Supply chain resilience via partial integration

Aroon Narayanan
MIT

Vishan Gandhi Nigam
MIT

April 16, 2026

Abstract

How do buyers influence supplier decisions without vertical integration? Using novel daily timeline data from Indian manufacturing supply chains, we document *partial integration*: targeted buyer interventions – monitoring, contingent contracts, and direct control of supplier decisions – that replicate the functions of vertical integration across firm boundaries. We first show that supplier underinvestment in key inputs, driven by working capital constraints and noncontractibility of input use, is a central contracting friction in these supply chains. In response, buyers exert control over supplier input decisions rather than merely advancing cash. This involvement intensifies as contracting frictions worsen: an unanticipated input cost shock leads to direct buyer control of inputs for the most constrained suppliers. We develop a model that rationalizes these organizational choices. Buyers control input decisions to prevent resource diversion, while leaving suppliers to retain control of production to preserve output market incentives. The model also predicts that relational buyers with low monitoring and sourcing costs enjoy a comparative advantage in sourcing from constrained suppliers.

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

Modern production relies on supplier decisions that are invisible beyond firm boundaries. A court can easily verify the quantity and price of iPhones that Foxconn delivers to Apple. It generally cannot verify whether Foxconn cuts corners on quality checks, hides production errors, starts orders late, or diverts Apple-provided financing and tooling to competing customers. In settings like these, how do buyers influence the internal decisions of their suppliers?

The classical answer is vertical integration. For highly specific transactions, Williamson (1979) argues, unified governance economizes on the costs of adapting to non-contractible contingencies relative to arms-length contracting; Grossman and Hart (1986) show that reallocating residual control rights through ownership mitigates holdup over relationship-specific investments.[1] But vertical integration is a costly and blunt instrument. Suppliers lose scale and specialization advantages, and buyers must acquire entire firms to address problems at a single production stage.

In practice, buyer-supplier relationships are often not fully integrated. For instance, about half of U.S. firm imports are from unrelated parties, including from many countries with weak contracting institutions and of customized goods for which supplier decisions are paramount. These non-integrated transactions can feature substantial buyer involvement: firms provide inputs, credit, and knowledge to suppliers, own capital equipment, and determine production decisions on supplier factory floors.[2] Despite their prevalence, we lack systematic knowledge regarding the forms and determinants of these intermediate arrangements. Studying them requires observing, within a single setting, both the internal supplier decisions at risk and the contractual and organizational tools that buyers deploy to shape them.

This paper documents and analyses this intermediate arrangement, which we call *partial integration* because it mirrors the activities of a vertically integrated firm, but is implemented across firm boundaries. In the spirit of Aghion and Tirole (1997)'s distinction between formal and real authority, partial integration is the acquisition and exercise of real authority – rights of observation, incentive provision, and direct control – over a bounded set of supplier decisions, without taking formal authority over the firm as a whole.

[1] A long empirical literature documents vertical integration in contexts with exactly these incentive problems, from Peruvian fishmeal (Hansman, Hjort, León-Ciliotta, et al., 2020) to U.S. cement (Hortacsu and Syverson, 2007) and the universe of U.S. imports (Alfaro et al., 2025).

[2] Agroprocessing firms, as in the well-known Birds Eye contract-farming case, often provide seed and closely direct production decisions on farms (Grant, 2002); Apple often owns critical equipment within supplier plants to control use and prevent knowledge diversion (McGee, 2025); and China's manufacturing growth was initially fueled by "processing trade" in which buyers provided key inputs (Manova and Yu, 2016).

We document partial integration using new administrative data from Indian metal fabrication – a sector supplying customized inputs for power plants, skyscrapers, and other critical infrastructure. The data come from a large manufacturing intermediary that brokers contracts between downstream multinational construction conglomerates and upstream suppliers, most of them small and medium enterprises. Critically, we observe stage-by-stage production timelines, from the sourcing and cutting of steel through welding, painting, and inspection, alongside the full text of contract terms and day-to-day buyer-supplier interactions at supplier plants.[3] This combination lets us see both the internal supplier decisions that matter for performance and the organizational tools the buyer uses to shape them.

Our empirical analysis proceeds in four steps. We begin by identifying the supplier decision that matters most for successful production in our setting. Decomposing total delays across production stages, we show that late sourcing of suppliers' own raw material, rather than slow physical production or failed quality checks, is the largest single source of production delays; and that sourcing delays are heavily concentrated in India's poorest regions. We then provide evidence that these sourcing delays are driven by short-run credit frictions. In particular, instrumenting for the timing of a supplier's cash inflows using a leave-out measure of delays at other suppliers on the same project, we show that procurement of raw material (steel) for new orders rises sharply when payment from past orders arrives, with larger effects for suppliers in poorer regions and for those that receive buyer financing. In other words, constrained suppliers behave as if hand-to-mouth: despite signing contracts to produce, suppliers do not always invest limited cash-on-hand in inputs for new orders.

With the key decision identified, we turn to the buyer's response. If the sole friction was imperfect supplier access to credit, buyers could simply advance cash and let the supplier buy steel when needed. However, no supplier in the contract data receives unconditional cash financing. Each intervention we observe instead is tied to the procurement decision itself. First, buyers increase physical *monitoring* of supplier sourcing decisions, as proxied by the frequency of status updates by buyer "relationship managers" who visit supplier sites. Monitoring intensifies precisely when payments from past orders arrive and input investment becomes feasible, which we establish using the same leave-out instrument as the credit constraint regression.

Two more direct forms of partial integration are directly visible from contract terms. The first is a *conditional advance*, in which the buyer reimburses the supplier only upon verified procurement of raw material. The second is *buyer-sourced raw material*, in which the buyer purchases steel on the spot market and delivers it directly to the supplier's plant. Both target

[3] In most trade and financial datasets, pre-delivery production stages and contract terms are considered trade secrets. A notable exception is Manova and Yu (2016) on Chinese processing trade.

the same upstream decision – suppliers' own input sourcing – via different means. Conditional advances are *incentives*, while buyer sourcing grants buyers direct *control* of supplier decision-making. Both are more common in poor regions (where sourcing delays are also more common). Both interventions are narrowly targeted: contracts do not involve control or incentives contingent on other production stages, such as the cutting or welding of steel. The presence of these interventions, and the absence of unconditional cash advances, rule out a pure financing story. Instead, there must be non-contractibility over the supplier's input *use*, not just access to credit.

Since partial integration is an organizational choice, its form should respond to the severity of the underlying frictions (credit and contracting). To test this, we exploit the sharp rise in global steel prices in late 2020, which raised working capital requirements for all suppliers in the sample. In a difference-in-differences (DID) design, we show that the buyer substituted from conditional advances to buyer sourcing – i.e. from *incentives* to *control* – exactly for suppliers with the greatest ex-ante credit needs. In particular, suppliers receiving a 10 p.p. larger conditional advance before the shock became 4.8 p.p. (28%) more likely to receive buyer-sourced steel afterward.

These results, combined with the fact that buyer sourcing and conditional advances coexist in the cross-section, imply that partial integration is relatively dynamic. The form of buyer involvement (incentives vs. control) is both customized to the supplier and adjusted over time in response to economic conditions. However, they leave several questions unanswered. First, why would a buyer choose to only partially integrate, i.e. to target certain stages of supplier production but not others? Second, within partial integration, why are the advantages of direct control relative to indirect incentives larger when credit frictions worsen? And third, which firms are best positioned to adopt this intermediate organizational form?

To answer these questions, we build a simple model of a supply chain in which integration (partial or full) is endogenous. Production has two stages. At the input stage, the supplier must procure raw materials using scarce working capital, which it can divert to competing uses, for example to another buyer's order – an input-diversion friction. At the delivery stage, by contrast, the supplier's effort is disciplined by the presence of other buyers: because production capacity is scarce, the only way for the supplier to start a lucrative outside job is to finish the current one. Thus, supplier processes at both stages are affected by external market incentives.

We derive three results from the model that speak directly to our empirics. First, we show that the buyer's choice to integrate partially is motivated by the desire to retain the favorable external market incentives at the delivery stage while tamping down on the harmful external market incentives at the input procurement stage. Partial integration dominates full integration when suppliers' opportunities to serve other buyers are sufficiently valuable, because ceding control of delivery preserves the external market incentive device that ownership would erase. Second, monitored conditional advances (i.e. *incentives*) dominate unconditional cash whenever the input friction is severe enough to warrant costly intervention, because cash is more fungible than the procured input. Third, there is a threshold on the financing gap (the share of cash for an order that suppliers cannot self-finance) above which the buyer switches from conditional advances to direct sourcing: once temptation rises high enough, even conditional advances are inadequate because suppliers are willing to divert physical inputs (not just cash) to other orders. In such a scenario, buyer must take the supplier's sourcing decision away entirely. The model also implies that "relational" buyers with lower monitoring and sourcing costs can serve a wider set of constrained suppliers, and that complete working capital markets would eliminate the need for buyer intervention altogether.

Taken together, these results imply that firm boundaries are flexible and *stage-specific*. Within a single buyer-supplier relationship, control rights over distinct production decisions can be allocated separately and coherently: real authority is taken at each stage according to where contracting frictions bind hardest, with autonomy preserved elsewhere precisely because the supplier's outside options already discipline it. This result refines the standard conceptualization of firm boundaries: the relevant organizational choice is not whether to integrate the firm, but which decision rights to integrate at which stages – and the efficient allocation depends on where non-contractibility bites, and on the structure of the supplier's outside options at adjoining stages.

The framework also suggests a comparative advantage for embedded domestic buyers, intermediaries, and firms practicing relational sourcing in the tradition of Japanese supplier management[4] over arms-length spot buyers in weak-contracting environments, because their lower monitoring costs and sourcing costs let them exercise control profitably at the most vulnerable stages. One implication is that efforts to integrate developing-country suppliers into global supply chains may be more effective when they attract firms with relational sourcing capabilities or develop short-term trade finance markets that substitute for buyer intervention.

Related literature. This paper contributes to several literatures. The first is on vertical integration. An established theoretical literature shows that governance structures minimize the

[4] See Cajal-Grossi, Macchiavello, and Noguera (2023), Heise et al. (2024), and Helper and Munasib (2022) on Japanese versus U.S. sourcing strategies. Intermediaries such as Li & Fung in global garments and our anonymous partner in Indian metal fabrication play exactly this role.

costs of maladaptation to contracting hazards (Williamson, 1979), that ownership reallocates residual control rights to prevent holdup (Grossman and Hart, 1986; Antràs, 2003; Antràs and Helpman, 2004; Baker, Gibbons, and Murphy, 2002), and that formal and real authority over decisions need not coincide with ownership (Aghion and Tirole, 1997). A related empirical literature documents that a large 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 show 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 evidence that the form of partial integration responds to the severity of contracting frictions.

A second literature documents how firms adopt "thick" contractual arrangements that relax constraints to production by blurring firm boundaries. In Chinese processing trade, foreign buyers own the inputs that domestic factories transform (Feenstra and Hanson, 2005); in Rwanda, coffee mills provide inputs and credit to farmers (Macchiavello and Morjaria, 2021); and in other agricultural supply chains, goods and credit are often bundled (Casaburi and Reed, 2022). Relative to this work, we study how buyers choose between distinct forms of thick arrangement – conditional incentives and direct control of input sourcing – and provide causal evidence on the determinants of that choice.

Our paper also relates to a literature on sourcing strategies used to avert supply chain disruptions. Most of this work focuses on long-term relationships, showing 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 discouraging resale of output on spot markets. We study a different margin of disruption – input diversion risk due to working capital constraints – and a 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), connecting our paper to the literature on organizational forms that enable Just-In-Time production.

Finally, a growing literature studies how financial frictions within supply chains shape the contractual arrangements between buyers and suppliers. Credit shocks to suppliers propagate to downstream customers through reduced trade credit and supply of goods (Costello, 2020), and firms actively manage liquidity shortfalls by adjusting credit extended to and received from supply chain partners (Amberg et al., 2021). In this environment, buyers adapt: they supply inputs in-kind to credit-constrained suppliers (Manova and Yu, 2016), and payment terms shift as contracting frictions worsen (Antràs and Foley, 2015). We study how buyers choose between two specific instruments – conditional advances and direct sourcing – and provide causal evidence on when each is chosen.

The remainder of the paper is organized as follows. Section 2 describes the setting and data. Section 3 identifies the key supplier decision and establishes the role of working capital constraints. Section 4 documents partial integration. Section 5 shows that the form of partial integration responds to the severity of the friction. Section 6 develops the model. Section 7 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 contracts into smaller orders, which are 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 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 order, and can be an order of magnitude larger than labor expenses (10-20%) and typical profit margins (5-10%). Because steel is a commodity sold on a spot market, suppliers must pay upfront for their primary input and then receive payment upon delivery for their customized product. 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 and measurement

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 have 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 observe delays at each stage, along with 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 single "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 is 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, according to buyer interviews, quality standards are similarly stringent across end customers.

Beyond production timelines, we observe contract terms agreed upon by the parties, including percentage advances (and the associated conditions, which relate to raw material

Figure 1: Production timeline structure

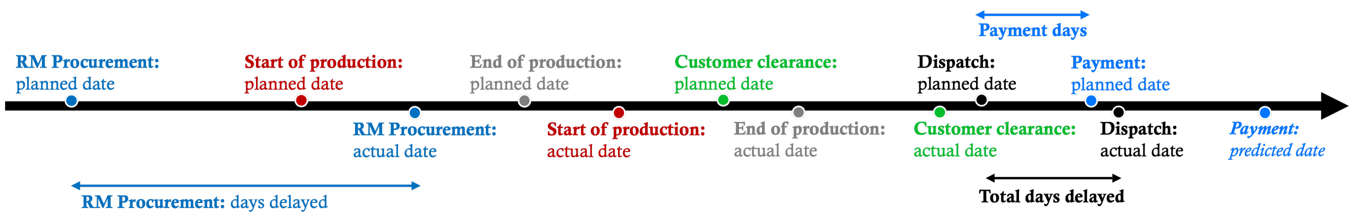

Notes: Figure shows production timeline data recorded by the buyer for a sample *lot* (or truckload) of fabricated metal. Planned dates by stage are set in advance, and actual dates of stage completion are verified by buyer visits to supplier production sites.

purchase), post-delivery payment terms (anywhere from 0 to 45 days post-delivery), and the supplier's responsibilities. 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.

Taken together, these data let us observe both the timing of production and the organizational arrangements buyers use to shape supplier behavior. We use them first to identify the key supplier decision for on-time delivery, and then to document how buyers intervene in that decision without taking ownership.

3 The key supplier decision: input sourcing

We begin by identifying the supplier decision that matters most for timely delivery. In our setting, the critical upstream choice is whether and when to invest scarce working capital in raw material for a particular order. The evidence below shows that this stage, rather than downstream production effort, is where delivery performance is most vulnerable.

Figure 2: Stagewide days delayed

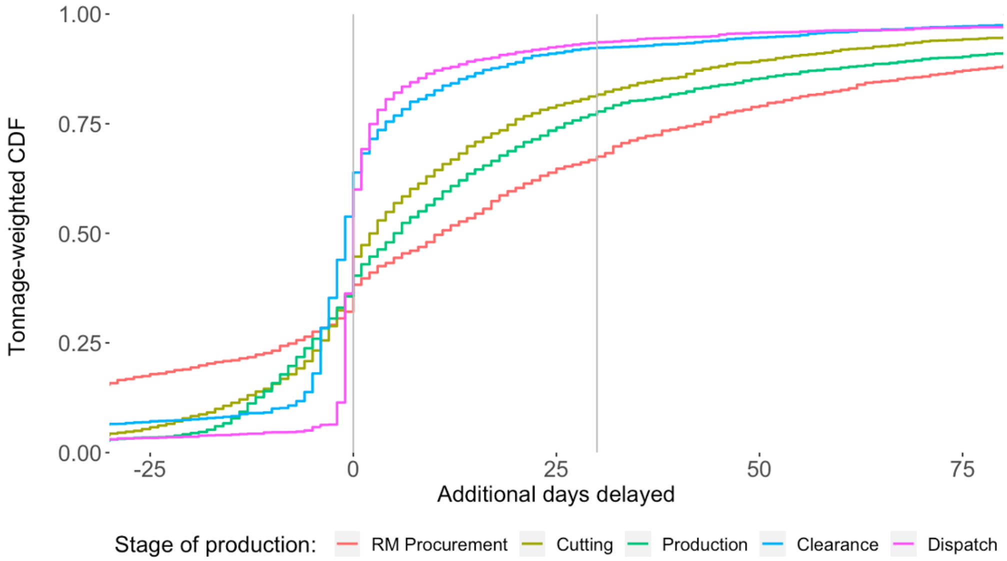

Notes: Figure shows CDF of additional delays by stage, across lots (truckloads) of output. Delays from stage s are the number of days behind schedule that stage s is completed (i.e. time between planned and actual completion dates), less the days behind schedule that stage $s-1$ is completed. RM Procurement 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.

3.1 Delays are concentrated at raw material procurement

We first decompose the total days delayed for each lot (truckload) into its constituent stages and plot the stage-wise CDF.[5]

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 due to insufficient labor or plant capacity, as well as delays in completing production due to worker absenteeism, weather disruptions, or rework after failed quality checks. The relative importance of procurement delays is especially striking because the primary input, steel sheet, is homogeneous, sold on a spot market, produced domestically, and widely available near supplier locations.

This pattern points to a decision problem rather than an access problem. Suppliers appear

[5]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 3: Cumulative delays by state at raw material and end-of-production stage

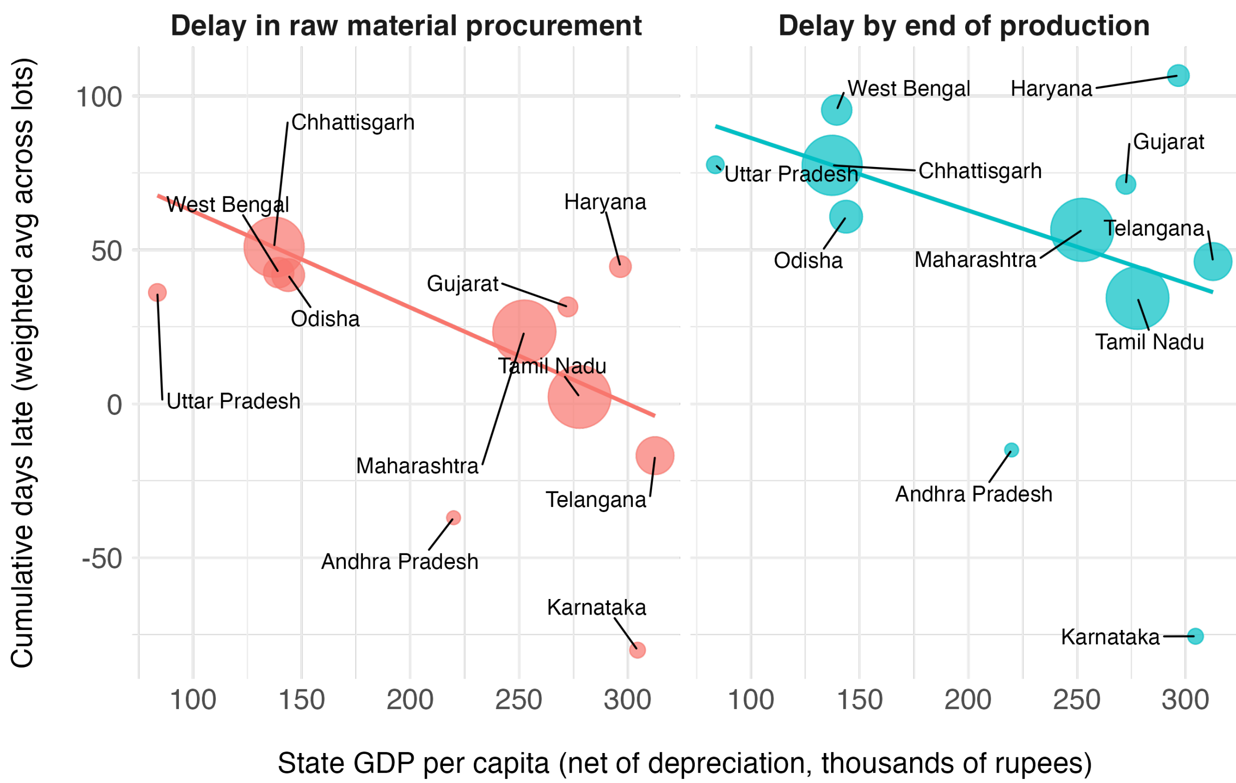

Notes: Figure shows association between state-level average delays (weighted average across all lots at suppliers in the state) and state-level GDP per capita. Left panel shows delay in sourcing raw material procurement, and right panel shows cumulative delays from sourcing raw material, starting orders, 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.

to struggle not because the relevant input is technologically complex or relationship-specific, but because they must decide when to commit scarce resources to that input. In field interviews, suppliers repeatedly attributed late sourcing to limited working capital: managers described waiting for payments from past orders before buying steel for new ones.[6]

3.2 The problem is more severe in poorer locations

We next show that delays in starting new orders are strongly associated with low economic development. Panel A of Figure 3 shows that steel is procured roughly on time in richer states but

[6]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, so suppliers delay procurement for new orders until payment arrives.

1.5-2 months late in poorer states. These patterns are largely driven by the poor performance of suppliers in the northern cluster in Bhilai, Chhattisgarh relative to clusters in the richer states of Maharashtra and Tamil Nadu. Panel B shows that the income gap is much smaller by the end of production. In other words, poor locations' inferior on-time performance arises primarily from late procurement of raw material rather than slower downstream transformation.

This association is informative for the organizational question because it again points to the start of the production process. If poorer locations were simply worse at fabrication, we would expect delays to widen during production. Instead, late starts are partly made up later on, which is more consistent with a financing or resource-allocation problem at the input stage.[7]

The descriptive evidence therefore points to a specific critical decision: timely investment in inputs for the order. We next test whether this decision is in fact shaped by short-run working capital availability.

3.3 Suppliers wait for cash to invest in inputs

To test whether working capital shapes this input decision, we ask whether procurement for order o responds to payment arrivals on the supplier's other orders. Detailed construction of variables and the leave-out instrument is in Appendix B; here we summarize the core design. If suppliers can freely finance raw material purchases, then the timing of payments from past orders should not affect when they source steel for new ones. If instead suppliers operate hand-to-mouth, then procurement should rise exactly when payment arrives.

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 $g = (\text{Steel}, \text{Dispatch}, \text{Payment}, \ldots)$. 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:

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

[7]One concern is that Figure 3 may reflect spatial heterogeneity in end-customer demand for on-time production rather than in suppliers' costs of JIT delivery. However, most construction projects physically take place in a different state from the supplier, and the end customers are firms that operate across India. As a result, suppliers in different states should face similar on-time premia.

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 total raw material sourced $Y_{ot,Steel}$ becomes $Steel_{ot}$. Our estimating equation is:

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

where we condition on planned sourcing and planned payment timing, so identification comes from shocks relative to plan rather than from the supplier's original production schedule.

Observed payments are potentially endogenous because the supplier can influence completion timing on other orders. We therefore instrument for observed payment timing using a leave-out measure of downstream project delays at other suppliers on the same project. Intuitively, if a construction project is postponed, payments to all suppliers on that project move together for reasons outside any one supplier's control. Additional details, the instrument construction, and validation exercises are in Appendix B.

Figure 4 provides intuition for the leave-out instrument: delays at other suppliers on the same downstream project strongly predict delays for the focal supplier, which is what lets us use common project shocks to isolate payment timing variation that is plausibly outside any one supplier's control.

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 OLS estimate; columns 2-4 implement the IV strategy. Our preferred estimate is in column 4: the elasticity of new steel sourcing with respect to instrumented realized

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.

Figure 4: Leave-out delays at other suppliers on common projects predict order-level delays

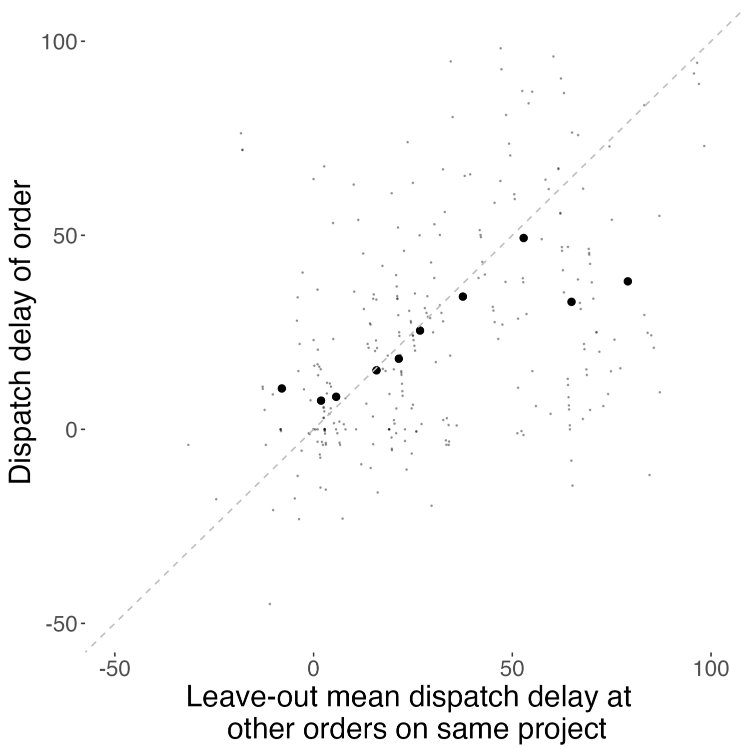

Notes: Figure shows relationship between delays (by dispatch) on an order with a particular supplier, and the average delays on orders at other suppliers that are part of the same downstream construction project. Scatterplot is shown in gray, binned means are shown in black, and the gray dashed line 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.

payments is 0.074, positive, statistically significant, and larger than the OLS estimate. Column 3 shows that the instrument is strong, with an F-statistic above 1600.

The interpretation is that some suppliers start new orders only when working capital is returned to them as liquid cash rather than remaining locked in receivables from past orders. Put differently, conditional on the plan, the supplier would have been less likely to procure steel in month t had payment not arrived in month t. That is the core evidence that the key supplier decision is constrained by short-run finance rather than being driven only by production sequencing.

The appendix develops the identifying assumptions and robustness checks in detail. Appendix Table 2 shows that the payment elasticity is larger for suppliers in poorer regions (0.125 in Central & North India vs. near-zero in the South and West) and for suppliers that receive buyer financing (0.091 vs. 0.038), consistent with credit constraints being the relevant friction. This completes the first step of the paper: the central supplier decision is investment in inputs for the order, and that decision is shaped by working capital availability.

4 Partial integration without ownership

Having identified input sourcing as the critical supplier decision, we next ask how buyers influence that decision. The answer in our setting is a bundle of organizational interventions that target the input stage: monitoring of procurement, incentive contracts tied to procurement, and direct buyer control over procurement itself. We refer to this bundle as *partial integration* because it replicates functions of vertical integration while leaving ownership boundaries formally unchanged.

4.1 Monitoring the critical stage

In standard contract theory, monitoring is one way to induce desired actions (Townsend (1979), Diamond (1984), Dye (1986)). In our setting, monitoring takes the form of visits by relationship managers, employees of the buyer who spend time at supplier sites, update internal timelines, and pressure plant managers when an order is falling behind. This gives the buyer visibility into a stage of production that would otherwise sit entirely inside the supplier.

Figure 5 shows that raw material-related site visits are concentrated precisely around the period before inputs are sourced, especially for late orders. For orders delayed by more than a month, there is a 10-15% chance of a buyer visit on any given day before sourcing occurs. The buyer is therefore not merely inspecting final output; it is closely involved in the specific upstream stage where disruptions originate.

Table 2 complements this time-series evidence using the same payment-based design as above. Monitoring rises when payments on other orders arrive, which is exactly when suppliers are newly able to invest in inputs for the focal order. The point is not only that buyers observe; it is that they monitor the critical stage when that stage becomes actionable.

Table 2: 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.

Figure 5: Probability of site visits before and after raw material is sourced

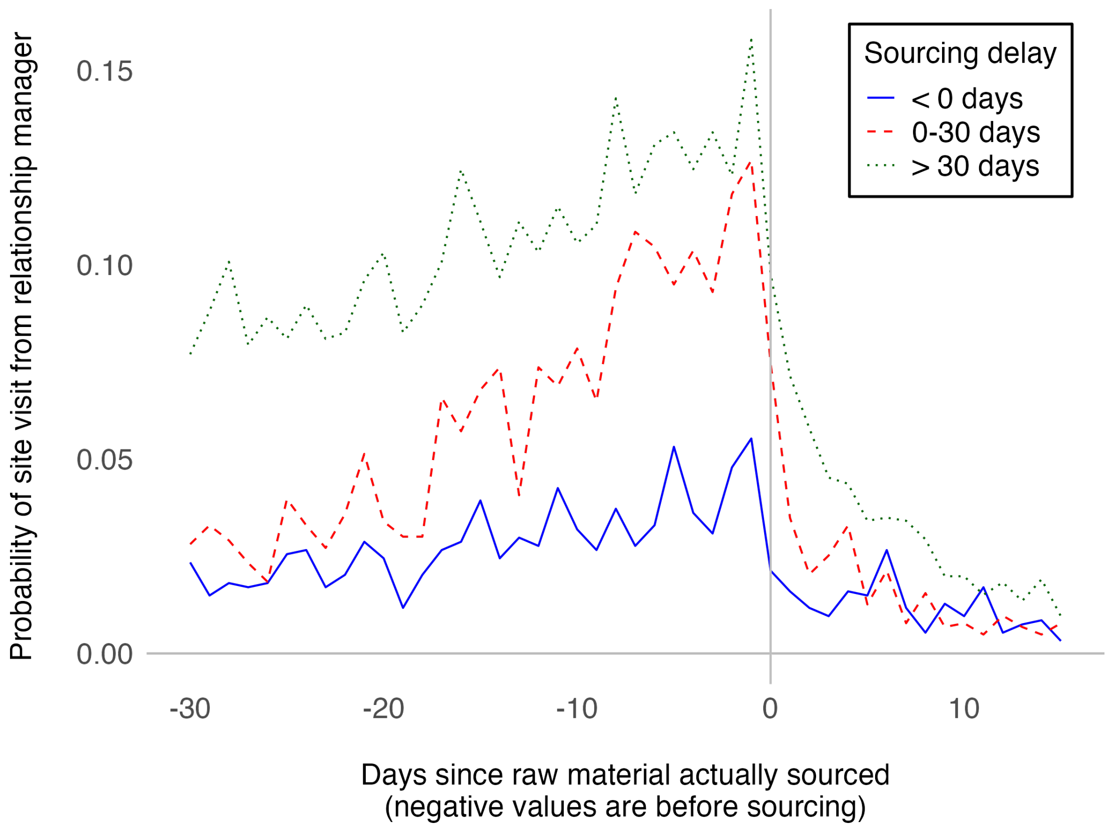

Notes: Plot shows the daily probability of a raw material-related site visit by a relationship manager as a function of time relative to when raw material is sourced. We divide lots (truckloads) into three groups: on-time sourcing (<0 days late), slightly late sourcing (0-30 days late), and late sourcing (>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.

4.2 Incentives and control in contracts

Figure 6 shows that monitoring is complemented by two formal contractual arrangements that also target the input decision. First, in 53 percent of contracts the supplier receives raw material financing from the buyer. Second, this financing takes two distinct forms. Under a *conditional advance*, the supplier sources and owns the steel, but the buyer reimburses part of the order value only after verifying that procurement has occurred. Under *buyer-sourced raw material*, the buyer procures steel itself, delivers it to the supplier, and retains property rights over that input.

Two features of these contracts are central for the paper's framing. First, suppliers never

Figure 6: Magnitude and form of buyer financing

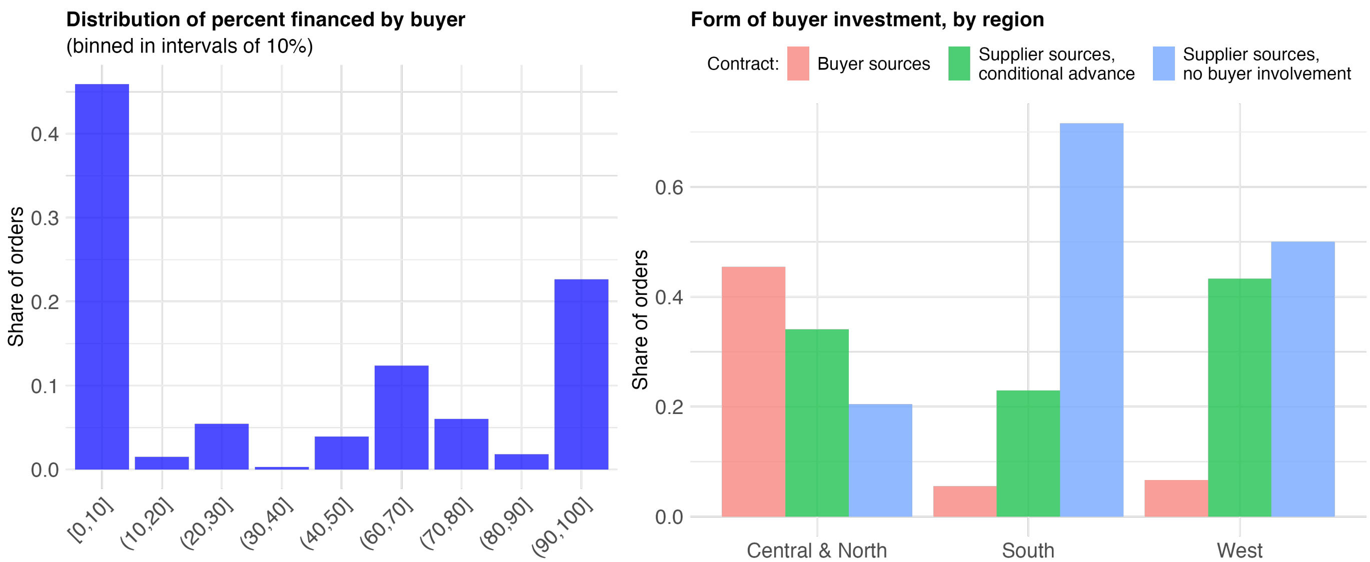

Notes: Left panel shows the distribution of input financing intensity across contracts. The percent financed is the share of order value that is paid to the supplier when steel for a lot is sourced (the remainder is paid after the completed lot is delivered to the downstream customer). This is specified 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 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 are from buyer contract terms with individual suppliers for 2019-2022.

receive unconditional advances. If the problem were only a lack of finance, the natural response would be a loan. Instead, all observed support is tied to the critical decision itself: procurement must either be verified or directly controlled. Second, the intervention is selective. The buyer does not take over fabrication, labor management, or delivery; it intervenes at the specific stage where the friction is most severe.

This is what we mean by *partial integration*. Timely sourcing of steel is sometimes incentivized, sometimes directly controlled, and in both cases monitored closely by the buyer. These are functions we typically associate with vertical integration, but here they are implemented across firm boundaries.

5 From incentives to control when the friction worsens

The previous section showed that buyers use both incentives and control to influence the key supplier decision. We now ask whether the form of partial integration responds to the severity of the underlying friction. The descriptive evidence already points in that direction: the right panel of Figure 6 shows that poorer regions use more buyer financing overall, and are also the only regions where direct buyer control of inputs is common.

To push beyond that descriptive pattern, we exploit a rapid increase in global steel prices in late 2020. This shock sharply increased the working capital required to start an order. If contractual frictions become more severe when the input decision becomes more expensive, then we should see buyers move from weaker intervention (conditional advances) to stronger intervention (direct buyer sourcing). Figure 7 presents the basic time-series evidence before turning to the formal difference-in-differences analysis.

Direct buyer sourcing is costlier than conditional advances because it requires the buyer to finance the entire input purchase and to handle the logistics of procurement. Its advantage is sharper control: the buyer owns the steel, can determine when it arrives at the supplier site, and can prevent the input from being redirected to another order. In organizational terms, direct sourcing is the strongest form of partial integration observed in the data.

We formalize the test for whether the form of partial integration responds to the severity of credit constraints using a difference-in-differences design. We estimate:

$$DirectSource_{it} = \delta_t + \alpha_{s(i)} + \beta^{DID}\left(HasAdvance_{s(i)} \cdot Post_t\right) + \epsilon_{it} \tag{4}$$

where $DirectSource_{it}$ takes value 1 if the buyer sources steel for lot i, and 0 otherwise; $HasAdvance_{s(i)}$ indicates whether the supplier ever received a conditional advance before the shock; and $Post_t$ indicates quarters after October 2020. The coefficient β^{DID} captures the additional post-shock increase in buyer sourcing for credit-constrained suppliers relative to unconstrained ones. The identifying assumption is parallel trends: absent the steel price shock, the probability that constrained firms receive buyer-sourced material would have evolved similarly to that for unconstrained firms. This assumption is supported by the fact that no buyer sourcing existed before the shock for *either* group – by analyzing internal records (email communications and manager interviews), we verified that the buyer had no plans to procure steel itself until prices rose in late 2020. Thus the counterfactual level of buyer sourcing is zero for both groups, and the counterfactual trend is indeed parallel. Additional specification details and alternative proxy measures are developed in Appendix B.

Figure 7: Time series of steel price and contract terms

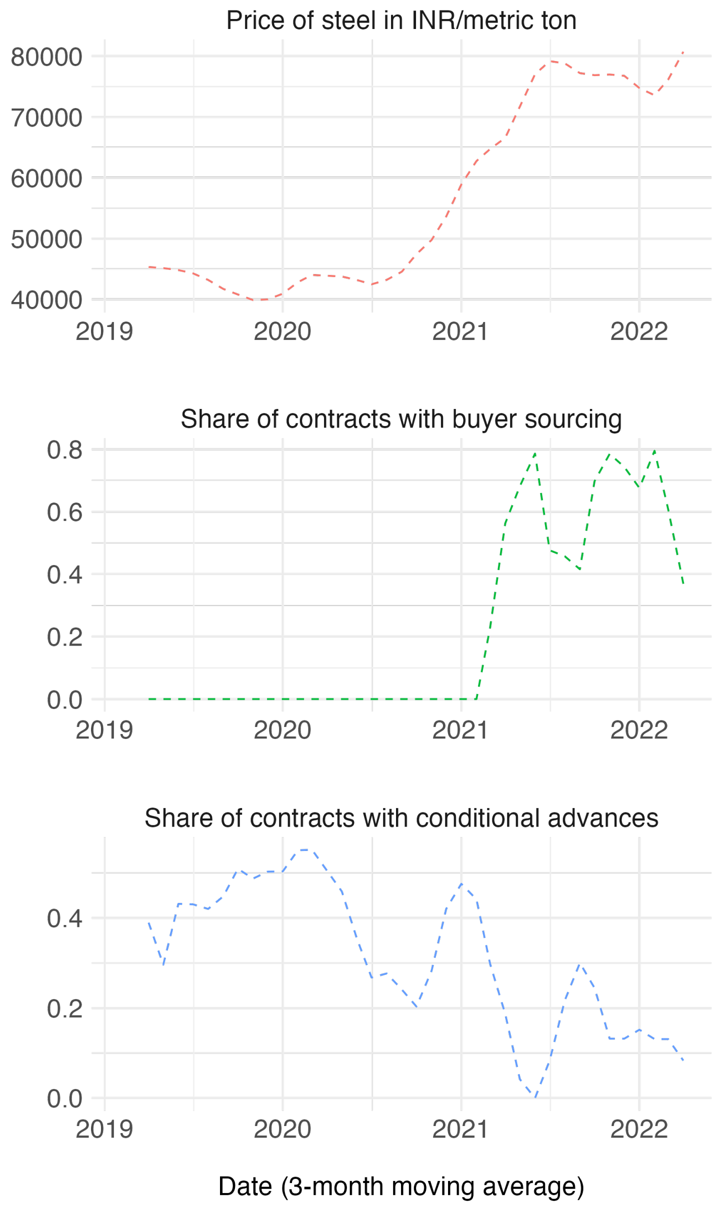

Figure shows three different time series from 2019-2022, all as three-month moving averages calculated from monthly data. Top panel shows Indian spot market steel prices, obtained from the Ministry of Steel, Government of India. Middle panel shows share of contracts (by date of contract signing) in which 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.

To understand *which suppliers* switch to buyer sourcing, we estimate an event-study version of equation (4), 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.

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

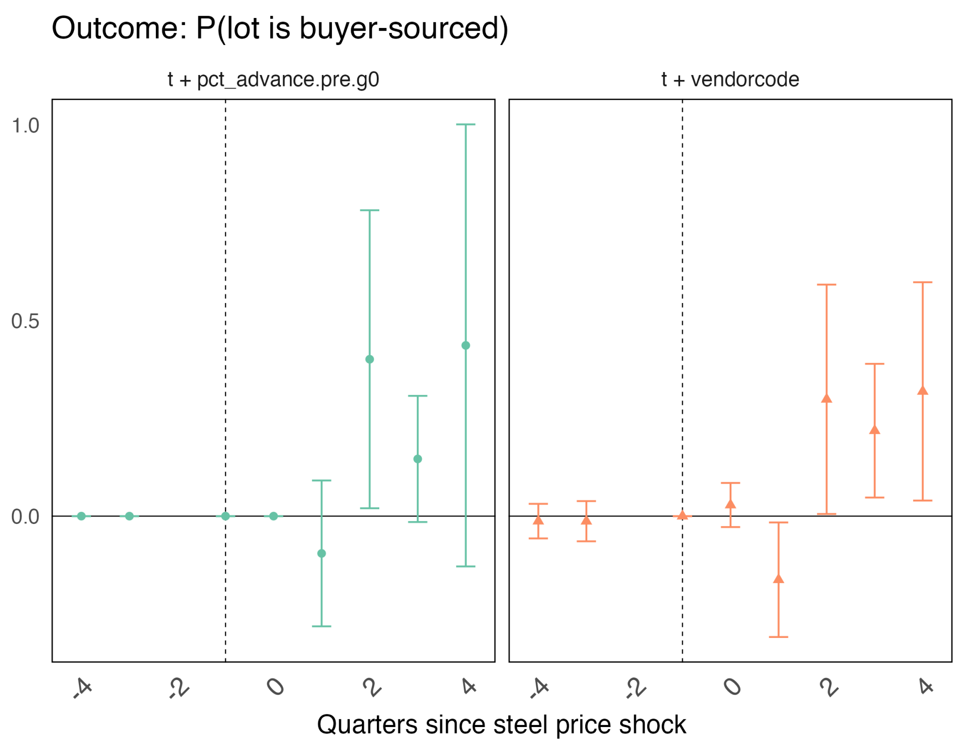

Notes: Figure shows event-study effects of supplier-level exposure to an input price shock on the probability of using buyer-sourced raw material. A supplier is exposed if it ever used conditional advances from the buyer *before* the price shock. Regressions are estimated at the supplier-quarter level. 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).

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

Table 3: 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, Table 3 shows that this switch is concentrated among suppliers that were already more constrained before the shock. Pre-shock use of conditional advances, the size of those advances, and pre-shock delays all predict post-shock buyer sourcing. The cleanest summary is in columns 2 and 5: a 10 percentage point higher conditional-advance share before the shock raises the probability of switching to buyer sourcing afterward by 4.8 percentage points. The form of partial integration therefore responds systematically to the severity of the friction.

6 A framework of partial integration

The empirical evidence raises three questions that empirics alone cannot answer. First, why does the buyer intervene selectively at the input stage rather than taking over production as a whole? Second, why are advances always monitored rather than unconditional? Third, why does the buyer switch from incentive-based to control-based intervention when the friction intensifies? This section builds a model to organize those facts, showing that the three questions have a common answer rooted in stage-specific noncontractibile decisions.

6.1 Setup

A supplier must procure an input and then produce output for a buyer. The working-capital friction is concentrated at the input stage, while production incentives matter downstream. Figure 9 summarizes the timing.

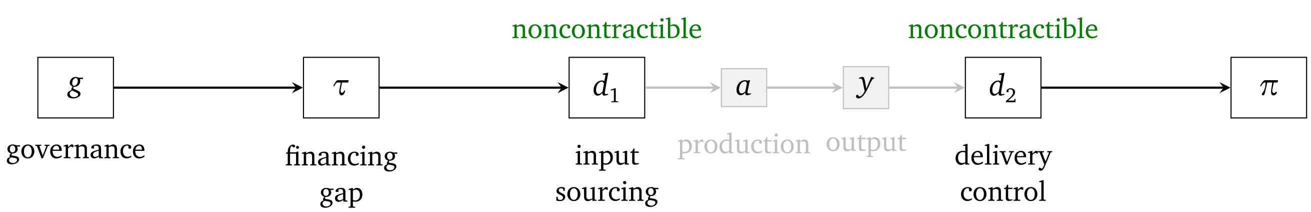

Figure 9: Sequential production with stage-specific control.

Technology and financing gap. One unit of output requires raw material worth m. The supplier has liquidity $l < m$, so the order has a financing gap

$$\tau \equiv m - l > 0.$$

Under governance structures U and A (cash advance regimes, defined below), the buyer advances τ to cover this gap. The advance is recovered through a deduction from the delivery payment when the order completes successfully, so in the no-diversion case the buyer bears no net input cost. Under buyer sourcing (B) or full integration (VI), the buyer procures the input directly at administrative cost F. The key state variable throughout is τ: it is the financing gap that creates the contracting problem.

At d_1, the critical decision is whether buyer-provided support is actually used to procure raw material for this order–or diverted. After inputs are secured, timely completion requires high supplier effort at private cost $c > 0$ and creates gross value $\bar{V} > 0$ for the buyer.

Downstream incentives. If the supplier retains downstream control d_2, a second order arrives with probability $q \in [0, 1]$. Completing the current order on time allows the supplier to pursue that outside opportunity, which yields private rent $R > 0$. We assume $R < c$: downstream market incentives sharpen but do not fully replace the need for incentive pay. Under U, A, or B, the buyer induces timely completion with a delivery bonus w satisfying $w + qR \geq c$. The minimum such bonus is $w^*(q) = c - qR$, so the buyer's continuation value after inputs are secured is

$$V(q) \equiv \bar{V} - w^*(q) = \bar{V} - c + qR.$$

If the buyer also takes d_2 under full integration, the supplier loses this market incentive. The buyer can still induce completion by paying c directly, but now obtains a direct downstream control benefit $G \geq 0$. The continuation value under full integration is therefore $\bar{V} - c + G$. We focus on the empirically relevant region where $V(q) > F > K > 0$.

Governance structures. The buyer chooses among four organizational forms, ordered by the degree of control over the supplier's decisions.

1. *Unconditional cash (U):* The buyer advances τ and the supplier retains both d_1 and d_2. No verification takes place.

2. *Conditional advance with monitoring (A):* The buyer advances τ and pays a monitoring cost $K > 0$ to verify that the supplier has procured the input, conditioning the advance on that verification. The supplier still controls d_1, but monitoring reduces the scope for diversion.

3. *Buyer-sourced input (B):* The buyer purchases the input directly, taking control of d_1 at administrative cost $F > K$, while leaving d_2 with the supplier. This corresponds to the "direct sourcing" contracts observed in the data.

4. *Full vertical integration (VI):* The buyer takes both d_1 and d_2, using the same input-sourcing technology as B (cost F) but also claiming downstream control.

Misallocation risk. When the supplier retains d_1, buyer-funded working capital can be diverted to other orders or otherwise misused. Under unconditional cash, diversion occurs with probability $\rho(\tau)$, where $\rho(0) = 0$, $\rho'(\tau) > 0$, and $\rho(\tau) < 1$. The assumption $\rho' > 0$ is central: a larger financing gap creates a larger temptation to divert, and this is how input costs affect organizational form. Under conditional advances, monitoring reduces diversion risk by a fraction $\mu \in (0,1)$, so the diversion probability falls to $(1-\mu)\rho(\tau)$. Monitoring helps but does not perfectly prevent misuse–the supplier can still redirect funds after procurement has been verified. Under buyer sourcing or full integration, the buyer owns the input decision, eliminating diversion at d_1.

Payoffs. If diversion occurs, the order does not start and the buyer loses the unrecovered advance τ. If diversion does not occur, inputs are secured and the buyer obtains the continuation value. The buyer can always walk away for payoff 0. Expected payoffs are therefore:

$$\begin{aligned}
\Pi_U(\tau) &= [1-\rho(\tau)]\,V(q) - \rho(\tau)\,\tau, \\
\Pi_A(\tau) &= [1-(1-\mu)\rho(\tau)]\,V(q) - (1-\mu)\rho(\tau)\,\tau - K, \\
\Pi_B(\tau) &= V(q) - F, \\
\Pi_{VI}(\tau) &= \bar{V} - c + G - F.
\end{aligned}$$

Note that Π_B and Π_{VI} do not depend on τ: once the buyer controls the input decision, the financing gap is irrelevant. By contrast, Π_U and Π_A are decreasing in τ because a larger gap raises both diversion risk and the loss when diversion occurs. This asymmetry drives the switch from incentives to control.

6.2 Implications for partial integration

We derive three results from the model, corresponding to the three questions that motivated the model. All proofs are in the appendix.

6.2.1 Why partial integration can beat vertical integration

Why does the buyer not simply acquire the supplier? In the model, the working-capital friction is at d_1, so both B and VI solve the upstream problem equally well. The difference is downstream: leaving d_2 with the supplier preserves market incentives worth qR, while taking d_2 gives the buyer direct control benefit G.

Proposition 1 (Partial integration versus full integration)**.** *For any financing gap τ,*

$$\Pi_B(\tau) - \Pi_{VI}(\tau) = qR - G.$$

Hence partial integration becomes more attractive as q or R rises and less attractive as G rises. If $G < R$, there exists a threshold $q^* = G/R \in (0,1)$ such that the buyer prefers full integration for $q < q^*$ and partial integration for $q > q^*$.

In our setting, the suppliers are small firms that serve multiple buyers and maintain ongoing relationships in local industrial clusters. Their downstream market incentives – repeat business, reputation, and the option to take other orders – are substantial, making the $qR > G$ region empirically relevant. The buyer accordingly intervenes at the input stage while leaving the supplier autonomous over production and delivery.

6.2.2 Why unconditional advances are absent

A second puzzle from the data is that the buyer never simply lends money to constrained suppliers. Every observed financial intervention is tied to input procurement: either the buyer verifies that procurement occurred (conditional advances) or the buyer procures the input itself (direct sourcing). In the model, the distinction between unconditional and conditional support is monitoring: conditional advances are unconditional cash plus a verification technology that reduces diversion risk at cost K.

Proposition 2 (Monitored advances dominate unconditional cash)**.** *If*

$$\mu\rho(\tau)(V(q) + \tau) \geq K,$$

then $\Pi_A(\tau) \geq \Pi_U(\tau)$, with strict inequality when the condition is strict.

The left-hand side is the expected reduction in diversion losses from monitoring: the probability reduction $\mu\rho(\tau)$ times the stakes $V(q) + \tau$. The right-hand side is the monitoring cost. Because the stakes grow with τ, the condition is easier to satisfy when the financing gap is large – precisely the region where financing constrained suppliers is worthwhile. So whenever the friction is severe enough to warrant intervention, monitoring more than pays for itself, and unconditional cash is dominated.

6.2.3 Why higher input costs shift the buyer from incentives to control

The final result addresses the steel-price shock evidence in Section 5. A rise in m or a fall in l raises $\tau = m - l$. Buyer sourcing (B) is unaffected, since the buyer owns the input decision regardless of τ. But the residual diversion risk under monitored advances (A) grows more costly.

Proposition 3 (Switch from conditional advances to buyer-sourced inputs). *Assume $F > K$, $\rho'(\tau) > 0$ for all $\tau > 0$, and there exists $\bar{\tau} > 0$ such that*

$$(1-\mu)\rho(\bar{\tau})(V(q)+\bar{\tau}) > F - K.$$

Then there is a unique threshold $\tau^* > 0$ such that monitored conditional advances are preferred to buyer sourcing for $\tau < \tau^*$, while buyer sourcing is preferred for $\tau > \tau^*$. Since $\tau = m - l$, a rise in m or a fall in l shifts the buyer from incentives to control.

This is the model's counterpart to the difference-in-differences result in Section 5. The 2020 steel-price shock raised m for all suppliers, pushing τ past τ^* for the most constrained firms. Suppliers that previously received conditional advances (high τ but below τ^*) are then the ones that switch to buyer-sourced inputs, consistent with the evidence in Figure 8.

6.3 Takeaways from the model

All three results from the model relate to the same logic: the buyer's problem is how deeply to intervene at the input stage, given that intervention at the production stage would sacrifice valuable supplier incentives. When the financing gap is moderate, monitoring procurement is sufficient (Proposition 2). When the gap widens, the buyer takes over procurement directly (Proposition 3). Full vertical integration is unnecessary whenever the downstream incentive qR exceeds the control benefit G (Proposition 1). Conditional advances and buyer-sourced inputs thus represent two points on a continuum of buyer control over a single decision, with τ determining which point is optimal.

Two extensions in Appendix D round out the picture. Monitoring and direct sourcing both require close engagement with the supplier, "relational" buyers with lower costs of engagement (K and F) can more profitably manage constrained suppliers (Proposition 4). And if outside finance closes the supplier's liquidity gap, the diversion problem at d_1 disappears, removing the rationale for buyer intervention entirely (Proposition 5). Partial integration in this setting exists because of a specific market failure; addressing its source with better working capital markets would eliminate it.

The broader point is that efficient integration should be selective when production involves multiple stages with distinct contracting problems. The buyer controls the stage where the friction binds, and leaves remaining stages to the supplier when doing so preserves useful incentives. Partial integration, in this view, is the organizational form implied by a stage-specific friction in a multi-stage production process.

7 Conclusion

This paper studies how buyers influence critical supplier decisions without vertical integration. Using detailed production timelines and contract data from Indian manufacturing supply chains, we show that the key supplier decision in this setting is timely investment in raw-material inputs, and that this decision is shaped by short-run working capital availability. Suppliers start new orders when cash from old ones arrives.

Buyers respond through *partial integration*: monitoring of the input stage, incentive contracts tied to procurement, and direct buyer sourcing of inputs. These arrangements look like vertical integration at one critical stage while leaving the rest of the supplier's production process outside the firm boundary. When the friction worsens, buyers move from incentives to control. A steel-price shock increases working capital needs and leads the buyer to substitute from conditional advances toward direct sourcing for the most constrained suppliers.

The model rationalizes this pattern as a selective allocation of control rights. The buyer wants tighter control over the noncontractible input decision, but leaves the rest of production to the supplier in order to preserve market incentives for timely completion. Partial integration emerges when upstream control matters more than downstream control. The results suggest that firm boundaries in supply chains are shaped by *which specific decision* requires control, and not merely by whether control is needed at all.

Because partial integration requires close involvement – monitoring of supplier sites, verification of input procurement, and sometimes direct sourcing – it follows that relational buyers with lower costs of engagement are better positioned to source from constrained suppliers. This logic implies a comparative advantage for embedded domestic buyers, intermediaries, and firms with Japanese-style supplier management over arms-length spot buyers in weak institutional settings. A policymaker seeking to promote supply chain integration with developing countries might therefore prioritize attracting firms with relational sourcing capabilities, or invest in intermediary institutions and short-term trade finance markets that reduce the costs of partial integration for all buyers.

References

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

Aghion, Philippe and Jean Tirole (1997). "Formal and Real Authority in Organizations". In: *Journal of Political Economy* 105.1, pp. 1–29.

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

Amberg, Niklas et al. (2021). "Curbing Shocks to Corporate Liquidity: The Role of Trade Credit". In: *Journal of Political Economy* 129.1, pp. 182–242.

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

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.

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.

Cajal-Grossi, Julia, Rocco Macchiavello, and Guillermo Noguera (2023). "Buyers' Sourcing Strategies and Suppliers' Markups in Bangladeshi Garments". In: *Quarterly Journal of Economics* 138.4, pp. 2391–2450.

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

Costello, Anna M. (2020). "Credit Market Disruptions and Liquidity Spillover Effects in the Supply Chain". In: *Journal of Political Economy* 128.9, pp. 3434–3468.

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.

Feenstra, Robert C and Gordon H Hanson (2005). "Ownership and control in outsourcing to China: Estimating the property-rights theory of the firm". In: *The Quarterly Journal of Economics* 120.2, pp. 729–761.

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.

Grant, Robert M. (2002). *Birds Eye and the U.K. Frozen Food Industry*. Case study. URL: https://www.blackwellpublishing.com/KOLB/downloads/private/grant/birdeye.pdf.

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.

Hortacsu, Ali and Chad Syverson (2007). "Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices". In: *Journal of Political Economy* 115.2, pp. 250–301.

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.

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

McGee, Patrick (2025). *Apple in China: The Capture of the World's Greatest Company*. New York: Scribner. ISBN: 9781668053379. URL: https://www.simonandschuster.com/books/Apple-in-China/Patrick-McGee/9781668053379.

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*

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 from the buyer. Columns 3, 4 and 5 subset to orders with suppliers in each region: Central & North (the poorest region), South (the wealthiest), and the West (intermediate). The payment elasticity is larger for orders with advances and for suppliers in poorer regions, consistent with credit constraints being the relevant friction.

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 the model makes clear, rationalizing the observed 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 $\sum_{o'} Payment_{o't}^{PLAN}$.[8]

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 4 provides some intuition as to why: at the lot level, delay

[8] 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:

$$\mathbb{E}[Payment_{ot}^{INST}] = Payment_{ot}^{PLAN} \tag{5}$$

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.

length at other suppliers s' (on the same construction project c) strongly predicts delay length 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, especially the prevalence of conditional advances and buyer-sourced inputs in Figure 6, 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.[9] 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. Under buyer sourcing, the buyer controls d_1 but leaves d_2 with the supplier. Since timely completion is then supported by downstream market incentives worth qR, the buyer only needs to pay the minimum delivery bonus $c - qR$. Hence

$$\Pi_B(\tau) = \bar{V} - (c - qR) - F = \bar{V} - c + qR - F.$$

Under full integration, the buyer also takes d_2. This removes the supplier's downstream market incentive, so the buyer must pay c directly to induce timely completion, but receives control benefit G. Thus

$$\Pi_{VI}(\tau) = \bar{V} - c + G - F.$$

Taking the difference gives

$$\Pi_B(\tau) - \Pi_{VI}(\tau) = qR - G.$$

This expression is increasing in q and R, and decreasing in G. If $G < R$, then $q^* = G/R \in (0,1)$ and $\Pi_B(\tau) > \Pi_{VI}(\tau)$ exactly when $q > q^*$. This yields the stated comparative static between partial and full integration. $\square$

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

Proof of Proposition 2. Using the payoff definitions in the text,

$$\begin{aligned}\Pi_A(\tau)-\Pi_U(\tau) &= [1-(1-\mu)\rho(\tau)]V(q)-(1-\mu)\rho(\tau)\tau-K \\ &\quad -([1-\rho(\tau)]V(q)-\rho(\tau)\tau) \\ &= \mu\rho(\tau)V(q)+\mu\rho(\tau)\tau-K \\ &= \mu\rho(\tau)(V(q)+\tau)-K.\end{aligned}$$

Hence $\Pi_A(\tau)\geq\Pi_U(\tau)$ if and only if

$$\mu\rho(\tau)(V(q)+\tau)\geq K.$$

The inequality is strict when the condition is strict. Thus the value of monitoring is exactly the reduction in expected diversion losses, and unconditional support is dominated whenever that benefit exceeds the monitoring cost. □

Proof of Proposition 3. Define the payoff difference between buyer sourcing and monitored advances by

$$D(\tau)\equiv\Pi_B(\tau)-\Pi_A(\tau).$$

Using the payoffs in the text,

$$\begin{aligned}D(\tau) &= (V(q)-F)-([1-(1-\mu)\rho(\tau)]V(q)-(1-\mu)\rho(\tau)\tau-K) \\ &= (1-\mu)\rho(\tau)(V(q)+\tau)+K-F.\end{aligned}$$

At $\tau=0$, this becomes

$$D(0)=K-F<0$$

because $F>K$ by assumption. So for low financing gaps, monitored advances are preferred.

Next, differentiate with respect to τ:

$$D'(\tau)=(1-\mu)\big[\rho'(\tau)(V(q)+\tau)+\rho(\tau)\big].$$

Since $1-\mu>0$, $\rho'(\tau)>0$, $V(q)+\tau>0$, and $\rho(\tau)\geq 0$, we have $D'(\tau)>0$ for all $\tau>0$. Thus $D(\tau)$ is strictly increasing in the financing gap.

By assumption, there exists $\bar{\tau}>0$ such that

$$(1-\mu)\rho(\bar{\tau})(V(q)+\bar{\tau})>F-K,$$

which is equivalent to $D(\bar{\tau}) > 0$. Since $D(0) < 0$, $D(\bar{\tau}) > 0$, and $D(\tau)$ is continuous and strictly increasing, there exists a unique threshold $\tau^* \in (0, \bar{\tau})$ with $D(\tau^*) = 0$. Therefore monitored advances are preferred for $\tau < \tau^*$, while buyer sourcing is preferred for $\tau > \tau^*$.

Finally, because $\tau = m - l$, a rise in input cost m raises τ while a rise in supplier liquidity l lowers it. Hence higher input costs and tighter supplier liquidity both move the buyer from incentives to control. □

D Additional results from the model

The main propositions characterize the buyer's optimal governance choice for a given supplier. Two natural extensions ask how the buyer's *type* and the supplier's *access to outside finance* affect the scope for partial integration.

D.0.1 Comparative advantage of relational firms

The main results imply a comparative advantage for buyers with lower costs of close supplier management – what we call "relational" buyers. In the model, such buyers have lower monitoring costs (K) and lower direct-sourcing costs (F).

Proposition 4 (Relational firms can sustain partial integration for a wider set of suppliers)**.** *Consider two buyers $i \in \{1, 2\}$ with $K_1 < K_2$ and $F_1 < F_2$. Then, for every financing gap τ,*

$$\max\{0, \Pi_{A,1}(\tau), \Pi_{B,1}(\tau)\} \geq \max\{0, \Pi_{A,2}(\tau), \Pi_{B,2}(\tau)\},$$

with strict inequality whenever buyer 2 serves the supplier using either A or B. Hence buyer 1 can profitably source from a weakly larger set of constrained suppliers.

Proof. For each τ,

$$\Pi_{A,1}(\tau) - \Pi_{A,2}(\tau) = K_2 - K_1 > 0$$

because the two buyers face the same diversion technology but buyer 1 has lower monitoring cost. Likewise,

$$\Pi_{B,1}(\tau) - \Pi_{B,2}(\tau) = F_2 - F_1 > 0.$$

Therefore buyer 1 gets weakly higher payoff from every intervention option and strictly higher payoff from any option that uses monitoring or direct sourcing. Taking the maximum over

$\{0, \Pi_A, \Pi_B\}$ preserves that ordering, so buyer 1 can profitably serve every supplier that buyer 2 can serve, and possibly more. $\square$

Thus, relational firms have a comparative advantage exactly when the environment requires partial integration. This connects to the conclusion's discussion of intermediaries and Japanese-style supplier management: firms that invest in the relational infrastructure to monitor and source on behalf of suppliers can profitably serve more constrained partners.

D.0.2 Completing working capital markets removes the need for partial integration

Finally, consider what happens when the supplier can access outside finance for input purchases. In the model, partial integration exists only because the buyer must solve a contracting problem at the input stage while funding a liquidity-constrained supplier. If an outside lender fills the financing gap, the contracting problem disappears and with it the need for buyer intervention.

Proposition 5 (Completing working capital markets removes the need for buyer intervention)**.** *Suppose an outside lender fills the financing gap at sufficiently low cost, so the supplier can procure the input without buyer support. If, in addition, $qR > G$, then market governance, under which the supplier keeps both stages and the buyer does not intervene, strictly dominates A, B, and VI.*

Proof. If outside finance closes the gap, the buyer no longer needs to provide cash support or source the input. In that case the supplier procures the input without buyer intervention, so market governance yields continuation value $V(q)$.

Under monitored advances, the buyer would obtain $V(q) - K$, because the diversion problem has already been removed but the monitoring cost would still be paid. Under buyer sourcing, the buyer would obtain $V(q) - F$. Under full integration, the buyer would obtain $\bar{V} - c + G - F$.

Since $K > 0$ and $F > 0$, we have

$$V(q) > V(q) - K \qquad \text{and} \qquad V(q) > V(q) - F.$$

Moreover, in the region where outside finance removes the need for input-stage control, market governance also dominates full integration whenever preserving the supplier's downstream market incentive is more valuable than direct downstream control, i.e. whenever $qR > G$. In that case

$$V(q) = \bar{V} - c + qR > \bar{V} - c + G - F = \Pi_{VI}.$$

Hence market governance strictly dominates all intervention options once outside finance removes the financing gap. $\square$