6] If the planned shipment does not arrive or if the shipment time needs to be modified, then the e-way bill is canceled and a new one is issued if the shipment will arrive at a later date. Data from Garg et al. (2023) reveals that our CV metric is highly correlated with the shipment cancellation rate.[7] This indicates that our measure successfully proxies cancellations and delays in shipments from suppliers.

If we observed the prices that firms charge, we could conduct additional checks. For instance, if reliability is a form of quality that buyers value, low-CV suppliers should charge higher prices for otherwise-equivalent goods. We cannot conduct this validation check in our sample because sellers do not systematically record prices and quantities in EBM machine data.

In summary, in purely mechanical terms the CV measure captures how much the time to next transaction deviates from its average. However, if buyers want to transact according to a repeated schedule (Assumption 1), and desired schedules involve equally-spaced payments (Assumption 2), variation in CV due to the supplier will capture the *unreliability* of the firm.

3.3. Removing demand-side and sector-level factors

Finally, to generalize the comparison in Fig. 1 to all Rwandan firms, we need to remove two sources of variation in the CV that are uninformative of seller reliability. The first is demand-side variation: a buyer j may simply have irregular demand across all its products. For

[6] This includes goods transported by road, air, railways, or water vessel. The law was introduced to increase tax compliance and reduce shipping times. Government officials have the authority to intercept any conveyance to verify the e-way bill or the e-way bill number for all inter and intra-state shipments. The penalty for non-compliance is Rs 10,000 ($ 141 USD) or the value of tax-evaded, whichever is greater. See Garg et al. (2023) for more details.

[7] Garg et al. (2023) computed our CV measure (over all shipments that were not canceled) and the cancellation rate (total cancellations/total e-way bills generated) for each buyer–supplier pair using the universe of e-way bills from April 1, 2018 to August 29, 2019 in Karnataka, India.

Table 3
Seller industries with highest and lowest CV of time between shipments.

Rank	Industry	Mean CV
Lowest CV		
1	Renting of personal and household goods	0.343
2	Other education	0.490
3	Legal activities	0.495
4	Telecommunications	0.501
5	Social work activities	0.579
6	Washing and (dry-)cleaning of textile	0.626
7	Production, transmission and distribution of electricity	0.629
8	Real estate activities on a fee or contract basis	0.633
9	Real estate activities with own or leased property	0.648
10	Publishing	0.656
Highest CV		
1	Agriculture, hunting and forestry	1.057
2	Manufacture of products of wood	0.991
3	Building completion	0.976
4	Manufacture of non-metallic mineral products	0.971
5	Other monetary intermediation	0.957
6	Wholesale of metals	0.953
7	Retail sale of hardware, paints and glass	0.948
8	Manufacture of dairy products	0.945
9	Restaurants and canteens	0.941
10	Library, archives, museums and other cultural activities	0.933

Notes: Table shows the 10 largest and 10 smallest supplier industries by mean CV using Rwanda Revenue Authority (RRA) industry definitions. The mean CV is obtained by calculating the coefficient of variation of the number of days between consecutive shipments between the same seller-buyer pair and then averaging over all pairs by seller industry. Sample consists of all VAT-recorded transactions in 2017 where sellers and buyers are registered firms. Data are from the Rwanda Revenue Authority (RRA).

instance, a buyer that fails to plan its own downstream production in advance will source *all* inputs irregularly, and the subsequent variation in CV should not be attributed to its suppliers. Thanks to the structure of our supply chain data, we observe many buyers per seller and vice versa, and thus can residualize out the average CV of a *buyer* across all of its sellers.

We also residualize out the average CV by *seller industry*. To see why, consider the seller industries with the 10 highest and lowest average CVs, as shown in Table 3. The lowest-CV industries are renting of goods, real estate, electricity, and telecommunications; all of which involve payments at regular intervals by definition regardless of whether a firm is reliable or not.[8] In contrast, the highest-CV industries feature industries like museums and construction, in which there is plausibly low demand for reliability, as well as dairy and wood manufacturing, where perishability or unforeseen errors could make reliability difficult to achieve. Differences in CV across firms in different industries are thus difficult to interpret, even if both firms are supplying the same buyer.

To remove buyer- and industry-driven variation in the CV, and identify characteristics of reliable firms (sellers), we estimate the following link-level regression:

$$CV_{ij} = X_i\beta + \gamma_{s(i)j} + \epsilon_{ij} \tag{11}$$

where the CV is calculated for each seller i-buyer j link in 2017, X are seller characteristics such as firm size, export status, and the share of sales to MNCs or exporters; $\gamma_{s(i)j}$ is a seller industry-buyer ID fixed effect, and ϵ is an error term. By including $\gamma_{s(i)j}$ we compare reliability across sellers in the same industry servicing the same buyer, analogous to Fig. 1. A negative β suggests that high-X_i firms have lower CV and are thus *more reliable*.

To understand whether reliable firms sort to each other or respond differentially to supply chain shocks, we also construct a measure of (residual) unreliability for each seller i via a fixed effects design. We estimate a specification analogous to (11) but replace βX_i with fixed effects α_i for each supplier:

$$CV_{ij} = \alpha_i + \gamma_{s(i)j} + \epsilon_{ij} \tag{12}$$

To ensure the estimated α_i are comparable, we subset to the largest connected set of suppliers and buyers, which comprises over 99% of links (see Abowd et al., 1999).[9]

Note that the CV should be interpreted as a measure of *output unreliability* for a particular *seller*. This is because, despite running the regression at the link level, we either correlate the CV with supplier-level variables or explicitly compute a supplier-level fixed effect after removing buyer-side variation. Thus, through the lens of our Section 2 model, high α_i in Eq. (12) means low r_y, and $\beta < 0$ in Eq. (11) implies that firms with high characteristic X are more reliable (choose higher r_y).

3.4. Data requirements

The reliability metric in Eq. (10) is fairly general, in that it does not require outside market- or industry-specific data. However, it does require time-stamped transaction data that is automatically recorded at the *relationship* (buyer–supplier) level. Most VAT systems fail this requirement because firms file a monthly return with all transactions, so the time between consecutive transactions is unknown; the Rwandan system we study is unique because EBMs are used at the time of sale to produce physical records for customers.

Several alternative datasets are more promising, though each is limited in coverage. In the international trade context, bill of lading data exist with precise arrival or departure dates (Liu et al., 2025). For domestic trade, many countries including India and Colombia collect freight records from trucks for tax purposes (Garg et al., 2023; Allen et al., 2024). And in many countries, payments are automatically recorded by credit card providers (Einav et al., 2021) and informal

[8] For instance, telecoms reliability is captured by how often a signal remains strong, not whether it is capable of charging a monthly phone bill.

[9] Bernard et al. (2019) show using Belgian VAT data that supplier and buyer fixed effects can be estimated cross-sectionally because, while employees each have one firm, firms have many trading partners and thus we do not rely on movers for identification.

Table 4
Larger firms supply more reliably.

Dependent Variables:	CV of days between consecutive transactions					CV (wks)	CV (qtr avg)
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variables							
Log interfirm sales	−0.013***	−0.022***	−0.021***	−0.013***	−0.014**	−0.013***	−0.006***
	(0.001)	(0.002)	(0.002)	(0.003)	(0.006)	(0.003)	(0.002)
Log other sales control		Yes	Yes	Yes	Yes	Yes	Yes
Min no. of trans.	10	10	10	10	20	10	10
Fixed-effects							
Buyer ID			Yes				
Seller Ind x Buyer ID				Yes	Yes	Yes	Yes
Fit statistics							
Observations	48,207	48,204	48,204	48,204	20,782	48,199	48,204
R^2	0.00310	0.00435	0.27276	0.73818	0.79194	0.71049	0.74907
Within R^2			0.00368	0.00142	0.00139	0.00344	0.00060

Clustered (Seller Ind x Buyer ID) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*
Notes: Table shows regressions at the link (seller-buyer) level of the CV (calculated separately by link) on the log of total interfirm sales of the seller (across all buyers), the log of other sales (across all buyers), and appropriate fixed effects. In columns 1-5, CV is the coefficient of variation of the number of days between consecutive transactions and in column 6 weeks are used in place of days, both for over all transactions in 2017. In column 7 the CV is computed separately by quarter in 2017and then averaged within a link to obtain the dependent variable. All data are from the Rwanda Revenue Authority (RRA).

digital payment providers (Houeix, 2025) in the process of regular operations.

4. Cross-sectional evidence for complementarities in reliability

4.1. Correlates of reliability

Having constructed our reliability metric for each formal firm in Rwanda, we now show that patterns of seller reliability across firms suggest a role for reliability in firm performance.

We first document that large firms are more reliable, which is consistent with Prediction 1 of the theoretical framework and suggests that the 'reliability premium' is positive. As shown in Column 1 of Table 4, without any controls a log point increase in total interfirm sales is associated with a decrease of 0.013 in the CV. However, Column 1 understates the magnitude of the reliability - formal sales coefficient due to a missing data problem: sales to firms and to households are positively correlated, and households (who do not have downstream buyers) are unlikely to pay the same premium for reliability as firms. Thus in Column 2 we add a control for log other sales. Conditional on any sales to consumers, a firm with an additional log point of interfirm sales (across all its links) has a 0.022 lower CV. Note that all results are qualitatively robust to not including this control.

The remainder of Table 4 shows that the negative CV-interfirm sales correlation is qualitatively robust to various specification choices. With buyer ID fixed effects or seller industry-buyer ID fixed effects (Columns 3 and 4) the coefficient is largely unchanged, implying that larger firms do not solely appear more reliable because they face differential demand or are in industries with low CV. Results are also unchanged when we subset to relationships with at least 20 transactions in 2017 (Column 5), consistent with the idea that reliability should not depend on how frequently transactions occur. Finally, if we define the CV using weeks between transactions to smooth over potential measurement error (where an order is filed 1–2 days early or late), or compute it separately for each quarter in 2017 and then average,[10] the result is qualitatively unchanged though smaller in magnitude.

The negative firm sales-CV correlation in Table 4 reflects two effects: that total sales will mechanically include a reliability premium, and that productive firms (who have more sales) choose higher reliability levels. To focus on the latter component, in Table 5 we examine how reliability varies with several other firm characteristics. By including seller industry-by-buyer ID fixed effects, we test if firms with a certain characteristic are more reliable than other sellers in the same industry that supply the *same* buyer (e.g., than other beer wholesalers supplying the same retailers). While the CV is itself unitless, to benchmark relative magnitudes, we compare all estimates to the coefficient on log interfirm sales, our preferred measure of firm size. As shown in Column 1 (Column 2), this coefficient is −0.006 without (−0.013 with) controlling for non-interfirm sales.

We begin with two common proxies for productivity: export status and being a multinational firm (Helpman et al., 2004). To do so, we classify firms into the mutually exclusive groups: multinational firms (MNCs) with some foreign ownership, local (non-MNC) firms that export, and domestic non-exporters.[11] Column 3 shows that local (non-MNC) Rwandan exporters are more reliable than non-exporters, with CV values that are on average 0.030 lower in magnitude. Because our regressions are estimated using domestic sales only, this estimate implies that firms that export also more reliably supply their domestic buyers, consistent with Prediction 2 from Section 2 that productive firms choose to be more reliable.[12] In contrast with the negative and significant estimate for export status, the coefficient on MNC status in Column 4 is statistically indistinguishable from zero. Limited variation may drive this null result: there are only 140 MNCs in our sample (2.3 percent of all formal sellers), restricting the set of cases in which an MNC and non-MNC in the same industry supply the same buyer. Consistent with this idea, the standard error of the MNC status coefficient in Column 4 is more than twice as large as for export status in Column 5; and after conditioning on the fixed effects, the share of remaining variation explained by MNC status (the within R^2) in Column 4 is two orders of magnitude lower than for any other covariate.

Next, we document that, along with choosing higher reliability levels, productive firms have more reliable *input suppliers*. To show this, we regress the CV on the share of each supplier's domestic (VAT-recorded) sales that are to an exporter (Column 5 of Table 5) or to an MNC (Column 6 of Table 5). Firms with a exporter sales share of 1 (who indirectly export all of their output) on average have a 0.266 lower CV than firms that only supply non-exporters; while firms that only supply

[10] This ensures we capture reliability in ongoing relationships, and not many-week gaps.

[11] All coefficients on export status are quantitatively similar if we pool MNC and non-MNC exporters, likely because only 18% of exporters have MNC status.

[12] Note that along with being more physically productive, exporters could also face higher reliability premia because they sell into international markets.

Table 5
Participants in interfirm trade are reliable and choose reliable suppliers.

Dependent Variable:	CV of days between consecutive transactions						
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variables							
Log interfirm sales	−0.006***	−0.013***					−0.007**
	(0.002)	(0.003)					(0.003)
Exporter			−0.034***				−0.035***
			(0.007)				(0.008)
MNC				0.012			0.0002
				(0.015)			(0.016)
Share of sales to exporters					−0.266***		−0.138**
					(0.050)		(0.060)
Share of sales to MNCs						−0.449***	−0.275**
						(0.120)	(0.129)
Log other sales control		Yes					Yes
Min no. of trans.	10	10	10	10	10	10	10
Fixed-effects							
Seller Ind x Buyer ID	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics							
Observations	48,207	48,204	48,207	48,207	48,207	48,207	48,204
R^2	0.73796	0.73818	0.73817	0.73783	0.73830	0.73816	0.73884
Within R^2	0.00054	0.00142	0.00134	4.24×10^{-5}	0.00185	0.00131	0.00392

Clustered (Buyer ID) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*
Notes: Table shows regressions at the link (seller-buyer) level of the CV (calculated separately by link) on covariates. CV is the coefficient of variation of the number of days between consecutive transactions. *Interfirm sales* is total recorded sales across all formal buyers and *Other sales* are sales to informal buyers or households. MNC is an indicator for foreign or joint ownership. *Exporter* takes value 1 if the seller is a domestic (non-MNC) Rwandan firm that records an export transaction. The share of sales to MNCs and to exporters are computed by summing sales to MNCs and to exporters, respectively, and dividing by total sales (interfirm + other). Sample consists of all transactions in 2017. All data are from the Rwanda Revenue Authority (RRA).

MNCs have a 0.449 lower CV than firms that only supply domestic Rwandan firms. In other words, both coefficients are statistically and economically significant, with larger effects on reliability of supplying MNCs.

These results suggest that productive firms and especially MNCs have higher returns to reliable supply, consistent with Prediction 3 of the model and therefore with the idea that reliable supply is necessary for reliable output. The coefficients on sales share to exporters and MNCs, as well as the coefficient on a seller's own export status, remain negative and significant even when we control for all covariates in Column 7.

In summary, cross-sectional patterns support a model in which investments in reliability increase sales, have higher returns for ex-ante productive firms, and require both internal production changes *and* reliable upstream suppliers.

4.2. Nonparametric evidence for correlated reliability

In the past section we used correlations with sales and exporting to suggest that reliable production mechanically depends on reliable supply. If this is true, and there are a limited number of reliable suppliers, then reliable firms should trade with each other. This correlation might arise mechanically, or because reliable firms may endogenously choose reliable suppliers.

To test for sorting, we estimate each supplier's residual unreliability by estimating Eq. (11) with a fixed effect α_i for each supplier in place of supplier characteristics. Then, for each supplier i, we compute the average residual unreliability across firms that directly supply i:

$$\bar{\alpha}_i^{upstream} = \mathbb{E}[\alpha_k | l_{ki} = 1]$$

as well as across firms that directly buy from i:

$$\bar{\alpha}_i^{downstream} = \mathbb{E}[\alpha_m | l_{im} = 1]$$

where $l_{ab} = 1$ if a supplies b. If reliable firms have reliable suppliers (buyers), then $\bar{\alpha}_i^{upstream}$ ($\bar{\alpha}_i^{downstream}$) will increase in α_i.

In practice, reliable firms do trade with each other. Fig. 2 plots linear regressions of $\bar{\alpha}_i^{upstream}$ and $\bar{\alpha}_i^{downstream}$ against α_i, in blue and red respectively, as well as binned scatterplots for 20 bins of α_i. Both lines slope upward; the correlation coefficient is about 0.20 with $\bar{\alpha}_i^{upstream}$ and 0.12 with $\bar{\alpha}_i^{downstream}$. Note that these correlations are *not* mechanical since α_i is always measured for a firm's behavior as a seller. Fig. 2 thus suggests that reliability is complementary across levels of a production chain, as in the Leontief production function we laid out in Eq. (1), or in an O-Ring-style model (Kremer 1993). The upshot is that, even if a firm is not an MNC or exporter, its ability to increase sales by selling reliably will depend on the presence of reliable input supply further up in the production chain – for instance, at customs offices or from the electricity grid.

5. Differential responses to a common supply chain shock

We next examine whether unreliable and reliable firms, as defined by our CV measure, respond differentially to a common shock to input supplies. The shock is the 2017 Kenya election, which due to a cessation of economic activity temporarily reduced imports to Rwanda from its second-largest trading partner. We find that, among firms equally exposed to the disruption, reliable firms with below-median CV saw smaller drops in *total* imports immediately after the election. In other words, the firms that we identified as reliable in the cross-section are the exact same firms that smooth over the election-induced import disruption. This natural experiment thus reveals another characteristic associated with reliable sellers: a secure supply of their own inputs.

5.1. Shock and identification: the 2017 Kenya presidential election

Kenya is Rwanda's second-largest source of imports after China, and since Rwanda is landlocked almost all goods must travel through western Kenya and Uganda by truck. However, every five years, Kenya holds an election that is typically accompanied by a temporary economic slowdown and threat of violence. The disputed 2007–08 election led to substantial rioting (see Macchiavello and Morjaria, 2015) and damage to goods. During and after subsequent elections, including the one we study, most traders have chosen not to make drives through Kenya.

To verify that imports to Rwanda declined around the August 8, 2017 election, we subset to firms importing from Kenya in a pre-period

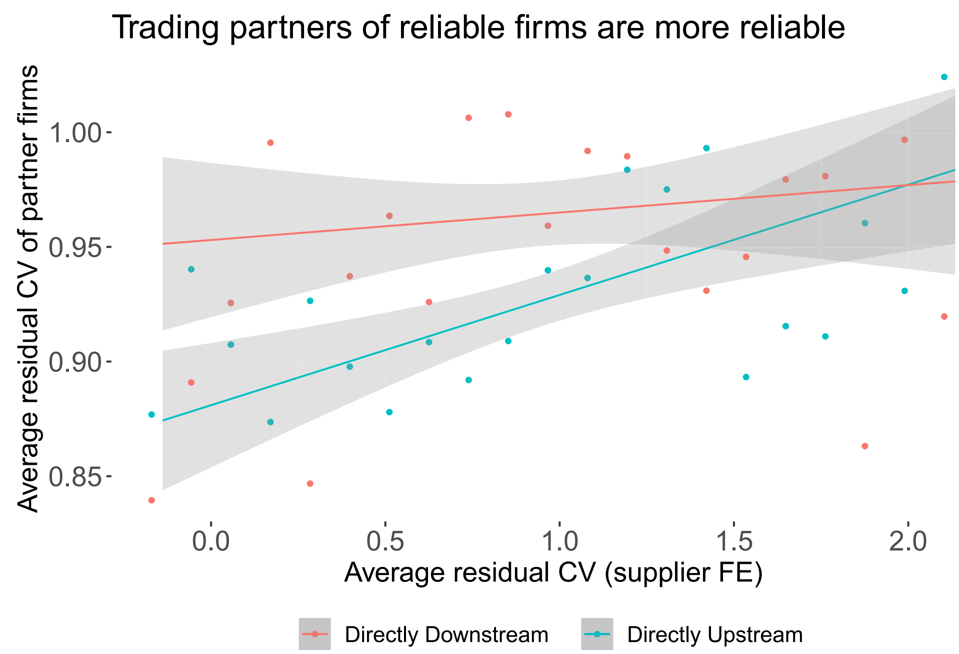

Fig. 2. Reliable firms sell to and buy from each other.
Notes: Figure shows each firm's residual unreliability $\hat{\alpha}_i$ on the horizontal axis and the average residual unreliability of that firm's direct suppliers (in blue) or direct customers (in red) on the vertical axis. Both binned scatterplots (for 20 bins of $\hat{\alpha}_i$) and the linear regression function are shown. For visibility and to remove outlier CV values we trim at the 2.5^{th} and 97.5^{th} percentiles of estimated $\hat{\alpha}_i$ before plotting. $\hat{\alpha}_i$ are supplier ID fixed effects obtained from a regression of CV on supplier ID fixed effects and supplier industry-buyer ID fixed effects, estimated at the link (supplier-buyer) level using all transactions in 2017 between the largest connected set of formal firms in the sample. CV is the coefficient of variation of the number of days between consecutive transactions. All data are from the Rwanda Revenue Authority (RRA).

(January to June 2017) and regress the inverse hyperbolic sine (IHS) of daily firm-level imports from Kenya on firm fixed effects and weeks-since-election dummies. Estimated coefficients on weeks-since-election dummies are plotted in Fig. 3. On average, firms with prior trade links with Kenya reduce imports by at least 50% during the election period.

To obtain variation across importing firms in exposure to the shock, we measure the share of each firm's pre-period imports that are from Kenya:

$$KenyaShare_i = \frac{\sum_{t=1/1/2017}^{6/30/2017} \sum_{c=Kenya} Imports_{ict}}{\sum_{t=1/1/2017}^{6/30/2017} \sum_{\forall c} Imports_{ict}} \tag{13}$$

and then estimate a differences-in-differences regression:

$$sinh^{-1}(Imports_{it}) = \delta_{id(t)} + \gamma_{s(i)t} + \sum_{\tau \neq -1} \beta_\tau KenyaShare_i \mathbb{1}[Week_t - Week^0 = \tau] + \epsilon_{it} \tag{14}$$

in which i indexes firm, c indexes the exporting country, t indexes the exact date, d is a day of the week, and s is an industry. $Imports_{it} = \sum_c Imports_{ict}$ are total imports from all countries made by the firm on date t. Due to there being zero imports on many firm-dates, we apply an inverse hyperbolic sine transformation in our main specification rather than a log transformation, and also show that a version of Eq. (14) with an indicator for nonzero importing as the outcome variable yields similar results. δ is a firm-weekday fixed effect used both for identification and to reduce noise from day-of-week effects, γ is an industry-date fixed effect, $Week_t$ is the calendar week in which date t falls, $Week^0$ is the week of the Kenya election, and ϵ is an error term. The coefficient of interest is β_τ, which is the additional change in imports, relative to 1 week before the election, due to a 1 percentage point (p.p.) increase in the pre-period share of imports from Kenya. A negative β_τ implies that the shock has a negative effect on the imports of exposed firms.

To test for heterogeneous by each firm's baseline reliability level, we first estimate Eq. (12) on all transactions observed in the base period (January 1 to June 30, 2017) to obtain the residual CV $\hat{\alpha}_i$ of each supplier. We then subset to suppliers with nonzero base-period imports (the sample for Eq. (14)), compute the median residual CV in this sample, and then re-estimate Eq. (14) separately for firms with above- and below-median values of $\hat{\alpha}_i$. If the CV is informative of a firm's unreliability as a seller, and reliable sellers choose higher input reliability levels, then firms with above-median $\hat{\alpha}_i$ (i.e. unreliable firms) should see larger declines in output due to the Kenya election shock. We thus test, both visually and formally, for differences in the time path of β_τ in these two samples.

Before moving to results we briefly discuss identification and our choice of sample. While we show two sets of difference-in-difference coefficients from separate subsamples, our approach is econometrically equivalent (in terms of point estimates) to estimating a version of (14) in which all covariates, including the time fixed effects, are interacted with an indicator for above-median CV. In other words, high- and low-CV firms are allowed separate non-parallel time trends, and the key identification assumption is that *within* each subsample, firms with high and low Kenya import exposure shares exhibit parallel trends in potential outcomes. The full sample comprises all Rwandan firms that had an import from any country in the pre-period; this is about a third of all Rwandan firms. While we focus on effects in the weeks around the election, the sample includes imports from January to November 2017 to reduce noise in the firm fixed effects.

Since our dataset is a firm-by-date panel and firm-level shares determine exposure to treatment, we cluster by firm in all regressions.

5.2. *Event study effects*

We now show that the average effect of the Kenya election disruption on imports is larger for less reliable firms with above-median CV than for reliable ones with below-median CV. Panel A of Fig. 4 plots

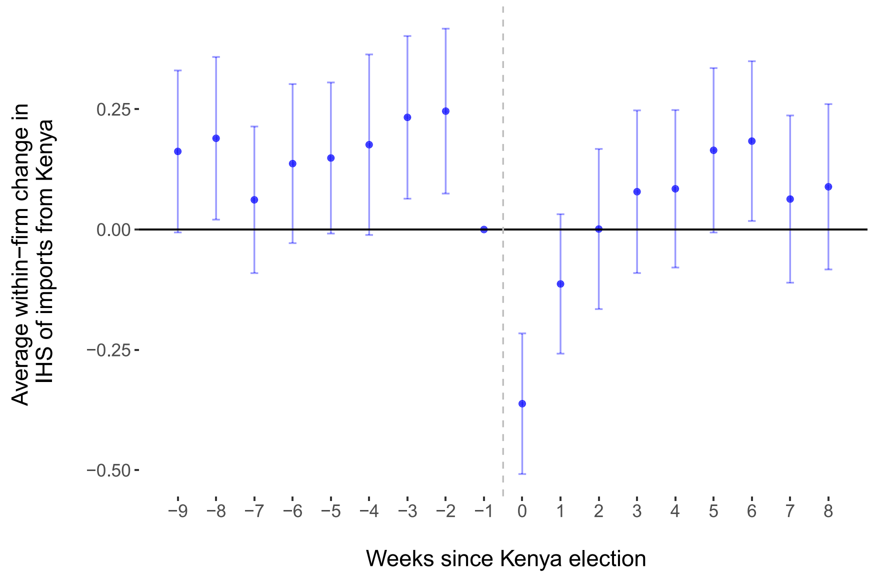

Fig. 3. Rwandan imports to Kenya decline in weeks around 2017 presidential election.
Notes: Figure shows estimated weeks-since-election coefficients from a regression of the inverse hyperbolic sine (IHS) of total imports on weeks-since-election dummies and firm ID-day of week fixed effects. Regression is estimated at the firm-by-date level using a panel of import transactions from January 1 to November 30, 2017. The gray dotted line separates the week ending on August 7, 2017 (the day before the election) from the week beginning on August 8, 2017 (the election date). The sample consists of all Rwandan firms who had at least one import *from Kenya* between January 1 and July 31, 2017. All data are from the Rwanda Revenue Authority.

estimate of β_τ for 9 weeks before to 6 weeks after the election. In the discussion that follows, we convert from effects on IHS to percent changes where appropriate,[13] and all effects should be interpreted as for a firm for which 100% of pre-period imports coming from Kenya.

For a high-CV firm (in blue) with full treatment exposure, we estimate that total imports are 80.2% lower during election week and 53.7% lower one week afterwards, both relative to the week before the election. With 95% confidence, we can rule out effects of smaller than 53.2% in the week of the election. These estimates suggest that high-CV firms were unable to either move goods through Kenya or substitute to alternative sources of imports during election week – in other words, they were unable to maintain reliable input supply during the disruption. In contrast, for a low-CV firm (in red), imports fall by 26.6% during election week and fully recover one week after the election, and we cannot reject zero effect throughout the shock period. When we estimate the fully interacted specification, the difference in treatment effects between the high- and low-CV firm samples is statistically significant in both weeks.

These results suggest that low-CV firms more strongly smooth over a shock to their input supply. To examine whether the differential effect is driven by lost transaction *dates* rather than lower value per shipment, we replace $sinh^{-1}(Imports)$ with an indicator for whether any import was observed on day t and then re-estimate Eq. (14). As shown in Panel B of Fig. 4, the time-path of effects mirrors the main specification. In the week after the election, the probability of a firm importing on any given day falls by 9.5 p.p. for high-CV firms (in blue) and 1.7 p.p for low-CV firms (in red); since the mean probability of importing on a day across firm-dates is 7 p.p., these effects imply that a fully-exposed high-CV firm essentially ceases importing, while an equally-exposed high-CV continues to receive goods from abroad. These differential extensive-margin responses imply that low-CV firms choose to transact on the same number of dates as before the shock, while high-CV firm temporarily shut down along this margin in response.

For two reasons, we interpret the effects of the Kenya election shock as a 'stress test' that reveals the superior quality of low-CV firms' existing supply chain. First, as discussed above, differential effects are largely due to the extensive margin, consistent with the idea that reliable firms need to stick to schedules and therefore need inputs to arrive on time regardless of conditions. Second, effects are concentrated in the election week and import volume fully recovers within three weeks of the election. Since importing to Rwanda involves manufacturing time *and* a trip with at least two border crossings, it is therefore unlikely that the short-run effects reflect post-shock changes in supply chains – for instance, seeking out new suppliers or placing extra orders with existing ones.

We briefly discuss a few additional results. First, to understand which goods drive the results in Fig. 4, in Appendix Figure A4 we separately estimate effects of the Kenya election shock on imports of above- and below-median values of upstreamness. We can do so because, unlike in the Rwandan VAT data, in the customs data we observe precise HS6-level product codes. We classify each product code as having above- or below-median values of upstreamness, as in Antràs and Chor (2018). Our motivation is that products with high measured upstreamness are mostly used as intermediates in production (rather than directly by households), and thus map most closely to Prediction 4 of our theoretical framework. We find that the differential effect between reliable and unreliable firms of the Kenya shock is larger in magnitude – and that negative effects on unreliable firms' inputs remain significant for an additional week – in the upstream sample. However, due to power reasons, we cannot formally reject that the differential effects by reliability are equal in both samples.

Second, in Appendix Figure A5 we separately estimate effects of the Kenya election shock on imports of customized and homogeneous inputs. We classify each product code as customized or homogeneous using the Rauch (1999) measure of input specificity, with a crosswalk to the HS6 level developed by Liao et al. (2020). While unreliable firms are differentially affected by the shock in both samples, the negative

[13] Since the IHS approximates logs we use % change $= e^{\beta_\tau} - 1$ when the outcomes is $sinh^{-1}(Imports)$.

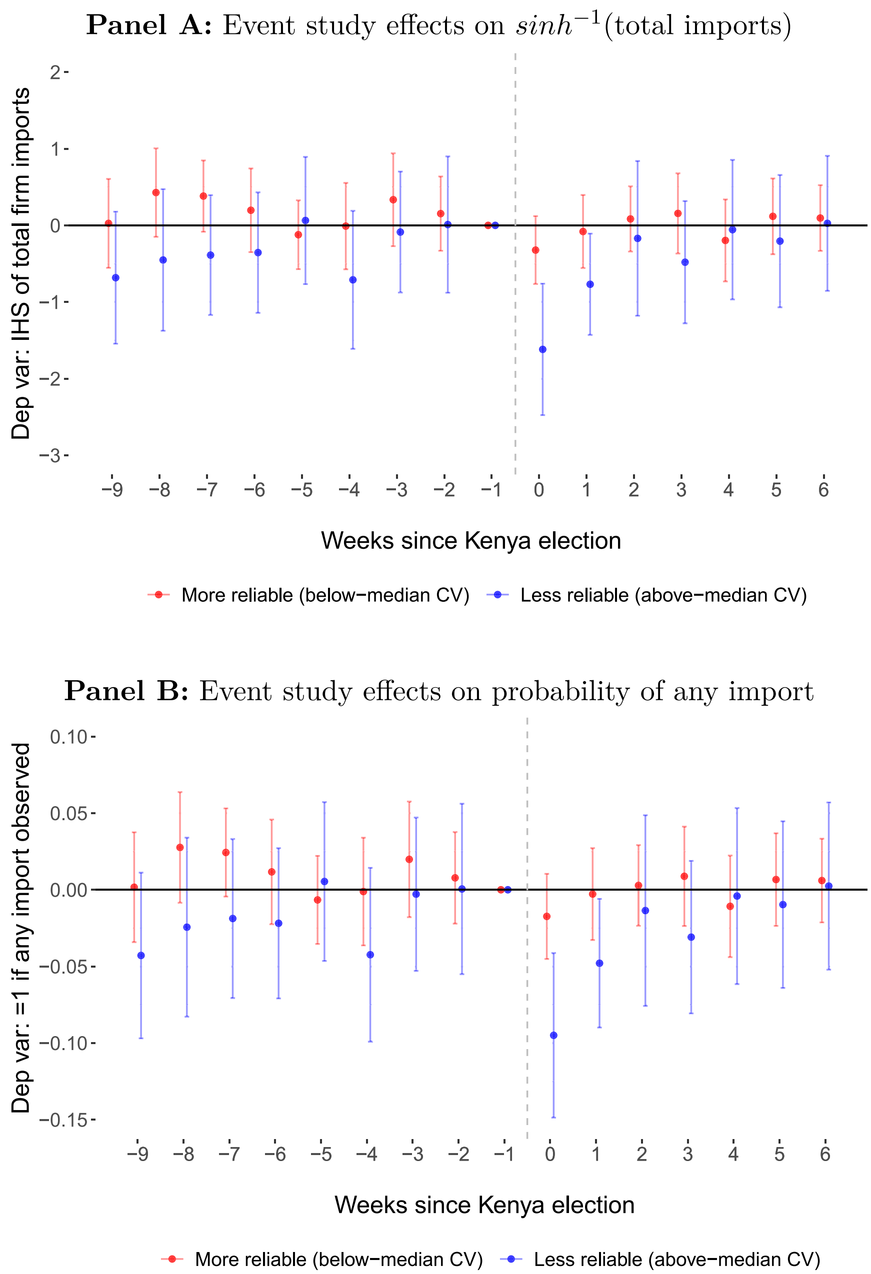

Fig. 4. Import declines are larger for unreliable (above-median CV) firms after Kenya election.
Notes: Figure shows estimated coefficients on interactions between week-since-election dummies and the share of a firm's pre-period imports that are from Kenya. All regressions are at the firm-by-date level. The dependent variable is either the inverse hyperbolic sine (IHS) of total imports (Panel A) or an indicator for whether any import transaction was observed for the firm-date cell. The regression includes the following covariates: (1) interactions between week-since-election dummies and the pre-period Kenya import share, where the election is August 8, 2017 and the pre-period is January 1 to July 31, 2017; (2) firm ID-day of week fixed effects; and (3) industry-date fixed effects. The full sample consists of all dates from January 1 to November 30, 2017 and all firms that made at least one import in the pre-period. Regressions are estimated separately for firms with below-median (in red) and above-median (in blue) residual unreliability $\hat{\alpha}_i$ for each supplier i. $\hat{\alpha}_i$ are the supplier fixed effects obtained from a first-stage link-level regression of the CV on supplier fixed effects and seller industry-buyer ID fixed effects. CV is the coefficient of variation of the number of days between consecutive transactions. The CV is calculated and the first-stage regression are estimated using all transactions between formal firms in the pre-period. All data are from the Rwanda Revenue Authority.

effects on unreliable firms' inputs is larger – and persists beyond the week of the election – in the customized input sample. This result is consistent with other evidence that disruptions to specific inputs are more likely to propagate downstream (Barrot and Sauvagnat, 2016); our contribution is to show that more unreliable buyers (according to our metric) face such propagation.

Third, in Appendix Figure A6 we separately show effects on imports by broad sector (agriculture vs. manufactures). The reduction in imports of unreliable Rwandan firms is driven entirely by imports of manufactures; agricultural imports are not affected. Since agricultural goods are typically less-customized, these results are consistent with Appendix Figure A5. In line with this, in Appendix Figure A3, we find that supplier and buyer reliability is highly correlated for manufacturing goods and services, while uncorrelated for agricultural goods.

Fourth, to test for substitution to other origin countries, in Panel A of Appendix Figure A7 we estimate effects on imports from all countries other than Kenya. Effects are not significant, and are close to zero in magnitude, so we cannot reject the null hypothesis of no substitution. This result implies that reliable firms maintain input supply in part by keeping imports flowing on existing routes, rather than identifying entirely new suppliers outside Kenya.

Fifth, to test if reliable firms also better maintain their *sales* volume during the shock, in Panel B of Appendix Figure A7 we estimate effects of the import shock on domestic sales. While *unreliable* exposed firms indeed reduce their sales (and significantly so two weeks after the shock), reliable firms do not sell any less. Note that due to a lack of power, we cannot reject that the event study coefficients in the high- and low-reliability regressions are the same. Nevertheless, the time-path of the sales reduction is consistent with the sequential nature of trade: missing inputs in the week of the shock leading to forgone domestic sales two weeks later.

The upshot is that firms which *sell* with high day-to-day reliably, as measured by our CV metric, continue to *buy* inputs during a disruption. The heterogeneous shock response is consistent with Prediction 4 in our theoretical framework: productive sellers maintain more reliable input supply in the short-run because the complementarity with management gives them higher returns to input reliability. Unlike in Fig. 2, where we proxied for input reliability with the average output reliability of a firm's upstream suppliers, here we learn input reliability directly from how firms respond in a situation where inputs are hard to obtain. We thus have additional suggestive evidence that, as suggested by a model with complementarities, reliable firms would particularly value improvements in input reliability – and thus benefit from improvements in customs processing or political stability that enable a steady and predictable flow of goods.

6. New suppliers to MNCs become more reliable

Along with sorting on the basis of reliability, there is a second potential reason for the high reliability of suppliers to MNCs and exporters: MNCs might causally alter supplier operations. For instance, reliability might involve knowledge of a particular production process which MNCs teach (unpriced) to their suppliers. Alternatively, if MNCs are willing to pay for reliability, in order to start supplying an MNC suppliers might pay the fixed costs of supplying all customers on a predictable schedule.

To examine potential causal effects of supply-chain linkages on reliability, we use an event-study framework. We split our 2015–2017 transaction-level dataset into 12 quarterly (3-month) intervals and compute the CV of time between shipments separately for each supplier, buyer, and *quarter*. This allows us to examine if, within its pre-existing set of relationships, a Rwandan supplier becomes more reliable in the months after supplying an MNC. The main difference-in-difference regression is:

$$CVDiff_{ijt} = \beta SuppliesMNC_{it} + \gamma_{ij} + \delta_{jt} + \epsilon_{ijt} \tag{15}$$

where $CVDiff$ is now computed separately for each quarter t. γ is a buyer x supplier fixed effect which implies all analysis is of changes within existing relationships. δ is a buyer x quarter fixed effect that allows us to flexibly control for buyer-specific demand shocks, so that changes in buyer ordering patterns are not incorrectly attributed to supplier behavior. Finally $SuppliesMNC$ is an indicator variable for whether supplier i, in any quarter from the beginning of our sample until time t, has a recorded sale to an MNC. Note that many firms in our sample never or always sell to MNCs; they are used to identify γ and δ but not the treatment effect.

To estimate Eq. (15), we require many events in which a multinational buyer adds a new supplier. Our data allow this because, while only 2.3 percent of our sample firms are MNCs as shown in Table 1, these MNCs are large and exhibit substantial churn in their supplier bases. Specifically, of the 140 MNCs in our sample, 124 add at least one new domestic supplier during our analysis period, leading to 656 new relationships, or an average of 5.3 new relationships per MNC. After restricting to events where a supplier first sells to an MNC, and to suppliers that continuously operate before and after the event (to ensure a sufficiently balanced panel), we are left with 181 suppliers that first sell to an MNC during the study period. While the anonymized nature of the data means we could not ask MNCs about specific sourcing events, conversations with the central Rwandan development agency (the Rwanda Development Board or RDB), suggested that there were no centralized government efforts to connect growing suppliers to MNCs during our sample period. Most relationships were instead initiated due to MNC buyers' business needs, often because existing suppliers were considered inadequate. In absence of treatment, it is unlikely that these MNC requirements are correlated with trends in our main outcome (supplier reliability in *pre-existing* relationships with non-MNC buyers).

Table 6 gives the DID average treatment effect on our unreliability metric, which is negative and significant in all specifications. This implies that new Rwandan suppliers to MNCs transact at more regular intervals within their existing relationships, net of any customer-(demand-)driven trends in reliability. If interpreted causally, these estimates suggest that demand for reliability by MNCs propagates upstream to chosen suppliers and downstream to that supplier's buyers.

We conduct several robustness checks. First, as shown in Columns 4 and 5, when we limit to relationships with greater transaction frequency we find larger and more strongly significant negative effects. This is intuitive, since reliability may be more important for firms that deliver frequently (in this case, approximately weekly rather than monthly). Second, as shown in Column 6, effects are robust to a sample of relationships in which both suppliers and buyers are not MNCs. This result is particularly suggestive of downstream propagation, since the non-MNC buyers which experience higher reliability could not have been responsible for the initial link.

We also estimate an event-study analogue to (3). Event-study coefficients are traced out for pairs of quarters, with 1–2 quarters before first supplying an MNC as the base period. Results for the full and domestic-only sample, analogous to Columns 2 and 6 of Table 6, are shown in Fig. 5. While pre-trends are close to zero and not significant, unreliability (the CV) declines rapidly in the first year after treatment, adding credibility to the results in Table 6.

Finally, note that analogous event-study regressions for the effect of supplying an *exporter* yield null effects. Thus our results suggest that MNCs have an effect on suppliers beyond that of (indirect) participation in international trade.

7. Conclusion

In this paper, we generate a measure of reliability for the universe of Rwandan formal firms and describe the characteristics of reliable firms. We use data on the dates of all VAT-eligible transactions to construct a metric of reliability, and use the network structure of the data to remove demand-driven transaction patterns and isolate the average reliability of each *seller*. Evidence from the cross-section of all Rwandan formal firms and from importers' responses to a supply chain disruption are consistent with an O-Ring-style model of reliability, that is, where the production of reliable output is complementary in both a firm's type and the reliability of that firm's suppliers. We also find that supplying an MNC is associated with increases in seller reliability even when servicing non-MNC buyers.

Our results imply that governments can use two complementary policy tools to improve supply chain reliability. The first are *reliability*

Table 6
Domestic firms become more reliable after selling to MNCs.

Dependent Variable:	CV of Time Between Shipments					
	All			> 8 trans.	> 10 trans.	Non-MNCs
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Has Supplied MNC	−0.0124***	−0.0135***	−0.0091*	−0.0128**	−0.0145*	−0.0114**
	(0.0039)	(0.0048)	(0.0047)	(0.0061)	(0.0086)	(0.0050)
Fixed-effects						
Calendar Month	Yes					
Supplier x Buyer	Yes	Yes	Yes	Yes	Yes	Yes
Buyer x Calendar Month		Yes	Yes	Yes	Yes	Yes
Transaction Count			Yes			
Fit statistics						
Observations	217,054	217,054	217,054	138,915	84,261	186,192
R^2	0.54460	0.68201	0.69402	0.71811	0.74981	0.68556
Within R^2	8.35×10^{-5}	9.61×10^{-5}	4.53×10^{-5}	9.37×10^{-5}	0.00013	7.55×10^{-5}

One-way (Supplier x Buyer) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*
Notes: The sample contains all buyer-seller-quarter cells from Q1 2015 to Q4 2017 with at least 5 transactions. In columns 4 and 5 we subset to cells with at least 8 and 10 transactions, respectively. In column 6 we subset to non-MNC buyers and sellers.

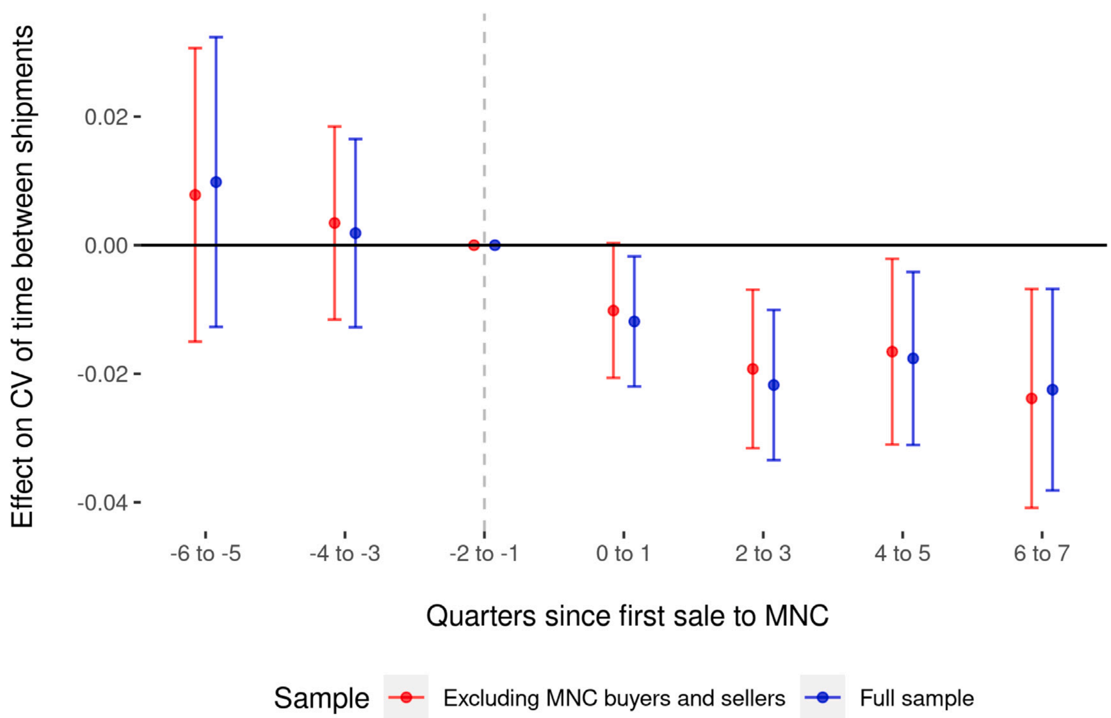

Fig. 5. New Rwandan suppliers to MNCs become more reliable in their existing relationships.
Notes: The outcome variable is the coefficient of variation of the time between consecutive shipments within a seller-buyer-quarter. The sample contains all seller-buyer-quarter cells in Rwanda from Q1 2015 to Q4 2017 with at least 5 transactions. Each coefficient represents two consecutive quarters following the base period (-2 to -1 quarters before the first recorded link to an MNC, as indicated by the gray line). The event-study specification corresponds to Column 3 (all relationships) and Column 6 (non-MNC buyers and sellers) of Table 5.

supply policies. Because improvements in reliability propagate downstream, interventions that target input reliability directly – such as express customs processing for repeat importers, 24/7 electricity for key manufacturers, and reductions in congestion on existing routes – may have high aggregate effects. To maximize gains, these interventions should systematically target direct and indirect suppliers of the most-productive downstream firms, such as MNCs and exporters. In contexts where total input supply is limited (for instance, by the number of customs officers or by total electricity generation capacity), our results support a policy of favoritism in which governments identify a set of reliability-sensitive firms to be prioritized for regular input supply. Such policies are already widespread in poor countries: to attract multinationals and their suppliers to a location, governments often designate an industrial park or special economic zone that will be prioritized for key inputs, especially electricity (Garg, 2024).

The second set of policies are *reliability demand* initiatives. In particular, subsidizing the entry of productive downstream firms, especially exporters and multinational firms, may also have magnified aggregate gains. This is because MNC entry can lead to chains of 'reliability upgrading' in poor countries, in which new suppliers to the MNC, followed by the trading partners of those suppliers, are induced to improve reliability as MNCs demand it. A promising area for future research is thus to quantify the effects on reliability (both directly and via spillovers through the supply chain) of the aforementioned demand-and supply-side policies; while many papers evaluate how these papers shape other outcomes, there has been little effect to quantify reliability effects until now due to a lack of data.

A final implication is that transaction data now automatically collected by many governments contains valuable information on the day-to-day behavior of individual firms. Many other features of trade, beyond the reliability metric that we construct, may now be easily observed at scale by poor-country policymakers. For instance, certain transaction patterns during non-crisis times may predict continued interaction during a shock, higher returns to trade credit, better management quality, or imminent firm growth. A promising avenue for future research is thus to investigate these transaction data and their potential to yield insights for a broader range of development questions.

CRediT authorship contribution statement

Vishan Gandhi Nigam: Software, Validation, Writing – review & editing, Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Resources, Visualization, Writing – original draft. **Brandon Joel Tan:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jdeveco.2025.103611.

References

Abowd, John M., Kramarz, Francis, Margolis, David N., 1999. High wage workers and high wage firms. Econometrica 67 (2), 251–333.

Alfaro-Urena, Alonso, Manelici, Isabela, Vasquez, Jose P., 2020. The effects of joining multinational supply chains: New evidence from firm-to-firm linkages. Available At SSRN 3376129.

Allcott, Hunt, Collard-Wexler, Allan, O'Connell, Stephen D., 2016. How do electricity shortages affect industry? Evidence from India. Am. Econ. Rev. 106 (3), 587–624.

Allen, Treb, Atkin, David, Cleves, Santiago Cantillo, Hernandez, Carlos Eduardo, 2024. The traveling trucker problem. In: AEA Papers and Proceedings, vol. 114, pp. 334–339.

Antràs, Pol, Chor, Davin, 2018. On the measurement of upstreamness and downstreamness in global value chains.

Asker, John, Collard-Wexler, Allan, De Loecker, Jan, 2014. Dynamic inputs and resource (mis) allocation. J. Political Econ. 122 (5), 1013–1063.

Atkin, David, Khandelwal, Amit K., Osman, Adam, 2017. Exporting and firm performance: Evidence from a randomized experiment. Q. J. Econ. 132 (2), 551–615.

Banerjee, Abhijit V., Duflo, Esther, 2000. Reputation effects and the limits of contracting: A study of the Indian software industry. Q. J. Econ. 115 (3), 989–1017.

Barrot, Jean-Noël, Sauvagnat, Julien, 2016. Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks *. Q. J. Econ. 131 (3), 1543–1592.

Bernard, Andrew B, Dhyne, Emmanuel, Magerman, Glenn, Manova, Kalina, Moxnes, Andreas, 2019. The origins of firm heterogeneity: A production network approach. http://dx.doi.org/10.3386/w25441, Working Paper Series 25441, National Bureau of Economic Research.

Bloom, Nicholas, 2009. The impact of uncertainty shocks. Econometrica 77 (3), 623–685.

Bloom, Nicholas, Van Reenen, John, 2007. Measuring and explaining management practices across firms and countries. Q. J. Econ. 122 (4), 1351–1408.

Boehm, Christoph E., Flaaen, Aaron, Pandalai-Nayar, Nitya, 2019. Input linkages and the transmission of shocks: Firm-level evidence from the 2011 tōhoku earthquake. Rev. Econ. Stat. 101 (1), 60–75.

Burgess, Robin, Greenstone, Michael, Ryan, Nicholas, Sudarshan, Anant, 2020. The consequences of treating electricity as a right. J. Econ. Perspect. 34 (1), 145–169.

Carvalho, Vasco M, Nirei, Makoto, Saito, Yukiko U, Tahbaz-Salehi, Alireza, 2021. Supply chain disruptions: Evidence from the great east japan earthquake. Q. J. Econ. 136 (2), 1255–1321.

Demir, Banu, Fieler, Ana Cecília, Xu, Daniel, Yang, Kelly Kaili, 2021. O-ring production networks. In: Working Paper Series 28433, National Bureau of Economic Research, http://dx.doi.org/10.3386/w28433, URL http://www.nber.org/papers/w28433.

Einav, Liran, Klenow, Peter J., Levin, Jonathan D., Murciano-Goroff, Raviv, 2021. Customers and retail growth. http://dx.doi.org/10.3386/w29561, Working Paper 29561, National Bureau of Economic Research.

Eissa, Nada, Zeitlin, Andrew, Karpe, Saahil, Murray, Sally, 2014. Incidence and impact of electronic billing machines for VAT in rwanda. Lond.: Int. Growth Cent..

Fieler, Ana Cecília, Eslava, Marcela, Xu, Daniel Yi, 2018. Trade, quality upgrading, and input linkages: Theory and evidence from Colombia. Am. Econ. Rev. 108 (1), 109–146.

Firth, John, 2017. I've been waiting on the railroad: The effects of congestion on firm production. Technical report, MIT Working Paper.

Fried, Stephie, Lagakos, David, 2020. Electricity and firm productivity: A general-equilibrium approach. Technical report, National Bureau of Economic Research.

Garg, Tishara, 2024. Can industrial policy overcome coordination failures? Theory and evidence. Job Market Paper, MIT Department of Economics, November 2024.

Garg, Shresth, Ghosh, Pulak, Tan, Brandon Joel, 2023. Within firm supply chains: Evidence from India. J. Int. Econ. 144, 103793.

Ghani, Tarek, Reed, Tristan, 2022. Relationships on the rocks: Contract evolution in a market for ice. Am. Econ. J.: Microeconomics 14 (1), 330–365.

Grossman, Gene M., Helpman, Elhanan, Lhuillier, Hugo, 2021. Supply chain resilience: Should policy promote diversification or reshoring?. http://dx.doi.org/10.3386/w29330, Working Paper 29330, National Bureau of Economic Research.

Hallak, Juan Carlos, Schott, Peter K., 2011. Estimating Cross-Country Differences in Product Quality*. Q. J. Econ. 126 (1), 417–474.

Helpman, Elhanan, Melitz, Marc J., Yeaple, Stephen R., 2004. Export versus FDI with heterogeneous firms. Am. Econ. Rev. 94 (1), 300–316.

Houeix, Deivy, 2025. Nationwide diffusion of technology within firms' social networks. Working paper, randomized experiment on digital payment adoption in Senegal.

Khandelwal, Amit, 2010. The Long and Short (of) Quality Ladders. Rev. Econ. Stud. 77 (4), 1450–1476.

Kremer, Michael, 1993. The O-ring theory of economic development. Q. J. Econ. 108 (3), 551–575.

Kremer, Michael, Lee, Jean, Robinson, Jonathan, Rostapshova, Olga, 2013. Behavioral biases and firm behavior: Evidence from Kenyan retail shops. Am. Econ. Rev. 103 (3), 362–368.

Kugler, Maurice, Verhoogen, Eric, 2011. Prices, Plant Size, and Product Quality. Rev. Econ. Stud. 79 (1), 307–339.

Liao, Steven, Kim, In Song, Miyano, Sayumi, Zhang, Hao, 2020. Concordance: Product concordance. R package version 2.0.0.

Liu, Ernest, 2019. Industrial Policies in Production Networks*. Q. J. Econ. 134 (4), 1883–1948, arXiv:https://academic.oup.com/qje/article-pdf/134/4/1883/32666171/qjz024.pdf.

Liu, Ernest, Liu, Yukun, Smirnyagin, Vladimir, Tsyvinski, Aleh, 2025. Supply chain disruptions, supplier capital, and financial constraints. Discussion Paper 2402R1. Cowles Foundation for Research in Economics, Yale University.

Macchiavello, Rocco, 2021. Relational contracts and development. Working Paper.

Macchiavello, Rocco, Morjaria, Ameet, 2015. The value of relationships: evidence from a supply shock to Kenyan rose exports. Am. Econ. Rev. 105 (9), 2911–2945.

Martin, Julien, Mejean, Isabelle, Parenti, Mathieu, 2020. Relationship stickiness, international trade, and economic uncertainty.

Mascagni, Giulia, Mukama, Denis, Santoro, Fabrizio, 2019. An analysis of discrepancies in taxpayers' VAT declarations in rwanda.

McMillan, John, Woodruff, Christopher, 1999. Interfirm relationships and informal credit in Vietnam. Q. J. Econ. 114 (4), 1285–1320.

Melitz, Marc J., 2003. The impact of trade on intra-industry reallocations and aggregate industry productivity. Econometrica 71 (6), 1695–1725.

Pomeranz, Dina, 2015. No taxation without information: Deterrence and self-enforcement in the value added tax. Am. Econ. Rev. 105 (8), 2539–2569.

Rauch, James E., 1999. Networks versus markets in international trade. J. Int. Econ. 48 (1), 7–35.

Schott, Peter K., 2004. Across-Product Versus Within-Product Specialization in International Trade*. Q. J. Econ. 119 (2), 647–678.

Verhoogen, Eric A., 2008. Trade, quality upgrading, and wage inequality in the mexican manufacturing sector. Q. J. Econ. 123 (2), 489–530.