Vertical Integration, Supplier Behavior, and Quality Upgrading among Exporters*

Christopher Hansman
Imperial College London

Jonas Hjort
Columbia University
& CEPR & NBER

Gianmarco León
Universitat Pompeu Fabra
& BGSE & IPEG & CEPR

Matthieu Teachout
Columbia University

June 26, 2018

Abstract

We study the relationship between firms' output quality and their choice of organizational structure. To do so, we use data on each step of the production and transaction chain that makes up Peruvian fishmeal manufacturing. We first show that quality upgrading is an important motive for vertically integrating. Firms integrate suppliers when the quality premium—the *relative* price of high quality output—rises for exogenous reasons, but not when average or low quality prices rise. The greater a firm's scope for shifting low to high quality production, the greater its integration response. We then show that integration changes suppliers' production behavior. A given supplier's actions are less geared towards increasing quantity and more geared towards maintaining input quality after the supplier is integrated and loses access to alternative pay-per-kilo buyers. Finally, we show that firms and individual plants that use integrated suppliers at the time of production produce a significantly higher *share* of high quality output. In sum, our results suggest that firms change their organizational structure when their output quality objectives change because controlling the incentives of independent suppliers facing a quantity-quality trade-off is difficult, as classical theories of the firm predict.

JEL codes: D2, O1

*_**Winner of the Oliver Williamson Award for Best Paper presented at the Society of Institutional and Organizational Economics' 2018 annual conference in Montréal**_. c.hansman@imperial.ac.uk, hjort@columbia.edu, gianmarco.leon@upf.edu, mrt2138@columbia.edu. We are especially grateful to Francesco Amodio, Bob Gibbons, Lorenzo Casaburi, Eric Verhoogen, and Chris Woodruff for detailed feedback. We also thank Charles Angelucci, David Atkin, Jie Bai, Jing Cai, Wouter Dessein, Dave Donaldson, Marina Halac, Oliver Hart, Amit Khandelwal, Rocco Macchiavello, Ameet Morjaria, Kaivan Munshi, Andrea Prat, John Van Reenen, Xiao Yu Wang, and Daniel Wolfenzon for very helpful comments and suggestions; seminar participants at Cambridge, Columbia, LSE, MIT, NBER, UPF, Zurich, the Markets and Organizations in Emerging Economies Conference, the Society of Institutional and Organizational Economics, the Summer School in Development Economics (SSDEV), and the IGC/Stanford SCID conference on Firms, Trade and Development; and Patrick Kennedy and Leonor Lamas for great research assistance and field work. The Peruvian Ministry of Production, Cesar Casahuamán, and IFFO (in particular Enrico Bachis and Jorge Mora) kindly shared access to data. León 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 grant ECO2011-25272.

1 Introduction

Why do so many of our economic transactions occur within firm boundaries (Coase, 1937; Gibbons, 2005a; Lafontaine & Slade, 2007)? Vertical integration occurs for many different reasons, and motives vary by context.[1] However, one potential motive—integrating in order to raise output quality—has gained increased relevance in recent decades. As global incomes rise and barriers to trade fall, many firms experience greater access to wealthier markets. With opportunities to sell to high willingness-to-pay buyers tend to come rising returns to quality upgrading.[2] However, to produce high quality output, firms need high quality inputs (Goldberg *et al.* , 2010; Kugler & Verhoogen, 2012; Halpern *et al.* , 2015; Bastos *et al.* , 2018; Amodio & Martinez-Carrasco, 2018), and input quality is hard to measure and contract over (Gibbons, 2005a; Lafontaine & Slade, 2007)—especially where institutions are weak (Woodruff, 2002; Nunn, 2007). Firms' organizational structure may therefore play a crucial role in their ability to meet demand for quality.

In this paper we test the hypothesis that firms vertically integrate in order to produce higher quality products. This hypothesis is inspired by classical theories characterizing how firm boundaries are expected to respond to output objectives (Baker *et al.* , 2001, 2002; Gibbons, 2005a,b) when suppliers multitask (Holmstrom & Milgrom, 1991). However, causal evidence on the extent to which—and reasons why—firms change their organizational structure to shift from producing one type of product to another has remained elusive.

The context we study in this paper enables progress. In the Peruvian fishmeal industry, 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] Unique data on the industry's entire chain of production, transactions, and output quality grades is available. In addition, fluctuations in other countries' demand for and supply of high versus low quality grades generate season-to-season variation in individual firms' incentive to quality upgrade.

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 if 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 who supply the same plant before and after being integrated (or sold). Finally, we investigate whether integration ultimately raises output quality.

Our model demonstrates how characteristics of the Peruvian fishmeal industry map directly to the existing theoretical work we build on. Fishmeal manufacturers face contracting challenges: the quality of the product's primary input—fish—is perishable and very difficult to observe and contract on, and the presence of outside options—other fishmeal manufacturers who may value input quality less—complicates controlling the incentives of an independent supplier (see also McMillan & Woodruff, 1999).[4] Holmstrom

[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 & Slade (2007); Bresnahan & Levin (2012) provide excellent overviews of the broader literature on firms' structure.

[2] See e.g. Sutton (1991, 1998); Hallak (2006); Verhoogen (2008); Manova & Zhang (2012).

[3] Fishmeal is a brown powder made by burning or steaming fish, and mostly used as animal feed. Peru's fishmeal industry—one of Latin America's largest industries—accounts for around 3 percent of GDP (Paredes & 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 Sub-section 2.2).

[4] It is in theory possible to imperfectly measure fish quality with chemical tests. As discussed in Section 2, industry insiders informed us that such tests were much too expensive and impractical to use during our data period. Alternatively, manufacturers could attempt

& Milgrom (1991) elegantly demonstrate how, when suppliers face a trade-off between producing inputs of 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 & Tirole, 1991; Holmstrom, 1999). In a variety of situations, the best or only way to do so may be to bring the suppliers inside the firm (Baker *et al.* , 2001, 2002; Gibbons, 2005a,b).

Testing these textbook theoretical insights empirically requires unusual forms of data and variation. To establish the existence of a relationship between the quality of a firm's output and the structure of its production process, appropriate measures of each dimension of output quality and firm-supplier transactions are needed.[5] Even if documented with ideal data, such a relationship may reflect third factors rather than an explicit organizational choice made *in order to* "climb" the quality ladder. It could be, for example, that productivity or demand affect both firms' choice of structure (Hortacsu & Syverson, 2007; Legros & Newman, 2013; Alfaro *et al.* , 2016) 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 (Ordover *et al.* , 1990; Hortacsu & Syverson, 2007; Macchiavello & Miquel-Florensa, 2016; Martinez-Carrasco, 2017)—but in the process coincidentally produce higher quality output. To isolate 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.

In the setting we study, the flow of goods can be tracked from suppliers, to manufacturers, to foreign buyers. This is because all transactions between fishmeal plants and suppliers—including within-firm transactions—are recorded in data from regulatory authorities, and all export transactions are recorded in customs data. Output quality is directly observed because firms are required to report each of their plants' production of fishmeal in the "prime" (high) quality and the "fair average" (low) quality range each month. Finally, researchers—but not manufacturers—directly observe both independent and integrated suppliers' behavior because fishing boats are required to transmit GPS signals to the regulatory authorities.[6]

Peru's fishmeal industry also features a suitably high frequency type of variation that allows us to causally test our hypothesis. Fluctuations in other countries' demand for and supply of high versus low quality grades generate season-to-season fluctuations in the quality premium. These fluctuations affect firms differently depending on their scope for upgrading quality and their export destinations.

We begin our empirical analysis by demonstrating 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 the average fine-grained quality grade of a firm's output, as inferred in exports and auxiliary price data. Our primary measure of integration is the fraction of inputs that are sourced from integrated suppliers ("Share VI")—a measure that is closely tied to our model, and to the hypothesis that integration and output quality are causally linked.[7]

to infer the range of input quality that could have resulted in a given level of output quality observed *post-production*. The range would be wide—input from multiple suppliers is, for technological reasons, typically used in a given batch of fishmeal, and other hard-to-measure factors also influence output quality realizations—but such inference may not be impossible. However, in the Peruvian contracting environment, punishing or rewarding a supplier for input quality inferred through such a strategy would require repeated interactions. Whether repeated interactions between formally independent parties suffice to quality upgrade through such a strategy—given the temptation to renege when alternative trade partners exist—is an empirical question. We address this question below.

[5]See Goldberg & Pavcnik (2007); Khandelwal (2010); Hallak & Schott (2011) for discussion of the indirect quality proxies used in the existing literature, which risk conflating quality with mark-ups and horizontal differentiation, and Atkin *et al.* (2017) for an example of direct measures of quality.

[6]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.

[7]If integration directly affects output quality, it presumably does so through firms' *production* process—that is, firms' *use* of inte-

Alternatively, we also consider the number of suppliers owned. The relationship we establish between a firm's output quality and its organizational structure at a given point in time holds across each of these measures.

Our first of three key pieces of evidence—and this paper's main finding—shows that vertical integration is used by firms *as a strategy* for increasing output quality. To demonstrate this causally, we construct an instrument for firms' incentives to upgrade quality. When the quality premium—the difference between the price of high and low quality grades—increases, firms producing mostly low quality output face strong incentives to shift low to high quality production. Firms already producing mostly high quality output face less strong incentives to do so. The opposite holds when the quality premium falls. We instrument the quality premium in Peru with either prices in or quantities produced by exporter countries (i) that each produce almost exclusively high or low quality grade fishmeal, and where (ii) the country's total production is determined by its regulatory fishing quota. We find that Peruvian manufacturers integrate when their incentives to upgrade quality rise, and vice versa.[8] In an alternative approach shown in the Appendix, 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 *do not* integrate when faced with income shocks and greater incentives to expand production of low or any-quality output—that is, higher average prices.

This first set 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. Our results indicate that quality upgrading itself is an important motive for integrating suppliers.[9]

Next we explore *why* firms use integration as a strategy for upgrading quality. Our second key piece of evidence demonstrates that getting integrated shifts suppliers' behavior towards quality-increasing actions. We proxy for actions that increase input quality—i.e., fish freshness (FAO, 1986)[10]—using GPS-based measures. We show that 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—both are shown in Appendix B—and with the fact that suppliers that de-integrate 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 sup-

grated versus independent suppliers (see also Breza & Liberman, 2017; Baker & Hubbard, 2003). Share VI is also a measure that is consistently defined for firms and individual plants, allowing us to make comparisons within firms later in the paper. It is important to note that Peruvian fishmeal manufacturers can generally increase the total amount of inputs they obtain from integrated suppliers in a given production season only by acquiring suppliers. The reason is that boats' total seasonal catch is governed by a quota, and boats almost always exhaust their quota over the course of a season. See Section 2 for details.

[8]The long-term trend is towards more vertical integration in the Peruvian fishmeal 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 de-integration during our data period—sales of boats from fishmeal firms to independent co-ops or captains, and from one fishmeal firm to another.

[9]Several of the most prominent alternative explanations are also inconsistent with other features of the context we study and auxiliary findings. We show, for example, that 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. Similarly, Peruvian fishmeal manufacturers appear to achieve supply assurance primarily through repeated interactions with independent suppliers (Martinez-Carrasco, 2017); the quantity supplied by a given supplier to a given firm is *lower* after integration (/before de-integration).

[10]"Freshness of raw material is important in its effect on the quality of the protein in the end product [fishmeal]. 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, sub-section 10.1.2).

plier's quantity-quality trade-off, but instead via associated knowledge transfers (Atalay *et al.* , 2014). We show 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*. Though we ultimately cannot rule out that other incentives emanating from organizational structure itself also help explain the impact of integration on supplier behavior, reconciling this finding with pure knowledge transfer theories is challenging.

Our third piece of evidence suggests that vertical integration in fact increases output quality. 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 level of individual plants. The explanation must thus operate at *sub-firm* level.[11] More direct evidence comes from instrumenting for a plant's Share VI using a leave-firm-supplied-out measure of the local presence of a particular type of independent supplier that is regulatorily prohibited from being integrated. The presence of such suppliers fluctuates due to natural variation in fish density, weather, and decisions made by their captains. The logic behind the instrument 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.[12] The IV estimates are very similar to OLS estimates.

When viewed through the lens of our model, the three pieces of evidence we present each follow logically from each other. We conclude that firms vertically integrate in order to produce higher quality products, and that the reason they do so appears to be that integration changes supplier behavior in a way that increases output quality.

A natural question is whether our results are specific to Peru's fishmeal industry. However, 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.[13] In the concluding section of the paper, we document a positive relationship between (a proxy for) the average quality of a given type of manufacturing product a country exports to the U.S. and the average degree of vertical integration among the exporters. This provides suggestive evidence that the relationship between firms' output quality and organizational structure we establish in this paper may hold also in other export manufacturing industries.

Our study bridges and advances the literatures on the boundaries of the firm and quality upgrading.[14] We make three contributions to the former. First, we identify an overlooked motivation for vertical integration. By showing that firms vertically integrate in order to raise output quality, we advance the body of work on the *causes* of organizational form (for pioneering empirical work, see Hart *et al.* (1997); Baker & Hubbard (2003, 2004); Forbes & Lederman (2009)). Existing studies convincingly demonstrate how firms

[11] The fact that the firm×season level relationship between Share VI and output quality holds also across plant×month observations challenges theories that explain the former by reference to firm level shocks—for example to demand—that are conjectured to affect both acquisitions/sales of suppliers and output quality without the two being causally related.

[12] In the first stage we find, as expected, that greater presence of independent-by-law suppliers (supplying other firms) makes a given firm more likely to use such suppliers. We also show that, beyond the presence of independent suppliers, there is no time-varying, geographical location level component of output quality. In combination with the estimated sign on the second stage—greater relative presence of independent suppliers *lowers* output quality—these two observations are hard to reconcile with plausible arguments against the exclusion restriction underlying our interpretation of the IV results. The IV results are very similar if we instrument with the total number of independent suppliers, rather than just the subset that is independent by law.

[13] There is a robust relationship between countries' input-output structure and their level of contract enforcement (Nunn, 2007; Boehm, 2016), and vertical integration is more common in developing countries (Acemoglu *et al.* , 2009; Macchiavello, 2011).

[14] Gibbons (2005a); Lafontaine & Slade (2007); Bresnahan & Levin (2012) provide great overviews of the literature on the boundaries of the firm. There is also a prominent literature studying the relationship between integration and international trade (see e.g. McLaren, 2000; Grossman & Helpman, 2002; Antràs, 2003; Nunn, 2007; Antràs & Staiger, 2012; Irarrazabal *et al.* , 2013; Antràs, 2014, 2016; Ramondo *et al.* , 2016). Our focus is on firms' domestic organizational structure.

change their relative use of integrated suppliers in response to changes in e.g. available contracts (Breza & Liberman, 2017) or monitoring technology (Baker & Hubbard, 2003). We instead study how firms change their organizational structure when their output *objectives* change.[15]

Second, and building on earlier evidence on the behavior of integrated and independent suppliers (Mullainathan & Sharfstein, 2001; Baker & Hubbard, 2003, 2004; Macchiavello & Miquel-Florensa, 2016), we provide what to our knowledge is the first evidence on how integration changes the quality-oriented behavior of a given supplier supplying a given firm.[16]

Finally, we show that vertically integrating raises output quality, which to our knowledge has not been done before. The simple, one-dimensional, nature of quality differentiation in our setting allows us to document this.[17] In general, there is little existing evidence on causal consequences of organizational structure for firm performance (see Gibbons & Roberts (2013), and Forbes & Lederman (2010) for a notable exception). Our results also imply that manufacturers often most effectively succeed at producing output in high volumes *rather than* of high quality by using independent suppliers. An especially unusual aspect of this paper is that the data and variation we exploit allow us to identify *both* the effectiveness of particular firm strategy and determinants of its use. We can therefore show that Peruvian fishmeal manufacturers vertically integrate when quality objectives indicate that they *should* do so.[18]

Both the friction—imperfect contracting over input quality—and the firm objective—producing high quality output—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 (2002)'s landmark study (Natividad, 2014; Macchiavello & Miquel-Florensa, 2016; Martinez-Carrasco, 2017).[19]

The literature on quality upgrading is larger. It is now well-documented that producers of high quality goods use high quality inputs (Goldberg *et al.*, 2010; Kugler & Verhoogen, 2012; Halpern *et al.*, 2015; Bastos *et al.*, 2018; Amodio & Martinez-Carrasco, 2018), more skilled workers (Verhoogen, 2008; Frías *et al.*, 2009; Brambilla *et al.*, 2012; Brambilla & Porto, 2016; Brambilla *et al.*, 2016), and export to richer destinations (Hallak, 2006; Verhoogen, 2008; Manova & Zhang, 2012; Bastos *et al.*, 2018). Firms with such a profile tend on average to be bigger, more productive, based in richer countries themselves, and to face foreign competition in low-quality segments (Schott, 2004; Hummels & Klenow, 2005; Baldwin & Harrigan, 2011; Johnson, 2012; Medina, 2017). We provide the first evidence linking quality upgrading to the boundaries of the firm.

[15] In a superficial sense, our finding that higher average fishmeal prices do not lead to more integration in the Peruvian fishmeal 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. However, we highlight the fact that this efficiency gain is not generic, but rather depends on firms' output quality objectives.

[16] We follow Atalay *et al.* (2014) in exploiting changes in integration within supplier-firm pairs. They focus on transfers of knowledge.

[17] In settings where product differentiation is multidimensional, an analysis like ours would be difficult. Like us, the innovative study by Forbes & Lederman (2010) exploits exogenous drivers of use of integrated suppliers, showing that routes airlines self-manage have fewer delays/cancellations (see also Gil *et al.*, 2016; Gil & Kim, 2016). Other important evidence on the *consequences* of organizational structure includes, among others, Novak & Stern (2008); Gil (2009); Kosová *et al.* (2013).

[18] Almost all existing studies of firms restrict attention to *either* the effectiveness of a strategy *or* the determinants of its use.

[19] Woodruff finds that forward integration is less common in the Mexican footwear industry when non-contractible investment by retailers is important, as the property rights framework predicts (Grossman & Hart, 1986; Hart & Moore, 1990). Macchiavello & 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 & Duflo, 2000; Macchiavello & Morjaria, 2015). We follow Natividad (2014) in studying organizational structure in the Peruvian fishmeal 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 *et al.*, 2011), which lead to an "Olympic race" for fish.

2 Background on Peru's Fishmeal Manufacturing Sector

In this section we provide an overview of Peru's fishmeal 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.

2.1 Sector profile

Fishmeal 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 percent of the world's fishmeal exports. During our data period, 2009 to 2016, around 95 percent of the country's total fishmeal production was exported. The three largest buyers are China, Germany, and Japan, but many other countries also import Peruvian fishmeal (see Appendix Table A1).

Fishmeal 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. 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 Appendix Figure A1, firm size correlates positively with average quality grade produced.

Plants receive inputs of raw fish from their suppliers. The suppliers may be large steel boats—which may be independent or owned by the firm that owns the plant—or smaller wooden boats. Regulations prohibit fishmeal firms from owning wooden boats. There are on average 812 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 (s.d. = 10 hours) and boats travel 76 kilometers away from the port of delivery (s.d. = 46 kilometers) on average. Changes in installed technology are observed in our data but rare both for boats and plants. Table 1 shows summary statistics, providing further detail on the sector.

2.2 Product differentiation and quality

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

Fishmeal's protein content depends crucially on input characteristics, namely the freshness and integrity of the raw fish that boats deliver (see e.g. FAO, 1986). Freshness at the time of delivery in turn depends on choices made by the boat's captain before and during a trip, such as the amount of ice brought on board, how tightly fish is 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 is processed as soon as possible after offload.

However, 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 such tests were too costly and timeconsuming to be of practical use in Peru during our data period. In addition to the fixed cost of adoption, this was due to the importance of processing fish immediately after offload.[20]

After offload, the fish is weighed, cleaned, and converted to fishmeal using one of two technologies: steam drying (hereinafter "High technology") or exposing the fish directly to heat (hereinafter "Low technology"). The technology used can matter for the protein content achieved.

Peru allows anchovy fishing for fishmeal production during two seasons each year and because of the need for fresh fish, fishmeal plants operate only during the fishing seasons. There were thus 14 fishing and fishmeal production seasons during our 2009-2016 study period. In theory fishmeal 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 Appendix Figure A2 and discussed below.

2.3 Organizational structure

Consistent with our hypothesis, both integration and average output quality have slowly increased over time. However, these long-term trends are not the source of the relationship between organizational structure and quality upgrading 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. On average, 28 percent of the boats that are active in a given season are integrated with a fishmeal firm. As seen in Panel A of Table 2, we observe 317 instances where ownership of a steel boat changes hands. In 103 of these instances, a fishmeal 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 fishmeal firm to an independent buyer, and 50 instances where a supplier is sold from one fishmeal firm to another.

In our data, we observe not only supplier *ownership* but also *deliveries* from integrated versus 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 fishmeal manufacturers' Share VI is on average 45 percent. Firms can generally increase or decrease the total amount of inputs that come from integrated suppliers only by buying or selling boats. The reason is that, in the regulatory system in place during the period we study, a boat's total catch in a season is governed by a quota, and each boat typically exhausts its quota. Of course, a firm may vary its Share VI also by increasing or decreasing its use of independent suppliers. As seen in Appendix Figure A3, and following the trend in ownership, Share VI slowly increased during our data period—by 2.9 percent from season to season. Approximately 77 percent of this growth came solely from increasing the amount of input coming from integrated suppliers, and the rest from lower total input purchases, as shown in Panel B of Table 2.

Crucially for our purposes, Share VI can be defined not just 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 different 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

[20] Note also that input from multiple suppliers is, for technological reasons, typically used in a given batch of fishmeal, and other hard-to-measure factors also influence output quality realizations. This complicates attempts at inferring input quality from output quality post-production.

varies considerably over time, and depends on variation in weather, fish density, and independent captains' decisions.

Figure 1 shows that integration and de-integration primarily represents a change in the formal status of the relationship of 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 bottom part of the figure restricts attention to the "switchers" we focus on in our empirical analysis of supplier behavior in Section 6. Suppliers that get integrated or sold deliver to the plant (within the acquiring/selling firm firm) they deliver to most frequently around 41 percent of the time *when independent* (i.e. before getting acquired or after getting sold), and around 45 percent of the time when integrated. Similarly, switchers deliver to the acquiring/selling firm around 63 percent of the time when independent and around 81 percent of the time when integrated.[21]

2.4 Contracting and supplier incentives

There is no centralized spot market for fish purchases: plants are spread out along the coast, in part because the fish move around. Where a boat makes a catch thus constrains the set of ports it can deliver to. 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 fishmeal industry associations, a major company's Chief Operating Officer, and others in the sector to gain a qualitative understanding of the characteristics of the contracts used and the incentives suppliers face. The interviewees reported that captains of boats owned by fishmeal firms generally are paid a fixed wage, in some cases with a bonus tied to some measure of performance.[22]

On the other hand, the interviewees reported that payments to independent suppliers—while agreed upon 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 percent of the price variation across transactions. We are not aware of formal contracts between independent suppliers and firms over when and where to deliver fish.

Our data on suppliers' behavior—discussed in Section 4—come from a map the regulators construct and 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. This is the primary reason why manufacturers and independent suppliers cannot contract over GPS-measured actions.

[21] The top part of the figure displays averages for all integrated and all independent suppliers. Integrated suppliers deliver to the firm they deliver to most often (i.e., the parent firm) about 90 percent of their trips, and the plant they deliver to most often 38 percent of their trips. Independent suppliers deliver to the firm they deliver to most often around 65 percent of their trips, and the plant they deliver to most often 45 percent of their trips.

[22] The fishmeal industry associations reported that payment schemes vary across firms; that some pay bonuses tied to measures of performance; but that these are on top of a fixed wage and usually small.

3 Theoretical Framework

3.1 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 tradeoff 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 & 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 over-incentivize 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 or 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 to understand the second point of intuition above. Firstly, it appears that contracts over when and where to deliver a catch are difficult to write ex-ante: independent suppliers retain their right to deliver their fish where they choose. Additionally, while some firms primarily produce high protein content fishmeal, others primarily produce low quality grades, and hence provide a (presumably less quality sensitive) alternative for suppliers to deliver their catch.[23]

With this in mind, a logic applies that is familiar from the models presented in Baker *et al.* (2001, 2002), based on the notion of integration as asset ownership that follows Grossman & 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 & Milgrom (1994), and Gibbons (2005a), among others.

3.2 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

[23] A question that our model abstracts from is why firms might want to produce different quality levels simultaneously. We return to this question in Sub-section 4.2.

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 $c(a_1, a_2) = \frac{1}{2}a_1^2 + \frac{1}{2}a_2^2$. These actions impact the surplus created 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 $\gamma, \delta \geq 0$.[24] 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. Q_0 is a baseline level of quality surplus.[25] 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: $\tilde{Q} = Q_0 - \gamma a_1 = Q - \delta a_2$.

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.

Independent suppliers

Although neither Q nor P is contractible[26], 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. 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:

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

[24] 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.

[25] 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.

[26] Alternatively, we could assume that only a portion of Q and P is non-contractible, and that we consider only this portion as in Baker *et al.* (2002).

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 valuable to note 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 impacts 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.

3.3 Discussion

The theoretical role of vertical integration is a contentious topic. Our framing follows Baker *et al.* (2001, 2002) in combining elements of the incentives based theories in the tradition Holmstrom & Milgrom (1991) and the property rights theories in the vein of Grossman & Hart (1986). Such a framing is not the only type of model that would produce a relationship between integration and output quality.[27] 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 our primary story, 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 in order 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 & 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 B, 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

[27]For example, integration might facilitate an efficient transfer of intangible inputs—such as information on the location of fish or fishing techniques—as in Hortacsu & Syverson (2007). Alternatively, adaptation may play an important role in integration, if high quality production requires a higher degree of strategic response to varying production conditions. We test whether knowledge transfer can help explain our findings in Section 6. As adaptation has seen a fair amount of focus in the empirical literature (see e.g. Forbes & Lederman, 2009), and because there is a close connection between adaptation models and the dynamic version of our model in Appendix B, we conduct a detailed exploration of the relationship between integration and the adaptation of firms and suppliers to exogenous changes in production conditions in Appendix D.

inputs are processed and output quality is measurable). The firm can then offer the supplier a (delayed) reward contingent on this surplus, but can only credibly promise to pay this reward if it interacts repeatedly 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 baseline 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. $P >> Q$), provides precisely this result.

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—for example due to increased demand for high quality grades—leads to more integration.

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.

4 Data, Variables, and the Relationship of Interest

4.1 Data

The primary datasets we use are the following:

Plant production. Administrative data on all plants' production come from Peru's Ministry of Production, which regulates the fishmeal industry. Every month plants are required to submit information on how much prime (high quality) and fair average (low quality) fishmeal they produce. Quality grade is thus directly reported in the plant production data, and subject to auditing by government inspectors. As discussed in Sub-section 2.2, the distinction between prime and fair average quality is based on the fishmeal'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 multi-plant fishmeal firm.[28] We also use this registry to link the production data to export data. We can do so for almost all firms, but not the smallest firms, which use intermediaries to export.

[28]The data contains information on the number of metric tons that can be produced per hour with respectively the installed Low and High technology. As very few firms in our sample only have 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 fishmeal 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 fishmeal, the value of the transaction, and the exporting firm.

Internal data from a large firm. One of the largest fishmeal 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, destination country, date of the contract, and the terms. Most importantly for our purposes, the specific plant that produced a given shipment of fishmeal is reported.

Supply transactions. The Ministry of Production records all transactions between the fishmeal plants and their suppliers of raw materials, i.e. 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.[29]

Boat GPS data. Peruvian fishing boats that supply fishmeal 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 dataset with us.[30]

4.2 Variables of interest

Our primary measure of an individual plant's output quality is the share of the fishmeal 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 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 output 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 fishmeal prices in detailed (week×export port×protein content level) data recorded by a fishmeal consulting company. We infer the protein content of each of a firm's export shipments by comparing the corresponding unit values to this price data. To construct a firm×season level measure, we average protein content across export shipments, weighting by quantity.[31] A priori, we have little reason to

[29] Information on engine power is only available for 2004-2006. However, changes in engine power are extremely rare in that period, so we treat this characteristic as fixed over time.

[30] Only about half of the observations in the Supply transactions dataset can be matched to a GPS recording. Some boat owners, for example, disappear from the GPS data for a complete calendar year. However, such missingness is unlikely to be of concern for within-boat analysis, the level at which we use the GPS data.

[31] The export transaction records do not report the specific plant that made the fishmeal so the inferred quality grade is only available at the firm level—except for data covering the fishmeal firm that shared internal data with us, including information on the plant that produced a given export shipment. One potential concern is that fishmeal 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 percent of total season production (see Appendix Figure A2)—is likely that many contracts are entered into before the production season starts (which helps the fishmeal 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 month level would require an assumption about how firms manage their inventories—for example, first-in-first-out versus first-in-last-out). A related concern is that firms that are about to end operations and close down might sell off their fishmeal, in which case a lower unit price might not reflect lower quality but rather a "going-out-of-business" discount. We thus

believe that this inferred protein content measure could be systematically biased.[32] 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 records of a firm that shared its data with us.

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") (see also Breza & Liberman, 2017; Baker & Hubbard, 2003). Recall from Sub-section 2.3 that firms generally use integrated suppliers to capacity over the course of a season, so the two measures are closely related.

Share VI is 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 upgrade quality, then it should matter not just if a firm owns suppliers, but the degree to which it sources inputs from integrated versus independent suppliers at the time of production. This observation holds both for firms and individual plants. We can construct Share VI in a consistent manner for firms and plants because we observe all transactions between *plants* and suppliers, allowing us to make within-firm comparisons in Section 7.[33] In addition, there is a useful link between Share VI and the intuition captured in the model in Section 3. In this 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 non-integrated suppliers, but choose an intermediate Share VI that depends on the relative importance of Q and P in the firm's *current* objective function.

4.3 Relationship of interest

In Section 5 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:

$$\text{Quality}_{it} = \alpha + \beta_1 \text{VI}_{it} + \beta_2 \text{HighTech}_{it} + \gamma_i + \delta_t + \varepsilon_{it} \tag{1}$$

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 how advanced the technology the firm uses to convert fish into fishmeal is[34], HighTech_{it}, and include firm 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 towards a strong baseline relationship between owning suppliers

exclude data from any firm×season observations that correspond to a firm's last season to produce and export fishmeal, 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.

[32] Fishmeal 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 Sub-section 2.2). This implies that pricing-to-market, bulk discounts, etc, are not a concern.

[33] The Share VI measure also automatically captures suppliers' size so we avoid making assumptions on "scale" effects (e.g., how the benefit of one big integrated supplier compares to two small ones).

[34] A firm's production technology is an important 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 output quality. The estimates in column 4 imply, for example, that a doubling of the number of suppliers owned is associated with an increase in protein content of about 3.4 percent of the range observed in Peru.[35]

In Panel B we show that, beyond simply owning suppliers, what matters for output quality is *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 50 percent higher, and an average protein content that is higher by about 20 percent of the range observed in Peru.

In Panel C we show that these results are not not driven by observable, time-varying supplier or firm characteristics. We control for the firm's share of total industry production and a series of supplier characteristics. This has little impact on the estimated coefficient.[36]

In Appendix B, 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 established for Share VI. In other words, it is integration itself, not the relationship, that co-varies with output quality. This is consistent with the predictions of a dynamic version of our model, also shown in Appendix B.

The relationship between a firm's organizational structure and its output quality documented in this sub-section is consistent with our hypothesis. However, this basic relationship is also consistent with the alternative theories discussed in the introduction, wherein the relationship between organizational structure and output quality is either not causal or not known to (or ignored by) firms. We rule out these alternative explanations in the next section.

5 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 firms make *in order to* "climb" the quality ladder. Specifically, Peruvian fishmeal manufacturers integrate suppliers when the returns they earn from upgrading quality rises for exogenous reasons. This finding provides empirical support for the prediction that a vertically integrated organizational structure is efficient for producing high quality output. It is difficult to reconcile with alternative theories.

5.1 Estimating how the overall price level co-varies with vertical integration

In the first column of Table 4 we show results from regressing the change in integration, $\text{VI}_{it} - \text{VI}_{it-1}$, on the average price in season t. We measure VI_{it} as Share VI. $\text{VI}_{it} - \text{VI}_{it-1}$ is thus the change from season $t-1$ to t in the share of inputs a firm obtains from integrated suppliers. Log(Average Price) is simply the average

[35] The range of protein content observed in Peru is approximately 63-68 percent.

[36] 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 share of inputs coming from VI suppliers essentially unchanged. Note that two of the supplier characteristics variables included—Share of inputs from high capacity boats and Share of inputs from boats with 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 also leaves the magnitude and significance of the coefficient on Share VI essentially unchanged.

fishmeal price in season t, i.e., the average price of high and low quality fishmeal, weighted by the overall market share of each. We control for firm fixed effects and cluster the standard errors at firm level.

The results suggest that vertical integration is not driven by pure income shocks or increases in the returns to producing fishmeal generally. Firms are not more likely to vertically integrate when the overall price level rises.

5.2 Estimating how the quality price premium affects vertical integration

To estimate how the quality premium—the *difference* between the price of high and low quality fishmeal—affects integration decisions, we first repeat the regression from the previous sub-section, now with the quality premium replacing the average price as the regressor of interest. The results in the second column of Table 4 demonstrate that firms overall *do* vertically integrate when the quality premium rises and de-integrate when the quality premium falls. This finding by itself does not allow causal conclusions—variation in the quality premium could correlate with other important factors that influence firms' organizational choices. But the *industry*-wide relationship between integration decisions and the quality premium established in column 2 underscores the general plausibility of the *firm level* causal relationship we establish below.

To isolate firm-specific incentives to upgrade quality, we categorize firms based on their *upgradeable share of production*, or the share of their production that is currently low quality and hence could potentially be upgraded—$\text{UpgradableShareOfProduction}_{it-1}$. Our conjecture is that firms that produce mostly low quality output should face a greater increase in their incentive to integrate than firms who already produce mostly high quality output when the quality premium rises. To estimate how integration decisions respond to changes in firm-specific incentives to upgrade quality, we thus run the following specification:

$$\begin{aligned} \text{VI}_{it} - \text{VI}_{it-1} = \alpha &+ \beta_1 \text{QualityUpgradingIncentives}_{it} \\ &+ \beta_2 \text{UpgradableShareOfProduction}_{it-1} + \gamma_i + \delta_t + \varepsilon_{it} \end{aligned} \tag{2}$$

where $\text{QualityUpgradingIncentives}_{it}$ is $\text{UpgradableShareOfProduction}_{it-1} \times \text{QualityPremium}_t$. Our hypothesis implies that $\beta_1 > 0$. We control for $\text{UpgradableShareOfProduction}_{it-1}$ itself and firm and production season fixed effects γ_i and δ_t and cluster the standard errors at firm level.

Our approach in (2) is a generalized difference-in-differences in which firms that are more versus less exposed to changes in quality upgrading incentives due to a high quality premium are compared in each of 13 different production seasons, and in each of these 13 seasons the quality premium may be relatively high or relatively low.[37] Additionally, the characteristic that defines a firm's exposure at a given point in time—$\text{UpgradableShareOfProduction}_{it-1}$—is itself controlled for.[38]

The results in column 3 of Table 4 show that firms with greater scope to shift from low to high quality production are more likely to vertically integrate when the quality premium rises, and vice versa when the quality premium falls—consistent with our hypothesis and the model in Section 3.

A potential concern is the possibility that the quality premium could be endogenous to individual Peruvian firms' output quality objectives and organizational structure. To address this concern, we first instrument the quality premium in Peru with fishmeal prices in other countries that produce only high quality

[37] In addition, firms' upgradeable share of production evolves over time so the high and low $\text{UpgradableShareOfProduction}_{it-1}$ firms being compared across one rise or fall in the quality premium may differ from those being compared across another rise or fall.

[38] This distinguishes our approach from traditional Bartik instrument approaches (Goldsmith-Pinkham *et al.*, 2017).

fishmeal—Chile, Denmark, and Iceland—using the following first stage:

$$\text{QualityUpgradingIncentives}_{it} = \text{UpgradableShareOfProduction}_{it-1} \sum_c \beta_c \text{ HighQualityPrice}_{ct} \qquad (3)$$
$$+ \beta_2 \text{UpgradableShareOfProduction}_{it-1} + \gamma_i + \delta_t + \varepsilon_{it}$$

where c is a high quality exporter country, and $\text{UpgradableShareOfProduction}_{it-1} \sum_c \beta_c \text{ HighQualityPrice}_{ct}$ are our excluded instruments. $\text{HighQualityPrice}_{ct}$ is the average price charged for high quality fishmeal exported from country c in season t.

High quality fishmeal prices in these three countries are highly correlated with those in Peru, and the first stage results—shown in Appendix Table A3—are thus strong.[39] The time series of the four prices are shown in Figure 3, which also highlights two other important facts. First, the quality premium in Peru is at most weakly correlated with average prices in Peru. It is thus not surprising that firms respond differently to the two, as columns 1 and 2 in Table 4 suggest. Second, while the long-term trend in the quality premium during our data period is positive, the quality premium fluctuates substantially from season to season, sometimes rising and other times falling.

The second stage results are shown in column 4 of Table 4. The IV estimate of $\hat{\beta}_1$ is of similar magnitude to the OLS estimate. The estimate implies that, when the quality premium rises by 10 percent, a firm with a high upgradeable share of production—one that produces only low quality output—increases its Share VI by 30 percent when compared to a firm producing only high quality output.

In theory, even prices in other countries could depend on actions taken by individual fishmeal firms in Peru. *Quantities* exported should not do so, however. The reason is that, in the countries we focus on, the total amount of fishmeal exported is a direct function of the aggregate fishing quota set by each country's regulatory authorities. In an alternative approach, we thus instrument the quality premium in Peru with quantities exported by the other four of the world's top five fishmeal exporters: Chile, Denmark, Iceland, and Thailand. We replace $\text{HighQualityPrice}_{ct}$ in (3) with HighQualityExports and LowQualityExports. The variation in HighQualityExports that we isolate is driven by Chile, Denmark, and Iceland,while the variation in LowQualityExports is driven by Thailand, which exports almost no high quality fishmeal.[40]

The second stage results from this approach are shown in column 5 of Table 4. The quantity-based IV estimate of $\hat{\beta}_1$ is of similar magnitude to that found in columns 3 and 4.

In the last column of Table 4, we show that firms' estimated response to $\text{QualityUpgradingIncentives}_{it}$ does not mask an underlying response to *firm-specific* average prices. We construct the firm-specific average price by weighting low and high quality grade prices by the firm's $t-1$ output shares. The results in column 6 underscore that firms do not vertically integrate in response to higher overall prices.

In the Appendix we exploit a different form of variation and find results consistent with those discussed here. 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 demand for high quality grades increases, and selling boats and decreasing Share VI when demand for high quality grades decreases. To do so we construct instruments for firm-specific

[39] On the other hand, high quality prices are not exactly the same across countries. This is because production seasons differ and the distance between countries would make it difficult to conduct arbitrage.

[40] Information on the price of (high quality) fishmeal in Chile, Denmark, and Iceland comes from IFFO, an industry association. Information on prices in Thailand was not available. Information on total exports from each country is available from BACI (see footnote 53). Conveniently, these other four of the world's top five exporters all specialize in either high or low quality fishmeal.

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.[41]

5.3 Interpretation

The results discussed in this section are consistent with the theoretical framework in Section 3. In our model, a firm 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 (i) easier to measure and (ii) valued by other buyers in the market. Alternatively, the results we have shown so far may also be explained by other, closely related forms of quality-relevant incentives embedded in the formal relationship between a manufacturer and its suppliers. We return to this possibility in Section 6. We now consider whether firms' decision to integrate suppliers when the benefits of quality upgrading rise can be explained by theories in which the relationship between organizational structure and output quality is either not causal or not known to (or ignored by) firms.

A first possibility is that firms simultaneously choose their organizational structure and output quality, and shocks—for example to demand (Legros & 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 fishmeal 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 fishmeal.

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 (Ordover *et al.* , 1990; Hortacsu & Syverson, 2007). If such a story explained our results, we should see manufacturers integrating suppliers also when the price of low (or any) quality fishmeal 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 & Miquel-Florensa, 2016; Martinez-Carrasco, 2017). One supply assurance story—namely that firms integrate suppliers to secure general access to inputs but in the process coincidentally produce higher quality output—cannot explain our findings. Such a story is inconsistent with the fact that manufacturers vertically integrate only when the relative price of high quality output increases. Another form of supply assurance—integrating to secure access to suppliers who deliver the high quality inputs that are needed to meet the demand for high quality output—*is* exactly the interpretation we favor.

We conclude that manufacturers vertically integrate when the quality premium rises for exogenous reasons *in order to* produce a higher share of high quality output.

6 Firms' Organizational Structure and Supplier Behavior

The model in Section 3 predicts that firms adopt a more integrated organizational structure when attempting to shift from low to high quality production for a specific reason. In particular, we expect to see suppliers

[41] We follow many fruitful applications of such an approach in the trade literature (see e.g. Park *et al.* , 2009; Brambilla *et al.* , 2012; Bastos *et al.* , 2018; Tintelnot *et al.* , 2017).

reduce behavior that increases quantity but is harmful to quality when integrated.

6.1 Estimating how vertical integration affects suppliers' quality-enhancing actions

We focus on three measures of behavior that capture the tradeoff between input quantity and quality: the total quantity supplied, the maximum distance travelled from the delivery port, and the total time the supplier spends at sea on a given trip. The first of these three we observe in supply transactions data, while the second 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, not stacking fish high on top of each other, etc—in order to bring back a high quantity of fish per trip. The maximum distance travelled 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 fishmeal. 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 [fishmeal]. 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, sub-section 10.1.2). Captains must thus balance traveling further 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 post-integration (or increase post-separation).

Our empirical strategy focuses primarily on "switchers". Switchers are suppliers that are either bought or sold by a fishmeal firm during our data period and observed supplying 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 de-integration).

As discussed in Section 2, we observe 103 instances in which a fishmeal firm acquires a supplier that is initially owned independently; 32 instances where a supplier is sold from a fishmeal firm to an independent buyer; and 50 instances where a supplier is sold from one fishmeal 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. The reason is that integrated suppliers sometimes supply other fishmeal firms.[42] 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 to boat characteristics[43], 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 percent of the time *before*

[42] This is because a firm's output objectives may vary across time within seasons, and because fish move around and the location of a catch constrains the set of plants a boat can deliver to. As seen in Figure 1, Panel (a), integrated suppliers on average deliver to other firms just over 10 percent of the time.

[43] As shown in Appendix Table A2, the characteristics of integrated suppliers unsurprisingly differ from the characteristics of independent suppliers. On observable features such as the size of the boat, the power of its engine, and whether or not it has a cooling system installed, the average switcher falls in between the average always-independent boat and the average always-integrated boat, but closer to the latter.

integration (or after de-integration): integration typically implies a simple change in the formal status of an already frequently used supplier.

We estimate the following regression:

$$B_{ijt} = \alpha + \beta \text{I}[\text{VI} \times \text{supplies owner firm}]_{ijt} + \gamma_{ij} + \delta_t + \varepsilon_{ijt} \quad (4)$$

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

The results in Table 5 corroborate our model's predictions. Column 1 of Panel A shows that, when integrated and supplying a parent plant, a boat delivers on average about ten percent less per trip compared to 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 five percent closer to the port of delivery, and spend on average three percent less time at sea on a trip when integrated with the plant supplied. These results suggest that, when integrated, suppliers reduce costly actions associated with long trips, and bring back fresher fish as a result.

In our model, integration is defined by asset ownership, as in Grossman & 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 Appendix Table B1, 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 shown in Appendix B. Thus, while repeated interactions help fishmeal 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.

6.2 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 due to 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 plausibility of such a story explaining 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 fishmeal 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.[44]

As seen in Panel B of Table 5, we find quite similar—even slightly larger—effects as compared to Panel A. If acquired, the supplier changes its behavior consistent with prioritizing quantity less—to the benefit of quality—while delivering to the acquiring firm, similarly to how previously independent "switchers" change their behavior once integrated. This evidence suggests that a story in which integration enables knowledge transfer from Peruvian manufacturers to their suppliers is unlikely to be the primary 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 *et al.* , 2014).

The results in Panel B of Table 5 also underscore that it is not the case that firms simply choose to integrate suppliers that have already begun changing their behaviors, providing 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 incentives, and that the behaviors we see might not generate any input quality benefits that manufacturers are aware of and "act on". Such a story is difficult to reconcile with this paper's central finding that firms integrate suppliers when the quality premium rises.[45]

7 Vertical Integration and Output Quality

In Section 5 we saw that firms vertically integrate when the benefits of shifting from low to high quality production rise. In Section 6 we saw that suppliers that get integrated take more input quality-increasing and less input quantity-increasing actions. In this section we show that plants' *output* quality responds to integrating suppliers in exactly the manner we expect if the integration-induced change in supplier behavior improves input quality. This provides empirical support for the model's third prediction, namely that vertical integration is an *effective* organizational strategy for producing high quality output.

We first show that there is a robust relationship between changes over time in the share of inputs *individual plants* obtain from integrated suppliers—Share VI—and changes in their output quality that goes beyond the firm level evidence discussed in Sub-section 4.3. We then show evidence from an IV approach that exploits geographic variation in the local concentration of a particular type of supplier that is regulatorily prohibited from being integrated suggesting that the Share VI-output quality relationship arises because integration increases output quality.

7.1 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 firm level should hold at *plant* level also. This is what we find in Table 6. We repeat regression (1) from sub-section 4.3, but now at plant (i) × month (t) level, the lowest level at which we directly observe output quality.

[44] To implement, we run the same specification as in Equation 4, but define I[VI×supplies owner firm] to be equal to one if the supplier is (i) always owned by a fishmeal firm, and (ii) currently delivering to its parent firm.

[45] Additionally, such a story would raise a conceptual question: if there is no known input quality benefit, and integration lowers input quantity, then why integrate at all?

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.[46] 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 percent 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 Appendix Table A5. There we use internal data provided to us by a single major firm.[47]

In combination with Table 3, the first two columns of Table 6 establish a positive, statistically significant, and quantitatively consistent association between Share VI and directly observed output quality at firm and plant level across across months and production seasons. The explanation for our findings must thus operate at plant (sub-firm)$\times$month level. One possibility is that plant specific shocks, for example to productivity[48], occur and *independently* affect the quality of a plant's output and the share of the plant's supply coming from integrated suppliers.

To investigate this possibility, we construct an instrument for a plant's use of integrated suppliers at a particular point in time. To do so, we use the presence of wooden fishing boats—which are, by law, independently owned—as a source of variation in a plant's Share VI. 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 firms will tend to decrease their Share VI—even holding their incentives to produce quality constant. With this in mind, we consider the number of suppliers in a port that are independent by law as a potential proxy for the relative cost of using independent suppliers. 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. The logic 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.

Of course, a plant's quality objectives may themselves influence independent suppliers' whereabouts. The plant may for example request deliveries from independent 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 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. The first stage, shown in Appendix Table A4, is strong: the number of independent boats supplying other plants in the port is highly correlated with the number of independent boats supplying 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 towards 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 of the same sign, statistical significance, and general magnitude as the corresponding OLS estimates, only slightly bigger. This holds whether we restrict attention to suppliers that are independent by law or include all independent suppliers. Additionally, the same is true in Appendix Table A5, where we use internal data from the firm

[46] As discussed in Section 4, we observe whether plants have any high technology installed so HighTech_{it} is now a dummy variable.

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

[48] Another example of a shock that may affect different plants within a firm differently is El Niño, which hit Peru in late 2009.

that shared its data with us.[49]

Might the composition of neighboring plants' suppliers correlate with the quality of a given plant's output for other reasons than having comparable access to independent suppliers? A time-varying, *port level* component of output quality that correlates with our instrument for other reasons than independent suppliers' inputs lowering output quality is a possibility.[50] However, beyond the presence of independent suppliers, we find no evidence of a relationship between changes in output quality across different plants within the same port.[51] This, and our result that presence of independent suppliers *decreases* output quality—so the explanation is not e.g. that better fishing conditions differentially attract independent boats—suggests that our instrument's exclusion restriction holds.

In this section we began by documenting that the firm level relationship between inputs coming from suppliers that are integrated at the time of production and output quality holds also at plant level, including within firms. We then showed that instrumenting for Share VI yields the same positive, estimated relationship with the quality of a plant's output as OLS regressions. We conclude that a higher share of supplies coming from integrated suppliers at the time of production appears to increase output quality.

7.2 Interpretation

Our interpretation of the results in this section is that vertical integration directly increases output quality because a manufacturer can incentivize suppliers to engage in more quality-increasing behavior once the suppliers are integrated, as the model in Section 3 predicts, and consistent with the results in Section 6. A priori, output quality may of course co-vary with organizational structure without necessarily reflecting a causal relationship. Perhaps the most plausible non-causal links between quality upgrading and integration—for example, that growing firms both produce higher quality output and acquire more suppliers—are ruled out as explanations of our findings by the simple OLS regressions in Table 6: output quality correlates with *use* of integrated suppliers at the time of production across *plants*, including within firms. The IV regressions take a step further by documenting that the same relationship holds when we restrict attention to fluctuations in the use of integrated suppliers that is driven by variation in the local presence of independent suppliers.[52] These results are most plausibly explained by a model along the lines of the one we present in Section 3. That conclusion is implied by the results in Section 5. There we showed that one of firms' motives for integrating suppliers is explicitly to produce a higher share of high quality output. It is thus not the case that higher output quality in vertically integrated Peruvian fishmeal manufacturers is simply an ignored by-product of integration decisions made for other reasons, nor that integration and output quality are causally unrelated in the "minds" of such firms. Instead, our evidence indicates that

[49] Remarkably, the instrumented coefficient on Share VI is statistically significant in one of the two columns in Appendix Table A5, despite there only being 66 observations in these single-firm regressions. In the other IV column, the estimated coefficient is very similar in magnitude, but not significant.

[50] It might also be that a plant's use of independent suppliers itself affects the number and share of independent suppliers supplying other plants in the port due to an "adding up" constraint, or that high fish density might simultaneously enable plants to produce higher quality fishmeal (as we show in Appendix D) and attract independent fishing boats. Both these scenarios would imply that the presence of independent suppliers in the port is positively related to Share VI in the first stage: we find negative signs.

[51] 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 respectively one and minus one. We find a coefficient of 0.04, with a standard error of 0.080.

[52] In Table 3 we also showed that the firm level relationship between vertical integration and output quality holds when we control for the firm's share of total industry output and supplier characteristics.

vertical integration increases output quality and that firms therefore integrate suppliers when the quality premium rises.

8 Conclusion

Guided by Holmstrom & Milgrom (1991)'s classical ideas and subsequent theories of the firm characterizing how we expect firm boundaries to respond to the multitasking nature of suppliers' work (Baker *et al.* , 2001, 2002; Gibbons, 2005a,b), this paper identifies an overlooked motivation for and consequence of vertical integration in incomplete contracts settings: downstream firms integrate to be able to produce output of high enough quality to sell to high-paying consumers abroad. Integration allows manufacturing firms to incentivize quality-increasing behavior from existing suppliers and better control input quality.

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 predicts 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 manufacturers produce in Peru's fishmeal industry. We show that, when firms' returns to shifting from low to high quality production rise for exogenous reasons, they acquire more of their suppliers. This strategy appears to be effective because fishing boats change their behavior in a way consistent with delivering fresher fish when they are acquired by the downstream firm they supply—which helps firms produce higher quality fishmeal. Finally, we show that firms in fact produce higher quality output when their organizational structure is more vertically integrated.

These results are inconsistent with alternative theories in which the integration-quality relationship reflects third factors that affect both firms' choice of structure and products produced without the two being directly related. They are also inconsistent with explanations in which 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. The evidence we present suggests that, while firms vertically integrate for many different reasons, in settings where output quality is vertically differentiated and contracts incomplete, one motive is quality upgrading. That is, integration is an explicit organizational choice made *in order to* "climb" the quality ladder.

A natural next question is the generality of this finding. In Figure 3, we plot a proxy for average quality that is available for most exporter countries—the average unit value of manufacturing products exported to the U.S.—against the share of those exports that is imported by "related party" downstream firms located in the U.S. (a measure of vertical integration). The figure shows clear evidence of an upward-sloping relationship between average unit values and related party import shares. The same relationship holds also within product categories.[53] This suggests that our findings reflect an association between vertical integration and

[53] We show this in Appendix Table A6. In Figure 3, the variable plotted on the y-axis is $\hat{\gamma}_c$ from the regression $\log(\text{unit value})_{cpt} = \alpha_{pt} + \gamma_c + \varepsilon_{cpt}$, where $\log(\text{unit value})_{cpt}$ is the average log unit value of products exported from country c, of HS6 code p, in year t to the U.S.; α_{pt} is a product×year fixed effect; and γ_c is an origin country fixed effect. This regression is estimated using COMTRADE data from BACI (See Gaulier & Zignago (2010) for a description of the data). The variable plotted on the x-axis is $\hat{\delta}_c$ from the regression $\text{Related party share of U.S. imports}_{cpt} = \beta_{pt} + \delta_c + \upsilon_{cpt}$, where $\text{Related party share of U.S. imports}_{cpt}$ is the share of products exported from country c, of NAICS code p, in year t to the U.S. that are imported by related parties (usually other units of the same

manufacturing output quality that tends to hold on average across countries and manufacturing industries. We find this unsurprising, as theory suggests that integration may address the contracting problems that are typical when producing high quality goods. Given this—and despite vertical integration *overall* being more common in developing countries (Acemoglu *et al.* , 2009; 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.

firm (Ruhl, 2015)); β_{pt} is a product×year fixed effect; and δ_c is an origin country fixed effect. This regression is estimated using data from the U.S. Census Bureau.

References

ACEMOGLU, D., AGHION, P., LELARGE, C., VAN REENEN, J., & ZILIBOTTI, F. 2007. Technology, Information and the Decentralization of the Firm. *Quarterly Journal of Economics*.

ACEMOGLU, D., AGHION, P., GRIFFITH, R., & ZILIBOTTI, F. 2010. Vertical Integration and Technology: Theory and Evidence. *Journal of European Economic Association*.

ACEMOGLU, DARON, JOHNSON, SIMON, & MITTON, TODD. 2009. Determinants of vertical integration: financial development and contracting costs. *Journal of Finance*.

ALFARO, LAURA, CONCONI, PAOLA, FADINGER, HARALD, & NEWMAN, ANDREW F. 2016. Do Prices Determine Vertical Integration? *Review of Economic Studies*, **83**, 855–888.

AMODIO, FRANCESCO, & MARTINEZ-CARRASCO, MIGUEL A. 2018. Input Allocation, Workforce Management and Productivity Spillovers: Evidence from Personnel Data. *Review of Economic Studies*.

ANTRÀS, P., & STAIGER, R. W. 2012. Offshoring and the Role of Trade Agreements. *American Economic Review*.

ANTRÀS, POL. 2003. Firms, Contracts, and Trade Structure. *Quarterly Journal of Economics*, **118**, 1375–1418.

ANTRÀS, POL. 2014. Grossman-Hart (1986) Goes Global: Incomplete Contracts, Property Rights, and the International Organization of Production. *Journal of Law, Economics, and Organization*.

ANTRÀS, POL. 2016. *Global Production: Firms, Contracts, and Trade Structure*. Princeton University Press.

ATALAY, ENGHIN, HORTACSU, ALI, & SYVERSON, CHAD. 2014. Vertical Integration and Input Flows. *American Economic Review*, **104**(4), 1120–48.

ATKIN, DAVID, KHANDELWAL, AMIT, & OSMAN, ADAM. 2017. Exporting and Firm Performance: Evidence from a Randomized Trial. *Quarterly Journal of Economics*.

AXBARD, SEBASTIAN. 2016. Income Opportunities and Sea Piracy in Indonesia: Evidence from Satellite Data. *American Economic Journal: Applied Economics*, **8**(2), 154–94.

BAKER, G, & HUBBARD, T. 2003. Make Versus Buy in Trucking: Asset Ownership, Job Design and Information. *American Economic Review*.

BAKER, G, & HUBBARD, T. 2004. Contractibility and Asset Ownership: On-Board Computers and Governance in US Trucking. *Quarterly Journal of Economics*.

BAKER, GEORGE, GIBBONS, ROBERT, & MURPHY, KEVIN J. 2001. Bringing the Market Inside the Firm? *American Economic Review, Papers and Proceedings*.

BAKER, GEORGE, GIBBONS, ROBERT, & MURPHY, KEVIN J. 2002. Relational Contracts and the Theory of the Firm. *The Quarterly Journal of Economics*, **117**(1), 39–84.

BALDWIN, R., & HARRIGAN, J. 2011. Zeros, Quality, and Space: Trade Theory and Trade Evidence. *American Economic Journal: Microeconomics*.

BANERJEE, A., & DUFLO, E. 2000. Reputation Effects and the Limits of Contracting: A Study of the Indian Software Industry. *Quarterly Journal of Economics*.

BASTOS, PAULO, SILVA, JOANA, & VERHOOGEN, ERIC. 2018. Export Destinations and Input Prices. *American Economic Review*.

BOEHM, JOHANNES. 2016. *The Impact of Contract Enforcement Costs on Outsourcing and Aggregate Productivity*. mimeo, Sciences Po.

BRAMBILLA, IRENE, & PORTO, GUIDO. 2016. High-Income Export Destinations, Quality and Wages. *Journal of International Economics*.

BRAMBILLA, IRENE, LEDERMAN, DANIEL, & PORTO, GUIDO. 2012. Exports, Export Destinations, and Skills. *American Economic Review*, **102**(7), 3406–38.

BRAMBILLA, IRENE, LEDERMAN, DANIEL, & PORTO, GUIDO. 2016. *Exporting Firms and the Demand for Skilled Tasks*. mimeo, UNLP.

BRESNAHAN, TIMOTHY F., & LEVIN, JONATHAN D. 2012. Vertical integration and Market Structure. *In:* GIBBONS, R., & ROBERTS, D.J. (eds), *Handbook of Organizational Economics*.

BREZA, EMILY, & LIBERMAN, ANDRES. 2017. Financial Contracting and Organizational Form: Evidence from the Regulation of Trade Credit. *Journal of Finance*.

COASE, R. 1937. The Nature of the Firm. *Economica*.

DE LA PUENTE, OSCAR, SUEIRO, JUAN CARLOS, HECK, CARMEN, SOLDI, GIULIANA, & PUENTE, SANTIAGO DE LA. 2011. *Evaluacion de los sistemas de gestion pesquera en el marco de la certificacion a cargo del marine. La pesqueria Peruana de anchoveta*. Stewardship Council, Cayetano Heredia University WP.

FAO. 1986. *The production of fish meal and oil*. FAO FISHERIES TECHNICAL PAPER: 142.

FLUCKIGER, MATTHIAS, & LUDWIG, MARKUS. 2015. Economic shocks in the fisheries sector and maritime piracy. *Journal of Development Economics*, **114**, 107 – 125.

FORBES, S. J., & LEDERMAN, M. 2009. Adaptation and Vertical Integration in the Airline Industry. *American Economic Review*.

FORBES, SILKE J., & LEDERMAN, MARA. 2010. Does vertical integration affect firm performance? Evidence from the airline industry. *The RAND Journal of Economics*, **41**(4), pp. 765–790.

FRÍAS, JUDITH A., KAPLAN, DAVID S., & VERHOOGEN, ERIC. 2009. *Exports and Wage Premia: Evidence from Mexican Employer-Employee Data*. mimeo Columbia University.

GAULIER, GUILLAUME, & ZIGNAGO, SOLEDAD. 2010 (Oct.). *BACI: International Trade Database at the Product-Level. The 1994-2007 Version*. Working Papers 2010-23. CEPII research center.

GIBBONS, R., & ROBERTS, D.J. 2013. *Handbook of Organizational Economics*. Princeton University Press.

GIBBONS, ROBERT. 2005a. Four formal(izable) theories of the firm? *Journal of Economic Behavior and Organization*, **58**(2), 200 – 245.

GIBBONS, ROBERT. 2005b. Incentives Between Firms (and Within). *Management Science*.

GIL, R. 2009. Revenue sharing distortions and vertical integration in the movie industry. *Journal of Law, Economics, and Organization*.

GIL, R., & KIM, M. 2016. *Does Competition Increase Quality? Evidence from the US Airline Industry*. mimeo, Carey School of Business.

GIL, R., KIM, M., & ZANARONE, G. 2016. *The Value of Relational Adaptation in Outsourcing: Evidence from the 2008 shock to the US Airline Industry*. mimeo, Carey School of Business.

GOLDBERG, PINELOPI K, & PAVCNIK, NINA. 2007. Distributional Effects of Globalization in Developing Countries. *Journal of Economic Literature*, **45**, 39–82.

GOLDBERG, PINELOPI KOUJIANOU, KHANDELWAL, AMIT KUMAR, PAVCNIK, NINA, & TOPALOVA, PETIA. 2010. Imported Intermediate Inputs and Domestic Product Growth: Evidence from India. *The Quarterly Journal of Economics*, **125**(4), 1727–1767.

GOLDSMITH-PINKHAM, PAUL, SORKIN, ISAAC, & SWIFT, HENRY. 2017. *Bartik Instruments: What, When, Why, and How*. Mimeo, Yale.

GROSSMAN, G. M., & HELPMAN, E. 2002. Integration Versus Outsourcing In Industry Equilibrium. *Quarterly Journal of Economics*.

GROSSMAN, SANFORD J., & HART, OLIVER D. 1986. The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration. *Journal of Political Economy*, **94**(4), pp. 691–719.

HALLAK, JUAN CARLOS. 2006. Product quality and the direction of trade. *Journal of International Economics*.

HALLAK, JUAN CARLOS, & SCHOTT, PETER. 2011. Estimating Cross-Country Differences in Product Quality. *Quarterly Journal of Economics*.

HALPERN, LASZLO, KOREN, MIKLOS, & SZEIDL, ADAM. 2015. Imported Inputs and Productivity. *American Economic Review*.

HART, OLIVER, & MOORE, JOHN. 1990. Property Rights and the Nature of the Firm. *Journal of Political Economy*, **98**(6), pp. 1119–1158.

HART, OLIVER, SHLEIFER, ANDREI, & VISHNY, ROBERT W. 1997. The proper scope of government: theory and an application to prisons. *The Quarterly Journal of Economics*, **112**(4), 1127–1161.

HOLMSTROM, B. 1999. The Firm as a Subeconomy. *Journal of Law, Economics, and Organization*.

HOLMSTROM, B., & MILGROM, P. 1994. The firm as an incentive system. *American Economic Review*.

HOLMSTROM, B., & TIROLE, J. 1991. Transfer pricing and organizational form. *Journal of Law, Economics, and Organization*.

HOLMSTROM, BENGT, & MILGROM, PAUL. 1991. Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, & Organization*, 24–52.

HORTACSU, ALI, & SYVERSON, CHAD. 2007. Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices. *Journal of Political Economy*.

HUMMELS, D., & KLENOW, P. 2005. The variety and quality of a nation's exports. *American Economic Review*.

IRARRAZABAL, ALFONSO, MOXNES, ANDREAS, & OPROMOLLA, LUCA DAVID. 2013. The Margins of Multinational Production and the Role of Intrafirm Trade. *Journal of Political Economy*.

JOHNSON, ROBERT C. 2012. Trade and Prices with Heterogeneous Firms. *Journal of International Economics*.

KHANDELWAL, AMIT. 2010. The Long and Short (of) Quality Ladders. *The Review of Economic Studies*, **77**(4), 1450–1476.

KOSOVÁ, RENÁTA, LAFONTAINE, FRANCINE, & PERRIGOT, ROZENN. 2013. Organizational Form and Performance: Evidence from the Hotel Industry. *Review of Economics and Statistics*.

KUGLER, MAURICE, & VERHOOGEN, ERIC. 2012. Prices, Plant Size, and Product Quality. *Review of Economic Studies*, **79**(1), 307–339.

LAFONTAINE, FRANCINE, & SLADE, MARGARET. 2007. Vertical Integration and Firm Boundaries: The Evidence. *Journal of Economic Literature*, **45**(3), 629–685.

LEGROS, PATRICK, & NEWMAN, ANDREW F. 2013. A Price Theory of Vertical and Lateral Integration. *Quarterly Journal of Economics*.

MACCHIAVELLO, ROCCO. 2011. Financial development and vertical integration: theory and evidence. *Journal of the European Economic Association*.

MACCHIAVELLO, ROCCO, & MIQUEL-FLORENSA, JOSEPA. 2016. *Vertical Integration and Relational Contracts: Evidence from Costa Rica's Coffee Chain*. Unpub. paper.

MACCHIAVELLO, ROCCO, & MORJARIA, AMEET. 2015. The Value of Relationships: Evidence from a Supply Shock to Kenya Rose Exports. *American Economic Review*.

MANOVA, KALINA, & ZHANG, ZHIWEI. 2012. Export Prices Across Firms and Destinations. *The Quarterly Journal of Economics*, **127**(1), 379–436.

MARTINEZ-CARRASCO, JOSE. 2017. *Coordination and Gains from Relational Contracts: Evidence from the Peruvian Anchovy Fishery*. Mimeo, Duke University.

MCLAREN, J. 2000. Globalization and Vertical Structure. *American Economic Review*.

MCMILLAN, JOHN, & WOODRUFF, CHRISTOPHER. 1999. Interfirm Relationships and Informal Credit in Vietnam. *Quarterly Journal of Economics*.

MEDINA, PAMELA. 2017. *Import Competition, Quality Upgrading and Exporting: Evidence from the Peruvian Apparel Industry*. Mimeo, University of Toronto.

MULLAINATHAN, S., & SHARFSTEIN, D. 2001. Do firm boundaries matter? *American Economic Review*.

NATIVIDAD, GABRIEL. 2014. Integration and Productivity: Satellite-Tracked Evidence. *Management Science*, **60**(7), 1698–1718.

NOVAK, SHARON, & STERN, SCOTT. 2008. How Does Outsourcing Affect Performance Dynamics? Evidence from the Automobile Industry. *Management Science*, **54**(12), pp. 1963–1979.

NUNN, NATHAN. 2007. Relationship-Specificity, Incomplete Contracts, and the Pattern of Trade. *The Quarterly Journal of Economics*, **122**(2), 569–600.

ORDOVER, JA, SALONER, G, & SALOP, SC. 1990. Equilibrium vertical foreclosure. *American Economic Review*.

PAREDES, E. CARLOS, & GUTIERREZ, MARIA E. 2008. *The Peruvian Anchovy Sector: Costs and Benefits. An analysis of recent behavior and future challenges*. IIFET 2008 Vietnam Proceedings.

PARK, ALBERT, YANG, DEAN, SHI, XINZHENG, & JIANG, YUAN. 2009. *Exporting and Firm Performance: Chinese Exporters and the Asian Financial Crisis*. NBER Working Papers 14632.

PIGOU, ARTHUR C. 1912. *Wealth and Welfare*. London: Macmillan.

RAMONDO, NATALIA, RAPPOPORT, VERONICA, & RUHL, KIM. 2016. Intrafirm Trade and Vertical Fragmentation in U.S. Multinational Corporations. *Journal of International Economics*.

RUHL, KIM J. 2015. How Well is U.S. Intrafirm Trade Measured? *American Economic Review: Papers and Proceedings*.

SCHOTT, P. 2004. Across-product versus within-product specialization in international trade. *Quarterly Journal of Economics*.

SUTTON, J. 1991. *Sunk Costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration*. MIT Press.

SUTTON, J. 1998. *Technology and Market Structure: Theory and History*. MIT Press.

TINTELNOT, FELIX, MOGSTAD, KEN KIKKAWA MAGNE, & DHYNE, EMMANUEL. 2017. *Trade and Domestic Production Networks*. Mimeo, Chicago University.

TVETERAS, SIGBJORN, PAREDES, CARLOS E., & TORRES, JULIO PENA. 2011. Individual Vessel Quotas in Peru: Stopping the Race for Anchovies. *Marine Resource Economics*, **26**, 225–232.

VERHOOGEN, ERIC A. 2008. Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector. *The Quarterly Journal of Economics*, **123**(2), 489–530.

WILLIAMSON, OLIVER. 1971. The Vertical Integration of Production: Market Failure Considerations. *The American Economic Review*, **61**(2), pp. 112–123.

WILLIAMSON, OLIVER. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications, a Study in the Economics of Internal Organization*. Free Press.

WILLIAMSON, OLIVER. 1985. *The Economic Institutions of Capitalism*. Free Press.

WOODRUFF, C. 2002. Non-contractible Investments and Vertical Integration in the Mexican Footwear Industry. *International Journal of Industrial Organization*.

TABLE 1: SUMMARY STATISTICS

		Mean	Sd
Firms	Total number of firms in sample	37	
	Export shipment (metric tons)	380	(351)
	Export Price ($/metric ton)	1454	(303)
	Number of destinations per season	7.05	(5.30)
	Number of export transactions per season	85	(99)
Plants	Total number of plants in sample	94	
	Has high technology	0.85	(0.36)
	High quality share of production	0.85	(0.35)
	Monthly production (metric tons)	3116	(3266)
	Processing capacity (metric tons/hour)	106	(54)
Boats	Number of boats operating per season	812	92
	Fraction owned by a downstream firm per season	0.28	(0.45)
	Fraction of boats made of steel per season	0.44	(0.50)
	Storage capacity (m3)	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)
	Max. distance from the plant's port (kms)	76	(46)

Notes: This table gives summary statistics over our sample period. *Has high technology* is a dummy equal to 1 if the plant is equipped with steam drying technology. Plants' *processing capacity* measures the total weight of fish that can be processed in an hour. *Steel* is a binary variable equal to 1 if a boat is a steel boat (which tend to be bigger, better suited for industrial fishing, and are 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. *Max. distance from the plant's port* is the maximum distance between the boat and the port it delivers to on any trip.

TABLE 2: SUMMARY STATISTICS ON INTEGRATION

Panel A: Boat purchases and sales

Total number of steel boats registered	741
Number of steel boat transactions	317
Number of transactions Indep. $\rightarrow$ VI	103
Number of transactions VI $\rightarrow$ Indep.	32
Number of transactions VI $\rightarrow$ VI	50
Number of transactions Indep. $\rightarrow$ Indep.	132

Panel B: Decomposition of the growth rate of Share of inputs from VI suppliers

$$\text{Growth (Share VI)}_{i,t} \approx \log\left(\frac{\text{Share VI}_{i,t+1}}{\text{Share VI}_{i,t}}\right) = \log\left(\frac{\frac{\text{VI}_{i,t+1}}{\text{Total}_{i,t+1}}}{\frac{\text{VI}_{i,t}}{\text{Total}_{i,t}}}\right) = \underbrace{\log\left(\frac{\frac{\text{VI}_{i,t+1}}{\text{Total}_{t+1}}}{\frac{\text{VI}_{i,t}}{\text{Total}_{t}}}\right)}_{A} - \underbrace{\log\left(\frac{\frac{\text{Total}_{i,t+1}}{\text{Total}_{t+1}}}{\frac{\text{Total}_{i,t}}{\text{Total}_{t}}}\right)}_{B}$$

	Total	A (Boats purchases or sales)	B (Buying less from Indep.)
Growth	2.1%	1.4%	0.7%
Relative Contribution		67%	33%

Notes: Panel A displays basic statistics on boat purchases and sales. In Panel B, the growth rate of "Share $\text{VI}_{i,t}$" – the share of the inputs sourced by firm i during production season t that comes from vertically integrated suppliers – can be decomposed as presented in the first row of this table. $\text{VI}_{i,t}$ and $\text{Total}_{i,t}$ is respectively the *amount* of inputs firm i sources from vertically integrated suppliers and in total during season t, and Total_t is the total amount of inputs sourced by the industry as a whole during season t. Term A can then be interpreted as the contribution to the growth rate of Share $\text{VI}_{i,t}$ that comes from increasing solely the (relative) amount of inputs coming from integrated suppliers. Term B can be interpreted as the contribution of a firm decreasing the (relative) amount of inputs sourced from all suppliers. The table gives the growth rate of "Share $\text{VI}_{i,t}$", Term A and Term B.

TABLE 3: OUTPUT QUALITY AND VERTICALLY INTEGRATED SUPPLIERS

Panel A: Output quality and number of suppliers owned

Dep. var:	High Quality share of prod.		Protein content	
	(1)	(2)	(3)	(4)
Asinh(Number of suppliers owned)	0.056	0.043	0.197**	0.167**
	(0.060)	(0.042)	(0.083)	(0.066)
High technology share of capacity	No	Yes	No	Yes
Season FEs	Yes	Yes	Yes	Yes
Firm FEs	Yes	Yes	Yes	Yes
Mean of Dep. Var.	0.75	0.75	65.7	65.7
N	275	275	208	208

Panel B: Output quality and Share of inputs from VI suppliers

Dep. var:	High Quality share of prod.		Protein content	
	(1)	(2)	(3)	(4)
Share of inputs from VI suppliers	0.377**	0.375**	0.982***	1.040***
	(0.166)	(0.164)	(0.293)	(0.283)
High technology share of capacity	No	Yes	No	Yes
Season FEs	Yes	Yes	Yes	Yes
Firm FEs	Yes	Yes	Yes	Yes
Mean of Dep. Var.	0.75	0.75	65.7	65.7
N	275	275	208	208

Panel C: Output quality and Share of inputs from VI suppliers

Dep. var:	High Quality share of prod.			Protein Content		
	(1)	(2)	(3)	(4)	(5)	(6)
Share of inputs from VI suppliers	0.313*	0.373**	0.313*	1.153**	0.936***	1.190**
	(0.159)	(0.148)	(0.159)	(0.457)	(0.260)	(0.461)
Share of inputs from steel boats	−0.098		−0.092	−1.152		−1.153
	(0.164)		(0.161)	(0.808)		(0.794)
Share of inputs from boats with high capacity	0.152		0.139	0.988		1.036
	(0.166)		(0.164)	(1.003)		(0.976)
Share of inputs from boats with cooling system	0.191		0.202	−0.232		−0.282
	(0.123)		(0.124)	(0.979)		(0.986)
Share of industry's production		−0.807	−0.748		1.799	1.889
		(0.925)	(0.891)		(3.869)	(3.930)
High technology share of capacity	Yes	Yes	Yes	Yes	Yes	Yes
Season FEs	Yes	Yes	Yes	Yes	Yes	Yes
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes
Mean of Dep. Var.	0.75	0.75	0.75	0.75	65.7	65.7
N	275	275	275	208	208	208

Notes: One observation is a firm during a production season. *Protein content* is the quantity weighted average of a measure of quality inferred from a database that provides weekly prices by quality. *Log(unit price)* is the log of the quantity weighted average unit price of exports during a season. *Share of inputs from VI suppliers* is the share of a firm's (or plant's) inputs that come from VI suppliers during a season. Steel boats tend to be bigger, 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 integrated cooling system use ice to keep fish fresh. *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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE 4: VERTICAL INTEGRATION AND THE QUALITY PRICE PREMIUM

Dep. var:	Share of inputs from VI suppliers (t) - Share of inputs from VI suppliers (t-1)					
	OLS	OLS	OLS	IV	IV	OLS
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Average Price)	0.030					
	(0.052)					
Quality Premium		0.265**				
		(0.114)				
Upgradable share of production (t-1) × Quality premium			3.014**	3.382***	2.591*	
			(1.091)	(1.050)	(1.393)	
Firm specific price (weighted by past production)						−1.225*
						(0.596)
Upgradable share of production (t-1)	No	No	Yes	Yes	Yes	No
Season FEs	No	No	Yes	Yes	Yes	Yes
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes
Mean of Dep. Var.	0.01	0.01	0.01	0.01	0.01	0.01
N	190	190	190	190	190	190
Kleibergen-Paap LM p-value (Under-id)				0.01	0.01	
Kleibergen-Paap Wald F statistic (Weak inst)				38.3	209.1	

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 5: SUPPLIER BEHAVIOR AND VERTICAL INTEGRATION

Panel A: Identified from all switchers (Idependent to VI, VI to Independent and VI to VI)

Dep. var:	Log(Quantity supplied)	Log(Max. distance from the plant's port)	Log(Total time spent at sea)
	(1)	(2)	(3)
I[VI × supplies owner firm]	-0.096^{***}	-0.054^{***}	-0.030^{*}
	(0.023)	(0.019)	(0.016)
Date FEs	Yes	Yes	Yes
Supplier × Plant FEs	Yes	Yes	Yes
N	315,442	137,278	159,724

Panel B: Idenfified only from VI switchers changing ownership (VI to VI)

Dep. var:	Log(Quantity supplied)	Log(Max. distance from the plant's port)	Log(Total time spent at sea)
	(1)	(2)	(3)
I[Always VI × supplies owner firm]	-0.147^{***}	-0.082^{***}	-0.073^{***}
	(0.027)	(0.026)	(0.023)
Date FEs	Yes	Yes	Yes
Supplier × Plant FEs	Yes	Yes	Yes
N	315,442	137,274	159,724

Notes: One observation is a boat during a fishing trip. *Quantity supplied* is the amount of fish the boat delivers to the plant per trip. *Max. distance from the plant's port* is maximum distance a specific boat is observed away from port. *Max. distance from the plant's port* can only be measured if the boat leaves from and arrives at the same port. *Total time at sea* is the amount of time the boat is away from port per trip. The number of observations varies from one column to the next as GPS variables for a given trip are sometimes missing.In panel A, we define I[VI×supplies owner firm] to be equal to one if the supplier is (i) currently vertically integrated (ii) currently delivering to its parent firm. In panel B, we define I[Always VI×supplies owner firm] to be equal to one if the supplier is (i) always owned by a fishmeal firm, and (ii) currently delivering to its parent firm. Because we include Supplier × Plant FEs, I[VI×supplies owner firm] and I[Always VI×supplies owner firm] are identified based only on suppliers who change ownership during our sample period. Standard errors clustered at the boat level are included in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 6: OUTPUT QUALITY AND SHARE OF INPUTS FROM VERTICALLY INTEGRATED SUPPLIERS

	Impact of Share of VI Inputs on Quality					
Dep. var:	High Quality Share of Production					
	OLS	OLS	IV: Ind. Boats		IV: Wooden Boats	
	(1)	(2)	(3)	(4)	(5)	(6)
Share of inputs from VI suppliers	0.102^{**}	0.064^{**}	0.213^{***}	0.169^{**}	0.160^{***}	0.142^{**}
	(0.038)	(0.030)	(0.080)	(0.068)	(0.060)	(0.060)
Has high technology	No	Yes	No	Yes	No	Yes
Month FEs	Yes	Yes	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes	Yes	Yes
Mean of Dep. Var.	0.85	0.85	0.85	0.85	0.85	0.85
N	2647	2647	2487	2487	2647	2647

Notes: Panel A includes data from all plants at the month level and uses the share of high quality production as a dependent variable—based on a directly observed dichotomous measure of quality that is available for all firms. Panel B focuses on a single firm for which more detailed plant level measures are available at the season level. *Log(unit price)* is the log of the quantity weighted average unit price of exports during a season. *Protein content* is the quantity weighted average of a measure of quality inferred with a database that provides weekly prices by quality. *Has high technology* controls for whether a plant is equipped or not with the steam drying technology. *Share of inputs from VI suppliers* is the share of a plant's inputs that come from VI suppliers in a given season. In IV specifications share of inputs from VI suppliers is instrumented by (a) the number of independent boats present in the plant's port in the season in question, excluding those that interact directly with the plant itself (formally, plants that belong to the firm that owns the plant in question, but one firm owning more than one plant in a given port is unusual), and (b) the ratio of the number of boats in (a) to the total number of boats in the plant's port in that season that do not interact with the plant itself. The first stage is shown in Appendix Table A4. Panel A shows standard errors clustered at the firm level in parentheses. Panel B shows robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

FIGURE 1: INTERACTIONS BETWEEN MANUFACTURERS AND INTEGRATED AND INDEPENDENT SUPPLIERS

All Suppliers

(a) Distribution of interactions across firms

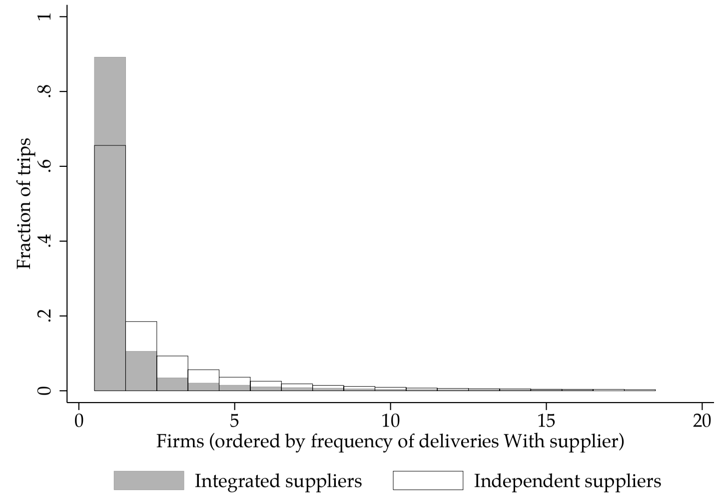

(b) Distribution of interactions across plants

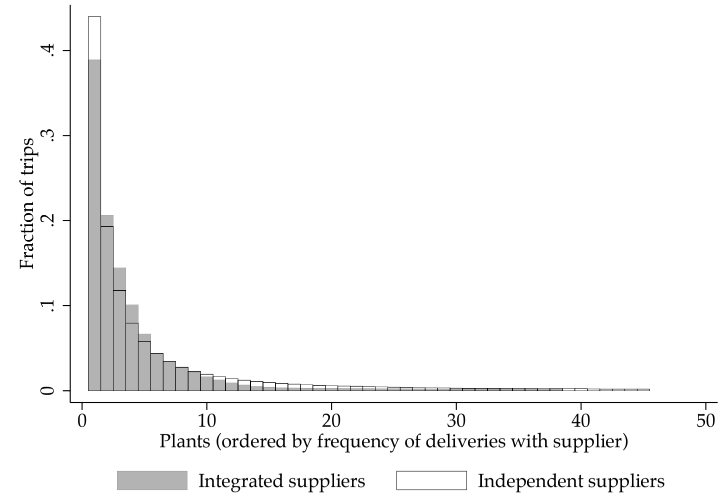

Switchers Only

(c) Distribution of interactions across firms

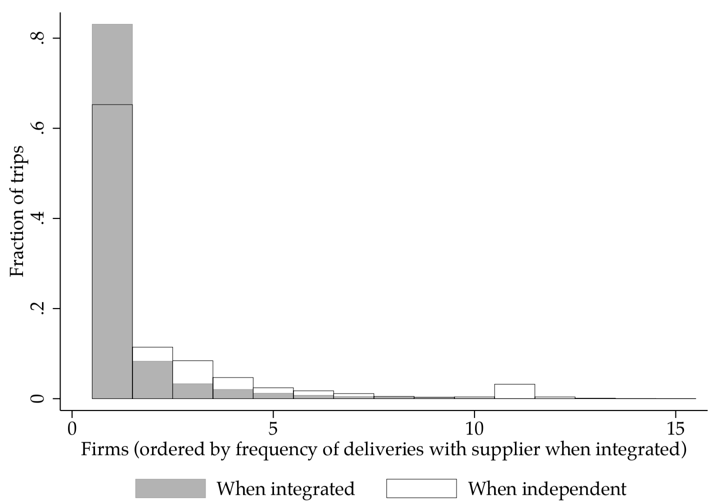

(d) Distribution of interactions across plants

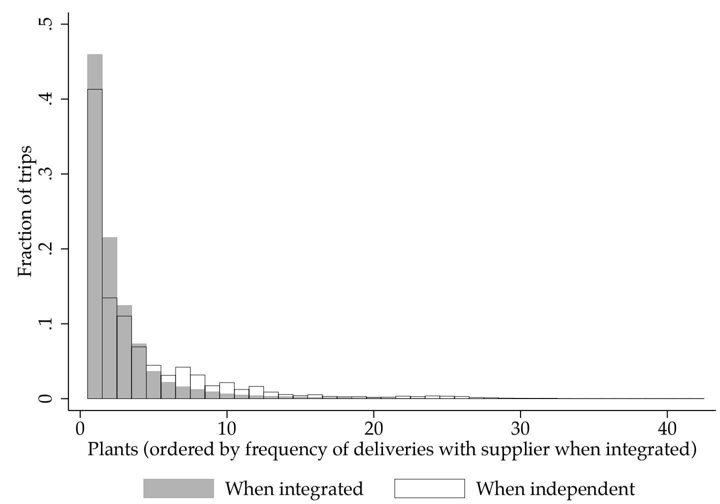

Notes: Figures above show the average fraction of deliveries at each firm and plant for independent and integrated suppliers. In figures (a) and (b), plants or firms are ordered based on the frequency of deliveries for each boat x boat owner pair: the plant or firm that receives the highest number of deliveries by the boat in question while owned by the owner in question is ranked one, the next highest ranked two, and so on. In figures (c) and (d) plants or firms are ranked based on the frequency of deliveries for each boat *while it is integrated*. Figures (a) and (b) include all suppliers, while figures (c) and (d) include only *switchers*: boats that were independent at one point and integrated at another point during our sample.

FIGURE 2: FISHMEAL PRICES

(a)

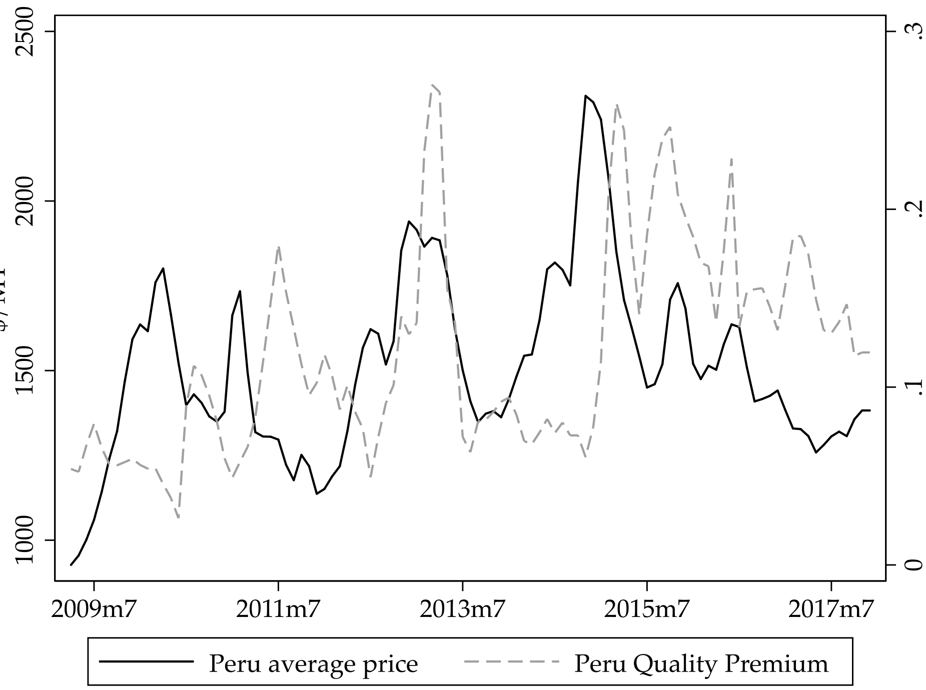

(b)

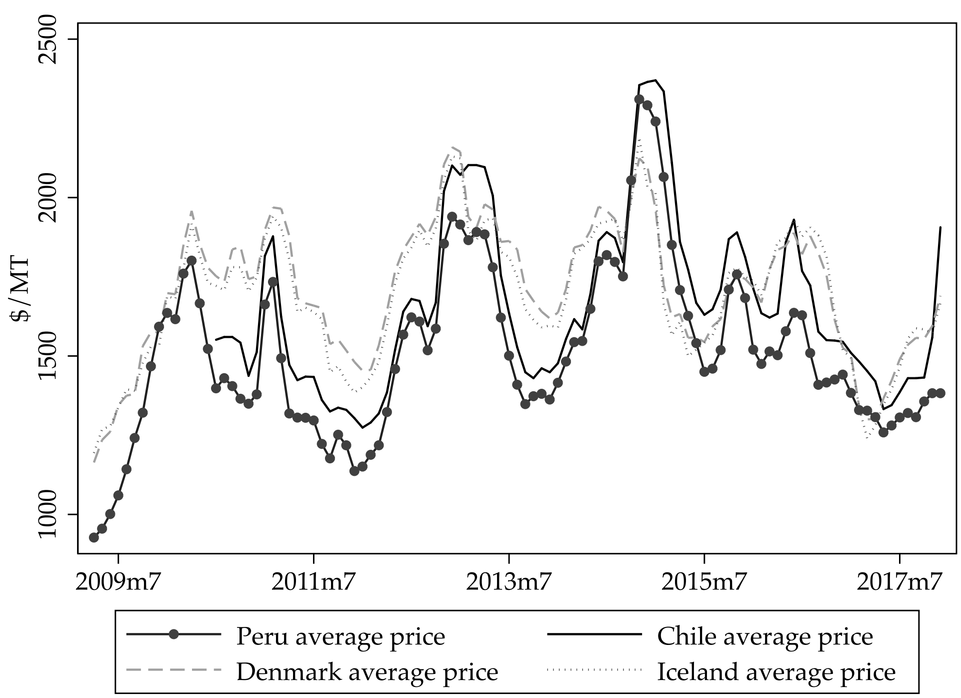

FIGURE 3: COUNTRIES' OUTPUT QUALITY AND VERTICAL INTEGRATION IN EXPORT MANUFACTURING

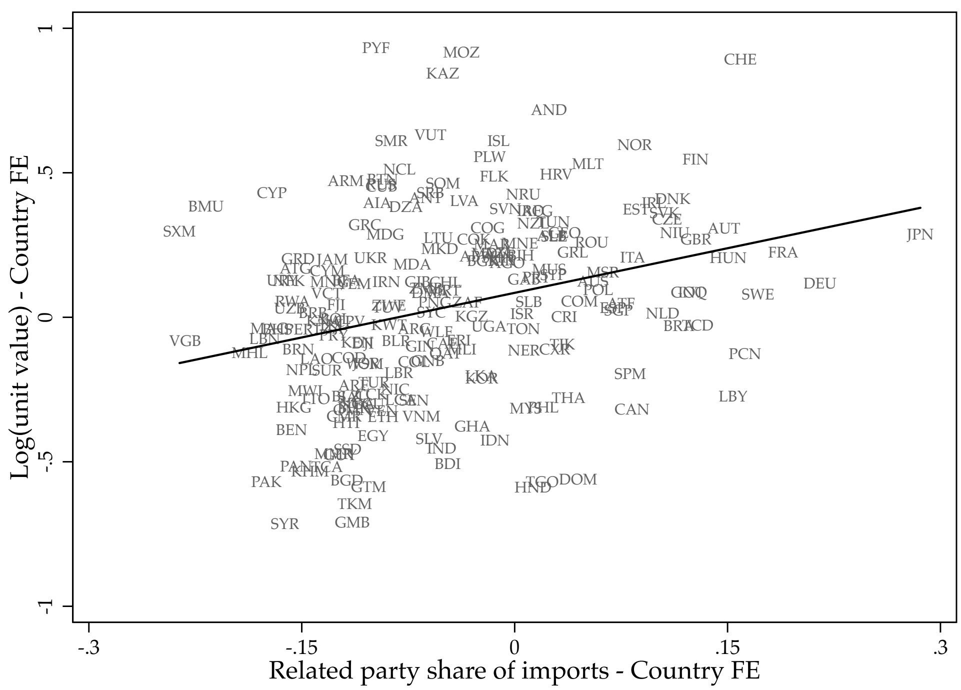

Notes: In this Figure, the variable plotted on the y-axis is $\hat{\gamma}_c$ from the regression log(unit value)$_{cpt} = \alpha_{pt} + \gamma_c + \varepsilon_{cpt}$, where log(unit value)$_{cpt}$ is the average log unit value of products exported from country c, of HS6 code p, in year t to the U.S.; α_{pt} is a product×year fixed effect; and γ_c is an origin country fixed effect. This regression is estimated using COMTRADE data from BACI (See Gaulier & Zignago (2010) for a description of the data). The variable plotted on the x-axis is $\hat{\delta}_c$ from the regression Related party share of U.S. imports$_{cpt} = \beta_{pt} + \delta_c + \upsilon_{cpt}$, where Related party share of U.S. imports$_{cpt}$ is the share of products exported from country c, of NAICS code p, in year t to the U.S. that are imported by related parties (usually other units of the same firm (Ruhl, 2015)); β_{pt} is a product×year fixed effect; and δ_c is an origin country fixed effect. Related party share of U.S. imports$_{cpt}$ is constructed using data from the U.S. Census Bureau. The data is from 2005 to 2014.

Appendix A Additional Tables and Figures

TABLE A1: MAIN IMPORTERS OF PERUVIAN FISHMEAL AND AVERAGE QUALITY IMPORTED

	Total Weight (1000 metric tons)	Average Protein content	Sd(Protein content)
CHINA	4266	66.06	1.60
GERMANY	972	65.42	1.62
JAPAN	545	66.12	1.69
CHILE	305	66.60	1.51
VIETNAM	277	65.91	1.59
TAIWAN	248	66.02	1.71
UNITED KINGDOM	147	65.26	1.62
TURKEY	128	64.91	1.52
INDONESIA	94	66.16	1.64
SPAIN	90	65.44	1.61
AUSTRALIA	85	66.06	1.80
CANADA	66	65.76	1.52
FRANCE	55	65.59	1.72
SOUTH KOREA	24	66.56	1.46
ITALY	21	64.97	1.52
BULGARIA	15	65.42	1.75
VENEZUELA	13	66.67	1.64
PHILIPPINES	12	64.92	1.47
BELGIUM	11	65.08	1.69
INDIA	10	65.17	2.03

Notes: This table reports the top 20 importers of Peruvian fishmeal, the total quantity imported over the whole period of our sample, the average quality imported and the standard deviation of the quality imported across all transactions.

TABLE A2: SUPPLIER CHARACTERISTICS

	Offload weight per trip (metric tons)	Cooling system	Capacity (m3)	Power engine (hp)	Max. Distance from the plant's port (kms)
Wooden	41.00	0.00	65.73	215.40	56.10
	(16.24)	(0.06)	(27.34)	(94.78)	(7.74)
Steel - Independent	104.03	0.09	219.30	412.31	81.15
	(40.77)	(0.28)	(84.35)	(189.82)	(13.43)
Steel - Switchers	148.88	0.25	301.18	616.30	92.25
	(0.43)	(0.444)	(129.92)	(328.51)	(15.37)
Steel - VI	181.62	0.34	382.00	769.96	97.29
	(68.13)	(0.47)	(137.11)	(352.52)	(12.62)

Notes: *Offload weight* is the amount fished on a trip. *Maximum distance from port* is the maximum distance at which a boat is from the port on a fishing trip. Steel boats are generally bigger, better suited for industrial fishing, and are subject to different regulations. Wooden boats cannot be owned by fishmeal firms. *Independent* boats are owned by an individual or a company that is not a fishmeal company. *Switchers* are boats that move from VI to Independent or from Independent to VI at some point in our data. *VI* are boats that remain vertically integrated during the whole sample of our data.

TABLE A3: VERTICAL INTEGRATION AND THE QUALITY PRICE PREMIUM – ROBUSTNESS CHECKS AND FIRST STAGE

First stage		
Dep. var:	Upgradable share of production (t-1) × Quality Premium	
	(1)	(2)
Upgradable share of production (t-1) × [Log(Chile Price)]	0.663^{***}	
	(0.090)	
Upgradable share of production (t-1) ×[Log(Denmark Price)]	-2.253^{***}	
	(0.558)	
Upgradable share of production (t-1) × [Log(Iceland Price)]	1.276^{***}	
	(0.426)	
Upgradable share of production (t-1) × [Log(Qty exp. by Chile)] (t-1)		-0.231
		(0.173)
Upgradable share of production (t-1) × [Log(Qty exp. by Denmark)] (t-1)		0.145^{*}
		(0.080)
Upgradable share of production (t-1) × [Log(Qty exp. by Iceland)] (t-1)		0.262^{**}
		(0.112)
Upgradable share of production (t-1) × [Log(Qty exp. by Thailand)] (t-1		0.059^{**}
		(0.022)
Upgradable share of production (t-1)	Yes	Yes
Season FEs	Yes	Yes
Firm FEs	Yes	Yes

Notes: XXXXX * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A4: OUTPUT QUALITY AND SHARE OF INPUTS FROM VERTICALLY INTEGRATED SUPPLIERS - FIRST STAGE

	First Stage			
Dep. var:	Share of Inputs From VI Suppliers			
	IV: Ind. Boats		IV: Wooden Boats	
	(1)	(2)	(3)	(4)
Number of Independent Boats in Port (Leave-Out)	-0.001^{***}	-0.001^{***}		
	(0.000)	(0.000)		
Number of Wooden Boats in Port (Leave-Out)			-0.002^{***}	-0.002^{***}
			(0.000)	(0.000)
Kleibergen-Paap LM p-value (Under-id)	0.011	0.011	0.010	0.010
Kleibergen-Paap Wald F statistic (Weak inst)	23.55	23.73	27.69	27.69
Anderson-Rubin Wald test p-value	0.01	0.01	0.01	0.01
Has high technology	No	Yes	No	Yes
Month FEs	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes
N	2487	2487	2647	2647

Notes: Results from the first stage of IV specifications reported in Table 6. *Share of inputs from VI suppliers* is instrumented by (a) the number of independent boats present in the plant's port in the season in question, excluding those that interact directly with the plant itself, and (b) the ratio of the number of boats in (a) to the total number of boats in the plant's port in that season that do not interact with the plant itself. The left two columns include standard errors clustered at the firm level in parentheses. The right two columns include robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A5: OUTPUT QUALITY AND SHARE OF INPUTS FROM VERTICALLY INTEGRATED SUPPLIERS - FIRST STAGE

	Impact of Share of VI on Quality			
Dep. var:	Protein Content			
	OLS: Ind. Boats		IV: Ind. Boats	
	(1)	(2)	(3)	(4)
Share of inputs from VI suppliers	1.369**	1.338**	1.469*	1.390
	(0.654)	(0.656)	(0.807)	(0.918)
Has high technology	No	Yes	No	Yes
Month FEs	Yes	Yes	Yes	Yes
Plant FEs	Yes	Yes	Yes	Yes
N	2487	2487	2647	2647
	First Stage			
Dep. var:	Share of Inputs from VI Suppliers			
	(1)	(2)	(3)	(4)
Number of Independent Boats in Port			−0.000	−0.000
			(0.000)	(0.000)
Share of Independent Boats in Port			−0.412**	−0.398*
			(0.200)	(0.207)
Kleibergen-Paap LM p-value (Under-id)			0.005	0.006
Kleibergen-Paap Wald F statistic (Weak inst)			3.61	3.06
Anderson-Rubin Wald test p-value			0.24	0.31

Notes: *Share of inputs from VI suppliers* is instrumented by (a) the number of independent boats present in the plant's port in the season in question, excluding those that interact directly with the plant itself, and (b) the ratio of the number of boats in (a) to the total number of boats in the plant's port in that season that do not interact with the plant itself. The left two columns include standard errors clustered at the firm level in parentheses. The right two columns include robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A6: COUNTRIES' OUTPUT QUALITY AND VERTICAL INTEGRATION IN EXPORT MANUFACTURING

Dep. var:	Log(unit value) - Residuals from HS6×Year FEs and Country FEs
	(1)
Related party share of imports - Residuals from HS6×Year FEs and Country FEs	0.038***
	(0.007)
N	208 024

Notes: In this table, the dependent variable is ε_{cpt} from the regression log(unit value)$_{cpt} = \alpha_{pt} + \gamma_c + \varepsilon_{cpt}$, where log(unit value)$_{cpt}$ is the average log unit value of products exported from country c, of HS6 code p, in year t to the U.S.; α_{pt} is a product×year fixed effect; and γ_c is an origin country fixed effect. This regression is estimated using COMTRADE data from BACI (See Gaulier & Zignago (2010) for a description of the data). The independent variable is υ_{cpt} from the regression Related party share of U.S. imports$_{cpt} = \beta_{pt} + \delta_c + \upsilon_{cpt}$, where Related party share of U.S. imports$_{cpt}$ is the share of products exported from country c, of NAICS code p, in year t to the U.S. that are imported by related parties (usually other units of the same firm (Ruhl, 2015)); β_{pt} is a product×year fixed effect; and δ_c is an origin country fixed effect. Related party share of U.S. imports$_{cpt}$ is constructed using data from the U.S. Census Bureau. Because the product level c (HS6) for the unit value residual is different from the product level p (NAICS) from the share of related party imports residuals, we compute the value weighted unit value residual at the p (NAICS) level using a HS6-NAICS conversion table. This regression includes data from 2005 to 2014. Robust standard errors in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

FIGURE A1: AVERAGE OUTPUT QUALITY AND FIRM SIZE

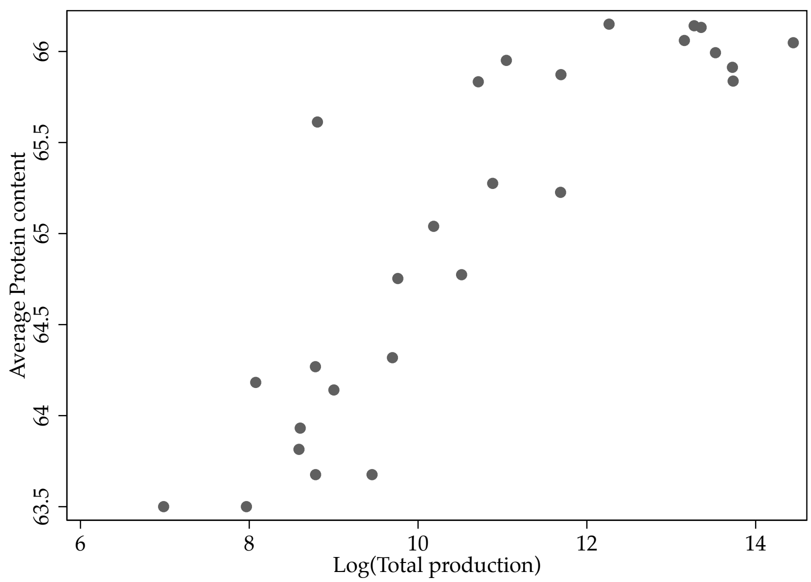

Notes: Each dot represents one fishmeal firm in our sample. Total production is the total weight of fishmeal the firm produced during our data period and average protein content is the quantity weighted average protein content of the firm's fishmeal exports.

FIGURE A2: DENSITY OF INVENTORIES

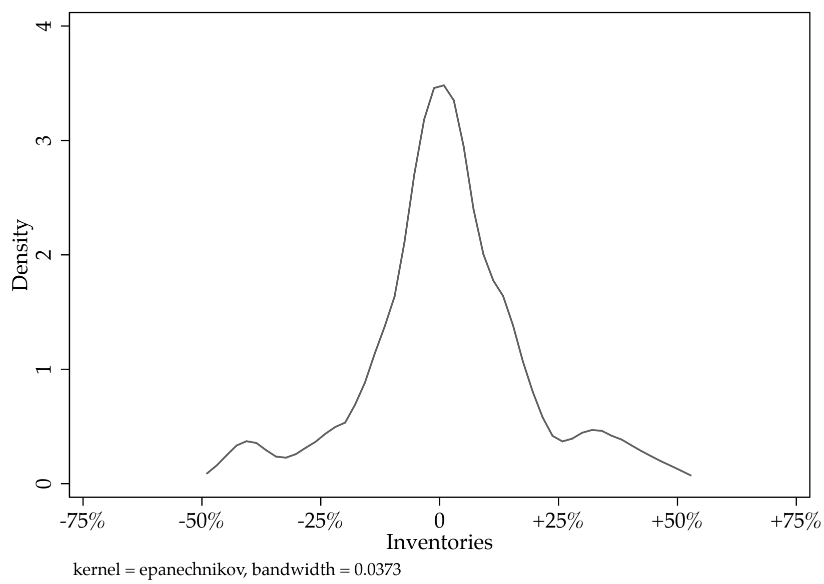

Notes: Kernel density of estimated inventories. Inventories are defined as the ratio of (Total Production - Total Exports) to Total Production, where Total Production is a firm's production during a given production season and Total Exports are the sum of exports that are shipped during the production season and the period directly following the relevant production season (before the next production season starts).

Figure A3: Evolution of the Vertically Integrated share of inputs Industry-wide

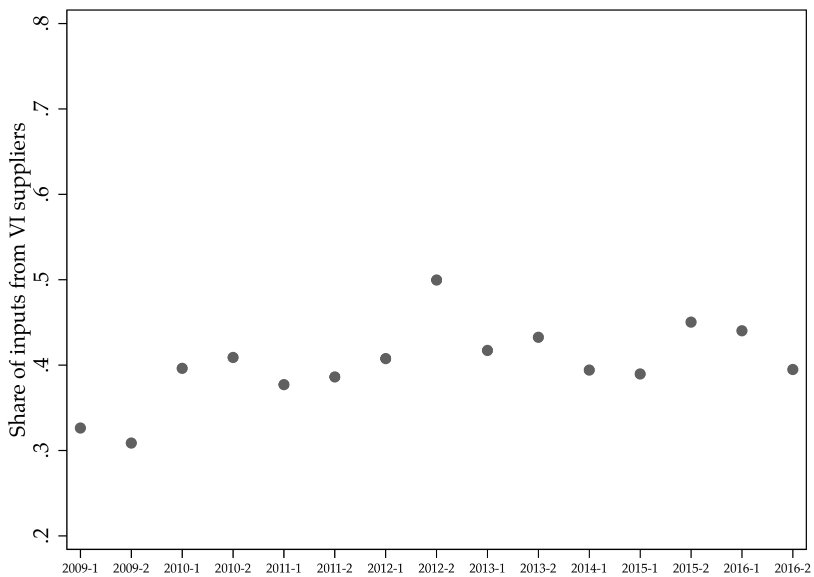

Notes: This graph shoes the evolution of the Peruvian fishmeal industry's share of inputs from integrated suppliers by production season. For every year, -1 is the first production season in the calendar year, in general from April to July, and -2 is the second production season, in general from November to January.

Appendix B Dynamic Theoretical Framework and Relational Contracts

Dynamic theoretical framework

The model presented in the main body of the paper assumes that all transactions are done on the spot market. This stylized version of the model results in the upstream party not taking any action when integrated and the absence of incentives to take a quality-increasing action ($a_2 = 0$). In this version of the model, we follow closely Baker *et al.* (2001, 2002) in allowing the downstream party to use relational contracts to incentivize the quality action.

We make the same assumptions for Q as before, but add a shock to the alternative use P:

$$P = a_1 + \epsilon$$
$$Q = Q_0 - \gamma a_1 + \delta a_2$$

where ϵ is orthogonal to any action taken by the upstream party[54]. We assume that $\epsilon = \bar{\epsilon}$ with probability $\frac{1}{2}$ and $\epsilon = -\bar{\epsilon}$ with probability $\frac{1}{2}$ and that ϵ is known by the upstream party at the time of delivery of the inputs. [55]

As in the main text model, we assume that both P and Q are not contractible. P -the quantity focused alternative use- is perfectly observable at the time of delivery of the inputs, but Q -the quality surplus- is only observed to the downstream party with some delay (e.g. once the inputs are processed)[56].To incentivize the quality-increasing action, the downstream party can offer a payment contingent on the realization of the surplus Q to the upstream party. However, since this payment can only be made after the inputs are delivered, the downstream party can only credibly promise to make this delayed payment through repeated interactions with the upstream party[57]. Note again that at the time of delivery of the inputs, since all parties know the value of Q_0, and because $P = a_1 + \epsilon$ is observable, Q has an observable portion (in expectation) at the time of delivery of the inputs: $\tilde{Q} = Q_0 - \gamma \mathbb{E}(a_1|P) = Q_0 - \gamma P$. Hence, a payment on the spot, proportional to $\tilde{Q}$ is still feasible.

As in Baker *et al.* (2002), we consider four possible organizational structures:

1. Spot Outsourcing (Nonintegrated Asset Ownership, Spot Governance Environment)
2. Relational Outsourcing (Nonintegrated Asset Ownership, Relational Governance Environment)
3. Spot Employment (Integrated Asset Ownership, Spot Governance Environment)
4. Relational Employment (Integrated Asset Ownership, Relational Governance Environment)

[54] We could also assume uncertainty over the realization of the Q surplus, but it would not change the intuition of the result below.

[55] As in the main text model, we assume that $0 \leq \delta \leq 1$ and $0 \leq \gamma \leq 1 - \alpha$. Also, note again that P could itself be the result of a bargaining process between the boat and a quantity focused firm.

[56] In our context, fish quality can hardly be assessed when the fish is offloaded at the factory. However, once the fish is processed in the factory, fishmeal quality can be measured.

[57] In the model, we suppose that this delay is shorter than a full time period, so the surplus Q is observed before the next period starts and the next transaction occurs. Thus, the downstream party does not discount the payment.

We write the relational compensation contract as $\{b(Q)\}$, where $b(Q)$ is a payment contingent on the observation of Q[58].

First Best

The first-best actions $\{a_1^*, a_2^*\}$ maximize the expected value of Q minus the cost of actions $c(a_1, a_2) = \frac{1}{2}a_1^2 + \frac{1}{2}a_2^2$. This gives $a_1^* = 0$ and $a_2^* = \delta$ and total surplus:

$$S^* = Q(a_1^*, a_2^*) - c(a_1^*, a_2^*) = Q_0 + \frac{1}{2}\delta^2$$

Spot Market

On the Spot Market, the supplier does not take the first best actions. In particular, under both Spot Employment and Spot Outsourcing $a_2 = 0$, because the downstream firm cannot credibly commit to rewarding the supplier's quality-focused actions.

Relational Contracts

Whether the upstream party is integrated with the downstream party or not, if she accepts the relational contract, she will choose actions a_1 and a_2 to solve:

$$\max_{a_1, a_2} = b\big(Q(a_1, a_2)\big) - c(a_1, a_2)$$

It is straightforward to see that the first best can only be achieved if the contract is of the form $b\big(Q(a_1, a_2)\big) = Q(a_1, a_2) - t$, where t is a transfer independant of the surplus Q. In the remainder of this section, we assume that the relational contract is written in such a way and that under relational employment (when the downstream party owns the supplier) or under relational outsourcing (when the supplier is independent), the suppliers take the first best actions $\{a_1^*, a_2^*\}$ [59].

This relational contract is self-enforcing if both parties choose to honor it for all possible realizations of P. We next explore the feasibility of the first best contract under employment and outsourcing and show that if the shock to the alternative use P is high enough, the first best contract is only self-enforceable under Relational Employment. We use superscripts {RE, SE, RO, SO} to indicate Relational Employment, Spot Employment, Relational Outsourcing and Spot Outsourcing and {U, D, S} to denote the upstream party, downstream party and overall surplus respectively.

Relational Employment

Since $S^{SE} > S^{SO}$, [60] if one of the two party reneges, the downstream party will retain ownership and earn D^{SE} in perpetuity, while the upstream party will earn U^{SE} in perpetuity. The upstream party reneges if

[58] Alternatively, we could consider a more general relational compensation contract of the form $\{s, b(Q)\}$ as in Baker *et al.* (2002), where salary s is paid by downstream to upstream at the beginning of each period and $b(Q)$ is a payment contingent on the realization of Q. Such an assumption would not change our results below.

[59] In particular, t must be such that $t \leq Q(a_1^*, a_2^*) - c(a_1^*, a_2^*) = Q_0 + \frac{1}{2}\delta^2$ so that the downstream party would accept the contract

[60] See the proof in the main text model.

she refuses to accept the promised payment $b(Q)$. Thus, the upstream party does not renege as long as:

$$b(Q) + \frac{1}{r}U^{RE} \geq \frac{1}{r}U^{SE} \tag{5}$$

Similary, the downstream party reneges if she takes the inputs and refuses to pay the bonus to the upstream party. The downstream party honors the contract as long as:

$$\frac{1}{r}D^{RE} \geq b(Q) + \frac{1}{r}D^{SE} \tag{6}$$

Summing (5) and (6), and noting that $S^X = U^X + D^X$, we get the following necessary condition:

$$S^{RE} \geq S^{SE} \tag{7}$$

(7) is actually sufficient as well as necessary, because a transfer t can always be chosen so that when (7) is statisfied, (5) and (6) are also satisfied [61].

As $S^{RE} = S^* = Q_0 + \frac{1}{2}\delta^2$ and $S^{SE} = S^* = Q_0$, (7) is satisfied, and so **the first best can always be enforced under Relational Employment**.

Relational Outsourcing

Since $S^{SE} > S^{SO}$, if one of the two party reneges, the upstream party will purchase the ownership right from the downstream party for some price π, after which the upstream and downstream parties will earn U^{SE} and D^{SE}, respectively, in perpetuity. If the upstream party reneges on the relational-outsourcing contract, she negociates to sell the good for the spot-outsourcing price of $(1-\alpha)P + \alpha\tilde{Q}$, where α is the supplier's bargaining coefficient and $\tilde{Q}$ is the observable portion of the surplus Q as in the main text model. Thus, the upstream party honors the contract as long as:

$$b(Q) + \frac{1}{r}U^{RO} \geq (1-\alpha)P + \alpha\tilde{Q} + \frac{1}{r}U^{SE} + \pi \tag{8}$$

The timing of reneging is slightly different for the downstream party. She has no incentives to renege at the time of delivery of the inputs as Q is unobservable. Instead, the downstream party reneges if she takes the inputs and refuses to pay the bonus to the upstream party. The downstream party does not renege as long as:

$$\frac{1}{r}D^{RE} \geq b(Q) + \frac{1}{r}D^{SE} - \pi \tag{9}$$

If (8) holds for all P and $\tilde{Q}$, then it must hold for the maximum value of $(1-\alpha)P + \alpha\tilde{Q}$. Summing (8) and (9) we get the following necessary condition:

$$\frac{1}{r}S^{RO} \geq \frac{1}{r}S^{SE} + \max\{(1-\alpha)P + \alpha\tilde{Q}\} \tag{10}$$

Evaluated at $\{a_1^*, a_2^*\}$, (10) is equivalent to:

[61] For both (5) and (6) to be satisfied and the supplier to accept the contract, it must be that $Q_0 + \frac{1}{2}\frac{r}{1+r}\delta^2 \leq t \leq Q_0 + \frac{1}{2}\delta^2$

$$(1 - \alpha\gamma - \alpha)\bar{\epsilon} \leq \frac{1}{2r}\delta^2 - \alpha Q_0 \tag{11}$$

Thus, if $\bar{\epsilon}$ is high enough, **the first best contract cannot be enforced under Relational Outsourcing**.

The intuition for why quality-oriented downstream firms may need to own upstream productive assets and hire the suppliers operating the assets as employees is as follows. Under any sort of outsourcing, suppliers are free to allocate the inputs produced to their alternative use. As a result, when the value of the input is high in its alternative use (e.g. if the supplier happens to get more fish or if there is less competition on a specific day in the quantity-focused sector), quality-oriented firms may be unable to prevent the suppliers they interact with from breaking their relationship and selling the goods for its alternative use. In contrast, under Relational Employment, the downstream firm has control over the inputs, and will choose to allocate them efficiently regardless of the value of the inputs in their alternative use.

A key testable prediction of this model in our context is that (1) independent suppliers under a relational contract should not adopt a behavior consistent with delivering higher quality inputs and (2) downstream firms should not produce higher quality output when they source more of their inputs from non-integrated suppliers with whom they have a relational contract.

Empirical evidence on relational contracts in the Peruvian fishmeal industry

We now test these predictions. We show results for two different, frequency-of-interacting based observable proxies for a supplier being engaged in a relational outsourcing contract with a downstream firm: specifically, (i) that the supplier delivers more than 80 percent of its fish to the same fishmeal firm (approx. the 75th percentile of the underlying distribution) for two consecutive production seasons, and (ii) that the supplier delivers to the same firm more than 10 times (approx. the 25th percentile of the underlying distribution) in a given production season and does so for three seasons in a row. We "turn on" the inferred contract at the start of the relevant period, not when the "cut-off" used in the proxy is reached.

In Appendix Table B1, which is analogous to Table 5, we show that relational outsourcing contracts appear not to be used to incentivize supplier quality-increasing actions in the Peruvian fishmeal industry, consistent with the dynamic version of our theoretical framework above. The results show that a supplier supplying a given plant does not deliver fresher fish when engaged in repeated interactions with the firm in question, relative to more isolated instances of supplying the same plant.

In Appendix Table B2, which is analogous to Table 6, we relate output quality not only to the share of inputs coming from integrated suppliers, but also to the share coming from suppliers under relational outsourcing contracts (as defined by the proxies described above). The estimated coefficients on the share of inputs coming from integrated suppliers remain positive and highly significant, while the estimated coefficients on the share coming from suppliers under relational outsourcing contracts are very small and insignificant. These results indicate that repeated interactions are not used to incentivize the delivery of high quality inputs in the Peruvian fishmeal sector, as the model above predicts.

In combination with the results in the body of the paper, the findings in tables 5 and 6 provide support for the idea that vertical integration enables downstream firms to incentivize specific supplier behaviors—and consequently the types of output associated with those behaviors—that other organizational structures do not.

Organizational structure and supplier behavioral response to plant input quality needs

The dynamic model with relational contracts presented above also predicts the following result. When the return on the quality surplus Q of the quality-increasing action is higher (when δ increases), integrated suppliers will choose a higher level of the that action ($a_2^* = \delta$ increases). We test this prediction below.

A change in the need for input quality arises when the plant aims to produce fishmeal of the high quality type (for example because of a change in demand). As in Section 6, we compare periods when the supplier is integrated with the plant supplied and periods when the supplier is independent from but supplies the same plant, but now differentially when the downstream plant produces a low or high quality output.

We first estimate the following equation:

$$\begin{aligned} \mathrm{B}_{ijt} = \alpha + \beta_1 \; & \mathrm{I[VI \times supplies\ owner\ firm]}_{ijt} \times \mathrm{I[Low\ Quality]}_{jt} \\ & + \beta_2 \; \mathrm{I[VI \times supplies\ owner\ firm]}_{ijt} \times \mathrm{I[High\ Quality]}_{jt} \\ & + \gamma_{ij} \times \mathrm{I[High\ Quality]}_{jt} + \gamma_{ij} \times \mathrm{I[Low\ Quality]}_{jt} + \delta_t + \varepsilon_{ijt} \end{aligned} \tag{12}$$

where $\mathrm{I[Low\ Quality]}_{jt}$ is a dummy equal to 1 when plant j—i.e. the plant supplier i supplies at t—produces comparatively low quality fishmeal in the month date t falls within (and conversely for $\mathrm{I[High\ Quality]}_{jt}$).[62] We include Supplier × Plant × Quality level fixed effects (that is, $\gamma_{ij} \times \mathrm{I[High\ Quality]}_{jt}$ and $\gamma_{ij} \times \mathrm{I[Low\ Quality]}_{jt}$) to focus on the supplier's *differential* response to the plant's input needs when integrated. The other variables are as defined in equation (4).

The marginal impact of the behavioral response of a single supplier on the output quality of the plant as a whole is likely to be limited. We thus interpret the coefficient of interest as the supplier's response to the plant's *intention* to produce higher quality output.

The results in Appendix Table B3 suggest that suppliers differentially adapt their quality behavior to the current needs of the downstream plant they supply when integrated. Column 1 shows that boats tend to deliver a lower quantity per trip when integrated with the plant supplied, regardless of whether the plant produces low or high quality at the time.[63] However, columns 2 and 3 show that, when integrated, boats adjust their behavior so as to deliver fresher fish when the plant supplied is producing high quality output. When integrated, boats fish about seven percent closer to port and spend about six percent less time at sea, when the plant supplied is producing fishmeal of the high quality type Overall, the evidence confirms the prediction from the relational model that integrated suppliers will provide more of the quality focused action when its return to the quality surplus is higher.

[62]We define this dummy variable using our directly observed measure of quality at plant level. The dummy is equal to 1 if the share of the plant's production that is of high quality type is higher than the median in our sample.

[63]The estimated decrease in quantity per trip when integrating with the plant being supplied is bigger when the plant is producing low quality fishmeal. This is surprising in light of our results in sections 7 and 5. A possible explanation is that independent suppliers face strong incentives to deliver high input quantities when the plant being supplied is attempting to produce high output quantities (and prioritizing output quality less) and that integrated suppliers do not.

Appendix B tables

TABLE B1: SUPPLIER BEHAVIOR AND RELATIONAL OUTSOURCING

Panel A: Relational outsourcing = 80% of offloads to the same firm for 2 consecutive production seasons			
Dep. var:	Log(Quantity supplied)	Log(Max. distance from the plant's port)	Log(Total time spent at sea)
	(1)	(2)	(3)
I[Relational × supplies relational firm]	0.010	0.016*	−0.000
	(0.007)	(0.009)	(0.006)
Date FEs	Yes	Yes	Yes
Supplier × Plant FEs	Yes	Yes	Yes
N	315,442	137,278	159,724

Panel B: Relational Outsourcing = more than 10 interactions with the same firm for at least 3 consecutive production seasons			
Dep. var:	Log(Quantity supplied)	Log(Max. distance from the plant's port)	Log(Total time spent at sea)
	(1)	(2)	(3)
I[Relational × supplies relational firm]	−0.009	0.026	0.002
	(0.020)	(0.022)	(0.015)
Date FEs	Yes	Yes	Yes
Supplier × Plant FEs	Yes	Yes	Yes
N	315,442	137,278	159,724

Notes: One observation is a boat during a fishing trip. *Quantity supplied* is the amount of fish the boat delivers to the plant per trip. *Max. distance from the plant's port* is maximum distance a specific boat is observed away from port. *Max. distance from the plant's port* can only be measured if the boat leaves from and arrives at the same port. *Total time at sea* is the amount of time the boat is away from port per trip. The number of observations varies from one column to the next as GPS variables for a given trip are sometimes missing. We define I[Relational ×supplies relational firm] to be equal to one if the supplier is (i) currently under a relational contract (ii) currently delivering to the firm it is under a relational contract with. In Panel A, we define an independent boat as being under a relational contract if the boat delivers more than 80% of its offloads (75th percentile) to the same fishmeal firm for 2 consecutive fishing seasons. In Panel B, we define an independent boat as being under a relational contract if the boat interacts more than 10 times (25th percentile) with the same firm during a fishing season and so, for at least 3 consecutive fishing seasons. Because use Boat × Plant FEs, I[Relational ×supplies relational firm] is identified from boats moving in and out of a relational contract. Standard errors clustered at the boat level are included in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE B2: OUTPUT QUALITY AND SHARE OF INPUTS FROM VERTICALLY INTEGRATED SUPPLIERS AND SUPPLIERS UNDER A RELATIONAL OUTSOURCING CONTRACT

Panel A: First definition of relational contracts

Dep. var:	High Quality share of prod.		Protein content	
	(1)	(2)	(3)	(4)
Share of inputs from VI suppliers	0.343**	0.372**	0.950**	1.039***
	(0.153)	(0.161)	(0.348)	(0.333)
Share of inputs from relational suppliers	−0.100	−0.009	−0.159	−0.003
	(0.093)	(0.065)	(0.520)	(0.430)
High technology share of capacity	No	Yes	No	Yes
Season FEs	Yes	Yes	Yes	Yes
Firm FEs	Yes	Yes	Yes	Yes
Mean of Dep. Var.	0.75	0.75	65.7	65.7
N	275	275	208	208

Panel B: Second definition of relational contracts

Dep. var:	High Quality share of prod.		Protein content	
	(1)	(2)	(3)	(4)
Share of inputs from VI suppliers	0.395**	0.395**	0.949***	0.997***
	(0.164)	(0.162)	(0.265)	(0.250)
Share of inputs from relational suppliers	−0.424***	−0.481***	0.809	1.076
	(0.141)	(0.131)	(2.015)	(1.812)
High technology share of capacity	No	Yes	No	Yes
Season FEs	Yes	Yes	Yes	Yes
Firm FEs	Yes	Yes	Yes	Yes
Mean of Dep. Var.	0.75	0.75	65.7	65.7
N	275	275	208	208

Notes: One observation is a firm during a production season. *Protein content* is the quantity weighted average of a measure of quality inferred from a database that provides weekly prices by quality. *Log(unit price)* is the log of the quantity weighted average unit price of exports during a season. *Share of inputs from VI suppliers* is the share of a firm's (or plant's) inputs that come from VI suppliers during a season. Steel boats tend to be bigger, 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 integrated cooling system use ice to keep fish fresh. *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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE B3: SUPPLIER BEHAVIOR, VERTICAL INTEGRATION AND OUTPUT QUALITY

	Panel A		
Dep. var:	Log(Quantity supplied)	Log(Max. distance from the plant's port)	Log(Total time spent at sea)
	(1)	(2)	(3)
I[VI × supplies owner firm] ×I[Plant producing low quality]	-0.133^{***}	0.017	-0.013
	(0.043)	(0.047)	(0.031)
I[VI × supplies owner firm] ×I[Plant producing high quality]	-0.066^{**}	-0.067^{***}	-0.042^{**}
	(0.029)	(0.026)	(0.019)
Date FEs	Yes	Yes	Yes
Supplier × Plant × High Quality FEs	Yes	Yes	Yes
N	314,383	136,538	158,918
p-val - Test: two coefficients equal	0.00	0.03	0.04

Notes: One observation is a supplier during a fishing trip. This table is similar to Table 5, but with I[VI × supplies owner firm] interacted with the quality produced by the downstream plant. *Quantity supplied* is the amount of fish the boat delivers to the plant per trip. *Max. distance from the plant's port* is maximum distance a specific boat is observed away from port. *Max. distance from the plant's port* can only be measured if the boat leaves from and arrives at the same port. *Total time at sea* is the amount of time the boat is away from port per trip. I[Plant producing high quality] is a dummy equal to one if the plant the supplier delivers to produces only high quality fishmeal. The number of observations varies from one column to the next as GPS variables for on given trip are sometimes missing. Standard errors clustered at the boat level are included in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix C how demand for output quality affects vertical integration

C.1 Estimating how demand for output quality affects vertical integration

In this section we show that firms choose to integrate their suppliers when they face increased demand for high quality fishmeal. To do so, we develop an IV strategy that exploits quality-differentiated firm-specific demand shocks. We find that these shocks cause firms to acquire suppliers and to increase their Share VI.

The logic behind our instruments for the quality grade of a firm's exports at a given point in time relies on two important facts about the Peruvian fishmeal sector. First, there is an exceptionally tight link between quality grade and export destination. This is apparent in the export transactions data, where some destination countries (e.g. Chile and Japan) consistently buy higher unit price and protein content fishmeal than other countries.[64] Sales records provided by a large firm drive home this connection. Country names are frequently used as a shorthand to represent different qualities—the quality column for exports is often simply filled in with the name of a country (e.g. "Thailand quality"). An increase in demand from high quality importers should thus increase the quality content of Peruvian fishmeal exports.

The second important fact about the Peruvian fishmeal sector is that the timing of sales contracts relative to production is typically such that a firm can integrate or sell suppliers in a given production season in response to high or low demand from particular importer countries. An industry association informed us that almost all contracts for a given season's production are negotiated either before the season starts, or early in the season.

To construct our demand shocks, we follow an approach similar to Bastos *et al.* (2018) (see also Park *et al.* (2009); Brambilla *et al.* (2012)). In the second stage, we estimate how acquisitions/sales of suppliers and firms' input mix respond to the quality grade produced:

$$\text{VI}_{it} = \alpha + \beta_1 \text{Quality}_{it} + \gamma_i + \delta_t + \varepsilon_{it} \tag{13}$$

We control for firm and production season fixed effects and cluster the standard errors at firm level as in Section 7. In the first stage, quality grade produced is instrumented by demand shocks from specific destinations as follows:

$$\text{Quality}_{it} = \gamma_i + \delta_t + \sum_j \beta_j \, (\text{I}^j_{i2008} \, S^j_{-i,t}) + \varepsilon_{it} \tag{14}$$

where j is an export destination country, and $\text{I}^j_{i2008} \, S^j_{-P,t}$ are our excluded instruments. I^j_{i2008} is a dummy variable equal to one if firm i exported to destination j in 2008, the year prior to our analysis period. $S^j_{-i,t}$ is the leave-firm-out share of Peru's fishmeal exports going to country j in season t, a proxy for the relative demand for firm i coming from destination j at a given point in time. Changes in j's demand should matter more for firms that previously exported to j, which we capture in the interaction between $S^j_{-i,t}$ and I^j_{i2008}.

[64] See Appendix Table A1 for a list of the main importers of Peruvian fishmeal and the average quality imported. Some of the countries that import comparatively high quality grades of fishmeal are rich—for example Canada, Chile, and Japan—while others are middle-income. Note that, as for humans, quantity and quality of feed (the latter here defined by protein content) are highly imperfect substitutes for the animals that consume fishmeal.

C.2 Variation in foreign demand for quality and vertical integration

We find that firms respond to positive shocks to demand for high quality fishmeal by acquiring more of their suppliers and by sourcing a higher share of their inputs from suppliers that have been integrated.[65]

The OLS and the second stage IV results are reported in Table 4. The estimates in Panel A indicate that a one percentage point increase in the average protein content demanded—about 20 percent of approx. 63-68 percent range observed in Peru—induces the firm to source 29 percent more of its supply from suppliers that have been integrated. A same magnitude increase in relative demand for high quality fishmeal would lead the firm to acquire 2.1 more suppliers, increasing its stock of integrated suppliers by nine percent on average, as shown in Panel B.

Our interpretation of the results in Table 4 is that firms vertically integrate *in order to be able to produce high quality output*. A potential alternative is that the liquidity that comes along with greater demand for quality (rather than the demand for quality itself) may affect firms' ability to integrate. That is, if firms' seasonal revenues are expected to be higher when relative demand for quality is high, they may be better able to access the capital necessary to vertically integrate, but actually integrate for other reasons than to satisfy the demand for high quality. While this is unlikely since not only acquisitions and sales of suppliers but also firms' actual input mix respond to quality demand shocks, we address the concern by including controls for total seasonal sales. This has little effect on the estimated coefficients.

In the first stage we use the 20 countries that import the most fishmeal from Peru (see Appendix Table A1). In Appendix Table A3 we show that our results are robust to instead using the 10 biggest importer countries and to using LASSO regressions to choose the importer countries whose demand fluctuations most affect quality grade exported.[66] The LASSO robustness check is in our view especially informative because the procedure picks the importer countries whose imports most affect the *specific* dimension of Peruvian fishmeal exports' characteristics we are interested in—their quality grade. The first stage results for the top 20, top 10 and LASSO are reported in Appendix Table **??** [67].

Since the existing literature that uses destination country demand shocks for identification often struggles with weak instruments, we compute the Kleibergen-Paap and Anderson-Rubin Wald test statistics. Comparing the statistics reported in Table 4 to the Stock-Yogo critical values[68], while we do not pass the Kleibergen-Paap under-identification test, we reject the null hypothesis that our instruments are weak (as the F-statistic surpasses the 10 percent critical value). We also reject the hypothesis that the coefficients on the excluded instruments are jointly zero when they are included in place of quality itself in the second stage regression using the Anderson-Rubin Wald test. It is additionally important to note that weak instruments would bias the IV coefficients *downward*, i.e., towards the OLS coefficients, rather than upward. See Bastos *et al.* (2018) for a lengthier discussion of this issue in the context of "demand pull" instruments.

[65] The IV coefficients in columns 3 and 4 are bigger than the OLS coefficients in columns 1 and 2. We believe this is in part to be expected because the relationship between output quality and vertical integration *at firm level* estimated in Table 3 partly reflects a causal effect of organizational structure on output quality and partly other mechanisms, as discussed at the end of that section. If the OLS estimates in that table are biased upwards, we would expect the OLS estimates here to be biased downwards, as we study the inverse relationship.

[66] LASSO (least absolute shrinkage and selection operator) is a regression analysis method that performs both variable selection and regularization in order to enhance the prediction accuracy and interpretability of the statistical model it produces, penalizing the model for including more regressors. LASSO selects eight importer countries.

[67] The sign of the coefficient for each instrument is broadly consistent with the relative average quality imported by each country (See Appendix Table A1).

[68] Though Stock-Yogo's critical values are computed for the homoskedastic case, it is standard practice to compare the Kleibergen-Paap Wald test statistics to these critical values even when one reports standard errors that allow for heteroskedasticity.

The strategic changes in organizational structure in response to changes in the composition of demand we have shown evidence of in this section are consistent with—and expected due to—the change in supplier behavior with integration we established in Section 6 and the resulting integration→quality relationship shown in Section 7. These results represent evidence of an overlooked determinant of firms' organizational structure. We conclude that Peruvian manufacturing firms are aware of, and act on, their greater ability to produce high quality grade output when their suppliers have been integrated.

Appendix D Supplier Behavioral Adaptation and Vertical Integration

Organizational structure and supplier behavioral response to variation in production conditions

In this appendix, we provide evidence that supplier "adaptation" depends on organizational structure (See Williamson (1975, 1985) for theoretical considerations and Forbes & Lederman (2009, 2010) for empirical tests). Specifically, we look at how independent and integrated suppliers differentially adjust their behavior to important variations in production conditions.

Plankton, the primary food source of Peruvian anchovies, is an important determinant of fishing conditions at a specific location (see also Axbard, 2016; Fluckiger & Ludwig, 2015). In the map in Panel (a) of Appendix Figure C1, we depict variation in plankton concentrations[69] along the coast of Peru on a randomly picked date. Fish density in the ocean outside of fishmeal plants located in different parts of Peru differed considerably on the date shown.[70]

On a specific day, around a specific fishmeal plant's port, low plankton concentrations should tighten the supplier's trade-off between quantity- and quality-increasing actions because (i) as fish follow their feed, low plankton concentration means less fish and so the fishermen would need to provide a higher effort to capture the same quantity of inputs; and (ii) a specific school of fish captured in an area with low plankton concentration is less fed and so the quality (protein content) of the fishmeal issued from that fish will be lower[71].

In the main text model, a tightening of the quality versus quantity trade-off corresponds to the returns to the alternative use of the quantity-increasing action being lower (P is now $P = \phi a_1$ with $\phi < 1$). Our model predicts that in that case, independent suppliers would adopt an even higher quantity-increasing behavior ($a_1 = 1 - \alpha\phi\gamma - \alpha\phi < 1 - \alpha\gamma - \alpha$), while the integrated suppliers' actions would be unaffected ($a_1 = 0$). In the remainder of this section, we test whether when plankton concentration is low (when production conditions are *difficult*), integrated suppliers adopt a more quality-increasing action relative to independent suppliers.

To define conditions under which the quality-quantity tradeoff is stronger, we take a split-sample approach. Specifically, we use 2015 data to identify the conditions that lead to availability of more and better fish, and thereafter exclude 2015 data from our regressions of interest. We first define good fishing conditions for a specific location. We match the plankton data with information on how much fishing takes place in a given grid-cell, as inferred from GPS measures of boats' movements [72].

[69] We use NASA chlorophyll concentration data from satellite images. This data allows scientists to measure how much phytoplankton is growing in the ocean by observing the color of the light reflected off the water. The data is available for each date and each 0.1°-latitude×0.1°-longitude (roughly 10 kilometer×10 kilometer) grid-cell. Phytoplankton contain a photosynthetic pigment called chlorophyll that lends them a greenish color. In the rest of this Appendix, we use the term "plankton concentration" when referring to chlorophyll concentration. The data is no longer available at the date level on the NASA website (only at the week or month level), but was still available in late 2015 when we scraped the data. See `http://neo.sci.gsfc.nasa.gov/view.php?datasetId=MY1DMM_CHLORA`. Because some data points are missing, we interpolate the missing data by taking the average of date and geographical interpolations.

[70] A dynamic version of the same map would show that the spatial distribution of plankton also varies extensively across time. Panel (b) of Appendix Figure C1 shows a map of plankton concentrations on the same date around the cluster of fishmeal plants in the town of Paracas. We see that boats concentrate their fishing in areas where plankton concentrations are highest.

[71] We provide evidence of (i) in the next paragraphs. Interviews with several actors in the fishmeal industry and the second row of Appendix Table C1 confirmed assumption (ii).

[72] Since we do not directly observe when and where a boat has its nets out, we construct an algorithm to infer fishing location and -time. The algorithm exploits the fact that a boat's speed is lower when searching for fish or actively fishing than when traveling back to port. Specifically, we follow Natividad (2014) and assume that a boat has its nets out if speed is below 2.9 kilometers/hour. The industry association IFFO confirmed to us that the method should provide fairly accurate results. We have also used two alternative

The top panel of Appendix Figure C2 shows that the higher the log plankton concentration, the higher the likelihood that the location is chosen by at least one boat. The bottom panel shows the total quantity fished by all boats in the grid-cell as a function of log plankton concentration, controlling for boat fixed effects. The graph shows a positive and approximately linear relationship. Overall, Appendix Figure C2 makes clear that a higher plankton concentration is associated with better fishing conditions. We thus define a grid-cell×date as *good for fishing* if the log plankton concentration is greater than the median as defined over all grid-cells where at least one boat fishes at some point in 2015.

Our objective is to define how good the fishing conditions in the area outside of a cluster of fishmeal plants (i.e., a fishmeal port) are on a specific date. To do so, we must aggregate the grid-cells around each port to construct a port-specific measure. We first construct the share of fishing locations around a cluster of plants that are *good for fishing* on the date in question.[73] We then define a port×date as having *difficult conditions* if the share of grid-cells surrounding the location that are *good for fishing* is lower than the 10th percentile in the distribution of port×dates. In this sense, our definition of *difficult conditions* corresponds to dates when it is challenging to find fish nearby a cluster of plants. Appendix Figure C3 shows that on the dates when upstream production conditions are *difficult*, supply of fish to plants is on average 5 percent lower.

With this measure in hand, we explore whether the benefits of vertical integration to firms attempting to produce high quality output are greater when suppliers' opportunity cost of delivering high quality inputs is high. We estimate the following equation:

$$\begin{aligned}\text{Quality}_{jt} &= \alpha + \beta_1 \text{VI}_{jt} + \beta_2 \text{Difficult conditions}_{jt} \\ &\quad + \beta_3 \text{VI}_{jt} \times \text{Difficult conditions}_{jt} + \beta_4 \text{HighTech}_{jt} + \gamma_j + \delta_t + \varepsilon_{jt} \qquad (15)\end{aligned}$$

where the firm×production season level continuous variable Difficult conditions$_{jt}$ is the average of port×date *difficult conditions* indicator variables for the locations where the firm's plants are located.

The results are presented in Appendix Table C1. The second row shows that if a downstream firm is subject to more *difficult conditions* upstream during a production season, the average quality grade of its fishmeal is significantly lower. We interpret this finding as evidence that when conditions are *difficult* according to our measure, it is more challenging for suppliers not only to deliver input quantity, but also quality.[74]

The third row of Appendix Table C1 shows that a firm can reduce the impact of *difficult conditions* on the quality of its output by integrating its suppliers. Since we normalize the *difficult conditions* variable to a mean of 0, the first row can be interpreted as the total correlation between the share of inputs coming from integrated suppliers and output quality. Comparing the first row of columns 1 and 2, and columns 3 and 4, we see that when we control for *difficult conditions* and its interaction with the VI share of inputs, the correlation between VI and output quality falls significantly.[75] This indicates that vertically integrating

algorithms for inferring fishing location and -time; these yield similar results.

[73] We use only the locations that are within 145 kilometers of the port, the 95th percentile of the maximum distance from the port of delivery at which boats are observed during fishing trips. Note that we do not focus on the conditions facing a specific boat at a specific location because the boat's choice of where to fish is endogenous to its objectives on the date in question.

[74] Greater plankton availability improves the fish's fatty acid profile, which in turn results in a fishmeal of higher protein content.

[75] We conducted similar regressions at the plant level (using the dichotomous measure of plant output quality available), and also when restricting the sample to the plants belonging to the fishmeal firm that shared its data with us. The results, available from the authors, are qualitatively very similar to those in Appendix Table C1.

allows firms to partially overcome the challenges to producing high quality output that arise when upstream production conditions are difficult. This accounts for part of the correlation between integration and output quality we established in Section 7 [76].

We next explore whether the ability of integrated suppliers to help downstream firms mitigate difficult production conditions upstream is explained by their behavior at such times. Since the focus is now on suppliers, we can again use Supplier×Plant×Date level data and estimate the following equation:

$$\begin{aligned} \text{B}_{ijt} = \alpha &+ \beta_1 \text{I[VI} \times \text{supplies owner firm]}_{ijt} \times \text{I[Not difficult conditions]}_{ijt} \\ &+ \beta_2 \text{I[VI} \times \text{supplies owner firm]}_{ijt} \times \text{I[Difficult conditions]}_{ijt} \\ &+ \gamma_{ij} \times \text{I[Difficult conditions]}_{ijt} + \gamma_{ij} \times \text{I[Not Difficult conditions]}_{ijt} + \delta_t + \varepsilon_{ijt} \end{aligned} \tag{16}$$

where I[Difficult conditions]$_{ijt}$ indicates that the fishing conditions around plant j's location are *difficult* on date t as defined above (and vice versa for I[Not difficult conditions]$_{ijt}$). Similar to the approach in Appendix B, we include Supplier×Plant×Difficult conditions fixed effects ($\gamma_{ij} \times$ I[Difficult conditions]$_{ijt}$ and $\gamma_{ij} \times$ I[Not Difficult conditions]$_{ijt}$) to focus on the supplier's differential response to production conditions when integrated. The other variables are as previously defined.

The results are in Appendix Table C2. Column 1 shows that a supplier tends to deliver a lower quantity of inputs on *difficult* production days when it is integrated with the plant supplied, relative to when it is not (though the estimate is not statistically significant). More importantly, boats fish 36 percent closer to port and spend 33 percent less time at sea on days when conditions are *difficult*, when integrated with the plant supplied relative to when not. Such changes in supplier behavior are likely to significantly affect the quality of the inputs available to the downstream firm. How suppliers adjust their behavior in response to an exogenous increase in the opportunity cost of quality-actions thus helps explain why it appears especially important for downstream output quality to use integrated suppliers when upstream production conditions are *difficult*[77].

Peruvian fish suppliers face a trade-off between taking quantity- and quality-increasing actions because of the technology they operate under. This trade-off is particularly pressing when production conditions are *difficult*. At such times, integrated suppliers seems to adopt their behavior to prioritize the quality of their inputs over the quantity even more. As in Forbes & Lederman (2009, 2010), this evidence suggests that vertical integration is a way for the downstream firm to insure that suppliers adopt the right (quality-increasing) behavior when there is important variation in production conditions upstream.

[76] We also checked these results are not sensitive to how we define difficult production conditions. The corresponding tables are available from the authors.

[77] These results are also not sensitive to the way we define *difficult* production conditions. The corresponding tables are available from the authors.

Appendix C tables

TABLE C1: OUTPUT QUALITY, VERTICALLY INTEGRATED SHARE OF INPUTS, AND DIFFICULT UPSTREAM PRODUCTION CONDITIONS

Dep. var:	High Quality share of prod.		Protein content	
	(1)	(2)	(3)	(4)
Share of inputs from VI suppliers	0.478***	0.470***	1.049**	1.022**
	(0.150)	(0.126)	(0.500)	(0.460)
Difficult conditions		−0.258**		−1.844*
		(0.105)		(0.961)
Share VI × Difficult conditions		0.771*		1.840**
		(0.416)		(0.860)
High technology share of capacity	No	Yes	No	Yes
Season FEs	Yes	Yes	Yes	Yes
Firm FEs	Yes	Yes	Yes	Yes
Mean of Dep. Var.	0.70	0.70	65.4	65.4
N	228	228	167	167

Notes: One observation is a firm during a production season. *Protein content* is the quantity weighted average of a measure of quality inferred from a database that provides weekly prices by quality. *Log(unit price)* is the log of the quantity weighted average unit price of exports during a season. *Share of inputs from VI suppliers* is the share of a firm's (or plant's) inputs that come from VI suppliers during a season. Steel boats tend to be bigger, 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 integrated cooling system use ice to keep fish fresh. *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.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

TABLE C2: SUPPLIER BEHAVIOR, VERTICAL INTEGRATION AND DIFFICULT UPSTREAM PRODUCTION CONDITIONS

Dep. var:	Log(Quantity supplied)	Log(Max. distance from the plant's port)	Log(Total time spent at sea)
	(1)	(2)	(3)
I[VI × supplies owner firm] ×I[Not difficult conditions]	-0.092^{***}	-0.039^{*}	-0.017
	(0.024)	(0.020)	(0.017)
I[VI × supplies owner firm] ×I[Difficult conditions]	-0.110	-0.355^{**}	-0.330^{***}
	(0.154)	(0.151)	(0.029)
Date FEs	Yes	Yes	Yes
Supplier × Plant × Difficult conditions FEs	Yes	Yes	Yes
N	223,698	12,627	141,412
p-val - Test: 2 coefficients equal	0.90	0.02	0.00

Notes: One observation is a supplier during a fishing trip. The number of observations is lower than in Table 5 as the year 2015 is excluded from the sample. (This year is used to define the plankton concentration threshold at which the production conditions can be considered as difficult). *Quantity supplied* is the amount of fish the boat delivers to the plant per trip. *Max. distance from the plant's port* is maximum distance a specific boat is observed away from port. *Max. distance from the plant's port* can only be measured if the boat leaves from and arrives at the same port. *Total time at sea* is the amount of time the boat is away from port per trip. The number of observations varies from one column to the next as GPS variables for a given trip are sometimes missing. I[*Difficult conditions*] is a dummy equal to 1 when the share of "good fishing locations" [Log(plankton concentration)>0.5] around a specific plant on a specific day is less than 5 percent (this corresponds to the bottom 10th percentile in the distribution of share of good fishing locations in our sample). The number of observations varies from one column to the next as GPS variables for on given trip are sometimes missing. Standard errors clustered at the boat level are included in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix C figures

FIGURE C1: MAP OF PHYTOPLANKTON CONCENTRATION ALONG THE COAST OF PERU

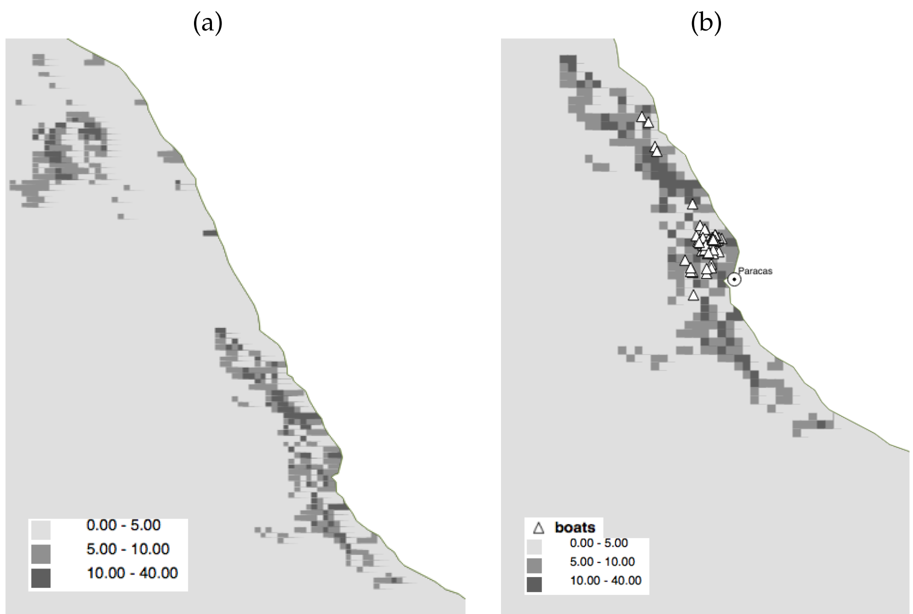

Notes: Panel (a) of this figure shows the distribution of plankton along the coast of Peru on December 10, 2012, as an example. A darker grey indicates a higher phytoplankton concentration (in mg/m^3). Panel (b) shows the same map zoomed around the port of Paracas, and the white triangles show where the boats offloading in Paracas last fished on a given trip. Fishing activity is proxied by the boat having a speed lower than than 2.9kms/hour maintained for at least half an hour as discussed in the text of Appendix D.

FIGURE C2: PLANKTON CONCENTRATION, FISHING LOCATIONS, AND QUANTITY SUPPLIED

(a)

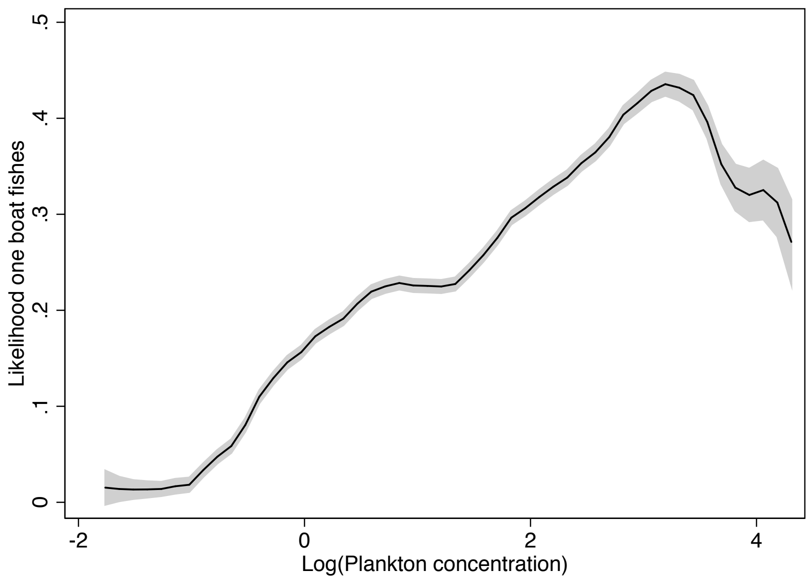

(b)

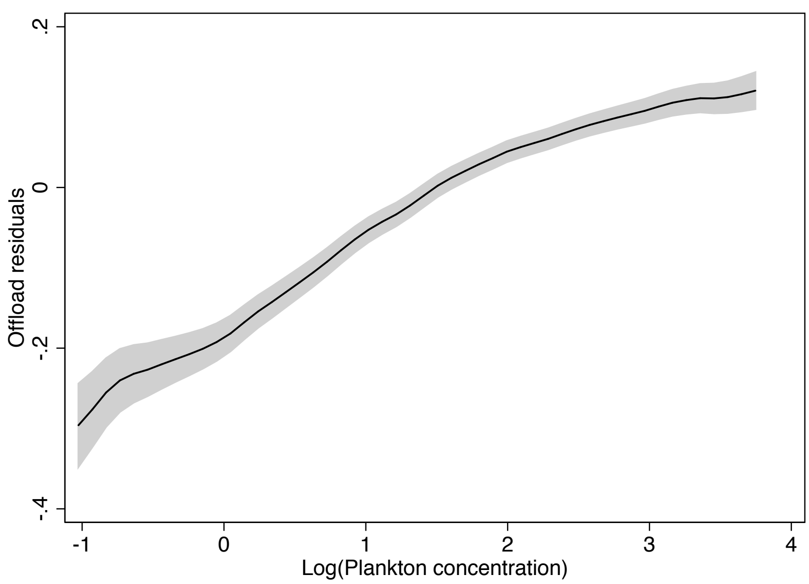

Notes: Panel (a) of this figure shows the likelihood that a boat fishes in a specific 0.1 degree×0.1 degree (roughly 10 kilometer×10 kilometer) grid-cell as a function of Log(phytoplankton concentration) at that location. Only locations were a boat fishes at least once during our data period and only the days when at least one boat goes out fishing are included. Panel (b) shows the residuals of a regression of quantity of fish caught in the grid-cell on boat fixed effects as a function of the Log(phytoplankton concentration). Catches are proxied by the boat having a speed lower than 2.9 kilometers/hour maintained for at least half an hour as discussed in the text of Appendix D.

FIGURE C3: DIFFICULT UPSTREAM PRODUCTION CONDITIONS AND QUANTITY SUPPLIED

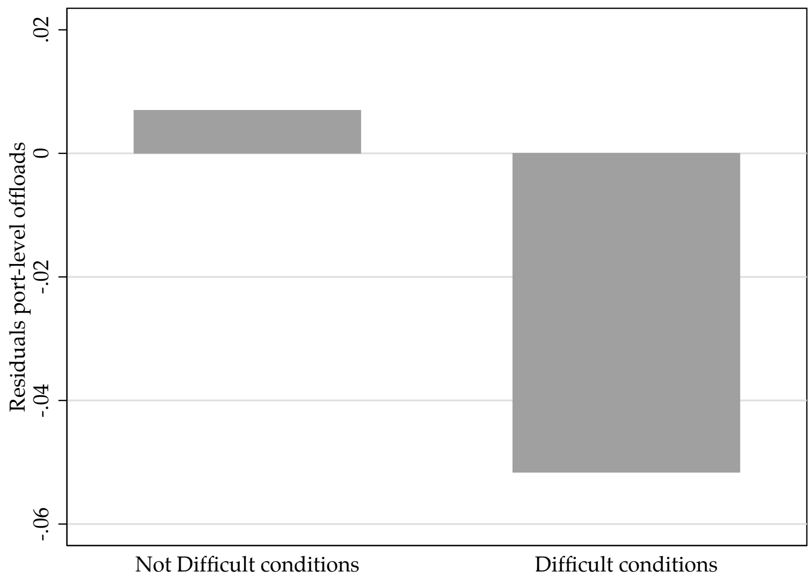

Notes: This graph shows how port residualized Log(fish offloads) vary with fishing conditions. *Difficult conditions* is defined in the text of Appendix D.