Partial integration and the boundaries of the firm

Aroon Narayanan
MIT

Vishan Gandhi Nigam
Yale

September 7, 2026

Abstract

How can buyers in supply chains influence upstream supplier decisions – to start orders on time, enforce quality, and avoid resource diversion – that are often noncontractible? The typical answer is vertical integration: buyers can forgo the market and bring suppliers inside the firm. Using detailed contracting and production data from a large Indian manufacturing intermediary, we document an alternative remedy, *partial integration* (PI), consisting of targeted buyer interventions across firm boundaries. In our setting, the critical decision is supplier investment – of cash-on-hand in inputs – which buyers influence via supplier site visits (*monitoring*), direct buyer ownership of suppliers' inputs (*control*), and contractual advance payments (*incentives*) linked to input investment. We establish via three distinct natural experiments that (1) incentives and control are substitutes, (2) monitoring increases following supplier cash inflows, and (3) expansions in supplier credit access reduce buyer control of suppliers' inputs. Interpreted through a stylized model of a supply chain, results (1)-(3) are consistent with PI as a response to noncontractibility (specifically, resource diversion risk). Other motives, such as efficient buyer specialization in input-related tasks, cannot explain these results. Finally, we show in India-wide production data that PI is prevalent, comprising nearly 30% of manufacturing VA for the average plant. Taken together, our results imply that – in contrast with classic make-or-buy models – firms in modern supply chains can address noncontractibilities within the broad framework of a market transaction.

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.nigam@yale.edu.

1 Introduction

Many decisions in modern supply chains are difficult to contract on. 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, starts orders late, mistreats its workers, releases pollution, or diverts Apple-provided financing and tooling to competing customers. How can buyers shape these fundamentally noncontractible supplier decisions?

The classic answer to this question is that firms can forego the market in favor of vertical integration (VI). For instance, Coase (1937), Williamson (1979), and Grossman and Hart (1986) respectively argue that VI removes the costs of specifying contracts, lets buyers directly implement and verify desired decisions, and mitigates hold-up.[1] A parallel empirical literature shows that firms vertically integrate to resolve incentive problems in contexts from Peruvian fishmeal (Hansman, Hjort, León-Ciliotta, et al., 2020) and Kenyan coffee (Macchiavello and Morjaria, 2021) to U.S. cement (Hortacsu and Syverson, 2007) and the universe of U.S. imports (Alfaro et al., 2025).

This paper shows that firms have an alternative: rather than forgoing the market, buyers can adopt an intermediate organizational form that consists of targeted interventions in supplier operations *across* firm boundaries. These interventions take several forms, including in-person monitoring, contractual incentives, and direct control of internal supplier decisions. We refer to them collectively as *partial integration* (PI) because they mirror activities of VI firms, but target specific sources of production failure within suppliers who remain formally independent. In documenting that buyers and suppliers blur firm boundaries in this way, we show that vertical integration is not necessary to address noncontractibilities.

We build our argument in three stages. First, we document partial integration using novel contracting and production data from an Indian manufacturing intermediary. Second, we isolate the key motive for PI: using a model and two distinct natural experiments, we show that PI is specifically a response to contracting frictions rather than a form of efficient buyer specialization in comparative-advantage tasks. Third, to generalize our results, we document the prevalence of PI and implement a third quasi-experimental test of the model using panel data from India's *Annual Survey of Industries*.

Our analysis requires that we observe both the internal supplier decisions at risk and the governance instruments buyers deploy to shape them. Since standard firm datasets exclude such

[1] A second class of remedies, relational contracts, also deviate from spot markets by requiring exclusivity.

information[2], we instead leverage novel administrative data on production timelines, contract terms (including decision rights and incentive structures), and buyer-supplier interactions from a large intermediary in the Indian metal fabrication sector. Using this data, we first establish that the key upstream decision at risk (i.e., the source of production failure) in our setting is that *suppliers fail to invest limited working capital in inputs for a buyer's order.* This is in part because suppliers are working capital-constrained: we show that input investment responds to cash inflows from past orders, including when we use common downstream payment delays as an IV for payment timing.

Our key empirical finding is that buyers intervene in this key supplier decision in costly ways that deviate from typical views of market transactions. Two such interventions are directly visible from contract terms: buyers *incentivize* input investment via conditional advance payments, released only when inputs physically reach supplier sites; and *control* investment by directly sourcing and owning suppliers' inputs. A third intervention occurs within each contract period: buyers *monitor* input investment via frequent visits to supplier factory floors, which occur multiple times a week and escalate when input investment is delayed. Arms-length buyers thus behave as if they are managing a vertically integrated firm, within which incentives, control, and monitoring of internal supplier decisions would be unsurprising.[3]

Having documented how buyers partially integrate their suppliers, we next separate between two plausible motives for PI with different economic implications. One natural interpretation is that PI is a form of *specialization* in which buyers take on tasks that they can perform more cheaply than their suppliers, such as the provision of inputs, financing, or managerial expertise. An alternative is that buyers intervene to stop suppliers from diverting cash or inputs to other orders – in other words, to internalize the *noncontractibility* of suppliers' resource allocation decisions. The consequences of PI differ significantly across these two interpretations. Under specialization, PI is an efficient division of labor with no bearing on firm boundaries. Under noncontractibility, PI is a previously undocumented alternative response to the same market friction that motivates vertical integration.

To test if PI is specifically a response to noncontractibility, we build a parsimonious model of contracting with sequential stages that rationalizes partial integration. At the input stage, the supplier can divert scarce resources (here, working capital) to orders for other buyers. At the delivery stage, new outside-market opportunities incentivize suppliers to complete existing or-

[2] A notable exception is Chinese trade data, which includes data on buyer- vs. supplier-owned inputs (see Manova and Yu (2016))

[3] A further resemblance to VI is that, in a process analogous to organizational design, buyers tailor interventions to each supplier, for instance by choosing between conditional advances (*incentives*) and raw materials (*control*), and by customizing the size of advance payments to match supplier credit needs.

ders, which discourages full vertical integration. Our theoretical framework can rationalize the mix of upstream interventions we see – site visits (monitoring), conditional advance payments (incentives), and direct ownership of supplier's inputs (control) – and predicts how those forms shift as resource requirements vary across suppliers and over time. To test the model, we derive and verify three predictions for PI form that hold if and only if resource allocation decisions are noncontractible.

The first prediction, which we validate using a differences-in-differences (DID) design, is that *two PI forms – ownership of supplier's inputs and advance payments – are substitutes.* Specifically, we study how contract form responded to a quasi-experimental shock in late 2020 to global steel prices, and thus to suppliers' working capital needs. We show that following the shock, buyers take direct ownership of suppliers' inputs (*control*) disproportionately for suppliers receiving conditional advances (*incentives*) before the shock.[4] This result reflects the classic insight (e.g., Aghion and Tirole, 1997) that control and incentives are direct and indirect methods respectively of influencing an agent's decision, which in our setting is supplier input investment.[5] In contrast, without noncontractibility, input cost shocks would change the magnitude but not the form of buyer intervention.

As supporting evidence, we verify descriptively that *all advance payments are conditional* on supplier decisions, paying out only after buyers physically verify that steel has reached supplier sites. Intuitively, when facing diversion risk in the presence of noncontractibility, buyers only pay after input procurement because raw material is harder to resell or redirect across buyers than cash.[6] In contrast, without diversion risk, buyers would simply provide unconditional cash advances to suppliers.

The second prediction is that, within a contract period, *buyer presence at supplier sites (monitoring) responds to the availability of internal resources*, which in our context means cash. Exploiting the same IV design we used to test for credit constraints – which predicts cash inflow timing using leave-out delays on common projects – we show that face-to-face visits to supplier sites rise when cash payments for past orders arrive at the supplier. This result, which implies that buyers monitor supplier sites in order to influence allocation decisions even when the resource (cash) is not buyer-provided, is consistent with the influencing of noncontractible decisions but not with alternative motives driven by buyer specialization in input procurement.

Having argued that noncontractibility motivates PI, we next examine its advantages relative

[4]This substitution across contract forms also holds *within-supplier*.

[5]In particular, while direct control lets buyers automatically implement desired investment, indirect incentives only succeed if the resale value of physical inputs is still low – something that is increasingly unlikely as input costs rise.

[6]This logic underpins many models of in-kind finance (see for instance Burkart and Ellingsen, 2004).

to full VI as a solution to contracting frictions. The model clarifies that PI may dominate VI because it allows buyers to *selectively preserve market incentives*. In our case, PI allows buyers to shut down input-stage market incentives, which can lead to resource diversion, while preserving delivery-stage market incentives, which encourage order completion once resources are sunk.[7]

The argument so far rests on a single supply chain, so we proceed by showing that partial integration is a much broader feature of Indian manufacturing. Using the *Annual Survey of Industries* (ASI) rounds from 2015-2019, which record one form of partial integration – direct ownership of suppliers' inputs – we calculate lower-bound PI shares of manufacturing value-added for each plant. 32% of plants earn some value added under partial integration, and among these plants the PI share averages 28%. PI shares are also higher in working capital-intensive industries, where cash diversion temptations may be higher.

Finally, to understand if widespread partial integration in the *ASI* is also a response to contracting frictions, we test a third model prediction using a policy-induced credit shock. If cash diversion risk motivates buyers to own and source suppliers' inputs, improved supplier credit access – which reduces resource needs – should reduce buyer ownership. In contrast, if buyers source suppliers' inputs due to comparative advantages in input sourcing, supplier credit access should not affect their involvement. Exploiting the differential expansion in firm credit access due to India's *2006 MSMED Act* (see Rotemberg, 2019), we show exposed firms reduce sales with buyer-owned inputs even while increasing other sales. As a result, the odds of manufacturing via PI falls significantly, consistent with our model of PI under noncontractibility.

Taken together, these economy-wide effects and the earlier evidence from Indian metal fabrication imply that the relevant boundary of authority in supply chains is decision-specific. Effective control can be allocated stage by stage according to where contracting frictions bind, deployed via several instruments – including monitoring, incentives, and direct control – and shifted across instruments as the friction intensifies. Arrangements that appear "non-integrated" in terms of formal ownership may therefore involve substantial buyer authority over specific supplier decisions.

Our results also suggest that – because effective control can also be exercised without ownership – contracting frictions are both less costly and more common than a focus on vertical integration alone would suggest. They are less costly because the "make-or-buy" framing assumes that resolving such frictions requires owning the supplier firm, whereas the existence of partial integration shows that buyers can secure targeted influence while preserving some

[7] In practice, VI may also be undesirable for other reasons, such as the costs of M&A, or the fact that outsourcing enables lower production costs through specialization or regulatory arbitrage.

of the cost advantages – from specialization, competition, and wage differences – of outsourcing. They are more common because the prevalence of contracting frictions is usually inferred from how often firms formally own upstream suppliers, as in Antràs (2003) and Boehm and Oberfield (2020), without accounting for partial integration responses.

This gap between formal ownership and effective control also connects to a puzzle in the empirical literature on vertical integration: ownership is only weakly related to physical flows of goods. Within U.S. multinationals, the input–output linkage between a parent's and an affiliate's industries does not predict intrafirm shipments (Ramondo, Rappoport, and Ruhl, 2016), and vertically integrated U.S. establishments often ship little to one another (Atalay, Hortaçsu, and Syverson, 2014). Our results suggest that this weak link need not indicate a limited role for vertical coordination, but rather may reflect instead the use of partial integration instruments that ownership data do not record.

Related literature. This paper contributes to several literatures. The first is on vertical integration. An established theoretical literature shows that integrated governance structures minimize the costs of maladaptation to contracting hazards (Williamson, 1979) and reallocate residual control rights to prevent holdup (Grossman and Hart, 1986; Antràs, 2003; Antràs and Helpman, 2004). A related empirical literature documents that a large share of trade takes place within integrated firms (Alfaro et al., 2025; Atalay, Hortaçsu, Li, et al., 2019; Garg, Ghosh, and Tan, 2023), that poor formal contract enforcement leads to integration (Boehm and Oberfield, 2020), and that acquisition leads supplier behavior to invest in meeting buyer requirements that markets do not reward, such as for quality (Hansman, Hjort, León-Ciliotta, et al., 2020). Relative to this work, we show that buyers can replicate functions of integration through targeted interventions – 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 governance arrangements that relax constraints to production by blurring firm boundaries. In Chinese processing trade, foreign buyers often own Chinese exporters' inputs (Feenstra and Hanson, 2005); in Rwanda, coffee mills provide inputs and credit to farmers (Macchiavello and Morjaria, 2021); in the U.K., agroprocessors such as Birds Eye supply seed and direct cultivation under contract farming (Grant, 2002); in Costa Rica, multinationals transfer knowledge to suppliers (Alfaro-Urena, Manelici, and Vasquez, 2022). Our contribution here is twofold: we document several forms of partial integration deployed by the same buyer (incentives, control, and monitoring of input sourcing), then formally isolate the friction (noncontractibility) that motivates their adoption, including

by providing causal evidence that buyers substitutes between arrangements as the underlying friction intensifies.

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 (Baker, Gibbons, and Murphy, 2002; 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 in upstream suppliers (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. While several case studies document such organizational responses, we both provide a formal theoretical rationalization for them and then use empirical tests to identify the contracting friction that such involvement addresses.

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 by supply inputs in-kind to credit-constrained suppliers (Manova and Yu, 2016), and adjust payment terms as contracting frictions worsen (Antràs and Foley, 2015). While this literature focuses on suppliers' financial constraints–something we formally document by showing that input investment causally responds to cash inflows–we diverge by emphasizing the noncontractibility of capital *allocation* as the friction buyers internalize. In particular, while credit constraints alone make a firm resource (cash) scarce, we argue it is the risk that suppliers divert that resource that leads buyers to intervene so directly in our setting.

The remainder of the paper is organized as follows. Section 2 presents our data and documents partial integration. Section 3 develops the model. Section 4 tests its first two predictions. Section 5 documents partial integration across Indian manufacturing and tests the third prediction. Section 6 concludes.

Figure 1: Metal fabrication in progress (left) and completed (right)





Notes: Photographs by authors.

2 Documenting partial integration

In this section, we use production and contracting data from Indian manufacturing to document how buyers intervene in a particular supplier decision (input investment) that leads to production failures.

2.1 Setting

Our empirical setting is the supply chain for fabricated metal in India. In this industry, small suppliers produce customized steel beams and rods according to customer-provided engineering drawings, as shown in Figure 1. Their customers are multinational conglomerates that design and assemble high-rise buildings, electrical transformers, power plant boilers, railway tracks, wind turbines, and other complex infrastructure. We highlight four features of this industry that are central to our analysis.

First, the production process is standardized. Every component passes through the same stages: sourcing, cutting, welding, grinding, blasting, painting, inspection, and delivery. This common structure lets us compare production outcomes across firms stage by stage.

Second, customers require just-in-time (JIT) delivery. Throughout the paper, we accordingly define production failure as the failure to deliver on time.

Third, production requires a large upfront investment of working capital (cash) in inputs. The primary input, raw steel sheet, accounts for 60-80% of total production cost and must be

paid for upfront on a spot market, with funds only recouped upon delivery. Suppliers must therefore finance the duration of their production (Kim and Shin (2012)). Importantly, this working capital is fungible: a supplier can divert it across orders, either by allocating cash to another order or by reselling or reallocating raw steel that it has already purchased.

Fourth, production decisions are not observable to third parties. Indian courts are congested, so contingencies on delivery timing are difficult to enforce (Boehm and Oberfield (2020)); and many upstream decisions, such as the allocation of supplier resources (cash, labor, and raw material) across orders, are invisible to outsiders altogether.

2.2 Data

To examine how buyers influence the decisions of their non-integrated suppliers, we require data on suppliers' realized production decisions, buyer requirements, and buyer involvement (beyond formal ownership) at production stages before final delivery. Standard production and trade datasets rarely record such information.[8] We instead turn to administrative data from India's largest manufacturing intermediary (henceforth "the buyer"), whose primary business is dividing large construction tenders from end customers into smaller contracts for many suppliers. Our data comprise 790 contracts with 224 suppliers over 2019-2022. We observe:

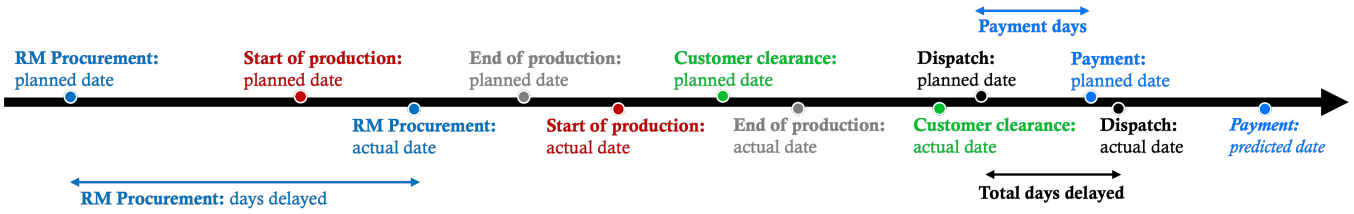
1. Production timelines. For each lot – roughly, a truckload – we observe the contract, the supplier, the end-customer project, and the planned and actual dates of each production stage (Figure 2). Planned final delivery dates are specified by end customers to align with downstream production schedules, and upstream plans are set by the buyer. Actual dates are then verified through buyer visits to supplier sites. Comparing planned to actual dates gives us delays at each stage. Summary statistics are in Appendix Table A1.

2. Contract terms. For each contract signed between the buyer and a supplier, we observe two sets of terms. The first is the *payment structure*: the share of order value paid upon completion of each production stage, and the lag between completion and payment. For instance, a supplier may be paid 20% after sourcing raw steel, with the remainder due 45 days after final delivery of the fabricated lot. The second is the *contract scope*: either "sourcing and conversion," in which the supplier sources and owns the raw steel, or "conversion only," in which the supplier is only responsible for transforming buyer-owned steel sheets.[9]

[8] A limited exception, which we exploit in Section 5, is that some datasets record buyer ownership of suppliers' inputs (for instance Feenstra and Hanson, 2005; Manova and Yu, 2016).

[9] We also observe supplier and downstream project identifiers. Supplier identifiers let us study endogenous changes in contract form for the same supplier, and project identifiers link contracts to downstream projects (a

Figure 2: 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.

3. Site visits. The buyer employs relationship managers, whose main role is to visit supplier factory floors multiple times per week. These managers provide production advice to suppliers, update production statuses to inform downstream customers, instruct suppliers to prioritize certain lots, monitor the use of provided resources, and apply pressure to start or finish orders. To quantify these interactions, we use the timing of scheduling updates that managers enter into the buyer's internal IT system.

Taken together, these data let us observe both upstream supplier decisions and the governance arrangements buyers use to shape them.

2.3 The key supplier decision: input investment

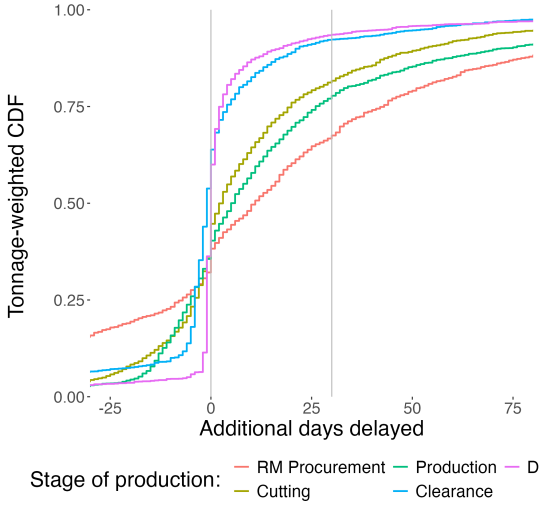
We first use these data to identify a specific source of production failure: suppliers fail to allocate a limited resource – working capital – to inputs for new orders. We build our argument in several steps.

Suppliers often source inputs late. In Figure 3 Panel A, we decompose the total days delayed for each completed lot into its constituent stages, and plot the distribution of stagewise delays. Panel A of Figure 3 shows that raw steel procurement has the largest median delay (15 days) and the largest variance of all stages. Median procurement delays exceed those in starting physical production (due to labor or capacity constraints), in completing production (which nests absenteeism, weather disruptions, or rework after quality checks), and in transportation.

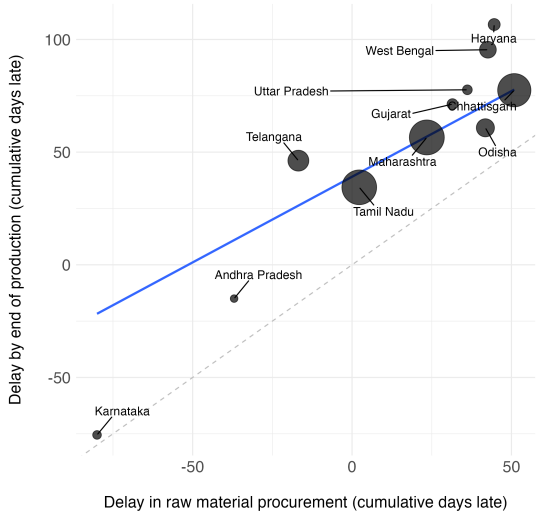
mean of 2.7 contracts per project), which we later use to construct payment shocks driven by common downstream project delays.

Figure 3: Delays are concentrated at input sourcing and persist to delivery

(a) Stage-wise days delayed



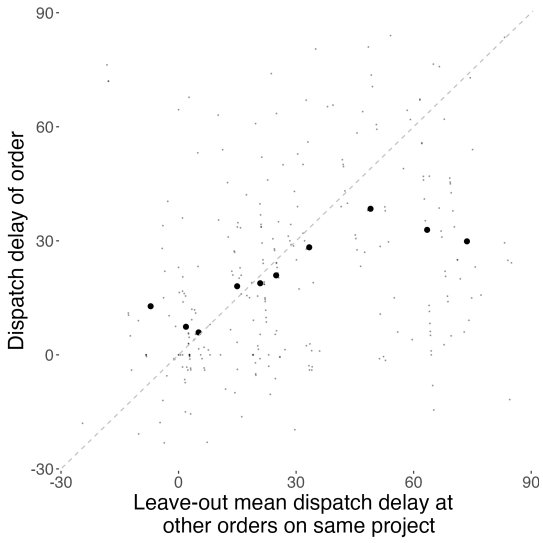
(b) Delivery vs. sourcing delays, by state



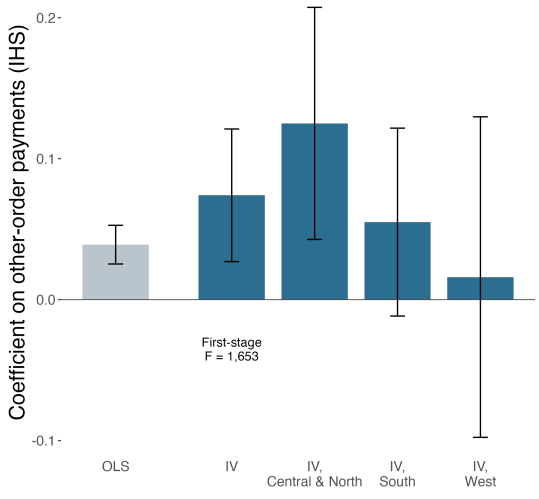
Notes: Panel A: tonnage-weighted CDF of additional delays by stage across lots, where a stage's additional delay is days behind schedule at its completion less days behind at the prior stage. Panel B: state-level mean cumulative delay at end of production vs. at raw material procurement; point size is state tonnage, the blue line is the tonnage-weighted OLS fit, and the dashed line is the 45-degree line. Data: buyer production timelines.

Figure 4: Input sourcing responds to payment inflows from other orders

(a) First stage: leave-out project delays



(b) Payment elasticity of input sourcing



Notes: Panel A shows the relationship between delays (by dispatch) on an order with a particular supplier and the average delay on orders at *other* suppliers within the same downstream construction project; the scatterplot is in gray, binned means in black, and the dashed line is the 45-degree line. We restrict to orders in contracts containing multiple orders with different suppliers. Panel B plots the coefficient on other-order payments from Equation (1): the gray bar is the OLS estimate and blue bars are IV estimates, for the full sample and by supplier region; whiskers are 95 percent confidence intervals based on standard errors clustered by order. Full estimates are in Appendix Tables A2 and A3. All data are from 2019-2022 buyer production timelines.

Late input sourcing predicts eventual production failure. Figure 3 Panel B verifies that input sourcing delays predict total delays at delivery, which is our definition of production failure, with slope slightly below one.[10] In other words, despite catching up slightly during production, locations that source inputs late still fail to deliver final output.

Input sourcing responds to exogenous payment inflows, consistent with input investment frictions. The prevalence of steel sourcing delays could imply that fabrication firms face significant input access constraints (Boehm and Oberfield 2020). However, raw steel is a widely-available commodity in India. We therefore test the alternative explanation that steel sourcing depends on the supplier's decision to allocate cash to an order.

To formally test this mechanism, we ask whether input sourcing responds to the timing of cash inflows from past orders.[11] Formally, denote order o, calendar month t, and supplier $s(o)$. We aggregate so that $Steel_{ot}^{OBS}$ is the raw material (in metric tons) actually sourced for order o in month t, and $Pay_{-o,t}^{OBS}$ is the total payment received that month on the supplier's *other* orders; $Steel_{ot}^{PLAN}$ and $Pay_{-o,t}^{PLAN}$ are the planned analogues. Writing IHS for the inverse hyperbolic sine, we estimate:

$$\text{IHS}\left(Steel_{ot}^{\text{OBS}}\right) = \alpha_o + \gamma_t + \beta^{\text{OBS}} \underbrace{\text{IHS}\left(Pay_{-o,t}^{\text{OBS}}\right)}_{\text{[observed payments]}} + \beta^{\text{PLAN}} \underbrace{\text{IHS}\left(Pay_{-o,t}^{\text{PLAN}}\right)}_{\text{[planned payments ctrl]}} + \delta \underbrace{\text{IHS}\left(Steel_{ot}^{\text{PLAN}}\right)}_{\text{[planned sourcing ctrl]}} + \epsilon_{ot} \tag{1}$$

Leave-out delays instrument. Because suppliers can influence completion timing on their own other orders, we instrument observed payments with the payments predicted when every order is delayed by the leave-out average delay on its downstream project. Formally, letting $\overline{Delay}_{-s,c}$ be the average dispatch delay at project c's *other* suppliers, and i index lots with weight MT_i and planned payment date $Date_i^{PLAN}$, predicted payments are

$$Pay_{-o,t}^{\text{INST}} = \sum_{\substack{o' \neq o \\ s(o')=s(o)}} \sum_{i \in o'} MT_i \cdot \mathbb{1}\left[\left(Date_i^{PLAN} + \overline{Delay}_{-s(o),\,c(o')}\right) \in t\right] \tag{2}$$

which serves as the excluded instrument in Equation (1). Identification therefore comes from leave-out average delays, conditional on plans: when a customer postpones a construction project, payments to every supplier on the project shift together for reasons outside any one supplier's control. Panel A of Figure 4 shows the resulting first stage: leave-out delays strongly

[10] We use state-level averages to be conservative given the possibility of mean reversion at the individual lot level.

[11] This test is analogous to classic investment-cash flow sensitivity tests in the finance literature (Fazzari, Hubbard, and Petersen, 1988).

predict own-order delays.

Panel B of Figure 4 summarizes the results, with full estimates in Appendix Tables A2 and A3. Input sourcing responds to cash inflows: the OLS elasticity of new steel sourcing with respect to other-order payments is 0.039, and the IV elasticity is 0.074, with a first-stage F-statistic above 1,600. Effects are larger for firms in Central and North India, which are poorer and may face greater financing constraints.

Given this evidence, we interpret input sourcing delays as (at least in part) failures to induce desired supplier input investment. The next subsection examines *how* buyers intervene in this decision.

2.4 Three forms of buyer responses

Having identified a key upstream decision (supplier input investment), we next show that buyers intervene in this decision without taking formal ownership of suppliers. These governance interventions take several forms: buyers pay advances conditional on supplier input sourcing, provide inputs in-kind to suppliers, and physically visit suppliers to observe procurement decisions. We refer to these interventions as *partial integration* (PI) because they mirror the behaviors but not the form of a vertically integrated firm.

Below we document these PI forms, then discuss two competing motives for their adoption: buyers may intervene either to internalize *noncontractibilities* in upstream supplier decisions, or to *specialize* in tasks (such as borrowing or input procurement) in which they enjoy cost advantages.

2.4.1 Contract-level interventions: conditional advances and input provision

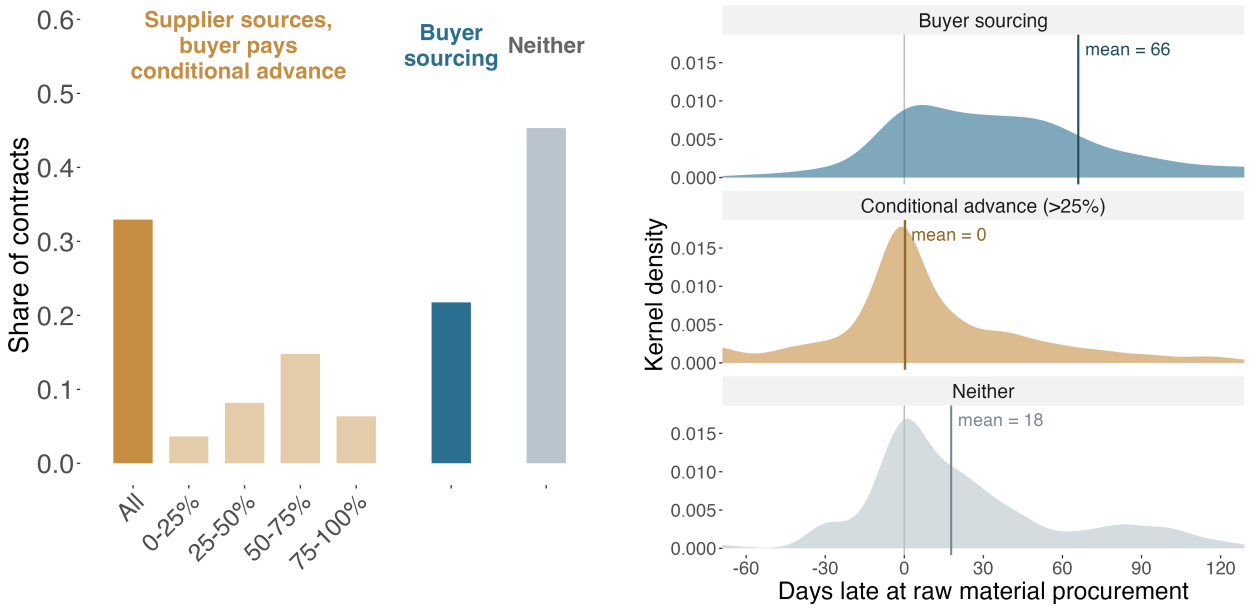
Using the text of the 331 contracts in our sample, Panel A of Figure 5 shows that two formal contractual arrangements target the input decision. The first is a *conditional advance*, in which the supplier sources and owns the raw steel, but the buyer advances a share of the order value after physically verifying that steel procurement has occurred. For instance, a contract may specify *"20% payment at RM [raw material] stage and 80% 45 days after dispatch [delivery]"*. 33% of contracts feature conditional advances, with a mean share paid of 60%. These advance payments vary across suppliers, consistent with heterogeneity in supplier financing needs.

In the second form of partial integration, *buyer sourcing*, the buyer procures steel itself from local traders and delivers it to the supplier. As a result, suppliers produce manufacturing services

Figure 5: Magnitude and form of buyer financing

(a) Contract forms and advance shares

(b) RM procurement delays by contract form



Notes: Panel A shows the share of all contracts (orders) by form. The solid orange bar is the total share with a conditional advance; the translucent orange bars split that total by the share of order value paid on verified sourcing. Blue is buyer-sourced raw material and gray is contracts with no input-related terms. Panel B shows kernel densities of days late at raw material procurement across lots, separately by the contract form of the lot's order; the conditional-advance group is restricted to advances above 25% of order value. Dark vertical lines mark group means, and the light gray line marks zero (on time). All data are from buyer contract terms and production timelines, 2020-2022 (N = 331 orders; 2,564 lots).

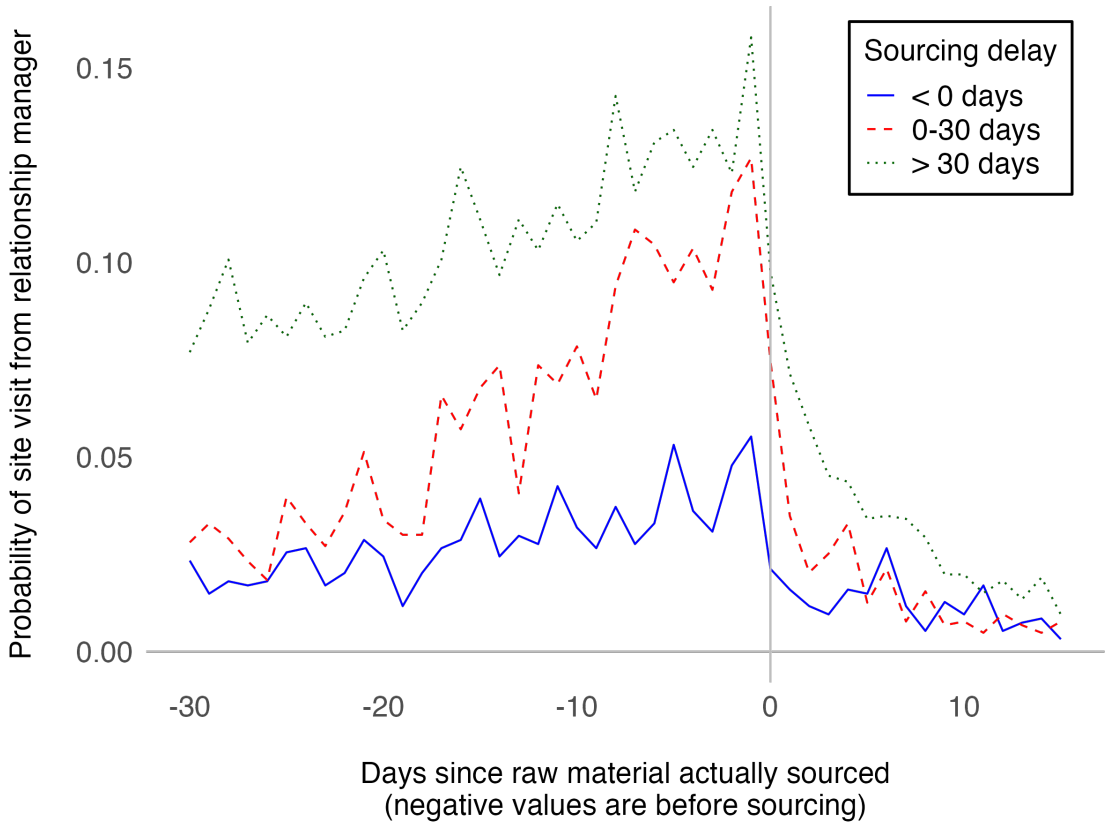
rather than fabricated steel goods. 22% of contracts involve buyer-sourced raw material.

Conditional advances and buyer sourcing represent different ways of intervening: buyers either pay for or directly implement upstream production tasks. An important similarity is that both require the buyer to know the supplier's production status with high precision – either to time input deliveries or to verify that procurement has occurred – both of which are unnecessary in arms-length contracts.

In Panel B of Figure 5, we show that conditional advances and buyer sourcing target firms with higher raw material procurement-stage delays – in other words, those with worse input sourcing behavior. This pattern supports the idea that partial integration is a response to suboptimal upstream decisions rather than a reward for good supplier performance.

Importantly, these interventions solely target the input sourcing decisions: buyers do not directly implement or pay conditional on completion of steel cutting, welding, painting, or any other stage before final inspection and delivery.

Figure 6: Probability of site visits before and after raw material is sourced



Notes: Daily probability of a raw material-related relationship-manager visit as a function of time relative to sourcing, split by sourcing lateness.

2.4.2 Within-contract interventions: site visits

Using separate administrative data on buyer-supplier interactions, we identify a third way that buyers involve themselves in supplier input investment: *site visits*. The buyer employs generalist 'relationship managers' who visit supplier shop floors, physically observe plant capacity and any in-progress orders, update expected delivery dates, provide productive advice, and pressure plant managers in conversation when an order falls behind.

Figure 6 shows that, when input sourcing falls behind the initial production plan, relationship managers visit more frequently until sourcing occurs: up to once every 8 days for minor delays of 0-30 days, and once every 6 days for major delays exceeding 30 days. These patterns of escalation and high visit frequencies together suggest that – rather than being passive check-ins with suppliers – site visits are linked to specific upstream production decisions.

2.4.3 Two interpretations: noncontractibility versus specialization

Subsections 2.4.1 and 2.4.2 provide revealed-preference evidence that partial integration can dominate both pure arms-length transactions, in which buyers do not intervene upstream; and vertical integration, after which upstream involvement is feasible but requires buyer responsibility for all production tasks. However, the underlying motive for adopting this intermediate organizational form – and thus the economic implications of PI – are ex-ante unclear. In Table 1, we lay out two potential motives that are observationally similar but have different economic implications.[12]

Table 1: Two interpretations of each buyer intervention (i.e., PI form)

Intervention	Noncontractibility (diversion risk)	Specialization (cost advantages)
Conditional advances: supplier sources raw steel, buyer pays	*Incentives* for input investment	Borrowing cost advantages
Buyer sourcing: buyer sources and owns raw steel, provides to supplier	*Control* of input investment	Input procurement cost advantages
Site visits: Relationship manager visits until steel sourcing occurs	*Monitoring* of supplier decision	Productive knowledge

The first motive is that buyers partially integrate suppliers in order to internalize the *noncontractibility* of supplier input investment – in other words, to ensure that suppliers do not divert fungible inputs (steel and working capital) to other orders. As shown in Table 1, under this hypothesis, conditional advances can be interpreted as *incentives* for input investment (which is less fungible than cash), buyer sourcing grants buyers *control rights* over specific decisions, and site visits are a form of *monitoring* of the upstream decision. Through this lens, partial integration is a 'lighter-touch' solution to contracting frictions than full vertical integration that preserves many benefits of market transactions: for instance, suppliers can compete with each other, and any supplier cost advantages in the physical transformation of inputs (such as hiring local labor) are preserved.

Under the noncontractibility motive, PI has several implications for theories of firm boundaries.

[12]Beyond these two motives, buyer-provided inputs could be rewards for suppliers with greater bargaining power. However, the fact that suppliers completing *smaller* orders receive buyer-sourced inputs is inconsistent with this motive.

First, the costs of firm boundaries are lower than classic make-or-buy models predict, because noncontractibilities can still be resolved within the broad framework of a market transaction. Second, because partial integration does not involve changes in formal firm ownership, its existence implies that contracting frictions are more prevalent than implied by typical empirical analyses of ownership data alone. A final implication of the noncontractibility interpretation is that – contrary to classic make-or-buy models – full vertical integration is not necessary to resolve contracting frictions.

The second motive, distinct from noncontractibility, is that partial integration is a form of buyer *specialization* in tasks for which buyers enjoy cost advantages vis-a-vis their suppliers. For instance, buyers may source inputs for suppliers simply because they have access to lower-cost raw material traders, not because of any desire to control the timing of input sourcing. Similarly, buyers may advance cash to suppliers simply because of better access to short-term credit. Through this lens, buyer interventions are efficient: they allow firms to reduce costs through specialization, rather than being a costly step that buyers take to make timely production feasible.

As shown in Table 1, both forces can generate – in the cross-section – the set of partial integration interventions that we document. Interpreted as responses to noncontractibility, conditional advances are *incentives* to invest cash in raw steel, which is less-easily diverted than cash; and buyer sourcing directly avoids diversion by granting buyers *control rights* over the decision. Under specialization, advance payments and buyer sourcing are instead ways for buyers to provide suppliers with inputs (working capital and materials, respectively) at lower cost than suppliers face. Similarly, under noncontractibility, supplier site visits are a form of *monitoring* of the internal supplier decision, whereas under specialization they are a way of delivering productive knowledge (through conversation).

Given the observational equivalence of these motives in the cross-section, our goal in the remainder of the paper is thus to formally test if the partial integration forms we see in Figures 5 and 6 are responses to noncontractibility. To do so, in Section 3 we develop a model that nests both hypotheses, and derive three predictions for shock-induced *changes* in contract form that hold if and only if the upstream decision is subject to diversion risk. We then provide causal tests of these predictions using three separate natural experiments in Sections 4 and 5.

3 A framework of partial integration

We present a model of the relationship between a buyer and a financially constrained supplier. The model nests the two motives for buyer intervention – noncontractibility of the supplier's input decision, and buyer specialization in finance, procurement, or knowledge – and delivers three predictions. In the presence of noncontractibility, the form of the intervention switches from incentives to control as the financing gap widens (Prediction 1); within a contract, monitoring responds to the supplier's cash-in-hand (Prediction 2); across credit regimes, control recedes as outside finance closes the gap (Prediction 3). A final result relates to the boundary with full integration. Because the null under each prediction is pure specialization, the three jointly separate the two hypotheses. Sections 4 and 5 take them to the data.

3.1 Setup

A supplier must procure an input and then produce output for a buyer. There is a working-capital financing gap at the input stage. The buyer can intervene in two ways. Ex ante, at the time of contracting, it can advance funds or take direct control of input sourcing. Ex post, after the contract is signed, it can monitor the supplier's input decision.

Technology and decisions. One unit of output requires raw material worth m. At the time of contracting the supplier has normal liquidity $\bar{\ell}$, so the order has an ex-ante financing gap $\tau \equiv \max\{m - \bar{\ell}, 0\}$, strictly positive for constrained suppliers. After the contract is signed but before sourcing is complete, realized liquidity is $\ell = \bar{\ell} + \varepsilon$, where ε is a mean-zero shock capturing the timing of cash inflows to the supplier. The gap τ measures how constrained the supplier is when the buyer chooses the contract, while the realization ℓ measures how much cash it has available at given point in time to allocate for orders.

Production has two stages, each with a noncontractible decision. At d_1, the input-sourcing decision is whether fungible resources are used to procure raw material for this order or diverted to another use. At d_2, the delivery decision is whether the supplier delivers the current order on time. Timely completion requires high supplier effort at private cost $c > 0$ and creates gross value $\bar{V} > c$ for the buyer.

Ex-ante governance structures. The buyer chooses among five governance structures:

1. *No input intervention (0):* no input-related contract term; the supplier retains d_1 and d_2.

2. *Unconditional advance (U):* the buyer advances τ with no conditions attached. The advance is recovered from the delivery payment, and lost if funds are diverted.

3. *Conditional advance (A):* the buyer advances τ after verifying at cost $K > 0$ that procurement has occurred. The material remains fungible, even though less than cash, and the supplier retains d_1. The advance is recovered from the delivery payment, and lost if material is diverted.

4. *Buyer-sourced input (B):* the buyer takes d_1 at net cost F – the incremental organizational and logistics cost of sourcing on the supplier's behalf, net of any procurement advantage in buying raw material.[13]

5. *Full vertical integration (VI):* the buyer takes both decisions.

Input-stage risk. When the supplier retains d_1 (under 0, U, and A), the input decision is noncontractible: with probability $\rho_1(\tau)$, the supplier's fungible resource is allocated away from this order, whether diverted to competing uses or absorbed by other obligations. The function ρ_1 is continuous and increasing, with $\rho_1(0) = 0$. Conditioning the advance reduces the risk to $(1-\mu)\rho_1(\tau)$, with $\mu \in (0,1)$: verification ensures that the supplier has converted cash into steel, but steel is still a fungible commodity, even if less fungible than cash. Once sourcing is complete and the material is committed to production, the input decision is closed under every supplier-retained regime. Under B and VI the buyer controls d_1 and the risk is zero.

Intervention motives. The model nests two motives for intervention. The *noncontractibility motive* is present when $\rho_1 > 0$ – a wider gap makes competing demands on liquidity more pressing – and absent when $\rho_1 \equiv 0$. The *specialization motive* operates through buyer cost advantages in credit financing, input sourcing, and production knowledge. To leverage these advantages, the buyer could fund working capital at a saving of $b \geq 0$ per rupee whenever it covers the gap; its net sourcing cost F may be negative,[14] when a procurement advantage exceeds the organizational cost; and its visits transfer production knowledge worth $\gamma \geq 0$ per unit of visit intensity, independent of the supplier's cash position.[15]

[13] The material cost m itself is common to all regimes and netted out throughout.

[14] The sourcing advantage is representing as a fixed cost advantage, rather than a scale advantage, because the order is assumed to require one unit of raw material. For the predictions of this model, which are tested within supplier with no change in order sizes, this is with no loss of generality.

[15] Steel in our setting is a homogeneous commodity sold on a spot market and produced in the same clusters that the suppliers locate in, leaving little knowledge about the act of procurement itself to transfer.

Ex-post monitoring. At any moment within a contract, the allocation decision is actionable – the supplier holds fungible funds it could invest in this order or allocate elsewhere – with probability $s(\ell)$, strictly increasing in liquidity: payment arrival puts divertible resources in the supplier's hands. After observing ℓ, the buyer can visit at intensity $h \in [0,1)$, deterring the fraction h of the diversion risk and transferring knowledge worth γh, at strictly convex cost $k(h)$, with $k(0) = k'(0) = 0$ and $k'(h) \to \infty$ as $h \to 1$. The temptation $\rho_1(\tau)$ and the deterrence technology are fixed at contracting, so liquidity moves the value of monitoring only through s.

Delivery at d_2. With probability $q \in [0,1]$, completing on time brings the supplier a further order with private rent $R \in (0,c)$, so when the supplier retains d_2 (under 0, U, A, or B) a delivery bonus of $w^*(q) = c - qR$ motivates effort. However, delivery too is noncontractible, so with probability $\rho_2 \in [0,1)$ the supplier diverts the finished output to another buyer instead of delivering on time. This risk falls with the specificity of the output: a customized product has few alternative takers, so ρ_2 is low, while a standardized one can be resold easily. The buyer's continuation value once inputs are secured is therefore $V \equiv (1-\rho_2)(\bar{V} - c + qR)$ under non-VI interventions. Under VI the buyer implements delivery directly, so there is no delivery risk, but the integrated supplier has no outside orders and thus the continuation value is $\bar{V} - c$.

Payoffs. If the input stage fails, the buyer loses V, and when the failure occurs through diversion of an advance, it also loses the advance τ.[16] Expected payoffs are therefore[17]

$$\begin{aligned}
\Pi_0(\tau) &= V - \rho_1(\tau)\, V, \\
\Pi_U(\tau) &= V + b\tau - \rho_1(\tau)\big[V + \tau\big], \\
\Pi_A(\tau) &= V + b\tau - (1-\mu)\,\rho_1(\tau)\big[V + \tau\big] - K, \\
\Pi_B(\tau) &= V + b\tau - F, \\
\Pi_{VI}(\tau) &= \bar{V} - c + b\tau - F.
\end{aligned}$$

Figure 7 summarizes the model's timing.

[16] Diverted funds are sunk in other uses and effectively irrecoverable. Idle funds, and funds embodied in verified procurement, can be netted against future payments. The results require only that diverted funds are harder to recover than idle ones.

[17] These payoffs hold visit probability at zero. Knowledge-motivated visiting is available under every governance structure, so its surplus $\max_h\{\gamma h - k(h)\}$ is a common constant that cancels from every margin below; we omit it.

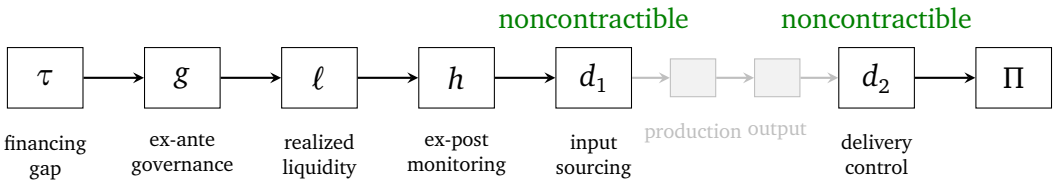


Figure 7: Sequential production with stage-specific control. The financing gap τ is known at contracting, the buyer chooses a governance structure $g \in \{0, U, A, B, VI\}$, liquidity ℓ realizes, the buyer monitors at intensity h, the supplier takes the input decision d_1, production and output follow (grey), the delivery decision d_2 determines completion, and payoffs Π realize.

3.2 Implications for partial integration

We first characterize governance as a function of the financing gap when noncontractibility is present, and then derive three predictions. Proofs are in Appendix C.

3.2.1 The governance choice

The key assumption that generates our predictions is that the contracting friction is present. Motivated by observed contracts, we also assume that conditional advance is optimal at some gap, and buyer sourcing beats no intervention at some gap.

Assumption 1. *(i) ρ_1 is strictly increasing; (ii) there exists a gap $\tau_0 > 0$ with*

$$\rho_1(\tau_0)\big[\mu V - (1-\mu)\tau_0\big] > K \qquad \text{and} \qquad (1-\mu)\,\rho_1(\tau_0)\big[V + \tau_0\big] < F - K;$$

(iii) there exists a gap τ_1 with $\rho_1(\tau_1)\,V > F$.

A larger financing gap raises the input-stage loss when the supplier controls d_1, while buyer control removes that loss at a fixed cost. Control is therefore reserved for large gaps.

Lemma 1 (The governance ladder)**.** *Suppose $qR \geq \rho_2(\bar{V} - c + qR)$ and Assumption 1 holds. Then: (i) in a neighborhood of $\tau = 0$, neither the conditional advance nor buyer sourcing is optimal; and (ii) there is a unique control threshold $\tau_C \in (\tau_0, \infty)$ such that buyer sourcing is optimal for $\tau > \tau_C$ and not optimal for $\tau < \tau_C$.*

Partial integration concentrates among constrained suppliers, and its form escalates from no intervention or incentives to control when the financing gap crosses the threshold τ_C.[18] Below

[18] The escalation mirrors Burkart and Ellingsen (2004), where input finance dominates cash lending because

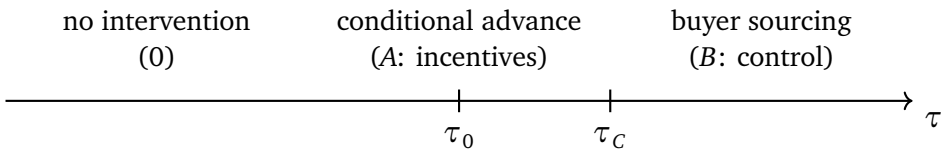


Figure 8: The governance ladder: input-stage governance as a function of the financing gap $\tau = m - \bar{\ell}$. No intervention is optimal at small gaps (Lemma 1(i)), the conditional advance is optimal at τ_0 (Assumption 1), and buyer sourcing is optimal above the control threshold $\tau_C > \tau_0$ (Lemma 1(ii)). Prediction 1 moves suppliers rightward (input costs raise m); Prediction 3 moves them leftward (credit access raises $\bar{\ell}$); Prediction 2 operates within a contract, at a fixed position on the ladder.

the control threshold, advances are used whenever the reduction in diversion risk justifies the verification cost.[19]

Remark 1 (Conditional versus unconditional advances)**.** *For any $b \geq 0$, conditioning an advance is worthwhile only against diversion: $\Pi_A - \Pi_U = \mu\rho_1(\tau)[V + \tau] - K$. With $\rho_1 \equiv 0$ this fails at every gap, so any advance the buyer wrote would be unconditional. The observation of conditional advances therefore indicates the presence of a contracting friction.*

3.2.2 Prediction 1: Substitution from incentives to control

The residual loss under a conditional advance rises with the financing gap. As the gap widens, the buyer switches to buyer sourcing, which removes the residual loss at a fixed additional cost.

Proposition 1 (Switch from incentives to control)**.** *Suppose $qR \geq \rho_2(\bar{V} - c + qR)$. Under Assumption 1, as input cost m rises, suppliers on conditional advances whose gaps cross τ_C are switched from incentives to control. Absent the friction ($\rho_1 \equiv 0$), conditional advances are never used and whether the buyer sources is independent of the gap, so no substitution occurs.*

The fact that incentives and control are substitutes is a key prediction of the model, with the substitution from the former to the latter occurring when the decision risk rises. We take this prediction to an unanticipated steel-price shock in Section 4.

inputs are harder to divert than money. Here a downstream buyer finances an upstream supplier, and conditioning the advance (μ) and sourcing the input (B) are two ways of recovering the divertibility wedge that unconditional cash lacks.

[19] For some set of primitives, an intermediate range of τ can exist in which diversion is too severe to risk a conditional advance but not severe enough to justify sourcing.

3.2.3 Prediction 2: Monitoring tracks realized liquidity

Monitoring is chosen after the contract is signed, when realized liquidity ℓ is known. A cash inflow increases ℓ, making the allocation decision more likely to be actionable and therefore more worth monitoring.

Proposition 2 (Monitoring responds to realized decision risk)**.** *Suppose the supplier retains d_1 (under 0, U, or A). When $\rho_1(\tau) > 0$, optimal monitoring h^* is strictly increasing in realized liquidity ℓ before sourcing is complete, and falls to the knowledge level $k'^{-1}(\gamma)$ once raw material is procured. Absent the friction, $h^* = k'^{-1}(\gamma)$ at every ℓ: visits occur whenever $\gamma > 0$ but are invariant to liquidity.*

Because other-order payments raise ℓ, buyer visits should rise when those payments arrive, and should fall to the knowledge level once procurement is complete. Under pure specialization in knowledge ($\rho_1 \equiv 0$, $\gamma > 0$), visits could occur but do not concentrate at the moments when fungible cash is in the supplier's hands. We take this timing prediction to relationship-manager visits in Section 4.

3.2.4 Prediction 3: Credit access leads to withdrawal of control

Consider a population of buyer-supplier pairs with heterogeneous gaps and cost positions, with at least some pairs satisfying Assumption 1.

Proposition 3 (Credit access reduces buyer control)**.** *Suppose $qR \geq \rho_2(\bar{V} - c + qR)$. For pairs satisfying Assumption 1, an improvement in outside credit access weakly lowers each gap: those whose gaps fall below the control threshold τ_C abandon buyer sourcing. For pairs with $F < 0$, buyer sourcing is optimal at every gap, for any diversion risk and any $b \geq 0$, and credit access leaves it unchanged.*

Thus, in the aggregate data, as long as there are at least some buyers choosing to source inputs due to contracting frictions, buyer sourcing should fall when there is an exogenous expansion of credit access for suppliers. When all buyers intervene only due to specialization in input sourcing, the prevalence of the intervention is unaffected by expanded supplier credit access. We take this prediction to the 2006 MSMED credit reform in Section 5.

3.2.5 The boundary with full integration

Our final result compares control-based partial integration with full integration. Both B and VI solve the upstream problem at net cost F, so the comparison relates to the delivery stage. Under B, the supplier has the temptation to divert the finished output to another buyer, but it also faces market incentives to complete the order and take on further order conditional on delivery. Integration removes the diversion risk but also kills market incentives.[20]

Proposition 4 (Buyer sourcing versus full integration)**.** *For any $b \geq 0$ and any financing gap τ,*

$$\Pi_B(\tau) - \Pi_{VI}(\tau) = qR - \rho_2\big(\bar{V} - c + qR\big).$$

Buyer sourcing becomes more attractive as q rises and, when $q > 0$, as R rises, and less attractive as ρ_2 rises. If $0 < \rho_2(\bar{V} - c) < (1 - \rho_2)R$, there exists a threshold $q^* = \frac{\rho_2(\bar{V}-c)}{(1-\rho_2)R} \in (0,1)$ such that the buyer prefers full integration for $q < q^*$ and buyer sourcing for $q > q^*$.

In the steel fabrication setting, the output is fully customized, so the scope for diverting it is small, and suppliers maintain many ongoing market relationships, so market incentives are strong. This rationalizes the fact that we observe buyer sourcing rather than widespread integration in our setting.

3.3 Summary of the predictions

The model delivers three predictions for PI contract form that hold iff partial integration is due to the noncontractibility of upstream decisions:

1. *Following an positive input-cost shock, buyers substitute from advance payments to input sourcing for their most-constrained suppliers.*

   Corollary: Advance payments are conditional.

2. *Within a contract, monitoring rises when the supplier holds divertible resources.*

3. *Increased supplier credit access reduces buyer ownership of suppliers' inputs.*

In contrast, the specialization motive for partial integration – in which buyers have cost advantages in upstream tasks – cannot generate these patterns. In the following sections we use three distinct natural experiments to test these predictions.

[20] Conditional advances can also dominate VI. Here, the comparison is made to B for simplicity.

4 Empirical tests of the framework

The model's first two predictions concern our specific empirical context: the form of partial integration should respond to the financing gap, and monitoring should respond to realized liquidity. Test 1 takes the substitution prediction to an unanticipated steel-price shock. Test 2 takes the monitoring prediction to the timing of cash inflows, reusing the payment-timing design of Section 2.3. We close with supplementary evidence on the form of advances. Prediction 3 requires economy-wide variation in credit access, and is tested in Section 5.

4.1 Test 1: Substitution from incentives to control

Prediction 1 is that as the financing gap widens, conditional advances should give way to direct buyer sourcing. Buyer sourcing is the costlier instrument because the buyer finances the entire input purchase and handles the logistics of procurement. Its advantage is more control: the buyer owns the steel, and can hence prevents the input from being redirected. The prediction is therefore that the costlier control-based instrument is reserved for the most constrained suppliers.

Cross-section. Panel B of Figure 5 is consistent with this logic: buyer sourcing is concentrated among the lots with the worst sourcing delays, while conditional advances concentrate among lots whose suppliers source close to on time.

Input-price shock. For causal variation we exploit a rapid increase in global steel prices in late 2020, which significantly raised the working capital required to start an order. The shock raised τ for every supplier; the model predicts that the most constrained suppliers will be moved from conditional advances to buyer sourcing. Figure 9 illustrates the effect of the shock - buyer sourcing was introduced as an instrument while the use of conditional advances went down.

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:

$$DirectSource_{it} = \delta_t + \alpha_{s(i)} + \beta^{DID}\left(HasAdvance_{s(i)} \cdot Post_t\right) + \epsilon_{it} \tag{3}$$

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

Figure 9: Time series of steel price and contract terms

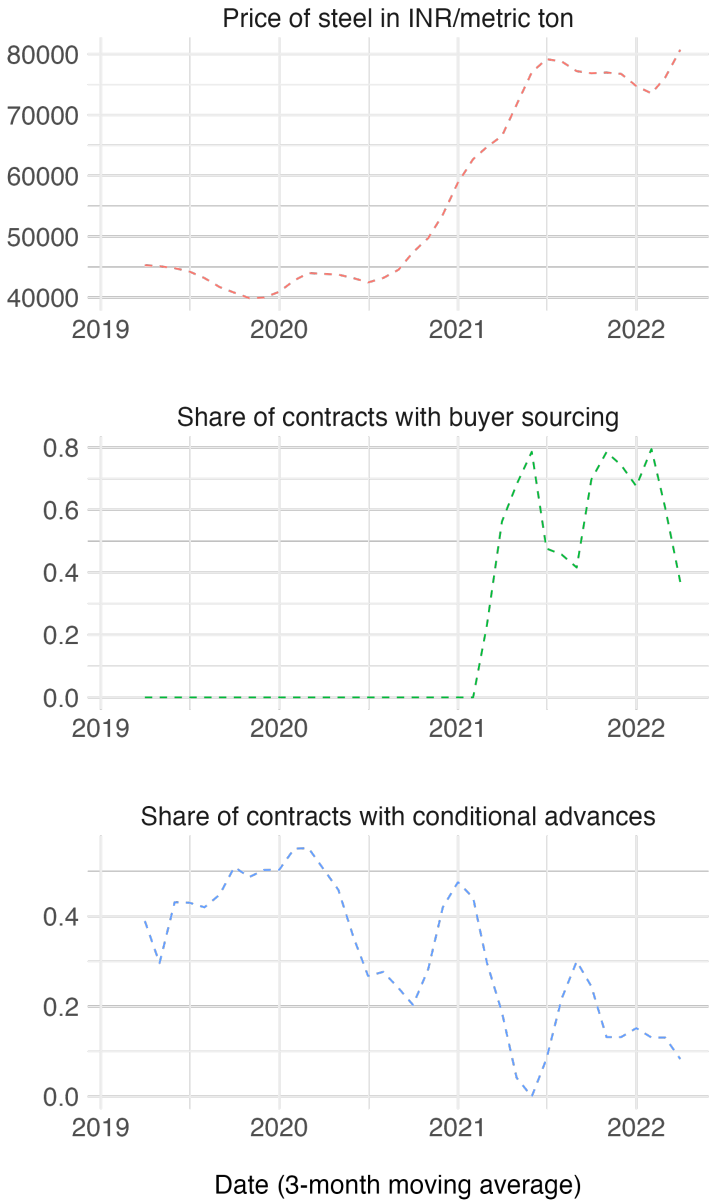


Figure shows three time series from 2019-2022, all as three-month moving averages calculated from monthly data. Top panel shows Indian spot-market steel prices, obtained from the Ministry of Steel, Government of India. Middle panel shows the share of contracts (by date of contract signing) in which the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

shock; and $Post_t$ indicates quarters after October 2020. We weight by lot tonnage and cluster standard errors by supplier. The coefficient β^{DID} captures the additional post-shock increase in buyer sourcing for suppliers that were more constrained before the shock, relative to suppliers with no pre-shock conditional advances. The identifying assumption is parallel trends: absent the steel price shock, buyer sourcing for suppliers with and without pre-shock advances would have evolved similarly. This assumption is supported by the fact that buyer sourcing was essentially absent before the shock for both groups – 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 close to zero for both groups, and the counterfactual trend is parallel. Additional specification details and alternative proxy measures are developed in Appendix B.

To understand *which suppliers* switch to buyer sourcing, we estimate an event-study version of equation (3), 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.

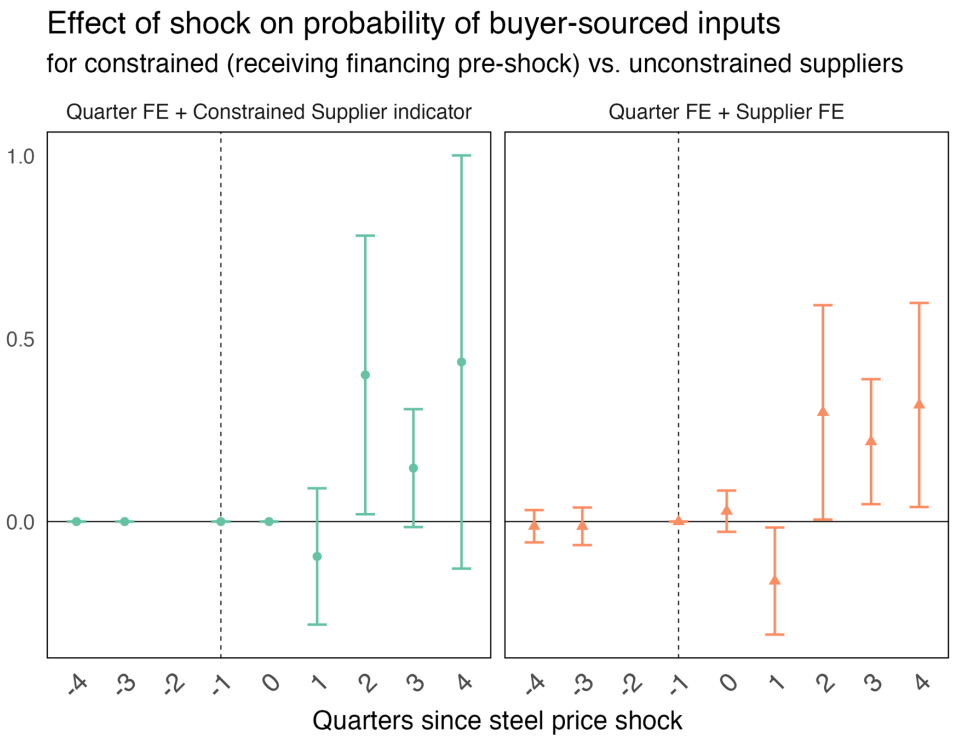
Figure 10 shows that the post-shock increase in direct buyer sourcing is concentrated among suppliers with pre-shock conditional advances. The result holds both in a pooled DID design comparing suppliers with and without pre-shock advances (left panel) and with supplier fixed effects (right panel). The dynamic path of effects is consistent with the time series in Figure 9, in which buyer sourcing was introduced suddenly. The zero pre-period effects reflect the fact that buyer sourcing was essentially not offered before the shock. Effects materialize from the second post-shock quarter because most production in quarters 0-1 is on contracts signed before the shock, leaving no scope to adjust contract terms.

To aggregate and examine robustness, we replace the event time dummies with a single indicator for being $\geq$2 quarters post-treatment (where effects materialize in Figure 10). These estimates are given in Table 2.

Column 1 shows that suppliers receiving conditional advances before the shock are 31 p.p. (405% of the sample mean) more likely to receive buyer-sourced raw material after the shock. In column 2, when we use the continuous (mean) percent advance a supplier received before the shock, effects remain positive but are no longer significant.

One alternative interpretation of these results, unrelated to contracting frictions, is that firms receiving conditional advances are smaller, and thus prefer to receive buyer-sourced inputs to avoid the fixed costs of raw material sourcing. To rule this out, in Columns 3 and 4 we include controls for a firm size proxy (mean order size in metric tons) interacted with the post-period

Figure 10: 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 on the lot-level panel at the raw material procurement stage, dated by the quarter of actual sourcing and weighted by lot tonnage; standard errors are clustered by supplier. In the left panel, the regression includes quarter FE and an indicator for being exposed (so there is a single treated and untreated group). In the right panel, the regression includes supplier and quarter FE.

dummy. This coefficient is negative, so small firms do take up buyer sourcing at higher rates. However, the effects of receiving conditional advances pre-shock (either the dummy in Column 3, or of the continuous share in Column 4), are both significant and larger in magnitude than in Columns 1-2. This means that firms receiving advance payments ex-ante are disproportionately switched to buyer sourcing *despite* them not being particularly small.[21]

These results show that buyer involvement deepens as credit needs rise. Initially, the buyer incentivizes raw-material sourcing through a conditional advance or provides no input-related support. As credit needs rise further, conditional advances are no longer sufficient for the most constrained suppliers, so the buyer takes the costlier step of directly sourcing inputs. In the

[21]Appendix Table A5 shows that results remain similar when, in lieu of the two-way FE design with supplier FE, we use a traditional parametric DID design that does not restrict to within-supplier comparisons.

Table 2: Substitution from incentives to control after the steel-price shock

Dependent Variable:	=1 if buyer sources raw material			
Model:	(1)	(2)	(3)	(4)
Variables				
Any advance × Post (≥2q)	0.314**		0.430***	
	(0.116)		(0.134)	
Percent advance × Post (≥2q)		0.126		0.598*
		(0.234)		(0.299)
Log mean order size × Post (≥2q)			-0.225**	-0.290**
			(0.100)	(0.123)
Fixed-effects				
Quarter	Yes	Yes	Yes	Yes
Supplier	Yes	Yes	Yes	Yes
Fit statistics				
Observations	968	968	968	968
Dependent variable mean	0.077	0.077	0.077	0.077

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

Notes: Table shows estimates of Equation (3). Lot-level panel at the raw-material procurement stage, weighted by lot tonnage to match Figure 10; standard errors clustered by supplier (31 clusters). Sample: suppliers active before and after October 2020; April 2020 dropped; event time winsorized to ±4 quarters. All columns include quarter and supplier fixed effects. Post (≥2q) indicates two or more quarters after the October 2020 steel-price shock. Percent advance is the supplier's pre-shock average advance share of order value, scaled zero to one, so its coefficient is the effect of moving from no advance to a full advance. All data are from buyer contract terms and production timelines, 2019-2022.

data, the shift is observed as the introduction of buyer sourcing among suppliers that were previously governed through conditional advances.

4.2 Test 2: Monitoring responds to realized cash inflows

Monitoring is valuable when the supplier has liquid resources to allocate. Other-order payments raise those resources, and so should raise monitoring.

To test this prediction, we re-estimate the payment-timing design of Section 2.3, replacing input sourcing with monitoring as the outcome:

$$\text{IHS}\big(Visits_{ot}\big) = \alpha_o + \gamma_t + \beta^{\text{OBS}}\,\text{IHS}\big(Pay^{\text{OBS}}_{-o,t}\big) + \beta^{\text{PLAN}}\,\text{IHS}\big(Pay^{\text{PLAN}}_{-o,t}\big) + \delta\,\text{IHS}\big(Steel^{\text{PLAN}}_{ot}\big) + \epsilon_{ot} \quad (4)$$

where $Visits_{ot}$ counts relationship-manager visits linked to order o in month t, measured either as the number of manager updates or as the number of unique managers visiting the plant. As in Section 2.3, we instrument observed payments with the leave-out predicted payments of Equation (2), so identification again comes from customer-driven project postponements rather than the supplier's own completion timing.

Table 3 shows that monitoring rises when cash arrives. Instrumented other-order payments raise both margins of monitoring: manager updates rise by 6.8 (s.e. 2.7) and the number of unique visiting managers by 0.05 (s.e. 0.01) per log-point increase in payments – roughly half and a third of the respective sample means of 13.9 updates and 0.14 managers. This result is consistent with site visits as a form of supervising supplier input investment: payments on other orders are the moment the supplier is newly able to invest in – or divert – cash for the focal order, and thus when relationship managers have strong incentives to visit. Other motives for site visits – to provide general technical knowledge, or gather information about production statuses – are entirely unrelated to suppliers' cash positions.

Combined with Figure 6, which shows that visits are concentrated before raw material is sourced, this result suggests the buyer visits external supplier sites to influence upstream supplier decisions – in the same way managers visit production sites to influence decisions within vertically-integrated firms.

4.3 Supplementary evidence: the absence of unconditional advances

The model also implies that if the input decision is noncontractible, conditional advances should exist and unconditional cash should not (Remark 1). The contract data confirm both sides: every advance in our sample pays out only on verified raw-material procurement, and no supplier receives an unconditional advance (Figure 5).

This pattern is hard to reconcile with finance alone. An unconditional advance would relax the supplier's working-capital constraint at lower verification cost, so a buyer concerned only with finance should prefer it. Conditional advances are valuable only because cash is fungible: once transferred, it can be redirected to another order or another use, and conditioning payment on verified procurement is what ties the transfer to the decision the buyer cares about.

The intervention is also specific. Payments are conditioned on raw-material procurement, not on downstream production stages, consistent with the contracting friction being concentrated at the input decision rather than spread across the production process.

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. Payment regressors are in inverse hyperbolic sine format; outcomes are counts in levels, so coefficients give the change in counts per log-point increase in payments.

5 Partial integration in Indian manufacturing

Our evidence so far comes from a single supply chain, in which we observe both the supplier decision at risk (input investment) and the instruments buyers use to shape it. We now ask whether the same organizational form appears *across* Indian manufacturing, using plant-level data from the *Annual Survey of Industries* (ASI). Of the three interventions we documented earlier – *incentives*, *monitoring*, and *control* – the ASI records the control intervention, in which the buyer owns the raw material and the supplier is paid only to transform it. We use this feature to measure control-based partial integration (PI) across Indian manufacturing, and

then we test Prediction 3: that PI recedes when suppliers gain access to outside credit.

5.1 Measuring buyer control

The ASI is an annual survey of registered Indian manufacturing plants, in which large plants are enumerated every year and smaller plants are sampled within state-industry strata (the census and sample schemes, respectively). Sampling multipliers are reported with each plant, which we use throughout, so that the shares we report describe registered manufacturing as a whole rather than the surveyed sample. We use two windows. For measurement, we use the 2015-16 through 2019-20 rounds, the only years in which the survey records receipts from manufacturing services separately from other services. For the credit experiment in Section 5.3, we use the 1998-2015 plant panel – a separate release from the public unit-level files – which carries a factory identifier and therefore supports within-plant designs.[22]

The ASI records buyer control because it records who owns the input. Consider two plants that perform the same physical operation: cutting and welding steel sheet into a fabricated assembly worth Rs 100, of which Rs 80 is raw steel, and Rs 20 the conversion work. The first plant buys the steel, transforms it, and sells the assembly, so it reports Rs 100 of output and Rs 80 of materials. The second plant receives steel owned by its customer, is paid only for the transformation, and therefore reports no product output and no materials at all, recording the Rs 20 instead as *receipts from manufacturing services*. The two plants differ not in what they physically do, but in *who owns the input* while production occurs. That is exactly the margin that separates buyer sourcing (*control*) from arms-length production in Section 2.4.1. We accordingly define the *conversion share*

$$CS_{ft} = \frac{\text{receipts from manufacturing services}_{ft}}{\text{value added}_{ft}} \tag{5}$$

for plant f in year t. We interpret $CS_{ft} > 0$ as evidence that some buyer holds control of the inputs to at least part of plant f's production.[23]

This measure has two limitations. First, it captures *control* alone, so the ASI covers one of the three instruments in Table 1, rather than partial integration as a whole. This gap highlights the value of the intermediary data in studying partial integration more comprehensively

[22] The public unit-level files carry only a within-year record serial, which cannot be linked across years.

[23] We build value added from the raw survey blocks (ex-factory output plus other receipts, less materials and other expenses), because the harmonized value-added variable double-counts the output block's grand total. We winsorize CS at one and drop plant-years with non-positive value added.

(however, the ASI is unusual in recording input ownership at all, which most production and trade datasets do not; see Feenstra and Hanson 2005 and Manova and Yu 2016). Second, the measure is recorded on the supplier's books – so the ASI tells us how often plants are partially integrated by *someone*, and not which buyer does the integrating.

Plants reporting conversion income are working with somebody else's steel. Plants that earn most of their value added from conversion work spend far less on materials than plants that earn none. The median materials share of gross production value falls from 55% among plants with $CS = 0$, to 39% among plants whose CS exceeds one half.

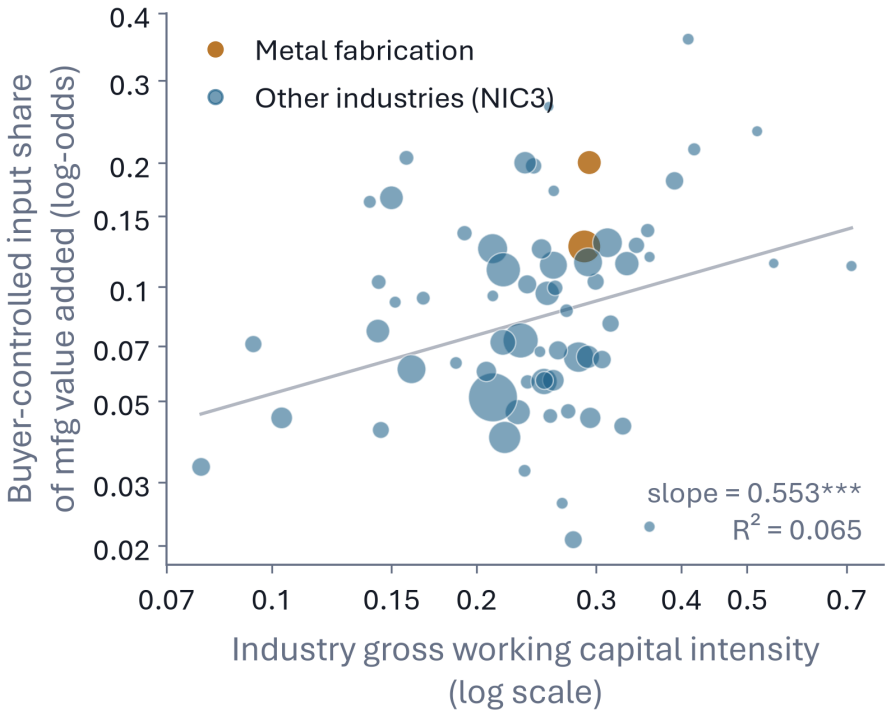
5.2 Buyer control is common, and concentrated where working capital is scarce

Roughly a third of manufacturing plants work on material they do not own. Pooling the 2015-2019 rounds, 32% of plant-years report positive conversion income, and among those plants, conversion work accounts for 28% of value added on average.[24] Control over a production stage is therefore not a rare arrangement, confined to particular buyers or to a handful of trades. It appears in a large minority of registered plants, and, where it appears, it accounts for a significant part of what the plant earns.

Buyer control is more common in industries where production ties up working capital. Our model predicts that partial integration will be more prevalent at large financing gaps, so the arrangement should concentrate in industries whose production requires suppliers to carry inventories and receivables for long stretches. In Figure 11, we plot each three-digit industry's mean conversion share against its gross working capital intensity, the latter measured over 2011-2014 (so that it is fixed before the outcome window). The relationship is positive: a doubling of working capital intensity raises the conversion share from 9% to about 13% at the sample mean, a slope of 0.553 in log-odds ($t = 2.6$). Metal fabrication – the industry of our intermediary's suppliers – has a working capital intensity of near 0.3, and its two constituent industries (NIC3 251 and 259) have conversion shares of 20% and 13%: both well above the manufacturing average.

[24] Weighting by the sampling multiplier; unweighted, the two figures are 34% and 25%. Our measurement sample keeps single-product plants, for comparability with the industry classification used below, and retaining multi-product plants raises the incidence to 34%. All figures cover core manufacturing (NIC2 10-32), which excludes repair and installation.

Figure 11: Buyer control across industries, by working capital intensity



Notes: Each point is a three-digit (NIC3) manufacturing industry, sized by its number of plant-years. The vertical axis is the multiplier-weighted mean conversion share over 2015-2019, on a log-odds scale with share labels; the horizontal axis is the industry's gross working capital intensity over 2011-2014, on a log scale, measured as the plant-weighted median ratio of finished goods, work in progress, stores and receivables to sales. Metal fabrication (NIC3 251 and 259) is in orange, and the line is the plant-count-weighted fit. Slope = 0.553 (heteroskedasticity-robust s.e. 0.212), $N = 66$ industries and 99,991 plant-years.

5.3 Credit access and the withdrawal of buyer control

The cross-industry pattern fits our framework, but it does not separate that framework from *specialization*: buyers may simply hold a procurement or borrowing advantage that matters most where inputs are expensive to carry. Prediction 3 separates the two hypotheses. If buyers take control of inputs because suppliers may divert them, then relaxing the supplier's financing constraint closes the gap that made control worthwhile, and control should recede. If instead buyers source inputs because they are cheaper at it, the arrangement survives at any level of supplier credit, and an expansion of supplier credit leaves it unchanged.

We test this using the 2006 Micro, Small and Medium Enterprises Development (MSMED) Act. The Act raised the ceiling on plant and machinery investment below which a firm qualifies as small. Lenders receive priority-sector credit against loans to firms under that ceiling, so the

reform extended cheaper formal finance to a band of mid-sized plants. The ceiling rose from Rs 1 crore to Rs 5 crore, so plants whose pre-reform investment fell in that interval became *newly eligible* for directed credit; plants above and below that band did not.[25] The reform therefore moves *supplier* credit access without moving buyer costs, which is the variation Prediction 3 requires.

Our outcomes separate the two ways a plant can earn value added. *Goods* value added is real gross output less materials, and *services* value added is receipts from manufacturing services. A plant that stops working on its customer's steel, and starts buying its own, therefore moves value added from the second category to the first. For plant f in state s, industry j and year t, we estimate

$$\mathbb{E}\left[Y_{ft}\right] = \exp\left(\alpha_f + \gamma_{st} + \delta_{jt} + \sum_{t \neq 2005} \beta_t \cdot NewlyEligible_f\right) \tag{6}$$

separately for each outcome. Estimation is by Poisson pseudo-maximum likelihood (PPML), with plant, state-by-year and industry-by-year fixed effects, clustering on the plant.[26] The estimating panel spans 2001-2010 and contains 252,334 plant-years on 62,459 plants, of which 8,848 are newly eligible.

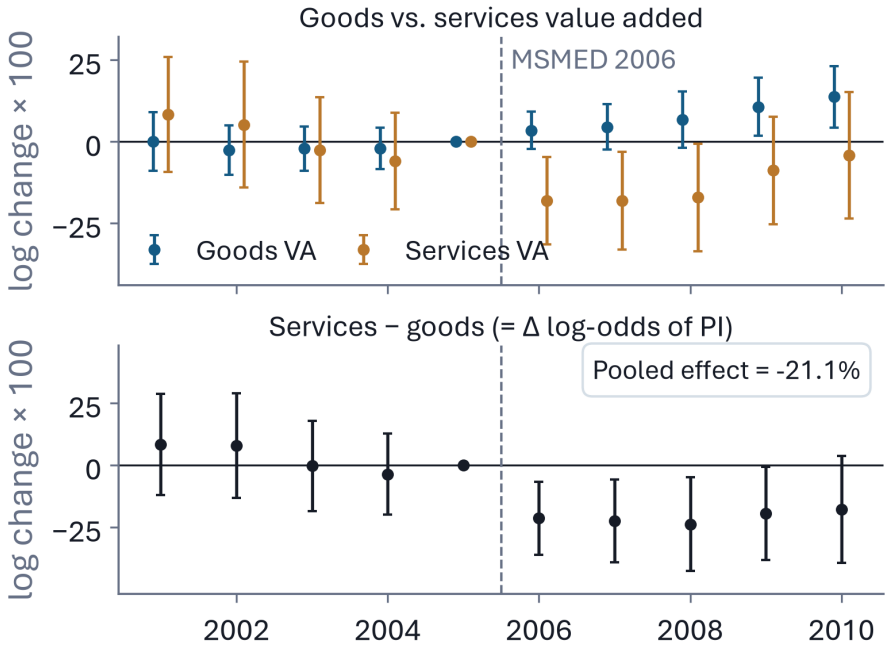
Buyer control falls when suppliers gain access to credit. Figure 12 shows that before the reform, newly eligible and other plants track each other on both margins, and no pre-period difference, in either series, is distinguishable from zero. After 2006 the two series separate immediately and then stay apart: goods value added at newly eligible plants rises steadily, standing 0.14 log points above its 2005 level by 2010, while services value added drops 0.18 log points in the first year and stays below its pre-reform level throughout. Their difference – which is the change in the log-odds of the conversion share – averages −0.211 over the five post-reform years (s.e. 0.076, $p = 0.006$), a decline of roughly a fifth.[27] Plants that gained access to formal credit shifted out of working on other firms' materials, and into buying their

[25] Banerjee and Duflo (2014) use an earlier movement in the same ceiling to show that these loans relax real borrowing constraints rather than substituting for other credit, and Rotemberg (2019) studies the 2006 change. We assign treatment from plant and machinery investment at the plant's last observation before 2006 (so it is fixed and pre-determined), and we exclude plants ever under government ownership.

[26] We estimate in levels rather than logs because roughly a fifth of plant-years report no conversion income at all. Plant-years in which the survey records the plant as closed enter as zeros, so the estimates include exit – which matters for the level of the output response (as Rotemberg 2019 finds) though not for the difference we report. Standard errors for the difference $\beta_t^S - \beta_t^G$ use the exact plant-clustered cross-covariance of the two coefficient vectors, recovered from the per-plant scores of the two fits, since the two are estimated on overlapping but not identical samples.

[27] The joint test that all five post-reform differences are zero gives $p = 0.085$, which reflects the imprecision of the individual year coefficients rather than of their average. Weighting each plant equally rather than each rupee of value added, the average difference is −0.145 (s.e. 0.037, $p < 0.001$), and the joint test gives $p < 0.001$. The response is thus smaller, but better estimated, among small plants.

Figure 12: Newly eligible plants withdraw from conversion work after the 2006 MSMED reform



Notes: Coefficients β_t from Equation (6), relative to 2005, on the ASI plant panel 2001-2010. The upper panel shows goods value added and services value added, estimated separately; the lower panel shows their difference, which is the change in the log-odds of the conversion share. Whiskers are 95 percent confidence intervals clustered on the plant, and the difference panel uses the exact cross-covariance of the two fits. The dashed line marks the reform, and the annotation is the equal-weighted mean of the five post-reform differences.

own.

This is the response that Prediction 3 attaches to noncontractibility, and it is not the response *specialization* predicts. If buyers held their suppliers' steel because they could procure or finance it more cheaply – the specialization account – then that advantage would be untouched by a reform that changes only what the *supplier* can borrow, so the arrangement would persist. What we see instead is that suppliers stop selling conversion services as soon as they can finance inventory themselves. That is what the model implies when control exists to keep a scarce, divertible resource (here, working capital) inside the order.

Taken together, these results establish that the arrangement we document in metal fabrication is a general feature of Indian manufacturing, rather than a peculiarity of one intermediary. Control over a production stage sits outside formal ownership in roughly a third of registered plants, concentrates in industries where production ties up working capital, and recedes when policy relaxes the underlying financing friction. Because none of this movement involves a

change in who owns whom, all of it would be invisible in the ownership data (as in Antràs 2003 and Boehm and Oberfield 2020) from which the prevalence of contracting frictions is usually inferred.

6 Conclusion

Buyers acquire effective control over the decisions inside their suppliers that they care most about, without acquiring the suppliers themselves. We document this margin in Indian metal fabrication, where the decision at risk is the supplier's allocation of scarce working capital to input procurement. Buyers respond through *partial integration*: visibility over the input stage, incentive contracts tied to procurement, and direct buyer sourcing of inputs. Advances are conditional on verified procurement, unconditional cash advances are absent, direct control expands when a steel-price shock raises working-capital needs, and ex-post monitoring rises when other-order payments arrive. These empirical patterns are consistent with non-contractibility over the supplier's use of liquid resources, rather than alternatives like buyer specialization in input procurement.

The same margin appears across Indian manufacturing: a control-based form of partial integration is common in the Annual Survey of Industries, falls when formal credit expands, rises where contract enforcement is weak, and is empirically distinct from vertical integration. The implication is that the relevant boundary of authority in supply chains is decision-specific. Effective control can be unbundled from ownership: buyers internalize particular noncontractible decisions without acquiring the supplier as a whole. Because doing so requires monitoring, verification, and sometimes direct sourcing, relational buyers with lower engagement costs are better positioned to source from constrained suppliers. Policies that expand short-term working-capital markets or support capable intermediaries may therefore substitute for some forms of buyer intervention while preserving the gains from supplier specialization. The costs of firm boundaries may be lower than standard models predict when governance can target the decisions where frictions bind.

A Appendix tables

Table A1: Summary statistics

	Mean	Median	SD	Min	Max	N
Panel A: lots						
Lot weight, metric tons	12.3	12.3	9.0	0.0	148.4	4,473
Delay at raw material procurement, days	26	11	90	-182	422	4,272
Delay at dispatch, days	53	22	86	-82	395	3,847
Panel B: contracts						
Contract size, metric tons	87	39	201	0	3,105	635
Lots per contract	7.0	2.0	16.6	1.0	283.0	635
Any conditional advance	0.37	0.00	0.48	0.00	1.00	778
Percent advance among advance contracts	48	40	24	3	100	286
Buyer sources raw material	0.13	0.00	0.33	0.00	1.00	790
Panel C: suppliers						
Contracts per supplier	3.5	2.0	4.5	1.0	39.0	224
Total production per supplier, metric tons	305	77	796	0	7,232	179

Notes: Panel A covers all lots with a raw material procurement record in the 2019-2022 production timelines. Delays are days between actual and planned completion of the stage, winsorized at the 1st and 99th percentiles. Contract size sums lot weights within the contract. Panel B covers contracts signed from 2019 onward with a linked supplier and recorded terms. Any conditional advance indicates a positive advance share, and percent advance is the share of contract value paid on verified sourcing. Panel C covers the suppliers of these contracts. All data are from the buyer's administrative records.

B Details on the empirical design

Our empirical analysis uses two quasi-experimental research designs. First, to test whether working capital constraints disrupt production over monthly 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 monitoring response. Second, to show how contracts adapt to credit frictions, we use a sudden increase in steel prices that raised working capital requirements. In a difference-in-differences design, we show that the steel price shock leads suppliers facing more severe credit constraints, proxied by pre-shock use of buyer-provided finance (conditional advances) and sourcing delays, to sign contracts in which the *buyer* directly sources steel on the supplier's behalf.

Table A2: 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.

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. We use past-order payments as a source of cash-flow variation for two reasons. First, unlike typical balance sheets, our data observe order completions at the daily or weekly time scale relevant for Just-In-Time production. Second, variation in order completion is partly outside the supplier's control. Using other-order completion timing is also theoretically

Table A3: 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.

appealing. As the model makes clear, rationalizing the observed contracts requires short-run uncertainty in cash-on-hand. Because contracts fix order size and price in advance, uncertain order timing is the empirical source of that uncertainty.

Construction of downstream delays instrument Note that $\sum_{o'} Payment^{INST}_{o't}$ 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 control for planned production $\sum_{o'} Payment^{PLAN}_{o't}$.[28]

[28]Borusyak and Hull (2020) show that one approach to unbiased estimation of (1) is to specify the data-

Table A4: 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*

Notes: Table shows extensive-margin analogues of Table A2: effects of an indicator for any other-order payment in a month on an indicator for any raw material sourcing for order o in that month. Column 1 provides OLS estimates using observed payments, conditional on indicators for planned sourcing and planned payment. Columns 2-4 give IV estimates (the reduced-form, first, and second stages) when we instrument for observed payment timing using predicted timing, as in Section 2.3. All regressions are at the order-month level. All data are from 2019-2022 production timelines from the buyer.

We now discuss IV relevance, exogeneity, and exclusion restrictions, in that order. The aggre-

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 there is no natural level at which to permute delays. 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 rather than in a nearby month. This control is imperfect because payments planned for month t may also be likely to arrive in months $t-1$ and $t+1$.

Table A5: Substitution from incentives to control: time-invariant covariates in lieu of supplier fixed effects

Dependent Variable:	=1 if buyer sources raw material			
Model:	(1)	(2)	(3)	(4)
Variables				
Any advance × Post (≥2q)	0.261**		0.399***	
	(0.122)		(0.133)	
Percent advance × Post (≥2q)		0.214		0.703**
		(0.343)		(0.276)
Log mean order size × Post (≥2q)			-0.297**	-0.375***
			(0.120)	(0.129)
Fixed-effects				
Quarter	Yes	Yes	Yes	Yes
Controls				
Any advance	Yes		Yes	
Percent advance		Yes		Yes
Log mean order size			Yes	Yes
Fit statistics				
Observations	968	968	968	968
Dependent variable mean	0.077	0.077	0.077	0.077

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

Notes: Table repeats the specifications of Table 2, replacing supplier fixed effects with the advance and size variables in levels as controls (coefficients suppressed). Lot-level panel weighted by lot tonnage; standard errors clustered by supplier; quarter fixed effects in all columns. See the notes to Table 2 for sample details.

gated order-time-level instrument has a strong first stage, with a coefficient of 0.269 and an F-statistic of 1653.2 (shown in Column 3 of Table A2). In other words, predicted variation in total payments due to customer-driven postponements strongly predicts realized variation in payments. Panel A of Figure 4 provides intuition for the first stage: at the order level, delay 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 $\sum_{o'} Payment_{o't}^{INST}$ 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 condition on planned production; future versions will explore alternative measures of delays relative to plan. For (ii), customer-driven delays $PredDelay_{ig}$ must be

common shocks outside any supplier's control. The instrument would be contaminated if a shock originating at supplier s pushed back the overall production timeline for the customer and therefore delayed supplier s'. Several features of the setting limit this concern. Suppliers in our sample are typically small, accounting for under 5% of any construction project, so even when contract delays reflect delayed inputs, the source of delays is unlikely to be in our sample. Furthermore, relationship managers cited several downstream delay sources beyond input shortages, including weather disruptions, regulatory holdups, and labor shortages due to pandemic-related migration restrictions.

Finally, the exclusion restriction would fail if predicted payments ($\sum_{o' \neq o} Payment_{o't}^{INST}$) affected raw material sourcing ($Steel_{ot}^{OBS}$) through channels *other* than observed payments ($\sum_{o' \neq o} Payment_{o't}^{OBS}$) The challenge is separating stages: since payments closely follow physical completion of the final product, a different scarce input could be driving the response. For instance, completion of goods for order o' might free up labor or plant space to start order o, and the supplier might therefore procure steel at the same time. If most lots were physically produced and paid for in different months, we could rule out this possibility by constructing an alternative $\sum_{o' \neq o} Production_{o't}^{INST}$ 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 about whether some scarce input binds, because $\beta^{OBS} > 0$ indicates that new orders start when past orders end. Several features of the setting suggest that credit is the relevant scarce input. First, in the industrial clusters we study, a 'gang' of workers can be hired on a spot market on short notice, so worker availability is rarely binding. Second, 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 5, suggests that buyers are aware of short-term supplier credit constraints.

To further test the credit constraints mechanism, we estimate heterogeneous elasticities in Equation (1) using proxies for pre-existing credit constraints. The first is firm-level use of partial financing (conditional advances) from the buyer. Conversations with the buyer suggest that the most constrained firms take up financing: advances are offered only to suppliers lacking independent connections with working-capital financing institutions, and the buyer charges suppliers above-market interest for the duration of the order. Thus orders with partial advances are the orders where we should expect responses to cash inflows.[29] The second proxy is supplier region: if credit constraints are relevant, then suppliers in Central and North India –

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

which generally has lower income levels, weaker contracting institutions, and thinner financial markets than West and South India – should be more responsive to the timing of payments.

C Proofs

Proof of Proposition 1. Define $D(\tau) \equiv \Pi_B(\tau) - \Pi_A(\tau)$. Using the payoffs,

$$D(\tau) \;=\; (1-\mu)\,\rho_1(\tau)\big[V+\tau\big] + K - F,$$

the financing surplus $b\tau$ cancelling since the buyer covers the gap under both forms. D is continuous and strictly increasing: ρ_1 is strictly increasing (Assumption 1(i)) and $V+\tau$ is increasing, both positive. By Assumption 1(ii), $D(\tau_0) < 0$. As $\tau \to \infty$, $\rho_1(\tau) \geq \rho_1(\tau_0) > 0$ while $V+\tau$ grows without bound, so $D(\tau) \to \infty$. Hence there is a unique $\tau_C > \tau_0$ with $D(\tau_C) = 0$: conditional advances are preferred below τ_C and buyer sourcing above it. Since $\tau = m - \bar{\ell}$, a rise in input cost m switches suppliers on conditional advances whose gaps cross τ_C from incentives to control. Absent the friction ($\rho_1 \equiv 0$), $\Pi_A - \Pi_U = \mu\,\rho_1(\tau)[V+\tau] - K = -K < 0$, so conditional advances are never used; and $D(\tau) = K - F$ is constant in τ, so whether the buyer sources is independent of the gap and no substitution occurs. □

Proof of Proposition 2. Fix a supplier-retained regime $g \in \{0, U, A\}$, with residual diversion risk $\rho_g(\tau)$ – equal to $\rho_1(\tau)$ under 0 and U, and $(1-\mu)\rho_1(\tau)$ under A – advance at stake a_g ($a_0 = 0$, $a_U = a_A = \tau$), and realized liquidity ℓ. Before sourcing is complete, the buyer chooses $h \in [0,1)$ to minimize

$$s(\ell)\,(1-h)\,\rho_g(\tau)\big[V + a_g\big] \;-\; \gamma h \;+\; k(h).$$

The objective is strictly convex in h, with derivative $k'(h) - s(\ell)\,\rho_g(\tau)[V+a_g] - \gamma$; since $k'(0) = 0$ and $k'(h) \to \infty$ as $h \to 1$, the minimizer is interior and satisfies

$$k'(h^*) \;=\; s(\ell)\,\rho_g(\tau)\big[V + a_g\big] + \gamma.$$

Because s and k' are strictly increasing, $h^*(\ell)$ is strictly increasing in realized liquidity ℓ whenever $\rho_1(\tau) > 0$. Once raw material is procured, the input decision is closed, the diversion term vanishes, and $k'(h^*) = \gamma$: monitoring falls to the knowledge level $k'^{-1}(\gamma)$. If $\rho_1 \equiv 0$, then $k'(h^*) = \gamma$ at every ℓ: visits occur whenever $\gamma > 0$ but are invariant to liquidity. □

Proof of Proposition 4. Under buyer sourcing the supplier retains d_2, so the buyer's continuation value is $V = (1-\rho_2)(\bar{V} - c + qR)$ and $\Pi_B(\tau) = (1-\rho_2)(\bar{V} - c + qR) + b\tau - F$. Under full

integration the buyer implements delivery directly, with no delivery risk but no outside orders: $\Pi_{VI}(\tau) = \bar{V} - c + b\tau - F$. Taking the difference, $b\tau - F$ cancels and

$$\Pi_B(\tau) - \Pi_{VI}(\tau) \;=\; (1-\rho_2)\big(\bar{V} - c + qR\big) - \big(\bar{V} - c\big) \;=\; qR - \rho_2\big(\bar{V} - c + qR\big).$$

This expression is increasing in q (derivative $(1-\rho_2)R > 0$) and, when $q > 0$, in R, and decreasing in ρ_2 (derivative $-(\bar{V} - c + qR) < 0$). Setting it to zero gives $(1-\rho_2)qR = \rho_2(\bar{V} - c)$, so $q^* = \frac{\rho_2(\bar{V}-c)}{(1-\rho_2)R}$, which lies in $(0,1)$ exactly when $0 < \rho_2(\bar{V} - c) < (1-\rho_2)R$; the buyer prefers full integration for $q < q^*$ and buyer sourcing for $q > q^*$. □

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 ex-post monitoring costs ($k(\cdot)$), lower verification costs (K), and lower direct-sourcing costs (F).

Proposition 5 (Relational firms can sustain partial integration for a wider set of suppliers)**.** *Consider two buyers $i \in \{1,2\}$ facing the same supplier and identical in all other primitives, where buyer 1 has uniformly lower engagement costs: $k_1(h) < k_2(h)$ for all $h > 0$, $K_1 < K_2$, 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. A lower ex-post monitoring cost weakly lowers expected input-stage losses under A: for every realized ℓ, the minimized loss

$$\min_{h\in[0,1]} \Big\{\big[1 - \sigma(x_A(\ell))\big]V(q) + \sigma(x_A(\ell))\,(1 - h\nu)\,\alpha_A\rho(\tau)\big[V(q) + \tau\big] + k_i(h)\Big\}$$

is weakly lower for buyer 1 than for buyer 2 because $k_1(h) < k_2(h)$ pointwise. Hence $\Lambda_{A,1}(\tau) \leq \Lambda_{A,2}(\tau)$. Therefore, for each τ,

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

whenever buyer 2 uses conditional advances. Likewise,

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

Therefore buyer 1 gets weakly higher payoff from every partial-integration option and strictly higher payoff from any option that uses verification, 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. □

Thus, relational firms can profitably sustain the incentive- and control-based forms of partial integration (*A* and *B*) for a wider set of constrained suppliers. 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 both fills the financing gap and stands ready against momentary shortfalls, the contracting problem disappears and with it the need for buyer intervention.

Proposition 6 (Completing working capital markets removes the need for buyer intervention)**.** *Suppose the supplier gains access to on-demand outside finance at no buyer-relevant cost – a working-capital line it can draw at any moment, leaving c, qR, and $V(q)$ unchanged – so that the input decision is always live and the financing gap is closed ($\sigma = 1$, $\tau = 0$). If, in addition, $qR > G$ and $F > \chi$, then market governance, under which the supplier keeps both stages and the buyer does not intervene, strictly dominates A, B, and VI.*

Proof. On-demand credit lets the supplier bridge any momentary shortfall, so the input decision is always live ($\sigma = 1$); with the gap closed, $\tau = 0$ and there is no diversion risk ($\rho(0) = 0$).

The input-stage loss under no intervention is zero and market governance yields continuation value $V(q)$.

Under conditional advances, the buyer obtains at most $V(q)-K$, because the input-stage problem has already been removed (and $b\tau = 0$ at $\tau = 0$) but the verification cost would still be paid. Under buyer sourcing, the buyer obtains $V(q)-(F-\chi)$. Under full integration, the buyer obtains $\bar{V} - c + G - F + \chi$.

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

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

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 net of sourcing costs, i.e. whenever $qR > G - (F - \chi)$, which holds since $qR > G$ and $F > \chi$. In that case

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

Hence market governance strictly dominates all costly intervention options once on-demand outside finance is available; an unconditional advance at $\tau = 0$ is vacuous, coinciding with market governance. □

Buyer sourcing that reflects a procurement advantage large enough that $\chi > F$ would survive complete working-capital markets. Improved credit access therefore eliminates partial integration driven by noncontractibility while leaving specialization-driven sourcing intact – a margin we use when interpreting the credit-expansion evidence in the ASI.

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.

Alfaro-Urena, Alonso, Isabela Manelici, and Jose P Vasquez (2022). "The effects of joining multinational supply chains: New evidence from firm-to-firm linkages". In: *The Quarterly Journal of Economics* 137.3, pp. 1495–1552.

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, Ali Hortaçsu, Mary Jialin Li, et al. (2019). "How Wide Is the Firm Border?" In: *Quarterly Journal of Economics* 134.4, pp. 1845–1908.

Atalay, Enghin, Ali Hortaçsu, and Chad Syverson (2014). "Vertical Integration and Input Flows". In: *American Economic Review* 104.4, pp. 1120–1148.

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.

Banerjee, Abhijit V. and Esther Duflo (2014). "Do firms want to borrow more? Testing credit constraints using a directed lending program". In: *Review of Economic Studies* 81.2, pp. 572–607.

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.

Burkart, Mike and Tore Ellingsen (2004). "In-Kind Finance: A Theory of Trade Credit". In: *American Economic Review* 94.3, pp. 569–590.

Coase, Ronald H. (1937). "The Nature of the Firm". In: *Economica* 4.16, pp. 386–405.

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.

Fazzari, Steven M., R. Glenn Hubbard, and Bruce C. Petersen (1988). "Financing Constraints and Corporate Investment". In: *Brookings Papers on Economic Activity* 1988.1, pp. 141–206.

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.

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.

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.

Ramondo, Natalia, Veronica Rappoport, and Kim J. Ruhl (2016). "Intrafirm Trade and Vertical Fragmentation in U.S. Multinational Corporations". In: *Journal of International Economics* 98, pp. 51–59.

Rotemberg, Martin (2019). "Equilibrium Effects of Firm Subsidies". In: *American Economic Review* 109.10, pp. 3475–3513.

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