Vertical Integration, Supplier Behavior, and Quality Upgrading among Exporters

---

Christopher Hansman

Imperial College London

Jonas Hjort

Columbia University, Center for Economic Policy and Research, and National Bureau of Economic Research

Gianmarco León-Ciliotta

Universitat Pompeu Fabra, Barcelona Graduate School of Economics, Institute of Political Economy and Governance, and Center for Economic Policy and Research

Matthieu Teachout

International Growth Centre

We study the relationship between firms' output quality and organizational structure. Using data on the production and transaction chain that makes up Peruvian fish meal manufacturing, we establish three results. First, firms integrate suppliers when the quality premium rises for exogenous reasons. Second, suppliers change their behavior to better maintain input quality when vertically integrated. Third, firms produce a higher share of high-quality output when weather and supplier availability shocks shift them into using integrated suppliers. Overall, our results indicate that quality upgrading is an important motive for integrating suppliers facing a quantity-quality trade-off, as classical theories of the firm predict.

An earlier version of this paper was the winner of the 2018 Oliver Williamson Award for best paper presented at the Society of Institutional and Organizational Economics' annual

Electronically published August 11, 2020

[*Journal of Political Economy*, 2020, vol. 128, no. 9]

© 2020 by The University of Chicago. All rights reserved. 0022-3808/2020/12809-0007$10.00

I. Introduction

Why do so many of our economic transactions occur within firm boundaries (Coase 1937; Gibbons 2005a; Lafontaine and Slade 2007)? Vertical integration occurs for many different reasons, and motives vary by context.[1] However, as global incomes rise and barriers to trade fall, one potential motive has gained increased relevance: firms integrating to improve product quality. Access to wealthier, quality-sensitive markets brings rising returns to output quality,[2] but producing high-quality output typically requires high-quality inputs (see, e.g., Kugler and Verhoogen 2012; Halpern, Koren, and Szeidl 2015; Amodio and Martinez-Carrasco 2018; Bastos, Silva, and Verhoogen 2018; Eslava, Fieler, and Xu 2018). Because input quality is often hard for firms to measure and contract over (Gibbons 2005a; Lafontaine and Slade 2007)—especially where institutions are weak (Woodruff 2002; Nunn 2007)—organizational structure may play a crucial role in firms' ability to meet demand for quality.

In this paper, we test the hypothesis that firms vertically integrate to produce higher-quality products. This conjecture is inspired by classical theories characterizing how firm boundaries are expected to respond to output objectives (Baker, Gibbons, and Murphy 2001, 2002; Gibbons 2005a, 2005b) when suppliers multitask (Holmstrom and Milgrom 1991). However, given the rarity of data on product quality and internal firm structure and the challenges of isolating firm strategy from confounding factors, causal evidence

---

conference. We are especially grateful to editor Chad Syverson, three anonymous referees, Francesco Amodio, Bob Gibbons, Amit Khandelwal, Kaivan Munshi, Eric Verhoogen, and Chris Woodruff for suggestions that significantly improved the paper. We also thank Charles Angelucci, John Asker, David Atkin, Jie Bai, Jing Cai, Lorenzo Casaburi, Wouter Dessein, Dave Donaldson, Marina Halac, Oliver Hart, Rocco Macchiavello, Ameet Morjaria, Andrea Prat, John Van Reenen, Xiao Yu Wang, and Daniel Wolfenzon for comments and advice; seminar participants at Cambridge, Columbia, London School of Economics, Massachusetts Institute of Technology, National Bureau of Economic Research, Universitat Pompeu Fabra, Zurich, the Markets and Organizations in Emerging Economies Conference, the Society of Institutional and Organizational Economics, the Summer School in Development Economics, and the International Growth Centre/Stanford Center for International Development Conference on Firms, Trade, and Development; Patrick Kennedy, Leonor Lamas, and Sebastián Sardón for great research assistance and fieldwork; and the interviewed executives in the Peruvian fish meal manufacturing industry for their time. The Peruvian Ministry of Production, Cesar Casahuamán, and the International Fishmeal and Fish Oil Organisation (in particular, Enrico Bachis and Jorge Mora) kindly shared access to data. León-Ciliotta acknowledges financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV2015-00563) and grants ECO2011-25272 and RYC-2017-23172. Data are provided as supplementary material online.

[1] Empirical work on the causes and consequences of firms' choice of organizational structure in developing countries began in earnest with Woodruff (2002). Gibbons (2005a), Lafontaine and Slade (2007), and Bresnahan and Levin (2013) provide excellent overviews of the broader literature on firms' structure.

[2] See, e.g., Sutton (1991, 1998); Hallak (2006); Verhoogen (2008); Manova and Zhang (2012); Atkin, Khandelwal, and Osman (2017); Bastos, Silva, and Verhoogen (2018).

on the extent to which firms change their organizational structure to upgrade quality has remained elusive.

The context we study—the Peruvian fish meal industry—enables progress. The structure of the sector is simple. Independent and integrated suppliers deliver inputs of hard-to-observe quality to manufacturers. Manufacturers convert these inputs into a vertically differentiated but otherwise homogeneous product.[3] Uniquely rich data on the sector's entire chain of production are available, including within-firm transactions and direct measures of output quality. Finally, there is substantial (and plausibly exogenous) variation in the quality premium—the price differential between high- and low-quality fish meal. This allows us to isolate explicit strategic responses to incentives to upgrade quality.

Our analysis proceeds in four steps. We first present a simple theoretical framework that describes how and why a firm's choice of organizational structure may depend on its output quality objectives. We then ask whether quality-upgrading motives are—empirically—a direct determinant of integration decisions—that is, whether a manufacturer is more likely to integrate its suppliers when its returns to shifting from low- to high-quality production are higher. Next, we explore the mechanisms through which output quality objectives may impact integration decisions. We estimate how organizational structure affects supplier behavior, focusing particularly on "switchers"—suppliers that supply the same plant before and after being integrated (or sold). To conclude, we investigate whether integration ultimately raises output quality.

There are several reasons why unique data are available on the fish meal industry in Peru. The regulatory authorities record all transactions between fish meal plants and suppliers and require firms to report each of their plants' production of fish meal in the "prime" (high) quality and the "fair average" (low) quality range each month, providing a direct measure of quality.[4] We can link these input and output quantities to all export transactions, which are recorded in customs data. Furthermore, researchers—but not manufacturers—directly observe both independent and integrated suppliers' behavior because fishing boats are required to transmit

---

[3] Fish meal is a brown powder made by burning or steaming fish and is mostly used as animal feed. Peru's fish meal industry accounts for around 3% of GDP (Paredes and Gutierrez 2008; De La Puente et al. 2011). Price differentials across shipments of a given quality level in a given time period are negligible (see sec. II.B). Our focus on a vertically differentiated but horizontally homogeneous product is inspired by influential earlier papers testing market power theories of integration in sectors producing homogeneous products (Syverson 2004; Hortaçsu and Syverson 2007; Foster, Haltiwanger, and Syverson 2008).

[4] See Goldberg and Pavcnik (2007), Khandelwal (2010), and Hallak and Schott (2011) for discussion of the indirect quality proxies used in the existing literature, which risk conflating quality with markups and horizontal differentiation, and Atkin, Khandelwal, and Osman (2017) for an example of direct measures of quality.

GPS signals to the regulatory authorities.[5] In combination, these data sources allow us to track the flow of goods—from suppliers to manufacturers to foreign buyers—and provide the measures of output quality and firm-supplier transactions necessary to establish a correlation between the quality of a firm's output and the organizational structure of its production chain.

However, even if documented with ideal data, such a correlation may reflect third factors rather than an explicit organizational choice made to climb the quality ladder. It could be, for example, that productivity or demand affect both firms' choice of structure and products produced without the two being directly related. It could also be that firms integrate for reasons other than quality—for example, to assure their own or restrict competitors' general access to inputs—but in the process coincidentally produce higher-quality output. To identify a direct relationship between output quality objectives and integration, exogenous variation in incentives to upgrade quality—the quality premium—that firms are differentially exposed to is needed.

The quality premium facing Peruvian fish meal firms varies considerably during the period we study. This allows us to test our causal hypothesis. Our empirical strategy exploits season-to-season variation in the quality premium that is due to fluctuations in the regulatory fishing quotas of countries other than Peru that produce high-quality fish meal. We construct an instrument for the quality premium and test whether these fluctuations affect firms differently depending on their scope for upgrading quality.

We begin our analysis with a stylized model, which demonstrates how characteristics of the Peruvian fish meal industry map directly to the existing theoretical work we build on. Fish meal manufacturers face two important contracting challenges. First, the quality of the product's primary input—fish—is difficult to observe and, because of its perishable nature, even harder to contract on. Second, the presence of outside options—other fish meal firms that may value input quality less—complicates controlling the incentives of an independent supplier (see also McMillan and Woodruff 1999).[6] Holmstrom and Milgrom (1991) elegantly demonstrate how, when suppliers face a trade-off between producing inputs of

---

[5] The regulators do not allow manufacturers to access data on the behavior of independent suppliers. This is the primary reason why manufacturers and independent suppliers cannot contract over GPS-measured actions.

[6] In theory, it is possible to imperfectly measure fish quality with chemical tests. As discussed in sec. II, industry insiders informed us that such tests were much too expensive and impractical to use during our data period. Alternatively, manufacturers and their suppliers could attempt to contract on plants' output quality. This would be difficult because of noise—input from multiple suppliers is, for technological reasons, typically used in a given batch of fish meal, and other hard-to-measure exogenous factors also influence output quality realizations—and, more importantly, because outside inspectors would need to be able to

high quality or in high volumes, weakening incentives over easier-to-measure quantity may be necessary to ensure that suppliers do not neglect quality (see also Holmstrom and Tirole 1991; Holmstrom 1999). In many situations, the best or only way to do so may be to bring the suppliers inside the firm (Baker, Gibbons, and Murphy 2001, 2002; Gibbons 2005a, 2005b).

To take these textbook theoretical insights to the data, we first demonstrate that output quality is in fact significantly positively correlated with integration. Our primary measure of quality is the share of a firm's output that is of high quality grade—which we directly observe in production data. We also consider a fine-grained (but not directly observed) measure of the average quality grade of a firm's output. Similarly, we consider two different measures of vertical integration, one based on use of integrated versus independent suppliers and the other based on supplier ownership. The relationship we establish holds across each of our measures of a firm's output quality and each of our measures of a firm's organizational structure.[7]

We then proceed to our central empirical analysis, which consists of three key pieces of evidence. The first—and this paper's main finding—shows evidence that vertical integration is used by firms as a strategy for increasing output quality. To demonstrate this, we develop an instrumental variable (IV) for firm-specific incentives to upgrade quality. We instrument for the quality premium—the difference between the price of high and low quality grades—using the regulatory fishing quotas of other top exporters. Because the other top exporters specialize in high quality grades and because their production quantities are driven by country-specific regulatory fishing quotas, these quotas generate plausibly exogenous variation in the premium. We test whether this variation differentially impacts Peruvian firms' integration decisions depending on their firm-specific scope for upgrading quality. Firms that are already producing

---

determine whether low output quality was due to poor input quality or due to actions taken by the fish meal plant itself during the production process. The dynamic version of our model in app. C (apps. A–D are available online) demonstrates why the presence of other fish meal manufacturers that value input quality less can make repeated-interaction solutions to these challenges infeasible.

[7] Our primary measure of integration is the fraction of inputs that are sourced from vertically integrated suppliers—"share VI"—a measure that is motivated by our hypothesis that the characteristics of the inputs actually delivered by a supplier change with integration so that integration and output quality are causally linked via firms' (and individual plants') production process. (On the use of integrated vs. independent suppliers, see also, among others, Baker and Hubbard [2003]; Atalay, Hortaçsu, and Syverson [2014]; Breza and Liberman [2017]). Alternatively, we also consider the number of suppliers owned. Note that in our setting, since boats' total seasonal catch is governed by a quota—and boats almost always exhaust their quota over the course of a season—Peruvian fish meal manufacturers can generally increase the total amount of inputs they obtain from integrated suppliers in a given production season only by acquiring suppliers (see sec. II).

mostly high-quality output have little scope to improve quality further and are hence unlikely to respond strategically to an increase in the quality premium. Conversely, firms producing mostly low-quality output have significant scope to upgrade. Thus, if firms use integration as a strategy for upgrading quality, then firms producing primarily or exclusively high-quality output will face weaker incentives to integrate when the quality premium rises.[8]

We find that Peruvian manufacturers integrate when their incentives to upgrade quality rise, and vice versa. The industry as a whole integrates when the quality premium increases for exogenous reasons, and the integration response is stronger for firms with greater scope for upgrading quality.[9] Firms adjust the fraction of their inputs that are sourced from integrated suppliers in response to exogenous variation in the quality premium almost entirely via the extensive margin of "share VI"—that is, by acquiring and selling suppliers. In an alternative approach shown in appendix D, we exploit a different form of variation to show that firms similarly integrate when faced with greater firm-specific relative demand for high-quality output. Finally and crucially, we show that firms' organizational response to the quality premium does not reflect associated income shocks or general incentives to expand production of fish meal of any quality: firms do not integrate suppliers when faced with higher average prices.

This first piece of evidence is hard to reconcile with alternative theories in which higher output quality is an unforeseen by-product of vertical integration driven by other motives and with stories wherein organizational structure and output quality are not causally linked in the minds of firms. Several of the most prominent alternative explanations—such as firms integrating suppliers for general supply assurance reasons but coincidentally achieving higher quality in the process or foreclosure motivations—are also inconsistent with other features of the context we study and auxiliary

---

[8] A potential criticism of this argument is that firm-specific scope for quality upgrading might also correlate with some unobservable related to the marginal cost of integration and/or quality upgrading. The most natural forms of such arguments—that firms producing high-quality output (*a*) have low marginal costs of further upgrading and are also more likely to integrate in general or (*b*) have low marginal costs of integration and are also more likely to upgrade in general—predict the opposite of our findings. Furthermore, as our primary measure of quality is a share (rather than a level), arguments such as *a* cannot hold for firms that are already producing exclusively high quality, as they mechanically have no scope to improve further. Our results considering these firms are similar.

[9] The long-term trend is toward more integration in the Peruvian fish meal industry, and the long-term trends in demand for quality and average output quality in Peru are also positive. These broad patterns are consistent with our hypothesis. However, it is higher-frequency variation around the long-term trends that we exploit to test our hypothesis. For example, we also observe deintegration during our data period—sales of boats from fish meal firms to independent co-ops or captains and from one fish meal firm to another.

findings.[10] While we cannot rule out that alternative mechanisms also play a role, our results indicate that quality upgrading itself is an important motive for integrating suppliers.

Next, we explore why firms use integration as a strategy for upgrading quality. Our second key piece of evidence shows that integration changes suppliers' behavior, causing them to shift toward quality-increasing actions. We proxy for actions that increase input quality—that is, fish freshness (FAO 1986)[11]—using GPS-based measures. We show that—as the managers in the industry we interviewed reported to us[12]—a given supplier supplying a given plant delivers lower total quantities but inputs whose quality has been better maintained when integrated with the plant. We also show that, in the context we study, it is integration per se—not repeated interactions—that influences a supplier's quantity-versus-quality behavior. This result is consistent with a dynamic version of our model—which, along with the result, is shown in appendix C—and with the fact that suppliers that deintegrate from a firm/plant supply that firm/plant almost as often after the change in status. Finally, we consider the possibility that integration affects behavior not via a supplier's quantity-quality trade-off but instead via associated knowledge transfers of the form that Atalay, Hortaçsu, and Syverson (2014) convincingly show occur in the United States after integration. We find that a given supplier behaves "as an integrated supplier" only when supplying its owner firm and not when owned by one firm but supplying another. Reconciling this finding with knowledge transfer theories would require such knowledge transfers to be useful only when supplying the parent firm. We ultimately cannot rule out that incentives emanating from organizational structure itself other than those that our model focuses on also help explain the impact of integration on supplier behavior, but our third piece of evidence suggests that any such alternative incentives would need to ultimately benefit downstream output quality.

Our third piece of evidence indicates that vertical integration in fact increases output quality, as managers in the industry reported to us. (In

---

[10] We show that (i) Peruvian fish meal manufacturers appear to achieve general supply assurance primarily through repeated interactions with independent suppliers (see also Martinez-Carrasco 2017)—the quantity supplied by a given supplier to a given firm is in fact lower after integration (or before deintegration)—and (ii) repeated interactions with suppliers do not enable firms to produce higher-quality output. Similarly, the relationship between output quality and integration holds when we control for a firm's share of the industry's total production, in contrast to what traditional foreclosure theories would predict.

[11] "Freshness of raw material is important in its effect on the quality of the protein in [quality of] the end product [fish meal]. The importance of minimizing the time between catching fish and processing, and of keeping the fish at low temperatures by icing [which reduces the amount of fish a boat can fit], has already been mentioned" (FAO 1986, sec. 10.1.2).

[12] In the words of a prominent executive of Peru's National Fisheries Society, "As a consequence of integration, they must adopt my rules. Things like saying, 'Hey, you must offload raw 24 hours after having caught it, at the maximum'" (authors' translation).

the words of the managing director of Pesquera Diamante, Peru's third-largest fish meal company, "Around 80% of my company's fish meal is high quality. If all my inputs came from integrated boats, around 95% of my fish meal output would be of high quality" [authors' translation].) We first show that the firm-level relationship between the share of inputs coming from integrated suppliers—share VI—and output quality holds also at the individual plant level, including within firms. We then instrument for a plant's share VI using both a "leave-firm-supplied-out" measure of the local presence of a particular type of supplier that is prohibited by regulation from being integrated and weather shocks that differentially affect integrated and independent suppliers' whereabouts.[13] The IV estimates are similar to each other but somewhat larger than ordinary least squares (OLS) estimates.

When viewed through the lens of our model, the three key pieces of evidence we present each follow logically from one another. We conclude that aspiring to produce higher-quality products can motivate firms to vertically integrate and that the reason appears to be that integration changes supplier behavior in a way that increases output quality. Because input quality is so frequently difficult to observe (and hence incentivize), the challenges we describe—while far from universal—are likely typical of industries producing vertically differentiated output, particularly in settings where contracts are difficult to enforce.[14]

Our study bridges and advances the literatures on the boundaries of the firm and quality upgrading. We make three contributions to the former. First, we identify an overlooked motivation for vertical integration. By showing that firms vertically integrate to raise output quality, we advance the body of work on the causes of organizational form (for seminal empirical work, see, e.g., Hart, Shleifer, and Vishny 1997; Baker and Hubbard 2003, 2004; Forbes and Lederman 2009). Existing studies convincingly demonstrate how firms change their relative use of integrated suppliers in response to changes in, for example, available contracts (Breza and Liberman 2017) or monitoring technology (Baker and Hubbard 2003).

---

[13] The presence of such suppliers fluctuates because of natural variation in fish density, weather, and decisions made by their captains. The logic behind this first instrument for share VI is simply that a plant—holding fixed output quality objectives—will be forced to source a higher share of its inputs from integrated suppliers when there happens to be a local scarcity of independent suppliers. Plausible arguments against the exclusion restriction underlying our interpretation of the results from this first IV approach would arguably require a positive sign on the first stage (negative correlation between the use of independent suppliers by different plants in a locality), a negative sign on the second stage (use of independent suppliers increasing output quality), and/or a time-varying, location-level component of output quality (that goes beyond the presence of independent suppliers)—none of which we find.

[14] There is a robust relationship between countries' input-output structure and their level of contract enforcement (Nunn 2007; Boehm 2018), and vertical integration is more common in developing countries (Acemoglu, Johnson, and Mitton 2005; Macchiavello 2011).

We instead study how firms change their organizational structure when their output objectives change.[15]

Second, and building on earlier studies of the behavior of integrated and independent suppliers (Mullainathan and Sharfstein 2001; Baker and Hubbard 2003, 2004; Macchiavello and Miquel-Florensa 2016), we provide what is to our knowledge the first evidence on how integration changes the quality-oriented behavior of a given supplier supplying a given firm. To do so, we follow Atalay, Hortaçsu, and Syverson (2014) and exploit changes in integration within supplier-firm pairs.

Finally, we show evidence that vertically integrating raises output quality, which to our knowledge has not been done before. The one-dimensional nature of quality differentiation in our setting allows us to document this.[16] In general, there is little existing evidence on causal consequences of organizational structure for firm performance (see Forbes and Lederman 2010, and see Gibbons and Roberts 2013 for a notable exception). Our results also imply that using independent suppliers is often efficient for producing output in high volumes rather than of high quality (see also Kosová, Lafontaine, and Perrigot 2013). An especially unusual aspect of this paper is that the data and variation we exploit allow us to identify both the effectiveness of a particular firm strategy and corresponding determinants of its use. We can therefore show that Peruvian fish meal manufacturers vertically integrate when quality objectives indicate that they should do so.

Both the friction—imperfect contracting over input quality—and the firm objective—producing high-quality output—that we focus on are especially relevant for poorer countries attempting to help meet growing global demand for quality. This connects our paper with a smaller empirical literature on the causes and consequences of firms' choice of organizational structure in the developing world that began with Woodruff's (2002) landmark study (see also Natividad 2014; Macchiavello and Miquel-Florensa 2016; Martinez-Carrasco 2017).[17]

---

[15] In a superficial sense, our finding that higher average fish meal prices do not lead to more integration in the Peruvian fish meal industry contrasts with the innovative work of Alfaro et al. (2016). We see our results as largely consistent with and complementary to theirs, however. Both their analysis and ours emphasize the impact of prices in the context of certain goods—in our case, high-quality products—where integration generates a gain in efficiency. We highlight that this efficiency gain is not generic but rather depends on firms' quality objectives, while they emphasize that efficiency gains can also depend on the need to coordinate production stages.

[16] In settings where product differentiation is multidimensional, an analysis such as ours would be difficult. Like this paper, the pioneering study by Forbes and Lederman (2010) exploits exogenous drivers of use of integrated suppliers, showing that routes that airlines self-manage have fewer delays/cancellations (see also Gil and Kim 2016; Gil, Kim, and Zanarone 2016). Other important evidence on the consequences of organizational structure includes, among others, Novak and Stern (2008) and Gil (2009).

[17] Woodruff finds that forward integration is less common in the Mexican footwear industry when noncontractible investment by retailers is important, as the property rights framework predicts (Grossman and Hart 1986; Hart and Moore 1990). Macchiavello and

The literature on quality upgrading is larger. It is now well documented that producers of high-quality goods use high-quality inputs (Kugler and Verhoogen 2012; Halpern, Koren, and Szeidl 2015; Amodio and Martinez-Carrasco 2018; Bastos, Silva, and Verhoogen 2018; Eslava, Fieler, and Xu 2018) and skilled workers (Verhoogen 2008; Frías, Kaplan, and Verhoogen 2009; Brambilla, Lederman, and Porto 2012, 2019; Brambilla and Porto 2016) and export to richer destinations (Hallak 2006; Verhoogen 2008; Manova and Zhang 2012; Atkin, Khandelwal, and Osman 2017; Bastos, Silva, and Verhoogen 2018). Firms with such a profile tend to be on average larger, more productive, and based in richer countries themselves and to face foreign competition in low-quality segments (Schott 2004; Hummels and Klenow 2005; Baldwin and Harrigan 2011; Johnson 2012; Medina 2017). We provide the first evidence linking quality upgrading to the boundaries of the firm.

II. Background on Peru's Fish Meal Manufacturing Sector

In this section, we provide an overview of Peru's fish meal manufacturing sector. We argue that three features are particularly salient for firms attempting to source high-quality inputs: input quantity is measurable at the time of delivery, but input quality is not, and formal contracts appear to be difficult to write.

A. *Sector Profile*

Fish meal is a brown powder made by burning or steaming fish (in Peru, the anchoveta) and is primarily used as feed for agriculture and aquaculture. Peru makes up around 30% of the world's fish meal exports. During our data period, 2009–16, around 95% of the country's total fish meal production was exported. The three largest buyers are China, Germany, and Japan, but many other countries also import Peruvian fish meal (see table A1; tables A1–A8, B1, B2, C1–C3, and D1 are available online).

Fish meal is produced in manufacturing plants located along the coast of Peru, of which there were 94 in 2009. These plants were in turn owned by 37 firms. The median number of plants per firm is two during our data period, while the 25th and 75th percentiles are one and five, respectively.

Miquel-Florensa (2016) convincingly show how supply assurance motives influence organizational structure in the Costa Rican coffee industry by relating measures of ex post reneging temptations to ex ante choice of structure (see also Banerjee and Duflo 2000; Macchiavello and Morjaria 2015). We follow Natividad (2014) in studying organizational structure in the Peruvian fish meal industry. He focuses on an earlier period when an unusual regulatory system—industry-wide fishing quotas—generated common pool incentives famously overshadowing other forms of supplier/plant incentives (see, e.g., Tveteras, Paredes, and Peña-Torres 2011), which lead to an "Olympic race" for fish.

There is heterogeneity in processing capacity, technology, and the share of production that is of high quality grade across both firms and individual plants in our sample. Firms differ considerably in their average number of export transactions per season and in the size and value of their shipments. As seen in figure A1 (figs. A1–A5 are available online), firm size correlates positively with average quality grade produced.

Plants receive inputs of raw fish from their suppliers. The suppliers are larger steel boats—which may be independent or owned by the firm that owns the plant—and smaller wooden boats. Regulations prohibit fish meal firms from owning wooden boats. There are on average 812 (wooden and steel) boats active in a given season and significant heterogeneity in boat characteristics, such as storage capacity, engine power, and average quantity caught per trip. Fishing trips last 21 hours (SD = 10 hours), and boats travel 76 km away from the port of delivery (SD = 46 km) on average. Changes in installed technology are observed in our data but are rare for both boats and plants. Table 1 shows summary statistics, providing further detail on the sector.

Since 2009, boats in Peru have operated under individual transferable quotas (ITQs), a common resource-management system used in fisheries and natural resource sectors worldwide. Individual boats are assigned a share of an industry-wide quota. We limit our analysis to the time period after ITQs were implemented to avoid any potential changes in quality production or integration driven by the quota system. Quota transfer is unlikely to explain firms' response to the time-varying and firm-specific quality upgrading incentives we focus on.

The long-term trend is toward more integration in the Peruvian fish meal industry, and the long-term trends in demand for quality and average output quality in Peru are also positive. These broad patterns are consistent with our hypothesis. However, it is higher-frequency variation around the long-term trends that we exploit to test our hypothesis. For example, we also observe deintegration during our data period—sales of boats from fish meal firms to independent co-ops or captains and from one fish meal firm to another.

B. Product Differentiation and Quality

An important feature of fish meal is that output quality is effectively captured by a single—measurable—dimension: protein content. Batches with protein content above a specified percentage are labeled prime quality, and plants report their production of fish meal of prime and fair average (below prime) quality to regulatory authorities each month. Price differentials across transactions for Peruvian fish meal of a given quality grade in a given time period are negligible, highlighting the horizontal homogeneity of the product.

TABLE 1
Summary Statistics

	Mean	Standard Deviation
Firms:		
Total number of firms in sample	37	
Export shipment (metric tons)	380	351
Export price ($/metric ton)	1,455	303
Number of destinations per season	7.09	5.35
Number of export transactions per season	87	100
Plants:		
Total number of plants in sample	94	
Has high technology	.85	.36
High-quality share of production	.85	.35
Monthly production (metric tons)	3,116	3,266
Processing capacity (metric tons/hour)	106	54
Boats:		
Number of boats operating per season	812	92
Fraction owned by a downstream firm per season	.28	.45
Fraction of boats made of steel per season	.44	.50
Storage capacity (m^3)	187	165
Power engine (hp)	432	343
Number of fishing trips per season	24.6	13.3
Number of delivery ports per season	3.49	1.90
Offload weight (metric tons) per trip	110	110
Time at sea per trip (hours)	20.85	9.96
Maximum distance from the plant's port (km)	76	46

Note.—This table gives summary statistics over our sample period. "Has high technology" is a dummy equal to one if the plant is equipped with steam-drying technology. Plants' "processing capacity" measures the total weight of fish that can be processed in 1 hour. "Steel" is a binary variable equal to one if a boat is a steel boat (which tends to be bigger and better suited for industrial fishing and is subject to different regulations). "Offload weight per trip" is the amount fished and delivered to a downstream firm on each trip. "Time at sea per trip" is the total time spent at sea on a fishing trip. "Maximum distance from the plant's port" is the maximum distance between the boat and the port it delivers to on any trip.

Fish meal's protein content crucially depends on input characteristics—namely, the freshness and integrity of the raw fish that boats deliver (see, e.g., FAO 1986). Freshness and integrity of the fish at the time of delivery in turn depend on choices made by the boat's captain before and during a trip, such as the amount of ice brought on board, the amount of fish packed on board, how tightly the fish are packed, and the time spent between a catch and delivery to a plant (FAO 1986). Because of the relationship between freshness and output quality, fish are processed as soon as possible after offload.

While it is easy to weigh and determine the quantity of fish a boat delivers, it is difficult to quantify or measure fish freshness directly. In theory, chemical tests of total volatile nitrogen content can be used to do so (imperfectly), but the managers in the industry we interviewed reported that such tests were too costly and time-consuming to be usable in Peru during our data period. In addition to the fixed cost of (the human and physical capital required for) adoption, this was because of high marginal cost of use and the value lost if fish were not processed immediately after offload. Footnote 6 discusses the extent to which input quality can be inferred from output quality after production.

After offload, the fish are weighed, cleaned, and converted to fish meal using one of two technologies: steam drying (hereafter, “high technology”) or exposing the fish directly to heat (hereafter, “low technology”). The technology used can matter for the protein content achieved.

Peru allows anchovy fishing for fish meal production during two seasons each year, and because of the need for fresh fish, fish meal plants operate only during the fishing seasons. There were thus 14 fishing and fish meal production seasons during our 2009–16 study period. In theory, fish meal can be stored for a short period of time, but we find that almost all is sold before the next production season begins, as shown in figure A2 and discussed below.

C. Organizational Structure

Consistent with our hypothesis, both integration and average output quality have slowly increased over time in the Peruvian fish meal industry. However, these long-term trends are not the source of the relationship between organizational structure and quality upgrading that we establish in this paper. This is because our empirical strategy exploits variation *around* the long-term trends for identification.

There is significant buying and selling of suppliers during our sample period: 741 steel boats (which can be vertically integrated) are registered during our data period. As seen in panel A of table 2, we observe 310 instances where ownership of a steel boat changes hands. In 103 of these instances, a fish meal firm acquires a supplier that is initially owned independently—that is, by a co-op or an individual captain. However, we also observe 32 instances where a supplier is sold from a fish meal firm to an independent buyer and 45 instances where a supplier is sold from one fish meal firm to another. On average, 28% of the boats that are active in a given season are integrated with a fish meal firm.

In our data, we observe not only supplier ownership but also deliveries from integrated and independent suppliers. We can therefore construct a measure of the vertical structure of firms’ production process—namely, the share of inputs coming from integrated suppliers (share VI). Peruvian fish meal manufacturers’ share VI is on average 43%. Firms can generally increase or decrease the maximum amount of inputs they can obtain from integrated suppliers only by buying or selling boats. The reason is that a boat’s total catch in a given season is governed by a regulatory quota, and each boat typically exhausts its quota. A firm may vary its share VI also by increasing or decreasing its use of the firm’s integrated suppliers or independent suppliers. As seen in figure A3, and following

TABLE 2
Summary Statistics on Integration

	(1)	(2)	(3)
	A. Boat Purchases and Sales		
Total number of steel boats registered		741	
Number of steel boat transactions		310	
Number of transactions, independent to VI		103	
Number of transactions, VI to independent		32	
Number of transactions, VI to VI		45	
Number of transactions, independent to independent		130	
	B. Variance Decomposition of Changes in Share of Inputs from VI Suppliers		
	VI Extensive Margin (Boat Purchases or Sales)	VI Intensive Margin (Intensity of Use of Own VI Suppliers)	Independent Margin (Buying Less/More from Independent Boats)
Contribution (%)	35	16	49

Note.—Panel A displays basic statistics on boat purchases and sales. Only steel boats can be owned by downstream firms. Panel B presents a variance decomposition of the three margins by which a firm can change its share of inputs from VI suppliers over time. "VI Extensive Margin" corresponds to the change in the variable because of the firm buying or selling boats. "VI Intensive Margin" corresponds to changes in the own use share of all the inputs produced by VI suppliers. VI suppliers sometimes supply other firms (around 10% of the time; see fig. 1), so a firm could decide to take in more or less of its VI suppliers' inputs. The last contribution is "Independent Margin," which corresponds to the firm buying more or less from independent boats than in the previous season, which mechanically impacts the share of inputs from VI suppliers. Appendix B provides the computational details of the decomposition and how the variance decomposition is done in practice.

the trend in ownership, share VI slowly increased during our data period. In panel B of table 2, we show a simple variance decomposition of changes in share VI (our key dependent variable in sec. V) to highlight the importance of the various margins by which firms may adjust share VI. We show that 35% of the overall variation in changes in share VI comes from firms buying and selling suppliers, 16% from the intensity with which firms make use of their integrated suppliers, and 49% from the intensity with which they make use of independent suppliers.

Importantly for our purposes, share VI can be defined not only for firms but also for individual plants within firms. A plant's share VI at a given point in time depends mostly on the organizational structure of the firm the plant belongs to, but there is significant variation across plants within the same firm. This variation depends both on the extent to which firm managers direct integrated suppliers to deliver to one plant over another and on the presence of independent suppliers near a given plant. The latter varies considerably over time and depends on variation in weather, fish density, and independent captains' decisions.

Figure 1 shows that integration and deintegration primarily represent a change in the formal status of the relationship between a firm/plant and a supplier engaged in frequent and continuing interactions. The figure displays the fraction of trips suppliers deliver to various firms and plants. The top two panels focus on all boats, while the bottom two panels restrict attention to the switchers we focus on in our empirical analysis of supplier behavior in section VI. These switchers—suppliers that get integrated or sold—deliver to the plant (within the acquiring/selling firm) they interact with most frequently around 41% of the time when independent (i.e., before getting acquired or after getting sold) and around 45% of the time when integrated. Similarly, switchers deliver to the acquiring/selling firm around 63% of the time when independent and around 81% of the time when integrated.[18]

D. Contracting and Supplier Incentives

There is no centralized spot market for fish purchases: plants are spread out along the coast, both because the fish move around and because of geography's influence on the location of ports. Similarly, the movements of boats are a complex function of fishing conditions, weather (winds, swell, etc.), and the captains' incentives. Because of the importance of fish freshness, independent captains typically begin contacting plants over the radio on their way to a port after fishing.

We interviewed fish meal industry associations, a major company's managing director, another major company's chief operating officer, union representatives, and others in the sector to gain a qualitative understanding of the characteristics of the contracts used and the incentives that suppliers face. The interviewees reported that there is variation in how the captains and crews of boats owned by fish meal firms are paid, with many firms using a primarily fixed wage system and some tying fishermen's pay to past weeks' average price of Peruvian fish meal. They also reported that captains and crews of integrated boats are generally paid in the same way, by the parent firm, regardless of whether they deliver to the parent firm or to other firms. Note that both the interviews we conducted and our supply

---

[18] In the top two figure panels, we see that, as a whole, integrated suppliers deliver to the firm they deliver to most often (i.e., the parent firm) about 90% of their trips and deliver to the plant they deliver to most often 38% of trips. Independent suppliers deliver to the firm they deliver to most often around 65% of trips and deliver to the plant they deliver to most often 45% of trips.

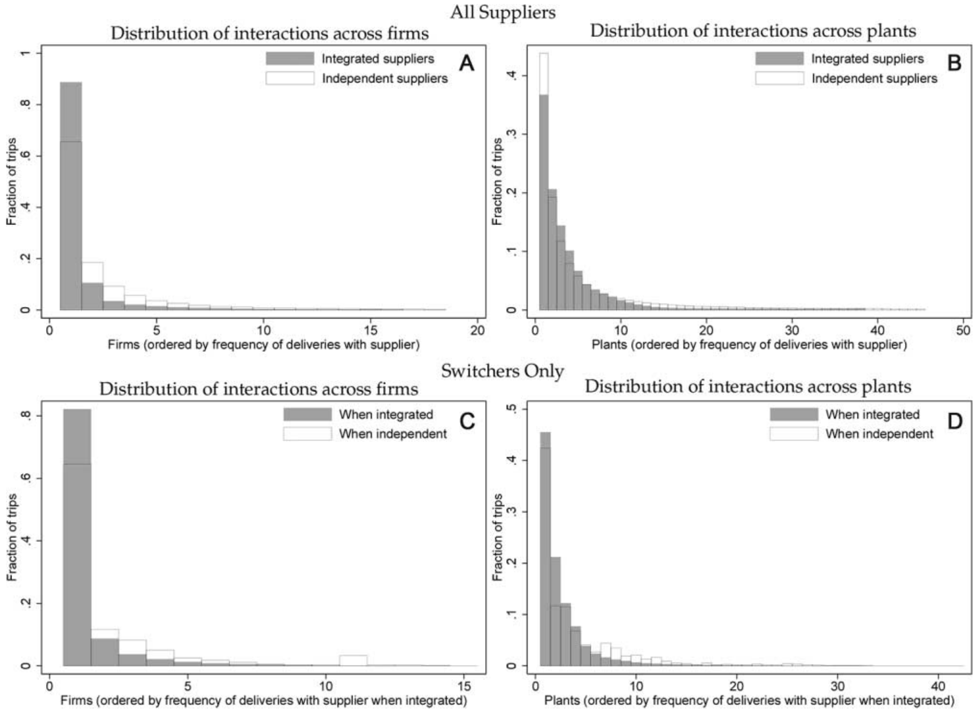

Fig. 1.—Interactions between manufacturers and integrated and independent suppliers. Panels show the average fraction of deliveries at each firm and plant for independent and integrated suppliers. *A*, *B*, Plants or firms are ordered based on the frequency of deliveries for each boat × boat owner pair: the plant or firm that receives the highest number of deliveries by the boat in question while owned by the owner in question is ranked one, the next-highest ranked two, and so on. *C*, *D*, Plants or firms are ranked based on the frequency of deliveries for each boat while it is integrated. Panels *A* and *B* include all suppliers, while panels *C* and *D* include only switchers—boats that were independent at one point and integrated at another point during our sample.

transactions data indicate that partial side selling by integrated suppliers is not an issue in the industry, primarily because of the significant time commitment necessary to dock and offload.

We are not aware of formal contracts between independent suppliers and firms over when, where, or what quality of fish to deliver. Interviewees reported that payments to independent suppliers—while agreed on case by case—are typically simply the quantity multiplied by a going price. We use internal data on payments to suppliers from a large firm to confirm this. These indicate that independent suppliers at a given point in time are paid a price per metric ton of fish delivered that is essentially fixed: port × date fixed effects explain 99% of the price variation across transactions.

Our data on suppliers' behavior—discussed in section IV—come from a map that the regulators update roughly every hour using the GPS signals all boats are required to transmit to authorities while fishing. Firms are allowed to access information on their integrated suppliers' whereabouts if they install the required technology but not the GPS data of independent suppliers or those owned by other firms. This is a primary reason why manufacturers and independent suppliers cannot contract over GPS-measured actions.

An industry association informed us that almost all contracts between fish meal firms and their foreign buyers for a given season's production are negotiated either before the season starts or early in the season. Firms can thus integrate or sell suppliers before a season starts in response to variation in the quality premium and/or high or low demand from particular importer countries.

III. Theoretical Framework

A. *Description*

In this section, we present a simple model to highlight how vertical integration may resolve the contracting issues facing downstream firms that aim to produce high-quality output. The intuition of the model is based on two insights. First, high-powered incentives to produce quantity can lead to actions that are wasteful and even harmful to quality. Second, the open market provides independent suppliers strong incentives to produce quantity, and in a setting where contracts are difficult to write, the only way to temper those incentives may be to integrate.

The first point of intuition above—the trade-off between quality and quantity—is one of the classic examples of the challenges of designing incentives in a multitask environment and in fact is used by Holmstrom and Milgrom (1991) to motivate their seminal work. This is for the simple reason that input quantity is typically straightforward to measure and reward, while quality is not. As a result, care must be taken not to overincentivize quantity to the detriment of quality.

Of course, the difficulty of determining quality is somewhat of a stereotype: there are goods for which quality depends on something like strength, size, or durability that is just as easy to measure as quantity. However, in our setting, this stereotype seems broadly accurate. While the quantity of fish that suppliers deliver is easily measured, the quality of that fish is difficult to ascertain for a purchasing manager examining several tons of anchoveta.

A few pieces of context are helpful for understanding the second point of intuition above. First, it appears that contracts are difficult to write ex ante: independent suppliers retain their right to deliver their catch where they choose. Additionally, while some firms primarily produce fish meal with high protein content, others primarily produce low quality grades and hence provide a (presumably less quality-sensitive) alternative for suppliers to deliver their catch.[19]

With this in mind, a logic applies that is familiar from the models presented in Baker, Gibbons, and Murphy (2001, 2002) based on the notion of integration as asset ownership that follows Grossman and Hart (1986). Even if a firm interested in sourcing high-quality inputs has no interest in high volumes, the fact that an independent supplier has the option to sell its inputs to an alternative downstream firm that values quantity creates powerful incentives. The independent supplier will then invest in producing quantity—although it may be wasteful or detrimental—if only to improve its bargaining position with the quality-focused firm. By acquiring the supplier, the manufacturer removes this outside option and hence any incentive for wasteful or harmful investment in quantity. In this sense, integration is valuable precisely because it mutes the power of market incentives, a notion that has been described by Williamson (1971), Holmstrom and Milgrom (1994), and Gibbons (2005a), among others.

B. Model Details

We consider a static game with two actors: suppliers and high-quality firms. Suppliers take costly actions to produce a good that is valuable both to the firms and in an alternative use. They may be integrated or independent. If the suppliers are integrated, the firms that own them have the right to the good after the actions are taken. If the suppliers are independent, they retain the right to the good. They bargain with the high-quality firms over whether to deliver the good or consign it to its alternative use.

We assume that suppliers have two potential actions $\{a_1, a_2\}$, with costs $c(a_1, a_2) = (1/2)a_1^2 + (1/2)a_2^2$. These actions impact the surplus created

---

[19] A question that our model abstracts from is why firms might want to produce different quality levels simultaneously. We return to this question at the end of sec. III.

by delivering their inputs to a downstream quality-focused firm. We denote this surplus by Q and refer to it as the quality surplus. Suppliers' actions also impact the surplus they receive by delivering the inputs to an alternative—quantity-focused—downstream firm. We denote this by P and refer to it as the quantity surplus. We assume that the good is specific, in the sense that $Q > P$. In particular, we define

$$P = a_1,$$

$$Q = Q_0 - \gamma a_1 + \delta a_2,$$

with $\gamma, \delta \geq 0$.[20] In this sense, a_1 is a quantity-focused action, while a_2 is a quality-focused action. While this is a simplified model, a_1 can be thought of along the lines of fishing for extended periods to catch the maximum amount, traveling long distances to find fish in high volumes, or packing the hold tightly with fish. On the other hand, a_2 can be thought of as carrying extra ice on board to keep the catch cool or taking care to ensure that the fish are not crushed. The baseline level of quality surplus is represented by Q_0.[21] Note also that a_1 enters negatively in Q to capture the notion that actions taken to increase the quantity caught, such as packing the hold tightly with fish, often adversely affect quality.

We assume that neither P nor Q is contractible but that P—the quantity surplus—is perfectly observable at the time of bargaining and Q—the quality surplus—is not. All parties know the value of Q_0, and because $P = a_1$ is observable, Q in effect has an observable portion: $\tilde{Q} = Q_0 - \gamma a_1 = Q - \delta a_2$.

1. Integrated Suppliers

If a supplier is integrated, the firm has rights to the supplier's catch. However, because the firm cannot write contracts over Q and P, it cannot credibly commit to rewarding the supplier's actions. As a result, the supplier chooses $a_1 = 0$ and $a_2 = 0$, and the total surplus is simply Q_0.

2. Independent Suppliers

Although neither Q nor P is contractible,[22] the firm and supplier may bargain ex post over the price of the delivery. We assume a Nash bargaining concept, with the supplier's bargaining coefficient equal to α.

---

[20] More specifically, we assume that $0 \leq \delta \leq 1$ and $0 \leq \gamma \leq 1 - \alpha$. Also, note that P could itself be the result of a bargaining process between the boat and a quantity-focused firm.

[21] This can be thought of as the amount that suppliers will catch before exerting any costly action or, perhaps more reasonably, as the result of some limited contractual agreement that we abstract from.

[22] Alternatively, we could assume that only a portion of Q and P is noncontractible and that we consider only this portion, as in Baker, Gibbons, and Murphy (2002).

Because the supplier can always deliver its catch to the alternative quantity-focused firm and receive P, the supplier must always receive at least P. The supplier additionally receives a share α of the observable portion of the surplus $\tilde{Q} - P$ that accrues to the firm: $\alpha(Q_0 - \gamma P - P)$. As a result, an independent supplier solves the problem:

$$\max_{a_1, a_2} \alpha Q_0 + (1 - \alpha\gamma - \alpha)a_1 - \frac{1}{2}a_1^2 - \frac{1}{2}a_2^2.$$

This gives $a_1 = (1 - \alpha\gamma - \alpha)$, $a_2 = 0$, and social surplus is

$$Q_0 - \gamma(1 - \alpha\gamma - \alpha) - \frac{1}{2}(1 - \alpha\gamma - \alpha)^2 < Q_0.$$

Because of the counterproductive actions to increase quantity ($a_1 > 0$) and the adverse effects of those actions on the quality surplus, the surplus is lower when the suppliers are independent. As a result, the more efficient organizational structure to produce quality is vertical integration.

It is worth noting that a number of assumptions made in this model are not strictly necessary to get this result. The relative efficiency of integration holds whether or not quantity-focused actions directly negatively impact the quality surplus (because of the inefficiency of quality actions) and would hold even more strongly if, for example, there were complementarities in the costs of quality and quantity actions.

C. Discussion

The theoretical role of vertical integration is a contentious topic. Our framing follows Baker, Gibbons, and Murphy (2001, 2002) in combining elements of the incentive-based theories in the tradition of Holmstrom and Milgrom (1991) and the property rights theories in the vein of Grossman and Hart (1986). Such a framing is not the only type of model that would produce a relationship between integration and output quality. In actuality, integration is a complex organizational change whose causes and consequences operate through multiple mechanisms. However, because the foundations of the model above depend on a series of salient features of our context—unobservable quality, observable quantity, and alternative buyers that are less concerned with quality—and because we are able to directly test the predictions of the model, we see these alternative theories as complementary to the mechanisms that our framework focuses on, rather than contradictory.

Our model presents a highly stylized and somewhat stark example to highlight a key intuition: that integration can act as a valuable tool for muting the incentives provided in the open market. We believe this starkness most simply portrays why firms in our context might want to integrate to produce high-quality output. That said, this oversimplification does have a few drawbacks, most notably the lack of incentive to take quality-focused actions and to take any actions at all when integrated. This is in some sense a strong version of what are sometimes called the drone employees (Gibbons 2005a) that appear in property rights theories of the firm that follow Grossman and Hart (1986). However, this feature may be easily remedied in more complex models that preserve the basic intuition and result. For example, assuming observability over Q induces quality-focused actions among independent suppliers and—for sufficiently small values of δ—does not affect the main result. Perhaps more realistically, introducing dynamics into the model, with long-term relationships between firms and suppliers, creates an environment in which the incentives of the downstream and the upstream parties can be aligned through repeated interactions.

In appendix C, we present and test the empirical implications of exactly such a dynamic model, in which we allow the downstream party to use relational contracts to incentivize the quality action. We posit that Q—the quality surplus—can be observed to the downstream party but only with some lag (e.g., once the inputs are processed and output quality is measurable). The firm can then offer the supplier a (delayed) reward contingent on this surplus but can credibly promise to pay this reward only if it repeatedly interacts with the upstream party. In this context, we show that the value of the relationship can incentivize the supplier to take the first best actions but that this sort of relational contract may be difficult to sustain if the supplier is independent. The intuition for this result is similar to our static baseline: independent suppliers own the rights over the inputs, and when the value of these inputs in their alternative use is high, they face incentives to renege on the relational contract and sell the goods in their alternative use.

Our model above also implicitly demonstrates the costs of integration. The market provides strong incentives for quantity, and for a low-quality firm that is aiming to produce quantity, integration would only interfere with and lessen the strength of these incentives. Accordingly, quantity-focused firms prefer independent suppliers. A similarly formulated model, with the roles of high- and low-quality firms switched (e.g., $P \gg Q$), provides precisely this result.

In our stylized model, firms are either quality oriented or not. In reality, a firm's output objectives are likely a combination of quality surplus and quantity surplus in which the weight attached to each depends on the demand the firm faces at a particular point in time. In this case, firms should not source all inputs from either integrated or nonintegrated suppliers but choose an intermediate organizational structure—that is, an intermediate level of integration—that depends on the relative importance of Q and P in the firm's current objective function.

The framework presented in this section motivates three empirical predictions that we test in the remainder of the paper:

1) Firms' organizational structure responds to variation in the relative profitability of producing high-quality output. An increase in the quality premium leads to more integration—that is, a higher share of inputs sourced from integrated suppliers.
2) The reason is that the actions of a supplier differ when the supplier is integrated. In particular, suppliers that get integrated reduce their effort to produce quantity, especially in ways that benefit quality.
3) As a result, the degree to which a firm or plant uses integrated suppliers affects output quality. Firms that use inputs from integrated suppliers produce higher-quality output.

IV. Data, Variables, and the Relationship of Interest

A. Data

The primary data sets that we use to test our three predictions are as follows.

Plant production.—Administrative data on all plants' production come from Peru's Ministry of Production, which regulates the fish meal industry. Every month, plants are required to submit information on how much prime (high-quality) and fair average (low-quality) fish meal they produce. Quality grade is thus directly reported in the plant production data and subject to auditing by government inspectors. As discussed in section II.B, the distinction between prime and fair average quality is on the basis of the fish meal's protein content. From these records, we construct each individual plant's and each firm's "high-quality share of production" in a given month or production season.

Plant registry.—We link the production data with an administrative plant registry that contains monthly information on each plant's (i) technological production capacity and (ii) owner, typically a multiplant fish meal firm.[23] We also use this registry to link the production data to export data. We can do so for almost all firms but not for the smallest firms, which use intermediaries to export.

---

[23] The data contain information on the number of metric tons that can be produced per hour with currently installed low and high technology. As very few firms in our sample have only the low technology, we define a high-technology firm as one for which the high-technology share of total processing capacity is higher than the median (0.67).

Export transactions.—Detailed data on the universe of fish meal exports at the transaction level come from Peru's customs authority. We observe the date of the transaction, the export port, the destination country, the weight of the fish meal, the value of the transaction, and the exporting firm.

Internal data from a large firm.—One of the largest fish meal firms in Peru shared its internal sales records with us. The firm owns many plants along the coast. The sales records include information on the shipment's packing, its free-on-board value, the price per metric ton, the buyer, the destination country, the date of the contract, and the terms. Most importantly for our purposes, the specific plant that produced a given shipment of fish meal is reported.

Supply transactions.—The Ministry of Production records all transactions between the fish meal plants and their suppliers of raw materials—that is, fishing boats. Information on the date of the transaction, the boat, the plant, and the amount of fish involved (though not the price) is included.

Boat registry.—We merge the supply transactions data with an administrative boat registry that provides information on a boat's owner, the material the boat is made of, its storage capacity and engine power, and whether it has a cooling system installed.[24]

Boat GPS data.—Peruvian fishing boats that supply fish meal plants are required to have a GPS tracking system installed and to continuously transmit their GPS signal to the Ministry of Production while at sea. The ministry stores the transmitted information—the boat's ID, latitude, longitude, speed, and direction—each hour on average and shared the resulting data set with us. Some observations are missing in the GPS data set (and therefore also in our measures of boat behavior), but the patterns of missingness appear arbitrary. Ninety-five percent of all boats and 99% of steel boats (which can be vertically integrated) that appear in the transactions data also appear in the GPS data.[25]

We merge many data sets and use a wide range of right- and left-hand-side variables across our empirical analysis. For this reason, the number of individual firms, plants, suppliers, and time periods available for each regression varies considerably. In the notes to our regression tables, we include information on the number of distinct i's and t's and, where

---

[24] Information on engine power is available only for 2004–6. However, changes in engine power are extremely rare in that period, so we treat this characteristic as fixed over time.

[25] The missingness seems to be driven by the regulator's data storage and maintenance and/or faulty equipment. In particular, there are no GPS data after June 2015, and before then, some boat owners disappear from the data for a complete calendar year. Such missingness in the GPS data is unlikely to drive our results since we focus on within-boat changes in behavior upon integration/deintegration.

relevant, an explanation of why some cannot be used in the relevant analysis.

B. *Variables of Interest*

Our primary measure of an individual plant's output quality is the share of the fish meal the plant produces in a given month that is of prime quality grades—a direct measure of quality whose interpretation requires no assumptions. We aggregate this measure up to firm or firm × season level to construct a corresponding measure of a firm's high-quality share of production.

We also construct a granular measure of the average quality grade—protein content—of the fish meal a firm produces. While we do not directly observe the exact protein content of each export shipment, we can go beyond simply using unit prices and approximate the precise quality grade. This is because we observe quality grade–specific fish meal prices in detailed (week × export port × protein content level) data recorded by a fish meal consulting company. We infer the protein content of each of a firm's export shipments by comparing the corresponding unit values with this price data. To construct a firm × season-level measure, we average protein content across export shipments, weighting by quantity.[26] A priori, we have little reason to believe that this inferred protein content measure should be systematically biased.[27] Empirically, it is highly correlated with the high-quality share of production directly observed for a firm's plants in production data and with the exact quality grade reported in the sales

---

[26] The export transaction records do not report the specific plant that made the fish meal, so the inferred quality grade is available only at the firm level—except for data covering the fish meal firm that shared internal data with us, including information on the plant that produced a given export shipment. One potential concern is that fish meal can be stored for a short period and hence firms could attempt to strategically time their export transactions. In practice, the product is almost always sold before the next production season starts. (The reason why inventories are small—between +10% and −10% of total season production [see fig. A2]—is likely that many contracts are entered into before the production season starts [which helps the fish meal manufacturers and their foreign buyers reduce demand/supply uncertainty] and because firms' ability to strategically "time" their sales is in actuality limited.) A shipment can thus be traced back to a specific production season (but not a specific production month; constructing the inferred protein content measure at the month level would require an assumption about how firms manage their inventories—e.g., first in, first out vs. first in, last out). A related concern is that firms that are about to end operations and close down might sell off their fish meal, in which case a lower unit price might not reflect lower quality but rather a "going out of business" discount. We thus exclude data from any firm × season observations that correspond to a firm's last season producing and exporting fish meal, but the results are robust to including these observations. These issues are not relevant for our directly observed high-quality share of production measure of output quality.

[27] Fish meal is a vertically differentiated but otherwise homogenous product, and price differentials across shipments of a given quality level (and across firms producing a given quality level) in a given time period are negligible (see sec. II.B). This implies that pricing to market, bulk discounts, etc., are not a concern.

records of a firm that shared its data with us. Nonetheless, we focus primarily on our directly observed measure of the high-quality share of production.

To quantify vertical integration, we consider both the number of suppliers a firm owns and the corresponding share of inputs used in its production process that come from integrated suppliers (share VI). Share VI is our preferred measure of integration for a number of reasons. First, because we observe all transactions between plants and suppliers, we can construct share VI in a consistent manner for both firms and individual plants. This allows us to move from the across-firm comparisons we make in section V to the within-firm comparisons we make in section VII. Second, we view share VI as the more relevant measure when asking whether organizational structure and output quality are causally related: if firms vertically integrate when the quality premium rises because doing so allows them to increase input quality, then it should matter not only whether a firm owns suppliers but the degree to which the firm as a whole and its individual plants actually source inputs from those suppliers at the time of production. Third and finally, share VI automatically captures suppliers' size, allowing us to avoid assumptions on "scale effects"—for example, how the benefit of one large integrated supplier compares with two small ones.

We also show the margins of adjustment driving changes in a firm's share VI. In addition to the decomposition shown in panel B of table 2, we show in section V and appendix B that firms' share VI response to variation in the quality premium is driven in large part by the extensive margin—acquiring and using new suppliers.

C. *Relationship of Interest*

In section V, we begin our analysis of how exogenous changes in incentives to quality upgrade affect integration decisions. Before doing so, we first demonstrate that the basic relationship predicted by our model holds empirically: integration and output quality are positively correlated. To do so, we estimate regressions of the form

$$\text{Quality}_{it} = \alpha + \beta_1 \text{VI}_{it} + \beta_2 \text{HighTech}_{it} + \gamma_i + \delta_t + \varepsilon_{it}, \qquad (1)$$

where Quality_{it} and VI_{it} respectively measure the quality of the output produced by firm i in season t and how vertically integrated the firm's organizational structure is in the same season. We control for the technology that the firm uses to convert fish into fish meal,[28] HighTech_{it}, and firm

---

[28] A firm's production technology is an important potential determinant of output quality and one that could plausibly correlate with organizational structure (Acemoglu et al. 2007, 2010). We thus control for installed HighTech_{it}, i.e., steam-drying (high) technology. At the firm level, HighTech_{it} is equal to the share of installed capacity that is of the high type.

and season fixed effects γ_i and δ_t. We thus estimate changes in output quality for those firms that vertically integrate in a given season relative to other firms that do not. We cluster the standard errors at the firm level.

The results in panel A of table 3 point toward a meaningful baseline relationship between owning suppliers and output quality. The estimate in column 1 implies, for example, that moving from the 25th to the 75th percentile of number of boats owned is associated with an increase in the high-quality share of production of around 17%, although this estimate is not statistically significant. The estimate for our secondary measure of output quality, shown in column 4, implies that moving from the 25th to the 75th percentile of number of boats owned is associated with an increase in protein content of just under 10% of the range observed in Peru. This estimate is significant at the 1% level.[29]

In panel B, we show that, beyond simply owning suppliers, what matters for output quality is share VI: the share of a firm's supplies coming from integrated suppliers at the time of production. The results imply that a firm that uses inputs coming entirely from integrated suppliers rather than inputs entirely from independent suppliers sees a share of high-quality output that is roughly 38%–41% (28–31 percentage points) higher and an average protein content that is higher by roughly 23% of the range observed in Peru.

As shown in table A2, the characteristics of integrated suppliers unsurprisingly differ from the characteristics of independent suppliers: integrated suppliers are on average larger and have more advanced equipment than independent suppliers, while the average switcher falls in between the average always-independent boat and the average always-integrated boat (but closer to the latter). However, in panel C we show that the results in panel B are not driven by observable, time-varying supplier or firm characteristics. We control for a series of supplier characteristics, as well as the firm's share of total industry production and the size of its production. Doing so has little impact on the estimated coefficient.[30]

---

[29] Firms in the 25th percentile own four boats, while firms in the 75th percentile own 36 boats, and the range of protein content observed in Peru is approximately 63%–68%. The Ministry of Production has not responded to our inquiries about the exact percentage cutoff they use to define prime (high-quality) fish meal (but our estimates do not make use of such a cutoff).

[30] Controlling for the share of inputs coming from steel boats, high-capacity boats, and boats with a cooling system leaves the magnitude and significance of the coefficient on the share of inputs coming from VI suppliers essentially unchanged. Note that two of the supplier characteristic variables included—share of inputs from high-capacity boats and share of inputs from boats with a cooling system—are significantly correlated with output quality in the cross section of firms. One reason why the coefficients on these characteristics are not significant is that we observe little change in these boat characteristics over time. Controlling for the firm's share of total industry production and log(production) also leaves the magnitude and significance of the coefficient on share VI essentially unchanged.

TABLE 3
OUTPUT QUALITY AND VERTICALLY INTEGRATED SUPPLIERS

	DEPENDENT VARIABLE: HIGH-QUALITY SHARE OF PRODUCTION			DEPENDENT VARIABLE: PROTEIN CONTENT		
	(1)	(2)	(3)	(4)	(5)	(6)
	A. Output Quality and Number of Suppliers Owned					
asinh(number of suppliers owned)	.057	.045		.227***	.171**	
	(.060)	(.041)		(.072)	(.080)	
High technology share of capacity	No	Yes		No	Yes	
Season fixed effects	Yes	Yes		Yes	Yes	
Firm fixed effects	Yes	Yes		Yes	Yes	
Mean of dependent variable	.75	.75		65.7	65.7	
Observations	275	275		208	208	
R^2	.75	.81		.79	.80	
	B. Output Quality and Share of Inputs from VI Suppliers					
Share of inputs from VI suppliers	.311**	.284**		1.157***	1.153***	
	(.124)	(.118)		(.276)	(.285)	
High technology share of capacity	No	Yes		No	Yes	
Season fixed effects	Yes	Yes		Yes	Yes	
Firm fixed effects	Yes	Yes		Yes	Yes	
Mean of dependent variable	.75	.75		65.7	65.7	
Observations	275	275		208	208	
R^2	.76	.81		.79	.80	
	C. Output Quality and Share of Inputs from VI Suppliers					
Share of inputs from VI suppliers	.269**	.302**	.288**	1.185***	1.146***	1.101***
	(.113)	(.121)	(.116)	(.350)	(.281)	(.289)
Share of inputs from steel boats	−.161		−.155	−.702		−.593
	(.205)		(.195)	(.962)		(.797)
Share of inputs from boats with high capacity	.254		.242	.847		.983
	(.218)		(.213)	(1.315)		(1.129)
Share of inputs from boats with cooling system	.090		.100	−.176		−.097
	(.120)		(.120)	(.845)		(.706)
Share of industry's production		−.405	−.324		−1.124	−1.032
		(1.069)	(1.024)		(3.663)	(3.684)
Log(production)		−.023	−.023		.322	.345
		(.032)	(.031)		(.301)	(.306)

TABLE 3 (*Continued*)

	C. Output Quality and Share of Inputs from VI Suppliers					
High technology share of capacity	Yes	Yes	Yes	Yes	Yes	Yes
Season fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Mean of dependent variable	.75	.75	.75	65.7	65.7	65.7
Observations	275	275	275	208	208	208
R^2	.82	.82	.82	.80	.81	.81

NOTE.—One observation is a firm during a production season. "High-Quality Share of Production" is the share of a firm's total production during a fishing season that is reported as high-quality (prime) output. This dependent variable is available for 34 firms over 14 fishing seasons from April 2009 to November 2016. "Protein Content" is the quantity-weighted average of a measure of quality inferred from a database that provides weekly prices by quality. This dependent variable is available for 23 firms over 15 fishing seasons from April 2009 to April 2017. We observe one more fishing season for this dependent variable than for "High-Quality Share of Production." The number of individual firms is smaller for the second dependent variable, as we do not observe export transactions for 11 firms. "Share of inputs from VI suppliers" is the share of a firm's inputs that come from VI suppliers during a season. Steel boats tend to be bigger and better suited for industrial fishing and are subject to different regulations. High-capacity boats are boats whose hold capacity is in the upper quartile of the distribution. Boats without an integrated cooling system must use ice to keep fish fresh. "Share of industry's production" is a firm's total production of fish meal during a given season divided by the industry's total production on that season. "Log(production)" is the logarithm of a firm's total production during a given season. "High technology share of capacity" controls for the share of the firm's total processing capacity (measured in metric tons per hour and averaged across all active plants within the firm) that uses steam-drying technology. Standard errors clustered at the firm level are included in parentheses.

** $p < .05$.

*** $p < .01$.

In appendix C, we consider whether the relationship between output quality and integration might be the result of long-term supplier-firm relationships rather than ownership per se. This does not appear to be the case, as we do not observe the association between quality and the share of inputs coming from suppliers in long-term relationships that we do for share VI. In other words, it is integration itself, not the relationship, that covaries with output quality. This is in line with the predictions of a dynamic version of our model, also shown in appendix C.

The relationship between a firm's organizational structure and its output quality that we established in this subsection is the starting point of our empirical analysis. This basic relationship is consistent with this paper's hypothesis. However, it is also consistent with the alternative theories discussed in the introduction. In these theories, a correlation between integration and output quality arises but the relationship is either not causal or not known to (or ignored by) firms. We show evidence that is hard to reconcile with such explanations in the next section.

V. The Quality Premium and Organizational Structure

We now show that the relationship between output quality and vertical integration we established in the previous section reflects an explicit organizational choice that firms make to climb the quality ladder. Specifically, Peruvian fish meal manufacturers integrate suppliers when the returns they earn from upgrading quality rise for exogenous reasons. This finding provides empirical support for the prediction that a vertically integrated organizational structure is efficient for producing high-quality output. We provide additional evidence that underscores the difficulty of reconciling this finding with alternative theories.

A. *Estimating How the Quality Price Premium Affects Vertical Integration*

We begin by showing that firms as a whole *do* vertically integrate when the quality premium is high. We quantify Integration_{it}—a firm's decision to integrate (deintegrate)—as a season-to-season increase (decrease) in the share of inputs the firm obtains from integrated suppliers (and show estimates for the level corresponding to this difference in appendix A, with qualitatively similar but less precisely estimated results).[31] We measure the quality premium as the difference between the (log) price of high- and low-quality fish meal.

We first show results from simple descriptive regressions of the form

$$\text{Integration}_{it} = \alpha + \beta\,\text{QualityPremium}_t + \eta_i + \varepsilon_{it}. \tag{2}$$

Here $\beta > 0$ indicates that firms vertically integrate when the quality premium is high and η_i represents a firm fixed effect.

Column 1 of table 4 shows the results from this specification, omitting firm fixed effects. There is a positive baseline correlation between the quality premium and firm integration. Column 2 replaces QualityPremium_t with two dummy variables: an indicator equal to one when the quality premium is above the mean and an indicator equal to one when the quality premium is below the mean (as these are exhaustive categories, the constant is suppressed). We see a positive and significant coefficient on the high-quality-premium indicator, suggesting that firms vertically integrate when

---

[31] Specifically, if VI_{it} is defined as share VI, we analyze $\text{VI}_{it} - \text{VI}_{it-1}$, where t indicates a season. In cols. 3–10 of table A3, we show all of our results using VI_{it} rather than the difference. Furthermore, in table A4 we show that our results are robust to a logit transformation of VI_{it} and VI_{it-1}—i.e., $\log(\text{VI}_{it}/(1 - \text{VI}_{it}))$.

the quality premium is high. We also see a negative coefficient on the low-quality-premium indicator, consistent with firms deintegrating when the quality premium is low, although the estimate is smaller and not statistically significant. Column 3 repeats the specification from equation (2), now including firm fixed effects. The estimated coefficient is extremely similar to that in column 1. In sum, a relationship between the quality premium and vertical integration is thus evident in the time series.

In column 4, we show that firms' response to the quality premium is not due to associated income shocks or general incentives to expand or reduce production. We repeat the regression from column 1 but with the log(average price)—the (equal-weighted) average price of high- and low-quality fish meal in season t—replacing the quality premium as the regressor of interest. The estimated coefficient is near zero and insignificant, indicating that firms are not more likely to vertically integrate when the overall price level is high. Figure 2 shows that the quality premium is only weakly correlated with average prices in Peru. It is thus not surprising that firms respond differently to the two.

B. Differential Responses by Firms with More versus Less Capacity to Upgrade

We next consider how the quality premium differentially impacts firms' integration decisions depending on their capacity to upgrade quality. Firms that are already producing exclusively high-quality output have no scope to upgrade the average quality they produce. For these firms, an increase in the quality premium cannot induce an increase in the high-quality output share and therefore should not lead to a change in organizational structure. Conversely, a firm that produces some mix of high- and low-quality output has the capacity to raise quality. If vertical integration indeed enhances output quality, we expect to see these firms integrate when the quality premium rises.

To investigate this differential response, we interact QualityPremium_t with measures of a firm's capacity to upgrade in season t. Our primary measure is a firm's "upgradable share of production"—that is, the share of low quality in its season $t - 1$ production. For example, a firm currently producing 25% high-quality output has a 75% upgradable share of production. As a secondary measure, we also consider a binary indicator, "low-quality producer," which is equal to one for firms that produced at least some low-quality output in season $t - 1$. We run specifications of the form

$$\begin{aligned}\text{Integration}_{it} = \alpha &+ \beta(\text{UpgradingCapacity}_{it-1} \times \text{QualityPremium}_t) \\ &+ \gamma \text{UpgradingCapacity}_{it-1} + \eta_i + \delta_t + \varepsilon_{it},\end{aligned} \tag{3}$$

TABLE 4
Vertical Integration and the Quality Price Premium

	Dependent Variable: Share of Inputs from VI Suppliers (t) − Share of Inputs from VI Suppliers $(t-1)$					
	A. OLS					
	(1)	(2)	(3)	(4)	(5)	(6)
Quality premium	.147* (.082)		.158* (.089)			
Quality premium is high		.014** (.006)				
Quality premium is low		−.007 (.010)				
Log(average price)				.057 (.043)		
Upgradable share of production $(t-1)$ × quality premium					2.709*** (.397)	
Low-quality producer $(t-1)$ × quality premium						.920** (.425)
Upgradable share of production $(t-1)$	No	No	No	No	Yes	No
Low-quality producer $(t-1)$	No	No	No	No	No	Yes
Season fixed effects	No	No	No	No	Yes	Yes
Firm fixed effects	No	No	Yes	Yes	Yes	Yes
Mean of dependent variable	.00	.00	.00	.00	.00	.00
Observations	191	191	191	191	191	191
R^2	.01	.02	.08	.07	.24	.21
Quality premium is high = quality premium is low (p)		.09				

	B. IV			
	(7)	(8)	(9)	(10)
Upgradable share of production $(t-1)$ × quality premium	2.536*** (.569)		2.712*** (.398)	
Low-quality producer $(t-1)$ × quality premium		.621 (.498)		.978** (.418)
Upgradable share of production $(t-1)$	Yes	No	Yes	No
Low-quality producer $(t-1)$	No	Yes	No	Yes
Season fixed effects	Yes	Yes	Yes	Yes
Firm fixed effects	Yes	Yes	Yes	Yes
Mean of dependent variable	.00	.00	.00	.00
Observations	191	191	191	191
R^2	.24	.20	.24	.21
First stage:				
Upgradable share of production $(t-1)$ × log(quota Chile) $(t-1)$	−1.153** (.457)			
Upgradable share of production $(t-1)$ × log(quota Denmark) $(t-1)$	−6.111** (2.775)			
Upgradable share of production $(t-1)$ × log(quota Iceland) $(t-1)$	−1.542** (.700)			
Low-quality producer $(t-1)$ × log(quota Chile) $(t-1)$		−.923*** (.196)		

TABLE 4 (*Continued*)

	B. IV			
	(7)	(8)	(9)	(10)
Low-quality producer $(t-1)$ × log(quota Denmark) $(t-1)$		−4.790***		
		(1.257)		
Low-quality producer $(t-1)$ × log(quota Iceland) $(t-1)$		−1.223***		
		(.317)		
Upgradable share of production $(t-1)$ × log(Chile price)			1.081***	
			(.135)	
Upgradable share of production $(t-1)$ × log(Denmark price)			−2.950***	
			(.740)	
Upgradable share of production $(t-1)$ × log(Iceland price)			1.469**	
			(.558)	
Low-quality producer $(t-1)$ × log(Chile price)				1.175***
				(.044)
Low-quality producer $(t-1)$ × log(Denmark price)				−2.158***
				(.534)
Low-quality producer $(t-1)$ × log(Iceland price)				.667
				(.486)
Kleibergen-Paap Lagrange multiplier *p*-value (underidentification)	.01	.07	.01	.06
Kleibergen-Paap Wald *F*-statistic (weak instruments)	37.6	22.1	30.4	471.2

NOTE.—One observation is a firm during a production season. “Share of Inputs from VI Suppliers (t) − Share of Inputs from VI Suppliers $(t-1)$” is the change between season $t-1$ and season t of the share of inputs sourced from integrated suppliers. Our sample includes 24 unique firms over 12 unique fishing seasons from November 2010 to November 2016. The number of seasons is smaller than in table 3, as the price data are available only from 2010 onward. The number of unique firms is also smaller, as several small firms die in 2009 and 2010. As shown in table 2, most of the variation in share VI is driven by acquisition or sales of suppliers. “Quality premium” is equal to log(high quality) − log(low quality), where high quality and low quality are the average price of prime and fair average fish meal, respectively, in the month preceding the current fishing season. We choose to take the month preceding the fishing season rather than the fishing season itself, as integration decisions are typically decided in the month preceding the season and integration within a season is extremely rare in the data. “Quality premium is high” (“Quality premium is low”) is equal to one if the quality premium is above (below) the sample average value. “Log(average price)” is the log of the average price of Peruvian fish meal, again computed in the month preceding the current fishing season. “Low-quality producer $(t-1)$” is equal to one if a firm’s share of low-quality output in the preceding season was at least 1%. “Upgradable share of production $(t-1)$” is the share of a firm’s production that was of low quality in the previous season. A firm that produces almost exclusively low-quality output has more potential to upgrade than a firm already producing almost exclusively high-quality output. In cols. 7 and 8, the instruments are interactions between “Low-quality producer $(t-1)$” or “Upgradable share of production $(t-1)$” and the fishing quota of top high-quality fish meal exporters. Standard errors clustered at the firm level are included in parentheses.

* $p < .1$.

** $p < .05$.

*** $p < .01$.

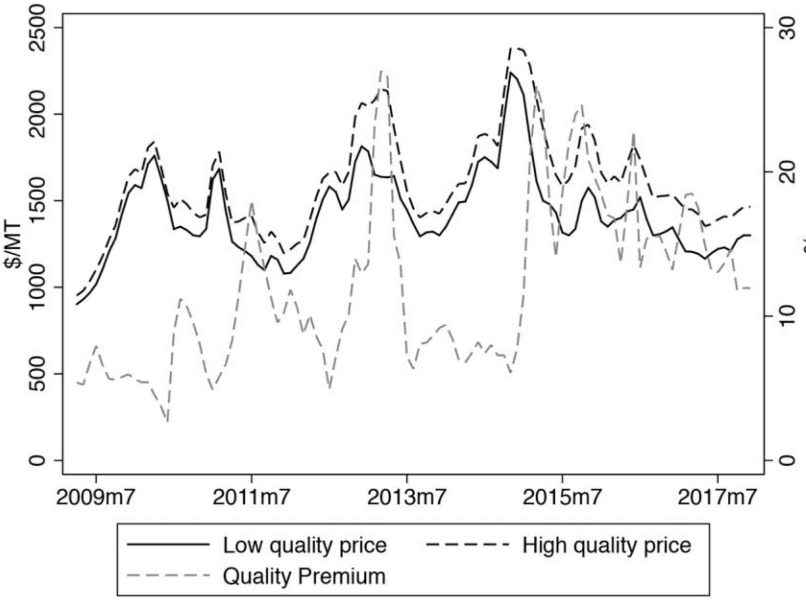

FIG. 2.—Average fish meal price and quality premium in Peru. This figure shows the evolution over time of the fish meal prices in Peru (the price of prime and fair average fish meal grades) and the quality premium in Peru. "Quality Premium" is equal to log(high quality) − log(low quality), where high quality and low quality are the average prices of prime and fair average fish meal grades, respectively.

where $\text{UpgradingCapacity}_{it-1}$ refers to either the upgradable share of production or a low-quality producer. This regression is a generalized difference in differences in which firms that are more versus less exposed to changes in quality upgrading incentives are compared in each production season, and in each of these seasons the quality premium may be relatively high or relatively low. We control for $\text{UpgradingCapacity}_{it-1}$ itself, firm and production season fixed effects (γ_i and δ_t), and cluster the standard errors at the firm level.

Before showing the results from estimating (3), we briefly discuss the variation we exploit. The season-to-season variation in the quality premium is shown in figure 2. While the long-term trend is weakly positive during our data period, the quality premium fluctuates substantially from season to season, sometimes rising and other times falling. Additionally, firms' upgrading capacity evolves over time.[32] For example, 19 of the 37 firms in our

---

[32] The fact that we control for the characteristic that defines a firm's exposure to a "treatment" variable that varies across time—here $\text{UpgradingCapacity}_{it-1}$—distinguishes our approach from traditional Bartik instrument approaches (see, e.g., Goldsmith-Pinkham, Sorkin, and Swift 2019).

sample are characterized as both a low-quality producer and a not-low-quality producer at some point in our sample. Of the remainder, about half are always low-quality producers and about half are never low-quality producers. Overall, 42% of all firm-season observations in our data are classified as low quality. A strength of our approach is thus that the firms with high and low $\text{UpgradingCapacity}_{it-1}$ being compared across seasons vary. Furthermore, the set of low-quality producers is not limited to a small number of firms or to small firms.[33]

Column 5 of table 4 shows OLS results from estimating (3) when using our primary measure of $\text{UpgradingCapacity}_{it-1}$, which is the upgradable share of production. We find that firms with greater scope to shift from low- to high-quality production are more likely to vertically integrate when the quality premium is high—consistent with our hypothesis and the model in section III. The same is true in column 6, where we define $\text{UpgradingCapacity}_{it-1}$ as a low-quality producer, our secondary measure. It is worth noting that if—not implausibly—firms producing a high share of high-quality output face a lower marginal cost of either quality upgrading or vertical integration, then such a countervailing force would, if anything, strengthen the support for our hypothesis implied by the results in columns 5 and 6 of table 4.[34]

A potential concern is that some omitted factor might influence both Peruvian firms' incentive to integrate and the quality premium (or alternatively, that the quality premium is itself influenced by integration decisions). To address this concern, we develop an IV strategy based on the regulatory fishing quotas of other top fish meal–producing countries. These aggregate fishing quotas, which directly constrain each country's fish meal production, are ideal instruments because they are set on the basis of sustainability considerations. Other countries' fishing quotas are thus unlikely to correlate with factors influencing Peruvian firms' integration decisions, except via their influence on market prices. Furthermore, because

---

[33] The mean number of plants active in firm-season observations that are classified as low quality is 2.8 (vs. 3.7 for not-low-quality firm-season observations). The 25th and 75th percentiles of number of plants are one and four for low-quality producers (vs. one and six for not-low-quality firm-season observations). Approximately 55% of firm × season observations have upgrading capacity equal to zero, while 17% have upgrading capacity equal to one.

[34] Suppose that firms producing a higher share of high-quality output, in addition to their mechanically lower scope for further quality upgrading, also face a lower marginal cost of quality upgrading. This would be a concern for our strategy only if this lower cost of upgrading was also related to those firms' ease of integration. (Similarly, a potential lower marginal cost of integration for firms producing a higher share of high-quality output would be a concern only if the lower cost of integration was also related to those firms' ease of upgrading.) The logic of our approach—the argument that firms with high scope for quality upgrading face stronger incentives to upgrade quality when the quality premium is high—would then need to not only hold but outweigh any such countervailing forces for this empirical strategy to yield evidence supporting our hypothesis.

the other top exporters all specialize in producing either high- or low-quality fish meal, shifts in these quotas are likely to impact the relative prices of high- and low-quality fish meal and hence the quality premium. While they may also affect the average price level, we know from column 4 of table 4 that such shifts are unlikely to impact integration decisions.

Specifically, our instruments are the seasonal quotas for the type of fish used to produce fish meal in each of the top five fish meal–exporting countries for which quota information could be found, excluding Peru: Chile, Denmark, and Iceland, all of which specialize in high quality grades.[35] We interact these quotas with $\text{UpgradingCapacity}_{it-1}$ to instrument for $\text{UpgradingCapacity}_{it-1} \times \text{QualityPremium}_t$ as follows:

$$\begin{aligned}\text{UpgradingCapacity}_{it-1} &\times \text{QualityPremium}_t \\ &= \alpha + \sum_c \beta_c \text{UpgradingCapacity}_{it-1} \times \text{FishingQuota}_{ct} \\ &\quad + \beta_2 \text{UpgradableShareOfProduction}_{it-1} + \gamma_i + \delta_t + \varepsilon_{it},\end{aligned} \tag{4}$$

where c is an exporter country and FishingQuota_{ct} is the log fishing quota of country c in season t.

The first stage, shown in columns 7 and 8 of panel B of table 4, is strong. The estimated coefficient on all three countries' quota is significant and negative, reflecting our proposed mechanism: because increases in the quotas of these high-quality-producing countries raise the global supply of high-quality fish meal, this decreases the quality premium.

The second stage with $\text{UpgradingCapacity}_{it-1}$ defined as the upgradable share of production is shown in column 7 of table 4. The IV estimate of $\hat{\beta}$ is of very similar magnitude to the OLS estimate in column 5 and is highly significant. The estimate in column 7 implies that, when the quality premium in Peru rises by 5%, a firm with a high upgradable share of production—one that produces only low-quality output—increases its share VI by about 30% when compared with a firm producing only high-quality output. The IV estimate with $\text{UpgradingCapacity}_{it-1}$ defined as a low-quality producer is shown in column 8; this estimate is similar in magnitude to the corresponding OLS estimate in column 6 but is slightly lower and not statistically significant.

In section II.C, we decomposed the overall variation in share VI into the components explained by each of the three margins through which

---

[35] Chile, Denmark, and Iceland have had aggregate fishing quotas in place for the relevant fish species throughout our sample period (IFFO 2014; Tanoue 2015; IRF 2017; European Commission 2018). Thailand—the fifth of the top five fish meal–exporting countries, which specializes in low quality grades—appears to have introduced such a system in 2015, but quota information for Thailand could not be found.

firms can integrate their production process. In table B2, we replace share VI itself with these three margins in IV regressions, similar to the one in column 7 of table 4 and corresponding OLS regressions, which are analogous to column 5. The results show that firms' share VI response to changes in firm-specific incentives to upgrade quality arising from exogenous variation in the quality premium is almost entirely driven by the extensive margin—that is, acquiring and selling suppliers. The intensive margin response—the intensity with which firms use their integrated suppliers—is also positive and significant but of much smaller magnitude. Use of independent suppliers appears not to drive manufacturers' share VI response to the quality premium: we find a small and negative but statistically insignificant coefficient with this margin on the right-hand side.

C. Robustness

Our IV strategy rests on the assumption that the quantities produced by other top fish meal–exporting countries affect integration decisions in Peru through their impact on market prices. If this argument is correct, we would expect these production volumes to manifest themselves in the price of high-quality fish meal locally, which in turn should impact the quality premium in Peru. In columns 9 and 10 of table 4, we repeat the regression in equation (4) but now instrument using the high quality grade–specific price in other top exporting countries rather than the fishing quotas themselves. High-quality fish meal prices in Chile, Denmark, and Iceland—shown in figure A4, along with Peru's price—are highly correlated with those in Peru.

For both definitions of UpgradingCapacity, the second-stage results closely match the OLS results in terms of magnitude and statistical significance. Of course, the price of high-quality output in, say, Chile may be responding to the price in Peru. Still, it is reassuring for our fishing quota–based IV strategy that variation in the quality premium that correlates with the prices in other top fish meal–producing countries is associated with integration in Peru. The signs in the first stage are slightly more mixed: positive for the prices in two—but not all three—of the other top exporters. However, it is not clear that the residual variation in the equilibrium price in Denmark, conditional on the price in Chile, will necessarily be predictive of the price in Peru. In table A5, we run pairwise versions of our first-stage regression. Indeed, when each country is included individually, we see the expected positive and significant correlation between the price in each country exporting large amounts of high-quality fish meal and the quality premium in Peru (interacted with our measure of UpgradingCapacity).

In columns 1 and 2 of table A3, we show two further robustness exercises for our primary, fishing quota–based IV specification.[36] To rule out the possibility that the estimated integration response to the quality premium is driven by changes in Peruvian firms' scale resulting from changes in other countries' fishing quotas, in column 1 we control for changes in total firm production. Specifically, we include $\log(\text{production})_t - \log(\text{production})_{t-1}$ as a control. The magnitude and significance of the estimated coefficient is effectively unchanged relative to column 7 of table 4.

In column 2 of table A3, we show that our approach is not confounded by the instrument's reliance on exporters that specialize in high quality grades. In this specification, we include Thailand—the fifth of the top five fish meal exporters, which specializes in low-quality fish meal—in the first stage. While information on Thailand's fishing quotas was not available, we use its realized production level. Again, the sign and significance of our second-stage estimate is effectively unchanged. Perhaps more notable is the first stage: the negative signs on the quotas of all three of the high quality–producing countries remain, but we see a positive and significant sign on the quantity exported by Thailand. This is consistent with the intuition underlying our instrument—that is, an increase in the quantity produced by a low-quality exporter causing the low-quality price to drop, increasing the quality premium.

In appendix D, we exploit a different form of variation and find results consistent with those discussed in this subsection. We show that manufacturers respond the same way to variation in firm-specific, quality-differentiated demand shocks as they do to analogous shocks to the quality price premium—integrating suppliers and increasing share VI when relative demand for high quality grades increases and selling boats and decreasing share VI when relative demand for high quality grades decreases. To do so, we construct instruments for firm-specific demand shocks that exploit the fact that each importer country tends to import very specific quality grades, that importer countries' relative demand fluctuates over time, and that changes in demand from a given country matter more for firms that previously exported to that country. We follow many fruitful applications of such an approach in the trade literature (see, e.g., Park et al. 2010; Brambilla, Lederman, and Porto 2012; Tintelnot, Mogstad, and Dhyne 2017); our implementation closely follows Bastos, Silva, and Verhoogen (2018).

D. Interpretation

The results discussed in this section are consistent with this paper's hypothesis and the theoretical framework in section III. In our model, a firm

---

[36] For brevity, we consider only our primary measure of UpgradingCapacity—UpgradableShareOfProduction—throughout this table.

integrates suppliers when its returns to upgrading quality rise because it is difficult to ensure that independent suppliers deliver high-quality inputs when the quantity they produce is valued by other buyers in the market. We now consider whether firms' decision to integrate suppliers when the benefits of quality upgrading rise can be explained by alternative theories.

A first possibility is that firms simultaneously choose their organizational structure and output quality, and shocks—for example, to demand (Legros and Newman 2013; Alfaro et al. 2016)—affect both without the two being directly related. Such a story is difficult to reconcile with the fact that Peruvian fish meal manufacturers integrate suppliers in response to increases in the relative price of high-quality output but not in response to increases in the average price of fish meal.

The same is true for a second possibility—namely, that firms, when the benefits of producing high-quality output rise, buy suppliers so as to restrict competitors' access to independent suppliers and thereby capture a higher share of a newly appealing market segment that happens to be the high-quality one (Ordover, Saloner, and Salop 1990). If such a story explained our results, controlling for the size of a firm's production should reduce the estimated coefficient on the quality premium, and we should see manufacturers integrating suppliers also when the price of low-quality (or any quality) fish meal rises—unless integrated suppliers are more useful when producing high-quality output (as we conjecture).

A third—and related—possibility is that the integration decisions we observe are driven by supply assurance motives (Macchiavello and Miquel-Florensa 2016; Martinez-Carrasco 2017). One potential argument along these lines is that an increase in the quality premium directly allows firms to satisfy a latent demand for supply assurance—for example, by releasing a credit constraint. However, such a story is inconsistent with the facts that (i) those producing primarily low quality—and hence benefiting the least from the quality premium in terms of cashflow—are the most likely to integrate when the quality premium rises and (ii) we see no integration response when the average price rises. An alternative supply assurance story is that an increase in the quality premium causes firms to match with new foreign buyers of fish meal that demand both high quality and supply assurance simultaneously. This and other supply assurance–based stories assume that an increase in our primary measure of vertical integration—share VI—allows firms to better assure supply in the context that we study. While we cannot rule out a role for supply assurance in the integration decisions of Peruvian fish meal producers, we find little empirical evidence that use of integrated suppliers improves their access to supply in our data.[37] Note,

---

[37] First, integration if anything appears to be associated with greater volatility of daily supply in our data. This can be seen from running firm × season–level regressions such as those in table 3 and equivalent plant × month–level versions shown in table 6 with the standard

however, that a firm objective of integrating to secure access to suppliers that are incentivized to deliver the high-quality inputs that are needed to meet the demand for high-quality output—exactly the interpretation we favor—could also be labeled supply assurance. We conclude that manufacturers vertically integrate when the quality premium rises for exogenous reasons at least in part to produce a higher share of high-quality output.

VI. Firms' Organizational Structure and Supplier Behavior

The model in section III predicts that integration is an efficient organizational structure for producing high-quality output for a specific reason: because integration weakens suppliers' incentives to maximize quantity in ways that might be detrimental to the quality of the inputs they produce. As a result, we expect to see suppliers reduce behavior that increases quantity but is harmful to quality when integrated.

A. Estimating How Vertical Integration Affects Suppliers' Quality-Enhancing Actions

We analyze three measures of behavior that capture the trade-off between input quantity and input quality in trip-level data: the total quantity supplied, the maximum distance traveled from the delivery port, and the total time the supplier spends at sea on a given trip. (In addition to these, we also show results for the total number of trips per season.) The first of these three we observe in supply transactions data, while the other two are constructed from boat GPS data. The total quantity supplied is a direct measure of actions taken by the supplier to increase quantity. However, this variable also relates to input quality. This is because the supplier may need to forego quality-increasing actions—such as bringing a lot of ice on board to keep it fresh, not stacking fish high on top of each other to prevent smashing it, and so on—to bring back a high quantity of fish. The maximum distance traveled and total time spent at sea are chosen because they explicitly capture quality-decreasing actions that will tend to increase quantity. Fish freshness—which depends on the time between catch and delivery—is paramount for the protein content of fish meal. As the Food and Agriculture Organization of the United Nations puts it, "Freshness of raw material is important in its effect on the quality of the protein in the end product [fish meal]. The importance of minimizing the time between catching fish and processing, and of keeping the fish

---

deviation of supply across days replacing output quality on the left-hand side. Second, constructing an approximation to the test used in Macchiavello and Morjaria (2015) in our setting reveals little evidence that integrated suppliers are more reliable when positive demand or negative supply shocks occur. These results are available from the authors upon request.

at low temperatures by icing [which reduces the amount of fish a boat can fit], has already been mentioned" (FAO 1986, sec. 10.1.2). Captains must thus balance traveling farther and longer to catch more fish against ensuring freshness. Because all three of these measures of behavior increase quantity but decrease quality, we expect them to decrease after integration (or increase after separation).

Our empirical strategy focuses on "switchers." Switchers are suppliers that are either bought or sold by a fish meal firm during our data period and observed supplying the same plant within the firm in question both before and after the change in status. We include supplier × plant fixed effects and hence compare the behavior of a specific supplier within a specific relationship before versus after integration (or deintegration).

As discussed in section II, we observe 103 instances in which a fish meal firm acquires a supplier that is initially owned independently, 32 instances in which a supplier is sold from a fish meal firm to an independent buyer, and 50 instances in which a supplier is sold from one fish meal firm to another. Conveniently, a subset of our qualifying switches—in which the supplier is observed supplying the firm in question both before and after the change in status—comes from this last set of firm-to-firm supplier transitions. This is because integrated suppliers sometimes supply other fish meal firms.[38] We exploit these transitions in which an always-integrated supplier's relationship with a specific firm changes below.

We do not observe any significant changes in suppliers' characteristics when switching in or out of integration with the plant supplied. Thus, while any average differences between the behavior of independent and integrated suppliers might be attributable in part to boat characteristics,[39] our analysis of within-supplier changes in behavior is unlikely to be influenced by these attributes. Recall also that we saw in figure 1 that suppliers that get integrated or sold deliver to the acquiring/selling firm 63% of the time before integration (or after deintegration): integration typically implies a simple change in the formal status of the relationship between a firm/plant and a supplier engaged in frequent and continuing interactions.

We estimate regressions of the following form:

$$B_{ijt} = \alpha + \beta \mathbb{1}[\text{VI} \times \text{supplies owner firm}]_{ijt} + \gamma_{ij} + \delta_t + \varepsilon_{ijt}, \qquad (5)$$

where B_{ijt} is a measure of the behavior of supplier i delivering to plant j on date t and $\mathbb{1}[\text{VI} \times \text{supplies owner firm}]_{ijt}$ is an indicator for the supplier

---

[38] A firm's output objectives may vary across time within seasons, and fish move around and the location of a catch constrains the set of plants a boat can deliver to. As a result, and as seen in fig. 1*A*, integrated suppliers on average deliver to other firms just over 10% of the time.

[39] As noted in sec. IV.B and shown in table A2, on observable features, the average switcher falls in between the average always-independent boat and the average always-integrated boat, but closer to the latter.

being integrated with the plant it delivers to on date t. We include date fixed effects (δ_t) to control for potential date-specific effects and supplier × plant fixed effects (γ_{ij}) to focus on how integration affects the behavior of a specific supplier supplying a specific plant. We cluster the standard errors at the boat level. We present our results in table 5 and show both the total number of observations and the total number of unique suppliers included in each column.

Column 1 of panel A of table 5 shows that, when integrated and supplying a parent plant, a boat delivers on average about 10% less per trip compared with when it supplies the same plant while independent. This result is clearly consistent with integration offering lower-powered incentives to produce quantity and also suggests that integrated suppliers dedicate more of their storage capacity to ice and/or are more concerned with crushing fish. Columns 2 and 3 show that boats fish approximately 5% closer to the port of delivery and spend on average 3% less time at sea on a trip when integrated with the plant supplied.

In panel C of table 5, we consider an element of boat behavior that is aggregated to the season level. In particular, we regress the log total number of trips per season on season and supplier fixed effects, as well as an indicator equal to one if the supplier is vertically integrated in season t. While our results in panel A show that suppliers deliver a lower quantity per trip once integrated, this does not appear to translate into less overall effort at the season level. In fact, our results indicate that a given supplier conducts significantly more trips each season when integrated. In other words, integration is associated with more frequent trips that are shorter and closer and result in lower quantities.

These results suggest that, when integrated, suppliers reduce costly actions associated with long trips and bring back fresher fish as a result—as the managers in the industry we interviewed reported to us. (In the words of a prominent executive of Peru's National Fisheries Society, "Independent boats prefer to extend their fishing trips [until] they are at full hold capacity, so as to maximize quantity, and this is not good for fish quality . . . as a consequence of integration, they must adopt my rules. Things like saying, 'Hey, you must offload raw 24 hours after having caught it, at the maximum'" [authors' translation].)

A potential concern is the possibility that these results reflect a correlation between the timing of boat purchases and trends in behavior. For example, a firm might choose to purchase a supplier because that supplier has recently begun prioritizing quality. To address this, figure A5 plots the trends in each of the three trip-level behaviors we study in the period just before and after integration as measured with the raw data. Specifically, for the set of supplier × plant pairs that transition from nonintegrated to integrated, we consider interactions in the 45 days (panel A) or 45 trips (panel B) before and after integration. These plots show no evidence that

TABLE 5

SUPPLIER BEHAVIOR AND VERTICAL INTEGRATION

	DEPENDENT VARIABLE: LOG(QUANTITY SUPPLIED) (1)	DEPENDENT VARIABLE: LOG(MAXIMUM DISTANCE FROM THE PLANT'S PORT) (2)	DEPENDENT VARIABLE: LOG(TOTAL TIME SPENT AT SEA) (3)
	A. Identified from All Switchers (Independent to VI, VI to Independent, and VI to VI)		
$\mathbb{1}$[VI × supplies owner firm]	−.096*** (.023)	−.053*** (.019)	−.030* (.016)
Date fixed effects	Yes	Yes	Yes
Supplier × plant fixed effects	Yes	Yes	Yes
Observations	319,827	140,365	163,165
Number of distinct suppliers	1,149	1,081	1,081
R^2	.62	.42	.34
	B. Identified Only from VI Switchers Changing Ownership (VI to VI)		
$\mathbb{1}$[always VI × supplies owner firm]	−.148*** (.026)	−.077*** (.025)	−.067*** (.022)
Date fixed effects	Yes	Yes	Yes
Supplier × plant fixed effects	Yes	Yes	Yes
Observations	319,827	140,365	163,165
Number of distinct suppliers	1,149	1,081	1,081
R^2	.62	.42	.34

C. Integration and Seasonal Effort

	Dependent Variable: Log(Number of Trips per Season)
$\mathbb{1}$[VI]	.190*** (.053)
Season fixed effects	Yes
Supplier fixed effects	Yes
Observations	12,992
Number of distinct suppliers	1,149
R^2	.61

NOTE.—For panels A and B, one observation is a boat during a fishing trip. For panel C, one observation is a boat during a fishing season. "Quantity Supplied" is the amount of fish the boat delivers to the plant per trip. "Maximum Distance from the Plant's Port" is the maximum distance a specific boat is observed away from port; this can be measured only if the boat leaves from and arrives at the same port. "Total Time Spent at Sea" is the amount of time the boat is away from port per trip. In panel A, we define $\mathbb{1}$[VI × supplies owner firm] to be equal to one if the supplier is (i) currently vertically integrated and (ii) currently delivering to its parent firm. In panel B, we define $\mathbb{1}$[always VI × supplies owner firm] to be equal to one if the supplier is (i) always owned by a fish meal firm and (ii) currently delivering to its parent firm. Because we include supplier × plant fixed effects, $\mathbb{1}$[VI × supplies owner firm] and $\mathbb{1}$[always VI × supplies owner firm] are identified based only on suppliers who change ownership during our sample period. In these panels, cols. 1, 2, and 3 reflect 1,749, 1,277, and 1,079 unique fishing days, respectively, covering the period from April 2009 to April 2017 for col. 1 and from April 2009 to November 2015 for cols. 2 and 3. In panel C, $\mathbb{1}$[VI] is identified from boats that switch from independent to VI or from VI to independent. This panel covers 16 unique fishing seasons from April 2009 to April 2018. Each column shows the total number of observations and total number of unique suppliers in the corresponding specification. The number of observations varies from one column to the next, as GPS variables for a given trip are sometimes missing. Note that most of this missingness is within supplier, as GPS data are available for nearly all boats for at least a portion of our sample. Standard errors clustered at the boat level are included in parentheses.

* $p < .1$.

*** $p < .01$.

quantity-oriented boat behaviors trend downward before integration. Furthermore, all three behaviors display an evident drop just after integration.

In our model, integration is defined by asset ownership, as in Grossman and Hart (1986). Indeed, suppliers' change in behavior appears to be the result of integration itself, as opposed to any long-term relationship that coincides with integration. In table C1, we show that—absent integration—repeated interactions with the same plant do not lead to a change in quality-increasing actions, consistent with the predictions of the dynamic version of our theoretical framework also shown in appendix C. Thus, while repeated interactions help fish meal manufacturers and independent suppliers exchange supply-and-demand assurance (Martinez-Carrasco 2017), they appear not to offer an alternative way to achieve the change in quality-conducive incentives associated with integration in the context we study.

B. Interpretation

In this section, we have seen that a given supplier supplying a given plant takes more quality-oriented and less quantity-oriented actions when the two are vertically integrated. Our interpretation is that integration dampens high-powered incentives to prioritize quantity over quality that suppliers face on the open market. Other changes in incentives that arise because of integration could also play a role. Perhaps the most plausible possibility is that what constrains suppliers' input quality is not their incentive to prioritize quality but their knowledge of how to do so. If so, firms may be reluctant to teach a supplier how to upgrade input quality if the supplier is independent (Pigou 1912). We can shed some light on the likelihood that such a story explains our results in this section by exploiting the fact that integrated suppliers occasionally deliver inputs to other firms. We analyze the behavior of suppliers that are always integrated with a fish meal firm (but sold from one firm to another during our sample period) and that supply a plant belonging to the acquiring and/or the selling firm both before and after the sale. We thus continue to focus on changes in supplier behavior within a supplier × plant pair.[40]

As seen in panel B of table 5, we find quite similar—even slightly larger—effects compared with panel A. If acquired, a supplier changes its behavior consistent with prioritizing quantity less—to the benefit of quality—while delivering to the acquiring firm. This pattern is identical to how previously independent switchers change their behavior once integrated, suggesting that a story in which integration enables knowledge transfer from Peruvian manufacturers to their suppliers is unlikely to be the primary

---

[40] To implement this analysis, we repeat the specification in eq. (5) but define $\mathbb{I}$[VI × supplies owner firm] to be equal to one if the supplier is (i) always owned by a fish meal firm and (ii) currently delivering to its parent firm.

explanation behind the difference in supplier behavior when integrated. In other contexts, such knowledge transfers may provide an additional—or the primary—motivation for vertical integration (see Atalay, Hortaçsu, and Syverson 2014).

The results in panel B of table 5 also underscore the patterns shown in figure A5—that it is not the case that firms simply choose to integrate suppliers that have already begun changing their behaviors. This provides further support for the parallel trends assumption that underlies a causal interpretation of the results in panel A.

Another alternative explanation of the change in supplier behavior when integrated is that our results simply reflect the fact that integrated suppliers face low-powered, fixed pay incentives and that the behaviors we see do not generate any input quality benefits that manufacturers are aware of and act on. Such a story is difficult to reconcile with the results in panel B of table 5—integrated suppliers behave "as independent suppliers" when supplying firms other than the parent firm. The intuition captured in the framework in section III and the corresponding evidence in section V suggest that it is likely in the parent firm's interest for its integrated suppliers to prioritize quantity over quality in such situations. A mechanism focusing on integrated suppliers' low-powered incentives in isolation is also difficult to reconcile with this paper's central finding that firms integrate suppliers when the quality premium rises for exogenous reasons but not when the general price level increases. Suppliers' pay system may affect their behavior, but the evidence discussed in this subsection suggests that, if so, it does so differently when combined with incentives that come from delivering to a parent firm.

VII. Vertical Integration and Output Quality

In section V, we saw that firms vertically integrate when the benefits of shifting from low- to high-quality production rise. In section VI, we saw that suppliers that get integrated behave in a more input-quality-increasing and less input-quantity-increasing manner. In this section, we show that plants' output quality responds to integration exactly how we would expect if the integration-induced change in supplier behavior improves input quality. That is, we show that plants with a greater share of vertically integrated inputs produce higher-quality output. This provides empirical support for our model's third prediction—namely, that vertical integration is an effective organizational strategy for producing high-quality output. It is also consistent with the reports provided by the managers in the industry we interviewed. For example, Ricardo Bernales Parodi—managing director of Pesquera Diamante, Peru's third-largest fish meal company—said the following: "From the boat to the factory, and to the commercialization, the flour has quality A, B, C, and D. If I only bought from my boats, I would make an effort so that 95% would be A and B, and only 5% of C and D. But when buying from third parties, I end up with 20% of C and D" (authors' translation).

We first show that there is a robust relationship between changes over time in the share of inputs that individual plants obtain from integrated suppliers (share VI) and changes in their output quality that goes beyond the firm-level evidence discussed in section IV.C. We then attempt to isolate shifts in a plant's share VI that occur for plausibly exogenous reasons. We show evidence from two IV approaches. The first exploits geographic variation in the local concentration of wooden boats—a particular type of supplier that is prohibited from being integrated by regulation—as well as in independent (nonintegrated) boats more broadly. The second exploits variation in weather patterns across the ports around which plants are clustered. Together, the results we present strongly suggest that the relationship between share VI and output quality arises because integration increases output quality.

A. *Estimating How Vertical Integration Affects Output Quality*

If integration increases output quality because integrated suppliers deliver higher-quality inputs, then the relationship between share VI and output quality we observe at the firm level should hold at the plant level as well. This is what we find in table 6. We repeat regression (1) from section IV.C but now at the plant (i) × month (t) level, the lowest level at which we directly observe output quality.

The sample consists of all 94 plants we observe across Peru. We include plant and month fixed effects and thus focus on variation in share VI across months within a given plant.[41] The results in columns 1 and 2 of table 6 imply that the share of a plant's output that is of the high quality type would be 8%–12% higher if its parent firm were to integrate all (relative to none) of the plant's suppliers. We also find the same integration-quality relationship across different plants within the same firm over time, as shown in table A6. There we use internal data provided to us by a single major firm.[42]

In combination with table 3, columns 1 and 2 of table 6 establish a positive, statistically significant, and quantitatively consistent association

---

[41] We observe whether each plant has any high technology installed, so HighTech_{it} is now a dummy variable.

[42] The firm's data report which plant produced the fish meal included in a given export shipment. In addition to "share high quality," for this firm's plants we can thus also measure output quality as the fine-grained quality grade inferred from export unit values and auxiliary price data, as we do for firms in cols. 4 and 5 of panel B of table 3. The magnitude and significance of the estimates are very similar to those in panel B of table 3.

TABLE 6
Output Quality and Share of Inputs from Vertically Integrated Suppliers

	Impact of Share of VI Inputs on Quality							
	Dependent Variable: High-Quality Share of Production							
	OLS		IV: Individual Boats		IV: Wooden Boats		IV: Weather	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Share of inputs from VI suppliers	.102**	.064**	.165***	.139**	.160***	.142**	.102	.094*
	(.037)	(.029)	(.056)	(.059)	(.055)	(.069)	(.069)	(.056)
Has high technology	No	Yes	No	Yes	No	Yes	No	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean of dependent variable	.84	.84	.84	.84	.84	.84	.83	.83
Observations	2,647	2,647	2,647	2,647	2,647	2,647	1,413	1,413
R^2	.58	.67	.58	.67	.58	.66	.59	.65

Note.—One observation is a plant in a particular month. “High-Quality Share of Production” is the share of a plant’s total production during a month that is reported as high-quality (prime) output. “Share of inputs from VI suppliers” is the share of a plant’s inputs that comes from VI suppliers during a month. “Has high technology” is a dummy variable equal to one if the plant in question has any steam-drying technologies installed. Columns 3 and 4 instrument for “Share of inputs from VI suppliers” with the number of independent boats present locally (in the plant’s port) in the month in question, excluding those that interact directly with the plant itself. Columns 5 and 6 instrument for “Share of inputs from VI suppliers” with the number of wooden boats present locally (in the plant’s port) in the month in question, excluding those that interact directly with the plant itself. Columns 7 and 8 instrument for “Share of inputs from VI suppliers” with the mean, median, 1st and 5th quintiles, and standard deviation of monthly port-level cloudiness and wind speed as well as all interactions between these variables. Because of data availability, weather instruments are available only for a subset of months. The sample in cols. 1–6 includes 94 unique plants in 79 unique months from April 2009 to June 2017. The sample in cols. 7 and 8 includes 88 unique plants in 62 unique months spanning April 2009 to December 2016. Standard errors are clustered at the port level.

* $p < .1$.
** $p < .05$.
*** $p < .01$.

between share VI and directly observed output quality at both the firm and the plant levels. Of course, the fact that these correlations hold for individual plants does not rule out noncausal interpretations. It could be that plant- or port-specific shocks—for example, to productivity—occur and simultaneously impact both the quality of a plant's output and the share of the plant's supply coming from integrated suppliers. Alternatively, there might be particular plant-level strategic choices that lead plants to simultaneously increase share VI and produce higher output quality without a causal link between the two. We now turn to two IV strategies that help us isolate plausibly exogenous variation in share VI at the plant level.

B. Instrumenting with the Presence of Wooden and Independent Boats

Our first IV strategy exploits the local presence of wooden fishing boats—which are, by law, independently owned—as a source of variation in a plant's share VI. These and other independent boats move up and down the coast as a function of weather, presence of fish, and other factors. The logic of our instrument is simply that, at times when there happens to be an abundance of independent suppliers in a given area for exogenous reasons, firms are more likely to use those suppliers. A plant's choice of suppliers is the result of a complex optimization process involving output quality objectives on the one hand and the relative cost of using integrated versus independent suppliers on the other. At times when input from independent suppliers is relatively cheap, optimizing plants will tend to decrease their share VI—even holding their incentives to produce quality constant. When independent suppliers are scarce, the cost of their inputs is likely to be high, and vice versa. This suggests that measures of the presence of independent suppliers may serve as instruments for a plant's share VI.[43]

With this in mind, we consider the number of wooden—and hence independent by law—suppliers active in a port (cluster of plants) in a given month as a proxy for the relative cost of using independent suppliers. For this to be a valid instrument, it must not impact or be correlated with output quality other than via share VI. Of course, a plant's quality objectives may themselves influence suppliers' whereabouts. The plant may, for example, request deliveries from particular wooden suppliers. We thus use a leave-firm-out measure of the presence of independent-by-law suppliers in a given port during a given period. In particular, our instrument for

---

[43] It is important to note that the use of wooden suppliers is not limited to small firms. The proportion of inputs sourced from wooden boats is relatively stable—and if anything slightly increasing—across the distribution of firm size. Single-plant firms source on average 23% of their inputs per season from wooden boats, while the largest firms (with eight or more plants) source on average 33% per season.

share VI is the number of wooden boats present, excluding any that supply the firm to which the plant in question belongs. We also show results for an analogous instrument using all independent suppliers, not restricting to wooden boats.

We believe the exclusion restriction to be valid given two key assumptions: (i) that plants are effective price takers in the local market (or more specifically, that an individual plant's actions that correlate with quality production do not influence the likelihood that independent boats serve other firms in the port) and (ii) that the presence of independent boats is not correlated with port-level shocks that influence the production of output quality at all plants in a port. Of course, it is possible that these assumptions do not hold. However, given our leave-out strategy and a robustness exercise provided below that exploits the lagged presence of wooden and independent suppliers, we believe assumption i to be quite plausible. Furthermore, two pieces of evidence indicate that assumption ii is reasonable. First, beyond the presence of wooden boats, we find no evidence of port-level correlations in the production of output quality.[44] That is, port-level shocks to output quality appear limited in general. Second, from our understanding of the industry, the types of factors that are attractive to independent fishing crews (who are paid per ton delivered)—for example, an abundance of fish close to a port—are ones we expect to be positively correlated with quality production. This stands in contrast to the argument underlying our IV, which predicts a negative relationship between independent boats and quality production (and which we confirm below).

The first stage, shown in table A7, is strong: the number of wooden (or independent) boats supplying other plants in the port is highly correlated with the share of integrated supply to the plant in question during the same period. The sign is negative, suggesting that—even using our leave-out proxy—the availability of independent suppliers influences share VI in the manner we expect. A plant substitutes toward integrated suppliers when independent suppliers are relatively scarce, and vice versa.

Results from the IV specifications are in columns 3–6 of table 6. The IV estimates are larger in magnitude but of the same sign and statistical significance as the corresponding OLS estimates. This holds whether we restrict attention to suppliers that are independent by law or include all independent suppliers. Additionally, the same is true in a similar but

---

[44] For example, consider a regression of the share of high-quality output at the plant level on the average share of high-quality output of other plants in the port, controlling for month and plant fixed effects as well as the presence of independent suppliers. If a given plant's output quality and that of other plants were perfectly positively or negatively correlated across time, the coefficient on the average share of high-quality output of other plants in the port would be $+1$ and -1, respectively. We find a coefficient of 0.04, with a standard error of 0.080.

not identical specification shown in table A6, which utilizes internal data from the firm that shared its data with us.

We also supplement this with a robustness exercise that exploits persistence in boat locations—that is, the tendency for boats to remain in the same port for multiple trips. Specifically, we instrument for share VI with the lagged presence of independent or wooden suppliers. To construct this measure for each port in month t, we first record the port of delivery for each boat in the last trip conducted in month $t - 1$.[45] We then count the total number of boats present in each port. Both the estimated first- and second-stage results, shown in columns 1 and 2 of table A8, are consistent with columns 4 and 6 of table 6, though the second-stage results are less precisely estimated. While this robustness exercise is subject to some of the same potential criticisms as our primary specification, this instrument will be valid even if our assumption i above is violated, so long as confounding plant-level actions are not autocorrelated. Similarly, this instrument will be valid even if assumption ii above is violated, so long as relevant port-level shocks are not autocorrelated.

C. *Instrumenting with Wind Speed and Cloudiness*

It is clearly possible to envision particular port-level shocks (or equilibrium responses to a given plant's actions) that cause the exclusion restriction of the IV presented in the previous subsection to fail. To address this, we conclude our analysis with an alternative IV strategy that exploits port-level weather conditions. Of course, weather variables are not inherently perfect instruments in this context. It is possible, for example, that some variation in weather might influence the relative share of VI suppliers at a given port while simultaneously directly influencing the quality of fish meal (e.g., by changing the location of the fish themselves). However, we focus specifically on two dimensions of weather conditions that plausibly influence the share of VI suppliers without impacting the underlying fish quality: wind speed and cloudiness. Our rationale is that the smaller wooden and independent boats might be less likely to fish in a port experiencing severe weather, as captured by these variables.

The National Oceanic and Atmospheric Administration (NOAA) provides monthly data at a 1 × 1-degree level on several aspects of the distribution of these two variables.[46] Given that there is no theoretical model predicting exactly how wind speed and cloudiness should correlate with a

---

[45] If no trip was conducted in $t - 1$, we select the first trip of month t. For brevity, we omit specifications that do not control for "Has high technology."

[46] Our weather data are drawn from the Monthly International Comprehensive Ocean-Atmosphere Data Set provided by the NOAA, with spatial coverage based on 1 × 1-degree boxes. Cloudiness refers to total cloudiness measured in oktas. Wind speed refers to scalar wind measured in 0.01 m/second.

plant's use of integrated suppliers, and given that the available moments of these variables individually have little predictive power, we take a predictive approach to the first stage. That is, we project share VI onto all available aspects of the distributions of cloudiness and wind speed, as well as their interactions. If $\mathbf{z}_i$ is a vector containing the mean, median, bottom and top sextiles, and standard deviation of both cloudiness and wind speed and $\otimes$ represents the Kronecker product, we estimate as a first stage (for plant i in month t)

$$\mathrm{VI}_{it} = \gamma_0 + (\mathbf{z}_i \otimes \mathbf{z}_i)'\Gamma + \gamma_1 \mathrm{HighTech}_{it} + \eta_i + \theta_t + \varepsilon_{it}. \tag{6}$$

Here η_i and θ_t are plant and month fixed effects and $\mathrm{HighTech}_{it}$ is an indicator equal to one if the plant has high technology, as above.

The results from our weather-based IV approach in columns 7 and 8 of table 6 show an increase in output quality with greater use of integrated suppliers that is very similar in magnitude to the OLS and independent boat–presence IV results in columns 1–6 and is marginally significant.[47] In table A8, we show basic versions of this approach to provide transparency on the strategy and to confirm that our results are not driven by our saturated approach to the first stage. In columns 3 and 4, we include first the full set of wind speed variables and then the full set of cloudiness variables. The second-stage results are consistent in magnitude with the range of our OLS and IV results. In column 5, we include the full set of wind speed and cloudiness variables. With this full set, we again see a coefficient that is consistent in magnitude with our original IV. Though these simpler versions of the weather-based instrument are weak instruments, the patterns in table A8 provide support for our interpretation of the full equation (6) estimates in columns 7 and 8 of table 6.

D. Interpretation

Our interpretation of the results presented in this section is that access to inputs from integrated suppliers directly increases output quality. In conjunction with the model in section III and the results in section VI, our analysis suggests that this is because a manufacturer can incentivize integrated suppliers to reduce behaviors that decrease quality.

Of course, output quality may in principle covary with organizational structure without an underlying causal relationship. Perhaps the most plausible noncausal link between quality upgrading and integration—that firm-level shocks or factors (e.g., overall growth) cause firms to simultaneously and independently produce higher-quality output and acquire more suppliers—is ruled out by the simple OLS regressions in tables 6 and

---

[47] We omit the first stage for cols. 7 and 8 given the large number of interactions.

A6: output quality correlates with the use of integrated suppliers at the time of production across plants, including within firms.

While each of our IV strategies is subject to potential criticism, the results go a step further by demonstrating that the same relationships hold when we restrict attention to (i) the local presence of independent or wooden suppliers, (ii) the lagged presence of independent or wooden suppliers, or (iii) port-level variation in wind speed and cloud cover. The consistency of our results across these approaches reinforces our basic OLS results.

Combining these findings with those found in section V, we conclude that it is not the case that higher output quality in vertically integrated Peruvian fish meal manufacturers is simply an ignored by-product of integration decisions made for other reasons or that integration and output quality are causally unrelated in the minds of the firms in our sample. Our evidence here and in table 3 suggests that integration is an effective strategy for producing high-quality output—that is, particular plants or firms that interact with integrated suppliers are able to produce higher-quality output. In section V, we showed that increasing output quality is an explicit motive for integration. In other words, our evidence indicates that vertical integration increases output quality and that, as a result, firms integrate suppliers when they intend to increase the quality of the goods they produce.

VIII. Conclusion

This paper identifies an overlooked motivation for—and consequence of—vertical integration in incomplete contracts settings—namely, downstream firms integrating to be able to produce output of high enough quality to sell to high-paying consumers abroad. We first present a simple theoretical framework that captures how suppliers and the downstream firms they supply are expected to behave in sectors where firms produce vertically differentiated goods and contracts are incomplete. The model motivates three predictions that follow logically from each other: how the quality premium—the difference between the price of high- and low-quality output—affects firms' choice of organizational structure, how suppliers' behavior changes with integration, and how integration consequently affects output quality.

We test these predictions using transaction-level data and direct measures of the quality grades that manufacturers produce in Peru's fish meal industry. We show that when the returns to shifting from low- to high-quality production rise for exogenous reasons, firms acquire more of their suppliers. This strategy appears to be effective because integration allows firms to incentivize input-quality-increasing behavior from their suppliers: fishing boats change their behavior in a way consistent with delivering fresher fish when they are acquired by the firm they supply—which is known to enable production of higher-quality fish meal. Finally, we show that firms ultimately produce higher-quality output when their organizational structure is more vertically integrated. The evidence we present thus suggests that—while firms vertically integrate for many different reasons—one motive for integration is quality upgrading. That is, in settings such as the one we study, integration is an explicit organizational choice made to climb the quality ladder.

A natural next question is the generality of this finding. Theory suggests that integration can help address contracting problems that are typical when input quality is difficult to observe (and hence incentivize), as is often the case. Despite vertical integration overall being common in developing countries (Acemoglu, Johnson, and Mitton 2005; Macchiavello 2011), it may thus be that the extent of vertical integration observed among firms in the developing world is actually suboptimally low, since upgrading output quality is essential for export-driven economic development. Of course, in a world with perfect contracting, there might be no need for integration. As such, our paper's results conversely imply that improvements in contract enforcement may reduce the need for firms to rely on organizational structure to align their suppliers' incentives.

References

Acemoglu, Daron, Philippe Aghion, Rachel Griffith, and Fabrizio Zilibotti. 2010. "Vertical Integration and Technology: Theory and Evidence." *J. European Econ. Assoc.* 8 (5): 989–1033.

Acemoglu, Daron, Philippe Aghion, Claire Lelarge, John Van Reenen, and Fabrizio Zilibotti. 2007. "Technology, Information, and the Decentralization of the Firm." *Q.J.E.* 122 (4): 1759–99.

Acemoglu, Daron, Simon Johnson, and Todd Mitton. 2005. "Determinants of Vertical Integration: Finance, Contracts, and Regulation." Working Paper no. 11424, NBER, Cambridge, MA.

Alfaro, Laura, Paola Conconi, Harald Fadinger, and Andrew F. Newman. 2016. "Do Prices Determine Vertical Integration?" *Rev. Econ. Studies* 83 (3): 855–88.

Amodio, Francesco, and Miguel A. Martinez-Carrasco. 2018. "Input Allocation, Workforce Management and Productivity Spillovers: Evidence from Personnel Data." *Rev. Econ. Studies* 85 (4): 1937–70.

Atalay, Enghin, Ali Hortaçsu, and Chad Syverson. 2014. "Vertical Integration and Input Flows." *A.E.R.* 104 (4): 1120–48.

Atkin, David, Amit Khandelwal, and Adam Osman. 2017. "Exporting and Firm Performance: Evidence from a Randomized Experiment." *Q.J.E.* 132 (2): 551–615.

Baker, George, Robert Gibbons, and Kevin J. Murphy. 2001. "Bringing the Market inside the Firm?" *A.E.R.* 91 (2): 212–18.

———. 2002. "Relational Contracts and the Theory of the Firm." *Q.J.E.* 117 (1): 39–84.

Baker, George P., and Thomas N. Hubbard. 2003. “Make versus Buy in Trucking: Asset Ownership, Job Design and Information.” *A.E.R.* 93 (3): 551–72.

———. 2004. “Contractibility and Asset Ownership: On-Board Computers and Governance in US Trucking.” *Q.J.E.* 119 (4): 1443–79.

Baldwin, Richard, and James Harrigan. 2011. “Zeros, Quality, and Space: Trade Theory and Trade Evidence.” *American Econ. J.: Microeconomics* 3 (2): 60–88.

Banerjee, Abhijit V., and Esther Duflo. 2000. “Reputation Effects and the Limits of Contracting: A Study of the Indian Software Industry.” *Q.J.E.* 115 (3): 989–1017.

Bastos, Paulo, Joana Silva, and Eric Verhoogen. 2018. “Export Destinations and Input Prices.” *A.E.R.* 108 (2): 353–92.

Boehm, Johannes. 2018. “The Impact of Contract Enforcement Costs on Value Chains and Aggregate Productivity.” Working paper, Dept. Econ., Sciences Po, Paris.

Brambilla, Irene, Daniel Lederman, and Guido Porto. 2012. “Exports, Export Destinations, and Skills.” *A.E.R.* 102 (7): 3406–38.

———. 2019. “Exporting Firms and the Demand for Skilled Tasks.” *Canadian J. Econ.* 52 (2): 763–83.

Brambilla, Irene, and Guido Porto. 2016. “High-Income Export Destinations, Quality and Wages.” *J. Internat. Econ.* 98 (1): 21–35.

Bresnahan, Timothy F., and Jonathan D. Levin. 2013. “Vertical Integration and Market Structure.” In *The Handbook of Organizational Economics*, edited by Robert Gibbons and John Roberts, 853–90. Princeton, NJ: Princeton Univ. Press.

Breza, Emily, and Andres Liberman. 2017. “Financial Contracting and Organizational Form: Evidence from the Regulation of Trade Credit.” *J. Finance* 72 (1): 291–324.

Coase, Ronald H. 1937. “The Nature of the Firm.” *Economica* 4 (16): 386–405.

De La Puente, Óscar, Juan Carlos Sueiro, Carmen Heck, Giuliana Soldi, and Santiago De La Puente. 2011. “Evaluación de los Sistemas de Gestión Pesquera en el Marco de la Certificación a Cargo del Marine Stewardship Council: La Pesquería Peruana de Anchoveta.” Working paper, Univ. Peruana Cayetano Heredia, Lima.

Eslava, Marcela, Cecilia Fieler, and Daniel Yi Xu. 2018. “Trade, Quality Upgrading, and Input Linkages: A Theory with Evidence from Colombia.” *A.E.R.* 108 (1): 109–46.

European Commission. 2018. “Fishing Quotas.” https://ec.europa.eu/fisheries/cfp/fishing_rules/tacs_en.

FAO (Food and Agriculture Organization of the United Nations). 1986. “The Production of Fish Meal and Oil.” FAO Fisheries Tech. Paper no. 142, Food and Agriculture Org. United Nations, Rome.

Forbes, Silke, and Mara Lederman. 2009. “Adaptation and Vertical Integration in the Airline Industry.” *A.E.R.* 99 (5): 1831–49.

———. 2010. “Does Vertical Integration Affect Firm Performance? Evidence from the Airline Industry.” *RAND J. Econ.* 41 (4): 765–90.

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2008. “Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?” *A.E.R.* 98 (1): 394–425.

Frías, Judith A., David S. Kaplan, and Eric A. Verhoogen. 2009. “Exports and Wage Premia: Evidence from Mexican Employer-Employee Data.” Working paper, Inst. Mexicano Seguro Soc., Mexico City.

Gibbons, Robert. 2005a. “Four Formal(izable) Theories of the Firm?” *J. Econ. Behavior and Org.* 58 (2): 200–245.

———. 2005b. "Incentives between Firms (and Within)." *Management Sci.* 51 (1): 1–150.

Gibbons, Robert, and John Roberts, eds. 2013. *The Handbook of Organizational Economics.* Princeton, NJ: Princeton Univ. Press.

Gil, Ricard. 2009. "Revenue Sharing Distortions and Vertical Integration in the Movie Industry." *J. Law, Econ., and Org.* 25 (2): 579–610.

Gil, Ricard, and Myongjin Kim. 2016. "Does Competition Increase Quality? Evidence from the US Airline Industry." Working paper, Carey Bus. School, Baltimore.

Gil, Ricard, Myongjin Kim, and Giorgio Zanarone. 2016. "The Value of Relational Adaptation in Outsourcing: Evidence from the 2008 Shock to the US Airline Industry." Working paper, Carey Bus. School, Baltimore.

Goldberg, Pinelopi Koujianou, and Nina Pavcnik. 2007. "Distributional Effects of Globalization in Developing Countries." *J. Econ. Literature* 45:39–82.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2019. "Bartik Instruments: What, When, Why, and How." Working paper, Yale Univ., New Haven, CT.

Grossman, Sanford J., and Oliver D. Hart. 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration." *J.P.E.* 94 (4): 691–719.

Hallak, Juan Carlos. 2006. "Product Quality and the Direction of Trade." *J. Internat. Econ.* 68 (1): 238–65.

Hallak, Juan Carlos, and Peter Schott. 2011. "Estimating Cross-Country Differences in Product Quality." *Q.J.E.* 126 (1): 417–74.

Halpern, László, Miklós Koren, and Adam Szeidl. 2015. "Imported Inputs and Productivity." *A.E.R.* 105 (12): 3660–703.

Hart, Oliver, and John Moore. 1990. "Property Rights and the Nature of the Firm." *J.P.E.* 98 (6): 1119–58.

Hart, Oliver, Andrei Shleifer, and Robert W. Vishny. 1997. "The Proper Scope of Government: Theory and an Application to Prisons." *Q.J.E.* 112 (4): 1127–61.

Holmstrom, Bengt. 1999. "The Firm as a Subeconomy." *J. Law, Econ., and Org.* 15 (1): 74–102.

Holmstrom, Bengt, and Paul Milgrom. 1991. "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design." *J. Law, Econ., and Org.* 7:24–52.

———. 1994. "The Firm as an Incentive System." *A.E.R.* 84 (4): 972–91.

Holmstrom, Bengt, and Jean Tirole. 1991. "Transfer Pricing and Organizational Form." *J. Law, Econ., and Org.* 7 (2): 201–29.

Hortaçsu, Ali, and Chad Syverson. 2007. "Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices." *J.P.E.* 115 (2): 250–301.

Hummels, David, and Peter J. Klenow. 2005. "The Variety and Quality of a Nation's Exports." *A.E.R.* 95 (3): 704–23.

IFFO (International Fishmeal and Fish Oil Organisation). 2014. "Chile Fishing Quotas Set and Landing Information Certification." http://www.iffo.net/node/544.

IRF (Iceland Responsible Fisheries). 2017. "Catch Quotas for 2017/2018 in Accordance with Scientific Advice." https://www.responsiblefisheries.is/news-and-media/news/catch-quotas-for-2017-2018-in-accordance-with-scientific-advice.

Johnson, Robert C. 2012. "Trade and Prices with Heterogeneous Firms." *J. Internat. Econ.* 86 (1): 43–56.

Khandelwal, Amit. 2010. "The Long and Short (of) Quality Ladders." *Rev. Econ. Studies* 77 (4): 1450–76.

Kosová, Renáta, Francine Lafontaine, and Rozenn Perrigot. 2013. "Organizational Form and Performance: Evidence from the Hotel Industry." *Rev. Econ. and Statis.* 95 (4): 1303–23.

Kugler, Maurice, and Eric Verhoogen. 2012. "Prices, Plant Size, and Product Quality." *Rev. Econ. Studies* 79 (1): 307–39.

Lafontaine, Francine, and Margaret Slade. 2007. "Vertical Integration and Firm Boundaries: The Evidence." *J. Econ. Literature* 45 (3): 629–85.

Legros, Patrick, and Andrew F. Newman. 2013. "A Price Theory of Vertical and Lateral Integration." *Q.J.E.* 128 (2): 725–70.

Macchiavello, Rocco. 2011. "Financial Development and Vertical Integration: Theory and Evidence." *J. European Econ. Assoc.* 10 (2): 255–89.

Macchiavello, Rocco, and Josepa Miquel-Florensa. 2016. "Vertical Integration and Relational Contracts: Evidence from Costa Rica's Coffee Chain." Working paper, London School Econ.

Macchiavello, Rocco, and Ameet Morjaria. 2015. "The Value of Relationships: Evidence from a Supply Shock to Kenya Rose Exports." *A.E.R.* 105 (9): 2911–45.

Manova, Kalina, and Zhiwei Zhang. 2012. "Export Prices across Firms and Destinations." *Q.J.E.* 127 (1): 379–436.

Martinez-Carrasco, Jose. 2017. "Coordination and Gains from Relational Contracts: Evidence from the Peruvian Anchovy Fishery." Working paper, Duke Univ., Durham, NC.

McMillan, John, and Christopher Woodruff. 1999. "Interfirm Relationships and Informal Credit in Vietnam." *Q.J.E.* 114 (4): 1285–320.

Medina, Pamela. 2017. "Import Competition, Quality Upgrading and Exporting: Evidence from the Peruvian Apparel Industry." Working paper, Univ. Toronto.

Mullainathan, Sendhil, and David Sharfstein. 2001. "Do Firm Boundaries Matter?" *A.E.R.* 91 (2): 1995–99.

Natividad, Gabriel. 2014. "Integration and Productivity: Satellite-Tracked Evidence." *Management Sci.* 60 (7): 1698–718.

Novak, Sharon, and Scott Stern. 2008. "How Does Outsourcing Affect Performance Dynamics? Evidence from the Automobile Industry." *Management Sci.* 54 (12): 1963–79.

Nunn, Nathan. 2007. "Relationship-Specificity, Incomplete Contracts, and the Pattern of Trade." *Q.J.E.* 122 (2): 569–600.

Ordover, Janusz A., Garth Saloner, and Steven C. Salop. 1990. "Equilibrium Vertical Foreclosure." *A.E.R.* 80 (1): 127–42.

Paredes, Carlos E., and Maria E. Gutierrez. 2008. "The Peruvian Anchovy Sector: Costs and Benefits—An Analysis of Recent Behavior and Future Challenges." Paper presented at the fourteenth biennial conference of the International Institute of Fisheries Economics and Trade, Nha Trang, Vietnam, July 23.

Park, Albert, Dean Yang, Xinzheng Shi, and Yuan Jiang. 2010. "Exporting and Firm Performance: Chinese Exporters and the Asian Financial Crisis." *Rev. Econ. and Statis.* 92 (4): 822–42.

Pigou, Arthur C. 1912. *Wealth and Welfare.* London: Macmillan.

Schott, Peter K. 2004. "Across-Product versus Within-Product Specialization in International Trade." *Q.J.E.* 119 (2): 647–78.

Sutton, John. 1991. *Sunk Costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration.* Cambridge, MA: MIT Press.

———. 1998. *Technology and Market Structure: Theory and History.* Cambridge, MA: MIT Press.

Syverson, Chad. 2004. "Product Substitutability and Productivity Dispersion." *Rev. Econ. and Statis.* 86 (2): 534–50.

Tanoue, Wataru. 2015. “Japan’s Total Allowable Catch Systems in Fishery Resource Management.” MMA thesis, School Marine and Environmental Affairs, Univ. Washington. https://digital.lib.washington.edu/researchworks/bitstream/handle/1773/34012/Tanoue_washington_0250O_14735.pdf.

Tintelnot, Felix, Ken Kikkawa, Magne Mogstad, and Emmanuel Dhyne. 2017. “Trade and Domestic Production Networks.” Working paper, Univ. Chicago.

Tveteras, Sigbjorn, Carlos E. Paredes, and Julio Peña-Torres. 2011. “Individual Vessel Quotas in Peru: Stopping the Race for Anchovies.” *Marine Resource Econ.* 26 (3): 225–32.

Verhoogen, Eric A. 2008. “Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector.” *Q.J.E.* 123 (2): 489–530.

Williamson, Oliver. 1971. “The Vertical Integration of Production: Market Failure Considerations.” *A.E.R.* 61 (2): 112–23.

Woodruff, Christopher. 2002. “Non-contractible Investments and Vertical Integration in the Mexican Footwear Industry.” *Internat. J. Indus. Org.* 20 (8): 1197–224.