Supply chain resilience via partial integration

Aroon Narayanan
MIT

Vishan Gandhi Nigam
MIT

September 28, 2025

Abstract

How do firms adapt to ensure supply chain resilience? The typical answer, vertical integration, is limited in practice by its high costs and inflexibility. This paper introduces *partial integration*, defined as targeted buyer interventions across firm boundaries, as an effective alternative. Using novel daily timeline data from Indian manufacturing supply chains, we show that supplier underinvestment in key inputs, stemming from working capital constraints and noncontractibilities in input use, is the primary driver of disruptions in these supply chains. To reduce the incidence of disruptions, buyers exert control over supplier processes—through in-person monitoring, contingent contracts, and direct sourcing of raw materials—rather than merely advancing cash. This buyer involvement escalates as disruption risk increases: an unanticipated input cost shock leads to direct buyer control of inputs for the most constrained suppliers. We develop a three-stage model that rationalizes these strategies. The model clarifies that buyers control input decisions to prevent resource diversion due to noncontractibility, while allowing suppliers to retain control of production in order to preserve output market incentives. It also predicts that relational buyers with low monitoring and sourcing costs enjoy a comparative advantage in fostering resilient trade with poor regions.

We are grateful to Charles Angelucci, David Atkin, Abhijit Banerjee, and Robert Gibbons for support and guidance throughout this project. We also thank Silke Forbes, Namrata Kala, Ameet Morjaria, Jorge Tamayo, and participants of the MIT Organization Economics lunch for helpful comments. We acknowledge support from the George and Obie Shultz Fund, the Jerry Hausman Graduate Dissertation Fellowship, the Sloan Africa Fellowship (Aroon), and the NSF Graduate Research Fellowship Program (Vishan). Aroon: aroon@mit.edu, Vishan: vishan@mit.edu.

1 Introduction

Supply chain disruptions are an increasing concern for firms around the world. Climate risks threaten production sites directly; geopolitical tensions affect the supply of critical inputs; and as production grows more complex, the costs of a single low-quality input grows exponentially (Kremer, 1993). These disruptions are exacerbated by the production technology of modern supply chains: thanks to Just-In-Time (JIT) methods, inputs must arrive in sufficient quantities on strict daily or hourly timelines, because any delays disrupt downstream assembly.

Vertical integration is typically considered the primary tool for achieving supply chain resilience. It allows buyers direct control over supplier operations, thereby minimizing disruptions and ensuring the timely flow of high-quality goods.[1] However, most modern supply chains are not vertically integrated. Beyond the challenges of acquiring and managing another firm, integration is also a costly and blunt instrument, requiring acquisition of an entire firm to address issues in even a single production stage. The result is that buyers and suppliers must achieve resilient production while remaining independent, and we have little knowledge of how they do so.

This paper shows that firms respond to supply chain risk via organizational interventions that target critical stages of production within upstream suppliers. We refer to these interventions as *partial integration* because they mirror the activities of a vertically integrated firm, but are implemented across firm boundaries. Leveraging novel administrative production timeline data from Indian manufacturing supply chains, we first establish that disruptions occur primarily because suppliers underinvest in inputs required for new orders. Two specific frictions drive input underinvestment: working capital constraints that limit supplier investment, and noncontractibilities that inhibit buyer ability to finance suppliers' inputs directly. Our primary result is that these frictions are mitigated by blurring firm boundaries: specifically, buyer control over critical supplier processes – in our case, input sourcing – enables timely delivery. These partial integration decisions take many forms: managing supplier production sites; making payments conditional on input purchase; and when disruption risks increase, directly sourcing and controlling suppliers' inputs.

To study supply chain risk, we require supplier delivery delays relative to plan, which are not recorded in standard trade datasets. We overcome this challenge by introducing novel administrative production timeline dataset from a large Indian manufacturing buyer. The buyer facilitates contracts between downstream multinational construction conglomerates and highly

[1] See, for instance, Hansman, Hjort, León-Ciliotta, et al. (2020) in Peru.

informal upstream suppliers. We focus on suppliers of fabricated metals, a key input for wind turbines, solar panels, skyscrapers, power plants, railways, and other modern infrastructure. Because downstream customers require Just-In-Time (JIT) delivery to production sites, the buyer tracks movement of each batch across production stages, from the sourcing and cutting of steel sheets to the welding, painting, and inspection of the finished good. As a result, we have essentially full visibility into batch-level production activities for over 200 firms that provide over 10% of industry capacity. This detail is both necessary to study JIT production and extremely rare: for confidentiality reasons, the existence of within-firm production data harmonized across multiple firms is highly infrequent.

Before documenting how buyers innovate to ensure on-time delivery, we first establish that supplier production delay occur because credit and contracting frictions inhibit supplier input investment. We first document descriptively that production delays are an investment problem: delays in sourcing raw material are the largest single determinant of delays, more than issues with labor availability, work speed, inspection failures, or transportation; and these sourcing delays are especially large in credit-poor regions. We then show that suppliers' purchases of raw material for new orders occur exactly when payment from past orders arrives, which suggests that the suppliers operate as-if hand-to-mouth[2] waiting to obtain working capital from previous orders to invest into future inputs. We tighten this link by conditioning on planned sourcing timelines and payment, and by instrumenting for delayed payment using a leave-out estimate of downstream customer delays on common contracts with other suppliers. The effect of payments on new sourcing is preserved in the IV regression, with larger effects on ex-ante constrained[3] suppliers.[4]

Our main result is that the buyer responds to supplier credit constraints by taking on management and sourcing responsibilities inside supplier plants that are exclusively linked to the procurement of raw material. These responsibilities first appear in observed buyer behavior: we show that visits by the buyer's relationship managers to production sites, as proxied by the frequency of internal updates to production statuses, increase substantially when pay-

[2]While this might be surprising in a rich country, inability to produce due to delayed payment from customers is often the first- or second-most important complaint of firms in the World Bank Enterprise Surveys, and has motivated substantial policy action by the Indian government in recent years, including a nationwide marketplace for factoring (short-term trade finance) and a ban on payment beyond 45 days after delivery by conglomerates to micro, small and medium enterprises (MSMEs).

[3]We use two proxies for constraints: supplier region (central & north India are less financially developed than the south and west), and whether the supplier is offered partial advance payment for raw material.

[4]Our results are conceptually similar to those of classic tests of investment-cash flow sensitivity: firms invest when the shadow interest rate falls due to an exogenous cash inflow. Our contribution is to show that cash flows matter for short-run input sourcing – in other words, for a firm's ability to produce Just-In-Time for a customer – and that production may be disrupted by the demands of other customers with similarly urgent production requirements.

ments arrives (using the same downstream customer delays instrument). We interpret this as a form of management via monitoring of supplier input investment: when cash arrives and can potentially be invested in a new order, the buyer engages with suppliers to ensure that inputs are actually procured.

The more striking form of buyer involvement in supplier operations – in other words, of of partial integration – is directly observable from contract terms. By analyzing provisions of individual contracts, we document that the buyer finances supplier input investment via two distinct contractual innovations. First, the buyer offers conditional advances (incentives) to selected suppliers, whereby the supplier is reimbursed conditional on the purchase of raw material. Second, the buyer sometimes chooses to directly source (control) the raw material herself, thereby gaining decision rights on when sourcing takes place, and mitigating any possibility of delayed sourcing of new orders. These terms are strongly suggestive of imperfect contract enforcement. In a setting with full enforcement, the buyer could directly relax the supplier's credit constraint buy advancing credit herself. The fact that we see alternative strategies suggests that such a contract can be hard to enforce, perhaps because credit can be diverted to inputs for other customers.

Finally, leveraging an unanticipated shock to supplier input costs, we show that increases in supply chain risk lead the buyer to own and control suppliers' inputs – in other words, to involve itself more directly – rather than rely on sourcing incentives. In a difference-in-difference framework, we show that in response to a large and persistent increase in supplier working capital needs, driven by a rise in global steel prices, the buyer switches from providing conditional advances (or no support) to providing the raw material herself to the supplier. The suppliers switched to direct sourcing are exactly those with larger ex-ante credit needs, as proxied by their use of conditional advances or propensities to delay before the steel price shock. In other words, when working capital needs (and thus risks of input sourcing delays) increase, the buyer invests in raw material on behalf of the most credit-constrained suppliers, but still does not participate in the actual transformation of raw steel into the customized final product.

Our empirical results lead to several questions: why do credit constraints lead to more buyer control, as opposed to simply providing cash upfront? How does the buyer decide the extent to which she exerts control over the supplier's production process? And given our results, which buyers are best positioned to reliably source from countries with weak contracting environments? To answer these questions, we build a sequential model of a supply chain in which production has three stages: input use, physical production, and allocation of output. The key friction is the noncontractibility of decisions, in particular that suppliers controlling input decisions may misallocate cash or physical inputs.

The model clarifies that the buyer's choice of how far to integrate along the supplier's processes is essentially a question of which market incentives to allow the supplier to retain. In our case, the buyer partially integrates the supplier because, without buyer involvement, supplier ownership of working capital creates perverse incentives to redirect cash and inputs towards other orders in times of a high shadow price. A similar mechanism explains why the buyer uses direct sourcing in lieu of conditional advances when input costs rise: while less fungible than cash, steel can still be diverted. Thus, when the deviation temptation rises as steel price rises, the buyers must directly control input use. At the same time, supplier ownership of the customized final product (in our case, a fabricated metal part) provides good incentives: since the supplier owns part of the surplus from the final product and also bears the shadow cost of idle capacity, she is incentivized to minimize delays once raw material is purchased. We can therefore rationalize the partial nature of integration that we observe in this setting in place of full vertical integration.

Our results suggest that the key to resilient production in settings with weak institutions is close engagement with suppliers: monitoring of production sites, incentives conditional on inputs arriving at supplier sites, and direct ownership and control of suppliers' inputs. A natural implication is that relational firms with lower costs of close engagement, because of their organizational structures or geographic location, can source more effectively in such settings. This logic implies, for instance, that relational buyers with Japanese-style sourcing strategies will more effective at resilient sourcing than American-style spot buyers who use arms-length contracts to enforce delivery. A policymaker in a weak institutional setting seeking to promote global supply chain integration might therefore aim to subsidize exports to (or FDI by) Toyota in lieu of General Motors, or Levi Strauss in lieu of JC Penney[5]. A more generalizable insight is that intermediaries and joint ventures with local firms may be critical for reliable sourcing from developing countries.[6]

Related literature This paper contributes to several literatures. The first is on vertical integration. An established theoretical literature shows that ownership reallocates residual control rights to prevent holdup and sharpen worker incentives (Antràs, 2003; Antràs and Helpman, 2004; Baker, Gibbons, and Murphy, 2002), and a related empirical literature shows that a large

[5]See **CajalGrossi2022**, Heise et al. (2024), and Helper and Munasib (2022) on Japanese vs. US sourcing strategies.

[6]Indeed, intermediaries such as Li and Fung in the global garment sector and our anonymous partner firm in Indian metal fabrication play a critical role in communicating and monitoring production timelines on behalf of their downstream buyers.

share of trade takes place within integrated firms and that trade frictions are lower within firm boundaries (Alfaro et al., 2025; Atalay et al., 2019; Garg, Ghosh, and Tan, 2023), as well as that buyers acquire suppliers when the returns to high-quality inputs increase (Hansman, Hjort, León-Ciliotta, et al., 2020). We contribute by showing that buyers can replicate many functions of integration through *partial integration* – monitoring, contingent incentives, and direct control of input sourcing – without actual ownership; and provide causal difference-in-difference evidence that a firm's chosen form of partial integration can change in response to disruption risk.

A second literature documents how firms adopt "thick" contractual arrangements that relax constraints to production by blurring firm boundaries. For instance, mills in Rwanda provide inputs and credit to farmers (Macchiavello and Morjaria, 2021), and in African agricultural supply chains goods and credit are often bundled (Casaburi and Reed, 2022). Relative to this work, we study a different industry (Just-In-Time construction supply chains) and specifically explore how buyers assist or take over sourcing of suppliers' upstream inputs to keep production on schedule.

Our paper also relates to a literature on the sourcing strategies used by firms specifically to avert supply chain disruptions. Most of this work focuses on long-term relationships, and specifically shows that promised future relationship value sustains trade during shocks (Macchiavello and Morjaria, 2015; Macchiavello and Miquel-Florensa, 2017; Ghani and Reed, 2022; Adhvaryu, Kala, and Nyshadham, 2024) by avoiding temptation to resell output on spot markets. We uncover a different margin of disruption – input diversion risk due to working capital constraints – and different set of remedies that rely on closer buyer involvement rather than the promise of future business. These practices mirror the "Toyota" model of buyer involvement (Helper and Henderson, 2014; Helper and Munasib, 2020; Jones, Womack, and Roos, 1990), and thus we also connect to the literature on organizational forms that enable Just-In-Time production.

Finally, a growing literature on trade amid credit frictions shows that high interest rates inhibit working capital-intensive production in poor countries (Antràs, 2023; Antràs and Tubdenov, 2025), that credit market imperfections further restrict trade (Manova, 2013), and that buyers supply inputs in-kind to credit-constrained Chinese exporters to overcome these challenges (Manova and Yu, 2016). Relative to this work, we explore how firms choose between two different financing tools – conditional advances and direct sourcing – and provide causal evidence on when each is chosen. In this sense, our paper resembles that of Antràs and Foley (2015), who explore how suppliers adapt to extend trade credit to buyers in weak legal environments.

The remainder of this paper proceeds as follows. Section 2 describes the empirical setting of Indian metal fabrication, our data, and several stylized facts. Section 3 explains our two empirical strategies, which test for credit-related input sourcing disruptions and contract changes in response to increased disruption risk. Section 4 contains our results. In Section 5, we build a model of the supply chain to rationalize the partial integration strategy we observe and discuss the key takeaways. Section 6 concludes.

2 Setting and Data

2.1 Setting

Our setting is Indian manufacturing, specifically the metal fabrication sector. Downstream customers, often construction companies, require custom-made steel parts that are produced according to engineering designs. For example, they may require steel beams and rods that go into the construction of a building, with specifications that are dictated by the specifics of the design of the building. The firms in our sample supply much of India's rapidly-growing critical infrastructure, including electrical transformers, power plant boilers, tracks for the Indian Railways, bases of wind turbines, and high-rise buildings in major cities. The vast majority of these projects are Just-In-Time: all parts are unique, and there is little room on-site to store deliveries, so the entire project gets delayed if delivery of any part is incomplete.

Our data comes from a single firm (henceforth "the buyer") that sources contracts from such downstream customers, and maintains relationships with small and medium manufacturing firms that have the capacity to produce the parts required by the contract. She then breaks up the contracts into smaller orders, which are then allocated to suppliers depending on availability and quoted prices.

While the buyer and its downstream customers are large multinational or at least India-wide conglomerates, suppliers are mostly small and medium enterprises (MSMEs), and hence operate with limited formality. This means that some level of legal enforcement is available, for example with respect to ownership of assets, while enforcement of more complicated legal activities, such as complex contracts that depend on delivery times, is often infeasible (Boehm and Oberfield (2020)). Suppliers also have little access to outside credit, especially for short-term working capital. Consistent with borrowing constraints, in interviews, several firm owners described their their capital stock as a fixed amount that is often 'locked' in as-yet-unpaid invoices for previous orders.

Critically, suppliers require working capital to invest in new orders, in particular to purchase raw materials such as steel. These purchases make up the bulk of the cost (approximately 50-70%) of fulfilling an orders, and can be an order of magnitude large than labor expenses (10-20%) and typical profit margins (5-10%). Because steel is a commodity and sold on a spot market, suppliers must pay upfront for their primary input (steel sheet) and then receive payment upon delivery for their customized product (fabricated metal). The result is that, like in many supply chain contexts (Kim and Shin (2012)), suppliers must finance the duration of their production.

2.2 Data description

Each order between the buyer and a supplier consists of many parts, where a part is a single unit, such as a single beam. These parts are delivered to the downstream customer in multiple lots, where a lot is usually a single truckload of parts. Our data has highly detailed part, lot and order-level observations for all orders between the buyer and her suppliers in a three year window from 2019 to 2022. This includes the weight of each part in metric tons (MT), its description, the lot it was a part of, and the supplier who worked on the order to which it belonged.

Each part must go through the following stages: raw material purchase, start of physical production (the first cutting of steel by labor), welding, blasting, grinding, painting, end customer clearance, and dispatch. Within each order, we see the planned, revised and actual dates for each stage in the production process. Thus we see the delay at each stage, along with the total cumulative delay. Figure 1 summarizes the production timeline data structure, where we pool all stages in which physical transformation takes place (cutting to painting) into a "production" stage for simplicity.

An important concern is that production-stage delays may be correlated with quality, for instance if firms compromise on technical standards to meet deadlines. In our setting, each part are inspected by the buyer's quality inspectors, and flawed parts need to be reworked. We include this post-inspection rework time in the pooled production stage. For reputational reasons (per buyer interviews), quality standards are similarly stringent across end customers.

Beyond production timelines, we observe contract terms agreed upon by the parties, including percentage advance (and the associated conditions, which relate to raw material purchase), post-delivery payment terms (anywhere from 0 to 45 days post-delivery), and the supplier's responsibilities. Contract terms in a bilateral relationship rarely change except for ex-

Figure 1: Production timeline structure

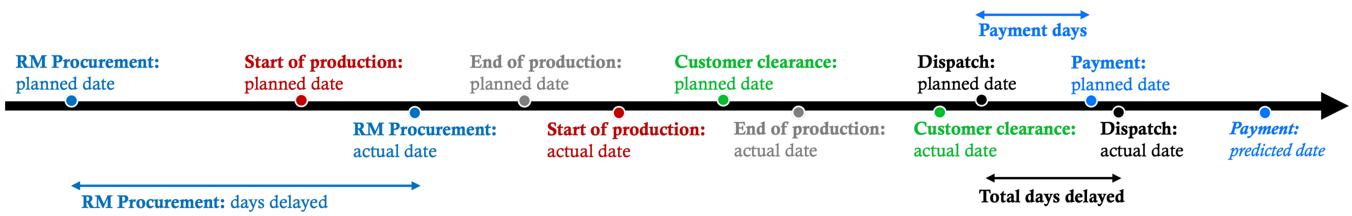

Notes: Figure shows production timeline data recorded by the buyer for a sample *lot* (or truckload) of fabricated metal. Planned dates by stage are set in advance, and actual dates of stage completion are verified by buyer visits to supplier production sites.

traordinary circumstances. We rely on one such circumstance for part of our empirical analysis i.e. a sudden jump in steel price post-COVID. The buyer also has employees termed relationship managers, whose main role is to spend time at the supplier site. They become part of the supplier's organization, directing and negotiating order progress ad-hoc. We observe their updates for each lot in each order on the buyer's portal. Finally, we also observe fixed characteristics of the supplier, notably the region and annual capacity.

2.3 Stylized facts

Leveraging that stagewise production information is recorded consistently by the buyer, we begin with stylized facts on the distribution of on-time delivery across suppliers.

Fact 1: Delays are concentrated in raw material investment, not physical production. We first decompose the total days delayed for each lot (truckload) into its constituent stages, and plot the stage-wise CDF.[7]

Figure 2 shows that raw material procurement has both the largest median delay (15 days) and the largest variance of all stages. Median procurement delays exceed those in starting physical production (cutting steel) due to insufficient labor or plant capacity; as well as all delays in completing production, including from worker absenteeism, weather disruptions, or rework of goods after a failed quality check. The relative importance of delays in raw

[7] While the stage-wise targets are not contractually binding, the intermediary uses them to determine when a supplier is 'behind' on a lot so that additional pressure can be applied.

Figure 2: Stagewide days delayed

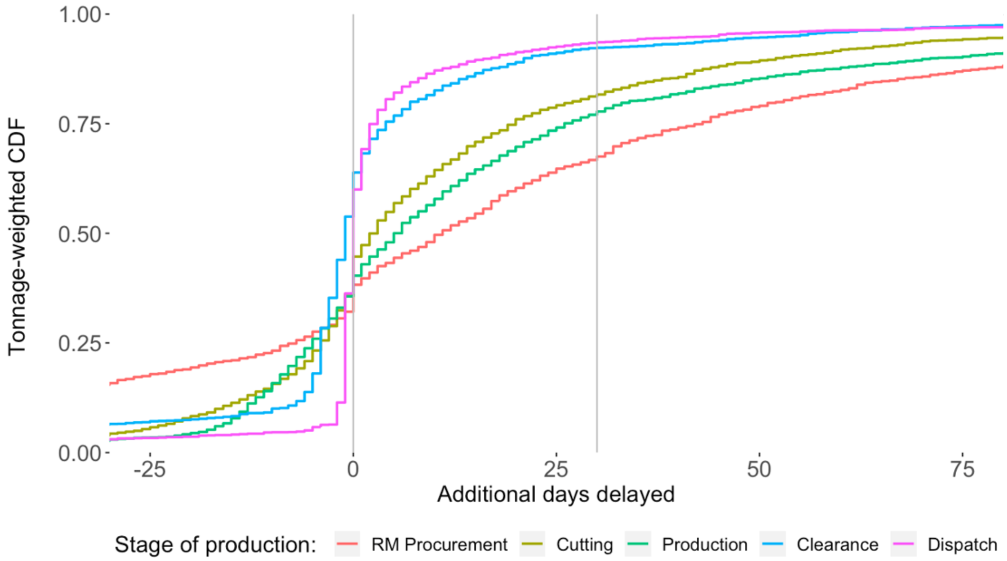

Notes: Figure shows CDF of additional delays by stage, across lots (truckloads) of output. Delays from stage s are the number of days behind schedule that stage s is completed (i.e. time between planned and actual completion dates), less the days behind schedule that stage $s-1$ is completed. RM Procurement refers to sourcing of steel, Cutting is the start of physical production (when steel is cut) sourcing; Production is the end of physical production; Clearance is when the downstream customer approves the lot for shipment, and Dispatch is the actual shipment date. All data are from production timelines for 2019-2022 pooled across suppliers.

material procurement is especially surprising given that the primary inputs – steel sheets – are homogeneous, sold on a spot market, produced domestically, and widely available near supplier locations.

What, then, inhibits on-time sourcing of steel? In our field interviews, suppliers attributed late sourcing of raw material to limited working capital. With limited cash-on-hand, managers could not always meet Just-In-Time buyers' strict timelines because they were left waiting for payments from past orders to buy steel for new ones. [8] Since other pattern could also explain the patterns in Figure 2 – for instance, we might attribute delays from late starts (for any reason) to raw material procurement simply because it is the first step – in Section 3 we

[8]Specifically, suppliers described procuring steel as an investment of cash-on-hand in inputs. Once raw material is procured, 3-6 weeks elapse before returns on that investment are realized through delivery of the final product. In the meantime cash is scarce (the shadow 'interest rate' facing the supplier remains elevated), which incentivizes suppliers to delay raw material procurement for new orders until payment arrives. Intuitively, the supplier operates like a hand-to-mouth household that cannot consume at the desired time; the difference is that a household's desired consumption timing is dictated by consumption-smoothing, whereas the firm's desired input sourcing timing is dictated by its buyer's Just-In-Time production schedule.

Figure 3: Cumulative delays by state at raw material and end-of-production stage

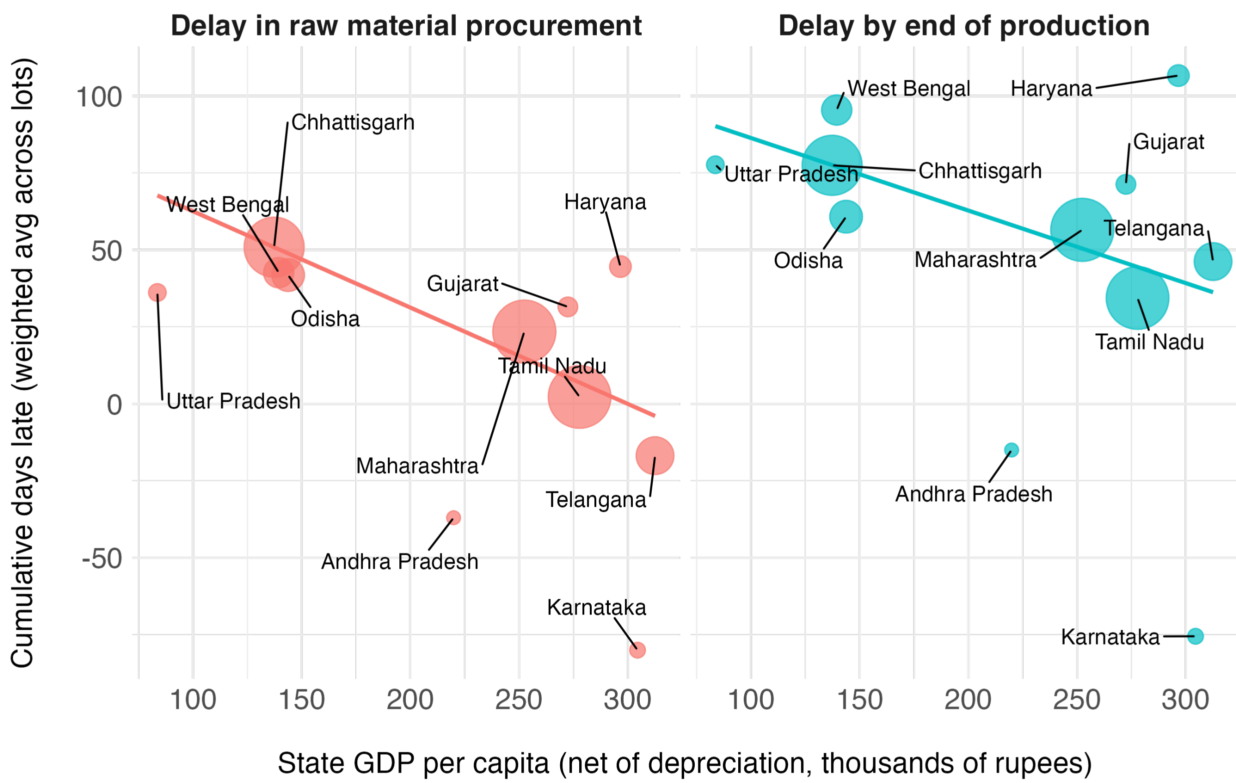

Notes: Figure shows association between state-level average delays (weighted average across all lots at suppliers in the state) and state-level GDP per capita. Left panel shows delay in sourcing raw material procurement, and right panel shows cumulative delays from sourcing raw material, starting orders, and completing orders. The overlaid lines give predicted values from an OLS regression that is weighted by total state-level production in metric tons. Average delays data from buyer production timelines for 2019-2022. Fiscal year 2022 GDP per capita data from the Reserve Bank of India.

explicitly test for the credit-constraint mechanism using exogenous variation in the timing at which previous orders are completed.

Fact 2: Consistent with credit constraints, poorer locations start orders late. We next show that delays in starting new orders are strongly associated with *low economic development*. Panel A of Figure 3 shows that on average, steel is procured on-time (1.5-2 months late) in richer (poorer) states. These patterns are largely driven by the poor performance of suppliers in the North's largest cluster in Bhilai, Chhattisgarh relative to two other clusters in the richer states of Maharashtra and Tamil Nadu. Panel B shows that, while this income gap persists but is smaller in magnitude by the end of production. In other words, poor locations' inferior on-time delivery rates are entirely due to late procurement of raw material.

Fact 2 is, to our knowledge, the first cross-firm evidence that timeliness is associated with levels of development (the within-India variation in per-capita income is roughly equivalent to that between Hungary and Rwanda). This positive association is not ex-ante obvious; we would expect the opposite, for instance, if timely production was more labor-intensive. Note that the association could either be due to a causal effect of location on timeliness, or by endogenous specialization of poor locations in time-insensitive orders.[9] In either case, Figure 3 suggests that poor locations more have a comparative disadvantage in on-time delivery.

The next two facts document how buyers involve themselves in supplier operations, especially in poor locations, to support timely supplier production.

Fact 3: The buyer physically monitors supplier operations, especially on late orders. In standard contract theory, increased monitoring of agent actions is one way to induce desired actions (Townsend (1979), Diamond (1984), Dye (1986)). In our context, monitoring takes the form of visits by relationship managers (employees of the buyer) to supplier sites, sometimes several times a week. During these visits, relationship managers observe the status of each truckload, update its internal timeline, verbally encourage managers to produce faster, and observe whether workers are allocated to the lot.[10] Thanks to our buyer's modern data systems, unlike in most informal trade contexts, we can directly observe the timing and frequency of these relationship-manager site visits.

Figure 4 shows that the frequency of raw material-related site visits before raw material is sourced, measured as the probability of a raw material-related timeline adjustment on a particular day, is substantially higher for orders on which raw material procurement is delayed. This pattern is particularly striking for orders delayed by over a month; on any given day there is a 10-15% chance that the buyer visits the plant. In other words, when sourcing for a particular part is behind schedule, the buyer frequently visits the supplier's factory until it is procured (which is why there is a mechanical drop at 0), inquires about its status, and sets a

[9] One concern is that Figure 3 may reflect spatial heterogeneity in end customer demand for on-time production by rather than in supplier's costs of JIT delivery. However, most construction projects physically take place in a different state from the supplier, and the end customers are themselves firms that operate across India (usually joint ventures between a foreign firm and an Indian conglomerate such as Tata and L&T). As a result, firms in different states should face similarly large on-time premia.

[10] This visibility into and management of, upstream production is an oft-cited reason for vertical integration; in our case, the intermediary achieves similar visibility into suppliers it does not own.

Figure 4: Probability of site visits before and after raw material is sourced

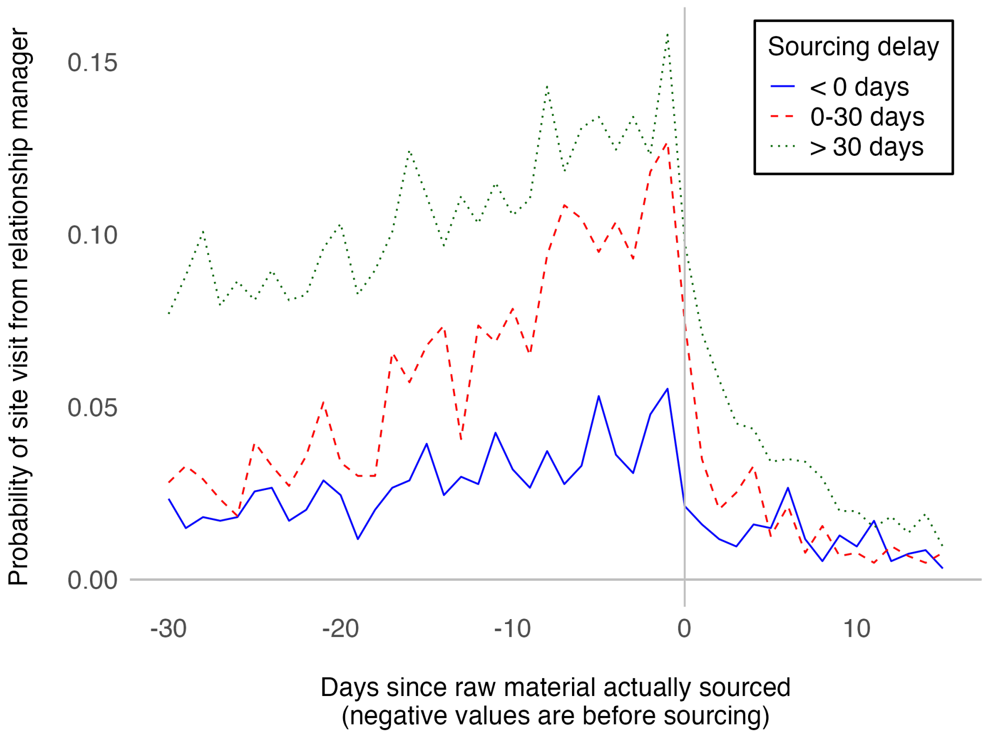

Notes: Plot shows the daily probability of a raw material-related site visit by a relationship manager as a function of time relative to when raw material is sourced. We divide lots (truckloads) into three groups: on-time sourcing (<0 days late), slightly late sourcing (0-30 days late), and late sourcing (>30 days late). Sourcing date and delay come from lot-level buyer production timelines. Relationship manager visits are inferred from timestamped production timeline updates in the buyer's internal IT system.

new expected delivery date.

The process shown in Figure 4 is essentially a form of management: relationship managers do not formally control production decisions, but they can check in, apply pressure, and agree on new 'expected' delivery dates with suppliers. Figure 4 thus provides evidence that relationship managers 'blur' firm boundaries monitoring resource allocation within supplier plants. In Section 4 we will show more formally that these management efforts are associated specifically with the relaxation of credit constraints, since they occur exactly when suppliers receive cash windfalls and therefore can make new investments.

Fact 4: the buyer invests in suppliers' raw material via two channels: incentives (advance payments conditional on sourcing) and control (direct sourcing of steel). Figures 2 and

Figure 5: Magnitude and form of buyer financing

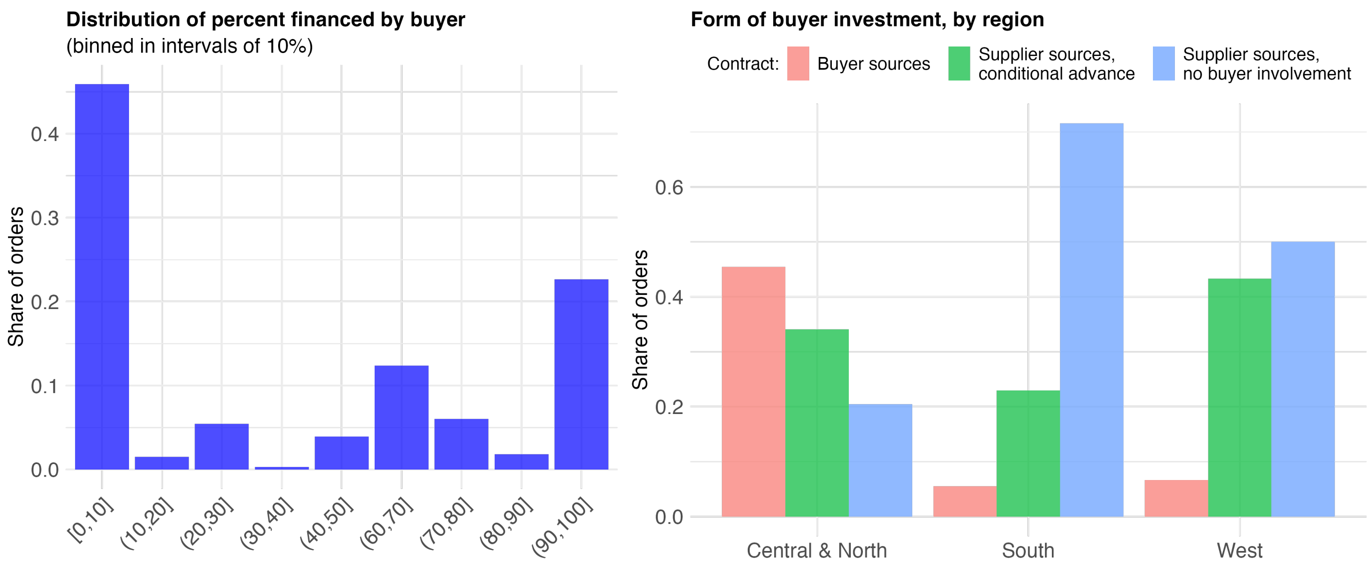

Notes: Left panel shows the distribution of input financing intensity across contracts. The percent financed is the share of order value that is paid to the supplier when steel for a lot is sourced (the remainder is paid after the completed lot is delivered to the downstream customer). This is specified in contract for conditional advances, and is 100 percent by default for orders with buyer-sourced raw material, and is 0 percent when the contract contains no input-related terms. The right panel shows the the share of contracts with each form of financing, separately by region. Central & North India is the poorest region, South India is the wealthiest region, and West India is in the middle. All data from buyer contract terms with individual suppliers for 2019-2022.

3 presented descriptive evidence from observed supplier performance consistent with short-term credit constraints that bind at the raw material stage, and Figure 4 showed that raw material decisions are heavily monitored through site visits. In Figure 5 below, we document two additional ways that buyers influence supplier sourcing decisions – incentives and control – using the formal contract terms agreed between suppliers and buyers. This is a relatively novel opportunity; in multinational supply chains, explicit contract terms are typically confidential.

First, as shown in the left panel of Figure 5, in 53 percent of contracts in 2019-2022, the supplier received raw material financing (either in cash or in kind) from the buyer. 30 percent of contracts were partially financed, with the buyer covering a mean of 60% of order value. The remainder were fully financed by the buyer, so the supplier requires no working capital investment. These contracts suggest that the buyer internalizes supplier credit constraints, as in other 'processing' trade contexts in which firms have high working capital needs, and thus that short-term credit constraints exist (Manova and Yu (2016)).[11]

[11]An alternative interpretation is that advances are offered as a reward to the highest-performing suppliers.

The right panel of 5 describes the specific mechanism that buyers use to finance supplier input purchases. Notably, suppliers never receive an unconditional advance (i.e. a loan), which is what one would expect if credit constraints were the only friction. Instead, we observe two forms of contracts. The first is a *conditional advance*, in which the supplier sources and owns the steel, but the buyer verifies that steel has been procured and then pays the advance before production takes place. Second, there are contracts with *buyer-sourced raw material* (23% of all contracts), in which the buyer procures steel from a mill, delivers it to the supplier, and has property rights over it. This is more costly, but allows the supplier to control buyer use of steel.

The common feature of both forms of involvement is that the buyer directly targets sourcing of raw material, a decision that in an arms-length transaction would be left to the supplier. The only variation is in the form of influence over sourcing: timely sourcing of steel is incentivized with a conditional advance and directly controlled with buyer sourcing. We collectively refer to these terms, along with the monitoring documented in Figure 4, as *partial integration* because they resemble the task-specific incentive contracts often deployed within integrated firms, but are being deployed by an outside buyer.

Variation in contract form across regions is broadly consistent with differences in income, and thus the severity of credit constraints. The right panel of Figure 5 shows that in South India (the richest region), under 30% of suppliers require financing, vs. 50% in the West and 80% in Central/North India (the poorest region). Central/North India is also the only region where the buyer favors direct control of inputs, rather than incentives to source inputs through conditional advances. These patterns suggest that buyer-provided financing is a response to some problem that is *worse* in poor locations with inferior on-time performance; if buyer financing was instead a reward or rent provided by buyers, we would expect to see larger advances in the South, where suppliers are larger, better-performing, and potentially have higher bargaining power. In Section 4 we will explicitly test if buyer involvement is a response to credit constraints by exploiting a shock to working capital needs.

2.4 Taking stock

The descriptive evidence presented above implies that, within Just-In-Time manufacturing supply chains, late deliveries are: (i) mainly due to late investment in supplier inputs, (ii) more severe in poor locations, (iii) addressed by buyers via actions that 'blur' firm boundaries, in-

However, in our sample, average lateness and percent financing are positively correlated, partially because poor locations are both later and receive more financing, suggesting that it is lower-type suppliers who receive the advances.

cluding monitoring of supplier input use, incentives to source inputs, and direct control control of supplier input sourcing.

In the next section we construct sharp empirical tests for two critical pieces of our story: that working capital constraints disrupt raw material sourcing, and that buyer involve themselves in – in other words, *partially integrate* – supplier operations exactly when credit conditions tighten, and thus risks of supplier delays worsen.

3 Empirical design

Our empirical analysis involves two distinct quasi-experimental research designs. First, to test if working capital constraints disrupt production in short (monthly) time intervals, we regress a supplier's investment in inputs (steel) for an order on predicted cash inflows from *past* orders, instrumenting for the timing of inflows using downstream postponements of common construction projects. We use the same instrument to document the buyer's short-run efforts to monitor input use via increased visits to supplier sites.

Second, to show how contracts adapt to credit frictions, we leverage a sudden increase in steel prices that raised working capital requirements. In a difference-in-differences-style design, we show that the steel price shock leads suppliers facing more severe credit constraints, as proxied by a supplier's pre-shock lateness and usage of buyer sourcing incentives (conditional advances), to sign contracts in which the *buyer* directly owns and sources steel on the supplier's behalf. Additional details beyond those in the text are provided in Appendix B.

3.1 IV test for credit constraints

To understand if insufficient working capital leads suppliers to deviate from production timelines (i.e. to delay), we test if supplier purchases of raw material respond to cash inflows from the completion of *other* orders. Our specification is analogous to classic tests of investment-cash flow sensitivity, with the focus on short-run investment in raw material.

Formally, denote order o and calendar month t. Each order o is completed by a supplier $s(o)$ for an end customer $c(o)$. There are multiple end customers because our partner ('the buyer') is itself an intermediary.

Within an order o, our raw data consists of many truckloads ('lots') i of goods, which are observed at many stages $g = (\text{Steel}, \text{Dispatch}, \text{Payment}, ...)$. Each lot i has a weight in metric tons (MT_i), a planned date to complete stage g ($Date_{ig}^{PLAN}$), and an actual date at which that stage is observably completed $Date_{ig}^{OBS}$. The lot-stage-wise delay is then:

$$Delay_{ig} = Date_{ig}^{OBS} - Date_{ig}^{PLAN} \tag{1}$$

We aggregate lot-level planned and observed outcomes to the order-month (o-t) level as follows:

$$\begin{aligned} Y_{otg}^{PLAN} &= \sum_{i \in o} MT_i \cdot \mathbb{1}[Date_{ig}^{PLAN} \in t] \\ Y_{otg}^{OBS} &= \sum_{i \in o} MT_i \cdot \mathbb{1}[Date_{ig}^{OBS} \in t] \\ &= \sum_{i \in o} MT_i \cdot \mathbb{1}[\left(Date_{ig}^{PLAN} + Delay_{ig}\right) \in t] \end{aligned} \tag{2}$$

and use g_{ot} in place of Y_{otg} for notational convenience, so that for instance total raw material sourced $Y_{ot,Steel}$ becomes $Steel_{ot}$. After aggregating, our estimating equation is:

$$\begin{aligned} \sinh^{-1}\left(Steel_{ot}^{\text{OBS}}\right) = \alpha_o + \gamma_t + \beta^{\text{OBS}} \sinh^{-1}\left(\sum_{\substack{o' \neq o \\ s(o')=s(o)}} Payment_{o't}^{\text{OBS}} \right) \\ + \beta^{\text{PLAN}} \sinh^{-1}\left(\sum_{\substack{o' \neq o \\ s(o')=s(o)}} Payment_{o't}^{\text{PLAN}} \right) + \delta \cdot \sinh^{-1}\left(Steel_{ot}^{\text{PLAN}}\right) + \epsilon_{et} \end{aligned} \tag{3}$$

where $Steel_{ot}^{OBS}$ is metric tons of steel actually sourced for order o in week t, and $Payment_{ot}^{OBS}$ is the quantity of order-o final output (in metric tons) for which payment is received in week t.[12] Payments are directly linked to past order completion: we take the observed date a lot is dispatched and add the number of credit days (often 0) specified in the written contract.[13] We exclude payments for the same order $o' = o$ on the RHS since those will mechanically occur around the time steel is sourced (the LHS). By including α_o and γ_t we examine only within-order changes over time. Since treatment variables vary by o, we cluster at the order level. Note that while the inverse hyperbolic sine operator allows for an elasticity interpretation, we show that estimates purely from the extensive margin, i.e. on any sourcing in response to any payment, are qualitatively similar.

[12] We measure payments in quantity terms because we do not observe prices per MT (these were kept confidential by our partner firm) and thus implicitly assume that prices are uncorrelated with timing of payment.

[13] If some share of payment is on delivery and the remainder is z days after delivery, we divide up total quantity accordingly.

The coefficient of interest β^{OBS} is the percent increase in metric tons sourced ($Steel_{ot}^{OBS}$) due to a log point increase in metric tons of production completed for other orders ($\sum_{o'} Payment_{o't}^{OBS}$). If firms can borrow freely at some fixed interest rate to finance steel purchases, the timing of past-order payment is irrelevant, and thus we should observe $\beta^{OBS} = 0$. If firms are working capital-constrained, then cash inflows should reduce the firm's (shadow) interest rate, implying $\beta^{OBS} > 0$.[14]

Importantly, in (3) we condition on planned values of both the independent and dependent variables, with coefficients β^{PLAN} and δ respectively. These regressors are constructed from original production schedules, before any postponements or modifications, as shown in equation (2). Controlling for plans explicitly gives our analysis a Just-In-Time production interpretation: conditional on planned sourcing, and planned cash inflows, is actual raw material investment responsive to actual cash inflows? In practical terms, these controls let us rule out any source of endogeneity that arises at the planning stage, including that firms intentionally sign orders in sequence to align cash flows or other scarce inputs.

Construction of downstream delays instrument The key issue in (3) is that, because the supplier can coordinate its day-to-day production, incoming payments may be endogenous to raw material sourcing intentions. For instance, a supplier might accelerate completion of order o' exactly so it has cash to invest in raw material for order o. To sharpen the causal interpretation of Equation (3), we next construct an instrument for planned payments $Payment_{o't}^{OBS}$. The instrument leverages variation in payment timing from downstream construction timeline changes, e.g. due to labor shortages, power outages, weather disruptions, local COVID lockdowns, or regulatory delays. Downstream timeline changes are unanticipated and exogenous to the supplier, which aids in identification.

To isolate the downstream component of delays, we leverage that multiple orders o and o', for different suppliers s and s', are often part of the same large downstream construction project c, and thus subject to common postponement shocks. This means we can use information on lots at suppliers s' to construct a leave-one-out estimator for cash inflows at s. To understand the intuition, suppose that a large wind farm requires 200 wind turbine towers, so that two suppliers s and s' are asked to complete 100 towers each. If a 45-day postponement of wind farm construction occurs, both suppliers will be equally affected, so the average 45-day

[14] Note that we do not necessarily require payment inflows to be random. The analogous result in the household finance literature, of a positive correlation between the timing of household income and consumption, is interpreted as a rejection of the permanent income hypothesis even without instrumenting for income.

observed delivery delay on order o' (completed by s') should estimate the delay on order o (completed by s).

To formalize this logic, we apply the definition of a delay from (1), so that our estimate of customer-driven delays for a lot i is

$$PredDelay_{ig} = \frac{\sum_{o':\substack{s(o')\neq s(o)\\ c(o')=c(o)}} \sum_{i'\in o'} MT_i \cdot Delay_{i'g}}{\sum_{o':\substack{s(o')\neq s(o)\\ c(o')=c(o)}} \sum_{i'\in o'} MT_i} \tag{4}$$

i.e. the weighted-average delay across all lots i' in other orders o' at other suppliers s' but for the same downstream construction project c. We then aggregate to the order-month level using $PredDelay_{ig}$ in place of the actual delay $Delay_{ig}$. Focusing on $g = Payment$, we define the instrument:

$$Payment_{ot}^{INST} = Y_{ot,Payment}^{INST} = \sum_{i\in o} MT_i \cdot \mathbb{1}[\left(Date_{ig}^{PLAN} + PredDelay_{ig}\right) \in t] \tag{5}$$

and use $sinh^{-1}(Payment_{ot}^{INST})$ as an instrument for observed $sinh^{-1}(Payment_{ot}^{OBS})$ in (3). As shown by comparing (2) to (5), the instrument is the counterfactual path of payments if all deviations from production plans were exactly equal to the leave-out average at other suppliers on the same contract.

Effect on monitoring of supplier If credit frictions lead to underinvestment in inputs, how do buyers respond to ensure on-time delivery? One margin is increased monitoring; recall we showed in Fact 3 that relationship managers regularly visit supplier sites, and these visits increase in frequency when raw material procurement is delayed. If monitoring affects supplier sourcing behavior – i.e. if site visits let buyers 'manage' input use by pressuring suppliers to invest in the appropriate order – then we should also observe increases in site visits for an order o exactly when payments from other orders o' arrive. In other words, the supplier's sourcing behavior should be monitored exactly when it has cash to invest, which is when the last order has ended.

We construct two measures of monitoring from the buyer's internal data system: the number of unique relationship managers that visit the supplier $NumExp_{ot}$, and the number of unique updates to production statuses $NumUpdates_{ot}$ made by all relationship managers during those visits (visits are order-specific and timestamped so they can be aggregated to the order-month level). We then re-estimate the OLS and instrumental variable versions of

Figure 6: Leave-out delays at other suppliers on common projects predict order-level delays

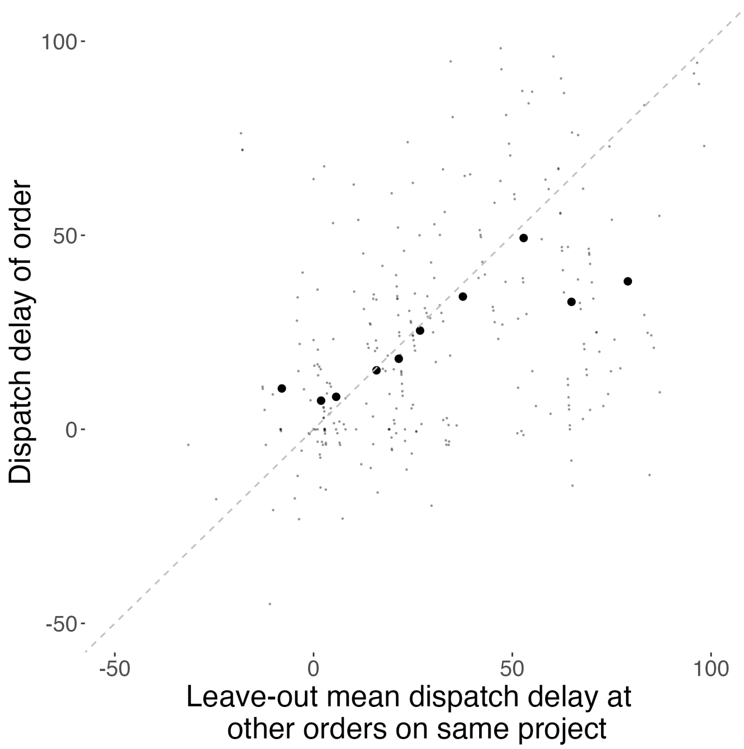

Notes: Figure shows relationship between delays (by dispatch) on an order with a particular supplier, and the average delays on orders at other suppliers that are part of the same downstream construction project. Scatterplot is shown in gray, binned means are shown in black, and the gray dashed line is the 45-degree line. Order-level delays are determined by the difference of planned and actual dispatch dates, averaged across lots within each order. All data from 2019-2022 buyer production timeline; we restrict to orders in which a larger contract contains multiple orders with different suppliers.

equation (3) with these outcomes in place of $Steel_{ot}^{OBS}$. With $NumExp_{ot}$ and $NumUpdates_{ot}$ as outcomes, $\beta^{OBS} > 0$ implies that monitoring of supplier's order-o production increases in the exact month when other orders o' are complete.[15]

3.2 DID changing of contracts

In Figure 5, we showed that the buyers engage in two forms of partial integration – incentives and control – that directly target supplier input sourcing. In this section, we examine the responsiveness of the form of contracting to working capital constraints. In particular we ask

[15]Note that, even if we find effects on monitoring, we do not know the exact stage of production that the relationship manager supervises. For instance, during a single visit the relationship manager may ask the buyer both to accelerate sourcing for lot i and welding of lot i'.

if buyers take on additional control rights when credit constraints worsen in order to ensure that suppliers start (and thus finish) their orders on-time. Such a response would occur if additional control overrides a supplier's increased temptation to misallocate working capital due to an increased shadow value of capital. For instance, if suppliers are tempted to either delay procuring inputs for an order when cash-on-hand is low or divert procured inputs to other uses, the buyer can procure and physically deliver inputs to the supplier site to ensure that production begins.[16]

We use a difference-in-differences design to construct a shifter of credit constraints. The natural experiment that we exploit is a rapid doubling in steel prices, from 40,000 to 80,000 Indian rupees, that occurred in fall 2020 (see Figure 7). Our partner firm and suppliers described the shock as sudden and unanticipated; according to local press, the shock was due to rapid rebound of industrial activity after the initial COVID lockdowns in India, combined with pandemic-related trade restrictions on imported steel.[17]

Steel prices rises affect working capital requirements (and thus, may lead credit constraints to bind) because they increase the amount of investment necessary to start an order. For instance, if an order with 1mn in value-added required Rs 2mn in input investment before the shock, it requires Rs 4mn after the shock. Our identification strategy relies on the idea that suppliers are differentially affected by the shock because they have different levels of pre-existing access to working capital finance. An unconstrained supplier faces a perfectly elastic credit supply curve at some interest rate r; for this supplier, the required investment (in rupees) has doubled, but remains feasible. In contrast, a constrained supplier has limited cash-on-hand; the increase in steel prices therefore could render the investment entirely infeasible without buyer support.

To capture this intuition, we construct three proxy measures of a supplier's pre-shock credit constraint. The first two measures leverages the fact that – to ensure that it does not finance inframarginal suppliers with substantial existing credit access – the buyer charges *interest* on conditional advances that is above the market rate.[18] As a result, conditional advances are not a fixed payment or reward offered to high-quality suppliers; consistent with this claim, suppliers in poor locations and with smaller order sizes have higher take-up of conditional advances. We therefore use *whether a supplier ever receives a conditional advance* as a proxy for the

[16]Ownership of supplier inputs is a well-known strategy for securing supplies of critical inputs during times of uncertainty; for instance, consider major automakers' plans to buy lithium mines to secure battery supplies, or Henry Ford's decision to operate rubber mines in Brazil for tires.

[17]To do: show that the shock tracked global prices and exceeded futures prices.

[18]For instance, if a supplier buys steel on January 1, receives a conditional advance of 100 rupees on January 2 (when the buyer verifies steel has arrived), and then delivers the final fabricated output on February 1, the buyer deducts $\frac{30}{365} \cdot r$ rupees from its post-delivery payment, where r is some interest rate.

existence of supplier credit constraints, and the *average share of an order that is conditionally advanced* as a measure of their severity.

The third measure relies on the assumption that suppliers that delay raw material purchases before the shock are exactly the credit-constrained suppliers. Under this assumption, we can use pre-shock average days delayed as a measure of working capital constraints. This final measure is of course imperfect because suppliers with higher vs. lower delays, and in richer vs. poorer regions, could differ for reasons entirely separate from credit constraints, so our preferred specifications use the first two conditional advance-based measures.

For our preferred specification, the unit of observation is a lot i, the panel unit is a supplier s, and the time unit is a quarter t (pooled over months for power). We estimate:

$$DirectSource_{it} = \delta_t + \alpha_s(i) + \beta^{DID}\left(HasAdvance_{s(i)} \cdot Post_t\right) + \epsilon_{it} \tag{6}$$

where $DirectSource_{lt}$ takes value 1 if the buyer sources steel for lot i, and 0 otherwise.[19] $HasAdvance_{s(i)}$ takes value 1 if the supplier of lot i *ever* has received a conditional advance for raw material from the buyer before the shock; and 0 if no pre-shock contract includes an advance (we drop suppliers with no pre-shock contracts). $Post_t$ takes value 1 for quarters after the steel shock begins (i.e. after October 2020). δ_t is a time (calendar quarter) fixed effects and the α_s is a supplier fixed effect. Finally, we cluster at the supplier level. The coefficient of interest β^{DID} then captures the additional post-shock increase in the probability of buyer sourcing of raw material for credit-constrained suppliers *relative* to unconstrained suppliers.

We estimate several alternative versions of equation (6). First, we replace $Post_t$ with a full set of calendar time indicators to trace out the dynamic effects in an event study. Second, we re-estimate (6) with a control for $HasAdvance_{s(i)}$ in place of the supplier fixed effect $\alpha_s(i)$. The reason is that suppliers rarely produce in all months, so (6) is a highly unbalanced panel. With a single $HasAdvance_{s(i)}$ control, we can average production within the unconstrained-supplier and constrained-supplier groups, which improves balance. A secondary benefit of doing so is power, as there are only 30 unique suppliers in the panel. Third, we estimate a regression in which we restrict to post-treatment period and drop the time fixed effect δ_t; because there is no direct sourcing in the pre-period (a fact we show in Section 4), this cross-sectional regression is mathematically equivalent to the pooled two-group DID regression. In this pooled regression, we explore the alternative proxies for credit constraints discussed above.

[19] To determine this, we look at the contract terms for the order o to which lot i belongs. Estimating the regression at the order o level yields qualitatively similar results, though the timing is slightly different when we use order signing dates instead of production dates.

The key identification assumption in all regressions is parallel trends: in absence of the steel price shock, the probability that credit-constrained firms use buyer-sourced raw material would have evolved similarly to the probability that unconstrained firms do so. This assumption would be violated if the buyer already planned to start sourcing raw material for the least well-capitalized firms in absence of the steel price shock. One plausible version of this story is that logistical constraints could have made it practically difficult for some suppliers to source raw material during the COVID shock. By analyzing the buyer's internal records (email communications and interviews with managers), we verified that the buyer had no plans to procure steel itself until steel prices, and thus working capital requirements, rose in late 2020. As a result, in absence of the shock, the counterfactual level of buyer sourcing, is zero for both groups, and thus the counterfactual trend is indeed parallel.

4 Results

We now describe our three sets of quasi-experimental results, summarized below:

1. **(The friction) Suppliers are short-term credit-constrained**. Suppliers source raw material for new orders exactly when payments from other orders arrive
2. **(Contractual solution I) Buyers monitor suppliers at the start of new orders.** When payments arrive from other orders (and thus a new order can begin), buyer site visits to the supplier increase.
3. **(Contractual solution II) Buyers source inputs for their most credit-constrained suppliers.** When working capital requirements rise due to a steel price shock, buyers differentially begin sourcing steel for suppliers that are more credit-constrained ex-ante.

4.1 Suppliers are short-term credit-constrained

Table 1 shows that conditional on planned sourcing for an order o and planned payments for all other orders o', actual sourcing of steel covaries with actual payments for other orders. Column 1 gives the results of OLS estimation of equation (3). The estimated elasticity of 0.04 is small but significant[20]: if the supplier was fully dependent on other-order payments we would expect an elasticity of 1. Nevertheless, applying the logic of standard investment-cash flow sensitivity

[20]Since both the LHS and RHS are in inverse hyperbolic sine terms, we can apply an elasticity interpretation.

Table 1: Effects of other-order payments on raw material sourcing for new orders

Dependent Variable:	**IHS(metric tons of raw material sourced)**			
IV stages	OLS	RF	1st	2nd
Model:	(1)	(2)	(3)	(4)
Variables				
RM sourced (plan, IHS)	0.419***	0.419***	0.025*	0.418***
	(0.024)	(0.024)	(0.015)	(0.024)
Other-order payment (plan, IHS)	0.004	0.017***	0.444***	-0.016
	(0.006)	(0.006)	(0.013)	(0.013)
Other-order payment (obs., IHS)	0.039***			0.074***
	(0.007)			(0.024)
Other-order payment (inst., IHS)		0.020***	0.269***	
		(0.007)	(0.011)	
Fixed-effects				
Order (within supplier)	Yes	Yes	Yes	Yes
Month	Yes	Yes	Yes	Yes
Fit statistics				
F-test (1st stage)			1,653.2	
R^2	0.23405	0.23192	0.49264	0.23156
Dependent variable mean	0.15894	0.15894	0.68989	0.15894
Observations	19,140	19,140	19,140	19,140

Clustered (Order (within supplier)) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Table shows effects of payments for other orders $o' \neq o$ on raw material sourcing for order o. Column 1 provides OLS estimates using observed payments ("obs"), conditional on planned sourcing and payment. Columns 2-4 give IV estimates (the reduced-form, first, and second stages) when we instrument for observed payment timing using predicted timing (assuming delays follow their leave-out average). All outcomes and regressors are in inverse hyperbolic sine (IHS) form to allow for percent-change interpretations. All regressions are at the order-month level. All data are from 2019-2022 production timelines from the buyer.

tests, the positive elasticity implies that suppliers are credit-constrained. As a sanity check, planned sourcing is strongly predictive of actual sourcing, with an elasticity of 0.41.

The key challenge in interpreting column 1 is that timing of other-order delivery (and thus of payment) is partially under the supplier's control. If the supplier is not cash-constrained but is simply completing orders in sequence, we might falsely estimate a positive elasticity 'of cash'. Our preferred strategy is thus to instrument for observed payments with the predicted payment instrument, which assumes payment delays relative to plan are exactly equal to their leave-out average at other suppliers on common downstream projects. As shown in column 2, the reduced-form elasticity of new steel sourcing to this predicted payment instrument is 0.02.

The IV specification results are shown in column 4. We estimate an elasticity of 0.074 with respect to (instrumented) realized payments, which is positive, significant, and larger than the OLS estimate. Column 3 shows that the instrument is strong (i.e. relevant) with an F-statistic of over 1600. The instrument is valid under the the exclusion restriction that predicted payments only shift new sourcing through actual payments. Under that assumption, the positive elasticity implies that credit constraints indeed disrupt Just-In-Time delivery, because some new orders are started only when working capital is 'returned' to the supplier as liquid cash instead of being locked as a receivable. The implication is that, even conditional on its plan, it would have been less likely to produce in a month t had payment not occurred. Stated formally, positive elasticity of raw material investment to (same-month) payments implies that the (shadow) price of working capital falls when payments arrive.

One potential interpretation is that firms are credit-constrained[21] (Banerjee and Duflo 2014) and thus have incentives to delay starting orders until payments arrive, interfering with stringent Just-In-Time delivery requirements. To examine this mechanism, we next estimate heterogeneous elasticities based on the intensity of credit constraints. If working capital is the relevant mechanism, then ex-ante constrained firms should have larger elasticities of sourcing to payment.

Table 2 shows heterogeneity in our 2SLS estimates by supplier-level proxies for credit constraints: region, based on Figure 3; and receipt of financing from the buyer (whether through conditional advances or direct sourcing), based on Figure 5.[22] Buyer financing is a valid proxy because our buyer aims to finance suppliers who lack pre-existing cash or bank lines. Columns 1 and 2 show that while the sourcing-payment elasticity is positive (0.091) and significant for orders with financing support from the buyer, the elasticity is smaller (0.038) and not signif-

[21]Formally, they at least face an upward-sloping credit supply curve instead of a constant interest rate.

[22]We are only using payments linked to order completion, not to advances, so there should not be any mechanical correlation between the timing of procurement and payment for firms receiving advances.

Table 2: Heterogeneity in effect of payments on sourcing by proxies for firm credit constraints

Dependent Variable:	RM sourced (IHS)				
Advance	Yes	No			
Region			Central & North	South	West
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Other-order payment (obs., IHS)	0.091**	0.038	0.125***	0.055	0.016
	(0.043)	(0.026)	(0.042)	(0.034)	(0.058)
RM sourced (plan, IHS)	0.373***	0.466***	0.320***	0.480***	0.480***
	(0.036)	(0.035)	(0.037)	(0.040)	(0.049)
Other-order payment (plan, IHS)	-0.025	-0.0004	-0.027	-0.031	0.022
	(0.020)	(0.018)	(0.021)	(0.022)	(0.026)
Fixed-effects					
Order (within supplier)	Yes	Yes	Yes	Yes	Yes
Month	Yes	Yes	Yes	Yes	Yes
Fit statistics					
F-test (1st stage)	610.05	905.59	486.06	849.04	334.52
R^2	0.20099	0.26681	0.16942	0.28064	0.31027
Dependent variable mean	0.14950	0.16451	0.16628	0.14768	0.15990
Observations	8,547	9,108	7,392	6,303	4,851

Clustered (Order (within supplier)) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Table shows IV second-stage effects of (instrumented) payments for other orders on sourcing for each order. Columns 1 and 2 respectively subset to orders that contain and do not contain an advance, respectively. Columns 3, 4 and 5 subset to orders with suppliers in each region: Central & North (the weakest), and the West (intermediate).

icant at conventional levels in orders without any financing support. Columns 3-5 show that responsiveness to cash is economically and significantly significant (magnitude 0.125) in the Central/North region than in the South and West regions, which have fewer input investment delays (recall Figure 3) and larger firm sizes, which enable better credit access. Overall the results suggest that, using conventional proxies, constrained firms are more responsive to cash inflows.

We briefly discuss identification concerns and related robustness checks. First, the magnitude of the estimated elasticity is difficult to interpret, because we only see orders from our buyer in the data. If suppliers disproportionately redirect cash from orders for one buyer to orders for other buyers (rather than to the buyer we work with) our elasticity will be biased.[23] A transparent alternative is to estimate an extensive margin-only regression, which we do in Appendix Table 1. We find qualitatively similar estimates: receiving any payment (vs. no payment) in month t for a past order predicts nonzero new sourcing in that month.

Second, the key threat to exogeneity is that our formula instrument for other-order payments is not fully exogenous. Realized payment timing is a function of both (i) planned payment timing and (ii) customer-driven deviations from plans. While we have argued that customer-driven deviations are as-if random, planned payments are not. For instance suppliers might accept orders in a sequence that (assuming no delays) allows them to immediately reinvest payments from one orders into the next's raw material. This means certain order-months may be predicted to receive more payments (i.e. be "treated" under any realization of plans).

The solution is to *condition explicitly on planned payment timing and planned sourcing* in all regressions. Unsurprisingly, both are highly predictive of actual sourcing. As shown in the reduced-form regression (Column 2 of 1), the elasticity of actual to predicted sourcing is 0.42 (vs. equal to 1 if sourcing always followed plans), and the elasticity of actual sourcing to planned payments is 0.017, which is similar in magnitude to the elasticity with respect to the instrument. Importantly, the results we present in 1 are conditional on plans.[24] Thus we show that suppliers source more raw material for an order *relative to plan* exactly when they receive more payments from other orders *than originally planned* in that month.

Third, the exclusion restriction is violated if predicted other-order payments shift sourcing through some mechanism other than actual payments. One plausible mechanism is that, because physical completion and payment occur at the same time for some orders, we may be detecting the effect of other orders being physically complete (enabling labor to start on new orders) rather than the effect of cash inflows. However, since payment and physical completion dates are largely in the same month, we are unable to separate relaxations of capacity and cash constraints. A weaker statement consistent with the results above is thus that *some* fixed factor exists that disrupts Just-In-Time delivery. As we argued in the methodology section, given that (i) buyer input financing is common in observed contracts, and (ii) credit-constrained firms have larger elasticities of sourcing with respect to (instrumented) past-order completion, it is

[23] The direction is unclear.

[24] Conditioning on planned payments means, in the Borusyak and Hull (2020) framework, that 'expected other-order payments' for each order-month are equal to planned payments. While not formally a statistical expectation, based on interviews with suppliers, the planned dates are used directly to form their production plans and judge if they can take on new orders without exceeding capacity. An alternative would be for us to permute the leave-out delay shocks from Figure **??** at some level and control for the empirical expectation. However, unlike in an RCT, it is unclear if these shocks are i.i.d. within a sufficiently broad cluster to allow for permutation (they are likely distributed differently across buyers, time periods, etc.) because the data-generating process is unclear.

highly plausible that short-term credit is one such fixed factor.

4.2 Buyers increase monitoring when suppliers receive payments

A natural implication of scarce working capital (or any other input) is that the supplier should be monitored once production starts, so that the buyer b can verify that cash (or workers) are actually allocated towards b's own order. In Table 3 we show that monitoring increases in response to payments from other end customers. To measure monitoring, we use administrative data from the buyer (i.e. our partner firm) on visits to supplier plant sites by relationship managers. These visits are frequent and involve conversations with plant managers about the stage-wise status of specific parts within an order. They are also entirely at the discretion of the supplier. We thus use our IV strategy to estimate effects on the number of updates made during such visits and the number of unique managers who visit, both of which are order-specific measures. The instrument and endogenous regressors (predicted and actual payments) remain the same as in Table 1.

Table 3 presents the IV results. The second-stage coefficients of interest are in columns 2 and 5. In column 2, a log point increase in other-order payments increase the number of unique updates by 6.81, and the number of unique managers visiting in the month by 0.05. Dividing by the outcome means give elasticities of 0.48 and 0.38, respectively. Note that these elasticities are substantially larger than the effects on actual raw material sourced. This result implies that monitoring (linked to order o) occurs even in cases where new sourcing does not happen, i.e. monitoring might not be immediately successful.

As a validation check on the data, note that in columns 2 and 5 of Table 3, there are large elasticities of monitoring with respect to *planned* sourcing (of around 1). These imply that monitoring occurs in the month of planned production – i.e. that relationship managers make more frequent visits to plants exactly when an order is supposed to begin.

Table 3: Effects of other-order payments on monitoring of new orders

Dependent Variables:	N of mgr updates			N of unique mgrs visiting plant		
IV stages	1st	2nd	RF	1st	2nd	RF
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Other-order payment (inst.)	0.269***		1.83**	0.269***		0.014***
	(0.011)		(0.716)	(0.011)		(0.003)
RM sourced (plan)	0.025*	14.6***	14.8***	0.025*	0.110***	0.111***
	(0.015)	(3.53)	(3.53)	(0.015)	(0.008)	(0.008)
Other-order payment (plan)	0.444***	-2.46	0.561	0.444***	-0.007	0.015***
	(0.013)	(1.54)	(0.856)	(0.013)	(0.007)	(0.003)
Other-order payment (obs.)		6.81**			0.051***	
		(2.66)			(0.013)	
Fixed-effects						
Order (within supplier)	Yes	Yes	Yes	Yes	Yes	Yes
Month	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
F-test (1st stage)	1,653.2			1,653.2		
R^2	0.49264	0.14125	0.14128	0.49264	0.21970	0.22370
Dependent variable mean	0.68989	13.946	13.946	0.68989	0.13589	0.13589
Observations	19,140	19,140	19,140	19,140	19,140	19,140

Clustered (Order (within supplier)) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Table shows IV estimates of payments for other orders (at the same supplier) on the total number of relationship manager updates for an order. Data on instrumented, planned, and observed payment timing come from production timelines; data on manager updates comes from internal buyer IT systems that track timeline changes. All outcomes and regressors are in inverse hyperbolic sine format to facilitate a percent change interpretation.

4.3 Buyers source inputs for the most-constrained suppliers (switch from incentives to control) when credit requirements rise

If a supplier is credit-constrained, but the buyer is a well-capitalized multinational firm, then one would expect the parties to adapt their contracts in response to worsening credit constraints. Adaptation could take multiple forms; recall that Figures 4 and 5 show that buyers engage in multiple forms of *partial integration* to influence supplier input use. In this section we show that exogenous increases in credit needs change the form of adaptation: rather than simply incentivizing supplier behavior through site visits and conditional contracts, the buyer takes on direct responsibility for sourcing raw material for the supplier, while leaving the supplier responsible for physical transformation and eventual delivery.

As a shifter of credit demand, we use the rapid increase in global steel prices in late 2020. Before we turn to the formal difference-in-difference analysis described in equation (6), in Figure 7 we document that the rise in global steel prices is associated with substitution between contract forms. Specifically, the buyer substitutes from conditional advances (essentially loans from buyers to suppliers) to direct sourcing of suppliers' inputs for the order.

What are the costs and benefits of direct buyer sourcing relative to conditional advances (or no support)? There are two primary costs: (i) first, buyer sourcing implies full investment by the buyer, whereas conditional advances in the data usually *partially* cover the value of the raw material (though this is not a hard constraint); and (ii) buyer sourcing places greater logistical demands on the buyer, who must have the local presence to procure and send steel to the supplier's site. The benefits are in the form of more control: buyer sourcing implies (iii) buyer ownership of steel, so the raw material cannot be resold or redirected towards another order; and (iv) buyer control over input arrival time, so the buyer can directly ensure that raw material is at the site well in advance. These benefits (iii) and (iv) imply buyer sourcing should be viewed as the most direct form of 'partial integration' – control of supplier input decisions – applying only to the sourcing stage.

To understand *which suppliers* switch to buyer sourcing, we estimate an event-study version of equation (6), interacting quarters since the shock with an indicator for pre-shock credit constraints. As a proxy for pre-shock constraints, we use whether a supplier was receiving conditional advances (vs. no support) before the shock.

The results are shown in Figure 8: *relative to unconstrained firms, constrained firms are 20-30% more likely to receive direct buyer sourcing after the shock.* These results hold both in a pooled DID design comparing all constrained to all unconstrained suppliers (left panel) and with supplier fixed effects (right panel). The dynamic path of effects is consistent with the time series in Figure 7, in which buyer sourcing was suddenly offered. The zero effects in the pre-period are because buyer sourcing was not offered.

Overall these results imply multiple stages of buyer involvement in sourcing as credit needs rise. Initially, the buyer incentivizes sourcing of raw material via a conditional advance, and otherwise provides nothing. As credit needs rise further, conditional advances are no longer sufficient for the most credit-constrained suppliers, so in the most severe cases the buyer takes the costly step of directly sourcing inputs for those suppliers. And since the most credit-constrained suppliers would have received advances before the shock, we should

Figure 7: Time series of steel price and contract terms

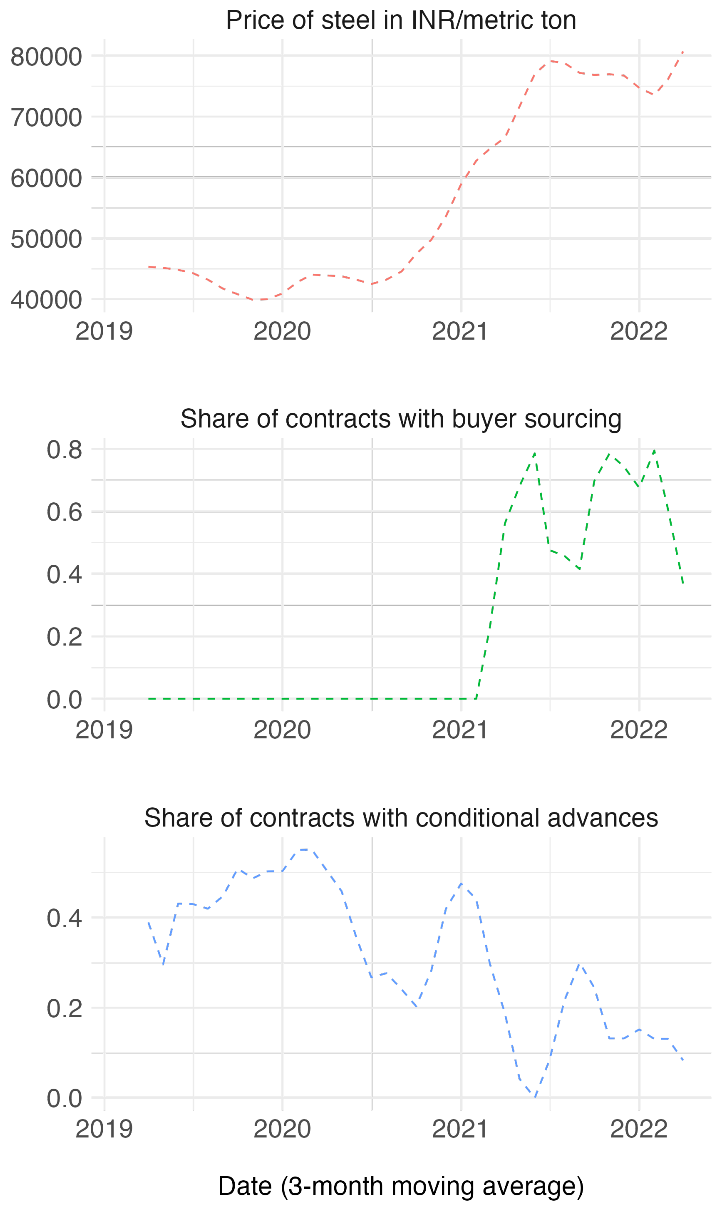

Figure shows three different time series from 2019-2022, all as three-month moving averages calculated from monthly data. Top panel shows Indian spot market steel prices, obtained from the Ministry of Steel, Government of India. Middle panel shows share of contracts (by date of contract signing) in which the buyer owns and controls suppliers' inputs. Bottom panel shows share of contracts in which buyer offers a conditional advance (partial payment once supplier sources its inputs, rather than after final delivery). Data on contract terms are from the buyer's administrative records.

Figure 8: Differential effect of steel price shock on buyer sourcing for credit-constrained firms

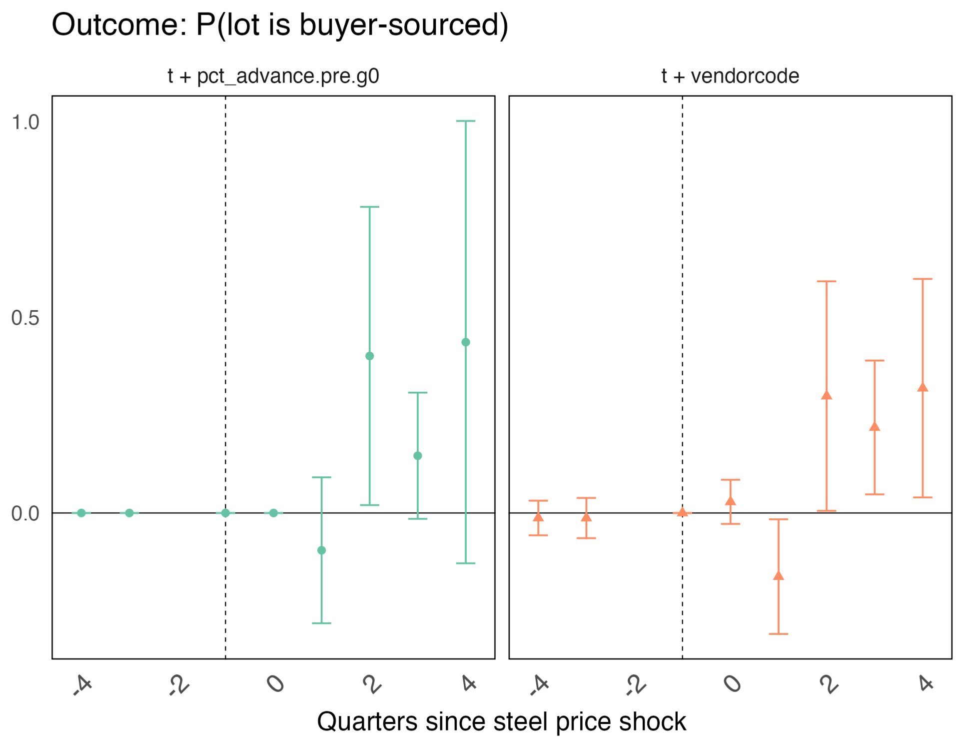

Notes: Figure shows event-study effects of supplier-level exposure to an input price shock on the probability of using buyer-sourced raw material. A supplier is exposed if it ever used conditional advances from the buyer *before* the price shock. Regressions are estimated at the supplier-quarter level. In the left panel, the regression includes supplier and quarter FE. In the right panel, the regression includes quarter FE and an indicator for being exposed (so there is a single treated and untreated group).

observe substitution from advances to direct sourcing, as is clear in Figure 7.

Table 4: Predictors of buyer sourcing: order-level regression (post-steel price shock)

Dependent Variable:	=1 if buyer sources raw material					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Constant	0.0294	0.0411	0.1117^{**}	0.3458^{**}	0.5806^{***}	0.1505
	(0.0293)	(0.0355)	(0.0432)	(0.1493)	(0.1875)	(0.1699)
=1 if any advance: pre-pd avg	0.2130^{***}			0.2580^{***}		
	(0.0608)			(0.0774)		
Pct advance, pre-pd avg		0.0048^{***}			0.0065^{***}	
		(0.0014)			(0.0017)	
Days delayed: pre-pd avg			0.0021^{*}			0.0021^{*}
			(0.0012)			(0.0012)
Log(mean order size)				-0.0851^{**}	-0.1438^{***}	-0.0092
				(0.0425)	(0.0520)	(0.0361)
Fit statistics						
Observations	100	100	98	100	100	98
Dependent variable mean	0.17000	0.17000	0.17347	0.17000	0.17000	0.17347
R^2	0.07216	0.12140	0.03880	0.09168	0.17131	0.03902

Heteroskedasticity-robust standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Table shows estimates from order-level regressions of an indicator for buyer control of raw material (rather than either incentives for raw material sourcing through conditional advances, or no intervention) on several proxies for credit constraints using data from the post-input price shock period. Information on buyer sourcing in the post-period and conditional advances in the pre-period come from individual order terms. Information on pre-period days delayed and order size, both averaged by supplier, come from lot-level production timelines.

Finally, to understand the aggregate magnitudes, in Table 4 we regress an indicator for buyer sourcing on various alternative measures of credit constraints using only post-period data. Columns (1)-(3) show that all three proxies for pre-period credit constraints – the indicator for a conditional advance, the size of the conditional advance (share of order), and average days delayed – are predictive of buyer sourcing of raw material in the post-period. Columns (4)-(6) show that even conditional on mean order size – our proxy for firm size – less well-capitalized firms receive buyer sourcing. There are 100 observations in most specifications because this is the number of unique *orders* for which we see both contract terms and pre-period supplier information.

Particularly of note are the effects of the average percent advance measure in columns 2 and 5, which is easily interpretable as the extent of the firm's pre-shock credit constraint (from 0 to 100%). We find that a 10 percentage point higher conditional advance share (before the shock) increases the probability of switching to buyer sourcing after the shock by 4.8 percentage points (28% of the mean). In summary, the buyer only directly owns and controls the inputs of its credit-constrained suppliers.

5 A framework of partial integration

We have shown via separate natural experiments that (i) suppliers invest in steel for new orders when past orders (exogenously) end, consistent with constraints on working capital or another fixed factor; (ii) the buyer provides conditional advances to suppliers and monitors supplier sites when a new order is starting; and (iii) the buyer takes direct control of raw material sourcing on behalf of its suppliers when credit constraints worsen.

The response of the buyer raises some natural questions, the most important relating to the mechanisms behind her governance choices. In particular, why does the credit constraint demonstrated in (i) lead to closer integration between buyer and supplier described in (ii) and (iii)? In this section, we build a theoretical model that rationalizes these high-engagement governance choices.

5.1 The model

We begin by developing a framework that embeds the key steps of the interaction in a sequential model of the supply chain. Then, we detail the governance options and the environmental details that the buyer faces. Both of these are enriched with the specifics of the context, in order to keep the model and results as close to our setting as possible. This then enables us to develop the results that shed light on the buyer's governance choices.

A sequential model of the supply chain Consider two agents interacting in a production process. The first agent is a buyer, denoted P, who is product-market facing. The second agent is a supplier, denoted S, who engages in physical production. Production requires the use of noncontractible decisions, which must be assigned by the buyer before production. Who gets these decisions rights has relevance because there is a stochastic payoff-relevant state that realizes after assignment. Figure 9 illustrates the standard timing setup that is used to capture this setup, especially in the incomplete contracting literature following on from Williamson (1979).[25]

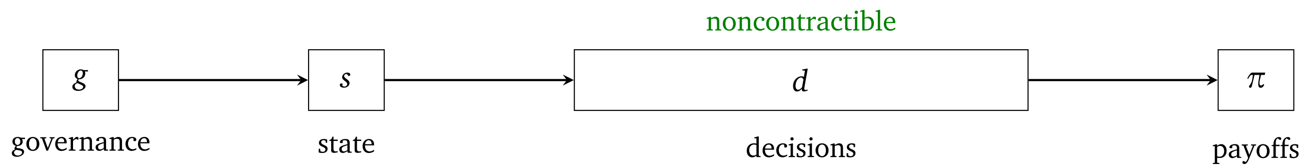

Figure 9: Standard timing diagram of the production process in an adaptation model.

We enrich this setup by expanding the production stage, represented by Figure 10. The original version can accommodate multiple decisions being taken simultaneously, but this version allows for some of these decisions to be taken sequentially, which is the key ingredient that leads to a partially integrated supply chain.[26] In this version, at $t = -1$, agent P determines the governance structure (g) by allocating decision rights and choosing additional governance instruments, such as cash support for input purchase and monitoring of supplier worksites. At $t = 0$, a second buyer P' exists with probability q. The interaction then unfolds in three key stages:

1. *Input Use Stage:* The decision on input use is made at time $t = 1$.
2. *Production Stage:* Agent S engages in production using costly private effort a.
3. *Output Allocation Stage:* The output y realizes on delivery date t_d, which is a function of effort and the start date t_s. The decision on the use of output is then taken, and payoffs π_P and π_S realize to both agents.

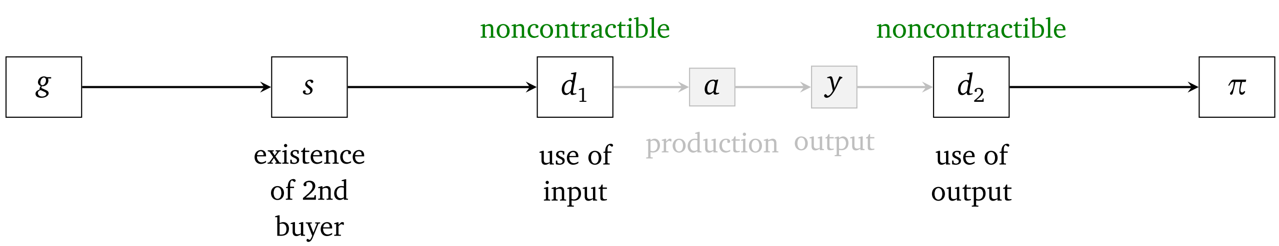

Figure 10: Modified timing diagram of the production process.

[25]See, e.g., Gibbons (2005).

[26]This version has aspects of both the adaptation model, in terms of the first half of the sequence, and the property rights model of Grossman and Hart (1986), in terms of the second half, which has both moral hazard and a decision right.

The governance options The right over decision d_1 provides the right to decide what to do with the inputs to the production process (in our case, money or steel) i.e. whether to use it for an order or to resell it, and which order to use it for if the former. The right over decision d_2 provides the right to decide what to do with the output of production[27] i.e. whether to sell it to the original buyer or sell it to the outside option. These rights are derived from ownership—owning the input provides decision rights over it, while owning the means of productions provides decision rights over the output. In that sense, the decision rights are noncontractible i.e. they cannot be moved via a contract.

Choosing g is about assigning these decision rights, so there are three main governance structures of interest:

1. *Vertical integration*: agent P owns both decision rights d_1 and d_2
2. *Partial integration*: agent P owns one of the decision rights, while agent S owns the other
3. *No integration*: agent S owns both decision rights d_1 and d_2

If the buyer owns the decision right d_2, then the supplier cannot take up second buyer's order[28]. If not, the supplier can choose to accept both. There is also a limited liability constraint - the supplier cannot be made to pay anything to the buyer.

The value of production In the production process, the supplier converts raw material (steel) into a final good (metal part of the required specifications). Producing one unit of the final good requires raw material of value m. We use v to denote value in of the final good for buyer P's order. It has value $v_t = \bar{v}_t > m$ to agent P when delivered at time t, and has value in alternate use $v_t = \underline{v}_t < \bar{v}_t$. Since production is Just-In-Time, the final good has higher value if delivered closer to the scheduled date, which we normalize to $t = 2$. Moreover, delivering closer to the scheduled date adds more value than delivering later. These are captured by the following assumption:

Assumption 1 (Just-In-Time restrictions). *The buyer's value satisfies the following restrictions:*

1. $v_{t-1} > v_t$

[27]The restriction to two decisions and a single outcome is made for the purpose of illustration. In reality, there will be many more stages and decisions, which can lead to richer insights depending on the setting. For us, the crux of the setup is that these decisions are sequential, and the fact that they are non contractible.

[28]This can be due to a number of reasons. One reason can be that the buyer wants to protect production capacity in case she needs another order fulfilled. Or it can be that the supplier needs to put in effort to acquire leads, which she will not do as effectively under integration.

2. $v_{t-1} - v_t > v_t - v_{t+1}$

The second buyer's value is analogously denoted with u, and satisfies the same restrictions as above.

We assume that it is efficient to do the first buyer's order first and deliver it early. This makes our results stronger, in that even when it is efficient to do the first buyer's order first, the supplier may choose to take up the second buyer's order first. We assume also that the second buyer values later delivery sufficiently. This captures the fact that there are significant external market incentives to finish pending orders. These assumptions are specified formally as follows.

Assumption 2 (First buyer values earlier delivery more). *The relative values of the two buyers satisfy:*

1. $v_2 - v_t > u_2 - u_t \;\; \forall\, t$, *and* $v_2 - c > u_3$
2. $\underline{u}_4 - \underline{u}_5 > \underline{v}_3$

Supplier actions during production The supplier can take private actions that can influence the delivery date of an order. In our empirical setting, this can correspond to when to start an order, or how quickly to finish production once it has been started. The higher action $a = 1$ leads to earlier delivery i.e. $t_d = t_s + 1 + \mathbb{1}_{a=0}$, with the higher action leading to private cost c to the agent S. We will assume that $c > \underline{u}_{t-1} - \underline{u}_t$ for all t i.e. the effort costs outweigh the marginal value of delivering the second order early for the supplier, and $\underline{v}_2 - \underline{v}_3 > c > \underline{v}_{t-1} - \underline{v}_t$ for all $t > 3$ i.e. marginal value of delivery early for P outweigh the effort costs when taken up first[29].

If the supplier owns the decision right d_1 over raw materials, any support provided by the buyer could be misallocated[30] with probability $\rho(t; f)$ depending on the temptation t[31] and form of support f. The form of support in our case will either be cash in the form of direct payment or steel in the form of conditional advances. The temptation t is measured by the level of support, so for example a cash payment of value t comes with a deviation temptation

[29] These assumptions make the algebra more straightforward, but are not necessary conditions for our analysis.

[30] We want to stress that misallocations need not always be malicious. A lack of proper management at the suppliers also leads to slippage if the materials are not tracked properly.

[31] The effective source of this temptation is the shadow value of having access to cash or cash-convertible objects for a cash-constrained agent.

of t. The probability of misallocation ρ is strictly increasing and strictly convex in t, with

$$0 < \rho(t, f = s) < \rho(t, f = c) = 1 \qquad \forall t > 0$$

i.e. misallocation of cash is easier than misallocation of steel[32]. For simplicity, the deviation temptation in the case where the decision right d_1 is owned by the buyer via ownership of steel is normalized to 0[33].

Finally, the supplier faces a hard credit constraint: he has access to cash/credit of value $l < m$, i.e. less than the amount required to purchase steel. This assumption ties in to our results in the previous section, which showed that the supplier often lives hand-to-mouth, waiting for payments from previous orders to invest in newer ones.

Additional governance instruments The buyer can pay the supplier either unconditionally or conditional on two commonly observed and verifiable objects - the procurement of raw material m, and the delivery date t_d. Note that since the supplier does the purchasing with conditional advances, the decision right d_1 is owned by the supplier. The agents face a shadow cost of capital r_c per period. Both conditional advances and direct sourcing by the buyer incur the shadow cost of capital for the buyer, while direct sourcing also incurs a fixed cost F of employees having to be engaged in the purchase.

The buyer can also provide management oversight of production. This corresponds to the relationship managers that we observe integrated into the supplier's organization. The manager technology converts manager effort $e \in [0,1]$ into monitoring, which successfully prevents deviations. At effort e, the deviation probability ρ is modified to $\hat{\rho}$ as:

$$\hat{\rho}(t; f, e) = (1-e)\rho(t; f)$$

i.e. with probability e any realized deviations are prevented. The buyer incurs a cost of monitoring effort of $c(e)$, which is strictly convex in e.

5.2 Model predictions

We derive results from the theoretical model that help contextualize our empirical results. First, we compare the difference governance structures, without any of the additional instruments.

[32] For simplicity, $\rho(t, f = c) = 1$, but all we need is that it is higher.

[33] In practice, it will just be lowered by ownership of the raw material.

This provides a sense of when each of those structures might be chosen by the buyer. Then, we derive a result on when the buyer would choose to use the additional instruments of advances and monitoring. We round off our analysis with a result on when the buyer would switch to tighter integration with the supplier. All proofs are in the appendix.

5.2.1 Comparing governance structures without advances and monitoring

Our analysis of the model begins with a comparison of decision right allocations. For this result, we shut down the contractual complications of advances and monitoring, in order to clarify the main force that guides the allocation of decision rights.

The key outcome here is that the buyer may choose to own the input decision right d_1. This happens more often when there is a high likelihood of a second buyer and when the supplier is highly credit constrained. The intuition behind this is that the supplier has a higher incentive to misappropriate the raw material when the second buyer exists and the shadow value of the raw material is high. Moreover, the buyer may also choose to allow the supplier to retain the output decision right d_2 when the second buyer is more likely to exist, since that provides the supplier an additional market incentive to finish the first buyer's order more quickly. This reinforces that the key idea is about tamping down bad market incentives while enabling good market incentives.

Proposition 1 (Partial integration)**.** *Suppose that l is low and q is high i.e. the supplier is highly credit constrained and the second supplier exists with a high probability. Then, the buyer chooses to own the decision right over inputs d_1. Moreover, the supplier is allowed to retain the decision right over output d_2.*

5.2.2 The role of conditional advances and monitoring

The effects of ownership of the decision right can be replicated via contract and management under some conditions, in line with what we observed empirically. This is what we illustrate next. First, suppose that the buyer can contract with the supplier to provide conditional advances, whereby a portion of the final payment is provided after the raw material is purchased. In such a contract, the buyer is able to prevent the deviation temptation induced by the shadow value of cash.

Proposition 2 (Conditional advances for highly constrained suppliers)**.** *When m is relatively low i.e. the cost of raw material needed for production is relatively low, the buyer can replicate the effects of ownership of decision rights over inputs d_1 via conditional advances.*

Although conditional advances reduce deviation temptations, they don't fully remove them. The usefulness of relationship management follows from the possible insufficiency of conditional advances. Since misallocation continues to exist even with conditional advances, the additional presence of a manager can add value for the buyer. It is straightforward to note that the buyer will choose a nonzero level of relationship manager deployment when choosing conditional advances.

Proposition 3 (Management effort is always nonzero)**.** *When the buyer provides conditional advances, her chosen level of management effort e is always nonzero.*

5.2.3 Increased buyer involvement

Our next result addresses the change we observe due to the steel price shock. As the price of steel increases, the cost of material m increases. As a result, the deviation temptation induced by the shadow value of the raw materials increases, increasing the likelihood of misallocation and hence also the cost of tamping down on it via manager effort. This prompts the buyer to shift to buyer sourcing, which only incurs the fixed cost of procurement. This result also rationalizes the observation that the buyer shifts the more constrained suppliers to buyer sourcing.

Proposition 4 (Switch from conditional advances to buyer-sourced inputs)**.** *Suppose the value of raw material required m increases. Then, the buyer shifts some contracts that were originally under conditional advances to buyer ownership of d_1.*

Moreover, the buyer is more likely to shift contracts with a lower l i.e. for the suppliers who are more credit-constrained.

5.3 Key learnings from the model

Taken together, the model and the theoretical results provide a sharp rationale for the nuanced governance choices we observe empirically. Its central contribution is to model the supply chain as a sequence of noncontractible decisions, moving beyond a monolithic view of production. This structure offers a novel lens that rationalizes the buyer's choice of partial integration as a two-fold strategic decision: she targets the raw material stage to mitigate the supplier's binding credit constraint in the presence of noncontractibilities, while deliberately leaving the output stage un-integrated to preserve the supplier's high-powered market incentives for timely production.

The framework rationalizes the various forms of buyer engagement and their heterogeneous application that we observe empirically. Conditional advances and monitoring represent a less integrated form of control, which suffices when the temptation to misallocate is low. As financial pressures caused by the rise in steel price intensify this temptation, direct ownership of inputs becomes the more efficient governance choice, explaining the shift toward deeper integration for the most constrained suppliers.

While calibrated to our setting, this insight is broadly generalizable. The framework suggests that in any supply chain characterized by sequential production and stage-specific frictions (be they financial, technological, or informational), the efficient form of integration will involve targeted control over the most vulnerable stage, rather than wholesale ownership of the entire process.

6 Conclusion

This paper introduces and analyzes partial integration—targeted buyer interventions across firm boundaries—as a powerful alternative to vertical integration. Using detailed daily data from Indian manufacturing suppliers, we show that supply chain disruptions primarily stem from supplier underinvestment in critical inputs, a direct consequence of binding working capital constraints. An instrumental variable analysis confirms that suppliers operate hand-to-mouth, their ability to start new production critically dependent on payments from prior orders.

We demonstrate that buyers mitigate these disruptions not with simple financing, but by blurring firm boundaries. This partial integration takes the form of direct control over key supplier processes: in-person monitoring, contingent contracts linking payments to input purchases, and, in some cases, direct sourcing of raw materials on the supplier's behalf. This involvement is dynamic, intensifying as disruption risk grows; when input prices rise, for example, the heightened risk of resource diversion prompts greater buyer control.

Our theoretical model clarifies why such interventions are necessary. In the presence of noncontractibility over input use, unconditional cash transfers are insufficient, as they invite misallocation. Partial integration resolves this friction by allowing the buyer to control the specific stage of production where noncontractibility bites hardest, while leaving the supplier with autonomy over output decisions to preserve high-powered incentives for timely completion.

This analysis highlights that relational buyers—those with lower monitoring and sourcing costs due to local embeddedness or superior management technology—possess a comparative advantage in fostering resilient trade with developing economies. For policymakers, the implication is twofold: focus on strengthening short-term credit markets to address the root cause of disruptions, and on attracting multinational firms equipped with the relational capabilities necessary to execute these sophisticated, partially integrated sourcing strategies.

References

Adhvaryu, Achyuta, Namrata Kala, and Anant Nyshadham (2024). "Negative Assortative Matching inside the Firm". Working Paper.

Alfaro, Laura et al. (2025). "Trade Within Multinational Boundaries". In: Harvard Business School Working Paper 25-058.

Antràs, Pol (2003). "Firms, Contracts, and Trade Structure". In: *Quarterly Journal of Economics* 118.4, pp. 1375–1418.

— (2023). "An Austrian Model of Global Value Chains". NBER Working Paper 30901.

Antràs, Pol and C. Fritz Foley (2015). "Poultry in Motion: A Study of International Trade Finance Practices". In: *Journal of Political Economy* 123.4, pp. 853–901.

Antràs, Pol and Elhanan Helpman (2004). "Global Sourcing". In: *Journal of Political Economy* 112.3, pp. 552–580.

Antràs, Pol and Vitalii Tubdenov (2025). "Time, Interest Rates, and Global Trade". Harvard University mimeo.

Atalay, Enghin et al. (2019). "How Wide Is the Firm Border?" In: *Quarterly Journal of Economics* 134.4, pp. 1845–1908.

Baker, George P., Robert Gibbons, and Kevin J. Murphy (2002). "Relational Contracts and the Theory of the Firm". In: *Quarterly Journal of Economics* 117.1, pp. 39–84.

Boehm, Johannes and Ezra Oberfield (2020). "Misallocation in the Market for Inputs: Enforcement and the Organization of Production". In: *The Quarterly Journal of Economics* 135.4, pp. 2007–2058.

Casaburi, Lorenzo and Tristan Reed (2022). "Competitive Subsidies and the Supply of Agricultural Inputs". In: *American Economic Review* 112.4, pp. 1246–1283.

Diamond, Douglas W. (1984). "Financial Intermediation and Delegated Monitoring". In: *Review of Economic Studies* 51.3, pp. 393–414. DOI: 10.2307/2297430.

Dye, Ronald A. (1986). "Optimal Monitoring Policies in Agencies". In: *RAND Journal of Economics* 17.3, pp. 339–350. DOI: 10.2307/2555718.

Garg, Shresth, Pulak Ghosh, and Brandon Joel Tan (2023). "Within-Firm Supply Chains: Evidence from India". In: *Journal of International Economics* 144, p. 103793.

Ghani, Ejaz and Tristan Reed (2022). "Relationship Lending and Supply Chain Resilience". In: *Journal of Development Economics* 156, p. 102836.

Gibbons, Robert (2005). "Four formal (izable) theories of the firm?" In: *Journal of Economic Behavior & Organization* 58.2, pp. 200–245.

Grossman, Sanford J and Oliver D Hart (1986). "The costs and benefits of ownership: A theory of vertical and lateral integration". In: *Journal of political economy* 94.4, pp. 691–719.

Hansman, Christopher, Jonas Hjort, Gianmarco León-Ciliotta, et al. (2020). "Vertical integration, supplier behavior, and quality upgrading among exporters". In: *Journal of Political Economy* 128.9, pp. 3570–3625.

Hansman, Christopher, Jonas Hjort, Gianmarco León-Ciliotta, et al. (2020). "Vertical Integration, Supplier Behavior, and Quality Upgrading among Exporters". In: *Journal of Political Economy* 128.9, pp. 3570–3625.

Heise, Sebastian et al. (Feb. 2024). *Tariff Rate Uncertainty and the Structure of Supply Chains*. Working Paper 32138. National Bureau of Economic Research. DOI: 10.3386/w32138. URL: http://www.nber.org/papers/w32138.

Helper, Susan and Rebecca Henderson (2014). "Management Practices, Relational Contracts and the Decline of GM". In: *Journal of Economic Perspectives* 28.1, pp. 49–76.

Helper, Susan and Abdul Munasib (2020). "Relational Contracting in the U.S. Auto Industry". NBER Working Paper 27838.

— (2022). *Economies of scope and relational contracts: exploring global value chains in the automotive industry*. US Department of Commerce, Bureau of Economic Analysis.

Jones, Daniel T., James P. Womack, and Daniel Roos (1990). *The Machine That Changed the World*. Free Press.

Kim, Se-Jik and Hyun Song Shin (2012). "Sustaining production chains through financial linkages". In: *American Economic Review* 102.3, pp. 402–406.

Kremer, Michael (1993). "The O-ring theory of economic development". In: *The quarterly journal of economics* 108.3, pp. 551–575.

Macchiavello, Rocco and Josepa Miquel-Florensa (2017). "Vertical Integration and Relational Contracts: Evidence from the Costa Rica Coffee Chain". IBREAD Working Paper 503.

Macchiavello, Rocco and Ameet Morjaria (2015). "The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports". In: *American Economic Review* 105.9, pp. 2911–2947.

— (2021). "Competition and Relational Contracts in the Rwanda Coffee Chain". In: *Quarterly Journal of Economics* 136.2, pp. 1089–1143.

Manova, Kalina (2013). “Credit Constraints, Heterogeneous Firms, and International Trade”. In: *Review of Economic Studies* 80.2, pp. 711–744. DOI: 10.1093/restud/rds036.

Manova, Kalina and Zhihong Yu (2016). “How Firms Export: Processing vs. Ordinary Trade with Financial Frictions”. In: *Journal of International Economics* 100, pp. 120–137.

Townsend, Robert (1979). “Optimal Contracts and Competitive Markets with Costly State Verification”. In: *Journal of Economic Theory* 21.2, pp. 265–293. DOI: 10.1016/0022-0531(79)90031-0.

Williamson, Oliver E (1979). “Transaction-cost economics: the governance of contractual relations”. In: *The journal of Law and Economics* 22.2, pp. 233–261.

A Tables and figures

Table 1: Effects of other-order payments on raw material sourcing: extensive-margin

Dependent Variables:	RM sourced >0		Payment >0	RM sourced >0
	OLS	Reduced form	Stage 1	Stage 2
Model:	(1)	(2)	(3)	(4)
Variables				
RM sourced >0 (plan)	0.414***	0.414***	0.019	0.412***
	(0.020)	(0.020)	(0.015)	(0.020)
Other-order payment >0 (plan)	0.008	0.022***	0.453***	-0.015
	(0.006)	(0.006)	(0.012)	(0.013)
Other-order payment >0	0.040***			0.081***
	(0.006)			(0.022)
Other-order payment >0 (inst.)		0.021***	0.264***	
		(0.006)	(0.012)	
Fixed-effects				
Order (within supplier)	Yes	Yes	Yes	Yes
Month	Yes	Yes	Yes	Yes
Fit statistics				
F-test (1st stage)			1,473.0	
R^2	0.21491	0.21239	0.47780	0.21136
Dependent variable mean	0.04169	0.04169	0.19096	0.04169
Observations	19,140	19,140	19,140	19,140

Clustered (Order (within supplier)) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

B Details on the empirical design

Our empirical analysis involves two distinct quasi-experimental research designs. First, to test if working capital constraints disrupt production in short (monthly) time intervals, we regress a supplier's investment in inputs (steel) for an order on predicted cash inflows from *past* orders, instrumenting for the timing of inflows using downstream postponements of common construction projects. We use the same instrument to document the buyer's short-run efforts to monitor input use via increased visits to supplier sites. Second, to show how contracts adapt to credit frictions, we leverage a sudden increase in steel prices that raised working capital requirements. In a difference-in-differences-style design, we show that the steel price shock leads suppliers facing more severe credit constraints, as proxied by pre-shock usage of buyer-provided finance (conditional advances) and lateness in sourcing, to sign contracts in which the *buyer* directly sources steel on the supplier's behalf.

B.1 IV test for credit constraints

Our credit constraints test is motivated by interviews with suppliers, who described a tightly managed cash cycle in which revenues from past orders are immediately invested in inputs for current orders. For several reasons, we leverage past-order payments as a source of variation in cash flows. The practical reasons are that, unlike typical balance sheets, we observe order completions at the daily or weekly time scale relevant for Just-In-Time production; and that variation in order completion is partly exogenous to the supplier. Using other-order completion timing is also theoretically appealing. As we clarify in Section 5, rationalizing the contracts requires short-run uncertainty in cash-on-hand. Because the contracts fix order size and price (and thus revenue) in advance, uncertain order timing is empirically the source of that uncertainty.

Construction of downstream delays instrument Note that $\sum_{o'} Payment_{o't}^{INST}$ could be high either because of a supplier's pre-existing plans to complete many orders in month t, or because customer-driven deviations shift many orders to be finished in month t. To isolate variation from the latter – i.e. from shocks to total payments relative to plan – we maintain our control for planned production $\sum_{o'} Payment_{o't}^{PLAN}$.[34]

We now discuss standard IV relevance, exogeneity, and exclusion restrictions, in that order. The aggregated order-time-level instrument has a strong first stage, with a coefficient of 0.269 and an F-statistic of 1653.2 (shown below in Column 3 of Table 1). In other words, predicted variation in total payments due to customer-driven postponements strongly predicts variation for all reasons. Figure ?? provides some intuition as to why: at the lot level, delay length at other suppliers s' (on the same construction project c) strongly predicts delay length

[34] Borusyak and Hull (2020) show that one approach to unbiased estimation of (3) is to specify the data-generating process for the shock, compute expected payments $\mu_{ot} = \mathbb{E}[Payment_{ot}^{INST}]$ over many permutations of the shock, and include μ_{ot} as a control. In our quasi-experimental setting, we do not know the DGP for $PredDelay_{ig}$, so the level at which to permute. In the Borusyak and Hull (2020) framework, controlling for $\sum_{o'} Payment_{o't}^{INST}$ is thus akin to assuming that:

$$\mathbb{E}[Payment_{ot}^{INST}] = Payment_{ot}^{PLAN} \tag{7}$$

i.e. that order-o payments are more likely to arrive at t *only* because they were planned to arrive in month t (not in a nearby month). This control is imperfect because payments could also be likely to arrive in months $t-1$ and $t+1$ because of planned arrival at t.

at supplier s, with a slope close to 1 for non-extreme delays.

The exogeneity of $\sum_{o'} Payment^{INST}_{o't}$ relies on (i) controlling for planned production, and (ii) the exogeneity of customer-driven delays relative to plan. For (i), as discussed above and in footnotes, we explicitly condition on planned production, and in future versions will explore alternative versions of delays relative to plan. For (ii), we must claim that customer-driven delays $PredDelay_{ig}$ are a common shock out of any supplier's control. If instead a shock originating at supplier s could push back the overall production timeline, for the customer and thus s', delays at our firm of interest s would contaminate the instrument. Based on features of the context, such contamination is of limited concern. The suppliers in our sample are typically small (under 5%) of any construction project, so even among contract delays are due to delayed inputs, the source of delays are likely not in our sample. Furthermore, in interviews relationship managers cited several other reasons for downstream delays beyond input shortages; these include weather disruptions, regulatory holdups, and labor shortages due to pandemic-related migration restrictions.

Finally, the exclusion restriction would fail if predicted payments ($\sum_{o' \neq o} Payment^{INST}_{o't}$) affected raw material sourcing ($Steel^{OBS}_{ot}$) through channels *other* than observed payments ($\sum_{o' \neq o} Payment^{OBS}_{o't}$). The challenge here is in separating stages: since payments closely follow physical completion of the final product, it could be that a different input than working capital is scarce. For instance, completion of goods for order o' might free up labor or plant space to start order o, and the supplier might therefore also procure steel in tandem. If most lots were physically produced and paid for for in different months, we could rule out this possibility by constructing an alternative $\sum_{o' \neq o} Production^{INST}_{o't}$ and showing it is less predictive of raw material sourcing.

In practice, for 69.6% of lots, *Production* and *Payment* occur in the same month, so we cannot separately identify the effects of delays at a particular stage. Nevertheless the regression remains informative of whether some scarce input exists, because $\beta^{OBS} > 0$ would indicate that new orders start when past orders end. Furthermore, several features of the setting suggest that credit is the relevant fixed factor. First, in the industrial clusters we study, a 'gang' of workers can always be hired on a spot market on short notice, so worker availability is rarely an issue. Second, to ensure that supplier (physical) capacity constraints are not met, the buyer knows total supplier capacity and never assigns orders that exceed 30% of supplier capacity. Third, evidence from contract terms (see Figure 5 and Section 3.2 below) suggests that buyers are aware of short-term supplier credit constraints.

To further test the credit constraints mechanism, we can estimate heterogeneous elasticities in Equation (3) using proxies for pre-existing credit constraints. The first is firm-level use of partial financing (conditional advances) from the buyer. Several conversations with the buyer suggested that the most-constrained firms take up financing: advances are only offered to suppliers lacking independent connections with working capital financing institutions, and the buyer charges suppliers above-market-rate interest for the duration of the order. Thus orders with partial advances are exactly those where we should anticipate responses to cash inflows.[35] The second proxy is supplier region: if credit constraints are relevant, then suppliers in Central and North India – which generally has inferior income levels, contracting institutions and financial markets to West and South India – should be more responsive to the timing of payments.

C Proofs

Proof of Proposition 1. We solve the model via backward induction. Denote:

- Market governance g_1: agent S owns both d_1 and d_2
- Partial integration g_2: agent P owns d_1, agent S owns d_2
- Full integration g_3: agent P owns both d_1 and d_2

Let us start with g_1. For a range of parameter values, the supplier will deliver the first buyer's order only in period 3 if the second buyer exists. To see this, suppose the buyer sets up a contract with the supplier for delivery in period 2. Due to credit constraints, she must pay $m-l$ upfront. Suppose she promises a payment p_1 for delivery at time 2.

Consider first what happens if the supplier starts the first buyer's order at $t=1$ and chooses $a=1$. She delivers at time $t=2$ and gets paid p_1 for the first order. Then, she chooses $a=0$ for the second order and delivers that at time $t=4$ (since, by assumption, $c>\underline{u}_3-\underline{u}_4$). For this, she gets at most $\underline{u}_4$, which is her outside option with the output. So her utility is

$$p_1+\underline{u}_4-c$$

Now suppose she starts the second buyer's order first. She chooses $a=0$ for both orders (since, by assumption, $c>\underline{u}_2-\underline{u}_3$ and $c>\underline{v}_4-\underline{v}_5$), and gets $\underline{u}_3$ for the second order and $\underline{v}_5$ for the first, but additionally, she also earns the shadow cost on the $m-l$ amount that the first

[35] Because the advances are partial, suppliers still need some cash from elsewhere to buy inputs.

buyer has provided to her[36]. Thus, her utility is

$$\underline{u}_3 + \underline{v}_5 + 2r_c(m-l)$$

To make the supplier choose to start the first order first, agent P must ensure

$$p_1 + \underline{u}_4 - c = \underline{u}_3 + \underline{v}_5 + 2r_c(m-l)$$

so that she is indifferent between the two. Now, note that the highest value of p_1 that the first buyer can credibly promise is $\underline{v}_2$. This is because once the supplier has produced the output, she cannot get more than the outside option for it. Thus, if we have

$$\underline{u}_3 + \underline{v}_5 + r_c(m-l) \geq \underline{v}_2 + \underline{u}_4 - c \tag{8}$$

$$\iff \underline{u}_3 - \underline{u}_4 + 2r_c(m-l) + c \geq \underline{v}_2 - \underline{v}_5 \tag{9}$$

then the first buyer cannot make the supplier take up her order first.

Assume Equation 8 holds, since l is assumed to be low. Then, if the second buyer exists and the supplier owns both decision rights, the first buyer's payoff is $\overline{v}_5 - \underline{v}_5 - 2r_c(m-l)$. If the second buyer does not exist, she gets $\overline{v}_2 - \underline{v}_2 - r_c(m-l)$. So her expected payoff from this governance structure is

$$V(g_1) = q[\overline{v}_4 - \underline{v}_4 - 3r_c(m-l)] + (1-q)[\overline{v}_2 - \underline{v}_2 - r_c(m-l)]$$

Next, suppose the buyer owns both d_1 and d_2. Then, agent S has to start working on her order at $t=1$, and the output is allocated to her own uses. To motivate the supplier to take $a=1$, she must pay him c for delivery to happen at period 2. Thus, her payoff is

$$V(g_3) = \overline{v}_2 - c - r_c(m-l) - F$$

If $c + F \geq \underline{v}_2$, we have that $V(g_3) \leq V(g_1)$ if $q = 0$. Moreover $V(g_3) > V(g_1)$ if $q = 1$, since l is low. Thus, by IVT, it must be that there is a $\hat{q}$ s.t. agent P chooses governance structure g_3 for $q \geq \hat{q}$. If $c + F < \underline{v}_2$, we have $\hat{q} = 0$.

Now consider the case where the supplier owns d_2 and the buyer owns d_1. Observe that now the $m-l$ term disappears from Equation 8 since the buyer now owns the input, so the

[36]One way to interpret this is that the second buyer pays her an amount $x \leq 2r_c(m-l)$ to use the first buyer's cash support for the second buyer's order, since otherwise the second buyer needs to incur $2r_c(m-l)$.

supplier does not earn the shadow value from it. Since by assumption

$$\overline{v}_{t-1} - \underline{v}_t > \overline{u}_{t-1} - \underline{u}_t$$

the supplier will not take up the second buyer's order first if $p_1 \geq \underline{u}_3 + \underline{v}_5 - \underline{u}_4$. Moreover, consider the payment p_1 that the buyer promises for delivery in period 2 i.e. $a = 1$. If the second buyer exists, the supplier gets $p_1 + \underline{u}_4 - c$ if he takes action $a = 1$, while he gets $\underline{v}_3 + \underline{u}_5$ if he takes action $a = 0$. Thus, $p_1 = c - (\underline{u}_4 - \underline{v}_3 - \underline{u}_5)$ is enough to make her take action $a = 1$. If the second buyer doesn't exist the supplier gets $p_1 - c$ if he takes action $a = 1$, while he gets $\underline{v}_3$ if he takes action $a = 0$. So suppose the buyer offers $p_1 = \max\{\underline{u}_3 + \underline{v}_5 - \underline{u}_4, c - (\underline{u}_4 - \underline{v}_3 - \underline{u}_5)\}$. Since $c > \underline{u}_3 - \underline{u}_4$, we have that $c + \underline{v}_3 - \underline{v}_5 > \underline{u}_3 - \underline{u}_4$, so $p_1 = c - (\underline{u}_4 - \underline{v}_3 - \underline{u}_5)$. Then she gets delivery in period 2 if the second buyer exists, and delivery in period 3 if she doesn't. So her total payoff is

$$V(g_2) = q[\overline{v}_2 - c - (\underline{u}_4 - \underline{v}_3 - \underline{u}_5)] + (1-q)[\overline{v}_3 - \underline{v}_3] - r_c(m-l) - F$$

Given $\underline{u}_4 - \underline{v}_3 - \underline{u}_5$, when $q = 1$, $V(g_2) > V(g_3)$. When $q = 0$, $V(g_2) < V(g_3)$. Thus, by the IVT, there is a $\underline{q}$ such that for $q > \underline{q}$, the buyer chooses g_2. Combining the comparisons between g_1 and g_3, and between g_2 and g_3, we see that the buyer owns d_1 when q is sufficiently high, and additionally when $\underline{u}_4 - \underline{v}_3 - \underline{u}_5$, the buyer allows the supplier to own d_2. This completes the proof. □

Proof of Proposition 2. Consider a contract which provides for a conditional advance instead of cash upfront. The $(m-l)$ term from Equation 8 gets replaced with $\rho(m-l,s)(m-l)$. The deviation temptation has been reduced since

$$\rho(m-l, f = s) < \rho(m-l, f = c) = 1$$

given that $m - l > 0$, and hence the first buyer's order can get taken up earlier even if the second buyer exists. Clearly the set of parameters for which this happens is smaller than the case of ownership of d_1 since

$$\rho(m-l, f = s) > 0$$

given that $m-l > 0$. In particular, we need m to be low enough given a specific ρ function. □

Proof of Proposition 4. For this result, we will compare the value of owning d_1 for the buyer versus the value of providing conditional advances with manager oversight.

The difference in these values is given by

$$\Delta V = (\overline{v}_2 - F) - [(1 - q(1-e)\rho(m-l;s))\overline{v}_2 + q(1-e)\rho(m-l;s)\overline{v}_5 - c(e)]$$

The first term follows from noting that ownership of d_1 removes deviation temptations, but at the cost F of doing the sourcing. The second term takes into account the fact that the supplier will have deviation temptations, which the buyer will optimally tamp down with manager effort. With the resulting probability of deviation, the buyer gets delivery in period 5 if the second buyer exists, while otherwise she gets delivery in period 3. The optimal effort level can be calculated as

$$e = \arg\max_{e'}[(1-(1-e')\rho(m-l;s))\overline{v}_2 + \rho(m-l;s)(1-e')\overline{v}_5 - c(e)]$$

Now, the derivative of the difference in values wrt m is given by

$$\frac{d\Delta V}{dm} = q(1-e)(\overline{v}_2 - \overline{v}_5)\rho'(m-l;s)$$

where we apply the Envelope Theorem to ignore the partial derivative wrt e. Since ρ is strictly increasing in $m-l$, this derivative is positive. Thus, increasing m can switch this difference from negative to positive, which implies that the regime switches from conditional advances to buyer sourcing.

Moreover, the derivative wrt l is given by

$$\frac{d\Delta V}{dl} = -q(1-e)(\overline{v}_2 - \overline{v}_5)\rho'(m-l;s)$$

By the same arguments as before, this derivative is negative. Starting out from a regime of conditional advances, ΔV must be negative for all observed l. Since the derivative wrt l is negative, it is more negative for higher l. Thus it is the lower l suppliers who cross over first to ΔV being positive. This can be seen from the cross partial

$$\frac{d^2\Delta V}{dldm} = -q(1-e)(\overline{v}_2 - \overline{v}_5)\rho''(m-l;s)$$

which is negative since ρ is convex. Thus, the suppliers who get switched to buyer sourcing are exactly those suppliers who were more constrained ex ante. □

D Additional results from the model

D.0.1 Comparative advantage of relational firms

The results so far showcase that closer involvement by the buyer in supplier processes is critical in settings with noncontractible decision rights. A natural implication is that firms that have a comparative advantage in a more integrated supplier management technology will do relatively better in more constrained scenarios.

Proposition 5 (Relational firms are able to get on-time delivery more often)**.** *Consider two types of agent P, with the first having lower cost of monitoring $c_1(e) < c_2(e)$ and lower cost of local raw material procurement $F_1 < F_2$. Then, the first type of agent P will be able to get on-time delivery more often.*

Proof. Consider first the case of conditional advances with management. Agent P's value from allowing the supplier to retain both decision rights is given by

$$V(g_1) = q[\overline{v}_4 - \underline{v}_4 - 3r_c(m-l)] + (1-q)[\overline{v}_2 - \underline{v}_2 - r_c(m-l)]$$

Her value from providing conditional advances with management is given by

$$V_{cm,i} = \max_e[(1-q(1-e)\rho(m-l;s))\overline{v}_2 + q(1-e)\rho(m-l;s)\overline{v}_5 - c_i(e)] = \max_e U_i(e)$$

Note that $U_2(e) < U_1(e)$ for all e. So we must have

$$\max_e U_2(e) = U_2(e_2^*) < U_1(e_2^*) \leq \max_e U_1(e)$$

i.e. the type with higher cost gets strictly lower value when choosing conditional advances and monitoring. Since $V(g_1)$ is the same for both, this means that the type with higher cost chooses g_1 more often, which means she settles for strictly later delivery more often. The same argument applies for the case of direct buyer sourcing as well. □

To complete the comparative advantage analysis, note that when $l = m$, both type of firms get on-time delivery under market transaction since the deviation temptation in *Equation* 8 goes away.

Proposition 6. *When suppliers are unconstrained i.e. $l = m$, both types of agent P get on-time delivery.*

Thus, relational firms have a comparative advantage in developing countries that necessitate a more integrated management style, since they do better when constraints worsen.

D.0.2 Completing working capital markets is a first-best policy solution

Finally, we discuss a policy solution that can restore the first-best in this setting. Note that the closer integration requirement means that major buyers need to either be good at closer, local involvement or be able to find intermediaries that are good it, like the buyer in our setting. This can often deter high value customers from entering a developing country's supply chain, and hence negatively impact the aspirations of that country's suppliers. Thus, while partial integration strategies offer a useful second-best from the agents' perspective, there is scope to improve further.

In our simplified setting, this improvement stems from relaxing the credit constraint. The crux of the inefficiency is that the buyer needs to provide some support for input procurement to the supplier, which can be misappropriated due to the noncontractibility of decision rights. Crucially, the buyer wants the support to be applied only to her order, since she gains value only in that case. If the supplier can get the same support from the bank instead, this inefficiency goes away since the bank only cares about being repaid, not about which order the support goes into.

Proposition 7 (Completing working capital markets is first-best)**.** *Suppose the supplier can obtain a working-capital loan from an external lender at (gross) interest rate* $(1+r)$*, where* r *is sufficiently low. Then:*

1. *If the supplier owns both decision rights, the supplier always completes the first buyer's order first under market governance.*
2. *As a result, the buyer* P *strictly prefers market governance.*

Proof. Suppose the supplier owns both decision rights. We use backward induction, comparing the two possible sequences: $S1 : (P, P')$ and $S2 : (P', P)$.

Sequence $S1$

Since r is low and $c > \underline{u}_t - \underline{u}_{t+1}$, if P' is second, the supplier chooses low effort ($a = 0$), finishing it in 2 periods at no effort cost. The supplier may use high effort ($a = 1$) or low effort ($a = 0$) for the first order.

- High effort on P: finishes P in 1 period, i.e. at time $t=1$, payoff $= \underline{v}_1 - c$, then P' runs from 1 to 3, so the net from P' is $\underline{u}_3 - m\big(1 + r \times 3\big)$. Thus her payoff is

$$(\underline{v}_2 - c) + (\underline{u}_4 - 3rm).$$

- Low effort on P: finishes P in 2 periods, i.e. at time $t=2$, payoff $= \underline{v}_2$, then P' runs from 2 to 4, net payoff from P' is $\underline{u}_4 - m\big(1+4r\big)$. So her payoff is

$$\underline{v}_3 + (\underline{u}_5 - 4rm).$$

Hence

$$\Pi_{S1} = \max\Big\{(\underline{v}_2 - c) + (\underline{u}_4 - 3rm),\ \ \underline{v}_3 + (\underline{u}_5 - 4rm)\Big\}.$$

Since $v_2 - v_3 > c$, we have

$$\Pi_{S1} = (\underline{v}_2 - c) + (\underline{u}_4 - 3rm).$$

Sequence $S2$

Since r is low, our assumption that $c > \underline{u}_{t-1} - \underline{u}_t$ for all t and $c > \underline{v}_{t-1} - \underline{v}_t$ for all $t > 3$ means that the order for P' takes 2 periods, finishing at time $t=3$. Then the order for P runs from 3 to 5. Hence

$$\Pi_{S2} \ = \ \underline{u}_3 \ + \ (\underline{v}_5 - 4rm).$$

We have that

$$\begin{aligned}\Pi_{S1} - \Pi_{S2} &= (\underline{v}_2 - c) + (\underline{u}_4 - 3rm) - \underline{u}_3 - (\underline{v}_5 - 4rm)\\ &= (\underline{v}_2 - c - \underline{u}_3) + (\underline{u}_4 - \underline{v}_5 + rm)\end{aligned}$$

Since $\underline{v}_2 - c > \underline{u}_3$ and $\underline{u}_4 > \underline{u}_4 - \underline{u}_5 > \underline{v}_3 > \underline{v}_5$ by assumption, the supplier chooses to produce the first buyer's order first. Since this is the best possible outcome for P, and since she incurs additional costs in all other governance forms, she chooses market governance.

$\square$