Unreliable Firms: Evidence from Rwanda*

Vishan Gandhi Nigam (MIT)† Brandon Joel Tan (IMF)‡

May 14, 2022

[PRELIMINARY AND INCOMPLETE]

Abstract

This paper examines if differences in *reliability* – whether firms execute transactions on-schedule – can explain low firm output in developing countries. To measure day-to-day reliability, we leverage transaction timing data for the universe of Rwandan formal firms. Reliable firms have larger interfirm sales, export more, supply exporters and multinationals, transact with other reliable firms, and maintain input supply during a supply chain disruption. These firm-level associations suggest that reliable output requires reliable inputs, and therefore that improving the reliability of a minority of firms could generate magnified differences in aggregate output.

Keywords: Trade, Firms and Development, Technological Choice

JEL Codes: F14, O12, O14

*We thank Pol Antras, David Atkin, Vittorio Bassi, Lauren Bergquist, Dave Donaldson, Esther Duflo, Marc Melitz, Anna Twum, seminar participants at NEUDC 2021, and the students of MIT 14.192 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.

†MIT Department of Economics (vishan@mit.edu)

‡International Monetary Fund (btan2@imf.org)

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 describe 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. Intuitively, and consistent with work in the supply chain literature, to meet production schedules firms require both know-how in supply chain and logistics as well as a network of upstream suppliers and the (Baldwin and Freeman 2021). 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. Prior work on the 'second moment' of firm production either models economy-wide uncertainty[1] or implicitly defines a firm's reliability or resilience using delivery after a specific type of external shock[2]. In contrast, we

[1] For instance, Bloom (2009) structurally models the impact of an uncertainty shock on productivity, and Asker, Collard-Wexler and De Loecker (2014) document substantial volatility of productivity within industries over time. In contrast, our focus is on uncertainty in the supply curve of each individual firm, which differs substantially 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.

[2] While some of the microeconomic evidence deals with one-time events, such as civil wars (Macchiavello and Morjaria 2015), earthquakes (Carvalho et al. 2021, Boehm, Flaaen, and Pandalai-Nayar 2019) and weather shocks (Barrot and Sauvagnat 2016), several other papers examine industry-specific shifters in the reliability of output. See

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. 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).[3]

Since the CV of time between transactions is technically a measure of transaction *regularity*, 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, since we observe transactions rather than shipments, we assume that recorded dates correspond to exchanges of goods. These assumptions are difficult to verify in existing data, so we plan to conduct a firm survey in Rwanda to directly identify industries where transaction dates are informative of delivery and regular deliveries are desired, analogous to the approach in Ghani and Reed (2022).

Since we do not actually observe desired transaction dates, our third assumption is somewhat stronger: that schedules consist of deliveries at *evenly-spaced intervals*, allowing us to summarize

Banerjee and Duflo (2000) on cost overruns in the IT industry, Allcott, Collard-Wexler, and O'Connell (2016) on power outages, Firth (2017) on freight congestion, and Ghani and Reed (2022) on ice deliveries.

[3] While we know of no similar exploration of firm-level reliability, to make our objective clear we briefly discuss how our work relates to existing work on quality in trade. First, 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 2012, 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), and Khandelwal (2010)). In most of this literature (Atkin, Khandelwal and Osman (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.

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 Herfindal-Hirschmann 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). In future versions of this paper we hope to relax the assumption of evenly spaced deliveries by using machine learning to flexibly predict transaction volume and timing and then classifying false positives as 'failures' to deliver.

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 an event study exercise. 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.

Finally, while we do not formally place our results in general equilibrium, we include a brief discussion of what our results mean for imply aggregate GDP through the lens of a Baqaee and Farhi (2019) framework. To do so, we think of fluctuations in delivery time as captured by the CV as analogous to transitory firm-specific TFP shocks. A high-CV firm is therefore a firm with a higher variance of productivity shocks.[4] We show that GDP losses from TFP shocks are magnified under two conditions: first, when the affected firm is a large or important supplier to other industries, as summarized by Domar weights, and second, when other firms in the economy are 'GE-complements' whose production more positively covaries with the TFP of the affected firm. This second mechanism highlights exactly why a low variance of deliveries is desirable. When trading partners are dependent on an affected firm, for instance as buyers or as co-suppliers in a near-Leontief production function, negative TFP shocks hurt more than positive TFP shocks help. In the extreme case of perfect complementarity, a missing or delayed input decreases trading partner output one-for-one while an additional or early input does nothing. [5]

Under such a framework, our empirical results suggest that small differences in our reliability metric may compound into larger changes in Rwandan GDP. This is because a high CV reduces GDP if, consistent with the reduced-form evidence in this paper, firms (particularly large, productive,

[4]Under this approach we do not need to argue outside the model whether our positive regularity metric implies normative 'reliability'; whether a low variance is good for GDP comes out of the model itself.

[5]The intuition is analogous to why income shocks are more damaging to households when the intertemporal elasticity of substitution is small.

and upstream ones) behave as if output is highly complementary across firms in the short run. In a future version of this paper, we will estimate the building blocks of the Baqaee-Farhi (2019) GDP approximation – firm responses to productivity shocks of key trading partners, output shares in the input-output network, and each firm's 'typical' variance of TFP shocks[6] – and combine them to calculate GDP under counterfactual distributions of firm reliability.

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 discusses how our results might matter for GDP in a Baqaee and Farhi (2019)-style framework. 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

[6] We may use our current CV measure, the CV of output *value* rather than timing, or the residuals from relationship-specific predictions constructed via machine learning.

and short-run substitution is impossible for many intermediates. This assumption plausibly holds for electricity (Fried and Lagakos 2020, Allcott and Collard-Wexler 2016, Burgess, Greenstone, Ryan and Sudarshan 2020) as well as for US importers of intermediate goods (Boehm, Flaaen, and Pandalai-Nayar 2019). Second, to take advantage of reliable supply, firms require capabilities in inventory software, 'just-in-time' production methods (Bloom and Van Reenen 2007), vertical integration of some inputs, or relationships with new suppliers (Grossman and Helpman 2021). 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.[7] 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:

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

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:

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

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

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

which means we will obtain total sales of:

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

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

and input use of:

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

Equations (3) and (6) gives 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.[8]

[8] Kugler and Verhoogen (2012) generate similar predictions for correlations between productivity and input *price*, which is their proxy for quality.

Prediction 4: $Cov(r_y^{buyer}, r_y^{supplier}) > 0$. *Reliable firms will buy from and sell to each other.* Prediction 4 follows from the fact that $r_y^{buyer} = r_s^{buyer}$ and the assumption that a buyer obtains high r_s by choosing a supplier with high r_y. This is a non-trivial assumption; as a counter-example, a buyer could improve the reliability of its input supply by sourcing from *more* suppliers, or suppliers with less correlated risks, even if all suppliers are equally reliable on average.[9]

Note that Predictions 3 and 4 have 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, since 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 (2018). We discuss these possibilities more fully in Section 6.

Since equations (3) and (6) holds in the long-run, we test for Predictions 1-4 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 (2012) 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 inputs to stay close to Section 5.

Prediction 5: 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

[9] Note that prediction 4 is a less informative test of our model since reliable firms might choose reliable suppliers even without heterogeneity in productivity A across firms.

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 equation (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-Aproducers will not. In Section 5, we test for Prediction 5 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 and Zeitlin 2014, Mascagni, Mukama and Santoro 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 usually when the buyer is physically handed the receipt. 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.[10] 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.[11]

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.

[10] Not used in this version of the paper, but can be used as an independent measure of firm size or profits.

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

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 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: VAT transaction date correspond to date that good was exchanged. The validity of this assumption depends on the sector – for some goods, contracting frictions may imply that money should be exchanged on sight, while for others (esp. goods sold in long-term relationships) payment may occur less frequently than goods are actually sent. We plan to conduct a firm survey to identify sectors where transaction dates correspond most closely to goods exchange.

Assumption 2: 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.

Under assumptions 1 and 2, we would ideally proxy for planned goods delivery schedules with relationship-specific transaction histories. With a sufficiently flexible model, we could estimate typical schedules (such as delivering every Monday and Wednesday in several consecutive weeks) and compute the variance of residuals of this model. Alternatively, we could use industry-specific knowledge; for instance Ghani and Reed (2022) use the fact that ice deliveries in Sierra Leone must happen daily. In absence of either, we make a second simplifying assumption:[12]

[12] In initial explorations, we found evidence for such schedules. Regular relationships involving one of Rwanda's

Assumption 3: 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 3 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:

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

The number of days between consecutive transactions is:

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

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, Figure 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

two largest dairy sellers tend to involve deliveries on the same days of the week ('Monday and Thursday' schedules for some buyers, 'Tuesday and Friday' schedules for others). However we have not yet formalized this or tried to identify schedules for a large sample of firms.

is more *reliable*.

To generate our reliability metric, we first compute the CV as in equation (10) for each pair with at least 10 transactions[13] 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 idiosyncracies 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, Mejean and Parenti 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 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

[13] 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.

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

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 2), and desired schedules involve equally-spaced payments (Assumption 3), 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 Figure 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 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 2. 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.[14] 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:

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

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

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 Figure 1. A negative β suggests that 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, Kramarz and Margolis 1999).[15]

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 equation (12) means low r_y, and $\beta < 0$ in equation (11) implies that firms with high characteristic X are more reliable (choose higher r_y).

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 3, without any controls a log point increase in total interfirm sales is associated

[15]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.

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 3 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,[16] the result is qualitatively unchanged though smaller in magnitude.

The negative firm sales-CV correlation in Table 3 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 4 we use multinational firm status and export status as separate proxies for productivity. The negative coefficient in Column 3 suggests that firms that export also reliably supply their domestic buyers, consistent with Prediction 2 from Section 2.[17] The coefficient on MNC status in Column 4 is indistinguishable from zero, perhaps because there are very few MNCs in each sector.

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) or to an MNC (Column 6). Both coefficients are statistically and economically significant: for instance, firms that 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. This effect suggests that productive firms have

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

[17] Note that along with being more physically productive, exports could also higher reliability premia because they sell on international markets.

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 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. It is worth briefly restating the development importance of these facts: if management and supplier reliability are complements, then unreliable upstream inputs (intermittent electricity, suppliers that break contracts, customs delays) will particularly constrain the output of an economy's most productive firms.

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, because reliable firms may endogenously choose reliable suppliers.

To test for sorting, we estimate each supplier's residual unreliability by estimating equation (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. Figure 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. Figure 2 thus suggests that reliability is complementary across levels of a production chain, as in the Leontief production function we laid out in equation (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 (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 Figure 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:

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

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 equation (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 equation (12) on all transactions observed in the base period (January 1 to June 30, 2017) to obtain the residual CV $\widehat{\alpha}_i$ of each supplier. We then subset to suppliers with nonzero base-period imports (the sample for equation (14)), compute the median residual CV in this sample, and then re-estimate equation (14) separately for firms with above- and below-median values of $\widehat{\alpha}_i$. 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 $\widehat{\alpha}_i$ (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 Figure 4 plots estimate of β_τ for 6 weeks before to 6 weeks after the election. In the discussion that follows, we convert from effects on IHS to percent changes where appropriate[18], 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 equation (14). As shown in Panel B of Figure 4, the time-path of effects mirrors the main specification. In the week after the election, the probability

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

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 effects on other outcomes which are not presented in the tables and figures here. First, when we estimate effects on imports from countries *other* than Kenya (to capture substitution to other sources), or estimate differential time trends by CV in imports *only* from Kenya (to capture direct declines on the affected route), declines are larger for high-CV firms but effects are not significant. Thus, due to reasons of statistical power, it is unclear whether reliable firms maintain input supply by finding new suppliers elsewhere or by somehow keeping imports flowing on the existing route. Second, to test if reliable firms also better maintain their *sales* volume during the shock, we estimate effects of the import shock on domestic sales. While exposed firms indeed significantly reduce their sales, we lack sufficient power to test for differential effects between high- and low-CV firms. Finally, we could construct a firm-product-date panel in the import data and estimate effects separately on imports in various categories – for instance, for specific rather than exchange-traded goods (Rauch 1999). Since supply chain interruptions are likely to be more disruptive for differentiated goods, we plan to estimate effects separately by HS code in future versions of the paper.

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 5 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 Figure 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 from improvements in customs processing or political stability that enable a steady and predictable flow of goods.

6 Framework for aggregate effects

Having characterized how reliability choices might be determined, and the consequences for firm-level output, we now examine how reliability choices interact with each other to determine aggregate output. The key conceptual leap we make is that a firm's reliability choice amounts to a choice of production technology. In particular, the firm *chooses the variance of TFP shocks it will face.* We can then apply results Baqaee and Farhi (2019) that characterize the change in GDP, up to the second order, due to technology shocks in an efficient network economy.

The core result is that variable production by any firm i will depress aggregate GDP when many firms j are GE complements with i, a condition that in practice means that i is 'critical' for j. This condition has wide relevance in practice: firms using specific inputs or in long-term relationships are GE complements of their suppliers, and most manufacturers are GE-complements of the electricity sector. Intuitively, under GE complementarity, negative shocks to i hurt j more than positive shocks to i help j. As a result, a high variance of TFP shocks at such a firm i will depress expected GDP.

We now express this result mathematically. Define the following:

1. **Shocks:** Over the course of day-to-day operations, firm i experiences TFP shocks $log(A_i) \sim N(0, \sigma_i^2)$. σ_i^2 captures the unreliability of a supplier.

2. **Importance:** Define Domar weights:

$$\lambda_i = \frac{p_i y_i}{\sum_j p_j c_j}$$

   λ_i is a sort of adjusted sales share that capture each supplier's importance in production. Large firms and suppliers of intermediates (i.e. for whom $y_i > c_i$) will tend to have large λ_i.

3. **Complementarity:** Define the *GE elasticity of substitution* $\rho_{ji} > 0$ between firms i and j such that:

$$\frac{d\log[(Y_i/Y_j)]}{d\log A_i} = -\frac{1}{\rho_{ji}} \quad \rightarrow \quad \frac{d\log(\lambda_i/\lambda_j)}{d\log A_i} = 1 - \frac{1}{\rho_{ji}}$$

   When ρ_{ji} is small, and in particular less than 1, an positive (negative) shock to firm i's productivity will tend to increase (decrease) firm j's adjusted sales share relative to that of the shocked firm i. This condition would hold if firms i and j are actually gross complements in the production of a downstream firm k, but also if j uses i as an input and no good short-run alternative exists.

Next, define sales shares $s_i = \frac{p_i y_i}{\sum_k p_k y_k}$ which sum to 1. Baqaee and Farhi (2019) show that, up to a second-order Taylor approximation, expected log of GDP (Y) relative to steady state with no productivity shocks $(\bar{Y})$ is:

$$\begin{aligned} \mathbb{E}\left[\log(Y/\bar{Y})\right] \approx \frac{1}{2}\Big[\sum_i \frac{d^2\log Y}{d\log A_i^2}\Big] &= \frac{1}{2}\Big[\sum_i \Big(\frac{\lambda_i}{\sum_k \lambda_k}\sum_{j\neq i}\lambda_j(1-\frac{1}{\rho_{ji}}) + \lambda_i\frac{d\log(\sum_i \lambda_i)}{d\log A_i}\Big)\sigma_i^2\Big] \\ &= \frac{1}{2}\Big[\sum_i \Big(p_i y_i \sum_{j\neq i} s_j(1-\frac{1}{\rho_{ji}}) + \lambda_i\frac{d\log(\sum_i \lambda_i)}{d\log A_i}\Big)\sigma_i^2\Big] \end{aligned} \tag{15}$$

Equation (15) tells us the lost GDP due to a particular level of unreliability of production, as summarized by the variance of TFP σ_i^2. Three factors characterize the extent of the fall (or rise) in GDP due to σ_i^2. To start, for GDP to fall, we require that $\rho_{ji} < 1$ for some firms j. In other words, some other firms must be GE complements with i, for instance if they depend on i for inputs and are likely to see sales fall if i has a negative shock.[19] Furthermore, those firms j must be large themselves, to the extent that the sales-weighted average extent of complementarity of all firms with j is negative: $\sum_j s_j(1-\frac{1}{\rho_{ji}}) < 0$. Finally, GDP will decline more when the shocked firm itself is large ($p_i y_i$ large), including through intermediate sales.[20]

Why do complementarities imply that unreliable production reduces expected GDP? The intuition is analogous to that of risk aversion when the intertemporal elasticity of substitution is small: due to exogenous technology (preferences), the aggregating firm (representative household) has diminishing marginal product (utility) across products (time periods), and therefore symmetric

[19] In contrast, if i has a positive TFP shock, j will be unhappy because p_i will fall and there will be GE substitution *away* from j.

[20] For now, we do not discuss the second term, which captures the extent to which the aggregate input-output multiplier responds to A_i; and assume it is small. For inputs like electricity this may not be the case.

shocks reduce expected output (expected utility). Our key point here is that **no actual household risk aversion is required** to make unreliable production damaging: even with risk-neutral households and firms, low GE elasticities of substitution makes the representative firm act risk-averse in its use of various inputs.

Moving forward, to calibrate equation (15) and quantify the GDP implications of unreliability we require three ingredients. The first is estimated (un)reliabilities σ_i^2, which we can estimate by isolating supplier i-specific deviations from production schedules or any other i-specific supply shock. The second is baseline sales shares λ_i, which we can read off the VAT data, with some care due to informality in our setting. The third and most important consists of estimates of firm-specific responses to productivity shocks ρ_{ji}: to estimate these we will need to construct firm-specific shocks, such as from sudden border closures (as in the Kenya case) that affect specific firms or from supplier-specific fluctuations in our data, and estimate effects on sales of trading partners. These elasticities will need to be estimated sector-by-sector, since it is unlikely that stochastic production is actually costly for all industries. At a high level, the key benefit of such an approach to aggregation is that we do *not* need to directly find exogenous shifters of firm-specific σ_i^2 for all firms i, which is useful since mean-preserving second-moment shocks are rare in practice. Instead we only require idiosyncratic shocks to A_i to identify ρ_{ji}, just as household income shocks can be used to estimate coefficients of risk aversion, as illustrated for instance in Ganong et al. (2020).

We briefly discuss what our reduced-form results in Sections 4 and 5 imply for GDP through the lens of equation (15). First, the fact that reliability levels are strongly correlated between suppliers and buyers suggests that ρ_{ji} is low: firms would only exhibit such behavior if they were unable to substitute in the short-run. Second, if we interpret the Kenya shock as a one-time supply shock (which is reasonable given that it only lasts three weeks), several firms exhibit little-to-no substitution to other routes, again consistent with low ρ_{ji} when i is a foreign supplier. Third, the fact that productive firms and exporters are more reliable suggests that their entry may improve aggregate output, especially if they sell downstream in the domestic market as well. Together these facts point to magnified negative effects of unreliable firms on expected GDP.

Finally, we note two key caveats in the application of Baqaee and Farhi (2019) to our setting. First, it is not obvious if the CV of time between shipments (or any measure of deviations from a production schedule) is analogous to a variance of TFP shocks. The CV is a production practice rather than residual output, and in addition, transactions will exhibit some variability (in timing and/or in value) due to transitory TFP shocks at other firms, even if the firm in question is not experiencing a shock. We leave the task of constructing a reliability metric that maps directly into our aggregation framework for future work. Second, the current framework takes σ_i^2 as exogenous, while much of our discussion concerned firms' endogenous choice of reliability levels. As a result there is no room in our framework for counterfactuals in which MNC entry or other changes in the *demand* for reliability might lead to 'reliability upgrading' through changes in the vector $\boldsymbol{\sigma^2}$.

7 Conclusion

In this paper we measure 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.

Interpreted through a production network model (Baqaee and Farhi 2019), these results imply that the reliability of a small number of firms may have economywide effects, and therefore that policymaker efforts to improve reliability – such as by attracting multinational corporations or investing in electricity or contract enforcement – can induce a series of interlocked improvements in the reliability of domestic firms elsewhere in the trading network. Thus, while the analysis thus far examines the *existing* distribution of reliability across firms, in future work we plan to evaluate reliability-improving policies explicitly and quantify the effects on aggregate output.

References

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

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

[3] Hunt Allcott, Allan Collard-Wexler, and Stephen D O'Connell. How do electricity shortages affect industry? evidence from india. *American Economic Review*, 106(3):587–624, 2016.

[4] Sam Asher and Paul Novosad. Rural roads and local economic development. *American economic review*, 110(3):797–823, 2020.

[5] John Asker, Allan Collard-Wexler, and Jan De Loecker. Dynamic inputs and resource (mis) allocation. *Journal of Political Economy*, 122(5):1013–1063, 2014.

[6] David Atkin, Amit K Khandelwal, and Adam Osman. Exporting and firm performance: Evidence from a randomized experiment. *The Quarterly Journal of Economics*, 132(2):551–615, 2017.

[7] Richard Baldwin and Rebecca Freeman. Risks and global supply chains: What we know and what we need to know. Working Paper 29444, National Bureau of Economic Research, October 2021.

[8] Abhijit V Banerjee and Esther Duflo. Reputation effects and the limits of contracting: A study of the indian software industry. *The Quarterly Journal of Economics*, 115(3):989–1017, 2000.

[9] David Rezza Baqaee and Emmanuel Farhi. The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem. *Econometrica*, 87(4):1155–1203, 2019.

[10] Jean-Noël Barrot and Julien Sauvagnat. Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks *. *The Quarterly Journal of Economics*, 131(3):1543–1592, 05 2016.

[11] Andrew B Bernard, Emmanuel Dhyne, Glenn Magerman, Kalina Manova, and Andreas Moxnes. The origins of firm heterogeneity: A production network approach. Working Paper 25441, National Bureau of Economic Research, January 2019.

[12] Andrew B Bernard, J Bradford Jensen, Stephen J Redding, and Peter K Schott. Firms in international trade. *Journal of Economic perspectives*, 21(3):105–130, 2007.

[13] Nicholas Bloom. The impact of uncertainty shocks. *Econometrica*, 77(3):623–685, 2009.

[14] Nicholas Bloom and John Van Reenen. Measuring and explaining management practices across firms and countries. *The quarterly journal of Economics*, 122(4):1351–1408, 2007.

[15] Christoph E Boehm, Aaron Flaaen, and Nitya Pandalai-Nayar. Input linkages and the transmission of shocks: Firm-level evidence from the 2011 tōhoku earthquake. *Review of Economics and Statistics*, 101(1):60–75, 2019.

[16] Robin Burgess, Michael Greenstone, Nicholas Ryan, and Anant Sudarshan. The consequences of treating electricity as a right. *Journal of Economic Perspectives*, 34(1):145–69, 2020.

[17] Julia Cajal-Grossi, Rocco Macchiavello, and Guillermo Noguera. International buyers' sourcing and suppliers' markups in bangladeshi garments. 2019.

[18] Vasco M Carvalho, Makoto Nirei, Yukiko U Saito, and Alireza Tahbaz-Salehi. Supply chain disruptions: Evidence from the great east japan earthquake. *The Quarterly Journal of Economics*, 136(2):1255–1321, 2021.

[19] Banu Demir, Ana Cecília Fieler, Daniel Xu, and Kelly Kaili Yang. O-ring production networks. Working Paper 28433, National Bureau of Economic Research, February 2021.

[20] Nada Eissa, Andrew Zeitlin, Saahil Karpe, and Sally Murray. Incidence and impact of electronic billing machines for vat in rwanda. *London: International Growth Centre*, 2014.

[21] Ana Cecília Fieler, Marcela Eslava, and Daniel Yi Xu. Trade, quality upgrading, and input linkages: Theory and evidence from colombia. *American Economic Review*, 108(1):109–46, 2018.

[22] Ana Cecília Fieler, Marcela Eslava, and Daniel Yi Xu. Trade, quality upgrading, and input linkages: Theory and evidence from colombia. *American Economic Review*, 108(1):109–46, January 2018.

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

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

[25] Stephie Fried and David Lagakos. Electricity and firm productivity: A general-equilibrium approach. Working Paper 27081, National Bureau of Economic Research, May 2020.

[26] Peter Ganong, Damon Jones, Pascal J Noel, Fiona E Greig, Diana Farrell, and Chris Wheat. Wealth, race, and consumption smoothing of typical income shocks. Working Paper 27552, National Bureau of Economic Research, July 2020.

[27] Shresth Garg and Brandon Tan. Within firm supply chains: Evidence from india. 2021.

[28] Tarek Ghani and Tristan Reed. Relationships on the rocks: Contract evolution in a market for ice. *American Economic Journal: Microeconomics*, 14(1):330–65, February 2022.

[29] Gene M Grossman, Elhanan Helpman, and Hugo Lhuillier. Supply chain resilience: Should policy promote diversification or reshoring? Working Paper 29330, National Bureau of Economic Research, October 2021.

[30] Juan Carlos Hallak and Peter K. Schott. Estimating Cross-Country Differences in Product Quality*. *The Quarterly Journal of Economics*, 126(1):417–474, 02 2011.

[31] David L Hummels and Georg Schaur. Time as a trade barrier. *American Economic Review*, 103(7):2935–59, 2013.

[32] Amit Khandelwal. The Long and Short (of) Quality Ladders. *The Review of Economic Studies*, 77(4):1450–1476, 10 2010.

[33] Michael Kremer. The o-ring theory of economic development. *The Quarterly Journal of Economics*, 108(3):551–575, 1993.

[34] Michael Kremer, Jean Lee, Jonathan Robinson, and Olga Rostapshova. Behavioral biases and firm behavior: Evidence from kenyan retail shops. *The American Economic Review*, 103(3):362–368, 2013.

[35] Maurice Kugler and Eric Verhoogen. Prices, Plant Size, and Product Quality. *The Review of Economic Studies*, 79(1):307–339, 11 2011.

[36] Ernest Liu. Industrial Policies in Production Networks*. *The Quarterly Journal of Economics*, 134(4):1883–1948, 08 2019.

[37] Rocco Macchiavello. Relational contracts and development. *Working Paper*, 2021.

[38] Rocco Macchiavello and Ameet Morjaria. The value of relationships: evidence from a supply shock to kenyan rose exports. *American Economic Review*, 105(9):2911–45, 2015.

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

[40] Giulia Mascagni, Denis Mukama, and Fabrizio Santoro. An analysis of discrepancies in taxpayers' vat declarations in rwanda. 2019.

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

[42] Dina Pomeranz. No taxation without information: Deterrence and self-enforcement in the value added tax. *American Economic Review*, 105(8):2539–69, 2015.

[43] James E. Rauch. Networks versus markets in international trade. *Journal of International Economics*, 48(1):7–35, 1999.

[44] Peter K. Schott. Across-Product Versus Within-Product Specialization in International Trade*. *The Quarterly Journal of Economics*, 119(2):647–678, 05 2004.

[45] Eric A Verhoogen. Trade, quality upgrading, and wage inequality in the mexican manufacturing sector. *The Quarterly Journal of Economics*, 123(2):489–530, 2008.

Tables and Figures

Figure 1: Reliability calculation example using two dairy manufacturers

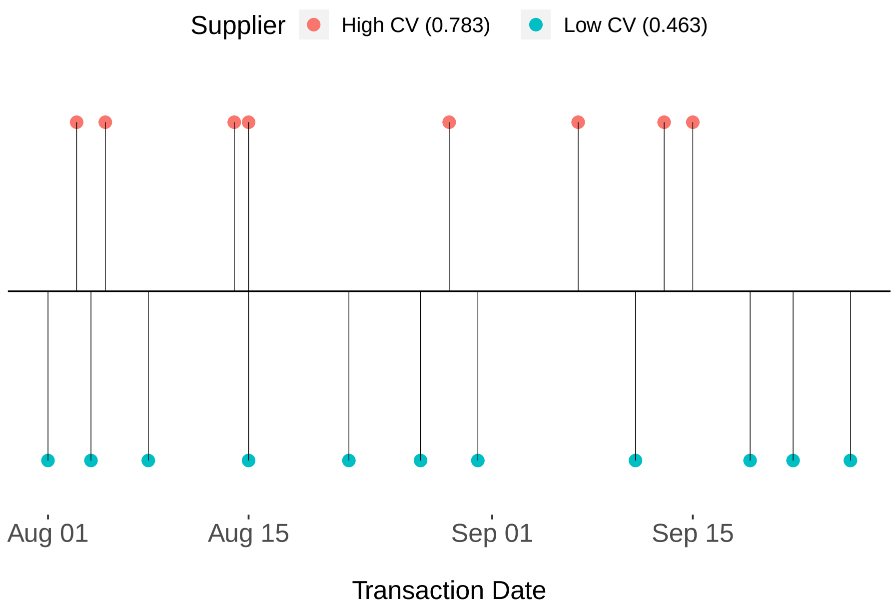

Notes: Figure shows transaction dates for two suppliers of dairy products from August 1 to September 30, 2017 to a common buyer. The CV shown is the coefficient of variation of the time, in days between consecutive shipments. Transaction dates are obtained from VAT electronic billing machine (EBM) records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

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: 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 activites	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).

Table 3: 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).

Table 4: 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			-0.030***				-0.035***
			(0.007)				(0.008)
MNC				0.012			0.029**
				(0.015)			(0.015)
Share of sales to exporters					-0.266***		-0.137**
					(0.050)		(0.060)
Share of sales to MNCs						-0.449***	-0.285**
						(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	0.73810	0.73783	0.73830	0.73816	0.73884
Within R^2	0.00054	0.00142	0.00108	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. *Exporter* is an indicator for whether an export transaction was recorded by the seller and MNC is a dummy for foreign or joint ownership. The share of sales to MNCs and to exporters are computed by summing sales to exporters / MNCs and dividing by total sales (interfirm + other). Sample consists of all transactions in 2017. All data are from the Rwanda Revenue Authority (RRA).

Figure 2: Reliable firms sell to and buy from each other

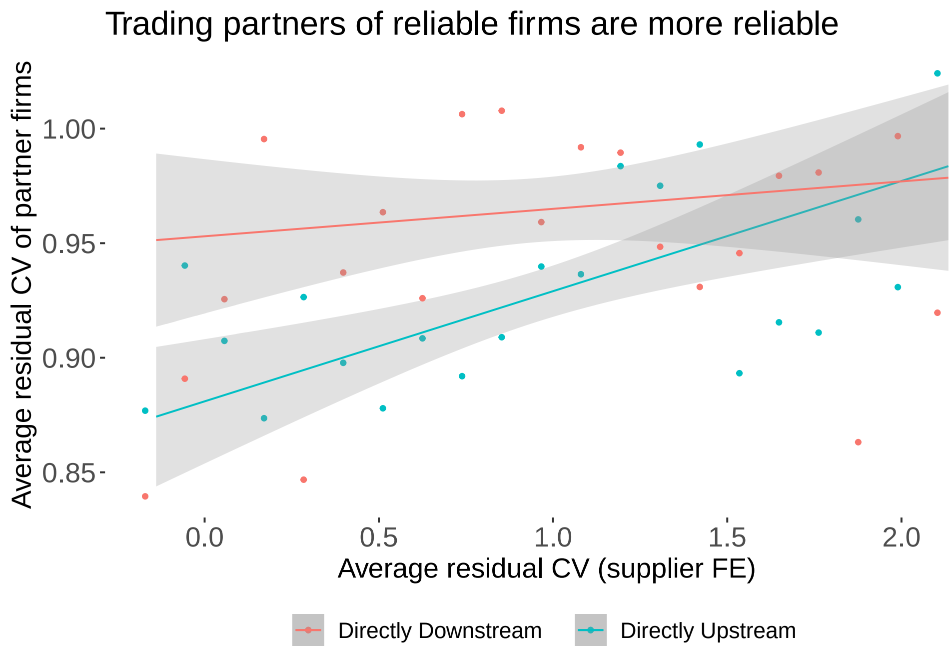

Notes: Figure shows each firm's residual unreliability $\widehat{\alpha}_i$ on the horizontal axis and the average residual unreliability of that firm's direct suppliers (in blue) or direct customers (in red) on the vertical axis. Both binned scatterplots (for 20 bins of $\widehat{\alpha}_i$) and the 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 $\widehat{\alpha}_i$ before plotting. $\widehat{\alpha}_i$ 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).

Figure 3: Rwandan imports to Kenya decline in weeks around 2017 presidential election

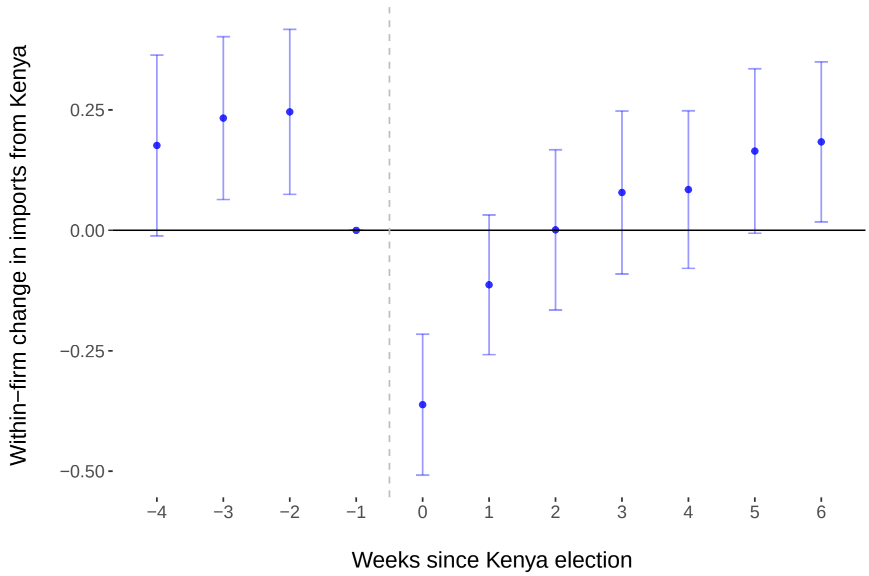

Notes: Figure shows estimated weeks-since-election coefficients from a regression of the inverse hyperbolic sine (IHS) of total imports on weeks-since-election dummies and firm ID-day of week fixed effects. Regression is estimated at the firm-by-date level using a panel of import transactions from January 1 to November 30, 2017. The grey 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.

Figure 4: Import declines are larger for unreliable (above-median CV) firms after Kenya election

Panel A: Event study effects on $sinh^{-1}$(total imports)

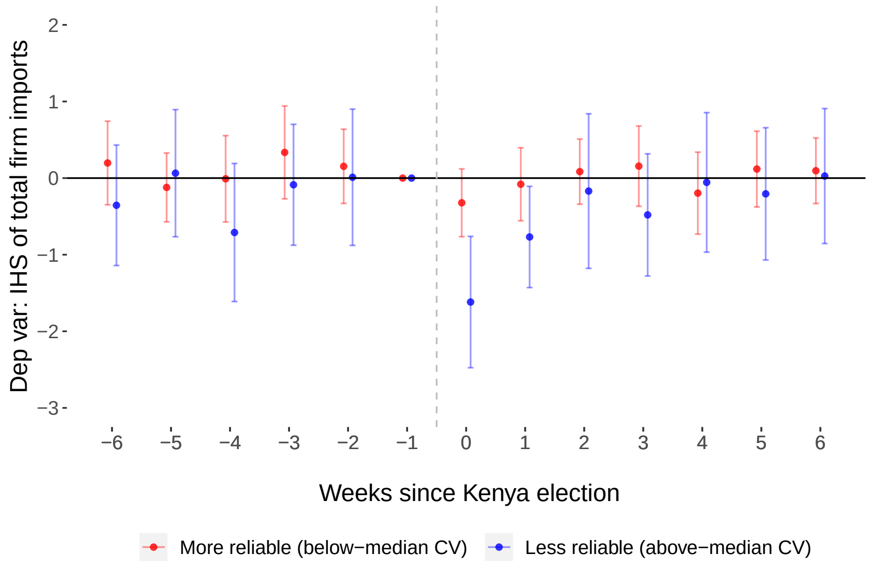

Panel B: Event study effects on probability of any import

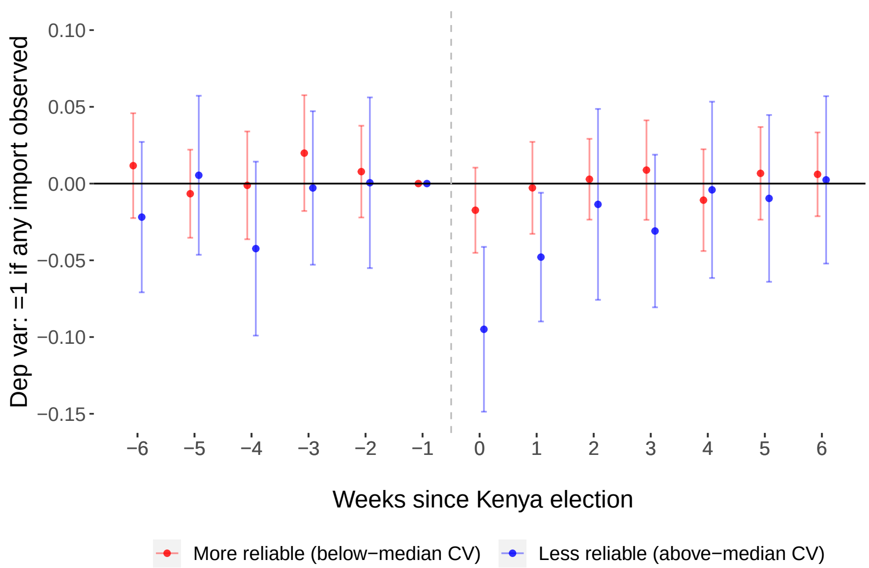

Notes: Figure shows estimated coefficients on interactions between week-since-election dummies and the share of a firm's pre-period imports that are from Kenya. All regressions are at the firm-by-date level. The dependent variable is either the inverse hyperbolic sine (IHS) of total imports (Panel A) 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 $\widehat{\alpha}_i$ for each supplier i. $\widehat{\alpha}_i$ 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.

Appendix

Figure A1: Low-CV firms concentrate transactions on selected weekdays or days of month

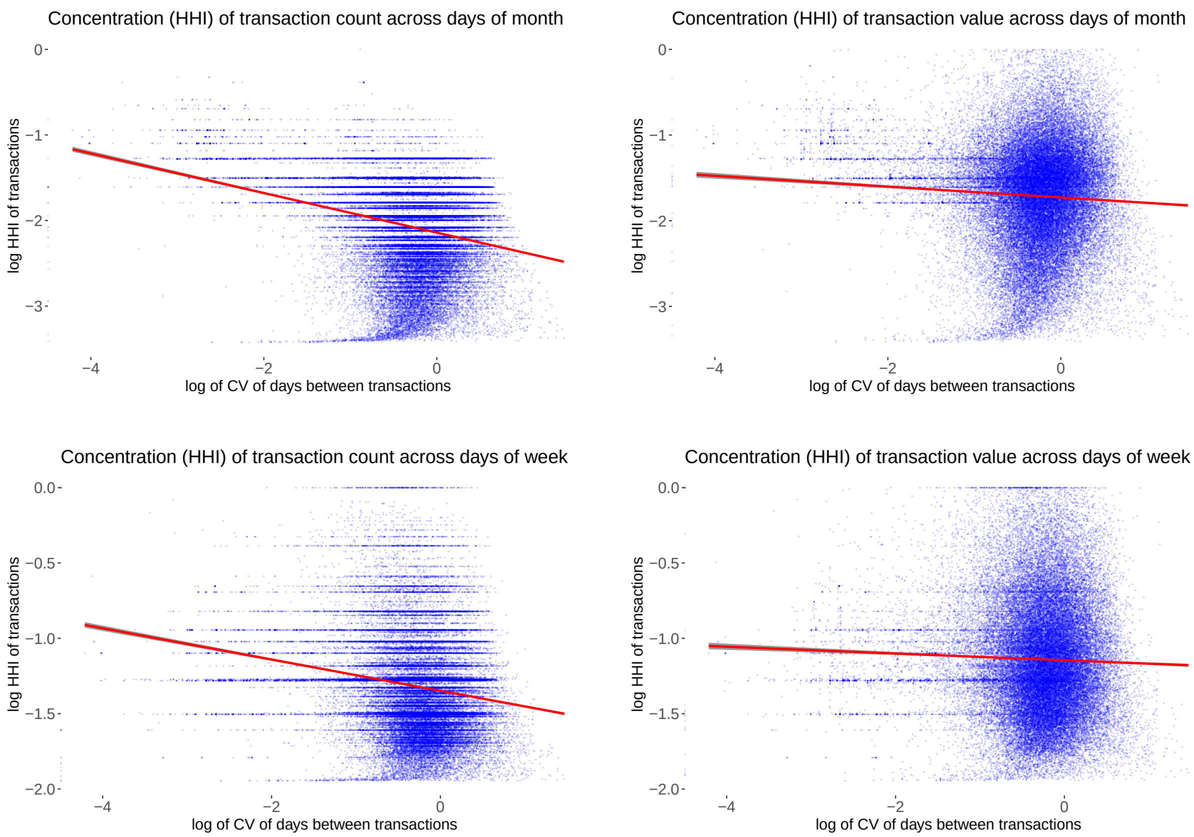

Notes: Sample consists of all relationships with at least 5 transactions from January 1 to June 30, 2017.

Figure A2: Daily sales and purchases for selected Bralirwa wholesalers

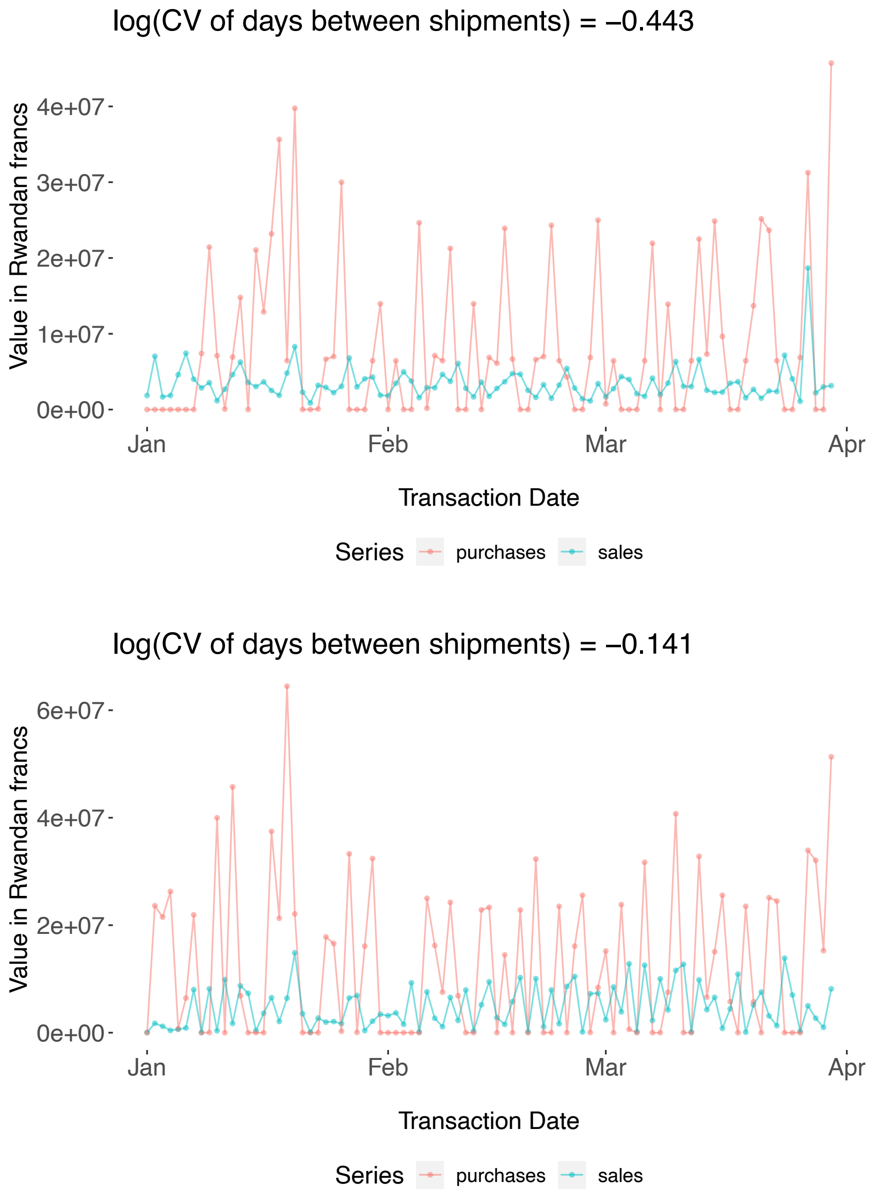

Notes: Purchases and sales are summed over all supplier and customers, including non-registered firms and households to get change in total inventory. Bralirwa is the largest manufacturer of beer and soft drinks in Rwanda.